Probabilistic Copeland Tournaments

Sam Saarinen
University of Kentucky
Lexington, KY
samuel.saarinen@uky.edu

Judy Goldsmith
University of Kentucky
Lexington, KY
goldsmit@cs.uky.edu

Craig Tovey
Georgia Tech
Atlanta, GA
ctovey@isye.gatech.edu

INTRODUCTION

We consider a probabilistic model of round-robin tournaments, or equivalently, Copeland voting, where candidates are the voters. We assume that the outcomes of each game or pairwise vote are jointly independent. In particular, we do not assume that votes arise from voters' ranked orderings of candidates. We can treat such games as pairwise preferences, without assuming any form of transitivity. We prove the #P-completeness of computing the probability of victory. As a consequence, it is #P-hard to manipulate a round-robin tournament by controlling the outcome of a subset of the games to raise the probability of winning above a particular threshhold. These results hold in the restricted case where all probabilities are zero, one half, or one.

According to Faliszewski, et al. [4], the notion of probabilistic Copeland elections go back to a 1929 paper by Zermelo [15] and more recently to Levin and Nalebuff [7]. In 2005, Konczak and Lang looked at Copeland elections with incomplete ballots, although they did not introduce probabilities [5]. Instead, they considered *possible* and *necessary winners* (probabilities > 0 and 1, respectively). These notions have been well studied (e.g., [3, 6, 11, 13, 14]).

Bachrach et al. [2] introduced a probabilistic interpretation of incomplete ordered ballots. In their interpretation, the ranked candidates are preferred to all candidates not mentioned, and all completions of the partial linear order are equally likely. This introduces correlations in the probabilities of individual pairings, which we do not assume. Bachrach et al. showed that computing the probability of a given candidate winning, in this setting, was #P-hard, using techniques that do not apply in the tournament setting.

Definition: A *Probabilistic Copeland Tournament (PCT)* is represented by an $n \times n$ nonnegative matrix, T, where $T_{i,j} + T_{j,i} = 1\ \forall i, j$. The n row and column indices of T represent n teams,[1] each distinct pair of which will play one game. Team i defeats team j with probability $T_{i,j}$, and game outcomes are jointly independent. Therefore, the probability of a set of game outcomes equals the product of the probabilities of the individual game outcomes. Equivalently, a PCT is represented by a complete, directed, simple graph (V, E), where V is the set of teams, with edge weights $w_e : e \in E$ such that $0 \leq w_e \leq 1\ \forall e \in E$. An edge (i, j) with weight w_{ij} has the same meaning as T_{ij}. Team i defeats team j with probability w_{ij} and loses to team j with complementary probability $1 - w_{ij}$. This graph representation employs only one edge between each distinct pair of vertices i and j.

[1]It is more common in AI to refer to "agents" or "candidates". We stick to "teams" for consistency.

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May, 4–8, 2015, Istanbul, Turkey.*

Figure 1: All four Eulerian orientations of a graph

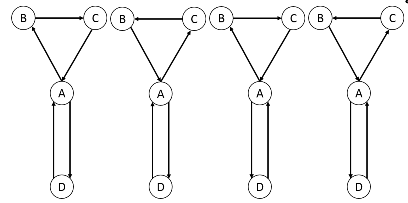

Definition: Let t be a distinguished agent of a PCT represented by matrix T. The *Probabilistic Tournament Problem (PTP)* is the problem of computing the probability, denoted $P_W(T, t)$, that t wins at least as many games as any other agent. If the PCT is represented by graph G, the probability is denoted $P_W(G, t)$.

Definition: Let T represent a PCT with distinguished agent t. The *Unique-winner Probabilistic Tournament Problem (UPTP)* is the problem of computing the probability, denoted $P_{UW}(T, t)$, that t wins strictly more games than any other agent. If the PCT is represented by graph G, the probability is denoted $P_{UW}(G, t)$.

The PTP was shown to be in #P [8], and a similar argument shows that the UPTP is in #P. However, no other previous hardness results are known. Work by Aziz, et al. has explored probabilistic knockout tournaments [1], but their results are not applicable to a PCT.

COMPLEXITY RESULTS

THEOREM 1. *The PTP is #P-complete, even if all probabilities are in* $\{0, \frac{1}{2}, 1\}$.

Given the previous result, which places the PTP in the complexity class #P, it suffices to show that the PTP is #P-hard. We prove the hardness of the PTP by a reduction from counting the number of Eulerian orientations of an Eulerian graph. An *Eulerian graph* is an undirected graph all of whose nodes have even degree. An *Eulerian orientation* of an Eulerian graph is a choice of orientation for each edge such that every vertex's indegree equals its outdegree. (See Figure 1.) The problem of counting the number of Eulerian orientations of an Eulerian graph has been shown to be #P-complete by a reduction from counting perfect matchings [9].

Assume we are given an Eulerian graph $H = (U, F)$. We will now construct a tournament graph $G = (V, E)$ with $U \subset V; F \subset$

E in which vertices $v \in V$ represent teams, edges corresponding to those in F have weight $\frac{1}{2}$, representing random game outcomes, and all other edges have weight 1, representing fixed game outcomes. We will set the fixed game outcomes so that a special team $z \notin U$ is a tournament winner in exactly the cases where every team in U wins exactly half of its random games. Thus, there is a one-one correspondence between Eulerian orientations of the edges of F and tournaments in which z is a co-winner; to obtain the number of Eulerian orientations, multiply z's probability of being a winner by $2^{|F|}$.

In the proof of the #P-completeness of #EULERIAN ORIENTATIONS [9], graphs can have parallel edges. Without loss of generality we can assume that H has none, by inserting a vertex into each parallel edge to transform them into paths of edge-length two. Since each inserted vertex has degree 2, every Eulerian orientation of the altered graph must make each of these paths behave like an edge. Hence the number of Eulerian orientations of the altered graph equals the number of Eulerian orientations of the original graph.

Define G to have vertex set $\{z\} \bigcup U \bigcup Y$. Label the vertices in U as $1, \dots n$. Set $|Y| = 4n + 1$, and label the elements of Y as $y_0, \dots, y_{4n}$. Let edge $(y_i, y_{(i+j \mod 4n+1)})$ have weight 1 for all $1 \le j \le 2n$, so each member of Y defeats exactly $2n$ other members of Y, and is defeated by the same number. For $i, j \in U$, set the weights of $(i, j) \notin F : i < j$ to 1 and make each member of U certain to defeat exactly enough members of Y so that the number of games it is certain to win plus half its degree in F equals $4n$. If i must defeat k members of Y let it defeat $y_1, \dots, y_k$. Set z to defeat all members of U (the weight of (z, i) is 1 for $1 \le i \le n$) and set z to defeat $y_1, \dots, y_{3n}$ but lose the other games. Now z is certain to win exactly $4n$ games. There are $5n+2$ teams. Every member of Y is certain to lose at least $2n$ games, hence can't win the tournament. Thus, z wins iff all members of U win exactly half of their random games. Therefore, the number of Eulerian orientations of H equals $2^{|F|}P_W(G, z)$. This completes the reduction.

COROLLARY 2. *The UPTP is #P-complete even if all probabilities are in* $\{0, \frac{1}{2}, 1\}$.

The hardness reduction is identical to that for the PTP, except that z gets one more win.

Hardness of PCT Manipulation

It is often important to evaluate the complexity of manipulation of a competition such as a round-robin tournament, or of a social choice mechanism as proposed earlier. As a consequence of Theorem 1, several versions of the manipulation question that we have examined are #P-hard. The exact complexity remains elusive, however. Consider the following manipulations.

PCTM: Given a rational number p, a PCT with distinguished team t, and a set of controllable games S, can the outcomes of the games in S be fixed so that $P_W(t) > p$? *PCTB:* Given rational number p, a PCT with distinguished team t, a set of games S that are bribable, integer bribery costs $c_e : e \in S$, and integer budget b, the *Probabilistic Copeland Tournament Bribery Problem (PCTB)* is the decision problem: is there a choice of games $B \subset S$ with total bribery cost $\sum_{e \in B} c_e \le b$ and a set of outcomes for the games in B such that $P_W(t) > p$? PCTM is a special case of PCTB where $c_e = 1 \forall e \in S$ and $b = |S|$. The unique-winner versions of these problems are defined analogously.

THEOREM 3. *PCTB and PCTM, restricted to all probabilities in* $\{0, \frac{1}{2}, 1\}$*, are #P-hard. Unique PCTB and Unique PCTM are also #P-hard under the same restriction.*

In the construction for Theorem 1, set $S = \phi$ and do binary search over the range $p = \frac{i}{2^{|F|}} : 0 \le i \le 2^{|F|}$. This search Turing-reduces PTP to PCTM. Hardness follows for the more general PCTB problem. The reduction for Unique PCTM works the same way.

OBSERVATION 4. *Straightforward simulation of tournament trials cannot estimate* $P_W(G, z)$ *to within a constant factor with high probability in polynomial time, because exponentially many trials would be needed to distinguish between the values* 0 *and* $1/(2^{n-2})$ *for team* 1 *for the following probabilities:* $T(1, j) = .5 : j = 3, \dots, n; T(1, 2) = 0; T(2, j) = 1; j = 3, \dots, n-1; T(2, n) = 1$ *and all games between* $3 \le j \le k$ *having fixed outcomes such that each of those teams wins between* $(k-5)/2$ *and* $(k-3)/2$ *of those games.*

On the other hand, it can be determined in polynomial time whether or not $P_W(G, z) = 0$ for any set of probabilities [10, 12].

Acknowledgments

This material is based upon work partially supported by the National Science Foundation under grants CCF-1215985 and CMMI-1335301. The content reflects the views of the authors and not necessarily those of NSF.

REFERENCES

[1] H. Aziz, S. Gaspers, S. Mackenzie, N. Mattei, P. Stursberg, and T. Walsh. Fixing a balanced knockout tournament. In *Proc. AAAI*, 2014.

[2] Y. Bachrach, N. Betzler, and P. Faliszewski. Probabilistic possible winner determination. In *AAAI*, volume 10, pages 697–702, 2010.

[3] N. Betzler and B. Dorn. Towards a dichotomy of finding possible winners in elections based on scoring rules. In *Proc. Mathematical Foundations of Computer Science*, pages 124–136. Springer, 2009.

[4] P. Faliszewski, E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and control. *arXiv preprint arXiv:0809.4484*, 2008.

[5] K. Konczak and J. Lang. Voting procedures with incomplete preferences. In *Proc. IJCAI*, 2005.

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

[7] J. Levin and B. Nalebuff. An introduction to vote-counting schemes. *The Journal of Economic Perspectives*, 9(1):3–26, 1995.

[8] N. Mattei, J. Goldsmith, and A. Klapper. On the complexity of bribery and manipulation in tournaments with uncertain information. In *Proc. FLAIRS*, 2012.

[9] M. Mihail and P. Winkler. On the number of Eularian orientations of a graph. In *Proceedings of the third annual ACM-SIAM Symposium on Discrete Algorithms*, pages 138–145. Society for Industrial and Applied Mathematics, 1992.

[10] B. L. Schwartz. Possible winners in partially completed tournaments. *SIAM Review*, 8(3):302–308, 1966.

[11] T. Walsh. Uncertainty in preference elicitation and aggregation. In *Proc. AAAAI*, 2007.

[12] K. D. Wayne. A new property and a faster algorithm for baseball elimination. *SIAM Journal on Discrete Mathematics*, 14(2):223–229, 2001.

[13] L. Xia and V. Conitzer. Determining possible and necessary winners under common voting rules given partial orders. In *Proc. AAAI*, pages 196–201, 2008.

[14] L. Xia and V. Conitzer. Determining possible and necessary winners under common voting rules given partial orders. *Journal of Artificial Intelligence Research*, 41(2):25–67, 2011.

[15] E. Zermelo. Die Berechnung der Turnier-Ergebnisse als ein Maximumproblem der Wahrscheinlichkeitsrechnung. *Mathematische Zeitschrift*, 29(1):436–460, 1929.