Manipulating Single-Elimination Tournaments in the Braverman-Mossel Model

Isabelle Stanton and **Virginia Vassilevska Williams**
Computer Science Department
UC Berkeley
{isabelle,virgi}@eecs.berkeley.edu

Abstract

We study the power of a tournament organizer in manipulating the outcome of a single elimination tournament by fixing the initial seeding. It is not known whether the organizer can efficiently fix the outcome of the tournament even if the match outcomes are known in advance. We generalize a result from prior work by giving a new condition such that the organizer can efficiently find a tournament bracket for which the given player will win the tournament. We then use this result to show that for most tournament graphs generated by the Braverman-Mossel model, the tournament organizer can (very efficiently) make a large constant fraction of the players win, by manipulating the initial bracket. This holds for very low values of the error probability, i.e. the generated tournament graphs are almost transitive. Finally, we obtain a trade-off between the error probability and the number of players that can efficiently be made winners.

Introduction

The study of election manipulation is an integral part of social choice theory. Results such as the Gibbard-Satterthwaite theorem [Gibbard, 1973; Satterthwaite, 1975] show that all voting protocols that meet certain rationality criteria are manipulable. The seminal work of [Bartholdi *et al.*, 1989; 1992] proposes to judge the quality of voting systems using computational complexity: a protocol may be manipulable, but it may still be good if manipulation is computationally expensive. This idea is at the heart of computational social choice.

The particular type of election manipulation that we study in this paper is called *agenda control* and was introduced in [Bartholdi *et al.*, 1992]: there is an election organizer who has power over some part of the protocol, say the order in which candidates are considered. The organizer would like to exploit this power to fix the outcome of the election by making their favorite candidate win. [Bartholdi *et al.*, 1992] focused on plurality and Condorcet voting, agenda control by adding, deleting, or partitioning candidates or voters. We study the balanced binary cup voting rule, also called a *single-elimination* tournament: the number of candidates is a power of 2; at each stage the remaining candidates are paired up and their votes are compared; the losers are eliminated and the winners move on to the next round, until only one candidate remains. The power of the election organizer is to pick the pairing of the players in each round. We assume that the organizer knows all the votes in advance, i.e. for any two candidates, he knows which candidate is preferred.

Single-elimination is prevalent in sports tournaments such as Wimbledon or March Madness. In this setting, a tournament organizer has some information, say from prior matches or from betting experts, about the winner in any possible player match. The organizer is to come up with a *seeding* of the players through which they are distributed in the tournament bracket. The question is, can the tournament organizer abuse this power to determine the winner of the tournament?

There is significant prior work on this problem. [Lang *et al.*, 2007] showed that if the tournament organizer only has probabilistic information about each match, then the agenda control problem is NP-hard. [Vu *et al.*, 2009; 2010] showed that the problem is NP-hard even when the probabilities are in $\{0, 1, 1/2\}$ and that it is NP-hard to obtain a tournament bracket that approximates the maximum probability that a given player wins within any constant factor. [Vassilevska Williams, 2010] showed that the agenda control problem is NP-hard even when the information is deterministic but some match-ups are disallowed. [Vassilevska Williams, 2010] also gave conditions under which one can fix the outcome of the tournament when the organizer knows each match outcome in advance. It is still an open problem whether one can always determine in polynomial time whether the tournament outcome can be fixed in this deterministic setting.

The binary cup is a complete binary voting tree. Other related work has studied more general voting trees [Hazon *et al.*, 2008; Fischer *et al.*, 2008], and manipulation by the players themselves by throwing games to manipulate single-elimination tournaments [Russell and Walsh, 2009].

The match outcome information available to the tournament organizer can be represented as a weighted or unweighted tournament graph, a graph such that for every two nodes u, v exactly one of (u, v) or (v, u) is an edge. An edge (u, v) signifies that u beats v, and a weight p on an edge (u, v) means that u will beat v with probability p. With this representation, the agenda control problem becomes a computational problem on tournament graphs. The tournament graph structure which comes from real world sports tournaments or from elections is not arbitrary. Although the graphs are not necessarily transitive, stronger players typically beat weaker ones. Some generative models have been proposed in order to study real-world tournaments. In this work, we study the Braverman-Mossel model [2008]. The basic idea is that there is an underlying total order of the players and the outcome of every match is probabilistic. There is some *global* probability $p << 1/2$ with which a weaker player beats a stronger player. This probability represents outside factors which do not depend on the players' abilities, such as weather or sickness.

[Vassilevska Williams, 2010] has shown that when $p \geq \Omega(\sqrt{\ln n/n})$, with high probability, the model generates a tournament graph where one can efficiently fix a single-elimination bracket for *any* given player. Two natural questions emerge. The first is can we still make almost all players win with a smaller noise value? The second is can we relax the Braverman-Mossel model to allow a different error probability for each pair of players? We address both questions.

Contributions We study whether one can compute a winning single-elimination bracket for a *king* player when the match outcomes are known in advance. A king is a player K such that for any other player a, either K beats a, or K beats some other player who beats a. We show that in order for a winning bracket to exist for a king, it is sufficient for the king to be among the top third of the players when sorted by the number of potential matches they can win. Before our work only much stricter conditions were known, e.g. that it is sufficient if the king beats half of the players. Our more general result allows us to obtain better results for the Braverman-Mossel model as well.

There are $\log n$ rounds in a single-elimination tournament over n players, so a necessary condition for a player to be a winner is that it can beat at least $\log n$ players. We consider a generalization of the Braverman-Mossel model in which the error probabilities $p(i, j)$ can vary but are all lower-bounded by a global parameter p. The expected outdegree of the weakest player i is $\sum_j p(i, j) \geq p(n-1)$, and it needs to be $\geq \log n$, so we focus on the case when p is $\Omega(\log n/n)$, as this is a necessary condition for all players to be winners.

Our results focus on this lower bound on the noise threshold. We improve previous results and show that when a tournament is generated with $\Omega(\log n/n)$ noise, we are able to fix the tournament for almost the top half of the players. We also show that there is a trade-off between the amount of noise and the number of players that can be made winners: as the level of noise increases, the tournament can be fixed for more and eventually all of the players. While this result does not answer the question of whether it is computationally difficult to fix a single-elimination tournament in general, it does show that for tournaments we might expect to see in practice, manipulation can be easy.

Notation
$N^{out}(a) = \{v \mid (a, v) \in E\}$
$N_X^{out}(a) = N^{out}(a) \cap X$
$N^{in}(a) = \{v \mid (v, a) \in E\}$, $N_X^{in}(a) = N^{in}(a) \cap X$
$out(a) = \lvert N^{out}(a)\rvert$, $out_X(a) = \lvert N_X^{out}(a)\rvert$
$in(a) = \lvert N^{in}(a)\rvert$, $in_X(a) = \lvert N_X^{in}(a)\rvert$
$\mathcal{H}^{in}(a) = \{v \mid v \in N^{in}(a), out(v) > out(a)\}$
$\mathcal{H}^{out}(a) = \{v \mid v \in N^{out}(a), out(v) > out(a)\}$
$\mathcal{H}(a) = \mathcal{H}^{in}(a) \cup \mathcal{H}^{out}(a)$
$E(X, Y) = \{(u, v) \mid (u, v) \in E, u \in X, v \in Y\}$

Table 1: A summary of the notation used in this paper.

Braverman-Mossel Model – Formal Definition

The premise of the Braverman-Mossel (BM) model is that there is an implicit ranking π of the players by intrinsic abilities so that $\pi(i) < \pi(j)$ means i has strictly better abilities than j. For clarity, we will call $\pi(i)$ i. When i and j play a match there may be outside influences so that even if $i < j$, j might beat i. The BM model allows that weaker players can beat stronger players, but only with probability $p < 1/2$. Here, p is a global parameter and if $i < j$, j beats i with probability $1 - p$. A random tournament graph generated in the BM model, a (*BM tournament*), is defined as: for every i, j with $i < j$, add edge (i, j) independently with probability $1 - p$ and otherwise add (j, i).

We give a generalization of the BM model, the GBM model, in which j beats i with probability $p(j, i)$, where $p \leq p(j, i) \leq 1/2$ for all i, j with $i < j$, *i.e.* the error probabilities can differ but are all lower-bounded by a global p. A random tournament graph generated in the GBM model (*GBM tournament*) is defined as: for every i, j with $i < j$, add edge (i, j) independently with probability $1 - p(j, i)$ and otherwise add (j, i).

Notation and Definitions Unless noted otherwise, all graphs in the paper are tournament graphs over n vertices, where n is a power of 2, and all single-elimination tournaments are balanced. In Table 1, we define the notation used in the rest of this paper. For the definitions, let $a \in V$ be any node, $X \subset V$ and $Y \subset V$ such that X and Y are disjoint. Given a player $\mathcal{A}$, unless otherwise stated, A denotes $N^{out}(\mathcal{A})$ and B denotes $N^{in}(\mathcal{A})$.

The outcome of a round-robin tournament has a natural graph representation as a tournament graph. The nodes of a tournament graph represent the players in a round-robin tournament, and a directed edge (a, b) represents a win of a over b.

We will use the concept of a *king* in a graph. Although the definition makes sense for any graph, it is particularly useful for tournaments, as the highest outdegree node is always a king.

Definition 1. *A* king *in $G = (V, E)$ is a node $\mathcal{A}$ such that for every other $x \in V$ either $(\mathcal{A}, x) \in E$ or there exists $y \in V$ such that $(\mathcal{A}, y), (y, x) \in E$.*

We also use the notion of a *superking*.

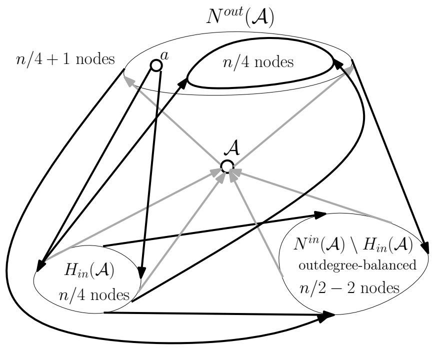

Figure 1: An example for which Theorem 1 does not apply, but for which Theorem 2 does apply.

Definition 2. *A* superking *in $G = (V, E)$ is a node $\mathcal{A}$ such that for every other $x \in V$ either $(\mathcal{A}, x) \in E$ or there exist $\log n$ nodes $y_1, \ldots, y_{\log n} \in V$ such that $(\mathcal{A}, y_i), (y_i, x) \in E$, $\forall i$.*

Kings that are also winners

A player being a king in the tournament graph is not a sufficient condition for it to also be able to win a single-elimination tournament. Consider that a player may be a king by beating only 1 player who, in turn, beats all the other players. [Vassilevska Williams, 2010] considered the question of how strong a king player needs to be in order for there to always exist a winning single-elimination tournament bracket for them.

Theorem 1. *[Vassilevska Williams, 2010] Let $G = (V, E)$ be a tournament graph and let $\mathcal{A} \in V$ be a king. One can efficiently construct a winning single-elimination tournament bracket for $\mathcal{A}$ if either*

$$\mathcal{H}^{in}(\mathcal{A}) = \emptyset, \text{ or } out(\mathcal{A}) \geq n/2.$$

We generalize the above result. The set $\mathcal{H}^{in}(\mathcal{A})$ represents all higher ranked nodes that beat the player $\mathcal{A}$. We show that it is sufficient for a player who is a king to only be as strong as the size of $\mathcal{H}^{in}(\mathcal{A})$.

Theorem 2 (Kings with High Outdegree). *Let G be a tournament graph on n nodes and $\mathcal{A}$ be a king. If $out(\mathcal{A}) \geq |\mathcal{H}^{in}(\mathcal{A})| + 1$, then one can efficiently compute a winning single-elimination bracket for $\mathcal{A}$.*

To see that the above theorem implies Theorem 1, note that if $out(\mathcal{A}) \geq n/2$, then $|\mathcal{H}^{in}(a)| \leq n/2 - 1 \leq out(\mathcal{A}) - 1$. Also, if $\mathcal{H}^{in}(\mathcal{A}) = \emptyset$ and $n \geq 2$, then $out(\mathcal{A}) \geq 1 \geq 1 + |\mathcal{H}^{in}(\mathcal{A})|$.

Theorem 2 is more general than Theorem 1. In Figure 1 we have an example of a tournament where Theorem 2 applies to the node $\mathcal{A}$ but not Theorem 1. Here, $|\mathcal{H}^{in}(\mathcal{A})| = \frac{n}{4}$, $|N^{out}(\mathcal{A})| = \frac{n}{4} + 1$ and the purpose of the node a is just to guarantee that $\mathcal{A}$ is a king. The example requires that each node in $N^{in}(\mathcal{A}) \setminus \mathcal{H}^{in}(\mathcal{A})$ have lower outdegree than $\mathcal{A}$ ($\frac{n}{4} + 1$) so we use an outdegree-balanced tournament for this set. This is a tournament where every vertex has outdegree equal to half the graph and it can be constructed inductively.

The intuition behind the proof of Theorem 2 is inspired by the results of [Stanton and Vassilevska Williams, 2011]. They show that a large fraction of highly ranked nodes can be tournament winners, provided a matching exists from the lower ranked to the higher ranked players. We are working with a king node so we are able to weaken the matching requirement. Instead, we carefully construct matchings that maintain that $\mathcal{A}$ is a king over the graph, while slowly eliminating the elements of $\mathcal{H}^{in}(\mathcal{A})$ until we reduce the problem to the case covered by Theorem 1.

We are now ready to prove Theorem 2. We will need a technical lemma from prior work relating the indegree and outdegree of two nodes. If a node $\mathcal{A}$ is a king then for every other node b, $N^{out}(\mathcal{A}) \cap N^{in}(b) \neq \emptyset$. This lemma is useful for showing a node is a king.

Lemma 1 ([Vassilevska Williams, 2010])**.** *Let a be a given node, $A = N^{out}(a), B = N^{in}(a)$, $b \in B$. Then $out(a) - out(b) = in_A(b) - out_B(b)$. In particular, $out(a) \geq out(b)$ if and only if $out_B(b) \leq in_A(b)$.*

Now we can prove Theorem 2.

Proof of Theorem 2: We will design the matching for each consecutive round r of the tournament. In the induced graph before the r^{th} round, let $\mathcal{H}_r$ be the subset of $\mathcal{H}^{in}(\mathcal{A})$ that is still live, A_r be the current outneighborhood of $\mathcal{A}$ and B_r be the current inneighborhood of $\mathcal{A}$. We will keep the invariant that if $B_r \setminus \mathcal{H}_r \neq \emptyset$, we have $|A_r| \geq |\mathcal{H}_r| + 1$, $\mathcal{A}$ is a king and the subset of nodes from the inneighborhood of $\mathcal{A}$ that have larger outdegree than $\mathcal{A}$ is contained in $\mathcal{H}_r$.

We now assume that the invariant is true for round $r - 1$. We will show how to construct round r. If $\mathcal{H}_r = \emptyset$ we are done by reducing the problem to Theorem 1, so assume that $|\mathcal{H}_r| \geq 1$. We begin by taking a maximal matching M_r from A_r to $\mathcal{H}_r$. Since $|A_r| \geq |\mathcal{H}_r| + 1$, $A_r \setminus M_r \neq \emptyset$ i.e. M_r can not match all of A_r. Now, let M'_r be a maximal matching from $A_r \setminus M_r$ to $B_r \setminus \mathcal{H}_r$.

If $A_r \setminus (M'_r \cup M_r) \neq \emptyset$, there is some node a' leftover to match $\mathcal{A}$ to. Otherwise, pick any $a' \in M'_r \cap A_r$. Remove the edge matched to a' from M'_r and match a' with $\mathcal{A}$. To complete the matching, create maximal matchings within $\bar{A}_r = A_r \setminus (M'_r \cup M_r) \setminus \{a'\}$, $\bar{B}_r = B_r \setminus \mathcal{H}_r \setminus M'_r$ and $\mathcal{H}_r \setminus M_r$. Either zero or two of $|\bar{A}_r|, |\bar{B}_r|, |\mathcal{H}_r \setminus M_r|$ can be odd and so there are at most 2 unmatched nodes. These can be matched them against each other. Let M represent the union of all of these matchings.

We will now show that the invariants still hold. Notice that $\mathcal{A}$ is still a king on the sources of the created matching M. Now, consider any node b from $B_r \setminus \mathcal{H}_r$ which is a source in M. We have two choices. The first is that b survived by beating another node of B_r so it lost at least one outneighbor from B_r. Since M'_r was maximal, b may have lost at most one of its inneighbors (a'). Hence we still have

$$out_{B_{r+1}}(b) + 1 \leq (out_{B_r}(b) - 1 + 1) \leq$$
$$\leq in_{A_r}(b) - 1 \leq in_{A_{r+1}}(b).$$

By Lemma 1 this means that $out(b) \leq out(\mathcal{A})$. The second choice is if b survived by beating a leftover node $\bar{a}$ from A_r. This can only happen if $A_r \setminus (M'_r \cup M_r) \neq \emptyset$. Thus, $\bar{a}$ was in $A_r \setminus (M'_r \cup M_r)$. However, since M'_r was maximal, $\bar{a}$

must lose to b, and so all inneighbors of b from A_r move on to the next round, and again $out(b) \leq out(\mathcal{A})$. Hence $\mathcal{A}$ has outdegree at least as high as that of all nodes in $B_{r+1} \setminus \mathcal{H}_{r+1}$.

Now we consider A_{r+1} vs $\mathcal{H}_{r+1}$. We have

$$|A_{r+1}| \geq \lfloor(|A_r| + |M'_r| + |M_r| - 1)/2\rfloor, \text{ and}$$

$$|\mathcal{H}_{r+1}| \leq \lceil(|\mathcal{H}_r| - |M_r|)/2\rceil = \lfloor(|\mathcal{H}_r| + 1 - |M_r|)/2\rfloor.$$

Since $|\mathcal{H}_r| \geq 1$ we must have $|M_r| \geq 1$. If either $|M_r| \geq 2$, $|A_r| \geq |\mathcal{H}_r|+2$, or $|M'_r| \geq 1$ then it must be that $|A_{r+1}| \geq \lfloor(|\mathcal{H}_r| + 2)/2\rfloor \geq |\mathcal{H}_{r+1}| + 1$. Also, if $|\mathcal{H}_r|$ is even then

$$|A_{r+1}| \geq |\mathcal{H}_r|/2 = 1 + \lfloor(|\mathcal{H}_r| - 1)/2\rfloor \geq |\mathcal{H}_{r+1}| + 1.$$

On the other hand, assume that $|M_r| = 1, |M'_r| = 0, |A_r| = |\mathcal{H}_r| + 1$ and $|\mathcal{H}_r|$ is odd. This necessarily implies that $|B_r \setminus \mathcal{H}_r| \leq 1$. Since $|A_r| = |\mathcal{H}_r| + 1$ is even, $|B_r|$ must be odd and so $|B_r \setminus \mathcal{H}_r|$ must be even. $|B_r \setminus \mathcal{H}_r|$ can only be 0. This means $|\mathcal{H}_r| = n_r/2 - 1$ (where n_r is the current number of nodes). We can conclude that $\mathcal{A}$ is a king with outdegree at least half the graph and the tournament can be efficiently fixed so that $\mathcal{A}$ wins by Theorem 1. □

Theorem 2 implies the following corollaries.

Corollary 1. *Let G be a tournament graph on n nodes and $\mathcal{A}$ be a king. If $|\mathcal{H}^{in}(\mathcal{A})| \leq (n-3)/4$, then one can efficiently compute a winning single-elimination tournament bracket for $\mathcal{A}$.*

Corollary 2. *Let G be a tournament graph on n nodes and $\mathcal{A}$ be a king in G. If $|\mathcal{H}(\mathcal{A})| \leq n/3 - 1$, then one can efficiently compute a winning single-elimination tournament bracket for $\mathcal{A}$.*

The proof of Corollary 1 follows by the fact that if $|\mathcal{H}^{in}(\mathcal{A})| = k$, then $out(\mathcal{A}) \geq (n-k)/3$. Corollary 2 simply states that any player in the top third of the bracket who is a king is also a tournament winner.

Proof of Corollary 2: Let $K = |\mathcal{H}(\mathcal{A})|$. Then the outdegree of $\mathcal{A}$ is at least $(n - K - 1)/2$. Let $h = |\mathcal{H}^{in}(\mathcal{A})|$. Then by Theorem 2, a sufficient condition for $\mathcal{A}$ to be able to win a single-elimination tournament is that $out(\mathcal{A}) \geq h+1$. Hence it is sufficient that $n-K-1 \geq 2h+2$, or that $2h+K \leq n-3$. Since $2h + K \leq 3K$, it is sufficient that $3K \leq n - 3$, and since $K \leq (n-3)/3$ we have our result. □

Braverman-Mossel Model

We can now apply our results to graphs generated by the Braverman-Mossel Model. From prior work we know that if $p \geq C\sqrt{\ln n/n}$ for $C > 4$, then with probability at least $1 - 1/\text{poly}(n)$, any node in a tournament graph generated by the BM model can win a single-elimination tournament. However, since p must be less than $1/2$, this result only applies for $n \geq 512$. Moreover, even for $n = 8192$ the relevant value of p is $> 13\%$ which is a very high noise rate. We consider how many players can be efficiently made winners when p is a slower growing function of n. We show that even when $p \geq C \ln n/n$ for a large enough constant C, a constant fraction of the top players in a BM tournament can be efficiently made winners.

Theorem 3 (**BM Model Winners for Lower** p)**.** *For any given constant $C > 16$, there exists a constant n_C so that for all $n > n_C$ the following holds. Let $p \geq C \ln n/n$, and let G be a tournament graph generated by the BM model with error p. Then with probability at least $1 - 3/n^{C/8-2}$, any node v with $v \leq n/2 - 5C\sqrt{n \ln n}$ can win a single-elimination tournament.*

This result applies for $n \geq 256$ and also reduces the amount of noise needed. For example, if $C = 17$ then when $n = 8192$, it is only necessary that $p < 2\%$, as opposed to $> 13\%$. This is a significant improvement. The proof of Theorem 3 uses Theorem 2 and Chernoff-Hoeffding bounds.

Theorem 4 (Chernoff-Hoeffding)**.** *Let $X_1, \ldots, X_n$ be random variables with $X = \sum_i X_i$, $E[X] = \mu$. Then for $0 \leq D < \mu$, $Pr[X \geq \mu + D] \leq exp(-D^2/(4\mu))$ and $Pr[X < \mu - D] \leq exp(-D^2/(2\mu))$.*

Proof of Theorem 3: Let C be given. Consider j. The expected of the number n_j of outneighbors of j in G is

$$E[n_j] = (1-p)(n-j)+(j-1)p = n(1-p)-p-j(1-2p).$$

This is exactly where we use the BM model. Our result is not directly applicable to the GBM model because this is only a lower bound on the expectation of n_j in that model. We will show that with high probability, all n_j are concentrated around their expectations and that all nodes $j \leq n/2$ are kings.

Showing that each n_j is concentrated around its' expectation is a standard application of the Chernoff bounds and a union bound. Therefore, $2/n^{C^2/4} < 1/n^C$ for $C > 16$ and $n > 2$. with probability at least $1 - 1/n^{C-1}$ for every j, $|E[n_j] - n_j| \leq C\sqrt{n \ln n}$.

We assume n is large enough so that $n >> \sqrt{n \ln n}$. We also assume that $p \leq 1/4$ so that $1 \geq (1-2p) \geq 1/2$. Now fix $j \leq n/2$. By the concentration result, this implies that

$$n_j \geq 3n/4 - 1 - j - C\sqrt{n \ln n} \geq$$

$$n/4 - 1 - C\sqrt{n \ln n} \geq \varepsilon n,$$

where $\varepsilon = 1/8$ works. The probability that j is a king is quite high: the probability that some node z has no inneighbor from $N^{out}(j)$ is at most

$$n(1-p)^{n_j} \leq n(1 - C \ln n/n)^{(n/(C \ln n)) \cdot C\varepsilon \ln n}$$

$$\leq 1/n^{\varepsilon C - 1}.$$

By a union bound, the probability that some node j is not a king is at most $1/n^{\varepsilon C-2}$. Therefore, we can conclude that the probability that all the n_j are concentrated around their expectations and all nodes $j \leq n/2$ are kings is at least $1 - (1/n^{C-1} + 1/n^{\varepsilon C-2})$.

We now need to upper bound $|\mathcal{H}^{in}(j)|$. We are interested in how many nodes with $i < j + 2C\sqrt{n \ln n}/(1-2p)$ appear in $N^{in}(j)$: if we have an upper bound on them, we can apply Theorem 2 to get a bound on j. First, consider how small $n_j - n_i$ can be for any i:

$$n_j - n_i \geq (i-j)(1-2p) - 2C\sqrt{n \ln n}.$$

So for $i \geq j + 2C\sqrt{n \ln n}/(1-2p)$, $n_j \geq n_i$ with high probability. The expected number of nodes $i < j$ that appear in $N^{in}(j)$ is $(1-p)(j-1)$. By the Chernoff bound, the probability that at least $(1-p)(j-1)+C\sqrt{j \ln n}$ of the $j-1$ nodes less than j are in $N^{in}(j)$ is $\leq exp(-C^2 j \ln n/4j) = n^{-C^2/4}$. Therefore, with probability at least $1-1/n^{C^2/4}$, the number of such i is at most $(1-p)(j-1)+C\sqrt{j \ln n}$. By a union bound, this holds for all j with probability at least $1-1/n^{C^2/4-1}$. Now, we can say with high probability that $|\mathcal{H}^{in}(j)|$ is at most

$$(1-p)(j-1)+C\sqrt{j \ln n}+2C\sqrt{n \ln n}/(1-2p) \leq$$

$$\leq (1-p)(j-1)+5C\sqrt{n \ln n}.$$

By Theorem 2, for there to be a winning bracket for j, it is sufficient that $\mathcal{H}^{in}(j) < n_j$ or that

$$(1-p)(j-1)+5C\sqrt{n \ln n} <$$

$$n(1-p)-p-j(1-2p)-C\sqrt{n \ln n}$$

. This is equivalent to

$$j < \frac{n(1-p)}{(2-3p)} + \frac{(1-2p)}{(2-3p)} - 6C\frac{\sqrt{n \ln n}}{(2-3p)}.$$

It is sufficient if

$$j < n/2 + \frac{pn}{(2(2-3p))} + \frac{(1-2p)}{(2-3p)} - 24C\sqrt{n \ln n}/5,$$

and so for all $j \leq n/2 - 5C\sqrt{n \ln n}$, there is a winning bracket for j with probability at least

$$1-(2/n^{C-1}+1/n^{\varepsilon C-2}) \geq 1-3/n^{C/8-2}.$$

□

Improving the result for the GBM model through perfect matchings.

Next, we show that there is a trade-off between the constant in front of $\log n/n$ and the fraction of nodes that can win a single-elimination tournament. The proofs are based on the following result by [Erdős and Rényi, 1964]. Let $B(n,p)$ denote a random bipartite graph on n nodes in each partition such that every edge between the two partitions appears with probability p.

Theorem 5 ([Erdős and Rényi, 1964]). *Let c_n be any function of n, then consider $G = B(n,p)$ for $p = (\ln n + c_n)/n$. The probability that G contains a perfect matching is at least $1-2/e^{c_n}$.*

For the particular case $c_n = \Theta(\ln n)$, G contains a perfect matching with probability at least $1-1/\text{poly}(n)$.

Lemma 2. *Let $C \geq 64$ be a given constant. Let $n \geq 16$. Let G be a GBM tournament for $p = C \ln n/n$. Then with probability at least $1-2/n^{C/32-1}$, G is such that one can efficiently construct a winning single-elimination tournament bracket for the node ranked 1.*

Proof. We will call the top ranked node s. We will show that with high probability s has outdegree at least $n/4$ and that every node in $N^{in}(s)$ has at least $\log n$ inneighbors in $N^{out}(s)$. This makes s a superking, and by [Vassilevska Williams, 2010], s can win a single-elimination tournament.

The probability that s beats any node j is $> 1/2$, the expected outdegree of s is $> (n-1)/2$. By a Chernoff bound, the probability that s has outdegree $< n/4$ is at most $exp(-(n-1)/16) << 1/n^{C/32-1}$. Given that the outdegree of s is at least $n/4$, the expected number of inneighbors in $N^{out}(s)$ of any particular node y in $N^{in}(s)$ is at least $(n/4)\cdot(C \ln n/n) = (C/4)\ln n$.

We can show that each node in $N^{in}(s)$ has at least $\log n$ inneighbors from $N^{out}(s)$ by using a Chernoff bound and union bound. By a Chernoff bound, the probability that y has less than $(C/8)\ln n$ inneighbors from $N^{out}(s)$ is at most $exp(-(C/32)\ln n) = 1/n^{C/32}$. By a union bound, the probability that some $y \in N^{in}(s)$ has less than $(C/8)\ln n$ inneighbors from $N^{out}(s)$ is at most $1/n^{C/32-1}$. Therefore, s is a superking is with probability at least $1-2/n^{C/32-1}$ where $n \geq 16$, $n/4 \geq \log n$, $C > 64$, and $(C/8)\ln n \geq \log n$. □

Lemma 2 concerned itself only with the player who is ranked highest in intrinsic ability. The next theorem shows that as we increase the noise factor, we can fix the tournament for an increasingly large set of players. As the noise level increases, we can argue recursively that there exists a matching from $\frac{n}{2}+1\ldots n$ to $1\ldots\frac{n}{2}$, and from $\frac{3n}{4}+1\ldots n$ to $\frac{n}{2}+1\ldots\frac{3n}{4}$ and so forth. These matchings form each successive round of the tournament, eliminating all the stronger players.

Theorem 6. *Let $n \geq 16$, $i \geq 0$ be a constant and $p \geq 64\cdot 2^i \ln n/n \in [0,1]$. With probability at least $1-1/poly(n)$, one can efficiently construct a winning single-elimination tournament bracket for any one of the top $1+n(1-1/2^i)$ players in a GBM tournament.*

Proof. Let G be a GBM tournament for $p = C2^i \ln n/n$, $C \geq 64$. Let S be the set of all $n/2^{i-1}$ players j with $j > n(1-1/2^{i-1})$. Let s be a node with $1+n(1-1/2^{i-1}) \leq s \leq 1+n(1-1/2^i)$. The probability that s wins a single-elimination tournament on the subtournament of G induced by S is high: there is a set X of at least $n/2^i-1$ nodes that are after s. By Lemma 2, s wins a single-elimination tournament on $X \cup \{s\}$ with high probability $1-\frac{2}{(n/2^i)^{C/32-1}}$.

In addition, by Theorem 5, with probability at least $1-\frac{2}{(n/2^i)^{C-1}}$, there is a perfect matching from $X \cup \{s\}$ to $S \setminus (X \cup \{s\})$. For every $1 \leq k \leq i-1$, consider

$$A_k = \{x \mid 1+n(1-1/2^k) \leq x\}, \text{ and}$$

$$B_k = \{x \mid 1+n(1-1/2^{k-1}) \leq x \leq n(1-1/2^k)\}.$$

Then $A_{k-1} = A_k \cup B_k$, $A_k \cap B_k = \emptyset$, and $|A_k| = |B_k| = n/2^k$. Hence $p \geq C \ln |A_k|/|A_k|$ for all $k \leq i-1$. By Theorem 5, the probability that there is no perfect matching from A_k to B_k for a particular k is at most $2/(n/2^k)^{C2^{i-k}-1}$. This value is maximized for $k = i$, and it is $2/(n/2^i)^{C-1}$. Thus by a union bound, with probability at least $1-2i/(n/2^i)^{C-1} =$

$1-1/\text{poly}(n)$, there is a perfect matching from A_k to B_k, for every k.

Thus, with probability at least $1-1/\text{poly}(n)$, s wins a single-elimination tournament in G with high probability, and the full bracket can be constructed by taking the unions of the perfect matchings from A_k to B_k and the bracket from S. □

For the BM model we can strengthen the bound from Theorem 3 by combining the arguments from Theorems 3 and 6.

Theorem 7. *There exists a constant n_0 such that for all $n > n_0$ the following holds. Let $i \geq 0$ be a constant, and $p = 64 \cdot 2^i \ln n/n \in [0,1]$. With probability at least $1-1/poly(n)$, one can efficiently construct a winning bracket for any one of the top $n(1-1/2^{i+1}) - (80/2^{i/2})\sqrt{n \ln n}$ players in a BM tournament.*

As an example, for $p = 256 \ln n/n$, Theorem 7 says that any of the top $7n/8 - 40\sqrt{n \ln n}$ players are winners while Theorem 6 only gives $3n/4+1$ for this setting of p in the GBM model.

Proof. As in Theorem 6, for every $1 \leq k \leq i$, consider $A_k = \{x \mid 1+n(1-1/2^k) \leq x\}$, and $B_k = \{x \mid 1+n(1-1/2^{k-1}) \leq x \leq n(1-1/2^k)\}$. Then $A_{k-1} = A_k \cup B_k$, $A_k \cap B_k = \emptyset$, and $|A_k| = |B_k| = n/2^k$. By the argument from Theorem 6, w.h.p. there is a perfect matching from A_k to B_k, for all k.

Consider A_i. By Theorem 3, with probability $1 - 1/\text{poly}(n/2^i) = 1-1/\text{poly}(n)$, we can efficiently fix the tournament for any of the first $n/2^{i+1} - 5 \cdot 16\sqrt{(n/2^i)\ln(n/2^i)}$ nodes in A_i. Combining the construction with the perfect matchings between A_k and B_k, we can efficiently construct a winning tournament bracket for any of the top

$$n - n/2^i + n/2^{i+1} - 80\sqrt{(n/2^i)\ln(n/2^i)} \geq$$

$$\geq n(1-1/2^{i+1}) - (80/2^{i/2})\sqrt{n \ln n} \text{ nodes.}$$ □

Conclusions

In this paper, we have shown a tight bound (up to a constant factor) on the noise needed to fix a single-elimination tournament for a large fraction of players when the match outcomes are generated by the BM model. As this model is believed to be a good model for real-world tournaments, this result shows that many tournaments in practice can be easily manipulated. In some sense, this sidesteps the question of whether it is NP-hard to fix a tournament in general by showing that it is easy on examples that we care about.

Acknowledgements

The authors are grateful for the detailed comments from the anonymous reviewers. The first author was supported by the National Defense Science and Engineering Graduate Fellowship, the National Science Foundation Graduate Fellowship and partially supported by the National Science Foundation Grant CCF-0830797. The second author was supported by the National Science Foundation under Grant #0937060 to the Computing Research Association for the CIFellows Project. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the Computing Research Association.

References

[Bartholdi *et al.*, 1989] J. Bartholdi, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice Welfare*, 6(3):227–241, 1989.

[Bartholdi *et al.*, 1992] J. Bartholdi, C. Tovey, and M. Trick. How hard is it to control an election. *Mathematical and Computer Modeling*, pages 27–40, 1992.

[Braverman and Mossel, 2008] M. Braverman and E. Mossel. Noisy sorting without resampling. In *Symposium on Discrete Algorithms (SODA)*, pages 268–276, 2008.

[Erdős and Rényi, 1964] P. Erdős and A. Rényi. On random matrices. *Publications of the Mathematical Institute Hungarian Academy of Science*, 8:455–561, 1964.

[Fischer *et al.*, 2008] F. Fischer, A. D. Procaccia, and A. Samorodnitsky. On voting caterpillars:approximating maximum degree in a tournament by binary trees. In *2nd International Workshop on Computational Social Choice (COMSOC)*, 2008.

[Gibbard, 1973] A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41, 1973.

[Hazon *et al.*, 2008] N. Hazon, P.E. Dunne, S. Kraus, and M. Wooldridge. How to rig elections and competitions. In *2nd International Workshop on Computational Social Choice (COMSOC)*, 2008.

[Lang *et al.*, 2007] J. Lang, M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Winner determination in sequential majority voting. In *The Eighteenth International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.

[Russell and Walsh, 2009] T. Russell and T. Walsh. Manipulating tournaments in cup and round robin competitions. In *Algorithmic Decision Theory*, 2009.

[Satterthwaite, 1975] M. A. Satterthwaite. Strategy-proofness and arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10, 1975.

[Stanton and Vassilevska Williams, 2011] I. Stanton and V. Vassilevska Williams. Rigging tournament brackets for weaker players. *The Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, 2011.

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

[Vu *et al.*, 2009] T. Vu, A. Altman, and Y. Shoham. On the complexity of schedule control problems for knockout tournaments. In *The Eighth International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2009.

[Vu *et al.*, 2010] 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.