Bypassing Combinatorial Protections: Polynomial-Time Algorithms for Single-Peaked Electorates*

Felix Brandt, Markus Brill, Edith Hemaspaandra, and Lane A. Hemaspaandra

Abstract

For many election systems, bribery (and related) attacks have been shown NP-hard using constructions on combinatorially rich structures such as partitions and covers. It is important to learn how robust these hardness protection results are, in order to find whether they can be relied on in practice. This paper shows that for voters who follow the most central political-science model of electorates—single-peaked preferences—those protections vanish. By using single-peaked preferences to simplify combinatorial covering challenges, we show that NP-hard bribery problems—including those for Kemeny and Llull elections—fall to polynomial time. By using single-peaked preferences to simplify combinatorial partition challenges, we show that NP-hard partition-of-voters problems fall to polynomial time. We furthermore show that for single-peaked electorates, the winner problems for Dodgson and Kemeny elections, though Θ_2^p-complete in the general case, fall to polynomial time. And we completely classify the complexity of weighted coalition manipulation for scoring protocols in single-peaked electorates.

1 Introduction

Elections are perhaps the most important framework for preference aggregation. An election (rule) is a mapping that takes as input the preferences of the voters with respect to the set of candidates (alternatives) and returns a set of "winners," which is some subset of the candidate set. Elections are central in preference aggregation among humans—in everything from political elections to selecting good singers on popular television shows. Elections are rapidly increasing in importance in electronic settings such as multiagent systems, and have been used or proposed for such varied tasks as recommender systems and collaborative filtering [23], web spam reduction and improved web-search engines [12], and planning [13]. In electronic settings, elections may have huge numbers of voters and alternatives.

One natural worry with elections is that agents may try to slant the outcome, for example, by bribing voters. Motivated by work from economics and political science showing that reasonable election systems always allow manipulations of certain types, starting in 1989, Bartholdi, Tovey, and Trick [3, 4] made the thrilling suggestion that elections be protected via complexity theory—namely, by making the attacker's task NP-hard. This line has been active ever since, and has resulted in NP-hardness protections being proven for many election systems, against such attacks as bribery (the attacker has a budget with which to buy and alter voters' votes [16]), manipulation (a coalition of voters wishes to set its votes to make a given candidate win [3]), and control (an agent seeks to make a given candidate win by adding/deleting/partitioning voters or candidates [4]). The book chapter [18] surveys such NP-hardness results, which apply to many important election systems such as plurality, single transferable vote, and approval voting.

In the past few years, a flurry of papers have come out asking whether the NP-hardness protections are satisfying. In particular, the papers explore the possibility that heuristic algorithms may do well frequently or that approximation algorithms may exist. The present paper questions the NP-hardness results from a completely different direction. In political science, perhaps the most

*This paper appeared in the proceedings of AAAI-2010.

"canonical" model of electorates is the unidimensional single-peaked model, in which the electorate has preferences over some one-dimensional spectrum (e.g., "very liberal through very conservative") along which the candidates are also located, and in which each voter's preferences (loosely put) have a peak, with affinity declining as one moves away from the peak. A brilliant paper by Walsh [26] recently asked whether NP-hardness protections against manipulation fall apart if electorates are single-peaked. For the case Walsh looked at, the answer he proved is "no"; he looked at a particular NP-hardness manipulation protection and proved it holds even for single-peaked societies. Faliszewski et al. [17], inspired by Walsh's work, looked at a range of election systems and came to the sharply differing conclusion that for many crucial cases, NP-hardness protections against manipulation and control vanish for single-peaked electorates.

Those two papers [17, 26] are the only two papers we know of that study the implications of single-peakedness on the complexity of manipulation and control. The present paper seeks to take this young line of research in new directions, and to improve one existing direction, via the following contributions:

(1) We show that checking who the winner is in Dodgson, Young, and Kemeny elections, which is Θ_2^p-complete in the general case, is in polynomial time for single-peaked electorates.

(2) We for the first time study the effect of single-peaked electorates on the complexity of *bribery*. We show that many NP-hardness protections against bribery in the general case vanish for single-peaked electorates. To show this, we give polynomial-time bribery algorithms for single-peaked electorates in many settings. Our polynomial-time algorithms apply to approval voting and to the rich range of "weak-Condorcet consistent" election systems and even to systems that are merely known to be weak-Condorcet consistent when the electorate is single-peaked, including weakBlack, weakDodgson, Fishburn, Kemeny, Llull, Maximin, Schwartz, Young, and two variants of Nanson elections.

The practical lesson is that we should be very skeptical about NP-completeness results if our electorate may have limitations (such as single-peakedness) on the ensembles of votes it produces.

(3) We for the first time study the effects of single-peaked electorates on the complexity of *control by partition of voters*, in which the voters are partitioned into two groups that vote on the candidates in "primary" elections, and only the winners of the primaries compete in the final election. This is one of the seven types of control introduced in the seminal control paper of Bartholdi et al. [4], but control by partition of voters has not been previously addressed for the single-peaked case. We show that some known NP-hardness protections against control-by-partition vanish for single-peaked electorates

The shared technical theme here and in the bribery case is that single-peakedness can be used to tame the combinatorial explosion (of partitions and covers) that in the general case protected elections from attack, and in particular single-peakedness yields polynomial-time attack algorithms.

(4) Our final contribution is a strong extension of an important result from Faliszewski et al. [17]. For the broad class of election systems known as scoring protocols, Faliszewski et al. gave a complete characterization of the computational complexity of the (weighted, coalition) manipulation problem in the case of single-peaked elections *with three candidates*. Such characterizations are important as they tell both which systems are manipulable and what it is about the systems that makes them manipulable. We extend this by providing, for single-peaked electorates, a complete characterization of easy manipulability of scoring protocols.

Proofs omitted due to space constraints can be found in the full version of this paper [6].

2 Preliminaries

Election Systems, Preferences, and weakCondorcet-Consistency An election system is a mapping from a finite set of candidates C and a finite list V of voter preferences over those candidates to a collection $W \subseteq C$ called the winner set. For all but one of the election systems we cover, each voter's preference is a linear order (by which we always mean a strict linear order: an irreflexive, antisymmetric, complete, transitive relation) over the candidates. For the election system called approval voting, each voter votes by a bit-vector, approving or disapproving of each candidate separately. Voter's preferences are input as a list of ballots (i.e., votes), so if multiple voters have the same preference, the ballot of each will appear separately in V.

We now very briefly describe the election systems considered in this paper. In *approval voting*, preferences are approval vectors, and each candidate who gets the highest number of approvals among the candidates belongs to the winner set. In all the other systems we use, voters will vote by linear orders. A candidate is said to be a *Condorcet winner* (respectively, *weak Condorcet winner*), if that candidate is preferred to each other candidate by a strict majority (respectively, by at least half) of the voters. In *Condorcet* voting the winners are precisely the set of Condorcet winners. In the election system *weakCondorcet*, the winners are precisely the set of weak Condorcet winners. It has been known for two hundred years that some election instances have neither Condorcet winners nor weak Condorcet winners [7]. And of course, no election instance can have more than one Condorcet winner, whereas there might be several weak Condorcet winners.

For a rational number $\alpha \in [0, 1]$, Copeland$^\alpha$ is the election system where for each pair of distinct candidates we see who is preferred between the two by a strict majority of the voters. That one gets one "Copeland point" from the pairwise contest and the other gets zero "Copeland points." If they tie in their pairwise contest (which can happen only when the number of voters is even), each gets α points. Copeland^1 is known as *Llull* elections, a system defined by the mystic Ramon Llull in the thirteenth century, and is known to be remarkably resistant, computationally, to bribery and control attacks [19].

An important class of elections is the class of *scoring protocols*. Each scoring protocol has a fixed number m of candidates and is defined by a *scoring vector* $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_m)$, $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_m$. Voters' votes are linear orders, and each voter contributes α_1 points to his or her most preferred candidate, α_2 points to his or her next most preferred candidate, and so on. Each candidate whose total number of points is at least as great as the totals of each other candidate is a winner. For example, m-candidate *plurality* voting is the scoring protocol defined by the scoring vector $\alpha = (1, 0, \ldots, 0)$. And m-candidate *Borda* voting is the scoring protocol defined by the scoring vector $\alpha = (m-1, m-2, \ldots, 0)$.

In *Black* elections (respectively, *weakBlack* elections), if there is a Condorcet winner (respectively, if there are weakCondorcet winners), then that defines the winners, and otherwise Borda's method is used to select the winners. Black elections were introduced by Black [5] and weakBlack elections (somewhat confusingly called Black elections there) were introduced by Fishburn [20]. In *Dodgson* elections (respectively, *weakDodgson* elections), whichever candidates can by the fewest repeated transpositions of adjacent candidates in voters' orders become Condorcet winners (respectively, weakCondorcet winners) are the winners. Dodgson elections were introduced in the 1800s by Dodgson and weakDodgson elections (somewhat confusingly called Dodgson elections there) were introduced by Fishburn [20]. In *Young* elections (respectively, *strongYoung* elections), whichever candidates can by the deletion of the fewest voters become weakCondorcet (respectively, Condorcet) winners are the winners. Young elections were introduced by Young and strongYoung elections (somewhat confusingly called Young elections there) were introduced by Rothe et al. [25].

Nanson elections are runoff methods based on Borda's scoring protocol. In Nanson's original definition, a series of Borda elections is held and all candidates who at any stage have not more than the average Borda score are excluded unless all candidates have identical Borda scores, in which case these candidates are declared the winners of the election. There exist two variants of Nanson due to Fishburn and Schwartz, which exclude candidates with the *lowest* Borda score and candidates whose Borda score is *less than* the average score, respectively. *Maximin* (a.k.a. *Simpson*) elections choose those candidates that fare best in their worst pairwise comparison against any other candidate. The remaining three election systems are based on the pairwise majority relation. In *Schwartz* elections (sometimes also called the *top cycle*), the winners are defined as the maximal elements of the asymmetric part of the transitive closure of the majority relation. The winners in *Fishburn* elections are the maximal elements of the Fishburn relation F, which is defined by letting $a\ F\ b$ if every candidate that beats a in a pairwise comparison also beats b and there exists a candidate that beats b but not a. Finally, *Kemeny* elections are based on the smallest number of reversals in the voters' pairwise preferences such that the majority relation becomes transitive and complete. The Kemeny winners are the maximal elements of such minimally modified majority relations.

An important notion in this paper is that of being weakCondorcet-consistent. An election system is said to be *weakCondorcet-consistent* (which we earlier wrote, equivalently, as *weak-Condorcet consistent*), if on every input that has at least one weak Condorcet winner, the winners of the election system are exactly the set of weak Condorcet winners. Some of our bribery results will hold for all election systems that are weakCondorcet-consistent, and even for all election systems that when restricted to single-peaked electorates are weakCondorcet-consistent on those.

Fishburn [20] has noted that the election systems weakBlack, weakDodgson, Fishburn, Maximin, and Young are weakCondorcet-consistent. We add to that the observation that Llull elections are easily seen from their definition to be weakCondorcet-consistent. We also make the (new) observation that the election systems Kemeny, Schwartz, and the two variants of Nanson are weakCondorcet-consistent when restricted to single-peaked electorates. (By Fishburn [20] and Niou, those systems are known not to be weakCondorcet-consistent in the general case.) We also observe that Black, Dodgson, strongYoung, the original version of Nanson, and for each $\alpha \in [0, 1)$, Copeland$^\alpha$ elections are not weakCondorcet-consistent even when restricted to single-peaked electorates.

Single-Peaked Preferences This paper's theme is that combinatorial protections crumble for the case of single-peaked electorates. We now briefly define what single-peaked preferences are and what their motivation is. The single-peaked preference model was introduced over half a century ago by Black [5] and has been influential ever since. The model captures the case where the electorate is polarized by a single issue or dimension, and each voter's utility along that dimension has either one peak or just rises or just falls. Candidates have positions (locations) along that dimension. And a voter's preferences (in the linear order model) simply order the candidates by utility (except with no ties allowed). Since the utility curves are very flexible, what this amounts to is that there is an overall societal ordering L of the candidates, and each voter can be placed in some location such that for all the candidates to his or her right the preferences drop off and the same to the left, although within that framework, the right and the left candidates can be interspersed with each other. A picture will make this clearer. Figure 1 shows an electorate with four voters and five candidates, in which society's polarization is on a (liberal-to-conservative) axis. From the picture, we can see that v_1 has preferences $c_5 > c_4 > c_3 > c_2 > c_1$, v_2 has preferences $c_1 > c_2 > c_3 > c_4 > c_5$, v_3 has preferences (note the interleaving) $c_2 > c_3 > c_1 > c_4 > c_5$, and v_4 has preferences $c_4 > c_5 > c_3 > c_2 > c_1$.

Formally, there are many equivalent ways to capture this behavior, and we use the following as our definition. A collection V of votes (each a linear ordering P_i of the candidates) over candidate set C is said to be *single-peaked* exactly if there exists a linear ordering L over C such that for each triple of candidates a, b, and c, it holds that $(a\,L\,b\,L\,c \vee c\,L\,b\,L\,a) \Rightarrow (\forall i)\,[a\,P_i\,b \Rightarrow b\,P_i\,c]$.

The single-peaked model has been intensely studied, and has both strengths and limitations. On the positive side, it is an excellent rough model for a wide range of elections. Votes on everything from American presidential elections to US Supreme Court votes to hiring votes within a CS department are often shockingly close to reflecting single-peaked preferences. It certainly is a vastly more reasonable model in most settings than is assuming that all voters are random and independent, although the latter model has been receiving a huge amount of study recently. In fact, a wide range of scholarly studies have argued for the value of the single-peaked model [5, 10, 24], and the model is one of the first taught to students in positive (i.e., theoretical) political science courses. On the

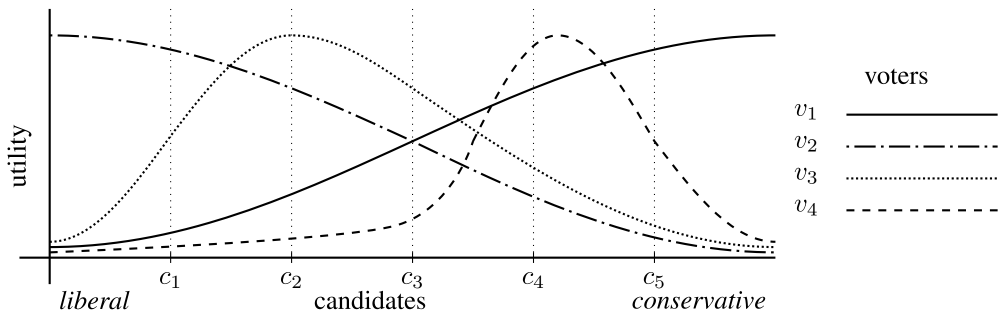

Figure 1: Example of a single-peaked electorate of four voters

other hand, some electorates certainly are driven by multidimensional concerns, and even a heavily unidimensional electorate may have a few out-of-the-box voters.

The single-peaked model also makes sense for approval voting [17]: There, a voter intuitively may be thought to have some utility threshold starting at which he or she approves of candidates. What this means is that each voter's "approved" candidates must be contiguous within society's linear order L.

Although we will assume that society's linear order is part of the input in our single-peaked winner, bribery, manipulation, and control problems, we mention in passing that given an election instance, one can in polynomial time tell whether the voters are single-peaked and when so can also in polynomial time compute a societal linear order instantiating the single-peakedness (Bartholdi and Trick [2] and Doignon and Falmagne [11] for linear-order preferences and Faliszewski et al. [17] for approval preferences).

3 Bypassing Winner-Problem Complexity

The main results sections of this paper study whether single-peakedness bypasses complexity-theoretic protections against attacks on elections. Before moving to those sections, we quickly present some results showing that single-peakedness also bypasses the complexity results some systems have for even telling who won. Unlike the "protection from attack" complexity-shield bypassings, which are in some sense bad news (for the security of the election systems), these "winner-hardness" complexity-shield bypassings are good news—taming the complexity of election systems such as Dodgson and Kemeny for the single-peaked case, despite the fact that they are known to have NP-hard winner problems in the general case.

For a given election system $\mathcal{E}$, the winner problem takes as input an election, (C, V), and a candidate $p \in C$, and asks if p is a winner in the election whose candidates are C and whose votes are V. When we speak of the single-peaked case of the winner problem, the input will also contain a linear order L relative to which the election is single-peaked. Note that the weakCondorcet winner problem is in P in the general case and thus certainly in the single-peaked case. Furthermore, something used often in our paper's proofs is the following standard fact about Condorcet voting and medians.

Fact 1. *Associate each voter with the candidate at the top of that voter's preference ordering. If we order the voters with respect to L in terms of that association, then if $\|V\|$ is odd, the weakCondorcet and Condorcet winner set is the top preference of the median voter, and if $\|V\|$ is even, the weakCondorcet winner set is the set of all candidates who in L fall in the range, inclusively, between the top preferences of the two median voters (and if those two coincide, then that candidate is the Condorcet winner and otherwise there is no Condorcet winner).*

An immediate consequence is the well-known fact that for single-peaked elections, there is always at least one weak Condorcet winner (we are tacitly here assuming $C \neq \emptyset$). Since we earlier noted that the winner problem is in P for weakCondorcet elections, the following holds.

Theorem 1. *For each election system $\mathcal{E}$ that is weakCondorcet-consistent when restricted to single-peaked electorates, the winner problem is in* P *when restricted to single-peaked elections.*

Of course, for many such systems the winner problem is obviously in P even in general. Yet we do get some interesting consequences from Theorem 1.

Corollary 1. *When restricted to single-peaked electorates, the winner problems for Kemeny, Young, and weakDodgson elections are in* P.

In contrast, the general-case Kemeny winner problem problem was proven by Hemaspaandra et al. [22] to be Θ_2^p-complete. And we prove in the full version of this paper that the general-case winner problems for Young and weakDodgson elections are Θ_2^p-complete as well. Thus, Theorem 1 implies sharp complexity simplifications for these three election systems.

The "identify with weakCondorcet" approach that just worked on Young and weakDodgson elections does not apply to Dodgson and strongYoung elections. However, we have constructed direct algorithms that solve their winner problems in polynomial time in the single-peaked case.

Theorem 2. *When restricted to single-peaked electorates, the winner problems for Dodgson and strongYoung elections are in* P.

Our algorithm that shows this for Dodgson elections is a good example of the general technical theme of this paper: That single-peakedness often precludes combinatorial explosion. In this particular case, single-peakedness simplifies the seemingly exponential-sized search space over "series of exchanges to provide upper bounds on Dodgson scores," and will allow us to instead search over a polynomial-sized possibility space related to a particular, simple set of exchanges happening and limited to at most two voters.

Both claims in Theorem 2 contrast directly with the known Θ_2^p-completeness of the general case Dodgson [21] and strongYoung [25] winner problems, and thus reflect a substantial complexity simplification that holds when electorates are single-peaked.

4 Bribery of Single-Peaked Elections

This section shows that single-peakedness undercuts many, although not all, NP-hardness protections for bribery problems.

All bribery notions presented here, except negative approval bribery, are from the paper that started the complexity-theoretic study of bribery [16]. Given an election system $\mathcal{E}$, the $\mathcal{E}$-bribery problem takes as input C, V, $p \in C$, and $k \in \{0, 1, 2, \ldots\}$, and asks if, by changing the votes of at most k members of V, p can be made a winner of this election with respect to $\mathcal{E}$. That is the basic bribery problem. And it can be modified by any combination of the following items: "$" means each voter has a price (belonging to $\{1, 2, 3, \ldots\}$) and the question is whether there is some set of voters whose total price is at most k such that by changing their votes we can make p be a winner. The intuition for prices is that some voters can be swayed more easily than others. "Weighted" means each vote has a weight (belonging to $\{1, 2, 3, \ldots\}$), and each weight w vote is bribed as an indivisible object, but when applying $\mathcal{E}$, is viewed as w identical "regular" votes. For the case where V consists of linear orders, by "negative" we mean that if we bribe a voter then after the bribe the voter must not have p as his or her top choice unless p already was the top choice before the bribe. The intuition is that in negative bribery one is trying to stay under the radar by not directly helping one's candidate. For approval-vector votes, by "negative" we mean that when you bribe a voter, his or her after-bribe vector can approve p only if his or her before-bribe vector approved p. By “strongnegative” we mean that when you bribe a voter the voter after being bribed cannot approve p. These can occur in any combination, e.g., we can speak of Llull-negative-weighted-\$bribery.

When we speak of the single-peaked case of any of the above, we require that all bribes must result only in votes that are consistent with the input societal order L.

4.1 Approval-Bribery Results

As our main result for approval-bribery, we prove that the bribery protection that complexity gives there fails on single-peaked electorates.

Theorem 3 (Faliszewski et al. [16])**.** *Approval-bribery is* NP*-complete.*

Theorem 4. *Approval-bribery is in* P *for single-peaked electorates.*

The specific technical reason we can obtain polynomial-time bribery algorithms is that the NP-hardness proofs were based on the combinatorially rich structure of covering problems (whose core challenge is the “incomparability” of voters), but we use single-peakedness to create a “directional” attack on covering problems that has the effect of locally removing incomparability.

By the same general approach—using a “directional” attack to in the single-peaked setting tame the incomparability challenges of covering problems—we can establish the following two additional cases in which NP-hard bribery problems fall to P for the single-peaked case.

Theorem 5.

1. *Approval-negative-bribery and approval-strongnegative-bribery are* NP*-complete.*
2. *For single-peaked electorates, approval-negative-bribery and approval-strongnegative-bribery are in* P*.*

4.2 Llull-Bribery and Kemeny-Bribery Results

We now state the following eight-case result. The P cases below are proved by direct algorithmic attacks using the connection between weakCondorcet and median voters, and the NP-complete cases are shown by using the problem to capture a partition instance.

Theorem 6. *For single-peaked electorates, weakCondorcet-weighted-\$bribery, weakCondorcet-negative-weighted-bribery, and weakCondorcet-negative-weighted-\$bribery are* NP*-complete, and the remaining five weakCondorcet bribery cases are in* P*.*

Theorem 6 is most interesting not for what it says about weakCondorcet elections, but for its immediate consequences on other election systems, since all weakCondorcet-consistent election systems coincide for single-peaked electorates due to the nonemptiness of the set of weakCondorcet winners.

Corollary 2. *Let $\mathcal{E}$ be any election system that is weakCondorcet-consistent on single-peaked inputs. Then the three* NP*-completeness and five* P *results of Theorem 6 hold (for single-peaked electorates) for $\mathcal{E}$.*

From our discussions earlier in the paper, Corollary 2 applies to the Llull, Kemeny, Young, weakDodgson, Maximin, Schwartz, weakBlack, Fishburn, and the two variants of Nanson election systems. In light of this, Corollary 2 is quietly establishing a large number of claims about NP-hardness shields failing for single-peaked electorates. For example, we have the following claims.

Theorem 7 (Faliszewski et al. [16])**.** *Llull-bribery, Llull-\$bribery, Llull-weighted-bribery, and Llull-weighted-\$bribery are each* NP*-complete.*

Theorem 8 (follows from Corollary 2)**.** *For single-peaked electorates, Llull-bribery, Llull-$bribery, Llull-weighted-bribery, and Llull-weighted-$bribery are each in* P.

To the best of our knowledge, bribery of Kemeny elections has never been studied. Note, however, that the winner problem for any election system $\mathcal{E}$ many-one reduces to each of the eight types of bribery problems mentioned in Theorem 6, except with "weakCondorcet" replaced by "$\mathcal{E}$." This holds because we can ask whether the preferred candidate wins given the bribe limit of 0, and this captures the winner problem. So, from the known Θ_2^p-completeness of the winner problem for Kemeny elections [22], we have the following result, which gives us eight contrasts of hardness (three between Θ_2^p-hardness and NP membership and five between Θ_2^p-hardness and P membership).

Theorem 9 (corollary, in light of the comments just made, to Hemaspaandra et al. [22])**.** *For Kemeny elections, all eight types of bribery mentioned in Theorem 6 are* Θ_2^p*-hard.*

Theorem 10 (follows from Corollary 2)**.** *For single-peaked electorates, Kemeny-weighted-$bribery, Kemeny-negative-weighted-bribery, Kemeny-negative-weighted-$bribery are* NP*-complete, and the remaining five types of bribery of Kemeny elections are in* P.

5 Control of Single-Peaked Electorates

The control problems for elections ask whether by various types of changes in an election's structure a given candidate can be made a winner. The types of control that were introduced by Bartholdi et al. [4], and that (give or take some slight refinements) have been studied in subsequent papers, are addition/deletion/partition of voters/candidates. However, the only previous paper that studied the complexity of control for single-peaked electorates, Faliszewski et al. [17], focused exclusively on additions and deletions of candidates and voters.

We for the first time study the complexity of partition problems for the case of single-peaked electorates. And we show that for a broad range of election systems the control by partition of voters problem is in P for single-peaked electorates. Among the systems we do this for are Llull and Condorcet elections, whose control by partition of voters problem is known to be NP-complete for general electorates. Our proofs again work by using single-peakedness to tame combinatorial explosion—in this case, the number of partitions that must be examined is reduced from an exponential number of partitions to a polynomial number of classes of partitions each of which can be checked as a block.

The control by partition of voters problem for an election system $\mathcal{E}$ takes as input an election instance (C, V) and a candidate $c \in C$ and asks whether there is a partition of votes (V_1, V_2) such that if the "appropriate candidates" move forward from the preliminary elections (C, V_1) and (C, V_2) to a final election in which those candidates are voted on by V, then c "wins." How one clarifies the quoted strings determines the precised type of voter control one studies. In particular, one can study the nonunique-winner model or the unique-winner model. And as to the "appropriate candidates" move forward means, one can study the Ties Promote (TP) model (all winners of the preliminary elections move forward) or the Ties Eliminate (TE) model (only unique winners move forward). Our results hold for all four combinations of these models.

We will briefly mention control results about adding and deleting voters and candidates. The definitions of those are just what one would expect, and we refer the reader to Faliszewski et al. [19] for those definitions. The following is our main result for this section.

Theorem 11. *For weakCondorcet elections, (constructive) control by partition of voters is in* P *for single-peaked electorates.*

The technical challenge here is the exponential number of partitions, and our algorithm circumvents this by using single-peakedness to allow us to in effect structure that huge number of partitions into a polynomial number of classes of partitions such that for each class we can look just at the class rather than having to explore each of its member partitions. Let us note some consequences of this theorem.

Corollary 3. *Let $\mathcal{E}$ be any election system that is weakCondorcet-consistent on single-peaked inputs. Then for election system $\mathcal{E}$, (constructive) control by partition of voters is in* P *for single-peaked electorates. In particular, this holds for the election systems Llull, Kemeny, weakDodgson, Maximin, Schwartz, weakBlack, Fishburn, and the two variants of Nanson.*

For Llull elections, this provides a clear contrast with the known NP-completeness for that same control type in the general case. We now state a result that will quickly give us a number of additional contrasts between general-case control complexity and single-peaked control complexity.

Theorem 12. *For weakCondorcet elections, (constructive) control by adding voters and (constructive) control by deleting voters are each in* P *for single-peaked electorates.*

The full version of this paper contains similar results for Condorcet elections.

6 Manipulation of Single-Peaked Electorates

Faliszewski et al. [17] completely characterized, *for three-candidate elections*, which scoring protocols have polynomial-time constructive coalition weighted manipulation problems and which have NP-complete constructive coalition weighted manipulation problems. We achieve a far more sweeping dichotomy theorem—our result applies to all scoring protocols, regardless of the number of candidates. In the constructive coalition weighted manipulation problem, the input is the candidate set C, the nonmanipulative voters (each a preference order over C and a weight), the manipulative voters (each just a weight), and a candidate $p \in C$, and the question is whether there is a way of setting the preferences of the manipulative voters such that p is a winner under the given election rule when all the manipulative and nonmanipulative voters vote in a weighted election.

Our extension of this three-candidate, single-peaked electorate result to the case of any scoring protocol over single-peaked electorates is somewhat complicated. Yet, since it is a complete characterization—a dichotomization of the complexities, in fact—it is in some sense simply reflecting the subtlety and complexity of scoring systems.

Theorem 13. *Let $\alpha = (\alpha_1, \alpha_2, \ldots, \alpha_m)$ be an m-candidate scoring protocol and consider the constructive coalition weighted manipulation problem for single-peaked electorates.*

- *If $\alpha_2 > \alpha_{\lfloor \frac{m-1}{2} \rfloor + 2}$ and there exist integers $m_1, m_2 > 0$, $i_1, i_2 > 1$ such that $m_1 + m_2 + 1 = m$, $i_1 \le m_1 + 1$, $i_2 \le m_2 + 1$, and $(\alpha_1 - \alpha_{i_1})(\alpha_1 - \alpha_{i_2}) > (\alpha_{i_1} - \alpha_{i_1+1})(\alpha_{i_2} - \alpha_{i_2+1})$, then the problem is* NP*-complete.*
- *If $\alpha_2 = \alpha_{\lfloor \frac{m-1}{2} \rfloor + 2}$ and $\alpha_1 > \alpha_2 > \alpha_m$ and ($\alpha_2 > \alpha_{m-1}$ or $\alpha_1 - \alpha_m > 2(\alpha_2 - \alpha_m)$), then the problem is* NP*-complete.*
- *In all other cases, the problem is in* P.

The "P" cases of Theorem 13's dichotomy align with our theme of single-peakedness often foiling combinatorial protections.

7 Related Work and Additional Discussion

The two papers most related to our work are Walsh [26] and Faliszewski et al. [17]. Walsh's paper first raised the issue of the effect of single-peaked electorates on manipulation, and for the particular case he looked at—weighted coalition manipulation under single transferable vote elections—he showed that manipulation remains hard even for single-peaked electorates. Faliszewski et al. showed cases where single-peakedness removes complexity shields against manipulation, and also opened the study of (nonpartition) control. Our paper in contrast with Walsh's stresses cases where single-peakedness removes combinatorial protections. And we go beyond Faliszewski et al. by for the first time studying bribery of single-peaked electorates and partition-control of single-peaked electorates.

Although [26] and [17] are by far the most related work, other work is much worth mentioning. Bartholdi and Trick [2], Doignon and Falmagne [11], and Escoffier et al. [14] provided efficient algorithms for finding single-peaked orderings. And Conitzer [8] studied the effect of single-peaked electorates on preference elicitation. Two of the papers just mentioned [14, 8] raise the issue of nearly single-peaked electorates, and we commend as a particularly important open issue the question of what effect nearly single-peaked electorates have on complexity.

The literature now contains many papers on the complexity (when single-peaked preferences are not assumed) of manipulation and control (as a pointer to those, see [18] and the citations therein), and contains a few papers on the younger topic of the complexity of bribery (e.g., Faliszewski et al. [16] and Faliszewski et al. [19]). Although the nonunique-winner model and the unique-winner model very typically have the same complexity results, Faliszewski et al. [15] (drawing also on Conitzer et al. [9]) show that this is not always the case.

A worry that comes immediate to the minds of social choice theorists can be expressed as follows: Since it is known that, for single-peaked electorates, "median voting" leaves voters with voting sincerely being an optimal strategy, single-peaked elections are not interesting in terms of other election systems, since median voting should be used. A detailed discussion of this worry would itself fill a paper. But we briefly mention why the above objection is not as serious as it might at first seem. First, the nonmanipulability claims regarding single-peaked elections and median voting are about manipulability, and so say nothing at all about, for example, control. Indeed, weakCondorcet in effect is a type of median voting on single-peaked electorates, and our partition of voters algorithm makes it clear that control can be exercised in interesting ways. Second, even if median voting does have nice properties, the simple truth is that in the real world, society—for virtually all elections and electorates—has chosen (perhaps due to transparency, comfort, or tradition) to use voting systems that clash sharply with median voting. The prominence of plurality voting is the most dramatic such case. So since in the real world we do use a rich range of election systems, it does make sense to understand their behavior. Third, one must be very careful with terms such as "strategy-proof." The paper people most commonly mention as showing that median voting is strategy-proof is Barberà [1]. But that paper's results are about "social choice *functions*" (election rules that always have exactly one winner), not—as this paper is—about election rules that select a set of winners. So one cannot simply assume that for our case median voting (say, weakCondorcet elections) never gives an incentive to misrepresent preferences. We should further stress that discussions of strategy-proofness typically assume that manipulators come in with complete preference orders, but in the Bartholdi et al. [3] model (which this paper and most complexity papers use when studying manipulation), the manipulative coalition is a blank slate with its only goal being to make a certain candidate p be a winner.)

8 Conclusions

The theme of this paper is that single-peaked electorates often tame combinatorial explosion. We saw this first for the case of the winner problem. In that case, this taming is good. It shows that for single-peaked electorates, election systems such as Kemeny have efficient winner algorithms, despite their Θ_2^p-hardness in the general case. But then for bribery and control (and in part, manipulation), we saw many cases where NP-hard problems fell to polynomial time for single-peaked electorates, via algorithms that bypassed the general-case combinatorial explosions of covers and partitions. Since those NP-hardness results were protections against attacks on elections, our results should serve as a warning that those protections are at their very core dependent on the extreme flexibility of voter preference collections the general case allows. For single-peaked electorates, those protections vanish.

Acknowledgments This work was supported in part by DFG grants BR-2312/3-2 and BR-2312/6-1, NSF grants CCF-0426761, IIS-0713061, and CCF-0915792, the European Science Foundation's EUROCORES program LogICCC, and Friedrich Wilhelm Bessel Research Awards to Edith Hemaspaandra and Lane A. Hemaspaandra. This work was done in part during visits of Edith Hemaspaandra and Lane A. Hemaspaandra to the Universities of Dusseldorf and Munich. We thank Piotr Faliszewski, Felix Fischer, Paul Harrenstein, Ariel Procaccia, Jörg Rothe, Hans Georg Seedig, Troels Sørensen, and the anonymous COMSOC and AAAI referees for enjoyable discussions and helpful suggestions.

References

[1] S. Barberà. An introduction to strategy-proof social choice functions. *Social Choice and Welfare*, 18(4): 619–653, 2001.

[2] J. Bartholdi, III and M. Trick. Stable matching with preferences derived from a psychological model. *Operations Research Letters*, 5(4):165–169, 1986.

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

[4] J. Bartholdi, III, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modeling*, 16(8/9):27–40, 1992.

[5] D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.

[6] F. Brandt, M. Brill, E. Hemaspaandra, and L. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. Technical Report TR-955, Dept. of Comp. Sci., Univ. of Rochester, Rochester, NY, 2010.

[7] J.-A.-N. de Caritat, Marquis de Condorcet. *Essai sur l'Application de L'Analyse à la Probabilité des Décisions Rendues à la Pluralité des Voix*. 1785. Facsimile reprint of original published in Paris, 1972, by the Imprimerie Royale.

[8] V. Conitzer. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research*, 35:161–191, 2009.

[9] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):Article 14, 2007.

[10] O. Davis, M. Hinich, and P. Ordeshook. An expository development of a mathematical model of the electoral process. *American Political Science Review*, 54(2):426–448, 1970.

[11] J. Doignon and J. Falmagne. A polynomial time algorithm for unidimensional unfolding representations. *Journal of Algorithms*, 16:218–233, 1994.

[12] 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, 2001.

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

[14] B. Escoffier, J. Lang, and M. Öztürk. Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence*, pages 366–370. IOS Press, 2008.

[15] P. Faliszewski, E. Hemaspaandra, and H. Schnoor. Copeland voting: Ties matter. In *Proceedings of the 7th International Conference on Autonomous Agents and Multiagent Systems*, pages 983–990. International Foundation for Autonomous Agents and Multiagent Systems, 2008.

[16] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. The complexity of bribery in elections. *Journal of Artificial Intelligence Research*, 35:485–532, 2009.

[17] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. In *Proc. of the 12th Conf. on Theoretical Aspects of Rationality and Knowledge*, pages 118–127. ACM, 2009.

[18] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. A richer understanding of the complexity of election systems. In S. Ravi and S. Shukla, editors, *Fundamental Problems in Computing: Essays in Honor of Professor Daniel J. Rosenkrantz*, pages 375–406. Springer, 2009.

[19] P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research*, 35:275–341, 2009.

[20] P. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.

[21] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Exact analysis of Dodgson elections: Lewis Carroll's 1876 voting system is complete for parallel access to NP. *Journal of the ACM*, 44(6):806–825, 1997.

[22] E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.

[23] D. Pennock, E. Horvitz, and C. Giles. Social choice theory and recommender systems: Analysis of the axiomatic foundations of collaborative filtering. In *Proceedings of the 17th National Conference on Artificial Intelligence*, pages 729–734. AAAI Press, 2000.

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

[25] J. Rothe, H. Spakowski, and J. Vogel. Exact complexity of the winner problem for Young elections. *Theory of Computing Systems*, 36(4):375–386, 2003.

[26] T. Walsh. Uncertainty in preference elicitation and aggregation. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, pages 3–8. AAAI Press, 2007.

Felix Brandt and Markus Brill
Institut für Informatik
Technische Universität München
85748 Garching bei München, Germany
Email: `{brandtf,brill}@in.tum.de`

Edith Hemaspaandra
Department of Computer Science
Rochester Institute of Technology
Rochester, NY 14623, USA
Email: `eh@cs.rit.edu`

Lane A. Hemaspaandra
Department of Computer Science
University of Rochester
Rochester, NY 14627, USA
Email: `lane@cs.rochester.edu`