Rigging Tournament Brackets for Weaker Players

Isabelle Stanton[*]
UC Berkeley
isabelle@eecs.berkeley.edu

Virginia Vassilevska Williams[†]
UC Berkeley
virgi@eecs.berkeley.edu

Abstract

The agenda control problem for balanced single-elimination tournaments is the following natural problem in game manipulation. Consider a tournament designer with full knowledge of the match outcomes between any possible pair of players. The organizer would like to create a bracket for a balanced single-elimination tournament so that their favorite player will win. Although this tournament fixing problem has been studied in the areas of voting and tournament manipulation, it is still unknown whether it can be tackled in polynomial time. We focus on identifying several general cases for which the tournament can always be fixed efficiently for the given player. We give constructive proofs that under some natural assumptions, if a player is ranked among the top K players in terms of the number of potential matches they could win, then a resource-bounded tournament organizer with full knowledge of all the match outcomes can fix the tournament for the given player, even when K is a constant fraction of the number of players.

1   Introduction

As a natural way to select a leader, competition is at the heart of life. Competition is intriguing, both for its participants, and its spectators. Society is riddled with organized competitions called *tournaments* where there are well-defined rules to select a winner from a pool of candidate players. Sports tournaments such as the FIFA World Cup and Wimbledon are immensely popular and generate huge amounts of revenue. Elections are another important type of tournaments: a leading party is selected according to some rules using votes from the population.

Given the importance of competition, a natural question is how can one increase one's chance of winning by manipulating the tournament. There are numerous examples of this kind of behavior in a non-malicious setting. Political parties are able to select candidates for each district's election and they carefully take into account the candidate's strengths and weaknesses against the opposition when making this decision. Major television networks compete for advertising dollars based on their viewership and, as a result, select their program schedules so that their strongest programs face the least competition from opposing networks. Finally, there are numerous examples of this in sports competitions. For example, boxers are very selective of the challenges that they accept, and must be in order to preserve any titles they may hold. The incentive to manipulate competitions is enormous, in part due to the positive external bonuses beyond simply winning. If the party wins the election, then they can implement important policy decisions. Similarly, the best athletes often receive multi-million dollar endorsement deals from companies.

We study the problem of manipulating the bracket of a *single-elimination* tournament in the favor of a given player. Single-elimination tournaments are played as follows. First, a permutation of the players, called the *bracket*, *schedule* or *agenda* is given. Then, the tournament is played in rounds, where for each consecutive round a tournament bracket is obtained by taking the sorted order of the remaining players according to the original bracket. According to the current bracket, the first two players are matched up, then the second pair of players etc. The winners of the matches move on to the next round. If the number of players is a power of 2, the tournament is balanced; otherwise, it is

---

unbalanced and some players advance to the next round without playing a match. In practice, these *byes*, as they are called in athletic tournaments, are usually granted in the first round.

Although the winner of a single-elimination tournament is always well-defined, the chance of a particular player winning the tournament can vary immensely depending on the tournament bracket. Arguably, this gives the tournament organizer a lot of power. The study of how much control a tournament organizer has over the outcome of a tournament is called *agenda control* [2, 11]. In this work, we assume that the tournament organizer has knowledge of the results of a *round-robin* tournament. In the round-robin tournament, every two players are matched up, and a player's score is how many matches they won. We refer to this score as the rank.

The best studied agenda control problem for balanced single-elimination tournaments involves finding a bracket which maximizes the probability that a given player will win the tournament [1, 10, 15]. The tournament organizer is given the probability that i will beat j for every pair of players i, j. This problem is NP-hard, even if all probabilities are in $\{0, 1, 1/2\}$ [12, 8, 13]. However, when the probabilities are all either 0 or 1, i.e. the match outcomes are predetermined, the agenda control problem is not well understood. It is an open problem in the area of computational social choice whether a tournament fixing bracket can be efficiently computed (e.g. [12]). Several variants are known to be NP-hard – when some players cannot be matched [16], when some players must appear in given rounds [14], or when some matches are more interesting and the most interesting tournament must be computed [12, 16].

The $\{0, 1\}$-agenda control problem for balanced single-elimination tournaments is also called tournament fixing (TFP). Besides its natural connection to tournament manipulation, TFP studies the relationship between round-robin and single-elimination tournaments. The decision version of TFP asks, given the results of a round-robin tournament and a player a, is a also the winner of some single-elimination tournament, given the same match outcomes?

TFP is studied in the context of voting: suppose all voters have voted and a chairman has to set-up an agenda so that a winning candidate is selected using single-elimination rules (binary cup). Can the chairman manipulate the agenda so that a particular candidate wins? In voting protocols, any candidate selection rule in which candidates are compared in pairs (and the candidate with fewer votes is eliminated) is called a *voting tree* rule. Voting tree manipulation has been studied extensively [3, 9, 12]. It is known that if the voting tree has no prescribed structure, then there is a polynomial time algorithm to decide whether there exists a voting tree for which a given candidate wins the election [12]. Finding a *balanced* voting tree in polynomial time is a major open problem. Fischer et al. [7] consider whether (potentially unbalanced) voting trees can be used to always elect a candidate preferred by close to the majority of voters; *i.e.* how well binary tree protocols approximate the Copeland winner. Fischer et al. give hardness results for deterministic trees, and show that by randomizing over voting trees one can obtain much better approximations.

In this work we investigate the following question: if we consider a round-robin tournament and a ranking produced from it by sorting the players according to their number of wins, how many of the top players can actually win some single-elimination tournament, given the same match outcomes? What conditions on the round-robin tournament outcome suffice so that one can efficiently rig the tournament outcome for many of the top players?

We consider the round-robin ranking according to the number of wins since it is known that this ranking approximates the optimum ranking within a constant factor [5]. Prior work has shown several intuitive results. For instance, if a is any player with a maximal number of wins in a round-robin tournament (e.g. the top player in the ranking), then one can efficiently construct a winning (balanced) single-elimination tournament bracket for the player [16]. Also, if a player a has won at least as many matches as any player that beat a, then a is also a winner of some efficiently constructable single-elimination tournament [16]. Our work extends and strengthens many of the prior results.

2 Motivation for Assumptions

Consider a transitive tournament graph G with nodes $v_1, \ldots, v_n$, where v_i beats all nodes v_j for $j > i$. Then v_1 is the winner of all single-elimination tournaments on G. Now, create any perfect matching from $\{v_{n/2+1}, \ldots, v_n\}$ to $\{v_1, \ldots, v_{n/2}\}$ by reversing the appropriate edges. This gives each node from the weaker half of G a win against some node from the stronger half. The outdegree ranking does not change, however now the top $n/2 - 1$ players can win a single-elimination tournament: every strong node a still beats at least $n/2$ other players, and moreover because of the back-edges of the matching each such a is also a *king*, a player who has a win-path of length at most 2 to every other player. Prior work showed that if a is a king and beats at least $n/2$ other players, then a can win some single-elimination tournament [16]. Thus, adding a matching to a transitive tournament can dramatically increase the set of winners.
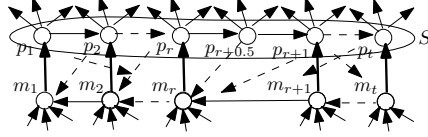
Figure 1: Here p_i only loses to p_j for $j < i$ and m_i. No matter how the other nodes are placed, since the p_i must beat everyone else and the m_i lose to everyone else, all single-elimination tournament winners are in S.

One of our goals is to understand the impact of the existence of matchings consisting of backedges with respect to the sorted order by outdegree, in general tournament graphs. In previous work, it was shown that a player with the highest score in the round-robin tournament can always win a single-elimination tournament, and a natural next question is to ask whether the second strongest player can as well. It is simple to show that this is true, provided the strongest player is beaten by at least one other person. This simple result leads to a larger question. If the second strongest player can always win a single-elimination tournament, provided the strongest node does not beat everyone, what about the third largest? The fourth? What are the necessary and sufficient conditions for the k^{th} node in the outdegree ranking to win a single-elimination tournament? A natural conjecture is that if there is a perfect matching onto the players that are both stronger than and beat a given player a, then a should be able to win.

In Figure 1 we give an example of a tournament graph and a subset S consisting of the top $t+1$ outdegree nodes of G. There is a matching of size t from $V \setminus S$ into S, but no matching of size $t+1$. In this graph, the only possible winners of single-elimination tournaments are contained in S. Since our goal is to give general conditions under which a node can win, the condition that there is a matching onto the stronger nodes that beat a is, in a sense, necessary.

While not all of our results assume the existence of this type of matching from the weaker to the stronger players, a natural question is how reasonable is the assumption of the existence of a matching from lower ranked players to higher ranked players? We give one example: One good model for generating tournament graphs is the Braverman-Mossel model [4]. In this model, one assumes an underlying ranking $v_1 \cdots v_n$ of the players according to skill and the tournament is generated by adding an edge (v_i, v_j) with probability p if $j < i$ and $1 - p$ if $j > i$ for $p < \frac{1}{2}$. This model can be viewed as the original transitive tournament we mentioned, with noise proportional to p on the back edges. A classic result of Erdős and Rényi [6] is that a bipartite graph with n nodes on each side with $2n \ln n$ edges selected uniformly at random contains a perfect matching with probability at least $1 - 2/n$. If a graph is generated by the Braverman-Mossel model with $p > \frac{4 \ln n}{n}$, then we expect there to be $n \ln n$ back edges from $v_{n/2} \cdots v_n$ to $v_1 \cdots v_{n/2-1}$, and a perfect back edge matching with high probability. Hence in almost all tournaments generated in this way, a back-edge matching exists. Because of this our results imply that for almost all graphs generated this way, one can fix the outcome of a single-elimination tournament for any one of the top $n/2 - O(\sqrt{n})$ players.

3   Contributions

Consider an ordering $\Pi = \{p_1, \ldots, p_n\}$ of the players in nondecreasing order of their number of wins in the given round-robin tournament. Our work considers conditions under which, for large K, the single-elimination tournament can be fixed efficiently for *any* of the first K players in Π. We are interested in natural and not too restrictive conditions under which a constant fraction of the players can be made to win. If the first player p_1 in Π beats everyone else (and is thus the Condorcet winner), then it is clear that p_1 wins all single-elimination tournaments. We show that if there exists a player who can beat p_1, then we can also fix the tournament for p_2. We then conjecture that if there is a matching from the bottom $n + 1 - K$ players $\{p_K, \ldots, p_n\}$ onto the top $K - 1$ players $\{p_1, \ldots, p_{K-1}\}$ in Π (so that there is a distinct bottom player that beats each top player), then we can efficiently find a bracket for which p_K wins. We prove the conjecture for tournaments larger than a fixed constant, and when K is at most 19% of the players.

It turns out that it is not strictly necessary that there should be a matching onto all of the top $K - 1$ players in order for p_K to always be able win. What is necessary, however, is that there exists a matching onto the set $H(p_K)$ of players in $\{p_1, \ldots, p_{K-1}\}$ which beat p_K; there are examples of round-robin tournaments and players p_K for which there is a matching into $H(p_K)$ of size $|H(p_K)| - 1$ and so that p_K is not a winner of any bracket. We prove that this necessary matching condition is often sufficient: Assuming a matching onto $H(p_K)$, we show that for large enough n, if $|H(p_K)| \leq n/7$, the bracket can always be rigged so that p_K wins. We note that for a general theorem of this form, a bound on $|H(p_K)|$ is necessary; e.g. a trivial bound is $|H(p_K)| \leq n/2$ since there is a matching onto $H(p_K)$.

We also strengthen and generalize results from prior work [16] on fixing the tournament for a player a that is a *king* in the sense that any player b that a cannot beat is beaten by some player that a can beat. Prior work showed that the tournament outcome can be efficiently fixed for a king a if either a beats at least half of the players, or if any player that a cannot beat is after a in the ranking Π. Using new techniques discussed below, we show that in fact, in order for us to be able to rig the bracket for a, it is sufficient for a to be able to beat $|H(a)| + 1$ players. This implies both of the previous results and also shows that if a is among the top third of the players and is a king, then a can always be made the winner. To summarize, we show the following positive results:

- If a is a king and beats more players than $|H(a)| + 1$, then a can win a single-elimination tournament. As a corollary, if a is a king and is in the top third of the players, then a can be made to win.

- If a is a king on $G \setminus T$ for a small subset T, and beats at least half the nodes in $V \setminus T$, then a can win provided there is a king k over T and at least $|T|$ players that beat k that a also beats.

- If a is not a king, but there is a perfect matching onto $H(a)$ then a can be made to win provided $|H(a)| \leq \frac{n-6}{7}$. As a corollary, this is always the case when a is in the top $\frac{6n+7}{31}$ players.

A single-elimination tournament can be viewed as a series of $\log n$ directed matchings, where the i^{th} matching is over the sources of the $(i-1)st$. In order to prove the above results, we develop several new techniques for finding matchings with special properties in a tournament graph. In particular, these include:

- We note several simple facts about the degree bounds of players. For example, if a is the k^{th} ranked player, then a must beat at least $\lfloor \frac{(n-k)}{2} \rfloor$ players in total. If a beats k higher ranked players, then k must beat at least $\frac{(n-k)}{3}$ players as well. These kind of bounds are used to determine when the following tools may be applied.

- A canonical matching for a player a that maximizes the number of players that a beats in the next round. In particular, if a is a king that beats at least half the other players, then after this matching, a is still a king who beats at least half the remaining players. Furthermore, a still beats at least half of the remaining players. We can guarantee the minimum size of the matching from $N_{out}(a)$ to $N_{in}(a)$ when we know the rank of a.

- The canonical matching involves some choices. We also provide a method for guaranteeing that the players that survive to the next round that beat a are covered by many short paths from a. This is achieved by using a greedy covering.

Ultimately, the tournament fixing problem is easiest when the problem is relatively unconstrained. The tournament fixing problem is identical to the problem of finding a binomial arborescence that is rooted at our given player. If there are many short paths from a to all the other players, then it becomes easier to greedily build a binomial arborescence in the graph, as we face fewer constraints of where certain players must appear based on previous choices.

References

[1] D. R. Appleton. May the best man win? *The Statistician*, 44(4):529–538, 1995.

[2] J. Bartholdi, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modeling*, 16(8/9):27–40, 1992.

[3] S. J. Brams and P. C. Fishburn. Voting procedures. *K. J. Arrow, A. K. Sen, and K. Suzumura, editors, Handbook of Social Choice and Welfare*, 2002.

[4] M. Braverman and E. Mossel. Noisy sorting without resampling. In *Proc. SODA*, pages 268–276, 2008.

[5] D. Coppersmith, L. Fleischer, and A. Rudra. Ordering by weighted number of wins gives a good ranking for weighted tournaments. In *Proc. SODA*, pages 776–782, 2006.

[6] P. Erdős and A. Rényi. On the existence of a factor of degree one of a connected random graph. *Acta. Math. Acad. Sci. Hungar.*, 17:359–368, 1966.

[7] F. Fischer, A.D. Procaccia, and A. Samorodnitsky. A new perspective on implementation by voting trees. In *Proc. EC*, pages 31–40, 2009.

[8] N. Hazon, P. E. Dunne, and M. Wooldridge. How to rig an election. In *Proc. BISFAI*, 2007.

[9] E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Anyone but him: The complexity of precluding an alternative. *Artif. Intell.*, 171(5/6):255–285, 2007.

[10] J. Horen and R. Riezman. Comparing draws for single elimination tournaments. *Operations Research*, 33(2):249–262, 1985.

[11] J. Bartholdi III, C. Tovey, and M. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice Welf.*, 6:157–165, 1989.

[12] J. Lang, M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Winner determination in sequential majority voting. In *Proc. IJCAI*, pages 1372–1377, 2007.

[13] T. Vu, A. Altman, and Y. Shoham. On the complexity of schedule control problems for knockout tournaments. In *Proc. AAMAS*, pages 225–232, 2009.

[14] T. Vu, N. Hazon, A. Altman, S. Kraus, Y. Shoham, and M. Wooldridge. On the complexity of schedule control problems for knock-out tournaments. *Submitted to JAIR*, 2010.

[15] T. Vu and Y. Shoham. Broadening the scope of optimal seeding analysis in knockout tournaments. In *Proc. AAMAS*, 1010.

[16] V. Vassilevska Williams. Fixing a tournament. In *Proc. AAAI*, 2010.