Manipulation of Copeland Elections

Piotr Faliszewski
AGH University of Science and Technology, Poland
faliszew@agh.edu.pl

Edith Hemaspaandra
Rochester Institute of Technology, USA
eh@cs.rit.edu

Henning Schnoor
Christian-Albrechts-Universität Kiel, Germany
schnoor@ti.informatik.uni-kiel.de

ABSTRACT

We resolve an open problem regarding the complexity of unweighted coalitional manipulation, namely, the complexity of Copeland$^\alpha$-manipulation for $\alpha \in \{0, 1\}$. Copeland$^\alpha$, $0 \leq \alpha \leq 1$, is an election system where for each pair of candidates we check which one is preferred by more voters (i.e., we conduct a head-to-head majority contest) and we give one point to this candidate and zero to the other. However, in case of a tie both candidates receive α points. In the end, candidates with most points win. It is known [13] that Copeland$^\alpha$-manipulation is NP-complete for all rational α's in $(0, 1) - \{0.5\}$ (i.e., for all the reasonable cases except the three truly interesting ones). In this paper we show that the problem remains NP-complete for $\alpha \in \{0, 1\}$. In addition, we resolve the complexity of Copeland$^\alpha$-manipulation for each rational $\alpha \in [0, 1]$ for the case of irrational voters.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent Systems*

General Terms

Algorithms, Theory

Keywords

preferences, computational complexity, multiagent systems

1. INTRODUCTION

The complexity of manipulation in various voting systems is one of the most thoroughly studied topics in computational social choice. Tremendous progress achieved in the last few years have left only very few voting systems for which the complexity of coalitional manipulation is unknown. One such system is Copeland$^\alpha$, for $\alpha \in \{0, 0.5, 1\}$. The idea of Copeland$^\alpha$ elections is that for each pair of candidates we compare which one is preferred by more candidates and we give one point to that candidate and zero points to the other one. However, if the result of such a head-to-head contest is a tie, both candidates receive α points. Faliszewski, Hemaspaandra, and Schnoor [13] have shown that coalitional manipulation is NP-hard for rational α's in $(0, 1) - \{0.5\}$. However, exactly the cases of $\alpha \in \{0, 0.5, 1\}$ are the really interesting ones. For $\alpha \in \{0, 0.5\}$ we get exactly the two systems that are known under the name Copeland,[1] and $\alpha = 1$ gives a system that is sometimes called Llull. Copeland voting is named after A. H. Copeland who proposed the method over 50 years ago in a lecture [6], but the method is very natural and its variants were often reinvented throughout history (e.g., by Jech and by Zermelo), with the earliest record going 700 years back, to Ramon Llull, a 13th century mystic and philosopher (see, e.g., [18]).

In this paper we resolve the complexity of coalitional manipulation in Copeland^1 and in Copeland^0. In addition, we show why the proof approach of Faliszewski, Hemaspaandra, and Schnoor [13] could not have succeeded for these cases. We also derive the exact complexity of manipulation in Copeland$^\alpha$ for the case of irrational voters.

Manipulation in elections means that some voter (or, a group of voters) decides to vote in a way that does not reflect his or her (their) true preference, but that does guarantee an outcome of the election that he or she (they) prefer to the one that their true votes would yield. Unfortunately, the famous Gibbard-Satterthwaite theorem [17, 25] says that essentially every practically useful voting system sometimes creates incentives for voters to attempt manipulation. As a response to this depressing result, Bartholdi, Tovey, Trick, and Orlin [2, 1] suggested that even though manipulation might be possible in principle, for some election systems finding a successful manipulation might be computationally so expensive as to prevent any possibility of finding a manipulative vote, short of luckily guessing one. They have set off to find such voting rules for the case of a single manipulator and they indeed showed that so-called second-order Copeland is NP-hard to manipulate by a single voter (see [2]) and so is STV (see [1]).

The issue of manipulation in elections is particularly relevant for researchers working on multiagent systems. A situation where agents need to make a joint decision often arises in multiagent systems and voting is one of the most natural ways of making such decisions. For example, Ephrati and Rosenschein [10] suggest a way of using voting in multiagent planning problems, Ghosh et al. [16] use voting to develop a recommender system, and Dwork et al. [8] show how voting is useful in designing a metasearch engine for the web. However, voters—in particular when they are software agents—

Cite as: Manipulation of Copeland Elections, P. Faliszewski, E. Hemaspaandra, H. Schnoor, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 367-374

[1] When Copeland is mentioned without the α argument, it typically refers to Copeland$^{0.5}$, but in some papers it means Copeland^0 (see, e.g., [24] and an early version of [11]).

may be capable of systematic analysis of a given voting situation and might attempt manipulation (thus, skewing the result of the election). The fact that the manipulation problem is computationally hard for a given election system may either be discouraging enough for agents so they choose to act honestly, or may prevent them from succeeding in manipulation attempts.

Related work. Currently, the complexity of manipulating elections is one of the most throughly researched subareas of computational social choice, a field of study that focuses on computational properties of group decision making. For example, Conitzer, Lang and Sandholm [5] considered coalitional manipulation in the case where voters might be weighted, that is, when votes might be worth different amounts (a typical situation, e.g., when stockholders vote) and Hemaspaandra and Hemaspaandra [19] classified the complexity of such weighted manipulation for all scoring protocols, an important family of voting rules (see also [5, 23]). In a different line of research, Elkind and Lipmaa [9] and Conitzer and Sandholm [4] showed universal constructions that make manipulation computationally hard for every election system, but at the price of modifying the system and possibly losing its desirable properties. Recent works of Faliszewski, Hemaspaandra, and Schnoor [13], and Xia et al. [27] consider coalitional manipulation in unweighted settings. Regarding various types of affecting the result of Copeland$^\alpha$ elections, Faliszewski et al. [11] study the complexity of so-called election control (where one attempts to affect the result by, e.g., adding/deleting candidates/voters) and the complexity of bribery for Copeland$^\alpha$ elections.

Almost all of the papers mentioned in the previous paragraph, as well as this paper, focus on the worst-case complexity analysis of manipulation in elections. That is, for each voting system they consider, they ask whether a given variant of manipulating elections is in P or if it is NP-hard. Recent work raises some criticism to this approach. The works of Friedgut, Kalai, and Nisan [14], Xia and Conitzer [26], and Dobzinski and Procaccia [7] (as well as, to some degree, paper [23]) take the frequency-of-hardness approach and show that assuming impartial culture (i.e., assuming the voters choose their votes independently and uniformly at random) for many voting systems trivial manipulation algorithms (e.g., choosing a random vote) succeed with small, but nonnegligible probability. Zuckerman, Procaccia, and Rosenschein [28] and Brelsford et al. [3] provide efficient approximation algorithms for many NP-hard manipulation problems, and Faliszewski et al. [12] consider single-peaked domains and show that there manipulation can sometimes be easy, even though it is hard on the unrestricted domain. Nonetheless, most researchers agree that establishing the worst-case complexity of manipulation for a given voting rule is a natural, important step in understanding computational properties of the rule.

Organization. The paper is organized as follows. In Section 2 we provide basic preliminaries and notation. In Section 3 we describe our (standard) model of elections, describe Copeland voting in detail, and provide lemmas that allow us to derive instances of the manipulation problem for Copeland^0 and Copeland^1 conveniently. In Section 4 we give the proof of our main result, that is, that manipulation in Copeland$^\alpha$ is NP-complete for $\alpha \in \{0, 1\}$, and we mention how our proof can be generalized to other values of α. In Section 4 we also argue why the proof approach from [13] could not have succeeded for the case of $\alpha \in \{0, 1\}$. Finally, in Section 5 we consider manipulation in the so-called irrational voter model, and in Section 6 we provide conclusions.

2. PRELIMINARIES

We write $\mathbb{Z}$ to denote the set of integers, $\mathbb{Z}_0^+$ to denote the set of nonnegative integers, and $\mathbb{R}$ to denote the set of real numbers. For each two real numbers a, b, $a \leq b$, by (a, b) we mean the set $\{x \in \mathbb{R} \mid a < x < b\}$ and by $[a, b]$ the set $\{x \in \mathbb{R} \mid a \leq x \leq b\}$. For each $x \in \mathbb{R}$, by $|x|$ we mean the absolute value of x. For each finite set S, be $\|S\|$ we mean the cardinality of S.

Graphs. Throughout this paper we will often speak of instances of elections in terms of the graphs they induce (see Example 3.1 below). A directed graph G is a pair $(V(G), E(G))$, where $V(G)$ is the set of vertices of the graph, and $E(G)$ is the set of directed edges, that is, of ordered pairs of two distinct vertices. Later on, we will overload letters E and V to mean elections and collections of voters. Our functional notation for the sets of vertices and sets of edges of a graph allows us to distinguish when we mean, say, an election and when we mean a set of edges. The appropriate meaning will always be clear from context.

Computational complexity. We assume that the reader is familiar with standard notions of complexity theory, such as the complexity classes P and NP, polynomial-time many-one reductions, NP-completeness, and NP-hardness. The book of Papadimitriou [22] is a good reference on the aspects of complexity theory relevant to this paper.

Our NP-hardness proofs follow via reductions from one of the standard NP-complete problems, X3C (exact cover by 3-sets); see the standard text of Garey and Johnson [15].

Definition 2.1. *In an X3C instance we are given a set $B = \{b_1, \ldots, b_{3k}\}$ and a family $\mathcal{S} = \{S_1, \ldots, S_n\}$ of 3-elements subsets of B. We ask if it is possible to pick k sets from $\mathcal{S}$ such that their union is exactly B.*

We assume that all our decision problems (X3C, manipulation problems) are encoded using the alphabet $\Sigma = \{0, 1\}$ in a natural, efficient manner.

3. COPELAND AND MANIPULATION

We use the standard model of elections, i.e., for us an election is a pair $E = (C, V)$, where $C = \{c_1, \ldots, c_m\}$ is the set of candidates and $V = (v_1, \ldots, v_n)$ is the collection of voters. Each voter is represented via his or her *preference order* (sometimes also called *preference list*), i.e., a strict total order over C. Preference order $c_{i_1} > c_{i_2} > \cdots > c_{i_m}$ indicates that the voter believes c_{i_1} is the best candidate, then c_{i_2}, and so on, until c_{i_m}, who—in the eyes of this voter—is the worst choice.

Let $E = (C, V)$ be an election, where $C = \{c_1, \ldots, c_m\}$ and $V = (v_1, \ldots, v_n)$. For each two distinct $c_i, c_j \in C$ by $N_E(c_i, c_j)$ we mean the number of voters in V who prefer c_i to c_j.

Given an election, we need a way to aggregate the votes (the preference orders of the voters) and decide who is a winner. A voting rule is a function that takes as input an election $E = (C, V)$ and outputs a subset W of C, the set of candidates who tie as the winners. Note that it is completely legal for a voting rule to output an empty set, a set with a single candidate (the unique winner), or any other subset of candidates. In practice, of course, we would use some tie-resolution rule to resolve ties that arise in a voting rule, but from the point of view of this research, as is typical, we assume that all candidates in W are winners. This assumption is sometimes referred to as the nonunique-winner model.

3.1 Copeland Voting

There is an abundance of voting rules in the world but in this paper we will focus on just a single one (or, more precisely, on a single family of voting rules), namely, on Copeland$^\alpha$ voting. Let us fix a rational number α, $0 \leq \alpha \leq 1$.

Given an election $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$, for each $c_i \in C$, we define Copeland$^\alpha$ score of c_i as follows:

$$\begin{aligned} \mathrm{score}_E^\alpha(c_i) \quad &= \quad \|\{c_j \mid c_i \neq c_j \wedge N_E(c_i, c_j) > N_E(c_j, c_i)\}\| \\ &\quad + \alpha \|\{c_j \mid c_i \neq c_j \wedge N_E(c_i, c_j) = N_E(c_j, c_i)\}\| \end{aligned}$$

Copeland$^\alpha$ winners of election E are simply the candidates whose Copeland$^\alpha$ scores are highest. Note that the parameter α—the value of ties in head-to-head contests—only matters for elections with an even number of voters. If there are an odd number of voters then no head-to-head ties can happen and all the election systems in the Copeland$^\alpha$ family are equivalent.

As easily seen from the definition, we do not really need to know the actual votes to compute score_E^α for an election $E = (C, V)$: It suffices to know the values of $N_E(c_i, c_j)$ for each ordered pair of candidates $c_i, c_j \in C$. Given an election E, slightly abusing notation, we will often refer to the function $M_E(c_i, c_j) = N_E(c_i, c_j) - N_E(c_j, c_i)$ as the *weighted majority relation* and present M_E visually as in the next example.

Example 3.1. *Let us fix a rational α, $0 \leq \alpha \leq 1$ and let $E = (C, V)$ be an election where $C = \{c_1, c_2, c_3, c_4\}$ and $V = (v_1, v_2, v_3, v_4)$. The voters have the following preference orders:*

$$\begin{aligned} v_1 \quad &: \quad c_1 > c_2 > c_3 > c_4, \\ v_2 \quad &: \quad c_1 > c_3 > c_2 > c_4, \\ v_3 \quad &: \quad c_1 > c_4 > c_2 > c_3, \\ v_4 \quad &: \quad c_1 > c_2 > c_4 > c_3. \end{aligned}$$

c_1 is the unique winner of this election, with 3 Copeland$^\alpha$ points. Below we list the values of the weighted majority relation for E

$$\begin{array}{lll} M_E(c_1, c_2) = 4, & M_E(c_1, c_3) = 4, & M_E(c_1, c_4) = 4, \\ M_E(c_2, c_3) = 2, & M_E(c_2, c_4) = 2, & M_E(c_3, c_4) = 0. \end{array}$$

We can represent this weighted majority relation as a weighted directed graph, which we will call the election graph. The candidates are vertices and for each two candidates $c_i, c_j \in C$ we have a weight $M_E(c_i, c_j)$ edge from c_i to c_j if $M_E(c_i, c_j) \geq 0$ (note that if $M_E(c_i, c_j) = 0$ then we have two directed edges of weight 0, one from c_i to c_j and one in the opposite direction; we will refer to such pairs of edges as undirected edges). The election graph for our example is given in Figure 1. It is easy to read off the Copeland$^\alpha$ scores of the candidates from Figure 1. We have $\mathrm{score}_E^\alpha(c_1) = 3$, $\mathrm{score}_E^\alpha(c_2) = 2$, $\mathrm{score}_E^\alpha(c_3) = \mathrm{score}_E^\alpha(c_4) = \alpha$.

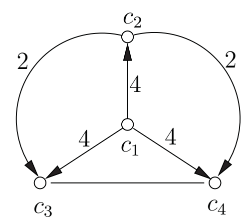

Figure 1: The election graph for the election from Example 3.1.

Given an election E, we write $G(E)$ to refer to its election graph (see the above example). Clearly, if we have an election $E = (C, V)$ it is easy to compute its election graph $G(E)$ and the weighted majority relation $M_E(\cdot, \cdot)$. In fact, it is also easy to go in the opposite direction. Given a candidate set $C = \{c_1, \dots, c_m\}$ and a set of even integers $\{m_{ij} \mid 1 \leq i, j \leq m \wedge i \neq j\}$ such that for each i, j ($1 \leq i, j \leq m, i \neq j$) it holds that $m_{ij} = -m_{ji}$, there is a collection V of $\sum_{1 \leq i \leq m} \sum_{1 \leq j \leq m, i \neq j} |m_{ij}|$ votes such that for each i, j, $1 \leq i, j \leq m, i \neq j$, it holds that $M_E(c_i, c_j) = m_{ij}$ (see [21]). In fact, this V is easily computable in polynomial time.

The fact that all weighted majority relations with even values can be implemented in votes is very useful. However, in our case, the next lemma (of Faliszewski et al. [11]) will be even more convenient.

Lemma 3.2 (Faliszewski et al. [11]). *Let $E = (C, V)$ be an election where $C = \{c_1, \dots, c_{n'}\}$, let α be a rational number such that $0 \leq \alpha \leq 1$, and let $n \geq n'$ be an integer. For each candidate c_i we denote the number of head-to-head ties of c_i in E by t_i. Let q be a positive integer and let $k_1, \dots, k_{n'}$ be a sequence of nonnegative integers such that for each k_i we have $0 \leq k_i \leq n^q$. There is an algorithm that in polynomial time in n outputs an election $E' = (C', V')$ such that:*

1. *$C' = C \cup D$, where $D = \{d_1, \dots, d_{2n'n^q}\}$,*
2. *E' restricted to C is E (that is, $G(E')$ restricted to C is $G(E)$),*
3. *the only ties in head-to-head contests in E' are between candidates in C,*
4. *for each i, $1 \leq i \leq n'$, $\mathrm{score}_{E'}^\alpha(c_i) = 2n'n^q - k_i + t_i\alpha$, and*
5. *for each i, $1 \leq i \leq 2n'n^q$, $\mathrm{score}_{E'}^\alpha(d_i) \leq n'n^q + 1$.*

3.2 Coalitional Manipulation for Copeland

Let us now formally define the (constructive) coalitional manipulation problem.

Definition 3.3. *Let R be a voting rule. In the R-manipulation problem we are given an election $E = (C, V)$, where $C = \{p, c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$, each voter with a preference order, and a list of voters $W = (w_1, \dots, w_k)$ without preassigned preference orders. We ask if it is possible to set the preference orders of the voters in W so that p is a winner of election $E' = (C, V + W)$, where $V + W$ is a concatenation of the lists V and W.*

We typically refer to voters in W as *the manipulative voters* or *the manipulators*. For the case of Copeland$^\alpha$, $0 \leq \alpha \leq 1$, it is known that Copeland$^\alpha$-manipulation is in P if there is only a single manipulator [2] but that the problems becomes NP-complete if $\alpha \in (0,1) - \{0.5\}$ and there are two manipulators [13]. We will show that for $\alpha \in \{0,1\}$ the two-manipulator case is also NP-complete. From now on, we will focus on the setting with exactly two manipulators.

As already suggested in the discussion before Lemma 3.2, it is not convenient to build Copeland$^\alpha$-manipulation instances vote by vote and a more robust mechanism is necessary. The next lemma gives such a mechanism.

DEFINITION 3.4. *A partial election graph is an election graph that (a) may contain pairs of vertices with no edge between them*[2]*, and (b) in addition to weights on edges, also has integer weights on the vertices (positive, negative, or zero).*

LEMMA 3.5. *Let us fix $\alpha \in \{0,1\}$. Let G be a partial election graph, let s be the weight function for its vertices (i.e., $s\colon V(G) \to \mathbb{Z}$), let m be the weight function for its edges (i.e., $m\colon E(G) \to \mathbb{Z}_0^+$), and for each $e \in E(G)$, let $m(e)$ be even. Let p be a candidate not in $V(G)$.*

There is a polynomial-time algorithm that given G, s, m, and p computes an instance (E, W, p) of Copeland$^\alpha$-manipulation where $E = (C, V)$, such that:

1. *$V(G) \cup \{p\} \subseteq C$, $W = (w_1, w_2)$,*
2. *for each $e \in E(G)$, $e = (c_i, c_j)$, $M_E(c_i, c_j) = m(e)$,*
3. *for each $c \in V(G)$, $score_E^\alpha(c) = score_E^\alpha(p) - s(c)$,*
4. *no candidate in $C - (V(G) \cup \{p\})$ is a winner of E, and none of them can become a winner after adding the votes of w_1 and w_2 (irrespective of what preference orders w_1 and w_2 pick),*
5. *for each two $c_i, c_j \in V(G)$, if there is no edge between c_i, c_j in $E(G)$ (i.e., $(c_i, c_j) \notin E(G)$ and $(c_j, c_i) \notin E(G)$) then $|M_E(c_i, c_j)| \geq 4$, and*
6. *irrespective of what votes the manipulators cast, p's score does not change.*

In short, the above lemma allows us to construct instances of Copeland$^\alpha$-manipulation, $\alpha \in \{0,1\}$, via fiddling with gadgets in partial election graphs. The lemma guarantees that we can translate partial election graphs to naturally corresponding manipulation instances.

Due to space constraints, we will not give the (easy) proof of Lemma 3.5, and instead we will just mention that the proof follows via a somewhat careful application of Lemma 3.2. In essence, we have to convert the partial election graph to an election graph (via adding arbitrarily directed edges, each with an even weight greater than 2, e.g., 4), compute how many dummy candidates we need, and convert vertex-weights to scores.

We mention in passing that our statement of Lemma 3.5 above is not the most general one. It is in fact possible to extend it any number of manipulators and any rational α, $0 \leq \alpha \leq 1$. However, the proofs presented in this paper only require the stated version of the lemma, and the more general version is more tricky to prove.

[2]Note that in an election graph there is at least one directed edge between any two vertices.

4. THE RESULTS

The main result of this paper is the following theorem.

THEOREM 4.1. *Copeland$^\alpha$-manipulation, $\alpha \in \{0,1\}$, is* NP*-complete, even for the case of two manipulators.*

Before we proceed with the proof, let us briefly discuss the proof of a similar theorem of Faliszewski, Hemaspaandra, and Schnoor [13], which shows that Copeland$^\alpha$-manipulation is NP-complete for all rational α's in $(0,1) - \{0.5\}$, even in the case of two manipulators.

In essence, Faliszewski, Hemaspaandra, and Schnoor build a partial election graph where the only head-to-head contests that can change are of the form "candidate c_i defeating candidate c_j by 2 votes" and the manipulators can either choose to turn these results into ties (i.e., in terms of election graphs, turn a $c_i \to c_j$ edge of weight 2 into a $c_i - c_j$ tie-edge of weight 0) or leave them as they are. The tricky part is that these choices are not independent and, via a clever choice of scores of various candidates, there is a limit on the number of swaps that the manipulators can make.[3] The "swaps of head-to-head victories to ties" are linked in such a way as to encode an instance of X3C or a certain variant of the satisfiability problem (depending whether $\alpha < 0.5$ or $\alpha > 0.5$). Faliszewski, Hemaspaandra, and Schnoor [13] convert their partial election graph into an instance of manipulation via a construction similar to our Lemma 3.5 (however, they do not explicitly state the lemma but rather describe how such a lemma can be derived). We will now show that constructions like theirs cannot possibly prove NP-hardness of manipulation in Copeland^0 and Copeland^1, unless P = NP.

THEOREM 4.2. *Let us fix $\alpha \in \{0,1\}$. There is a polynomial-time algorithm that solves Copeland$^\alpha$-manipulation on instances (E, W, p), where $E = (C, V)$, such that (a) $\|W\| = 2$, (b) either for each $c_i, c_j \in C - \{p\}$ it holds that $|M_E(c_i, c_j)| \geq 2$ or for each $c_i, c_j \in C - \{p\}$ it holds that $M_E(c_i, c_j) = 0 \vee |M_E(c_i, c_j)| > 2$.*

PROOF. Let us fix $\alpha \in \{0,1\}$ and let (E, W, p), where $E = (C, V)$, be an instance of Copeland$^\alpha$-manipulation. We will now give our algorithm.

Let us rename the candidates in C so that $C = \{p, c_1, \ldots, c_m\}$ and let us assume that $W = (w_1, w_2)$. Our algorithm, of course, always chooses that w_1 and w_2 rank p first. It remains to decide how to rank $c_1, \ldots, c_m$.

We consider four cases. First, assume $\alpha = 0$ and that for each $c_i, c_j \in C - \{p\}$ it holds that $M_E(c_i, c_j) = 0 \vee |M_E(c_i, c_j)| > 2$. In this case we set w_1's preference order to $p > C - \{p\}$ and w_2's preference order to $p > \overleftarrow{C - \{p\}}$ (i.e., both manipulators rank p first, w_1 ranks the remaining candidates in an arbitrary order, and w_2 ranks the remaining candidates in exactly the reverse of that order). We accept if this makes p a winner of election $(C, V + W)$ and reject otherwise. This choice of w_1's and w_2's preference orders is optimal because, in the $\alpha = 0$ case, the more ties among the candidates in $C - \{p\}$ there are, the better p's situation is.

[3]To be strictly correct, in their $\alpha \in (0.5, 1)$ part of the proof, Faliszewski, Hemaspaandra, and Schnoor for several candidates also leave the possibility of turning a tie into a victory, but there are only four such candidates. It is possible to truth-table reduce their construction to a scenario where there are no head-to-head ties that the manipulators can turn into head-to-head victories.

By assumption, it is impossible to introduce any more new ties via the choice of preference orders of w_1 and w_2, so we at least maintain the ones that already exist. (Also, since both manipulators rank p first, p gets as many extra points as possible). By an analogous argument, if $\alpha = 1$ and for each $c_i, c_j \in C - \{p\}$ it holds that $|M_E(c_i, c_j)| \geq 2$, we also set w_1's preference order to $p > C - \{p\}$ and w_2's preference order to $p > \overleftarrow{C - \{p\}}$ (in this case the manipulators are best off if no head-to-head ties among $C - \{p\}$ are introduced).

Let us now consider the case where $\alpha = 0$ and for each $c_i, c_j \in C - \{p\}$ it holds that $|M_E(c_i, c_j)| \geq 2$. Informally put, in this case the manipulators want to maximize the number of head-to-head ties. We will show that the manipulators can safely choose to cast identical votes, which means that we can compute their votes in polynomial time [2].

Assume that we have set preference orders of w_1 and w_2 in such a way that p is a winner of $(C, V + W)$. Let $W' = (w'_1, w'_2)$ be two voters that both have a preference order identical to that of w_1. We claim that p is a winner of $(C, V + W')$: For any two $c_i, c_j \in C - \{p\}$, if $M_{(C,V+W)}(c_i, c_j) = 0$ then $M_{(C,V+W')}(c_i, c_j) = 0$ as well. This is so because $M_{(C,V+W)}(c_i, c_j) = 0$ if and only if $M_E(c_i, c_j) = 2$ and both w_1 and w_2 prefer c_j to c_i, however for any such c_i, c_j, by definition, both w'_1 and w'_2 prefer c_j to c_i as well. This implies that the set of pairs of candidates that tie their head-to-head contests in $(C, V + W')$ is a superset of that set for $(C, V + W)$ and so, if p was a winner in $(C, V + W)$ then he or she certainly is a winner in $(C, V + W')$.

Now that we know that both w_1 and w_2 can safely cast identical votes, it is easy to find appropriate preference order for them by running the greedy manipulation algorithm of Bartholdi, Tovey, and Trick [2].

The case where $\alpha = 1$ and for each $c_i, c_j \in C - \{p\}$ it holds that $M_E(c_i, c_j) = 0 \vee |M_E(c_i, c_j)| > 2$ can be handled in the analogous manner as the previous case. The proof is completed. $\square$

The point of Theorem 4.2 is not necessarily to give a practically useful algorithm for manipulation (even though it succeeds in this task for the limited cases it considers) but rather to point us in the direction of a proof that does handle the cases of $\alpha \in \{0, 1\}$. In particular, by Theorem 4.2 we know that a reduction of the two-manipulator case with $\alpha \in \{0, 1\}$ needs to have the property that it constructs instances (E, W, p), $E = (C, V)$, of Copeland$^\alpha$-manipulation where for some $c_i, c_j \in C$ we have $M_E(c_i, c_j) = 2$ and for some other $c_i, c_j \in C$ we have $M_E(c_i, c_j) = 0$. Otherwise the constructed instances could be solved in polynomial time.

The remainder of this section is devoted to the proof and the discussion of the proof of Theorem 4.1. We first handle $\alpha = 1$, then $\alpha = 0$, and then discuss how our proofs can be generalized.

4.1 Tie-Value One

We now prove that Copeland^1-manipulation is NP-complete.

Lemma 4.3. *Copeland^1-manipulation is* NP-*complete.*

Proof. It is easy to see that the problem is in NP. To show NP-hardness we give a reduction from X3C. Let $(B, \mathcal{S})$ be an instance of X3C, where $B = \{b_1, \ldots, b_{3k}\}$ and $\mathcal{S} = \{S_1, \ldots, S_n\}$. Our goal is to build an instance (E, W, p), $E = (C, V)$, $\|W\| = 2$, of Copeland^1-manipulation such that p can become a winner of $(C, V + W)$ if and only if $(B, \mathcal{S})$ is a "yes"-instance of X3C. We will do so via Lemma 3.5, and so we first derive our partial election graph G. In our description of G below we use the same notation as in the statement of Lemma 3.5 (in particular, recall that the function s gives weights of candidates that translate to the difference of points between the candidate in question and the preferred one, and function m gives weights of edges).

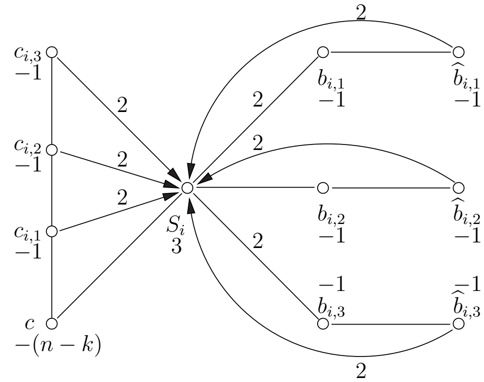

Figure 2: The gadget for set $S_i = \{b_{i,1}, b_{i,2}, b_{i,3}\}$ in the proof of Lemma 4.3. As in Figure 1, numbers next to directed edges indicate the value of the weighted majority relation. The numbers next to vertices indicate the value of the $s(\cdot)$ function.

The partial election graph G involves the following candidates.

1. We have a single candidate c, $s(c) = -(n - k)$.
2. For each i, $1 \leq i \leq n$, we have candidate S_i, $s(S_i) = 3$.
3. For each i, j, $1 \leq i \leq n$, $1 \leq j \leq 3$, we have candidate $c_{i,j}$ with $s(c_{i,j}) = -1$.
4. For each j, $1 \leq j \leq 3k$, we have candidate b_j with $s(b_j) = -1$ and we have candidate $\widehat{b}_j$ with $s(\widehat{b}_j) = -1$.

Let $\widehat{B} = \{\widehat{b}_1, \ldots, \widehat{b}_{3k}\}$ and $K = \{c_{i,j} \mid 1 \leq i \leq n, 1 \leq j \leq 3\}$.

Before we proceed with the rest of the proof, we should warn the reader about the notation we use throughout the proof. Symbols such as, e.g., S_i, $1 \leq i \leq n$, here may refer to two entities. S_i might either refer to a set (in the world of our input X3C instance) or may refer to a candidate (in the world of the manipulation instance that we build). We always make sure that it is clear from context which meaning we have in mind. Occasionally, when describing preference orders we will include names of sets. For example, we might write $S_i > c$ or $\mathcal{S} > c$. We set the following convention: Whenever some S_i appears in a preference order, it refers to *the candidate* S_i appearing in the preference order. Whenever $\mathcal{S}, B, \widehat{B}$, or K (or some subset of either of these sets) appears in a preference order, it refers to listing *the members* of the set in the preference order in an arbitrary order.

The partial election graph consists of a bunch of gadgets, one for each set S_i. Figure 2 presents the gadget for S_i, $1 \leq i \leq n$. With our partial election graph in hand, we simply invoke Lemma 3.5 to obtain instance (E, W, p) of Copeland^1-manipulation, with $W = (w_1, w_2)$, and where p can become a winner if and only if it is possible to choose the preference orders for w_1 and w_2 such that for each candidate d in the partial election graph:

1. If $s(d) \geq 0$ then the score of d increases by at most $s(d)$ points, and
2. if $s(d) < 0$ then the score of d decreases be at least $-s(d)$ points.

This reduction clearly works in polynomial time and it remains to show that it is correct.

However, before jumping into the proof, let us explain intuitively how our gadgets work. Let us fix an integer i, $1 \leq i \leq n$, and let $S_i = \{b_{i,1}, b_{i,2}, b_{i,3}\}$. Note that if p is to become a winner, each candidate $c_{i,1}, c_{i,2}, c_{i,3}$ has to lose one point. This happens only if both of the manipulators rank $c > c_{i,1} > c_{i,2} > c_{i,3}$. Thus, we can assume that both manipulators rank candidates in $K \cup \{c\}$ in this way. This implies that if both manipulators rank $S_i > c$ then S_i's score increases by 3, but c's score decreases by 1. For p to become a winner, c's score has to go down by $(n-k)$ points.

Let us fix an integer $j \in \{1,2,3\}$. If p is to become a winner, then the scores of each of $b_{i,j}, \widehat{b}_{i,j}$ have to go down by one point. This happens, e.g., if both manipulators rank $S_i > b_{i,j} > \widehat{b}_{i,j}$. However, such ranking increases the score of S_i by one. Thus, for each set S_i, the manipulators have a choice. They can either rank $S_i > c$ and decrease c's score, or, for each $j \in \{1,2,3\}$, they can rank $S_i > b_{i,j} > \widehat{b}_{i,j}$, but they cannot do both, because then S_i's score goes up by more than 3 and p is not a winner.

Let us now formally prove that the reduction is correct. First, let us assume that $(B, \mathcal{S})$ is a "yes" instance. Let $\widehat{\mathcal{S}}_c$ be a subset of $\mathcal{S}$ that constitutes an exact cover of B, and let K_c be a set containing those $c_{i,j} \in K$ for which $S_i \in \widehat{\mathcal{S}}_c$. We can verify that if the manipulators cast the following votes then p indeed becomes a winner (that is, all scores that need to go down, do go down by an appropriate amount and all scores that can go up, go up by no more than an allowable amount; as described by function $s(\cdot)$):

$$\begin{aligned} w_1 &: \mathcal{S} - \widehat{\mathcal{S}}_c > c > K > \mathcal{S}_c > B > \widehat{B} \\ w_2 &: \mathcal{S}_c > B > \widehat{B} > \mathcal{S} - \widehat{\mathcal{S}}_c > c > K \end{aligned}$$

For the opposite direction, assume that w_1 and w_2 have preference orders such that p is a winner of election $E' = (C, V+W)$. Let us assume that the following claim holds.

Claim 1. *Let w_1 and w_2 have preference orders such that p is a winner of $(C, V+W)$. For each i, $1 \leq i \leq n$, and each $b_j \in S_i$, it is not the case that both w_1 and w_2 rank $S_i > c$ and $S_i > b_j$.*

We show how to complete the proof of the theorem, using Claim 1. For at least $n-k$ candidates S_i in $\mathcal{S}$, it holds that both w_1 and w_2 rank $S_i > c$. This is so because, compared to E, in E' candidate c has to have at least $n-k$ points less. The only way for c to lose a point is to have both manipulators rank some candidate S_i as more preferred than c. However, if p is a winner, then also each member of B has to lose a point. That is, for each b_j, $1 \leq j \leq 3k$, there is a candidate S_i such that (a) $b_j \in S_i$, and (b) both manipulators rank $S_i > b_j$. Due to Claim 1, there are only k candidates S_i left for this task. Since there are $3k$ candidates in B, the candidates S_i who are not ranked as preferred to c by both manipulators must form an exact cover of B.

To complete the proof, it remains to prove Claim 1. Let us assume that preference orders of w_1 and w_2 are set so that p is a winner of $(C, V+W)$ and yet there exist two candidates, S_i and b_j, such that (a) both manipulators rank $S_i > c$ and $S_i > b_j$, and (b) b_j is a member of the set S_i. However, having both manipulators rank $S_i > c$ means that S_i's score goes up by 3 (compared to E) and having both manipulators rank $S_i > b_j$ means that S_i's score goes up by one additional point. Thus, if p is to be a winner, the remaining parts of the manipulators' preference orders must ensure that S_i's score goes down by at least one point. However, this can only happen if there is some candidate $b_{j'} \in B$ such that both manipulators rank $b_{j'} > S_i$. Since $b_{j'}$'s score (by the construction of our manipulation instance) also needs to go down by one point, there is some candidate $S_{i'}$ such that (a) both manipulators rank $S_{i'} > b_{j'}$, and (b) $b_{j'}$ is a member of the set $S_{i'}$. Thus, by transitivity of preference orders, we now have that both manipulators rank $S_{i'} > b_{j'} > S_i > b_i$ and $S_{i'} > c$. That is, now we have the problem that $S_{i'}$ has gained 4 points as opposed to his or her score in E while he or she is only allowed to gain at most 3. Via repeating the same reasoning for $S_{i'}$ as we have just done for S_i (and further candidates in $\mathcal{S}$) we, eventually, reach the conclusion that both manipulators rank $S_i > S_i$, which is a contradiction. Thus, the claim is proved and so is the whole lemma. ☐

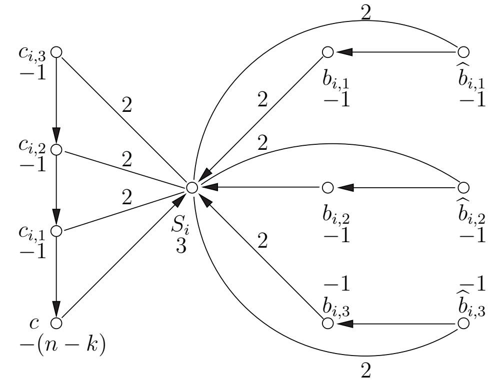

Figure 3: The gadget for set $S_i = \{b_{i,1}, b_{i,2}, b_{i,3}\}$ for Copeland^0. The notation is the same as in Figure 2.

4.2 Tie-Value Zero

Let us now turn to Copeland^0.

Lemma 4.4. *Copeland^0-manipulation is* NP*-complete.*

The proof of Lemma 4.4 is very similar to the proof of Lemma 4.3. Thus, instead of giving the formal proof let us simply discuss the differences. We maintain the same notation as in the proof of Lemma 4.3. Our reduction for Copeland^0 is identical to the one employed for Copeland^1 except that now, given a set $S_i = \{b_{i,1}, b_{i,2}, b_{i,3}\}$ we use a gadget with "reversed" edges (i.e., directed edges turn into head-to-head ties and head-to-head ties turn into edges directed in the appropriate way). The gadget for a set S_i is given in Figure 3. The proof that the modified reduction is correct is, in essence, identical to the correctness proof for Lemma 4.3.

4.3 Generalizations

The construction given in the proofs of Lemmas 4.3 and 4.4 is quite different from the proofs given by Faliszewski, Hemaspaandra, and Schnoor [13] for Copeland$^\alpha$-manipulation for $\alpha \in (0,1) - \{0.5\}$. However, it turns out that our constructions here can be extended to cover the cases of rational α, $\alpha \in [0,1] - \{0.5\}$. The general structure of our proofs remains the same, only we need to somewhat extend our gadgets. Due to space constraints we will not give the extended proofs.

5. IRRATIONAL VOTERS

Up to now, throughout the paper, we have considered only the standard model of elections, where the voters are assumed to be rational (i.e., are assumed to have preferences that can be represented as strict, linear orders over the candidate set). Let us now deviate from this and consider an interesting preference model, where voters are allowed to be irrational. Here, irrationality does not refer to any sort of deficiency on the side of the voters, but simply means that instead of providing a single, transitive preference order, for each pair of candidates, say c and d, the voter specifies whether he or she prefers c ($c > d$) or d ($d > c$). In this model, given three candidates, c, d, e, the voter can legally specify $c > d$, $d > e$, and $e > c$. Thus, the preferences of an irrational voter are described not by a preference order, but by a *preference table* which specifies for each pair of candidates which one, given the choice among the two, the voter prefers. Irrational voters, in the context of computational social choice, were introduced in paper [11].

The irrational voter model, in spite of its discouraging name, is quite natural. Irrational preference tables arise, e.g., if a voter makes a decision according to multiple criteria, or if each voter itself represents a result of voting (compare with the election graph in Section 3.1).

For each rational α, $0 \leq \alpha \leq 1$, the definition of Copeland$^\alpha$ naturally extends to irrational voters. We can similarly extend the definition of the manipulation problem.

Definition 5.1. *Let R be a voting rule defined for irrational voters. In the R-irrational-manipulation problem we are given an election $E = (C, V)$, where $C = \{p, c_1, \ldots, c_m\}$ and $V = (v_1, \ldots, v_n)$ (each voter with a preference table), and a list of voters $W = (w_1, \ldots, w_k)$ without preassigned preference tables. We ask if it is possible to set the preference tables of voters in W so that p is a winner of election $E' = (C, V + W)$.*

For the case of irrational voters, manipulation in Copeland$^\alpha$, $0 \leq \alpha \leq 1$ shows some interesting behavior, rather different from the one that we know of (expect of) for the rational case.

Theorem 5.2. *Let α be a rational number, $0 \leq \alpha \leq 1$. Copeland$^\alpha$-irrational-manipulation is in* P *if $\alpha \in \{0, 0.5, 1\}$ and is* NP*-complete otherwise.*

The proof relies on three results already known in the literature, and what we focus on is adapting these three results to work together. Because of this, and due to space constraints, we will discuss the proof informally instead of giving full details.

First, let us consider the case of $\alpha \in \{0, 0.5, 1\}$. Faliszewski et al. [11] have shown the following result: Given an irrational Copeland$^\alpha$ election with $\alpha \in \{0, 0.5, 1\}$, if we are allowed to switch each entry in each voter's preference table at unit cost, then finding a lowest-cost set of switches that ensure that some given candidate p is a winner is solvable in polynomial time. (They called the problem of finding the right switches *microbribery*.) In fact, we can easily modify their algorithm to also work for the case where each switch in each preference table is associated with its own cost (possibly bigger than 1, possibly 0). With this result at hand, it is easy to solve Copeland$^\alpha$-irrational-manipulation instances, $\alpha \in \{0, 0.5, 1\}$, by assigning prohibitively large costs to switching entries in the preference tables of nonmanipulators, assigning the cost 0 of switching the entries in the manipulators' preference tables, and then running the modified microbribery algorithm.

Let us now turn to the proof that for each rational α, $\alpha \in (0,1) - \{0.5\}$, Copeland$^\alpha$-irrational-manipulation is NP-complete. Let us fix one such α. The membership of Copeland$^\alpha$-irrational-manipulation in NP is obvious, so we focus on NP-hardness. Kern and Paulusma [20] study the following problem, which in our terms would be called $SC(0, \alpha, 1)$. In this problem we are given an undirected graph, $G = (V(G), E(G))$ and a function c such that for each $v \in V$, $c(v)$ is of the form $i + j\alpha$, for some integers i, j. The question is if we can orient (some of) the edges of G so that the following constraint is satisfied: For each vertex v the number of edges directed from v to some other vertex v' plus α times the number of undirected edges adjacent to v is at most $c(v)$. While Kern and Paulusma defined their problem to model sport tournaments, given the connection between election graphs and Copeland$^\alpha$ voting, it is easy to see that $SC(0, \alpha, 1)$ is, almost, Copeland$^\alpha$-irrational-manipulation.

We could reduce $SC(0, \alpha, 1)$ to Copeland$^\alpha$-irrational-manipulation as follows. Given input $G = (V, E), c$ for $SC(0, \alpha, 1)$, we form an instance (E, W, p) of the irrational-manipulation problem where $W = (w_1, w_2)$, election $E = (C, V)$, where $C = V(G) \cup \{p\}$, and V is set so that for each edge $\{v, v'\}$ in $E(G)$, the result of the head-to-head contest among candidates v, v' in E is a tie. We set the remaining results of head-to-head contests in such a way that if both w_1 and w_2 indicate in their preference tables that they prefer p to every other candidate then p is a winner if and only if it is possible to set the remaining parts of preference tables of w_1 and w_2 such that no candidate $v \in V(G)$ obtains more than $c(v)$ points from the head-to-head contests among the candidates to which v is adjacent in G. This reduction would be correct, if it was possible to construct a set of voters V that satisfies the requirements of the reduction. Unfortunately, this seems impossible.

Nonetheless, it is possible to save the idea of the reduction. Faliszewski, Hemaspaandra, and Schnoor [13] show how to build instances of (rational) manipulation with voter sets V that do satisfy the requirements of our reduction. We have verified that the construction of Faliszewski, Hemaspaandra, and Schnoor [13] also works for the case of irrational-manipulation. Thus, we can give a correct reduction from $SC(0, \alpha, 1)$. This completes this (rough) sketch of the proof of Theorem 5.2.

6. CONCLUSIONS

We have established the complexity of the constructive coalitional manipulation problem for Copeland^0 and Copeland^1. As a result, the only variant of Copeland voting for which the complexity of this problem remains unknown is Copeland$^{0.5}$.

Our proofs have several nice properties. For example, the reduction that we use for Copeland^0 (or, more specifically, the graph that we build in this reduction) is, in a sense, a mirror image of the reduction we have devised for Copeland^1 (or, more specifically, of the graph that we build in that reduction). More importantly, our reductions can be modified (though, due to space constraints, we have not described this precisely) to prove that Copeland$^\alpha$-manipulation is NP-complete for each rational α in $[0, 1] - \{0.5\}$, hence our proofs also allow to obtain the results of [13].

Finally, we have shown a full dichotomy theorem for the complexity of Copeland$^\alpha$ manipulation for the case of irrational voters. It is quite interesting that for the irrational voter case, Copeland$^\alpha$ is in P for $\alpha \in \{0, 0.5, 1\}$, and in the rational case it is either NP-complete ($\alpha \in \{0, 1\}$) or we suspect it to be NP-complete ($\alpha = 0.5$).

Acknowledgments. We would like to thank the anonymous AAMAS referees. Piotr Faliszewski was supported in part by AGH University of Science and Technology grant 11.11.120.865, by Polish Ministry of Science and Higher Education grant N-N206-378637, and by the Foundation for Polish Science's Homing/Powroty. Edith Hemaspaandra was supported in part by grant NSF-IIS-0713061 and by a Friedrich Wilhelm Bessel Research Award.

7. REFERENCES

[1] J. Bartholdi, III and J. Orlin. Single transferable vote resists strategic voting. *Social Choice and Welfare*, 8(4):341–354, 1991.

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

[3] E. Brelsford, P. Faliszewski, E. Hemaspaandra, H. Schnoor, and I. Schnoor. Approximability of manipulating elections. In *Proceedings of AAAI-08*, pages 44–49, July 2008.

[4] V. Conitzer and T. Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proceedings of IJCAI-03*, pages 781–788, Aug. 2003.

[5] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):1–33, 2007.

[6] A. Copeland. A "reasonable" social welfare function. Seminar on Mathematics in Social Sciences, University of Michigan, 1951.

[7] S. Dobzinski and A. Procaccia. Frequent manipulability of elections: The case of two voters. In *Proceedings of WINE-08*, pages 653–664, Dec. 2008.

[8] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th International World Wide Web Conference*, pages 613–622, Mar. 2001.

[9] E. Elkind and H. Lipmaa. Small coalitions cannot manipulate voting. In *Proceedings of the 9th International Conference on Financial Cryptography and Data Security*, pages 285–297, Feb./Mar. 2005.

[10] E. Ephrati and J. Rosenschein. A heuristic technique for multi-agent planning. *Annals of Mathematics and Artificial Intelligence*, 20(1–4):13–67, 1997.

[11] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and constructive control. *JAIR*, 35:275–341, 2009.

[12] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. In *Proceedings of TARK-09*, pages 118–127, July 2009.

[13] P. Faliszewski, E. Hemaspaandra, and H. Schnoor. Copeland voting: Ties matter. In *Proceedings of AAMAS-08*, pages 983–990, May 2008.

[14] E. Friedgut, G. Kalai, and N. Nisan. Elections can be manipulated often. In *Proceedings of FOCS-08*, pages 243–249, Oct. 2008.

[15] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.

[16] S. Ghosh, M. Mundhe, K. Hernandez, and S. Sen. Voting for movies: The anatomy of recommender systems. In *Proceedings of the 3rd Annual Conference on Autonomous Agents*, pages 434–435, 1999.

[17] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41(4):587–601, 1973.

[18] G. Hägele and F. Pukelsheim. The electoral writings of Ramon Llull. *Studia Lulliana*, 41(97):3–38, 2001.

[19] E. Hemaspaandra and L. Hemaspaandra. Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1):73–83, 2007.

[20] W. Kern and D. Paulusma. The new FIFA rules are hard: Complexity aspects of sports competitions. *Discrete Applied Mathematics*, 108(3):317–323, 2001.

[21] D. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.

[22] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[23] A. Procaccia and J. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *JAIR*, 28:157–181, 2007.

[24] A. Procaccia, J. Rosenschein, and G. Kaminka. On the robustness of preference aggregation in noisy environments. In *Proceedings of AAMAS-07*, pages 416–422, May 2007.

[25] M. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10(2):187–217, 1975.

[26] L. Xia and V. Conitzer. A sufficient condition for voting rules to be frequently manipulable. In *Proceedings of EC-09*, pages 99–108, July 2008.

[27] L. Xia, V. Conitzer, A. Procaccia, and J. Rosenschein. Complexity of unweighted manipulation under some common voting rules. In *Proceedings of IJCAI-09*, pages 348–353, July 2009.

[28] M. Zuckerman, A. Procaccia, and J. Rosenschein. Algorithms for the coalitional manipulation problem. *Artificial Intelligence*, 173(2):392–412, 2009.