

Refining Tournament Solutions via Margin of Victory

Markus Brill

Technische Universität Berlin
Chair of Efficient Algorithms
brill@tu-berlin.de

Ulrike Schmidt-Kraepelin

Technische Universität Berlin
Chair of Efficient Algorithms
u.schmidt-kraepelin@tu-berlin.de

Warut Suksompong

University of Oxford
Department of Computer Science
warut.suksompong@cs.ox.ac.uk

Abstract

Tournament solutions are frequently used to select winners from a set of alternatives based on pairwise comparisons between alternatives. Prior work has shown that several common tournament solutions tend to select large winner sets and therefore have low discriminative power. In this paper, we propose a general framework for refining tournament solutions. In order to distinguish between winning alternatives, and also between non-winning ones, we introduce the notion of *margin of victory (MoV)* for tournament solutions. MoV is a robustness measure for individual alternatives: For winners, the MoV captures the distance from dropping out of the winner set, and for non-winners, the distance from entering the set. In each case, distance is measured in terms of which pairwise comparisons would have to be reversed in order to achieve the desired outcome. For common tournament solutions, including the top cycle, the uncovered set, and the Banks set, we determine the complexity of computing the MoV and provide worst-case bounds on the MoV for both winners and non-winners. Our results can also be viewed from the perspective of bribery and manipulation.

1 Introduction

A number of practical choice scenarios involving pairwise comparisons of alternatives can be modeled using tournaments. For instance, the pairwise comparisons could represent match outcomes when alternatives are teams in a round-robin sports competition, or the results of pairwise majority comparisons when the alternatives are candidates in an election. In order to select the set of “winners” from a tournament, several methods, known in the literature as *tournament solutions*, have been proposed. Over the past decades, many of these tournament solutions have been extensively studied from both the axiomatic and the computational point of view (Laslier 1997; Brandt, Brill, and Harrenstein 2016). Due to their generality and wide range of applications, the study of tournament solutions and their properties has attracted considerable attention from the multiagent systems research community in recent years (e.g., Brandt, Brill, and Harrenstein, 2018; Aziz et al., 2015; Mnich, Shrestha, and Yang, 2015)

Although tournament solutions provide a rich supply of procedures for choosing tournament winners according to

various criteria, they often exhibit low discriminative power because the chosen winner sets tend to be large. Indeed, previous work has shown that common tournament solutions such as the top cycle, the uncovered set, the Banks set, and the minimal covering set almost never exclude any alternative in a random tournament (Fey 2008; Scott and Fey 2012), while the bipartisan set includes on average half of the alternatives in the winner set (Fisher and Ryan 1995).¹ This naturally raises the question of how tournament solutions can be refined in order to differentiate among the winners of a given tournament.

In this paper, we propose a general framework for refining tournament solutions and for distinguishing among the winners—as well as among the non-winners—of a tournament. We introduce the concept of *margin of victory (MoV)* for tournaments, which captures how close a winner is to dropping out of the winner set, and by symmetry how close a non-winner is to entering the winner set. For a given tournament and weights on the tournament edges, the MoV of a winner is defined as the minimum total weight of edges whose reversals take it out of the winner set. Analogously, the MoV of a non-winner is defined as the negative of the minimum total weight of edges whose reversals bring it into the winner set. An important special case is when the edges are unweighted: in this case, the problem reduces to finding the minimum number of edges to be reversed.

The edge weights in our MoV framework can be interpreted in a number of different ways. Generally speaking, they represent the strength of the edges or the cost that one incurs by reversing them. In an election, a weight may reflect the proportion of voters who agree with the corresponding pairwise comparison, while in a sports competition, it may indicate the gap between the two teams in the match result. Alternatively, our refinements can also be viewed through the lens of bribery and manipulation. In this context, the weights express the amount of bribe that a manipulator needs to pay in order to reverse a pairwise comparison; the recipients of the bribe are voters in the case of an election and teams or referees in the case of a sports competition.

¹These results assume that tournaments are chosen from the uniform distribution. Brandt and Seedig (2016) and Saile and Suksompong (2018) relax this assumption and study the discriminative power of tournament solutions when tournaments are generated according to different stochastic models.

	MoV for Winners		MoV for Non-Winners		Bounds on MoV	
	unweighted	weighted	unweighted	weighted	lower bound	upper bound
Copeland (CO)	P (3)	P (3)	P (8)	P (8)	$-(n-2)$ (16)	$\lfloor n/2 \rfloor$ (14)
Top Cycle (TC)	P (5)	P (5)	P (12)	P (12)	-1 (15)	$\lfloor n/2 \rfloor$ (14)
Uncovered Set (UC)	P (5)	P (5)	$n^{O(\log n)}$ (9)	NP-h (10)	$-\lceil \log_2 n \rceil$ (17)	$\lfloor n/2 \rfloor$ (14)
3-kings	P (5)	P (5)	P (15)	NP-h (11)	-1 (15)	$\lfloor n/2 \rfloor$ (14)
k -kings (for $k \geq 4$)	NP-h (6)	NP-h (6)	P (15)	NP-h (11)	-1 (15)	$\lfloor n/2 \rfloor$ (14)
Banks set (BA)	NP-h (7)	NP-h (7)	NP-h (13)	NP-h (13)	$-\lceil \log_2 n \rceil$ (17)	$\lfloor n/2 \rfloor$ (14)

Table 1: Overview of our results, with n denoting the number of alternatives in the tournament. The computational results for Copeland (first four entries) also follow from Faliszewski et al. (2009); for completeness, we give proofs tailored to our setting. The numbers in parentheses refer to the corresponding theorem or corollary numbers.

1.1 Our Results

We study the computational complexity of the MoV with respect to several common tournament solutions, including the Copeland set, the top cycle, the uncovered set, and the Banks set. For each tournament solution, we determine the complexity of computing the MoV for both winners and non-winners, in both the unweighted and weighted setting. In addition, we derive tight or asymptotically tight lower and upper bounds on the MoV for all of the considered tournament solutions in the unweighted setting. An overview of our results can be found in Table 1.

1.2 Related Work

Kruger and Airiau (2017) considered refinements of several tournament solutions based on their *binary tree representations*. This approach can only be applied to tournament solutions that admit such a representation, and different representations may yield different refinements.

While our work is the first to consider a MoV concept for tournament solutions (to the best of our knowledge), a related notion with the same name has been extensively explored in the context of voting. Unlike in our setting where the MoV serves the purpose of distinguishing among alternatives, in voting the MoV is typically used to measure the robustness of election outcomes (Cary 2011; Magrino et al. 2011; Xia 2012; Dey and Narahari 2015). As such, the notion there is defined for election outcomes as a whole rather than for individual alternatives. The same holds for the robustness measure of Shiryaev, Yu, and Elkind (2013).

A long line of work has investigated various forms of bribery and manipulation in tournaments. This includes manipulating the tournament bracket to help a certain candidate win the tournament (Vu, Altman, and Shoham 2009; Vassilevska Williams 2010; Kim, Suksompong, and Vassilevska Williams 2017; Aziz et al. 2018) and bribing players to intentionally lose matches (Russell and Walsh 2009; Kim and Vassilevska Williams 2015; Mattei et al. 2015; Konicki and Vassilevska Williams 2019). In particular, Russell and Walsh (2009) considered a model where only a given subset of edges can be reversed, while other edges are assumed to be fixed. This constitutes a special case of our weighted setting, with sufficiently high weights on fixed edges.

In the context of bribery in voting, Faliszewski et al. (2009) considered a “microbribery” setting in which voters can be bribed to change individual pairwise comparisons between candidates, even if this results in intransitive preferences of the voter. This corresponds to our weighted setting, with weights given by pairwise majority margins.

Finally, a closely related problem is that of finding possible (resp., necessary) winners of partially specified tournaments: Given a tournament with some missing edges, the goal is to determine whether a certain alternative can be a winner for some (resp., all) completions of the tournament (Aziz et al. 2015). We observe that both problems can be reduced to computing the MoV in the weighted setting, by considering an arbitrary completion of the partial tournament and making the original edges prohibitively expensive to reverse.

2 Preliminaries

A tournament $T = (V, E)$ is a directed graph such that there is exactly one directed edge between every pair of vertices. The vertices of a tournament T , denoted $V(T)$, are often referred to as *alternatives*. Let $n = |V(T)|$. The set of directed edges of T , denoted $E(T)$, represents an asymmetric and connex *dominance relation* on the set of alternatives. An alternative x is said to *dominate* another alternative y if $(x, y) \in E(T)$ (i.e., there is a directed edge from x to y). When the tournament is clear from the context, we often write $x \succ y$ to denote $(x, y) \in E(T)$. By definition, for each pair x, y of distinct alternatives, either x dominates y ($x \succ y$) or y dominates x ($y \succ x$), but not both.

For a given tournament T and $x \in V(T)$, the *dominion* of x , denoted by $D(x)$, is defined as the set of alternatives y such that $x \succ y$. Similarly, the set of *dominators* of x , denoted by $\bar{D}(x)$, is defined as the set of alternatives y such that $y \succ x$. An alternative $x \in V(T)$ is said to be a *Condorcet winner* in T if it dominates every other alternative (i.e., $D(x) = V(T) \setminus \{x\}$), and a *Condorcet loser* in T if it is dominated by every other alternative. See Figure 1 for an example tournament.

The dominance relation can be extended to sets by writing $X \succ Y$ if $x \succ y$ for all $x \in X$ and all $y \in Y$. A set $X \subseteq V(T)$ is called a *dominating set* in T if every alternative outside of X is dominated by at least one alternative in X .




Figure 1: Tournament T with $V(T) = \{a, b, c, d, e, f\}$. All omitted edges are assumed to point from right to left (e.g., $D(f) = \{a, b, d, e\}$ and a is a Condorcet loser in T).

For $U \subseteq V(T)$, $T|_U$ denotes the restriction of T to U , and T_{-x} is short for $T|_{V(T)\setminus\{x\}}$. For an edge $e = (x, y)$, we let \bar{e} denote its reversal, i.e., $\bar{e} = (y, x)$. Similarly, for a set of edges $R \subseteq E(T)$, we define $\bar{R} = \{\bar{e} : e \in R\}$.

A *tournament solution* is a function that maps each tournament to a nonempty subset of its alternatives, usually called the set of *winners* or the *choice set*. The set of winners of a tournament T with respect to a tournament solution S is denoted by $S(T)$. The tournament solutions considered in this paper are as follows:

- The *Copeland set* (CO) is the set of alternatives with the largest dominion, i.e., $CO(T) = \arg \max_{x \in V(T)} |D(x)|$.
- The *top cycle* (TC) is the (unique) smallest set B of alternatives such that $B \succ V(T) \setminus B$. Equivalently, the top cycle is the set of alternatives that can reach every other alternative via a directed path.
- The *uncovered set* (UC), is the set of alternatives that are not “covered” by any other alternative. An alternative x *covers* another alternative y if $D(y) \subseteq D(x)$. Equivalently, UC is the set of alternatives that can reach every other alternative via a directed path of length at most two.
- The set of k -*kings*, for an integer $k \geq 3$, is the set of alternatives that can reach every other alternative via a directed path of length at most k .
- The *Banks set* (BA) is the set of alternatives that appear as the Condorcet winner of some inclusion-maximal transitive subtournament.²

All of the above tournament solutions satisfy *Condorcet-consistency*: Whenever a Condorcet winner exists, it is chosen as the unique winner.

It is clear from the definitions that UC (the set of “2-kings”) is contained in the set of k -kings for any $k \geq 3$, which is in turn a subset of TC . Moreover, both CO and BA are contained in UC (Laslier 1997).

3 Margin of Victory in Tournaments

We define the margin of victory (MoV) for a winning (resp., non-winning) alternatives in terms of sets of edges whose reversals result in the alternative becoming a non-winner (resp., winner). Edge sets with this property will be called destructive (resp., constructive) reversal sets. To formally define these notions, we need some notation. For a tournament T and a set $R \subseteq E(T)$ of edges, we let T^R denote

²A transitive subtournament is *inclusion-maximal* if it is not contained in any other transitive subtournament. If an alternative x dominates all alternatives in a transitive subtournament T' , we say that x *extends* T' .

	a	b	c	d	e	f
$\text{MoV}_{UC}(x, T)$	-2	-1	1	1	1	2
min DRS/CRS	fa, da	fb	cf	dc	ed	fe, fb

Table 2: MoV values and minimal reversal sets with respect to UC for the tournament T in Figure 1 (unweighted setting). For improved readability of reversal sets, we omit set braces and use xy to denote edge (x, y) .

the tournament that results from T when reversing all edges in R , i.e., $V(T^R) = V(T)$ and $E(T^R) = (E(T) \setminus R) \cup \bar{R}$.

Fix a tournament solution S and consider a tournament T . An edge set $R \subseteq E(T)$ is called a *destructive reversal set* (DRS) for $x \in S(T)$ if $x \notin S(T^R)$. Analogously, R is called a *constructive reversal set* (CRS) for $x \in V(T) \setminus S(T)$ if $x \in S(T^R)$.³

In general, destructive and constructive reversal sets are not unique, and finding *some* DRS or CRS is usually easy. For example, for all Condorcet-consistent tournament solutions S , a straightforward CRS for an alternative $x \notin S(T)$ is given by $R = \{(y, x) : y \in \bar{D}(x)\}$. This is because x is a Condorcet winner in T^R .

We furthermore assume that we are given a weight function $w : E(T) \rightarrow \mathbb{R}^+$ that assigns a positive weight $w(e) > 0$ to each edge $e \in E(T)$.⁴ The weight of an edge can be thought of as the *cost* that is incurred by reversing the edge. The cost of a set $R \subseteq E(T)$ is $w(R) = \sum_{e \in R} w(e)$.

A natural special case is the setting in which reversing is equally costly for all edges. In this *unweighted* setting, we assume $w(e) = 1$ for all $e \in E(T)$, and finding a minimum cost reversal set reduces to finding a reversal set of minimum cardinality.

We are now ready to define the main concept of this paper.

Definition 1. For a tournament solution S and a tournament T , the *margin of victory* of $x \in S(T)$ is given by

$$\text{MoV}_S(x, T) = \min\{w(R) : R \text{ is a DRS for } x \text{ in } T\},$$

and for an alternative $x \in V(T) \setminus S(T)$, it is given by

$$\text{MoV}_S(x, T) = -\min\{w(R) : R \text{ is a CRS for } x \text{ in } T\}.$$

By definition, $\text{MoV}_S(x, T)$ is positive if $x \in S(T)$, and negative otherwise. In the unweighted setting, all MoV values are (positive or negative) integers.

Example 2. Consider the tournament T in Figure 1. It can be easily verified that $UC(T) = \{c, d, e, f\}$. For the unweighted setting, Table 2 gives the MoV values for this tournament with respect to the uncovered set, together with examples of minimal destructive or constructive reversal sets.

³The terms “destructive” and “constructive” are borrowed from the literature on control and bribery in voting (e.g., Faliszewski and Rothe, 2016), where the goal is either to prevent a given candidate from winning (destructive control/bribery) or to make a given candidate a winner (constructive control/bribery).

⁴We forbid zero-weight edges for technical reasons. Their existence can be imitated by setting their cost to a small $\epsilon > 0$.

Note that minimal reversal sets are generally not unique, and that a minimal reversal set for an alternative x may exclusively consist of edges *not* incident to x (e.g., $\{(f, e)\}$ is a minimal CRS for b in Example 2).

4 Computing the MoV for Winners

We now study the computational complexity of computing the MoV for winners. We are given a tournament T , a weight function $w : E(T) \rightarrow \mathbb{R}^+$, a tournament solution S , and an alternative $x \in S(T)$; the task is to compute $\text{MoV}_S(x, T)$. Clearly, a polynomial-time algorithm for the weighted setting also applies to the unweighted setting, while a hardness result in the unweighted setting implies one for the weighted setting. In all cases where we provide a polynomial-time algorithm (i.e., entries “P” in Table 1), our algorithm not only determines the MoV value, but also finds a minimum DRS (or CRS when considering non-winners). Omitted proofs can be found in the full version of this paper (Brill, Schmidt-Kraepelin, and Suksompong 2019).

4.1 Copeland

The MoV for Copeland has already been studied (under different names) in slightly different settings (Faliszewski et al. 2009; Russell and Walsh 2009). In particular, Theorem 3.7 of Faliszewski et al. (2009) implies that the MoV for Copeland winners can be computed efficiently whenever the weights correspond to pairwise majority margins resulting from a preference profile. For completeness, we provide a (simpler) proof tailored to our setting.⁵

Theorem 3. *Computing the MoV of a CO winner in the weighted setting can be done in polynomial time.*

4.2 Uncovered Set, k -Kings and Top Cycle

The problems of computing the MoV for UC , k -kings and TC are not only closely related to each other but also to the theory of network flows. Since UC can be interpreted as 2-kings and TC as $(n-1)$ -kings, we only refer to k -kings and assume that k can be chosen from $\{2, \dots, n-1\}$. A DRS for x is then an edge set R such that x has distance greater than k to at least one alternative y in T^R .

Finding a minimum DRS is closely related to finding ℓ -length bounded s - t -cuts of minimum capacity. In the latter problem, we are given a directed network $G = (V, E)$ with a capacity function $u : E \rightarrow \mathbb{R}^+$, two distinguished nodes $s, t \in V$ and a length bound $\ell \in \mathbb{N}$. An edge set $C \subseteq E$ is an ℓ -length bounded s - t -cut if all s - t -paths in $(V, E \setminus C)$ have length greater than ℓ . The set C is a *minimum ℓ -length bounded s - t -cut* if it minimizes the sum of the capacities of edges in C . When $\ell \geq |V(G)| - 1$, the problem is equivalent to the standard minimum cut problem and can be solved via linear programming due to the well known max-flow min-cut theorem (Ford and Fulkerson 1956). However, for general $\ell \in \mathbb{N}$, Adámek and Koubek (1971) showed that a generalization of this theorem does not hold. More recently, it

⁵In the *unweighted* case, a greedy approach suffices to compute the MoV of a Copeland winner. However, this case is not particularly interesting, as it can be easily verified that $\text{MoV}_{CO}(x, T) = 1$ for all $x \in CO(T)$ whenever $|CO(T)| > 1$.

was shown by Baier et al. (2010) that finding a minimum ℓ -length bounded s - t -cut is NP-hard for $\ell \in \{4, \dots, n^{1-\epsilon}\}$ for fixed $\epsilon > 0$, even if capacities are uniform. By contrast, for $\ell \leq 3$, Mahjoub and McCormick (2010) showed that there exists a polynomial-time algorithm which reduces the problem to a standard cut problem. In the following, we show how we can adjust and apply these results to our setting.

Despite its similarity to our problem (which can be observed by setting $G = T$, $u(e) = w(e)$, $\ell = k$, and $s = x$), the problem described above differs in three ways from the problem under consideration. First, the node which should be disconnected, in this case t , is specified; second, edges are removed instead of reversed; and third, the graph is not restricted to be a tournament. For ease of presentation, we define a new problem which lies in between MoV for k -kings winners and minimum ℓ -length bounded s - t -cuts.

For a network $G = (V, E)$, we say that $C \subseteq E$ is an ℓ -length bounded s -cut if it is a ℓ -length bounded s - t -cut for some $t \in V \setminus \{s\}$. We say that C is a *minimum ℓ -length bounded s -cut*, if it is a minimum ℓ -length bounded s - t -cut and capacity-minimizing among all $t \in V \setminus \{s\}$. Computing a minimum ℓ -length bounded s -cut can be reduced to computing a minimum ℓ -length bounded s - t -cut by iterating over all $t \in V \setminus \{s\}$.

The following lemma formalizes the connection between length bounded cuts and DRSs for k -kings. Though intuitive, note that the statement is not obvious because reversing the edges of a cut may create new paths of bounded length.⁶

Lemma 4. *A set $R \subseteq E(T)$ is a minimum DRS for x w.r.t. k -kings iff R is a minimum k -length bounded x -cut in T .*

Since there exist polynomial-time algorithms for computing minimum ℓ -length bounded s - t -cuts for $\ell \leq 3$ and $\ell = n - 1$ (Mahjoub and McCormick 2010; Ford and Fulkerson 1956), Lemma 4 immediately yields polynomial-time algorithms for the minimum ℓ -length bounded s -cut problem for $\ell \in \{2, 3, n - 1\}$.

Corollary 5. *Computing the MoV of a UC winner, a 3-king or a TC winner in the weighted setting can be done in polynomial time.*

The following result is obtained by carefully adjusting the proof of Baier et al. (2010) showing that approximating minimum ℓ -length bounded cuts for $\ell \geq 4$ is NP-hard. We give the entire proof in the full version of this paper, where we also point out deviations from the original construction.

Theorem 6. *For any constant $k \geq 4$, computing the MoV of a k -king in the unweighted setting is NP-hard. For any constant $\epsilon > 0$, the problem is still NP-hard when we restrict to non-constant $k \geq n^{1-\epsilon}$.*

Proof sketch. We reduce from vertex cover; see Figure 2 for the construction for $k = 4$. Lemma 4 implies that determining the MoV of node x with respect to 4-kings is equivalent to computing the cost of a 4-length bounded minimum x -cut. The key part of the proof is to show that, for any $c \leq |V(G)|$,

⁶In the full version of this paper, we give an example showing that Lemma 4 does not hold for cuts that are not minimal.




Figure 2: Illustration of the construction used in the proof of Theorem 6 for the case $k = 4$. For any graph G (left image), a tournament T is constructed by introducing node gadgets and edge gadgets as follows. A *node gadget* N_v consists of four nodes v_1, v_2, v_3, v_4 and three supernodes $\bar{v}_1, \bar{v}_2, \bar{v}_3$, where the latter are tournaments themselves. The center image shows the node gadget for node v . An *edge gadget* for $e = \{u, v\}$ consists of two nodes e_1, e_2 and edges connecting the node gadgets of u and v ; see the right image. Nodes x and y are connected to all node gadgets as illustrated. All omitted edges point “backwards” (from right to left) and the direction of vertical edges, if not specified, can be chosen arbitrarily.

there exists a vertex cover in G of size c iff there exists a 4-bounded x -cut in T of size $c + |V(G)|$. For the direction from left to right, a vertex cover U can be translated to a 4-bounded x - y -cut by including edges ℓ_v and r_v (depicted by red dashed edges) whenever $v \in U$ (depicted by a red dashed node), and m_v otherwise. For the other direction, we argue that any 4-bounded x -cut of size $c + |V(G)|$ can be translated to a 4-bounded x - y -cut which includes only edges of the type ℓ_v, r_v and m_v and is of no greater size. Reversing the previously described transformation gives us a vertex cover of size c . The proof is then extended to $k > 4$. \square

4.3 Banks Set

Deciding whether an alternative x is contained within the Banks set of a tournament T , and hence deciding whether $\text{MoV}_{BA}(x, T) > 0$, is NP-complete (Woeginger 2003). Our next result shows that determining $\text{MoV}_{BA}(x, T)$ is computationally intractable *even* if we know that x is a Banks winner in tournament T .

Theorem 7. *Computing the MoV of a BA winner in the unweighted setting is NP-hard.*

Proof. We reduce from the NP-hard problem of determining whether an alternative is contained in the Banks set (Woeginger 2003). Take any instance of that problem, which consists of a tournament T and one of its alternatives x . Add two new alternatives $y, z \notin V(T)$ so that y dominates only $\bar{D}(x) \cup \{z\}$, and z dominates only $D(x)$. Call the resulting tournament T' (see Figure 3). Observe that $x \in BA(T')$: the transitive subtournament $T'_{\{x,y\}}$ cannot be extended, since no alternative dominates both x and y . We claim that $\text{MoV}_{BA}(x, T') = 1$ if and only if $x \notin BA(T)$.

First, assume that $x \notin BA(T)$. We show that $R = \{x, y\}$ is a DRS for x . Consider any transitive subtournament in $T'' = (T')^R$ with Condorcet winner x . This tournament cannot include y , but may include z . Since $x \notin BA(T)$, there exists an alternative w in T that dominates all alternatives in the subtournament. In particular, since $w \in \bar{D}(x)$,

w also dominates z . Hence the transitive subtournament can be extended by w , implying that $x \notin BA(T'')$.

If $x \in BA(T)$, we claim that $\text{MoV}_{BA}(x, T') > 1$. Since $x \in BA(T)$, there exists a transitive subtournament in T with Condorcet winner x that cannot be extended by any alternative in T . Moreover, since x dominates both y and z , this subtournament cannot be extended by y or z . Unless we reverse an edge in T or the edge (x, z) , this subtournament still cannot be extended. If we reverse the edge xz , the transitive subtournament $T'_{\{x,y\}}$ cannot be extended. Else, if we reverse an edge in T , the transitive subtournament $T'_{\{x,y,z\}}$ cannot be extended. Hence, there is no DRS for x of size one, as claimed. \square




Figure 3: Illustration of the tournament T' constructed in the proof of Theorem 7.

5 Computing the MoV for Non-Winners

We now turn to computing the MoV for non-winners.

5.1 Copeland

Similarly to the winner case, the results by Faliszewski et al. (2009) already imply that the MoV for non-winners can be computed in polynomial time. For completeness, we remark that a greedy algorithm suffices for our unweighted setting, and present an easy network flow approach for the weighted case.

Theorem 8. *Computing the MoV of a CO non-winner in the weighted setting can be done in polynomial time.*

5.2 Uncovered Set, k -Kings, and Top Cycle

To get x into the uncovered set, we need its dominion to be a dominating set in T_{-x} . Since a tournament with n vertices always has a dominating set of size $\lceil \log_2 n \rceil$ (Megiddo and Vishkin 1988), we do not need to flip more than $\lceil \log_2 n \rceil$ edges. This also means that there exists an $n^{O(\log n)}$ algorithm for finding the minimum number of necessary edge reversals, as we can try all combinations of at most $\lceil \log_2 n \rceil$ vertices to add to the dominion of x . Megiddo and Vishkin (1988) also proved that the problem of finding a dominating set of minimum size in a tournament, which we call MINIMUM DOMINATING SET, is unlikely to admit a polynomial-time algorithm: the existence of such an algorithm would have unexpected implications on the satisfiability problem. We present a reduction from MINIMUM DOMINATING SET to the problem of computing the MoV for UC non-winners, which means that the latter problem is also unlikely to admit an efficient algorithm.

Theorem 9. *Computing the MoV of a UC non-winner in the unweighted setting is at least as hard as MINIMUM DOMINATING SET.*

Proof. Consider an instance of MINIMUM DOMINATING SET given by a tournament T . Define a new tournament T' by adding an alternative $x \notin V(T)$ to T , and by making x a Condorcet loser in T' . We claim that $-\text{MoV}_{UC}(x, T')$ is equal to the minimum size of a dominating set in T . For any dominating set in T , we obtain a constructive reversal set for x in T' consisting of the edges between x and all members of the set. On the other hand, consider a CRS R for x in T' . Suppose that R contains an edge (z, y) with $x \notin \{z, y\}$, such that $y \succ z$ in T'^R . The only alternative that this reversal can help x reach in two steps is z . In this case, we can instead include (z, x) in R and maintain the property that x can reach all other alternatives in at most two steps. Hence there is always a minimal CRS that only contains edges incident to x . The alternatives involved in this CRS besides x form a dominating set in T . \square

In the unweighted setting, minimum CRSs w.r.t. k -kings ($k \geq 3$) are single edges (see Theorem 15 in Section 6.2) and hence can be found efficiently. In the weighted setting, we show hardness for UC and k -kings and tractability for TC.

Theorem 10. *Computing the MoV of a UC non-winner in the weighted setting is NP-hard.*

Theorem 11. *For any constant $k \geq 3$, computing the MoV of a non- k -king in the weighted setting is NP-hard. For any constant $\epsilon > 0$, the problem is still NP-hard when we restrict to non-constant $k \geq (1 - \epsilon)n$.*

Theorem 12. *Computing the MoV of a TC non-winner in the weighted setting can be done in polynomial time.*

5.3 Banks Set

For Banks non-winners, we present an analogous result as in the winner case: even if we know that x has a negative MoV in tournament T , determining $\text{MoV}_{BA}(x, T)$ is intractable.

Theorem 13. *Computing the MoV of a BA non-winner in the unweighted setting is NP-hard.*

6 Bounds on the Margin of Victory

In this section, we consider the unweighted setting and establish bounds on the MoV values for winners and non-winners. There are at least two insights that one could draw from these bounds. First, tournament solutions with a low absolute value of MoV bound are easily manipulable: indeed, if the absolute value of the MoV bound is low, then a manipulator can always obtain the desired outcome by reversing a small number of edges regardless of the tournament instance. Second, knowing these bounds is useful for interpreting MoV values for specific tournaments. For example, one can calculate the *normalized MoV* by dividing the actual MoV value by the bound. The resulting ratio provides a relative measure of how far away an alternative is from winning or losing; in contrast to the standard MoV measure, the normalized MoV enables us to make comparisons between tournaments of different sizes.

6.1 Upper Bounds for Winners

We show that for all considered tournament solutions, one may need to reverse up to $\lfloor n/2 \rfloor$ edges to take a winner out of the winner set, but no more.

Theorem 14. *Let $S \in \{CO, TC, UC, BA, k\text{-kings}\}$, where $k \geq 3$. For any tournament T and any $x \in S(T)$, we have $\text{MoV}_S(x, T) \leq \lfloor n/2 \rfloor$. Moreover, this bound is tight.*

Proof. Since all of the tournament solutions considered are contained in TC, an upper bound for TC carries over to the other solutions as well. By analogous reasoning, it suffices to show the tightness of the bound for BA and CO.

We first prove the upper bound. Let y be an arbitrary Copeland winner in T_{-x} . Since T_{-x} consists of $n - 1$ alternatives, y dominates at least $\lceil (n - 2)/2 \rceil = \lceil n/2 \rceil - 1$ other alternatives. Hence, we can make y a Condorcet winner in T by reversing at most $(n - 1) - (\lceil n/2 \rceil - 1) = \lfloor n/2 \rfloor$ edges. Since TC is Condorcet-consistent, $\lfloor n/2 \rfloor$ edge reversals suffice to take x out of TC.

Next, we show the lower bound for BA. Assume first that n is even, say $n = 2\ell$. Besides x , suppose that T contains alternatives $y_1, \dots, y_{2\ell-1}$, which are placed around a circle in clockwise order. Each alternative dominates the $\ell - 1$ following alternatives in clockwise order (e.g., y_1 dominates y_2, \dots, y_ℓ), and all $2\ell - 1$ alternatives are dominated by x . We claim that taking x out of the Banks set requires at least $\lfloor n/2 \rfloor = \ell$ edge reversals. Consider the $2\ell - 1$ sets

$$\{y_1, y_\ell\}, \{y_2, y_{\ell+1}\}, \dots, \{y_\ell, y_{2\ell-1}\}, \\ \{y_{\ell+1}, y_1\}, \dots, \{y_{2\ell-1}, y_\ell\}.$$

Note that each y_i is contained in exactly two of these sets. For each set, we say that it is ‘good’ if x is the only alternative that dominates both of the alternatives in the set, and ‘bad’ otherwise. Note that the existence of a good set implies that x is a Banks winner, as the transitive subtournament consisting of the good set and x cannot be extended. Initially, all $2\ell - 1$ sets are good. A reversal involving x and y_i can turn at most two good sets into bad sets (i.e., the two sets containing y_i). Similarly, a reversal involving y_i and y_j , where y_j dominates y_i after the reversal, can make at most

two good sets bad (i.e., the two sets containing y_i). So after at most $\ell - 1$ reversals, at least one set is still good. This implies that there is no DRS of size at most $\ell - 1$. Hence, $\text{MoV}_{BA}(x, T) \geq \ell = \lfloor n/2 \rfloor$.

The case where n is odd can be handled similarly. Let $n = 2\ell - 1$. Construct a tournament with alternatives $x, y_1, \dots, y_{2\ell-1}$ as before, and remove $y_{2\ell-1}$. We claim that taking x out of the Banks set in this tournament requires at least $\lfloor n/2 \rfloor = \ell - 1$ edge reversals. Consider $2\ell - 3$ sets, starting with the $2\ell - 1$ sets above and removing the two sets that contain $y_{2\ell-1}$. Each y_i is contained in at most two of these sets. The previous argument can be applied to show that $\text{MoV}_{BA}(x, T) \geq \ell - 1 = \lfloor n/2 \rfloor$.

To conclude the proof, we show that the same tournaments as constructed above also imply the tightness of the bound $\lfloor n/2 \rfloor$ for CO . In order to make x a non-winner, we must reverse edges so that another alternative y has a larger dominion than x . If n is even, then initially x dominates $n - 1$ alternatives while y dominates $n/2 - 1$ alternatives, so $|D(x)| - |D(y)| = (n - 1) - (n/2 - 1) = n/2$. Each edge reversal decreases this difference by at most 1, except for the reversal of the edge (x, y) , which reduces the difference by 2. Hence, in order to make the difference negative, we need at least $n/2$ reversals. A similar argument applies for the case where n is odd, since we have $|D(y)| \leq (n - 1)/2$, and therefore $|D(x)| - |D(y)| \geq (n - 1)/2 = \lfloor n/2 \rfloor$. \square

6.2 Lower Bounds for Non-Winners

Next, we turn our attention to non-winners. For TC and k -kings with $k \geq 3$, it is clear that reversing one edge suffices to make any alternative a winner. Indeed, we can simply reverse the edge between x and an arbitrary alternative in the uncovered set of T_{-x} . This ensures that x can reach every other alternative via a directed path of length at most three.

Theorem 15. *Let $S \in \{TC, k\text{-kings}\}$, where $k \geq 3$ is arbitrary. For any tournament T and any $x \in V(T) \setminus S(T)$, we have $\text{MoV}_S(x, T) = -1$.*

For CO , as many as $n - 2$ edge reversals may be required.

Theorem 16. *For any tournament T and any $x \in V(T) \setminus CO(T)$, we have $\text{MoV}_{CO}(x, T) \geq -(n - 2)$. Moreover, this bound is tight.*

Proof. With a budget of $n - 2$ reversals, we can make x dominate at least $n - 2$ alternatives. Moreover, if the tournament initially contains a Condorcet winner, one of these reversals can be used to make x dominate it, meaning that every alternative dominates at most $n - 2$ alternatives after the reversals. Hence x becomes a Copeland winner.

To show tightness, consider a tournament where x is a Condorcet loser and there is a Condorcet winner y . We have $|D(y)| - |D(x)| = n - 1$. Each edge reversal reduces this difference by at most 1, except for the reversal of the edge (x, y) , which reduces the difference by 2. In order for x to be a Copeland winner, this difference must be nonpositive. It follows that we need at least $n - 2$ reversals, as claimed. \square

Finally, we show that for UC and BA , reversing $O(\log n)$ edges can bring any alternative into the winner set.

Theorem 17. *Let $S \in \{UC, BA\}$. For any tournament T and any $x \in V(T) \setminus S(T)$, we have $\text{MoV}_S(x, T) \geq -\lceil \log_2 n \rceil$. Moreover, this bound is asymptotically tight.*

Proof. Since $BA \subseteq UC$, it suffices to establish the bound for BA and the tightness for UC . We first prove the bound for BA , by considering a tournament T and iteratively constructing a CRS for an alternative $x \notin BA(T)$. Let T' be a transitive subtournament of T that initially contains only the alternative x , and let B be the set of alternatives that dominate all alternatives in T' . Let $\ell = |B|$, and let y be a Copeland winner of the tournament $T|_B$. Note that y dominates at least $\lceil (\ell - 1)/2 \rceil$ other alternatives in B as well as x . We reverse the edge between x and y , insert y into the transitive tournament T' at the position after x , and update the set B . Since y is added to T' , y and all alternatives dominated by y are no longer in B . Also, no new alternative is added into B . Hence the size of B reduces to at most $\ell - 1 - \lceil (\ell - 1)/2 \rceil = \lfloor (\ell - 1)/2 \rfloor$. Since $|B| \leq n - 1$ at the beginning, the size of B becomes 0 after at most $\lceil \log_2 n \rceil$ reversals, at which point $x \in BA(T)$.

To show the asymptotic tightness for UC , assume that x is a Condorcet loser and T_{-x} is a tournament for which any dominating set has size $\Omega(\log n)$; such a tournament is known to exist (Erdős 1963; Graham and Spencer 1971). Let $R \subseteq E(T)$ be a CRS for x with respect to UC . Observe that if there is an edge $(y, z) \in R$ such that $x \notin \{y, z\}$, then by replacing (y, z) with (y, x) (or simply removing (y, z) if (y, x) already belongs to R), the resulting set R' is still a CRS for x . Moreover, $|R'| \leq |R|$. Therefore we may assume that all edges in R are incident to x ; let these edges be $(y_1, x), \dots, (y_{|R|}, x)$. Since $x \in UC(T^R)$, the set $\{y_1, \dots, y_{|R|}\}$ necessarily forms a dominating set in T_{-x} . It follows that $|R| \geq \Omega(\log n)$, as desired. \square

7 Discussion

In this paper, we have proposed a new framework for refining tournament solutions based on the notion of margin of victory (MoV). We have determined the complexity of computing the MoV, as well as worst-case bounds on the MoV, for several common tournament solutions. Besides the tournament solutions that we have considered, it would be interesting to study the MoV with respect to other tournament solutions such as the bipartisan set, the minimal covering set, the tournament equilibrium set, and the Markov set.

Viewing the MoV as a robustness measure, one could aim to obtain more comprehensive information about the space of all (not necessarily minimum) reversal sets. For example, one may ask *how many* reversal sets of cost at most c exist for a given alternative. Investigating the complexity of computing these numbers is an appealing future direction; similar counting questions have been considered in the context of knockout tournaments (Aziz et al. 2018).

In particular, one could use the number of minimum reversal sets as a tie-breaker for alternatives with equal MoV. Indeed, note that for some tournaments, especially small ones, the MoV in the unweighted setting may not distinguish between all winners (or non-winners). An example is the tournament in Figure 1, where three of the four UC winners

have a MoV of 1. A natural way to differentiate between alternatives with the same MoV is to consider the number of minimal reversal sets for each of them—for the example above, c has two minimal reversal sets ($\{(c, f)\}, \{(f, d)\}$), d has four ($\{(d, c)\}, \{(d, b)\}, \{(c, f)\}, \{(b, e)\}$), and e has three ($\{(c, f)\}, \{(e, d)\}, \{(e, c)\}$).

Understanding the counting problem is also relevant for settings in which there is uncertainty regarding the pairwise comparisons (e.g., say that the direction of each edge is incorrect with some fixed probability $p < 0.5$). In such a scenario, the number of reversal sets of a given size can be used to compute the winning probabilities of alternatives.

Acknowledgments

This work was partially supported by the Deutsche Forschungsgemeinschaft (DFG) under grant BR 4744/2-1 and by the European Research Council (ERC) under grant 639945 (ACCORD). We would like to thank Felix Brandt and Piotr Faliszewski for helpful discussions and the anonymous reviewers for valuable feedback.

References

- Adámek, J., and Koubek, V. 1971. Remarks on flows in network with short paths. *Commentationes Mathematicae Universitatis Carolinae* 12(4):661–667.
- Aziz, H.; Brill, M.; Fischer, F.; Harrenstein, P.; Lang, J.; and Seedig, H. G. 2015. Possible and necessary winners of partial tournaments. *Journal of Artificial Intelligence Research* 54:493–534.
- Aziz, H.; Gaspers, S.; Mackenzie, S.; Mattei, N.; Stursberg, P.; and Walsh, T. 2018. Fixing balanced knockout and double elimination tournaments. *Artificial Intelligence* 262:1–14.
- Baier, G.; Erlebach, T.; Hall, A.; Köhler, E.; Kolman, P.; Pangrác, O.; Schilling, H.; and Skutella, M. 2010. Length-bounded cuts and flows. *ACM Transactions on Algorithms* 7(1):4:1–4:27.
- Brandt, F., and Seedig, H. G. 2016. On the discriminative power of tournament solutions. In *Selected Papers of the International Conference on Operations Research, OR2014*, Operations Research Proceedings. Springer. 53–58.
- Brandt, F.; Brill, M.; and Harrenstein, P. 2016. Tournament solutions. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press. 57–84.
- Brandt, F.; Brill, M.; and Harrenstein, P. 2018. Extending tournament solutions. *Social Choice and Welfare* 51(2):193–222.
- Brill, M.; Schmidt-Kraepelin, U.; and Suksompong, W. 2019. Refining tournament solutions via margin of victory. Technical report, arXiv:1911.06289.
- Cary, D. 2011. Estimating the margin of victory for instant-runoff voting. In *Proc. of 2011 EVT/WOTE Conference*, 3:1–3:16.
- Dey, P., and Narahari, Y. 2015. Estimating the margin of victory of an election using sampling. In *Proc. of 24th IJCAI*, 1120–1126. AAAI Press.
- Erdős, P. 1963. On a problem in graph theory. *The Mathematical Gazette* 47(361):220–223.
- Faliszewski, P., and Rothe, J. 2016. Control and bribery in voting. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*. Cambridge University Press. 146–168.
- Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 2009. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research* 35:275–341.
- Fey, M. 2008. Choosing from a large tournament. *Social Choice and Welfare* 31(2):301–309.
- Fisher, D. C., and Ryan, J. 1995. Tournament games and positive tournaments. *Journal of Graph Theory* 19(2):217–236.
- Ford, L. R., and Fulkerson, D. R. 1956. Maximal flow through a network. *Canadian Journal of Mathematics* 8:399–404.
- Graham, R. L., and Spencer, J. H. 1971. A constructive solution to a tournament problem. *Canadian Mathematical Bulletin* 14(1):45–48.
- Kim, M. P., and Vassilevska Williams, V. 2015. Fixing tournaments for kings, chokers, and more. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Kim, M. P.; Suksompong, W.; and Vassilevska Williams, V. 2017. Who can win a single-elimination tournament? *SIAM Journal on Discrete Mathematics* 31(3):1751–1764.
- Konicki, C., and Vassilevska Williams, V. 2019. Bribery in balanced knockout tournaments. In *Proc. of 18th AAMAS Conference*, 2066–2068. IFAAMAS.
- Kruger, J., and Airiau, S. 2017. Refinements and randomised versions of some tournament solutions. In *Proc. of 16th AAMAS Conference*, 1584–1586. IFAAMAS.
- Laslier, J.-F. 1997. *Tournament Solutions and Majority Voting*. Springer-Verlag.
- Magrino, T. R.; Rivest, R. L.; Shen, E.; and Wagner, D. 2011. Computing the margin of victory in IRV elections. In *Proc. of 2011 EVT/WOTE Conference*, 4:1–4:16.
- Mahjoub, A. R., and McCormick, S. T. 2010. Max flow and min cut with bounded-length paths: complexity, algorithms, and approximation. *Mathematical Programming* 124(1–2):271–284.
- Mattei, N.; Goldsmith, J.; Klapper, A.; and Mundhenk, M. 2015. On the complexity of bribery and manipulation in tournaments with uncertain information. *Journal of Applied Logic* 13(4):557–581.
- Megiddo, N., and Vishkin, U. 1988. On finding a minimum dominating set in a tournament. *Theoretical Computer Science* 61(2–3):307–316.
- Mnich, M.; Shrestha, Y. R.; and Yang, Y. 2015. When does Schwartz conjecture hold? In *Proc. of 24th IJCAI*, 603–609. AAAI Press.
- Russell, T., and Walsh, T. 2009. Manipulating tournaments in cup and round robin competitions. In *Proc. of 1st ADT*, 26–37.
- Saile, C., and Suksompong, W. 2018. Robust bounds on choosing from large tournaments. In *Proc. of 14th WINE*, 393–407. Springer.
- Scott, A., and Fey, M. 2012. The minimal covering set in large tournaments. *Social Choice and Welfare* 38(1):1–9.
- Shiryayev, D.; Yu, L.; and Elkind, E. 2013. On elections with robust winners. In *Proc. of 12th AAMAS Conference*, 415–422.
- Vassilevska Williams, V. 2010. Fixing a tournament. In *Proc. of 24th AAAI Conference*, 895–900.
- Vu, T.; Altman, A.; and Shoham, Y. 2009. On the complexity of schedule control problems for knockout tournaments. In *Proc. of 8th AAMAS Conference*, 225–232.
- Woeginger, G. J. 2003. Banks winners in tournaments are difficult to recognize. *Social Choice and Welfare* 20(3):523–528.
- Xia, L. 2012. Computing the margin of victory for various voting rules. In *Proc. of 13th ACM-EC Conference*, 982–999. ACM.