

Query Complexity of Tournament Solutions

Palash Dey

Indian Institute of Science, Bangalore

Abstract

A directed graph where there is exactly one edge between every pair of vertices is called a *tournament*. Finding the “best” set of vertices of a tournament is a well studied problem in social choice theory. A *tournament solution* takes a tournament as input and outputs a subset of vertices of the input tournament. However, in many applications, for example, choosing the best set of drugs from a given set of drugs, the edges of the tournament are given only implicitly and knowing the orientation of an edge is costly. In such scenarios, we would like to know the best set of vertices (according to some tournament solution) by “querying” as few edges as possible. We, in this paper, precisely study this problem for commonly used tournament solutions: given an oracle access to the edges of a tournament \mathcal{T} , find $f(\mathcal{T})$ by querying as few edges as possible, for a tournament solution f . We first show that the set of Condorcet non-losers in a tournament can be found by querying $2n - \lfloor \log n \rfloor - 2$ edges only and this is tight in the sense that every algorithm for finding the set of Condorcet non-losers needs to query at least $2n - \lfloor \log n \rfloor - 2$ edges in the worst case, where n is the number of vertices in the input tournament. We then move on to study other popular tournament solutions and show that any algorithm for finding the Copeland set, the Slater set, the Markov set, the bipartisan set, the uncovered set, the Banks set, and the top cycle must query $\Omega(n^2)$ edges in the worst case. On the positive side, we are able to circumvent our strong query complexity lower bound results by proving that, if the size of the top cycle of the input tournament is at most k , then we can find all the tournament solutions mentioned above by querying $\mathcal{O}(nk + n \log n / \log(1 - 1/k))$ edges only.

Introduction

Many scenarios in multiagent systems can often be modeled and analyzed using tournaments (Mou86; BBFH14). An important example of such scenarios is *voting* where we have a set of alternatives and a set of votes which are linear orders over the set of alternatives. A important tournament to consider in this context is defined by the *majority relation* induced by the set of votes. In the majority relation, an alternative x is preferred over another alternative y if x is preferred over y in a majority of the votes. Indeed, many

important voting rules, like Copeland for example, are defined using the tournament induced by the majority relation of the input set of votes. Other than voting, tournaments have found many applications in multi-criteria decision analysis (AR⁺86; BMP⁺06), zero-sum games (FR95; DLB96), coalitional games (BH10), argumentation theory (Dun95; Dun07), biology (CH07), etc.

Formally, a *tournament* is defined as a set of alternatives along with an *irreflexive, antisymmetric, and total relation*, called *dominance relation*, on the set alternatives. An equivalent and often more convenient view of a tournament is as a *directed* graph on the alternatives where, between every pair of vertices (which corresponds to the alternatives), there is exactly one edge. A tournament solution takes a tournament as input and outputs a subset of the vertices. Tournament and tournament solutions are important mathematical tools in any general situation where we have to make a choice from a set of alternatives solely based on pairwise comparisons.

Motivation. We often have situations where the input tournament is given “implicitly” – the vertices of the tournament are given explicitly and we have to query for an edge to know its orientation. Moreover, knowing the orientation of an edge of the tournament can often be costly. For example, we can think of an application where we have a set of drugs for a particular disease and we want to know the “best” (according to some tournament solution) set of drugs. A natural dominance relation in this context would be to define a drug x to dominate another drug y if the probability that the drug x cures the disease is more than the corresponding probability for the drug y . Since these probabilities are often not known a priori, estimating them often requires extensive lab experiments as well as clinical trials. Hence, we would like to make as few queries as possible to know the best set of drugs. More generally, we can think of any *tournament based voting rules* like Copeland in an election scenario. A tournament based voting rule chooses winners solely based on the tournament induced by the pairwise majority relation between the alternatives. However, in many applications of voting in multiagent systems, recommender systems (PHG00) for example, the number of voters is huge and consequently, learning the majority relation is costly. Hence, we would like to learn the set of most popular items (according to the tournament solution under consideration)

with the smallest number of queries possible. Motivated by these applications, we study, for a tournament solution f , the problem of finding $f(\mathcal{T})$ of an input tournament \mathcal{T} by querying smallest number of edges possible. A *query* for an edge, in our model, reveals the orientation of the edge in the input tournament.

Finding the query complexity of various graph properties has drawn significant attention in literature. In the most general setting, the input is a directed graph on n vertices and one has to find whether the input graph satisfies some property we are concerned with, by asking a minimum number of queries. A query is a question of the form: “Is there an edge from a vertex x to another vertex y ?” The *query complexity* of a property is the worst case number of queries that must be made to know whether the input graph has that property. A graph property, in this context, is called *evasive* if its query complexity is $n(n-1)$, that is, one has to query all possible edges of the input digraph to test the property in the worst case. Karp conjectured that every monotone and nontrivial graph property is evasive. A graph property is *monotone* if it continues to hold even after adding more edges and *nontrivial* if it holds for some but not all graphs. A substantial amount of research effort has provided increasingly better query complexity lower bounds for monotone and nontrivial properties, although Karp’s conjecture remained open (KSS84; Kin88; RV76; Ros73; CKS01; Kul13; KT10). In the case of tournaments (where there is exactly one edge between every pair of vertices), Balasubramanian et al. (BRS97) showed (rediscovered by Procaccia (Pro08)) that a Condorcet winner – a vertex with $n-1$ outgoing edges – can be found, if it exists, with $2n - \lfloor \log n \rfloor - 2$ queries and this query complexity upper bound is tight in the worst case. This further motivates us to study the query complexity of other commonly used tournament solutions.

Our Contribution. In this paper, we prove tight bounds on the query complexity of commonly used tournament solutions. Our specific contributions in this paper are as follows.

- We show that the query complexity of the problem of finding the set of Condorcet non-losers is $2n - \lfloor \log n \rfloor - 2$ [Observation 1].
- We show that any algorithm for finding the Copeland set, the Slater set, and the Markov set in a tournament has query complexity $\binom{n}{2}$ [Theorem 1]. We remark that Goyal et al. (GJR17) also discovered this result independently (and in parallel) around same time.
- We prove that any algorithm for finding the bipartisan set [Theorem 2], the uncovered set [Theorem 3], the Banks set [Theorem 4], and the top cycle [Theorem 5] has query complexity $\Omega(n^2)$.
- We complement the our strong query complexity lower bounds above by showing that, if the tournament \mathcal{T} has a top cycle of size at most k , then the Copeland set, the Slater set, the Markov set, the bipartisan set, the uncovered set, the Banks set, and the top cycle of a tournament \mathcal{T} can be found using $\mathcal{O}(nk + n \log n / \log(1-1/k))$ queries [Theorem 6 and 7].

Related Work. The work of Balasubramanian et

al. (BRS97) (rediscovered by Procaccia (Pro08)) is the closest predecessor of our work where he shows that the query complexity of Condorcet winner in tournaments is $2n - \lfloor \log n \rfloor - 2$. There have been several other works in the literature which study the query complexity of various problems in tournaments. For example, Noy and Naor use comparison-based sorting algorithms to find feedback sets and Hamiltonian paths in tournaments (BNN90). There have been works in computational social choice theory on communication complexity of different voting rules (CS02) and query complexity of preference elicitation in various domains (Con09; DM16a; DM16b; Dey16b). However, the querying model in the above works is completely different from ours and consequently, neither the results nor the techniques involved in these works are directly applicable to our work; a query in the works above asks: “Does a voter v prefer an alternative x over another alternative y ?”

Preliminaries

For a positive integer ℓ , we denote the set $\{1, 2, \dots, \ell\}$ by $[\ell]$. For a finite set \mathcal{X} and a positive integer ℓ , we denote the set of all subsets of \mathcal{X} of size ℓ by $\mathfrak{P}_\ell(\mathcal{X})$ and the set of all probability distributions on \mathcal{X} by $\Delta(\mathcal{X})$.

Tournaments. A tournament $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ is a directed graph on a set \mathcal{V} of n vertices such that, for any two vertices $u, v \in \mathcal{V}$, either $(u, v) \in \mathcal{E}$ or $(v, u) \in \mathcal{E}$ but not both. If not mentioned otherwise, we denote the number of voters by n . We call any subgraph of a tournament a *partial tournament*. We call a tournament *balanced* if the in-degree and out-degree of every vertex are the same. A tournament \mathcal{T} defines a relation $>_{\mathcal{T}}$ on the set of vertices \mathcal{V} : $u >_{\mathcal{T}} v$ if $(u, v) \in \mathcal{T}$. Alternatively, any irreflexive, antisymmetric, and total relation $>$ on a set \mathcal{V} defines a tournament $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ where $(u, v) \in \mathcal{E}$ if $u > v$. When there is no possibility of confusion, we drop \mathcal{T} from the subscript of $>$. We call the relation $>_{\mathcal{T}}$ associated with a tournament \mathcal{T} the *dominance relation* of \mathcal{T} . We say a vertex u *dominates* or *defeats* another vertex v if $u >_{\mathcal{T}} v$. Let us define the *dominion* $D(v)$ of a vertex v as $D(v) = \{u \in \mathcal{V} : v > u\}$ and $\overline{D(v)}$ is called the *dominators* of v . Given a tournament \mathcal{T} , let $\mathfrak{A}(\mathcal{T})$ be its adjacency matrix. We call the matrix $\mathfrak{G}(\mathcal{T}) = \mathfrak{A}(\mathcal{T}) - \mathfrak{A}(\mathcal{T})^t$ the *skew-adjacency matrix* of \mathcal{T} , where $\mathfrak{A}(\mathcal{T})^t$ is the transpose of $\mathfrak{A}(\mathcal{T})$. A vertex v is called the *Condorcet winner* of a tournament if the out-degree of v is $n-1$; alternatively, a vertex v is the Condorcet winner if $v > u$ for every $u \in \mathcal{V} \setminus \{v\}$. Given a set \mathcal{V} , we denote the set of all tournaments over \mathcal{V} by $\mathfrak{T}(\mathcal{V})$. A *tournament solution* $f : \cup_{|\mathcal{V}| \geq 0} \mathfrak{T}(\mathcal{V}) \rightarrow 2^{\mathcal{V}} \setminus \{\emptyset\}$ is a function that takes a tournament as input and selects a nonempty set of vertices as output. Examples of commonly used tournament solutions are as follows (MBC⁺16, Chapter 3).

Condorcet non-loser: The Condorcet non-loser set of a tournament is the set of all vertices which has at least one outgoing edge.

Copeland set: The Copeland set of a tournament is the set of vertices with maximum out-degree.

Slater set: Given a tournament $\mathcal{T} = (\mathcal{V}, >)$, let us denote the maximal element of \mathcal{V} according to a strict linear order

$>$ on \mathcal{V} by $\max(>)$. The Slater score of a strict linear order $>$ over \mathcal{V} with respect to tournament $\mathcal{T} = (\mathcal{V}, >)$ is $|\mathcal{V}| - |\mathcal{S}(\mathcal{T})|$ where $\mathcal{S}(\mathcal{T}) = \{(x, y) \in \mathcal{V} \times \mathcal{V} : x > y, x > y\}$. A strict linear order is a Slater order if it has maximum Slater score. Then the Slater set is defined as $\text{SL}(\mathcal{T}) = \{\max(>) : > \text{ is a Slater order for } \mathcal{T}\}$.

Markov set: Given a tournament $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, we define a Markov chain $\mathcal{M}(\mathcal{T})$ as follows. The states of the Markov chain $\mathcal{M}(\mathcal{T})$ are the vertices of \mathcal{T} and the transition probabilities are determined by the dominance relation: in every step, stay in the current state v with probability $|\mathcal{D}(v)|/|\mathcal{T}|-1$, and move to state u with probability $1/|\mathcal{T}|-1$ for all $u \in \mathcal{D}(v)$. The Markov set is the set of vertices that have maximum probability in the unique stationary distribution of $\mathcal{M}(\mathcal{T})$. Formally, the transition matrix of the Markov chain is defined as.

$$Q = \frac{1}{|\mathcal{T}|-1} (\mathcal{A}(\mathcal{T})^t + \text{diag}(\text{CO}))$$

where $\text{diag}(\text{CO})$ is the diagonal matrix of the Copeland scores. Markov set $\text{MA}(\mathcal{T})$ of a tournament \mathcal{T} is then given by $\text{MA}(\mathcal{T}) = \arg \max_{a \in \mathcal{V}} \{p(a) : p \in \Delta(\mathcal{V}), Qp = p\}$.

Bipartisan set: Bipartisan set generalizes the notion of Condorcet winner to lotteries over the vertices of the tournament. Interestingly, for every tournament $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, there exists a unique *maximal lottery* (MBC⁺16). That is, there exists a probability distribution $p \in \Delta(\mathcal{V})$ such that, for the skew-adjacency matrix $\mathcal{G}(\mathcal{T}) = (g_{ab})_{a,b \in \mathcal{V}}$ of \mathcal{T} , $\sum_{a,b \in \mathcal{V}} p(a)q(b)g_{ab} \geq 0$ for all $q \in \Delta(\mathcal{V})$ which is, by convexity, equivalent to the following condition.

$$\sum_{a \in \mathcal{V}} p(a)g_{ab} \geq 0 \text{ for all } b \in \mathcal{V} \quad (1)$$

Let $p_{\mathcal{T}}$ denote the unique maximal lottery of \mathcal{T} . Then the bipartisan set $\text{BP}(\mathcal{T})$ of \mathcal{T} is defined as the support of $p_{\mathcal{T}}$. That is,

$$\text{BP}(\mathcal{T}) = \{a \in \mathcal{V} : p_{\mathcal{T}}(a) > 0\}$$

Uncovered set: Given a tournament $\mathcal{T} = (\mathcal{V}, >)$, we say a vertex $v \in \mathcal{V}$ covers another vertex $u \in \mathcal{V}$ if $\mathcal{D}(u) \subseteq \mathcal{D}(v)$ and is denoted by $v \text{Cu}$. We observe that $v \text{Cu}$ implies $v > u$ and is equivalent to $\mathcal{D}(v) \subseteq \mathcal{D}(u)$. The uncovered set $\text{UC}(\mathcal{T})$ of a tournament \mathcal{T} is given by the set of maximal elements of the covering relation \mathcal{C} . That is,

$$\text{UC}(\mathcal{T}) = \{v \in \mathcal{V} : u \text{Cv for no } u \in \mathcal{V}\}$$

Banks set: A sub-tournament of a tournament $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ is an induced subgraph of \mathcal{T} . The Banks set $\text{BA}(\mathcal{T})$ is the set of maximal elements of all the maximal transitive sub-tournaments of \mathcal{T} .

Top cycle: A non-empty subset of vertices $\mathcal{B} \subseteq \mathcal{V}$ is called dominant in a tournament $\mathcal{T} = (\mathcal{V}, >)$ if $x > y$ for every $x \in \mathcal{B}$ and $y \in \mathcal{V} \setminus \mathcal{B}$. Dominant sets are linearly ordered via set inclusion and the top cycle returns the unique smallest dominant set.

A tournament solution is called *neutral* if the output does not depend on the names of the input set of vertices. All the tournament solutions discussed above are neutral.

Essential States in a Markov Chain. A state i in a finite Markov chain is called *essential* if for every state j that is

reachable from i , the state i is also reachable from j . A state is called inessential if it is not essential. A well known fact from probability theory is that, $\pi(i) = 0$ for every inessential state i , where π is a stationary distribution of the Markov chain. Hence, every vertex whose corresponding state in the Markov chain is inessential, does not belong to the Markov set of that tournament. We refer the reader to (Bré13) for a more detailed discussion on Markov chains.

Query Model. Given a tournament $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ on n vertices, a query for a pair of vertices $\{x, y\} \in \mathcal{P}_2\{\mathcal{V}\}$ reveals whether $(x, y) \in \mathcal{E}$ or $(y, x) \in \mathcal{E}$. The *query complexity of an algorithm* is the maximum number of queries the algorithm makes in the worst case. The *query complexity of a tournament solution* f is the minimum query complexity of any algorithm for computing f .

Query Complexity Lower Bounds of Tournament Solutions

We begin with the following observation which gives us the query complexity of the Condorcet non-loser set of tournaments.

Observation 1 *The query complexity of the Condorcet non-loser set in tournaments is equal to the query complexity of the Condorcet winner in tournaments. Hence, the query complexity of the Condorcet non-loser set in tournaments is $2n - \lfloor \log n \rfloor - 2$.*

Proof: Given a tournament $\mathcal{T} = (\mathcal{V}, \mathcal{E})$, let us define another tournament $\overline{\mathcal{T}} = (\mathcal{V}, \overline{\mathcal{E}})$, where $\overline{\mathcal{E}} = \{(x, y) : (y, x) \in \mathcal{E}\}$. Now the result follows from the observation that a vertex v is a Condorcet non-loser in \mathcal{T} if and only if v is not the Condorcet winner in $\overline{\mathcal{T}}$. Now the $2n - \lfloor \log n \rfloor - 2$ query complexity of the Condorcet non-loser set follows from the $2n - \lfloor \log n \rfloor - 2$ query complexity of the Condorcet winner in tournaments (BRS97). \square

We next consider the query complexity of the Slater set of tournaments. The following result provides a necessary condition for a vertex to belong to the Slater set of a tournament. We will use it to prove query complexity lower bound of the Slater set. In the interest of space, we defer the proof of Lemma 1 to the full version of the paper (Dey16a).

Lemma 1 *Suppose the out-degree of a vertex $v \in \mathcal{V}$ in a tournament $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ on n vertices is strictly less than $(n-1)/2$. Then v does not belong to the Slater set of \mathcal{T} .*

Let us consider the tournament $\mathcal{T}_{\text{reg}} = (\mathcal{V}, \mathcal{E})$ on n vertices $\mathcal{V} = \{a_i : i \in [n]\}$. We assume n to be an odd integer. In \mathcal{T}_{reg} , vertex a_i defeats $a_{i+j \pmod n}$ for every $i \in [n]$ and $j \in [(n-1)/2]$. We will use the tournament \mathcal{T}_{reg} crucially in our proofs of the query complexity lower bounds of the Copeland set, the Slater set, and the Markov set. The following result is immediate from the definition of neutral tournament solutions.

Lemma 2 *Given the tournament \mathcal{T}_{reg} as input, every neutral tournament solution outputs the set of all vertices in \mathcal{T}_{reg} .*

Suppose there exists an edge from a vertex u to another vertex v in \mathcal{T}_{reg} . Let $\mathcal{T}_{\text{reg}}^{u,v}$ be the tournament which is the same as \mathcal{T}_{reg} except the edge from u to v is reversed, that is, $\mathcal{T}_{\text{reg}}^{u,v} = \mathcal{T}_{\text{reg}} \cup \{(v, u)\} \setminus \{(u, v)\}$. The following lemma will be used crucially in our proofs of the query complexity lower bounds of the Copeland set, the Slater set, and the Markov set.

Lemma 3 *The Copeland set, the Slater set, and the Markov set of $\mathcal{T}_{\text{reg}}^{u,v}$ do not contain u .*

Proof: Copeland set: For the Copeland set, the result follows from the observation that u is no longer a vertex with highest out-degree in $\mathcal{T}_{\text{reg}}^{u,v}$.

Slater set: For the Slater set, the result follows immediately from Lemma 1 since the out-degree of u is strictly smaller than $(n-1)/2$ in $\mathcal{T}_{\text{reg}}^{u,v}$.

Markov set: If possible, let us assume that the stationary distribution of the Markov chain $\mathfrak{M}(\mathcal{T}_{\text{reg}}^{u,v})$ associated with the tournament $\mathcal{T}_{\text{reg}}^{u,v}$ be π such that $\pi(a) = p$ for every $a \in [n] \setminus \{u, v\}$. Then we have the following.

$$p + \pi(u) \frac{2}{n+1} = \pi(v) \quad (2)$$

$$p \frac{n-3}{n-1} = \pi(u) \quad (3)$$

$$\pi(u) + (n-2)p + \pi(v) = 1 \quad (4)$$

We first observe that, since the Markov chain $\mathfrak{M}(\mathcal{T}_{\text{reg}}^{u,v})$ is ergodic, it has a unique stationary distribution. Now, since there exists a unique solution to the equations above, the stationary distribution π of $\mathfrak{M}(\mathcal{T}_{\text{reg}}^{u,v})$ is indeed of the form we assumed (that is, $\pi(a) = p$ for every $a \in [n] \setminus \{u, v\}$). We observe that equation 3 shows that $\pi(u) < p$ and equation 2 shows that $p < \pi(v)$. Hence, the Markov set of $\mathcal{T}_{\text{reg}}^{u,v}$ is $\{v\}$. \square

We now prove that any algorithm for finding the Copeland set, the Slater set, and the Markov set must query every edge in the input tournament in the worst case.

Theorem 1 *Any algorithm for finding the Copeland set, the Slater set, and the Markov set of tournaments has query complexity $\binom{n}{2}$.*

Proof: Let us consider the tournament \mathcal{T}_{reg} . We observe that, due to Lemma 2, the Copeland set, the Slater set, and the Markov set of \mathcal{T}_{reg} is the set of all vertices. Let us now consider the oracle which answers the queries according to \mathcal{T}_{reg} . We claim that any algorithm for finding the Copeland set, the Slater set, and the Markov set of tournaments must query all the $\binom{n}{2}$ edges of \mathcal{T}_{reg} . Suppose not, then there exists an edge from u to v in \mathcal{T}_{reg} for some vertices u and v that the algorithm does not query. Let $\hat{\mathcal{T}}$ be the partial tournament of \mathcal{T}_{reg} containing the edges that the algorithm queries. If the output of the algorithm does not contain u , then the oracle completes $\hat{\mathcal{T}}$ to \mathcal{T}_{reg} and thus the algorithm does not output correctly since the output should contain all the vertices. On the other hand, if the output of the algorithm contains u then the oracle completes $\hat{\mathcal{T}}$ by directing the edge between u and v from v to u and directing the rest of

the edges as in \mathcal{T}_{reg} . Again the algorithm outputs wrongly due to Lemma 3. \square

We now present our query complexity lower bound for the bipartisan set of tournaments. Before embarking on the query complexity lower bound of the bipartisan set, let us prove a few structural results for the bipartisan set which we will use crucially later. The following result for the bipartisan set is well known (GLM93; FR95).

Lemma 4 *In a tournament $\mathcal{T} = (\mathcal{V}, >)$, suppose a vertex u covers another vertex v , that is, $u > w$ whenever $v > w$ for every $w \in \mathcal{V}$. Then v does not belong to the bipartisan set of \mathcal{T} .*

The following result shows that, in the tournaments where every vertex has the same number of incoming edges as the number of outgoing edges, the bipartisan set is the set of all vertices.

Lemma 5 *Let \mathcal{T} be a tournament on n vertices where the in-degree and out-degree of every vertex is $(n-1)/2$. Then the only maximal lottery of \mathcal{T} is the uniform distribution over the set of all vertices of \mathcal{T} and thus the bipartisan set of \mathcal{T} is the set of all vertices.*

Proof: We observe that the uniform distribution over all the vertices satisfy Equation (1) for the tournament \mathcal{T} , since the average of the entries in every column of the skew-symmetric matrix \mathfrak{G} of the tournament \mathcal{T} is 0. Now the result follows from the fact that the maximal lottery in every tournament is unique (MBC⁺16). \square

In the next lemma, we formalize the intuition that, if a $(\mathcal{A}, \mathcal{V} \setminus \mathcal{A})$ cut in a tournament has a majority of its edges from $\mathcal{V} \setminus \mathcal{A}$ to \mathcal{A} and \mathcal{A} is balanced, then the bipartisan set of the tournament must include at least one vertex from $\mathcal{V} \setminus \mathcal{A}$.

Lemma 6 *Let $\mathcal{T} = (\mathcal{V}_1 \cup \mathcal{V}_2, \mathcal{E})$ be a tournament such that there exist at most $|\mathcal{V}_1||\mathcal{V}_2|/2 - 1$ edges from \mathcal{V}_1 to \mathcal{V}_2 and the in-degrees and out-degrees of all the vertices in the sub-tournament $\mathcal{T}[\mathcal{V}_1]$ of \mathcal{T} induced on \mathcal{V}_1 are exactly $|\mathcal{V}_1|-1/2$. Then the bipartisan set of \mathcal{T} must include at least one vertex from \mathcal{V}_2 .*

Proof: Let $p^* \in \Delta(\mathcal{V}_1 \cup \mathcal{V}_2)$ be the maximal lottery of \mathcal{T} and $\mathcal{V} = \mathcal{V}_1 \cup \mathcal{V}_2$. If possible, let us assume that $p^*(v) = 0$ for every $v \in \mathcal{V}_2$. Let $q \in \Delta(\mathcal{V}_2)$ be the uniform distribution on \mathcal{V}_2 and $\mathfrak{G} = (g_{ab})_{a,b \in \mathcal{V}}$ the skew-adjacency matrix of \mathcal{T} . We first claim that p^* cannot be the uniform distribution on \mathcal{V}_1 . Indeed, otherwise we have $\sum_{a,b \in \mathcal{V}} p^*(a)q(b)g_{ab} < 0$ (since a strict majority of the edges between \mathcal{V}_1 and \mathcal{V}_2 are from \mathcal{V}_2 to \mathcal{V}_1) which contradicts the fact that p^* is the maximal lottery of \mathcal{T} . We now consider the sub-tournament $\mathcal{T}[\mathcal{V}_1]$ of \mathcal{T} induced on \mathcal{V}_1 . Due to Lemma 5, the only maximal lottery of $\mathcal{T}[\mathcal{V}_1]$ is the uniform distribution on \mathcal{V}_1 . Hence, since p^* is not the uniform distribution over \mathcal{V}_1 , there exists a distribution $q' \in \Delta(\mathcal{V}_1)$ such that $\sum_{a,b \in \mathcal{V}_1} p^*(a)q'(b)g_{ab} < 0$. However, this contradicts our assumption that p^* is the maximal lottery of \mathcal{T} . Hence, the bipartisan set of \mathcal{T} must include at least one vertex from \mathcal{V}_2 . \square

We now have all the ingredients to present our query complexity lower bound result for the bipartisan set.

Theorem 2 *The query complexity of the bipartisan set of a tournament is $\Omega(n^2)$.*

Proof: Let n be an odd integer. We consider a partial tournament $\mathcal{T} = (\mathcal{A} \cup \mathcal{B}, \mathcal{E})$ where $\mathcal{A} = \{a_i : i \in [n]\}$, $\mathcal{B} = \{b_i : i \in [n]\}$, and $\mathcal{E} = \{(a_i, a_{i+j \pmod n}), (b_i, b_{i+j \pmod n}) : i \in [n], 1 \leq j \leq (n-1)/2\}$. The oracle answers the queries of the algorithm as follows. If a query comes for an edge between vertices a_i and a_j or b_i and b_j for any $i, j \in [n]$, then the oracle answers according to \mathcal{T} . If a query comes for an edge between a_i and b_j for any $i, j \in [n]$, then the oracle says that the edge is oriented from a_i to b_j . We now claim that the algorithm must query at least $n^2/2$ edges between \mathcal{A} and \mathcal{B} . Suppose not, then, if the output of the algorithm contains any vertex from \mathcal{B} , then the oracle orients every edge between \mathcal{A} and \mathcal{B} , from \mathcal{A} to \mathcal{B} , thereby ensuring that the output of the algorithm is wrong due to Lemma 4. On the other hand, if the output of the algorithm does not contain any vertex from \mathcal{B} , then the oracle orients all the edges between \mathcal{A} and \mathcal{B} that are not queried, from \mathcal{B} to \mathcal{A} , thereby ensuring that the output of the algorithm is again wrong due to Lemma 6. Hence, the algorithm must query at least $n^2/2$ edges between \mathcal{A} and \mathcal{B} and thus the query complexity of the bipartisan set is $\Omega(n^2)$. \square

We now show that the query complexity of the uncovered set is $\Omega(n^2)$.

Theorem 3 *The query complexity of the uncovered set of a tournament is $\Omega(n^2)$.*

Proof: Consider a partial tournament $\mathcal{T} = (\mathcal{A} \cup \mathcal{B} \cup \{x\}, \mathcal{E})$ where $\mathcal{A} = \{a_i : i \in [n]\}$, $\mathcal{B} = \{b_i : i \in [n]\}$ and $\mathcal{E} = \{(a_i, x), (x, b_i) : i \in [n]\}$. Let us consider the following oracle. Let us define the partial tournament \mathcal{T}' to be the graph on $\mathcal{A} \cup \mathcal{B} \cup \{x\}$ consisting of all the answers of the oracle till now. Hence, to begin with, \mathcal{T}' does not have any edge. The oracle answers the queries for any edge in \mathcal{T} according to \mathcal{T} . For any query of the form $\{a_i, a_j\}$ or $\{b_i, b_j\}$, the oracle answers arbitrarily but consistently. For a query $\{a_i, b_j\}$ for some $i, j \in [n]$, if a_i defeats b_k for every $k \in [n] \setminus \{j\}$ in \mathcal{T}' , then the oracle answers that b_j defeats a_i ; otherwise oracle answers that a_i defeats b_j . We claim that any algorithm for finding the uncovered set of tournaments must query for the pair $\{a_i, b_j\}$ for every $i, j \in [n]$. Suppose not, then there exists a pair $\{a_i, b_j\}$ which the algorithm does not query. Notice that, by the design of the oracle, a_i defeats b_k in \mathcal{T}' for every $k \in [n]$ such that $\{a_i, b_k\}$ has been queried by the algorithm. For every pair $\{a_k, b_\ell\}$ with $k, \ell \in [n]$, $k \neq i$ and $\{a_k, b_\ell\}$ has not been queried, the oracle orients the edge from b_ℓ to a_k . The oracle also orients all the edges from a_i to b_k for every $k \in [n] \setminus \{j\}$. Now, if the output of the algorithm contains x , then the oracle orients the edge $\{a_i, b_j\}$ from a_i to b_j . Notice that, in this case, x is covered by a_i and thus x should not be in the uncovered set and hence the output of the algorithm is wrong. On the other hand, if the output of the algorithm does not contain x , then the oracle orients the

edge $\{a_i, b_j\}$ from b_j to a_i . In this case, x is not covered by any other vertex and thus x belongs to the uncovered set. Hence, the algorithm outputs incorrectly in both the cases thereby proving the result. \square

Next we consider the Banks set and show its query complexity to be $\Omega(n^2)$.

Theorem 4 *The query complexity of the Banks set of a tournament is $\Omega(n^2)$.*

Proof: Consider a partial tournament $\mathcal{T} = (\mathcal{A} \cup \mathcal{B} \cup \{x\}, \mathcal{E})$, where $\mathcal{A} = \{a_i : i \in [n]\}$, $\mathcal{B} = \{b_i : i \in [n]\}$, and $\mathcal{E} = \{(a_i, x), (x, b_i), (b_i, b_j) : i, j \in [n], i > j\}$. Let us consider the following oracle. Let us define the partial tournament \mathcal{T}' to be the graph on $\mathcal{A} \cup \mathcal{B} \cup \{x\}$ consisting of all the answers of the oracle till now. Hence, to begin with, \mathcal{T}' does not have any edge. The oracle answers the queries for any edge in \mathcal{T} according to \mathcal{T} . For any query of the form $\{a_i, a_j\}$ or $\{b_i, b_j\}$, the oracle answers arbitrarily but consistently. For a query $\{a_i, b_j\}$ for some $i, j \in [n]$, if a_i defeats b_k for every $k \in [n] \setminus \{j\}$ in \mathcal{T}' , then the oracle answers that b_j defeats a_i ; otherwise oracle answers that a_i defeats b_j . We claim that the algorithm must query for the pair $\{a_i, b_j\}$ for every $i, j \in [n]$. Suppose not, then there exists a pair $\{a_i, b_j\}$ which the algorithm does not query. Notice that, by the design of the oracle, a_i defeats b_k in \mathcal{T}' for every $k \in [n]$ such that $\{a_i, b_k\}$ has been queried by the algorithm. For every pair $\{a_k, b_\ell\}$ with $k, \ell \in [n]$, $k \neq i$ and $\{a_k, b_\ell\}$ has not been queried, the oracle orients the edge from a_k to b_ℓ . The oracle also orients all the edges not in \mathcal{T}' between a_i and b_k from a_i to b_k for every $k \in [n] \setminus \{j\}$. Now if the output of the algorithm contains x , then the oracle orients the edge between a_i and b_j from a_i to b_j . We claim that x can not be the maximum vertex of any maximal transitive sub-tournament \mathcal{T}'' of \mathcal{T} . To see this, we first observe that the sub-tournament \mathcal{T}'' must have all the vertices in \mathcal{B} and no vertex from \mathcal{A} . Indeed, otherwise either x is not the maximum vertex of \mathcal{T}'' (if any vertex from \mathcal{A} is there in \mathcal{T}'') or \mathcal{T}'' is not a maximal transitive sub-tournament (if any vertex from \mathcal{B} is not there in \mathcal{T}''). However, such a sub-tournament is not a maximal transitive sub-tournament since a_i can be added to \mathcal{T}'' without violating transitivity. Hence x does not belong to the Banks set of the resulting tournament and thus the algorithm's output is wrong. On the other hand, suppose the output of the algorithm does not contain x . Then the oracle orients the edge between a_i and b_j from b_j to a_i . In this case, the sub-tournament of \mathcal{T} induced on $\mathcal{B} \cup \{x\}$ is a maximal sub-tournament where x is the maximum vertex and thus x belongs to the Banks set of the resulting tournament. Hence, the algorithm outputs incorrectly in both the cases thereby proving the result. \square

We now show that the query complexity of the top cycle of tournaments is $\Omega(n^2)$.

Theorem 5 *The query complexity of the top cycle of a tournament is $\Omega(n^2)$.*

Proof: We consider a partial tournament $\mathcal{T} = (\mathcal{A} \cup \mathcal{B}, \mathcal{E})$ where $\mathcal{A} = \{a_i : i \in [n]\}$, $\mathcal{B} = \{b_i : i \in [n]\}$, and

$\mathcal{E} = \{(a_i, a_{i+1 \pmod n}), (b_i, b_{i+1 \pmod n}) : i \in [n]\}$. The oracle answers the queries of the algorithm as follows. If a query comes for the edge between vertices a_i and a_j or b_i and b_j for any $i, j \in [n]$, then the oracle answers according to \mathcal{T} if the edge is present in \mathcal{T} , and arbitrarily but consistently otherwise. If a query comes for an edge between a_i and b_j for any $i, j \in [n]$, then the oracle says that the edge is oriented from a_i to b_j . Now we claim that the algorithm must query all the n^2 edges between \mathcal{A} and \mathcal{B} . Suppose not, then there exist a_i and b_j for some $i, j \in [n]$ such that the algorithm has not queried for the edge between a_i and b_j . Now if the output of the algorithm does not contain any vertex from \mathcal{B} , then the oracle orients the edge between a_i and b_j from b_j to a_i . Notice that, in this case the top cycle of the resulting tournament \mathcal{T} contains at least one vertex $b_j \in \mathcal{B}$ and thus the algorithm does not output correctly in this case. On the other hand, if the output of the algorithm contains any vertex from \mathcal{B} , then the oracle orients all the edges between \mathcal{A} and \mathcal{B} from \mathcal{A} to \mathcal{B} . In this case, the top cycle of the resulting tournament is \mathcal{A} and thus the algorithm again fails to output correctly. Hence, the algorithm must make $\Omega(n^2)$ queries. \square

Results for Tournaments with Small Top Cycle

It turns out that, if we a priori know that the size of the top cycle in the input tournament is at most k , then there is an algorithm for finding the top cycle with much less number of queries.

Theorem 6 *Suppose we know that the top cycle of the input tournament \mathcal{T} is of size at most k . Then there exists an algorithm for finding the top cycle of \mathcal{T} with query complexity $\mathcal{O}(nk + n \log n / \log(1 - 1/k))$.*

Proof: We first partition the set of vertices \mathcal{V} into $\lceil n/2k \rceil$ subsets $\mathcal{V}_i, i \in [n/2k]$, each of size at most $2k$. For each partition, we query all pair of vertices. We notice that, in each set \mathcal{V}_i of the partition, there must exist at least one vertex v_i with in-degree at least k (for every k larger than some small constant) and consequently v_i does not belong to the top cycle of \mathcal{T} for every $i \in [n/2k]$. We delete the vertex v_i from \mathcal{V}_i for every $i \in [n/2k]$, thereby deleting $\lceil n/2k \rceil$ vertices in total. The now iterate the same process on the remaining set of vertices. The first iteration takes $\mathcal{O}((n/k)k^2) = \mathcal{O}(nk)$ queries. From the next iteration onwards, we can manage with only n queries per iteration by partitioning the vertices cleverly: since we have deleted exactly one vertex from each set of the partition we can add one vertex to each set of the partition by “breaking” some of the sets from the partition. We now observe that, in every set, we now need to compare the new vertex with the rest of the vertices thereby requiring at most $n - k$ queries in total. Since, each iteration decreases the size of the tournament by a factor of $\Omega(1 - 1/k)$, after $\mathcal{O}(\log n / \log(1 - 1/k))$ iterations, we have $\mathcal{O}(k)$ vertices in the tournament where we can find the top cycle using $\mathcal{O}(k^2) = \mathcal{O}(nk)$ queries. Hence, the query complexity of our algorithm is $\mathcal{O}(nk + n \log n / \log(1 - 1/k))$. The correctness of the algorithm follows from the fact that whenever we remove a vertex v from the tournament, v does

not belong to the top cycle of \mathcal{T} . \square

The following result gives relationship between the top cycle of a tournament and other tournament solutions like the Copeland set, the Slater set, the Markov set, the bipartisan set, the uncovered set, and the Banks set.

Lemma 7 *Let \mathcal{T} be a tournament whose top cycle is \mathcal{C} . Then the Copeland set, the Slater set, the Markov set, the bipartisan set, the uncovered set, and the Banks set of \mathcal{T} are the same as the corresponding solutions for the tournament $\mathcal{T}(\mathcal{C})$.*

Proof: **Copeland set, bipartisan set, uncovered set:** Follows from the observation that every vertex in \mathcal{C} covers every vertex in $\mathcal{V} \setminus \mathcal{C}$ and Lemma 4.

Markov set: All the states corresponding to the vertices in $\mathcal{V} \setminus \mathcal{C}$ are inessential and thus do not belong to the Markov set of \mathcal{T} .

Slater set: We observe that, in the Slater order $>$ of the tournament \mathcal{T} , every vertex in \mathcal{C} must be preferred over every vertex in $\mathcal{V} \setminus \mathcal{C}$. If not, then let there be a vertex $a \in \mathcal{C}$ and $b \in \mathcal{V} \setminus \mathcal{C}$ such that a immediately follows b in $>$. Then by swapping the positions of the vertices a and b in $>$, we can strictly decrease the disagreement of $>$ with \mathcal{T} thereby contradicting that $>$ is a Slater order of \mathcal{T} .

Banks set: Follows from the fact that every maximal element of every maximal sub-tournament of \mathcal{T} belongs to \mathcal{C} . \square

Lemma 7 and Theorem 6 immediately give the the following query complexity upper bound for the Copeland set, the Slater set, the Markov set, the bipartisan set, the uncovered set, and the Banks set when we a priori know that the size of the top cycle of the input tournament is at most k .

Theorem 7 *Suppose we know that the top cycle of the input tournament \mathcal{T} is of size at most k . Then there exists an algorithm for finding the Copeland set, the Slater set, the Markov set, the bipartisan set, the uncovered set, and the Banks set of \mathcal{T} with query complexity $\mathcal{O}(nk + n \log n / \log(1 - 1/k))$.*

Proof: We first find the top cycle \mathcal{C} of \mathcal{T} using Theorem 6. This step requires $\mathcal{O}(nk + n \log n / \log(1 - 1/k))$ queries. Next, we query for all the pair of vertices in \mathcal{C} and output the corresponding solution of $\mathcal{T}(\mathcal{C})$. The correctness of the algorithm follows immediately from Lemma 7. Since the second step requires $\mathcal{O}(k^2) = \mathcal{O}(nk)$ queries, the query complexity of our algorithm is $\mathcal{O}(nk + n \log n / \log(1 - 1/k))$. \square

Conclusion and Future Directions

We have shown that, for finding many common tournament solutions, one has to query, in the worst case, almost entire set of edges in the tournament. On the positive side, we have exhibited an important structural property, in terms of the top cycle of the tournament being small, which substantially reduces the query complexity of common tournament solutions. An immediate future direction of research is to study query complexity for other tournament solutions like the minimal covering set, the minimal extending set, the minimal TC-retentive set, the tournament equilibrium set, etc.

References

- Kenneth J Arrow, Hervé Raynaud, et al. Social choice and multicriterion decision-making. *MIT Press Books*, 1, 1986.
- Felix Brandt, Markus Brill, Felix Fischer, and Paul Harrenstein. Minimal retentive sets in tournaments. *Soc. Choice Welf.*, 42(3):551–574, 2014.
- Felix Brandt and Paul Harrenstein. Characterization of dominance relations in finite coalitional games. *Theor. Decis.*, 69(2):233–256, 2010.
- Denis Bouyssou, Thierry Marchant, Marc Pirlot, Alexis Tsoukiàs, and Philippe Vincke. *Evaluation and decision models with multiple criteria: Stepping stones for the analyst*, volume 86. Springer Science & Business Media, 2006.
- Amotz Bar-Noy and Joseph Naor. Sorting, minimal feedback sets, and hamilton paths in tournaments. *SIAM J. Discrete Math.*, 3(1):7–20, 1990.
- Pierre Brémaud. *Markov chains: Gibbs fields, Monte Carlo simulation, and queues*, volume 31. Springer Science & Business Media, 2013.
- R. Balasubramanian, Venkatesh Raman, and G. Srinivasaragavan. Finding scores in tournaments. *J. Algorithms*, 24(2):380–394, 1997.
- Irène Charon and Olivier Hudry. A survey on the linear ordering problem for weighted or unweighted tournaments. *4OR*, 5(1):5–60, 2007.
- Amit Chakrabarti, Subhash Khot, and Yaoyun Shi. Evasiveness of subgraph containment and related properties. *SIAM J. Comput.*, 31(3):866–875, 2001.
- Vincent Conitzer. Eliciting single-peaked preferences using comparison queries. *J. Artif. Intell. Res.*, 35:161–191, 2009.
- Vincent Conitzer and Tuomas Sandholm. Vote elicitation: Complexity and strategy-proofness. In *Eighteenth National Conference on Artificial Intelligence (AAAI)*, pages 392–397, 2002.
- Palash Dey. Query complexity of tournament solutions. *CoRR*, abs/1611.06189, 2016.
- Palash Dey. Recognizing and eliciting weakly single crossing profiles on trees. *CoRR*, abs/1611.04175, 2016.
- John Duggan and Michel Le Breton. Dutta’s minimal covering set and shapley’s saddles. *J. Econ. Theory*, 70(1):257–265, 1996.
- Palash Dey and Neeldhara Misra. Elicitation for preferences single peaked on trees. In *Proc. Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 215–221, 2016.
- Palash Dey and Neeldhara Misra. Preference elicitation for single crossing domain. In *Proc. Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 222–228, 2016.
- Phan Minh Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artif Intel*, 77(2):321–357, 1995.
- Paul E Dunne. Computational properties of argument systems satisfying graph-theoretic constraints. *Artif Intel*, 171(10):701–729, 2007.
- David C Fisher and Jennifer Ryan. Tournament games and positive tournaments. *J. Graph Theory*, 19(2):217–236, 1995.
- Dishant Goyal, Varunkumar Jayapaul, and Venkatesh Raman. Elusiveness of finding degrees. *To appear in Proc. Third International Conference on Algorithms and Discrete Applied Mathematics (CALDAM), February 13 - 18, 2017, Goa, India.*, 2017.
- Laffond G., Jean-Francois Laslier, and Le Breton M. The bipartisan set of a tournament game. *GEB*, 5(1):182–201, 1993.
- Valerie King. Lower bounds on the complexity of graph properties. In *Proc. twentieth annual ACM symposium on Theory of computing*, pages 468–476. ACM, 1988.
- Jeff Kahn, Michael Saks, and Dean Sturtevant. A topological approach to evasiveness. *Combinatorica*, 4(4):297–306, 1984.
- Torsten Korneffel and Eberhard Triesch. An asymptotic bound for the complexity of monotone graph properties. *Combinatorica*, 30(6):735–743, 2010.
- Raghav Kulkarni. Evasiveness through a circuit lens. In *Proc. 4th Conference on Innovations in Theoretical Computer Science, ITCS ’13*, pages 139–144, New York, NY, USA, 2013. ACM.
- Hervé Moulin, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.
- H Moulin. Choosing from a tournament. *Soc. Choice Welf.*, 3:271–291, 1986.
- David M. Pennock, Eric Horvitz, and C. Lee Giles. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proc. Seventeenth AAAI, July 30 - August 3, 2000, Austin, Texas, USA.*, pages 729–734, 2000.
- Ariel D Procaccia. A note on the query complexity of the condorcet winner problem. *Inform. Process. Lett.*, 108(6):390–393, 2008.
- Arnold L Rosenberg. On the time required to recognize properties of graphs: A problem. *ACM SIGACT News*, 5(4):15–16, 1973.
- Ronald L Rivest and Jean Vuillemin. On recognizing graph properties from adjacency matrices. *Theor Comput Sci*, 3(3):371–384, 1976.