Copeland Voting: Ties Matter

Piotr Faliszewski
University of Rochester
Rochester, NY 14627, USA
pfali@cs.rochester.edu

Edith Hemaspaandra
Rochester Institute of Technology
Rochester, NY 14623, USA
eh@cs.rit.edu

Henning Schnoor
Rochester Institute of Technology
Rochester, NY 14623, USA
hs@cs.rit.edu

ABSTRACT

We study the complexity of manipulation for a family of election systems derived from Copeland voting via introducing a parameter α that describes how ties in head-to-head contests are valued. We show that the thus obtained problem of manipulation for unweighted Copeland$^{\alpha}$ elections is NP-complete even if the size of the manipulating coalition is limited to two. Our result holds for all rational values of α such that $0 < \alpha < 1$ except for $\alpha = \frac{1}{2}$. Since it is well known that manipulation via a single voter is easy for Copeland, our result is the first one where an election system originally known to be vulnerable to manipulation via a single voter is shown to be resistant to manipulation via a coalition of a constant number of voters. We also study the complexity of manipulation for Copeland$^{\alpha}$ for the case of a constant number of candidates. We show that here the exact complexity of manipulation often depends closely on the α: Depending on whether we try to make our favorite candidate a winner or a unique winner and whether α is 0, 1 or between these values, the problem of weighted manipulation for Copeland$^{\alpha}$ with three candidates is either in P or is NP-complete. Our results show that ways in which ties are treated in an election system, here Copeland voting, can be crucial to establishing complexity results for this system.

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

Every democratic society faces the problem of making decisions that accommodate the needs, desires, and goals of all its members, or, at least, making decisions that take the needs of all its members into account. A natural way to arrive at such decisions is to perform an election where members of the group express their preference regarding possible alternatives and some formalized procedure, an election system, is used to aggregate these preferences and declare which alternative wins. Traditionally we think of elections in terms of human democracies but artificial societies of software agents use elections as well. Examples of elections in multiagent systems include planning scenarios where agents vote on the next step of the plan [11], aggregating search results for the web [9], and others. We point out that often the voters in an election are weighted, that is, have different voting powers.

Unfortunately, all reasonable deterministic election systems share the problem that there always exists a scenario where the voters are motivated to vote dishonestly, i.e., differently than their true preference. Consider candidate set $\{a, b, c\}$. If a voter prefers a to b and b to c, but knows that a has no chances of winning, the voter might vote for b instead. A fundamental theorem of Gibbard and Satterthwaite shows that this problem cannot be avoided: Any election system of practical value sometimes gives incentive to vote insincerely.

A situation where a coalition of voters (or a single voter) cast their votes strategically, that is, in such a way that guarantees the best outcome from their point of view, is called manipulation. Via Gibbard-Satterthwaite Theorem we know that theoretical possibility of manipulation is unavoidable, but we can at least hope to make it as hard as possible. Particularly, in elections among software agents each voter has limited computational power and limited time to determine his or her vote. Thus, if computing the optimal vote is computationally challenging then we may hope that agents would choose to vote sincerely. This line of thought inspired Bartholdi, Orlin, Tovey, and Trick to study computational properties of elections [2, 1], in search of an election system that in addition to properties desirable from the point of view of social choice theory would also be computationally hard to manipulate. Interestingly, they have also realized that some very attractive election systems are practically unusable because the problem of determining their winners is computationally intractable [3]. (This line of study later on yielded very interesting complexity results regarding parallel access to NP.) Here we focus on election systems for which determining winners can be done in polynomial time.

In their papers, Bartholdi, Orlin, Tovey, and Trick showed that two election systems, STV and second-order Copeland, have the property that manipulation via a single unweighted

Cite as: Copeland Voting: Ties Matter, P. Faliszewski, E. Hemaspaandra, H. Schnoor, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp.983-990.

voter is NP-complete for them. Later, in an important sequence of papers, Conitzer, Sandholm, and Lang [5] showed that many election systems are difficult to manipulate even if one is limited to a certain constant number of candidates, provided the coalition of manipulators includes many weighted voters. Conitzer and Sandholm [6] and Elkind and Lipmaa [10] studied general techniques for modifying existing election systems as to guarantee resistance to manipulation. Hemaspaandra and Hemaspaandra [13] showed that a broad class of election systems, including veto and Borda count, is computationally resistant to manipulation by coalitions of weighted voters. Recent years also brought several papers that call for more detailed study of the complexity of manipulation via arguing that NP-completeness results may not be sufficient [7, 16, 20].

In this paper, following [12], we focus on a family of election systems derived from Copeland voting via introducing value-of-tie parameter α. Intuitively, Copeland$^\alpha$ voting works as follows: For each pair of candidates we ask each voter which one he or she prefers; the candidate preferred by the majority receives one point and the other one receives no points. In case of a tie both candidates receive α points. Apparently, Copeland defined his system with $\alpha = \frac{1}{2}$ in mind, but we study it for all rational α's between 0 and 1.[1] We also mention that some papers define Copeland voting to be what we call Copeland^0 [17, 12].

There are many reasons why Copeland voting is attractive. For example, if there is a candidate who is preferred by the majority of voters to every other candidate, so-called Condorcet winner, then this candidate is a winner of Copeland election. This indicates that, in some sense, Copeland voting is fair. Also, as opposed to many other systems that share this property the set of winners for Copeland$^\alpha$ can easily be computed in polynomial time. This means that Copeland voting is practical in settings where multiple elections happen every second.

Among other results, Conitzer, Sandholm, and Lang [5] showed that manipulation is computationally difficult for Copeland$^{0.5}$ for the case of four candidates and a large coalition of weighted voters. They also showed that it is easy if we have three candidates or less. Our results complement and enrich those of Conitzer, Sandholm and Lang. In particular, we show that the problem of manipulation for unweighted Copeland$^\alpha$ elections is NP-complete even if the size of the manipulating coalition is limited to two. Our result holds for all rational values of α such that $0 < \alpha < 1$ except for $\alpha = \frac{1}{2}$. Since it is well known that single-voter manipulation is easy for Copeland$^\alpha$ (this follows easily from [2]), our result is the first one where an election system originally known to be vulnerable to manipulation via a single voter is shown to be resistant to manipulation via a coalition of a constant number of unweighted voters. The result is interesting also because there are very few natural voting systems for which it is known that unweighted manipulation is hard. Two famous exceptions are STV [1] and second-order Copeland [2], both of which are hard to manipulate even by a single voter.

It is somewhat disappointing, yet fascinating, that our hardness-of-manipulation result for Copeland$^\alpha$ does not hold for $\alpha \in \{0, \frac{1}{2}, 1\}$. Each of these values of α is special in a slightly different way. For example, for $\alpha = 1$ we have a situation where each tied head-to-head contest gives one point to both candidates involved and thus the manipulators have very limited incentive to try to arrange for such ties. Yet, all known hardness-of-manipulation results for Copeland work via creating situations where the manipulators need to carefully seek possibilities for head-to-head ties. Similar singularities, but of a different nature, happen for $\alpha = 0$ and $\alpha = \frac{1}{2}$ (see Section 3 for further discussion).

The fact that our proofs fail for $\alpha \in \{0, 1\}$ can also be explained in the context of the results of Faliszewski et al. [12] regarding so-called microbribery. Consider Copeland$^\alpha$ election where voters represent their votes via irrational preference tables, that is, for every pair of candidates each voter specifies, independently, which one among the two he or she prefers. Note that irrational preference tables do not necessarily yield transitive relations. In bribery a briber tries to ensure that his or her favorite candidate wins via modifying certain votes. Microbribery is a form of bribery where the briber can flip each entry of each preference table at unit cost. In their papers, Faliszewski et al. showed that microbribery is solvable in polynomial time for Copeland$^\alpha$ with $\alpha \in \{0, 1\}$. This is significant for our proofs which, in essence, work via constructing microbribery instances where flips of preference have to be achieved using only two *rational* voters. Alternatively, one would hope that our hardness-of-manipulation results could yield hardness-of-microbribery results and we have made some preliminary steps in that direction, but the transition is, in fact, much more difficult than it may appear.

The focus of this paper is on unweighted voters, but we also study the complexity of manipulation for Copeland$^\alpha$ for weighted voters and a constant number of candidates. We show that the exact complexity of manipulation often depends closely on the winner model as well as on the parameter α: Depending whether we try to make our favorite candidate a winner or a unique winner and whether α is 0, 1 or between these values, the problem of weighted manipulation for Copeland$^\alpha$ with three candidates is either in P or is NP-complete. The only previously known result regarding the complexity of manipulation in three candidate Copeland elections [5] was for $\alpha = \frac{1}{2}$ and the unique winner model and indicated vulnerability.

We believe that our results are interesting in their own right, but we feel that they also indicate a more general trend: We show that exact ways in which ties are handled in an election system, here Copeland voting, can be crucial to establishing complexity results for this system. In our case it matters both how we handle ties at the local level of head-to-head contests as well as how we handle global ties, i.e., whether we seek winners or unique winners. Since it is often tempting to disregard tie-handling issues as uninteresting, we feel that our paper's theme, ties matter, is of importance.

2. PRELIMINARIES

We model an election as a pair $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ is a set of candidates and $V = \{v_1, \dots, v_n\}$ is a multiset of voters. Each voter v_i is represented via a description of his or her preference over the set of candidates and a nonnegative integer w_i, his or her weight. For the purpose of counting votes we treat each weight w_i voter as w_i weight 1 voters. We say that voters are unweighted if all their weights are equal to 1.

[1] Copeland's manuscript [8] is the standard reference for Copeland voting. However, this manuscript is hard to obtain and we were unable to locate it. Thus, we cannot speak with certainty how Copeland defined his system.

There are many ways of representing voters' preferences but for the purpose of this paper we take the standard model where each voter's preference is a strict linear order over the set of candidates. For example, if we have candidates c_1, c_2, and c_3 then a voter who likes c_1 best and really dislikes c_2 would represent his or her preference via order $c_1 > c_3 > c_2$. We have already mentioned irrational preference tables as an alternative representation of preferences [12] and other representations exist as well. We mention the work of Boutilier et al. [4] as a particularly interesting example of a different approach to preference specification, and a paper of Xia et al. [19] as an example of applying ideas from that work in the context of computational social choice theory.

There are many rules used to aggregate votes and define the winners. Plurality rule says that the candidates ranked first by most voters are winners. For a scoring protocol $(\alpha_1, \ldots, \alpha_m)$, where m is the number of candidates and $\alpha_1, \ldots, \alpha_m$ is a nonincreasing sequence of integers, each candidate receives α_i points for each voter that ranks him or her i'th; the candidates with most points are the winners. Natural examples of scoring protocols include plurality, veto (with vectors $(1, \ldots, 1, 0)$), and Borda count (with vectors $(m-1, m-2, \ldots, 0)$). STV is a system where an election over m candidates is performed in up to m rounds where in each round a group of candidates ranked first by the least number of voters is removed; the candidates removed in the last round are the winners.

A different approach to designing election rules, dating back to Ramon Llull and, independently, marquis de Condorcet (see [15]), is based on the analysis of head-to-head majority contests between the candidates. Let $E = (C, V)$ be an election and let c_i and c_j be two distinct candidates in C. By $\mathrm{pref}_E(c_i, c_j)$ we mean the number of voters in V that prefer c_i to c_j. We define $\mathrm{vs}_E(c_i, c_j)$ to be $\mathrm{pref}_E(c_i, c_j) - \mathrm{pref}_E(c_j, c_i)$. If $\mathrm{vs}_E(c_i, c_j)$ is positive then we say that c_i wins his or her head-to-head contest with c_j and if it is negative then we say that c_i loses this contest. If $\mathrm{vs}_E(c_i, c_j) = 0$ then we say that c_i and c_j tie in their head-to-head contest. By a slight abuse of notation we sometimes refer to $\mathrm{vs}_E(c_i, c_j)$ as the result of head-to-head contest between c_i and c_j. (Note that sometimes, when it is clear from the context, we will mean this phrase to simply indicate who had won a particular head-to-head contest.) We define $\mathrm{wins}_E(c_i)$ to be the number of candidates that c_i defeats in their head-to-head contest and $\mathrm{ties}_E(c_i)$ as the number of candidates with whom c_i ties. Following [12], we define Copeland$^\alpha$.

Definition 2.1. *Let $E = (C, V)$ be an election and let α be a rational number, $0 \le \alpha \le 1$. We define Copeland$^\alpha$ score of a candidate c_i from election E as $\mathrm{score}_E^\alpha(c_i) = \mathrm{wins}_E(c_i) + \alpha \cdot \mathrm{ties}_E(c_i)$. The candidates that have highest Copeland$^\alpha$ scores are the winners of the given Copeland$^\alpha$ election.*

Let us now formally define the problem of manipulation for a given election system $\mathcal{E}$, $\mathcal{E}$-manipulation.

Given: An election $E = (C, V \cup W)$, where voters in V have fixed preference lists over C and W are the manipulative voters, whose preference lists we are to set, and a distinguished candidate $p \in C$.

Question: Is it possible to set preference lists of voters in W such that p is a winner in $\mathcal{E}$ election E?

We define the problem of manipulation for the case of weighted voters, $\mathcal{E}$-weighted-manipulation, analogously. Originally, Bartholdi, Tovey, Trick, and Orlin [2, 1] studied the case where only a single voter attempts manipulation. Later on, Conitzer, Sandholm, and Lang [5] introduced the above coalitional manipulation problems both for the weighted and unweighted cases, both in the constructive and destructive settings. (In the destructive case, not studied in this paper, we seek not to make p a winner but to prevent him/her from winning.)

The main goal of this paper is to show hardness of manipulation for Copeland$^\alpha$. We do so via showing that appropriate manipulation problems are NP-complete, via reductions from several well-known NP-complete problems. Below we define these problems.

In the exact-cover-by-3-sets problem (X3C) we are given a finite set $B = \{b_1, \ldots, b_{3k}\}$ and some collection of its 3-element subsets $\mathcal{S} = \{S_1, \ldots, S_n\}$ and we ask if there are k sets $S_{a_1}, \ldots, S_{a_k} \in \mathcal{S}$ such that $\bigcup_{i=1}^{k} S_{a_i} = B$. In the Partition problem we ask whether a sequence of nonnegative integers can be broken into two parts that each sum up to the same value.

Both Partition and X3C are standard, well-known, NP-complete problems.The following problem is also standard, but perhaps a little less widely known. In the 1-in-3-Sat problem we are given a collection of clauses of the form 1-in-3(x, y, z) and we are to determine if there is a truth assignment to the variables such that for each clause 1-in-3(x, y, z), exactly one of the variables x, y or z is true. In the context of 1-in-3-Sat we will refer to collections of clauses as 1-in-3-formulas. For the purpose of this paper we define a restricted version of 1-in-3-Sat where we only allow 1-in-3-formulas with each variable appearing exactly in 4 clauses, and each clause containing three distinct variables. The question is whether there exists a solution to the formula making exactly $\frac{1}{3}n$ variables true, where n is the number of clauses, and the formulas have $\frac{3}{4}n$ variables. We call thus restricted version of the problem 1-in-3-Sat$'$. We have shown that 1-in-3-Sat$'$ is NP-complete.

Lemma 2.2. 1-in-3-Sat$'$ *is* NP*-complete.*

Our proofs often involve fairly complicated instances of elections and we now provide some results that simplify crafting these instances. We start with a useful notation convention. Consider an election with candidate set C and some voter v in this election. Let D be some subset of C. If in the preference list of v we include D then this means: candidates from the set D, listed in some arbitrary but fixed order. If the preference list include $\overleftarrow{D}$ then it means: candidates from D listed in reverse order.

Each election $E = (C, V)$ induces a directed graph $G(E)$ with edges labeled with nonnegative integers. Vertices of $G(E)$ are exactly the candidates of E and edges correspond to the results of head-to-head contests between candidates. That is, for each two distinct candidates $c_i, c_j \in C$ we have an edge in $G(E)$ coming from c_i to c_j with label k if and only if $\mathrm{vs}_E(c_i, c_j) = k$ and $k > 0$. The following lemma, due to McGarvey [14] (see also the work of Stearns [18]) says that each directed, antisymmetric graph with edges labeled by nonnegative even integers is induced by an election.

Lemma 2.3. *For each antisymmetric directed graph G with edges labeled with nonnegative even integers there exists* *an election E such that $G = G(E)$, and E can be computed in polynomial time in the size of G and the largest label.*

For a given positive integer n, let Pad_n be the election (C, V) with $C = \{0, 1, \ldots, 2n\}$ and where for each pair of candidates i, j we have that i defeats j in their head-to-head contest if and only if $(i - j) \bmod 2n \leq n$. It is easy to see that $\|C\| = 2n + 1$ and that each candidate in this election has Copeland$^\alpha$ score n. We use this election as padding in other constructions, e.g., in the next lemma which simplifies designing complicated elections.

LEMMA 2.4. *Let $E = (C, V)$ be an election where $C = \{c_1, \ldots, c_n\}$, and let α be a rational number such that $0 \leq \alpha \leq 1$. For each candidate c_i, let $t_i = \mathrm{ties}_E(c_i)$. For each positive integer q and a sequence of nonnegative integers $k_1, \ldots, k_n$ such that for each k_i we have $0 \leq k_i \leq n^q$ there is an election $E' = (C', V')$ such that: (a) $C' = C \cup D$, where $D = \{d_1, \ldots, d_{2n^{q+1}}\}$; (b) for each i, $1 \leq i \leq n$, $score^\alpha_{E'}(c_i) = 2n^{q+1} - k_i + \alpha t_i$; (c) for each i, $1 \leq i \leq 2n^{q+1}$, $score^\alpha_{E'}(d_i) \leq n^{q+1} + 1$; (d) head-to-head contests between candidates from C in E' have the same results as in E.*

PROOF. We build E' via taking a disjoint union of the input election E with a specifically crafted election $F = (D, W)$, where $D = \{d_1, \ldots, d_{2n^{q+1}}\}$, and setting the result of head-to-head contests between the candidates in C and in D appropriately.[2] F is $\mathrm{Pad}_{n^{q+1}}$ with one arbitrary candidate removed. For each $d_i \in D$ we have $score^\alpha_F(d_i) \leq n^{q+1}$.

We now describe the head-to-head contests between candidates in C and D. Each such contest is won exactly by 2 votes (see Lemma 2.3) and we arrange them as follows. We split D into n groups, $D_1, \ldots, D_n$, each with exactly $2n^q$ candidates. Let c_i be some candidate in C. For each j such that $1 \leq j \leq n$ and $i \neq j$, candidate c_i wins all the head-to-head contest with members of D_j. Regarding the head-to-head contests between c_i and the members of D_i, we set them so that in the final election $score^\alpha_{E'}(c_i) = 2n^{q+1} - k_i + t_i\alpha$. This is easy to do, as—without counting the points from interacting with members of D_i—each c_i has a score between $2n^{q+1} - 2n^q + t_i\alpha$ and $2n^2 - 2n^q + n - 1 + t_i\alpha$ and via interacting with members of D_i, c_i can get any arbitrarily chosen number of points between 0 and $2n^q$. (Recall that for each k_i we have $0 \leq k_i \leq n^q$. Also, c_i already has the $t_i\alpha$ points from the ties within E.) It is easy to see that each $d_i \in D$ also has Copeland$^\alpha$ score as required by the lemma. □

Note that in the proof of Lemma 2.4 we never introduce head-to-head contest ties other than those already present in the election E.

We conclude with the following observation. Let $E = (C, V)$ be an election and let c_i and c_j be two candidates. We can add two voters, v and v', one with preference order $c_i > c_j > C - \{c_i, c_j\}$, and the other with preference order $\overleftarrow{C - \{c_i, c_j\}} > c_i > c_j$, so that we do not change the result of any of the head-to-head contests except the one between c_i and c_j, where we give c_i two votes of advantage. Thus, we can build elections using Lemma 2.4 and then amplify the results of specific head-to-head contests as we please. (We will soon see how this ability is crucial in our proofs.)

[2]One can think of a disjoint union of two elections in terms of a disjoint union of their underlying election graphs.

3. UNWEIGHTED MANIPULATION

This section is dedicated to proving Theorem 3.1.

THEOREM 3.1. *Let α be a rational number such that $0 < \alpha < 1$ and $\alpha \neq \frac{1}{2}$. Copeland$^\alpha$-manipulation is NP-complete.*

Both the case with exactly one manipulator and the case with unweighted voters and a bounded number of candidates are solvable in polynomial time via greedy algorithms (see, e.g., [2]); thus, our result is in some sense optimal.

We split the proof into two lemmas below, one for $0 < \alpha < \frac{1}{2}$ and the other for $\frac{1}{2} < \alpha < 1$. These two parts of the proof, while sharing some common infrastructure, are fairly different. One of the reasons for this diversity is that the dynamics of manipulation differ depending on whether α is above or below $\frac{1}{2}$. In our proofs the manipulators often need to vote in such a way as to change a result of the head-to-head contest between some candidates c_i and c_j from one of the winning to them tieing. If originally c_i wins then changing the result to a tie means that c_i loses $1 - \alpha$ points and c_j gains α points. It often matters which one of these two values is larger.

Infrastructure for the Proof

Both our reductions for Copeland$^\alpha$-manipulation follow the same general structure: We are given some instance I of the problem we reduce from (X3C, or 1-in-3-Sat′) and we build an election $E = (C, V \cup \{v, v'\})$, where candidates in C together with the nonmanipulative voters in V in some way correspond to the structure of I and where the two manipulators, v and v', are trying to ensure that a distinguished candidate $p \in C$ is a winner. Later, for each of the reductions, we describe precisely how this correspondence is realized; now we present the underlying mechanisms we use when designing our elections.

In particular, we specify our elections via listing some candidates $c_1, \ldots, c_n$ together with their Copeland$^\alpha$ scores expressed relative to our designated candidate p's score and together with results of those head-to-head contest between candidates $c_1, \ldots c_n$ that the manipulators have a chance of changing. The goal of the discussion below is to show that we can build elections specified this way (though, of course, besides candidates $c_1, \ldots, c_n$ our elections will have multiple padding candidates, but we will argue that their presence, aside from contributing to c_i's scores, can be essentially ignored).

Since the manipulators' goal is to ensure p's victory, we assume, w.l.o.g., that they always rank p first.

We design our election E to have an even number of voters. Most of the head-to-head contests in E are either won or lost by more than two votes so that the manipulators are too few to affect them. The remaining head-to-head contests are either won or lost be exactly two voters or are tied; the manipulators are in power to modify these results to either a tie (in the former case) or to a victory for one of the involved candidates (in the latter case) via casting appropriate votes. We assume that all head-to-head results for which we do not indicate otherwise are won by more than two votes.

In our constructions of E we ensure that, not counting the manipulators' votes, p obtains Copeland$^\alpha$ score K, where K is some fairly large integer. We design E so that p receives all of these points from victories in head-to-head contests. (Note that the exact value of K is not significant as long as we can build the election E in polynomial time and the scores of all other candidates, expressed relative to K, are maintained.) p loses all remaining head-to-head contests, and—except for one—p loses these contests by more than 2 votes. In this one singled out head-to-head contest between p and a special candidate t, p loses by exactly 2 votes. Since both manipulators rank p first, after including their votes we have that p's final Copeland$^\alpha$ score in E is $\ell = K + \alpha$. We ensure that t has Copeland$^\alpha$ score lower than K and we never use candidate t for purposes other than this in our construction. (As the reader may point out, we will still have to assign results of head-to-head contests between t and all other candidates, but this will be done automatically via an invocation of Lemma 2.4.)

By the above paragraph we know that in our instances of manipulation p ends up with Copeland$^\alpha$ score ℓ (provided both manipulators rank p first, but, as we said, we can assume that w.l.o.g.). The core of our constructions is designing a correspondence between a given instance I of the problem we reduce from and the election E. We build this correspondence via introducing a group of candidates $c_1, \ldots, c_n$ that either correspond to objects in I or whose goal is to enforce consistency constraints of I. For each such candidate c_i we specify Copeland$^\alpha$ score, relative to K (equivalently, relative to ℓ), that we want him/her to have before manipulators' votes are included and the results of those head-to-head contest between c_i's that we want the manipulators to be able to affect (i.e., head-to-head contests that are either tied or won/lost by two votes). We will call those head-to-head contests flexible. The remaining, nonflexible, head-to-head contests can be set arbitrarily, provided the c_i's have Copeland$^\alpha$ scores as specified.

For each candidate c_i, $1 \leq i \leq n$, let f_i and t_i be two nonnegative integers such that our construction requires candidate c_i to have prior-to-manipulation Copeland$^\alpha$ score $K + \alpha t_i \pm f_i$. In addition, we do not want the manipulators to be able to change this score in any way other than via affecting the flexible head-to-head contests we mentioned above. In particular, this means that given a candidate c_i we cannot get his or her αt_i points via simply adding padding candidates with whom c_i would tie in their head-to-head contests. The two manipulators could change the result of such a contest to a victory for one of the candidates involved. On the other hand, it is easy to see that if all t_i's are 0, then an election fulfilling all our criteria is fairly easy to build. All we have to do is build an election E with candidates $p, t, c_1, \ldots, c_n$, with results of head-to-head contests matching the requirements described above (i.e., regarding p and t, and regarding candidates $c_1, \ldots, c_n$) and apply Lemma 2.4 to E with a high enough value of q so that all scores can be implemented. Of course, this also involves computing the exact value of K, but it is fairly easy, given all f_i's. Applying Lemma 2.4 will, of course, involve adding many padding candidates, but it is easy to ensure that each padding candidate has score lower than K (and, thus, lower than ℓ). Finally, via the discussion below Lemma 2.4 it is easy to see how we can ensure that the results of head-to-head contest that are to be won by more than 2 votes are won by exactly 4 votes.

However, limiting ourselves to having each $t_i = 0$ is not good enough for our constructions. So, how do we implement the α part of the scores? Let $T = \Sigma_{i=1}^{n} t_i$. In both our constructions T is polynomial in the size of I, the instance we reduce from. We introduce T candidates $e_1, \ldots, e_T$ and we require that their Copeland$^\alpha$ scores, not counting the manipulators' votes, are exactly $K + 1$. We also stipulate that for each c_i, exactly t_i distinct candidates from the set $\{e_1, \ldots, e_T\}$ win their head-to-head contests with c_i, each by exactly two votes. With this modification we build our election as before, as if the α parts of scores were empty.

How does the above construction help? Before we "begin the manipulation" it does not help at all. However, it is easy to see that if p is to become a winner then both manipulators, v and v', have to guarantee that each candidate e_j in $\{e_1, \ldots, e_T\}$ ties the head-to-head contest with the candidate c_i that e_j used to defeat by 2 votes, thus giving each c_i the additional αt_i points. The reason is that p can at best have Copeland$^\alpha$ score $\ell = K + \alpha$ and, had v and v' not ensured all the ties we mention then at least one of $e_1, \ldots, e_T$ would have Copeland$^\alpha$ score $K + 1 > K + \alpha$ (recall that $0 < \alpha < 1$) and p would not be a winner of the election. v and v' can ensure that all these ties happen via listing each of $c_1, \ldots, c_n$ before any of $e_1, \ldots, e_T$ in their votes.

Thus, using the above described logic on top of the whole construction we can, in effect, assume that we can specify the α parts of the scores of c_i's as well. This completes the description of how in our NP-completeness proofs we can claim that particular instances of manipulation can be built. We quickly mention that in our specifications of Copeland$^\alpha$ scores for candidates we can also use expressions $f_i - \alpha t_i$: For each rational α in $(0, 1)$ there are two nonnegative integer constants, s_1 and s_2, such that $\alpha s_1 = s_2 - \alpha$.

We handle the cases $0 < \alpha < \frac{1}{2}$ and $\frac{1}{2} < \alpha < 1$ separately.

The case $0 < \alpha < \frac{1}{2}$

LEMMA 3.2. *Let α be a rational number such that $0 < \alpha < \frac{1}{2}$. Copeland$^\alpha$-*manipulation *is* NP*-complete.*

PROOF. It is clear that the problem is in NP and so it remains to show NP-hardness. We do so via a reduction from X3C. Let $I = (B, \mathcal{S})$ be an instance of X3C, where $B = \{b_1, \ldots, b_{3k}\}$ and $\mathcal{S} = \{S_1, \ldots, S_n\}$ is some family of 3-element subsets of B. We describe an election E where two manipulative voters, v and v', can ensure a distinguished candidate p's victory if and only if I is a *yes*-instance of X3C. Note that, following the long discussion between Theorem 3.1 and Lemma 3.2 we will only describe significant candidates and omit the padding ones. Similarly, we will express scores that our candidates have before manipulators' votes are counted in the form $\ell + f + \alpha t$, where ℓ is the (essentially fixed) number of points that p obtains. From now on when describing E we will use the word "candidates" to refer only to the significant candidates, but one should keep in mind that of course the padding ones are there as well.

Given $(B, \mathcal{S})$, we build our election E to have the following candidates (together with their Copeland$^\alpha$ scores and results of some of their head-to-head contests.

p. The distinguished candidate whose victory we want to ensure. By the discussion below Theorem 3.1, after manipulation p has exactly ℓ Copeland$^\alpha$ points.

$b_1, \ldots, b_{3k}$. For each $i \in \{1, \ldots, 3k\}$ we have a single candidate b_i with Copeland$^\alpha$ score $\ell - \alpha$.

$S_1, \ldots, S_n, z_{1_1}, \ldots z_{n_3}$. For each set S_i we have a single candidate S_i with score $\ell + 3 - 3\alpha$. Each S_i defeats in their head-to-head contests exactly those b_j's that are members of S_i (these victories are by 2 votes each).

c. The *counter* candidate; has Copeland$^\alpha$ score $\ell-(n-k)\alpha$.

$z_{1_1},\ldots,z_{n_3}$. Candidates z_{i_1}, z_{i_2}, and z_{i_3}, $i \in \{1,\ldots,n\}$ are responsible for implementing a certain consistency gadget. For each $i \in \{1,\ldots,n\}$ we have that z_{i_1} wins by two votes the head-to-head contest with z_{i_2}, z_{i_2} wins by 2 votes the head-to-head contest with z_{i_3}, and z_{i_3} wins by 2 votes the head-to-head contest with c. Also, each S_i defeats, by two votes, each of the candidates z_{i_1}, z_{i_2}, and z_{i_3}. The form of Copeland$^\alpha$ scores of candidates z_{i_t} depends on α and we specify it later.

Aside from head-to-head contests mentioned above, all other head-to-head contests are either won or lost by more than 2 votes.

Each candidate S_i, $i \in \{1,\ldots,n\}$, corresponds to a set in $\mathcal{S}$. We refer to the members of that set, as well as to corresponding candidates, as b_{i_1}, b_{i_2}, and b_{i_3}. Note that each candidate S_i, $i \in \{1,\ldots,n\}$ has a surplus of $3-3\alpha$ Copeland$^\alpha$ points that we have to remove in order to ensure p's victory. For each S_i we can do so via enforcing that S_i ties with at least three of $z_{i_1}, z_{i_2}, z_{i_3}, b_{i_1}, b_{i_2}, b_{i_3}$. Later we show how to specify scores of candidates $z_{i_1}, z_{i_2}, z_{i_3}$, $i \in \{1,\ldots,n\}$ in such a way that in every manipulation that guarantees p's victory, if S_i ties with at least one of z_{i_1}, z_{i_2}, or z_{i_3} then c ties with some candidate that he or she used to lose to (i.e., for each S_i that ties with at least one of z_{i_1}, z_{i_2}, or z_{i_3}, c's Copeland$^\alpha$ score increases by α). We now show that this implies that any manipulation that ensures p's victory has to guarantee that each S_i either ties with all three of $b_{i_1}, b_{i_2}, b_{i_3}$ or with neither of them.

Let us assume that there is a way for v and v' to cast their votes in such a way that p is a winner. This means that all other candidates can have scores at most ℓ. For each $j \in \{0,1,2,3\}$, let K_j be the number of candidates S_i that, including votes v and v', tie with exactly j of $b_{i_1}, b_{i_2}, b_{i_3}$. Since there are exactly n candidates S_i, we have that

$$K_0 + K_1 + K_2 + K_3 = n. \tag{1}$$

Each S_i that is not accounted for in K_3 has to tie with at least one z_{i_t}, $t \in \{1,2,3\}$, and for each such S_i the counter candidate c gets extra α points. Thus,

$$K_0 + K_1 + K_2 \le n-k, \tag{2}$$

as $(n-k)\alpha$ is the largest number of points c can accept without having his or her score over ℓ (recall our gadget connecting z_{i_t}'s and c). Finally, since there are exactly $3k$ candidates $b_1,\ldots,b_{3k}$, and each of them can tie with at most one S_i, we have $3K_3 + 2K_2 + 1K_1 \le 3k$. If we sidewise add to it inequality (2) multiplied by 3 then we obtain $3K_0 + 3K_1 + 3K_2 + 3K_3 + 2K_2 + 1K_1 \le 3n$. Since, via (1), $3K_0 + 3K_1 + 3K_2 + 3K_3 = 3n$, we have that $2K_2 + K_1 \le 0$. Since K_1 and K_2 are nonnegative integers, it implies that K_1 and K_2 are 0; each S_i either ties with all his or her members or with none of them.

Thus, if p is to be a winner, at most k of candidates S_i can tie with candidates corresponding to the members of S_i (because $3K_3 + 2K_2 + K_1 \le 3k$ and both K_2 and K_1 are 0). Since, at most $n-k$ of S_i's can tie with their associated candidates $z_{i_1}, z_{i_2}, z_{i_3}$, it is easy to see that those S_i's that tie with the corresponding candidates $b_{i_1}, b_{i_2}, b_{i_3}$ constitute exactly an exact-3-cover of B. To finish this direction of the proof it remains to show that for any manipulators' votes that ensure p's victory it really is the case that for each S_i that ties with at least one of $z_{i_1}, z_{i_2}, z_{i_3}$ (including manipulators' votes) candidate c's score increases by α.

We set the scores of candidates $z_{1_1},\ldots,z_{n_3}$ depending on the value of α. We first handle the case when $\frac{1}{3} \le \alpha < \frac{1}{2}$. In this case, we declare that each of z_{i_t}, $i \in \{1,\ldots,n\}$, $t \in \{1,2,3\}$, has score exactly $\ell + 1 - 3\alpha$. (Note that since $\alpha \ge \frac{1}{3}$, $1-3\alpha \le 0$.) It is easy to see that if any of z_{i_t} obtains extra α (or more) points from tieing either with S_i or $z_{i_{t-1}}$ (provided $t>0$ for the latter) then we need to ensure that this z_{i_t} also "unloads" these extra points somewhere. The only way to decrease z_{i_t}'s score is via ensuring that he or she ties with $z_{i_{t+1}}$ (or c, if $t=3$). Since $\alpha < \frac{1}{2}$, the amount of points z_{i_t} loses this way balances all the points z_{i_t} might obtain due to manipulation and ensures that his or her score is at most ℓ. Also, due to z_{i_t} tieing with $z_{i_{t+1}}$, $t \in \{1,2\}$, $z_{i_{t+1}}$ obtains extra α points he or she needs to unload. This way the effect of S_i tieing with either one of z_{i_t}'s ($t \in \{1,2,3\}$) propagates to eventually increasing c's score by α. Also, the reader can easily verify that if each of $z_{i_1}, z_{i_2}, z_{i_3}$ ties with S_i, z_{i_1} ties with z_{i_2}, z_{i_2} ties with z_{i_3} and z_{i_3} ties with c then each of $z_{i_1}, z_{i_2}, z_{i_3}$ has Copeland$^\alpha$ score at most ℓ. Such tieing can be implemented by manipulators v and v' if in both their votes they prefer c to z_{i_3} to z_{i_2} to z_{i_1} to S_i.

For the case of rational α such that $0 < \alpha \le \frac{1}{3}$ it is easy to see that the same arguments work provided that each candidate z_{i_t}, $i \in \{1,\ldots,n\}$, $t \in \{1,2,3\}$ starts with Copeland$^\alpha$ score equal to ℓ.

We now show that if I is a yes-instance of X3C then p can become a winner of our election. Let $\mathcal{S}_C$ be a set of all candidates S_i that correspond to some exact-3-cover of B and let Z_C be the set of their corresponding z_{i_t} candidates. Let Z be the set of all the z_{i_t} candidates in the election. It is easy to check that voters v and v' can ensure p's victory via casting votes as follows.

$$v \;:\; p > c > Z_C > \mathcal{S}_C > B > \mathcal{S} - \mathcal{S}_C > Z - Z_C > \cdots$$
$$v' \;:\; p > \mathcal{S} - \mathcal{S}_C > Z - Z_C > c > Z_C > \mathcal{S}_C > B > \cdots$$

The ellipsis means the padding candidates, listed in arbitrary order. It is easy to verify that with these votes all candidates end up with Copeland$^\alpha$ score of at most ℓ and that p gets exactly ℓ points, becoming a winner. □

The case $\frac{1}{2} < \alpha < 1$

Lemma 3.3. *Let α be a rational number such that $\frac{1}{2} < \alpha < 1$. Then Copeland$^\alpha$-manipulation is* NP*-complete.*

Proof. The problem is clearly in NP; we show hardness via a reduction from 1-in-3-Sat′ (see Lemma 2.2). Let $\varphi = \bigwedge_{i=1}^{n}$ 1-in-3(a_i, b_i, c_i) be a problem instance. We construct an election E where each variable from φ is a candidate; there are also candidates $F_1,\ldots,F_4$, and for each $i \in \{1,\ldots,n\}$, we have a candidate x_i. Additionally, we introduce candidates $K_1,\ldots,K_4, J_1,\ldots,J_4, B_1,\ldots,B_4$, and a preferred candidate p.

Let all F_i tie among each other, each F_i win against each candidate from VAR(φ) with two votes. Additionally, let x_i lose against a_i, b_i and c_i with 2 votes, every variable lose against $K_1,\ldots,K_4$ by two votes and win against $J_1,\ldots,J_4$ by two votes. For $i = 1,\ldots,4$, F_i and B_i are tied, B_i wins against J_i by 2 votes. All other relationships cannot be changed by two voters. The scores are set up as follows (note that $\frac{n}{4}$ is a natural number):

$$
\begin{array}{ll}
p & : \ \ell \\
F_i & : \ \ell + (1-\alpha)(\frac{n}{4} + i - 5) + (i-1)\alpha \\
\mathrm{VAR}(\varphi) & : \ \ell - 4(2\alpha - 1) \\
x_i & : \ \ell - \alpha \\
B_i & : \ \ell + \alpha \\
J_1, \ldots, J_4 & : \ \ell - \frac{1}{2}n\alpha \\
K_1, \ldots, K_4 & : \ \ell + \frac{1}{2}n(1-\alpha)
\end{array}
$$

As in the discussion at the beginning of the proof of Lemma 3.2, we use the same techniques to construct the election with these relevant candidates and scores in polynomial time. As in the previous proof, we only mention the relevant candidates, and ignore the padding candidates.

We show that φ has a solution I making exactly $\frac{n}{4}$ variable true if and only if p can be made winner in E with two additional votes. First assume that φ has such a solution I. We view I as the set of its true variables, and denote the set of false variables with $\overline{I}$. Let B denote the set $\{B_1, \ldots, B_4\}$, J the set $\{J_1, \ldots, J_4\}$, K the set $\{K_1, \ldots, K_4\}$, and let X denote the set $\{x_1, \ldots, x_n\}$. If the manipulators vote:

$$
\begin{array}{ll}
V_1 & : \ p > X > I > F_1 > \cdots > F_4 > B > J > \overline{I} > K \\
V_2 & : \ p > J > \overline{I} > K > X > I > F_1 > \cdots > F_4 > B,
\end{array}
$$

then it is easy to verify that all candidates have exactly ℓ points, hence p is a winner.

For the converse, assume that p can be made a winner with two additional voters. Let E' be the manipulated election; the score of each candidate must be at most ℓ in the election E' for p to win.

We say that some group G of candidates *wins* (or *loses*) t points against a group H if the changes in the relationships between candidates from G and H make the total score of candidates in G increase (decrease) by t.

The main idea of the proof is the following: We want to obtain a truth assignment for the variables in φ such that exactly one variable is true in each clause. The "true" variable among a_i, b_i, c_i will be the one that x_i ties with instead of losing. Due to the score of x_i, it is obvious that there can be at most one such variable, and it is easy to set up the election in a way that each x_i must tie with at least one of the variables in its clause in order for all points to get down to ℓ. However, variables must behave consistently: A variable cannot be true in one clause and false in another. Ensuring consistency is the trickiest part of our proof. Due to space reasons we cannot present our complete argument and instead we list facts that our proof establishes. The proofs for these are based on the above-mentioned observation that each candidate can have at most ℓ points. By considering the possibilities for each group of candidates to lose or gain points by changing their win/lose/tie-relationships with candidates in other groups, we can prove the following:

FACT 1. *In E', F_i wins against B_i for all $i \in \{1, \ldots, 4\}$.*

FACT 2. *In E', the group $\mathrm{VAR}(\varphi)$ wins at least $2n\alpha$ points against the group K, and loses at most $2n(1-\alpha)$ points against the group J.*

FACT 3. *In E', each x_i ties with at most one candidate from $\mathrm{VAR}(\varphi)$. The group $\mathrm{VAR}(\varphi)$ loses at most $n(1-\alpha)$ points against the group X.*

FACT 4. *In E', all ties among the group F are broken, and there are exactly n ties between F and $\mathrm{VAR}(\varphi)$. Also, there are at least $2n$ ties between the groups J and $\mathrm{VAR}(\varphi)$.*

FACT 5. *For each $i \in \{1, \ldots, 4\}$, there is a manipulator vote where B_i is voted ahead of J_i in E'.*

FACT 6. *In E', there are exactly $\frac{n}{4}$ variables which tie against all of the F_i. We denote this set with I. The other variables tie with none of the F_i. The candidates in I tie with all of their related x_i and those in $\mathrm{VAR}(\varphi) - I$ win against all of the related x_i. Each x_i ties with some candidate in I.*

We regard I as a truth assignment, setting exactly those variables true which are elements of I. Since $|I| = \frac{n}{4}$ due to Fact 6, it remains to prove that I satisfies φ. Let 1-in-3(a_i, b_i, c_i) be a clause in φ. We need to show that exactly one of the a_i, b_i, c_i is an element of I. Fact 6 implies that there is a variable $v \in I$ such that x_i ties against it. Since x_i can only tie against a_i, b_i, or c_i, this implies that one of these is an element of I. Now assume that at least two of these are elements of I. Due to Fact 6, both of them must tie against x_i. This is a contradiction, since due to Fact 3, each x_i can tie with at most one candidate from $\mathrm{VAR}(\varphi)$. Therefore, exactly one of a_i, b_i, c_i is an element of I as required. Thus I satisfies every clause 1-in-3(a_i, b_i, c_i), and thus it satisfies the whole formula φ. □

4. WEIGHTED MANIPULATION

We turn our focus to the weighted manipulation problem for Copeland$^\alpha$ elections with three candidates. We study both the regular weighted manipulation problem and *unique manipulation problem*, Copeland$^\alpha$-weighted-unique-manipulation, where we ask if there is a way to make our designated candidate the *only* winner of the election. As we indicate in the introduction, depending whether α is 0, 1, or in between, and whether we consider winners or unique winners, the complexity of weighted manipulation for 3-candidate Copeland$^\alpha$ elections can vary greatly. The table below summarizes our results.

	$\alpha = 0$	$0 < \alpha < 1$	$\alpha = 1$
manipulation	NP-c	NP-c	P
unique manipulation	NP-c	P	P

We skip parts of the proofs of the following theorems but we mention that the parts that we skip in each are very similar to those that we keep in the others.

We first consider our regular manipulation problem, where the designated candidate is to become one of the winners, but not necessarily the only winner. For three candidates this problem appears to be difficult if $\alpha < 1$ and easy if $\alpha = 1$. The results differ in the unique-winner case.

THEOREM 4.1. *Let α be a rational number such that $0 \leq \alpha < 1$. Then Copeland$^\alpha$-weighted-manipulation is* NP-*complete when considering elections with exactly* 3 *candidates. For $\alpha = 1$ the same problem is in* P.

PROOF. We skip the NP-completeness proof due to space constraints and give the polynomial-time algorithm for Copeland^1. We are given three candidates, a, b and p, set V of weighted voters, the sequence of manipulators' weights, and we are to decide if there is a way to set manipulators' votes as to ensure p's victory in this Copeland^1 election.

Our algorithm works as follows: We let each manipulator rank p first and we compute p's score (which by now is fully determined). Accept if it is 2 and reject if it is 0. If p's score is 1, then, w.l.o.g., let a be the candidate which beats p in their head-to-head contest and let each manipulator vote $p > b > a$. Accept if and only if p wins this election. Correctness and efficiency of the algorithm are clear. □

THEOREM 4.2. *For* $0 < \alpha \leq 1$, *Copeland*$^{\alpha}$*-weighted-unique-manipulation can be solved in polynomial time for* 3 *candidates. For* $\alpha = 0$ *the same problem is* NP*-complete.*

PROOF. Due to space reasons we only give a sketch of the NP-completeness part of the proof, a reduction from Partition. Let $s_1, \ldots, s_n$ with $\sum_{i=1}^{n} s_i = 2k$ be an instance. We construct an election with candidates p, a, b, two voters, each having weight k, one voting $a > p > b$ and the other voting $b > a > p$, and with n manipulators, with weights $s_1, \ldots, s_n$. It is easy to see that if partition is possible then the manipulators can ensure that p is a unique winner. Conversely, assume that p can be made a unique winner. W.l.o.g. we assume that p is ranked first by all manipulators and so p's score is 1 If p is a unique winner, then the scores of both a and b must be zero, ensuring that a and b are tied in their head-to-head contest. This is only possible if there is a partition of $s_1, \ldots, s_n$. □

5. OPEN PROBLEMS

The most interesting problem left open by this paper is to determine the complexity of Copeland$^{\alpha}$-manipulation for $\alpha \in \{0, \frac{1}{2}, 1\}$. Other directions for future research include studying Copeland$^{\alpha}$-weighted-manipulation for more than 3 candidates.

6. ACKNOWLEDGMENTS

Supported in part by NSF grants CCF-0426761 and IIS-0713061, a Friedrich Wilhelm Bessel Research Award, the Alexander von Humboldt Foundation's TransCoop program, and the DAAD postdoc program. We thank the anonymous referees for their very helpful comments.

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] J. Bartholdi, III, C. Tovey, and M. Trick. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare*, 6(2):157–165, 1989.

[4] C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.

[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] V. Conitzer and T. Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, pages 781–788. Morgan Kaufmann, Aug. 2003.

[7] V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proceedings of the 21st National Conference on Artificial Intelligence*. AAAI Press, July/August 2006.

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

[9] 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. ACM Press, Mar. 2001.

[10] E. Elkind and H. Lipmaa. Hybrid voting protocols and hardness of manipulation. In *The 16th Annual International Symposium on Algorithms and Computation, ISAAC 2005*, pages 206–215. Springer-Verlag *Lecture Notes in Computer Science #3872*, December 2005.

[11] E. Ephrati and J. Rosenschein. Multi-agent planning as a dynamic search for social consensus. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, pages 423–429. Morgan Kaufmann, 1993.

[12] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting broadly resist bribery and control. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 724–730. AAAI Press, July 2007.

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

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

[15] I. McLean and A. Urken. *Classics of Social Choice*. University of Michigan Press, 1995.

[16] A. Procaccia and J. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 28:157–181, Feb. 2007.

[17] A. Procaccia, J. Rosenschein, and G. Kaminka. On the robustness of preference aggregation in noisy environments. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM Press, 2007. To appear, 2007.

[18] R. Stearns. The voting problem. *The American Mathematical Monthly*, 66(9):761–763, 1959.

[19] L. Xia, J. Lang, and M. Ying. Stringly decomposable voting rules on multiattribute domains. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 776–781. AAAI Press, July 2007.

[20] M. Zuckerman, A. Procaccia, and J. Rosenschein. Algorithms for the coalitional manipulation problem. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 277–286, Jan. 2008.