

Llull and Copeland Voting Broadly Resist Bribery and Control *

Piotr Faliszewski

Dept. of Computer Science
University of Rochester
Rochester, NY 14627, USA

Edith Hemaspaandra

Dept. of Computer Science
Rochester Institute of Technology
Rochester, NY 14623, USA

Lane A. Hemaspaandra

Dept. of Computer Science
University of Rochester
Rochester, NY 14627, USA

Jörg Rothe

Inst. für Informatik
Heinrich-Heine-Univ. Düsseldorf
40225 Düsseldorf, Germany

Abstract

Control of elections refers to attempts by an agent to, via such actions as addition/deletion/partition of candidates or voters, ensure that a given candidate wins (Bartholdi, Tovey, & Trick 1992). An election system in which such an agent's computational task is NP-hard is said to be resistant to the given type of control. Aside from election systems with an NP-hard winner problem, the only systems known to be resistant to all the standard control types are highly artificial election systems created by hybridization (Hemaspaandra, Hemaspaandra, & Rothe 2007b). In this paper, we prove that an election system developed by the 13th century mystic Ramon Llull and the well-studied Copeland election system are both resistant to all the standard types of (constructive) electoral control other than one variant of addition of candidates. This is the most comprehensive resistance to control yet achieved by any natural election system whose winner problem is in P. In addition, we show that Llull and Copeland voting are very broadly resistant to bribery attacks, and we integrate the potential irrationality of voter preferences into many of our results.

Introduction

Elections have played an important role in human societies for thousands of years. For example, elections were of central importance in the democracy of ancient Athens. There, citizens typically could only agree (vote *yes*) or disagree (vote *no*) with the speaker, and simple majority-rule was in effect. Mathematical study of elections, give or take a few discussions by the ancient Greeks, was until recently thought to have been initiated only a few hundred years ago, namely, in the breakthrough work of Borda and Condorcet—later in part reinvented by Dodgson. One of the most interesting results of this early work is Condorcet's observation that if one conducts elections with more than two alternatives then even if all voters have rational (i.e., transitive) preferences, the society in aggregate can be irrational (indeed, can have cycles of strict preference). Based on his observations, Condorcet suggested that if there exists a candidate c such that c defeats any other candidate in a head-to-head contest then that candidate should win the election. Such a candidate is

*Supported in part by Humboldt-{TransCoop, Bessel Award}, DFG-RO-1202/9-3, and NSF-{CCR-0311021, CCF-0426761}. Copyright © 2007, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

called a Condorcet winner. Clearly, there can be at most one Condorcet winner in any election and there might be none.

This understanding of history has been shattered during the past few decades, as it has been rediscovered that the study of elections was in fact considered deeply as early as the thirteenth century (see Hägele and Pukelsheim (2001) and the citations therein regarding Ramon Llull and the fifteenth century figure Cusanus, esp. the citations that there are numbered 3, 5, and 24–27). Ramon Llull (b. 1232, d. 1315), a Catalan mystic, missionary, and philosopher, developed an election system that (a) has an efficient winner-determination procedure and that (b) elects a Condorcet winner whenever one exists and otherwise elects candidates that are, in some sense, closest to being Condorcet winners. Llull's motivation for developing an election system was to obtain a method of choosing the abbesses, abbots, bishops, and perhaps even the pope. His ideas never gained public acceptance in medieval Europe and were long forgotten.

Llull's election system is remarkably similar to what is now known as Copeland elections. We introduce a parameterized notion, Copeland $_{\alpha}$, broad enough to capture both systems and more.

Definition 1 Let α be a rational number such that $0 \leq \alpha \leq 1$. We define the election system Copeland $_{\alpha}$ as follows. For each pair of distinct candidates, all voters report which among the two they prefer. The winner of such a head-to-head majority contest receives one point and the loser receives zero points; in case of a tie both candidates receive α points. The candidate (candidates, if tied) with the most points is the winner.

We will refer to Copeland $_1$ as Llull and to Copeland $_0$ as Copeland.¹ (Note that the term Copeland is sometimes used for what in our notation is Copeland $_{0.5}$. In the computational

¹Page 23 of Hägele and Pukelsheim (2001) indicates in a way we find deeply convincing (namely by a direct quote of Llull's in-this-case-very-clear words from his *Artifitium Electionis Personarum*—which was rediscovered by those authors in the year 2000) that at least one of Llull's election systems was Copeland $_1$, and so in this paper we refer to the both-candidates-score-a-point-on-a-tie variant as Llull voting.

In some settings Llull required the candidate and voter sets to be identical and had an elaborate two-stage tie-breaking rule ending in randomization. We disregard these issues here and cast his system into the modern idiom for election systems.

social choice literature, there is no complete agreement on the definition of Copeland voting. For example, Procaccia, Rosenschein, and Kaminka (2007) define Copeland as Copeland₀ and Conitzer and Sandholm (2002a) use a definition that in effect is equivalent to Copeland_{0.5}. We mention that the Copeland_{0.5} approach is taken far more frequently.) The definitional difference between Lull and Copeland might seem trivial but, in fact, it can make the dynamics of Lull's system quite different from those of Copeland's. Proofs of results for Lull differ considerably from those for Copeland. Sometimes analogous problems for Lull and Copeland have very different complexities (the full version will give an example).

It is interesting to note that Lull in fact allowed voters to have *irrational* preferences. Given three candidates, *c*, *d*, and *e*, it was perfectly acceptable for a voter to prefer *c* to *d*, *d* to *e*, and *e* to *c*. On the other hand, in modern studies of voting and election systems each voter's preferences are most typically modeled as a *strict* linear order over all candidates. Yet irrationality is a very tempting and natural concept. Consider Bob, who likes to eat out but is often in a hurry. Bob prefers diners to fast food because he is willing to wait a little more to get better food. Also, given a choice between a fancy restaurant and a diner he prefers the restaurant; again because he is willing to wait somewhat longer to get better quality. However, given the choice between a fast-food place and a fancy restaurant Bob might reason that he is not willing to wait so long for the fancy restaurant and will choose fast food instead. Thus, regarding catering options, Bob's preferences are irrational in our sense. Similar irrationalities easily come up when voters make their choices based on multiple criteria—a very natural scenario both among humans and software agents.

The goal of this paper is to study Lull's and Copeland's election systems from the point of view of computational social choice theory, in the setting where voters are rational and in the setting where the voters are allowed to have irrational preferences. (Note: When we henceforth say "irrational voters," we mean that they may have irrational preferences, not that they each must.)

In general it is impossible to design a perfect election system. In the 1950s Arrow famously showed that there is no social choice system that satisfies a certain small set of requirements, and later Gibbard, Satterthwaite, and Duggan and Schwartz showed that any natural election system can be manipulated by strategic voting, i.e., by a voter who reveals different preferences than his or her true ones in order to affect an election's result in his or her favor. Also, no natural election system that has a polynomial-time winner-determination procedure has yet been shown to be resistant to all types of control via procedural changes.

These obstacles are very depressing, but the field of computational social choice theory grew in part from the realization that computational complexity provides a tool to partially circumvent them. In particular, around 1990 Bartholdi, Tovey, and Trick (1989; 1992), and Bartholdi and Orlin (1991) brilliantly observed that while we might not be able to make manipulation (i.e., strategic voting) and control of elections impossible, we can at least try to make such

manipulation and control so difficult computationally that neither voters nor election organizers will attempt it. For example, if there is a way for a committee's chair to set up an election within his or her committee in such a way that his or her favorite option is guaranteed to win but the chair's computational task would take a million years, then for all practical purposes we may assume that the chair is prevented from finding such a set-up.

So since the seminal work of Bartholdi, Orlin, Tovey, and Trick a large body of research has naturally been dedicated to the study of computational properties of election systems. Some topics that received much attention are the complexity of manipulating elections (Conitzer & Sandholm 2002a; 2003; 2002b; 2006; Conitzer, Lang, & Sandholm 2003; Elkind & Lipmaa 2005; Hemaspaandra & Hemaspaandra 2007; Procaccia & Rosenschein 2006; Procaccia, Rosenschein, & Zohar 2007) and of controlling elections via procedural changes (Hemaspaandra, Hemaspaandra, & Rothe 2007a; 2007b; Procaccia, Rosenschein, & Zohar 2007). Recently, Faliszewski, Hemaspaandra, and Hemaspaandra (2006) studied the complexity of bribery in elections. Bribery shares some features of manipulation and some of control. In particular, the briber picks the voters he or she wants to affect (as in voter control problems) and asks them to vote as he or she wishes (as in manipulation).

In this paper we study Lull and Copeland elections with respect to the difficulty of bribery and procedural control. We believe that the computational complexity of procedural control within elections is an important topic in multiagent systems since these elections can be used as a tool to solve many practical problems. As just a few of examples we mention the work of Ephrati and Rosenschein (1993) where elections are used for the purposes of planning, and the work of Dwork et al. (2001) where elections are used to aggregate results from multiple web-search engines. In a multiagent setting we might have hundreds of elections happening every minute and we cannot hope to carefully check if in each case the party that organized the elections did not attempt some procedural changes to skew the results. However, if it is computationally hard to effectuate such procedural changes then we can hope it is impossible for the organizers to undertake them.

A standard technique for showing that a particular elections-related problem (e.g., the problem of deciding whether the chair can make his or her favorite candidate *p* a winner by influencing at most *k* voters not to cast their votes) is computationally difficult is to show that it is NP-hard. This approach is taken in almost all of the papers on computational social choice theory cited above and it is the approach that we take in this paper. One of the justifications for using NP-hardness as a measure of difficulty is that in multiagent settings manipulation and control of elections is conducted by computationally bounded software agents that have neither human intuition nor the computational ability to solve NP-hard problems.

Recently, such papers as (Procaccia & Rosenschein 2006; Conitzer & Sandholm 2006; Homan & Hemaspaandra 2006) have studied the frequency (or sometimes, probability weight) of correctness of heuristics for voting prob-

lems. We view worst-case study as a natural prerequisite to a frequency-of-hardness attack: After all, there is no point in seeking frequency-of-hardness results if the problem at hand is in P to begin with. Also, current frequency results have debilitating limitations (for example, being locked into specific distributions; depending on unproven assumptions; adopting “tractability” notions that declare undecidable problems tractable and that are not robust under even linear-time reductions). Although frequency of hardness is a fascinating direction, these models are not ready for prime time and, contrary to some people’s impression, fail to imply average-case polynomial runtime claims. Erdélyi et al. (2007) provides a discussion of some of these issues.

Although space doesn’t allow a full presentation of our proof approaches here, our bribery results and some of our control results are proven by a method we introduce of controlling the relative performances of certain voters in such a way that, if one sets up other restrictions appropriately, the legal possibilities for control/change actions are sharply constrained. We call our approach “the UV technique,” since it is based on dummy candidates u and v . The proof of Lemma 5 is a particular case of the method. We feel that the UV technique will be useful even beyond the bounds of this paper, and we now provide an example showing that. In particular, the authors have recently noticed that the proof of Bartholdi, Tovey, and Trick (1992, Theorem 12) of an important result of theirs related to this paper (namely, that Condorcet voting is resistant to constructive control by deleting voters) is invalid. The invalidity is due to the proof centrally using nonstrict voters, in violation of Bartholdi, Tovey, and Trick’s (1992) (and our) model, and the invalidity seems very hard to fix with the proof approach taken there. However, using the UV technique one indeed can easily obtain a correct proof, and we have done so (our in-preparation full version will present the proof). Thus the UV technique this paper introduces has applications even beyond those in this paper, and in particular Theorem 12 of (Bartholdi, Tovey, & Trick 1992) holds. We noticed that Theorem 14 of the same paper also has a similar flaw and we have validly proven it also, again using the UV technique. The UV technique in some sense provides a more uniform approach for (certain) voter-related election problems.

This paper is organized as follows. In the Preliminaries section we formalize the notions of Llull and Copeland elections and introduce our notation. In the Bribery section we show that Llull and Copeland are perfect from the point of view of resistance to bribery, both in the case of rational voters and in the case of irrational voters. On the other hand, we show that if one changes the bribery model to allow “micro-bribes” of voters, their resistance in the irrational case is pierced by the existence of a polynomial-time algorithm. In the Control section we present our results on the procedural control of Llull and Copeland elections. What this section shows is, in effect, that both Llull and Copeland have resistance to more types of (constructive) control than has been shown for any other known natural system with a polynomial-time winner problem.

Preliminaries

An election consists of a candidate set $C = \{c_1, \dots, c_n\}$, a collection V of voters, and a rule that specifies the winners. As is standard, each voter is represented (individually, except later when we discuss succinct inputs) via his or her preferences. We consider two ways in which voters can express their preferences. In the rational case, each voter’s preferences are a strict linear order over the set C , e.g., each voter has a preference list $c_{i_1} > c_{i_2} > \dots > c_{i_n}$, with $\{i_1, \dots, i_n\} = \{1, \dots, n\}$. In the irrational case, each voter’s preferences are represented as a table that for every unordered pair of distinct candidates c_i and c_j in C indicates whether the voter prefers c_i to c_j (i.e., $c_i > c_j$) or c_j to c_i (i.e., $c_j > c_i$). Procaccia, Rosenschein, and Kaminka (2007) discuss encoding *rational* preferences in a way that is very similar to such preference tables.

Some well-known election rules for the case of rational voters include plurality, Borda count, and Condorcet. Plurality elects the candidate(s) ranked highest by the largest number of voters. Borda count elects the candidate(s) receiving the most points, where each voter v_i gives each candidate c_j as many points as the number of candidates c_j defeats on v_i ’s preference list. A Condorcet winner is a candidate c_i such that for every other candidate c_j it holds that c_i is preferred to c_j by a strict majority of the voters. Note that all of these election systems can easily be adapted to work with irrational voters. Here, however, we focus on Llull and Copeland (see Definition 1).

Let us now describe the computational problems that we study in this paper. Our problems come in two flavors: constructive and destructive. In the constructive version the goal is to test whether, via either bribery or control, it is possible to make a designated candidate a winner of the election. In the destructive case we seek to guarantee that a despised candidate is not a winner of the election.

The (constructive) bribery problem for the Copeland system with rational voters, which we call Copeland-bribery, is defined as:

Given: A set C of candidates, a collection V of voters specified via their preference lists, a distinguished candidate p , and a nonnegative integer k .

Question: Is it possible to make p a winner of the Copeland election by modifying the preference lists of at most k voters?

The same problem for Copeland with irrational voters is called Copeland_{Irrational}-bribery, and the corresponding problems for Llull are called Llull-bribery for the rational case and Llull_{Irrational}-bribery for the irrational case.

Bribery problems seek to change the outcome of elections via modifying the preferences of some of the voters. On the other hand, control problems seek to change the outcome of elections by modifying their structure, e.g., via adding candidates (AC), deleting candidates (DC), partitioning candidates (PC), partitioning candidates with run-off (RPC), adding voters (AV), deleting voters (DV), and partitioning voters (PV). The name of a control problem is formed by concatenating the name of the election system with CC or

DC, for constructive control and destructive control respectively, and the acronym of the control type that we have in mind. For the case of partitioning candidates and voters we also distinguish two subcases, ties-eliminate (TE) and ties-promote (TP). In the ties-eliminate subcase, if a subelection is won by more than one candidate then that subelection doesn't promote any of its candidates to the next round. In the ties-promote subcase, all the winners within a given subelection are promoted to the next round. We below give a formal description of just a few of the above control types, and via their names have above described the rest of the control types informally; interested readers can find the missing control-type definitions in Hemaspaandra, Hemaspaandra, and Rothe (2007a).

Copeland destructive control via deleting voters for the case of rational voters, Copeland-DCDV, is defined as:

Given: A set C of candidates and a collection V of voters represented via preference lists over C , a distinguished candidate c , and a nonnegative integer k .

Question: Is it possible to by deleting at most k voters ensure that c is not a winner of the resulting Copeland election?

Similarly, $\text{Llull}_{\text{Irrational-CCAC}_u}$, the problem of constructive control of Llull elections with irrational voters via adding candidates is defined as:

Given: Disjoint sets C and D of candidates, a collection V of voters specified via their (possibly irrational) preference tables over the candidates in the set $C \cup D$, and a distinguished candidate p .

Question: Is it possible to choose a subset E of D such that p is a winner of the Llull election with voters V and candidates $C \cup E$?

The above definition of $\text{Llull}_{\text{Irrational-CCAC}_u}$ is based on that introduced by Bartholdi, Tovey, and Trick (1992). In contrast with the other control problems involving adding or deleting candidates or voters, in the adding candidates problem Bartholdi, Tovey, and Trick did not introduce a nonnegative integer k that bounds the number of candidates (from the set D) the chair is allowed to add. We feel this asymmetry is unjustified and thus we define a with-change-parameter version of the control by adding candidates problems, which we denote by AC_1 (where the "1" stands for a limit on the number of candidates that can be added, in contrast with the model of Bartholdi, Tovey, and Trick (1992), which we denote AC_u with the "u" standing for an unlimited number of added candidates). The with-parameter version is the long-studied case for AV, DV, and DC, and we in this paper will use AC as being synonymous with AC_1 , and will thus use the notation AC for the rest of this paper when speaking of AC_1 . We suggest this as a natural regularization of the definitions and we hope this version will become the "normal" version of the problem for further study. However, we caution the reader that in earlier papers AC is used to mean AC_u . So, turning to what we mean by AC (equivalently, AC_1), in \mathcal{E} -CCAC we ask whether it is possible to make the distinguished candidate p a winner of election \mathcal{E} by adding at most k candidates from the spoiler candidate set D . (Note

that k is part of the input.) We define the destructive version, \mathcal{E} -DCAC, analogously.

Note that our problems—both those regarding bribery and those regarding control—speak of nonunique winners. Nonetheless, we have proven all our control results both in the case of nonunique and (to be able to fairly compare them with existing control results, which except for the interesting "multiwinner" model of Procaccia, Rosenschein, and Zohar (2007) are for the unique winner case) unique winners.

Not all election systems can be affected by each control type. For example, if a candidate c is a Condorcet winner then it is impossible to prevent him or her from being a Condorcet winner by deleting other candidates (i.e., Condorcet is *immune* to this type of control). However, for the case of Llull and Copeland systems it is easy to see that for each standard type of control there is a scenario where the outcome of the election can be changed via conducting the control action (i.e., Llull and Copeland are *susceptible* to each standard control type).

We say that an election system (Llull, Copeland, etc.) is *resistant* to a particular attack (be it a type of control or of bribery) if the appropriate computational problem is NP-hard and (only for the control case) there is a scenario where this type of attack can change the winners of the election in the appropriate constructive/destructive way (note: for Llull and Copeland the second clause, needed for some systems with complex evaluation problems, is superfluous). On the other hand, if the computational problem is in P and there is a scenario where this type of attack can change the winners of the election, then we say the system is *vulnerable* to this attack. For all resistance claims in this paper, NP-membership is clear, and so NP-completeness holds. The notions of resistance and vulnerability (and of immunity and susceptibility) for control problems in election systems were introduced by Bartholdi, Tovey, and Trick (1992).

Bribery

Our main result regarding bribery is that the Llull and Copeland systems are resistant to bribery regardless of voters' rationality and our mode of operation (constructive versus destructive).

Theorem 2 *The Llull and Copeland systems are resistant to both constructive and destructive bribery in both the rational-voters case and the irrational-voters case.*

Theorem 2 follows by an application of our UV technique. For reasons of space we don't include the proof here, but this application of the UV technique is very similar to that in Lemma 5, whose proof we include in the next section.

There is a natural yet different and more micro-scale way of defining bribery for the case of irrational voters. The way our bribery problems are defined, we can choose up to k voters and modify their preferences freely. However, what if, for the case of irrational voters, instead of modifying the whole preference table of some selected voters we now pay unit cost each time we flip a single entry in a voter's table? Call this version of the problem bribery'. Theorem 2 notwithstanding, the bribery' problems are easy.

Theorem 3 *The Llull and Copeland systems with irrational voters are vulnerable to constructive and destructive bribery’.*

The destructive cases follow via greedy algorithms. On the other hand, our algorithms for the constructive cases are rather involved. Very briefly put, the main idea is that we model Llull/Copeland elections via a network flow problem, where the units that flow in the network are Llull/Copeland points. This allows us to use min-cost flow problem algorithms to model even very complicated interactions among candidates in the election.

Bribery’ problems are somewhat reminiscent of Procaccia, Rosenschein, and Kaminka’s (2007) work on the robustness of election systems, which studies the impact of a small number of random switches within rational preference lists on the outcome of the election.

Control

In this section we focus on control in Llull and Copeland elections, and we compare our results to plurality elections.

Theorem 4 *Llull and Copeland elections are resistant and vulnerable to control types as indicated in Table 1. The same results, regarding Llull and Copeland, hold for the case of irrational voters.*

Control type	Llull		Copeland		Plurality	
	CC	DC	CC	DC	CC	DC
AC _u	V	V	V	V	R	R
AC	R	V	R	V	R	R
DC	R	V	R	V	R	R
RPC-TP	R	V	R	V	R	R
RPC-TE	R	V	R	V	R	R
PC-TP	R	V	R	V	R	R
PC-TE	R	V	R	V	R	R
PV-TE	R	R	R	R	V	V
PV-TP	R	R	R	R	R	R
AV	R	R	R	R	V	V
DV	R	R	R	R	V	V

Table 1: Comparison of control results for Llull and Copeland elections and for plurality-rule elections. R means resistance to a particular control type and V means vulnerability. The results for plurality come from Bartholdi, Tovey, and Trick (1992) and Hemaspaandra, Hemaspaandra, and Rothe (2007a).

Note that both Llull and Copeland have a higher number of constructive resistances, by two, than even plurality, which was before this paper the reigning champ. (Although the results regarding plurality in Table 1 regard the unique winner version of control, for all the table’s Llull and Copeland cases our results hold both in the cases of unique winners and of nonunique winners, thus allowing an apples-to-apples comparison to hold.) Admittedly, plurality does perform better with respect to candidate destructive control problems, but still our study of Copeland and Llull makes significant steps forward in the quest for a fully

control-resistant natural election system with an easy winner problem: If one views AC as replacing AC_u in the standard type list, then Llull and Copeland each provide resistance for all ten standard constructive control types. In addition, the hybrid (in the sense of (Hemaspaandra, Hemaspaandra, & Rothe 2007b)) of plurality with either of Llull or Copeland is resistant to each standard type of constructive and destructive control (via Theorem 4 and the results of Hemaspaandra et al. (2007b)). And, unlike the hybrid system constructed by Hemaspaandra et al. (2007b), this hybrid uses only natural systems as its constituents.

Table 1 separates its results into two main groups. In the first group we have results regarding procedural control where the chair affects candidate structure. In the second group the chair affects the voters. All our resistance results regarding candidate control follow via reductions from vertex cover, and all our vulnerability results follow via greedy algorithms. Our resistance results for the case of control by modifying voter structure follow from reductions—often employing our UV technique—from variants of the exact-cover-by-3-sets (X3C, for short) problem. Below we include two proof sketches somewhat indicative of our approaches and techniques.

Let us quickly define a useful notation. Within every election we fix some arbitrary order over the candidates, and any occurrence of a subset D of candidates in a preference list means the candidates from D are listed with respect to that fixed order. Occurrences of \vec{D} mean the same except that the candidates from D are listed in the reverse order.

Lemma 5 (subcase of Theorem 4) *Copeland is resistant to destructive control via deleting voters.*

Proof. Let X3C-odd be a special case of the exact-cover-by-3-sets problem. We are given a set $B = \{b_1, \dots, b_{3k}\}$ and a family of sets $\mathcal{S} = \{S_1, \dots, S_n\}$ such that each S_i is a size-3 subset of B . It is guaranteed that k is odd. The question is whether it is possible to pick k members of \mathcal{S} in such a way that their union is equal to B . It is easy to see that this version of the X3C problem is NP-complete. Our proof follows by a reduction from X3C-odd to Copeland-DCDV.

Let (B, \mathcal{S}) be an X3C-odd instance with $\|B\| = 3k$ and $k = 2q + 1$. We assume that the union of all S_i ’s is B as otherwise (B, \mathcal{S}) is trivially not in X3C-odd. We construct the following Copeland election. The candidate set is $C = \{p, u, v\} \cup B$, where u is the despised candidate whom we want to prevent from being a winner. The voter collection contains $2n(k + 1)$ voters; for each set S_i , we introduce $2k + 2$ voters each of whom has a preference list selected from the following types:

- (i) $u > v > S_i > p > B - S_i$,
- (ii)-1 $\overleftarrow{B - S_i} > p > u > v > \overrightarrow{S_i}$,
- (ii)-2 $\overleftarrow{B - S_i} > u > p > v > \overrightarrow{S_i}$,
- (ii)-3 $\overleftarrow{B - S_i} > v > p > u > \overrightarrow{S_i}$.

For each i , $k + 1$ of S_i ’s voters are of type (i), and each of the remaining $k + 1$ voters for S_i is of some variety of type (ii). We select the type (ii) voters in such a way that globally (i.e., summed over all the voters) there are exactly

q voters of type (ii)-2, q voters of type (ii)-3, and all the remaining $n(k+1) - 2q$ type (ii) voters are of type (ii)-1. It is irrelevant how we allocate the type (ii) voters among S_i 's voters as long as these constraints are met.

We claim that u can be prevented from being a winner of this election via deleting at most k voters if and only if (B, \mathcal{S}) is in X3C-odd.

Before any voters are deleted, we have the following results of head-to-head contests between each pair of candidates. Both u and v defeat p by $2q = k - 1$ votes, u defeats every candidate other than p by more than k votes, and v defeats every candidate other than p and u by more than k votes. For any i , $1 \leq i \leq 3k$, p is tied with b_i . For any i and j , $1 \leq i < j \leq 3k$, b_i is tied with b_j . Thus it holds that u has Copeland score $3k + 2$, v has Copeland score $3k + 1$, and p and all members of B have Copeland score 0.

Since u has more than k votes of advantage over each candidate other than p , via deleting k voters u can lose at most one vote, and this happens only if p defeats u after deleting k voters. Since the highest possible Copeland score in our election is $3k + 2$, to ensure that u is not a winner after deleting k voters we need to guarantee that some candidate other than u defeats everyone else in head-to-head contests. The only candidate that can possibly achieve this is p .

Thus we need to show that p can defeat all other candidates in our election via deleting k voters if and only if (B, \mathcal{S}) is in X3C-odd. Before any deletions, p loses by $k - 1$ votes to both u and v . Thus any deletion of at most k voters that guarantees p defeating both u and v necessarily involves deleting only voters that rank both u and v above p . The only such voters are those of type (i).

Thus if u can be prevented from being a winner then there exists a set V' , $\|V'\| \leq k$, of type (i) voters such that deleting the voters in V' guarantees that p defeats every b_i , $1 \leq i \leq 3k$, by at least one vote. However, regarding B candidates, deleting a single type (i) voter associated with set S_i gives p a one-vote advantage over exactly the members of S_i . Thus V' defines an exact cover by 3-sets over B . On the other hand, if the X3C-odd instance is positive then p clearly can be made a unique winner by deleting k type (i) voters that correspond to a cover. \square

The above proof illustrates some interesting points. For example, even though the proof is matched with a theorem that speaks of destructive control, it in fact can be seen to simultaneously prove that Copeland is resistant to constructive control via deleting voters in the unique winner case. Also, since in this problem the chair is not at liberty to change the preference lists (or tables) of the voters, we automatically get the analogous results for the case of irrational voters.

The proof of Lemma 5 is one of the simplest applications of the UV technique. The idea is that we introduce dummy candidates u and v and set up voters in such a way that to make a particular candidate win, say p , we need to delete (or otherwise change, as in the applications of the UV technique to, e.g., bribery) only the voters who rank p below both u and v . This allows us to introduce whatever padding candidates we need as long as we guarantee that u and v do not both defeat p on the padding voters' preference lists.

The next lemma illustrates a different technique.

Lemma 6 (subcase of Theorem 4) *Copeland is resistant to constructive control via adding a bounded number of candidates.*

Proof. The theorem follows via a reduction from the vertex cover problem to Copeland-CCAC. In the vertex cover problem, our input is an undirected graph $G = (V, E)$ and a nonnegative integer k . The question is if there exists a subset W of V such that $\|W\| \leq k$ and each edge $e \in E$ has at least one of its endpoints in W . We will show how this question can be reduced to asking whether within a certain Copeland election candidate p can be made a winner by adding at most k candidates.

W.l.o.g., we assume that $V = \{v_1, \dots, v_n\}$ and that $E = \{e_1, \dots, e_m\}$ and $k \geq 1$. We construct the following Copeland-CCAC instance: $C = \{p\} \cup E$ is the set of candidates that definitely participate in the election and $D = V$ is the set of spoiler candidates, i.e., the ones that the chair can convince to participate.

We set up the voters' preference lists over $C \cup D$ in such a way that: (a) each e_i defeats p , (b) each e_i defeats all v_i that e_i is not incident to in G , (c) p defeats each v_i , and (d) all the other pairs of candidates are tied in head-to-head contests. It is easy to construct such an election using at most polynomially (in the input size) many voters. For any two candidates c' and c'' and a set of remaining candidates K , we introduce two voters, one with preference list $c' > c'' > K$ and the other one with preference list $\bar{K} > c' > c''$. These two voters give c' two votes advantage over c'' without affecting the relations between any other pair of candidates.

Before adding any of the spoiler candidates each candidate e_i has Copeland score 1 and p has Copeland score 0. Adding a single candidate v_i has the following effect: (a) p gets one additional Copeland point, (b) each e_j candidate that is not incident with v_i gets one additional Copeland point, and (c) v_i has Copeland score 0. Thus the only possibility that p has no less Copeland points than any other candidate in the election is that while adding candidates v_i , each e_j happened to be incident with at least one of the added vertex-candidates. However, this means that the added candidates constitute a vertex cover of G . Thus if p can be made a winner of this election by adding at most k candidates then G has a vertex cover of size at most k . The converse is trivial and the theorem is proven. \square

Resistance to control is generally viewed as a desirable property in system design. However, suppose one is in the role of someone trying to solve resistant control problems. Is there any hope? Though space does not permit a full discussion here, we have obtained a broad range of efficient algorithms for resistant-in-general control problems for the case when the number of candidates or voters is bounded. For example, each of the 16 problems regarding control via voters (Lull or Copeland; constructive or destructive; AV or DV or PV-TE or PV-TP) is in FPT (is fixed-parameter tractable, i.e., is not merely in P but indeed the family of P algorithms have degrees that do not depend on the value of the fixed number of voters or candidates) when the number of candidates is bounded, and also when the number of voters is bounded, and this claim holds even under the succinct

input model (in which the voters are input via (preference-list, binary-integer-giving-frequency-of-that-preference-list pairs). We prove some of these results using Lenstra's algorithm for bounded-variable-cardinality integer programming. On the other hand, for irrational voters, all of the resistant candidate control problems remain resistant even for two voters.

Conclusions

We have shown that from the computational point of view the election systems of Lull and Copeland are broadly resistant to bribery and procedural control, regardless of whether the voters must have rational preferences. It is rather charming that Lull's 700-year-old system shows perfect resistance to bribery and more resistances to (constructive) control than any natural system (even far more modern ones) with an easy winner-determination procedure—other than Copeland—is known to possess, and this is even more remarkable when one considers that Lull's system was defined long before control of elections was even explicitly studied. Copeland voting matches Lull's perfect resistance to bribery and Lull's broad resistance to (constructive) control.

A natural open direction would be to study the complexity of control for additional election systems. Particularly interesting would be to seek existing, natural voting systems that have polynomial-time winner determination procedures but that are resistant to all standard types of both constructive and destructive control. Also extremely interesting would be to find single results that classify, for broad families of election systems, precisely what it is that makes control easy or hard, i.e., to obtain dichotomy meta-results for control (see Hemaspaandra and Hemaspaandra (2007) for some related discussion regarding manipulation).

References

- Bartholdi, III, J., and Orlin, J. 1991. Single transferable vote resists strategic voting. *Social Choice and Welfare* 8(4):341–354.
- Bartholdi, III, J.; Tovey, C.; and Trick, M. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* 6(3):227–241.
- Bartholdi, III, J.; Tovey, C.; and Trick, M. 1992. How hard is it to control an election? *Mathematical and Computer Modeling* 16(8/9):27–40.
- Conitzer, V., and Sandholm, T. 2002a. Complexity of manipulating elections with few candidates. In *Proc. of AAAI-02*, 314–319. AAAI Press.
- Conitzer, V., and Sandholm, T. 2002b. Vote elicitation: Complexity and strategy-proofness. In *Proc. of AAAI-02*, 392–397. AAAI Press.
- Conitzer, V., and Sandholm, T. 2003. Universal voting protocol tweaks to make manipulation hard. In *Proc. of IJCAI-03*, 781–788. Morgan Kaufmann.
- Conitzer, V., and Sandholm, T. 2006. Nonexistence of voting rules that are usually hard to manipulate. In *AAAI-06*. AAAI Press.
- Conitzer, V.; Lang, J.; and Sandholm, T. 2003. How many candidates are needed to make elections hard to manipulate? In *Proc. of the 9th Conf. on Theoretical Aspects of Rationality and Knowledge*, 201–214. ACM Press.
- Dwork, C.; Kumar, R.; Naor, M.; and Sivakumar, D. 2001. Rank aggregation methods for the web. In *Proc. of the 10th Intl. World Wide Web Conf.*, 613–622. ACM Press.
- Elkind, E., and Lipmaa, H. 2005. Small coalitions cannot manipulate voting. In *Proc. of the 9th Intl. Conf. on Financial Cryptography and Data Security*, 285–297. Springer-Verlag *Lecture Notes in Computer Science* #3570.
- Ephrati, E., and Rosenschein, J. 1993. Multi-agent planning as a dynamic search for social consensus. In *Proc. of IJCAI-93*, 423–429. Morgan Kaufmann.
- Erdélyi, G.; Hemaspaandra, L.; Rothe, J.; and Spakowski, H. 2007. On approximating optimal weighted lobbying, and frequency of correctness versus average-case polynomial time. Technical Report TR-914, Department of Computer Science, University of Rochester, Rochester, NY.
- Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. 2006. The complexity of bribery in elections. In *AAAI-06*, 641–646. AAAI Press.
- Hägele, G., and Pukelsheim, F. 2001. The electoral writings of Ramon Lull. *Studia Lulliana* 41(97):3–38.
- Hemaspaandra, E., and Hemaspaandra, L. 2007. Dichotomy for voting systems. *Journal of Computer and System Sciences* 73(1):73–83.
- Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 2007a. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence* 171(5-6):255–285.
- Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 2007b. Hybrid elections broaden complexity-theoretic resistance to control. In *Proc. of IJCAI-07*, 1308–1314. AAAI Press.
- Homan, C., and Hemaspaandra, L. 2006. Guarantees for the success frequency of an algorithm for finding Dodgson-election winners. In *Proceedings of the 31st International Symposium on Mathematical Foundations of Computer Science*, 528–539. Springer-Verlag *Lecture Notes in Computer Science* #4162.
- Procaccia, A., and Rosenschein, J. 2006. Junta distributions and the average-case complexity of manipulating elections. In *Proc. of the 5th Int. Joint Conf. on Autonom. Agents and Multiagent Systems*, 497–504. ACM Press.
- Procaccia, A.; Rosenschein, J.; and Kaminka, G. 2007. On the robustness of preference aggregation in noisy environments. In *Proc. of the 6th Int. Joint Conf. on Autonom. Agents and Multiagent Systems*. ACM Press. To appear, 2007.
- Procaccia, A.; Rosenschein, J.; and Zohar, A. 2007. Multi-winner elections: Complexity of manipulation, control, and winner-determination. In *Proc. of IJCAI-07*, 1476–1481. AAAI Press.