Combinatorial Voter Control in Elections*

Jiehua Chen[1], Piotr Faliszewski[2], Rolf Niedermeier[1], and Nimrod Talmon[1]

[1] Institut für Softwaretechnik und Theoretische Informatik, TU Berlin, Germany
{jiehua.chen@tu-berlin.de,rolf.niedermeier}@tu-berlin.de,
nimrodtalmon77@gmail.com
[2] AGH University of Science and Technology, Krakow, Poland
faliszew@agh.edu.pl

Abstract. Voter control problems model situations such as an external agent trying to affect the result of an election by adding voters, for example by convincing some voters to vote who would otherwise not attend the election. Traditionally, voters are added one at a time, with the goal of making a distinguished alternative win by adding a minimum number of voters. In this paper, we initiate the study of combinatorial variants of control by adding voters: In our setting, when we choose to add a voter v, we also have to add a whole bundle $\kappa(v)$ of voters associated with v. We study the computational complexity of this problem for two of the most basic voting rules, namely the Plurality rule and the Condorcet rule.

1 Introduction

We study the computational complexity of control by adding voters [2, 18], investigating the case where the sets of voters that we can add have some combinatorial structure. The problem of election control by adding voters models situations where some agent (e.g., a campaign manager for one of the alternatives) tries to ensure a given alternative's victory by convincing some undecided voters to vote. Traditionally, in this problem we are given a description of an election (that is, a set C of alternatives and a set V of voters who decided to vote), and also a set W of undecided voters (for each voter in $V \cup W$ we assume to know how this voter intends to vote which is given by a linear order of the set C; we might have good approximation of this knowledge from preelection polls). Our goal is to ensure that our preferred alternative p becomes a winner, by convincing as few voters from W to vote as possible (provided that it is at all possible to ensure p's victory in this way).

Control by adding voters corresponds, for example, to situations where supporters of a given alternative make direct appeals to other supporters of the alternative to vote (for example, they may stress the importance of voting, or

* JC was supported by the Studienstiftung des Deutschen Volkes. PF was supported by the DFG project PAWS (NI 369/10). NT was supported by the DFG Research Training Group "Methods for Discrete Structures" (GRK 1408). This work has been partly supported by COST Action IC1205 on Computational Social Choice.

help with the voting process by offering rides to the voting locations, etc.). Unfortunately, in its traditional phrasing, control by adding voters does not model larger-scale attempts at convincing people to vote. For example, a campaign manager might be interested in airing a TV advertisement that would motivate supporters of a given alternative to vote (though, of course, it might also motivate some of this alternative's enemies), or maybe launch viral campaigns, where friends convince their own friends to vote. It is clear that the sets of voters that we can add should have some sort of a combinatorial structure (e.g., a TV advertisement appeals to a particular group of voters and we can add all of them at the unit cost of airing the advertisement).

The goal of our work is to formally define an appropriate computational problem modeling a combinatorial variant of control by adding voters and to study its computational complexity. Specifically, we focus on the Plurality rule and the Condorcet rule, mainly because the Plurality rule is the most widely used rule in practice, and it is one of the few rules for which the standard variant of control by adding voters is solvable in polynomial time [2], whereas for the Condorcet rule the problem is polynomial-time solvable for the case of single-peaked elections [14]. For the case of single-peaked elections, in essence, all our hardness results for the Condorcet rule directly translate to all Condorcet-consistent voting rules, a large and important family of voting rules. We defer the formal details, definitions, and concrete results to the following sections. Instead, we state the high-level, main messages of our work:

- Many typical variants of combinatorial control by adding voters are intractable, but there is also a rich landscape of tractable cases.
- Assuming that voters have single-peaked preferences does not lower the complexity of the problem (even though it does so in many election problems [6, 9, 14]). On the contrary, assuming single-crossing preferences does lower the complexity of the problem.

We believe that our setting of combinatorial control, and—more generally—combinatorial voting, offers a very fertile ground for future research and we intend the current paper as an initial step.

Related Work. In all previous work on election control, the authors always assumed that one could affect each entity of the election at unit cost only (e.g., one could add a voter at a unit cost; adding two voters always was twice as expensive as adding a single voter). Only the paper of Faliszewski et al. [15], where the authors study control in weighted elections, could be seen as an exception: One could think of adding a voter of weight w as adding a group of w voters of unit weight. On the one hand, the weighted election model does not allow one to express rich combinatorial structures as those that we study here, and on the other hand, in our study we consider unweighted elections only (though adding weights to our model would be seamless).

The specific combinatorial flavor of our model has been inspired by the seminal work of Rothkopf et al. [23] on *combinatorial auctions* (see, e.g., Sandholm [24] for additional information). There, bidders can place bids on combinations of items. While in combinatorial auctions one "bundles" items to bid on, in our scenario one bundles voters.

In the computational social choice literature, combinatorial voting is typically associated with scenarios where voters express opinions over a set of items that themselves have a specific combinatorial structure (typically, one uses CP-nets to model preferences over such alternative sets [5]). For example, Conitzer et al. [10] studied a form of control in this setting. In contrast, we use the standard model of elections (where all alternatives and preference orders are given explicitly), but we have a combinatorial structure of the sets of voters that can be added.

2 Preliminaries

We assume familiarity with standard notions regarding algorithms and complexity theory. For each nonnegative integer z, we write $[z]$ to mean $\{1, \ldots, z\}$.

Elections. An election $E := (C, V)$ consists of a set C of m alternatives and a set V of $|V|$ voters $v_1, v_2, \ldots, v_{|V|}$. Each voter v has a linear order $\succ_v$ over the set C, which we call a *preference order*. We call a voter $v \in V$ a *c-voter* if she prefers $c \in C$ the best. Given a set C of alternatives, if not stated explicitly, we write $\langle C \rangle$ to denote an arbitrary but fixed preference order over C.

Voting Rules. A voting rule $\mathcal{R}$ is a function that given an election E outputs a (possibly empty) set $\mathcal{R}(E) \subseteq C$ of the (tied) election winners. We study the Plurality rule and the Condorcet rule. Given an election, the *Plurality score* of an alternative c is the number of voters that have c at the first position in their preference orders; an alternative is a Plurality winner if it has the maximum Plurality score. An alternative c is a *Condorcet winner* if it beats all other alternatives in head-to-head contests. That is, c is a *Condorcet winner* in election $E = (C, V)$ if for each alternative $c' \in C \setminus \{c\}$ it holds that $|\{v \in V \mid c \succ_v c'\}| > |\{v \in V \mid c' \succ_v c\}|$. Condorcet's rule elects the (unique) Condorcet winner if it exists, and returns an empty set otherwise. A voting rule is *Condorcet-consistent* if it elects a Condorcet winner when there is one (however, if there is no Condorcet winner, then a Condorcet-consistent rule is free to provide any set of winners).

Domain Restrictions. Intuitively, an election is *single-peaked* [4] if it is possible to order the alternatives on a line in such a way that for each voter v the following holds: If c is v's most preferred alternative, then for each two alternatives c_i and c_j that both are on the same side of c (with respect to the ordering of the alternatives on the line), among c_i and c_j, v prefers the one closer to c. For example, single-peaked elections arise when we view the alternatives on the standard political left-right spectrum and voters form their preferences based solely on alternatives' positions on this spectrum. There are polynomial-time algorithms that given an election decide if it is single-peaked and, if so, provide a societal axis for it [1, 13]. *Single-crossing* elections, introduced by Roberts [22], capture a similar idea as single-peaked ones, but from a different perspective. This time we assume that it is possible to order the voters so that for each two alternatives a and b either all voters rank a and b identically, or there is a single point along this order where voters switch from preferring one of the alternatives to preferring the other one. There are polynomial-time algorithms that decide if an election is single-crossing and, if so, produce the voter order witnessing this fact [12, 7].

Combinatorial Bundling Functions. Given a voter set X, a combinatorial bundling function $\kappa : X \to 2^X$ (abbreviated as *bundling function*) is a function assigning to each voter a subset of voters. For convenience, for each subset $X' \subseteq X$, we let $\kappa(X') = \bigcup_{x \in X'} \kappa(x)$. For $x \in X$, $\kappa(x)$ is called x's *bundle* (and for this bundle, x is called its *leader*). We assume that $x \in \kappa(x)$ and so $\kappa(x)$ is never empty. We typically write b to denote the maximum bundle size under a given κ (which will always be clear from context). Intuitively, we use combinatorial bundling functions to describe the sets of voters that we can add to an election at a unit cost. For example, one can think of $\kappa(x)$ as the group of voters that join the election under x's influence.

We are interested in various special cases of bundling functions. We say that κ is *leader-anonymous* if for each two voters x and y with the same preference order $\kappa(x) = \kappa(y)$ holds. Furthermore, κ is *follower-anonymous* if for each two voters x and y with the same preference orders, and each voter z, it holds that $x \in \kappa(z)$ if and only if $y \in \kappa(z)$. We call κ *anonymous* if it is both leader-anonymous and follower-anonymous. The swap distance between two voters v_i and v_j is the minimum number of swaps of consecutive alternatives that transform v_i's preference order into that of v_j. Given a number $d \in \mathbb{N}$, we call κ a *full-d bundling function* if for each $x \in X$, $\kappa(x)$ is exactly the set of all $y \in X$ such that the swap distance between the preference orders of x and y is at most d.

Central Problem. We consider the following problem for a given voting rule $\mathcal{R}$:

> $\mathcal{R}$ Combinatorial Constructive Control by Adding Voters ($\mathcal{R}$-C-CC-AV)
> **Input:** An election $E = (C, V)$, a set W of (unregistered) voters with $V \cap W = \emptyset$, a bundling function $\kappa : W \to 2^W$, a preferred alternative $p \in C$, and a bound $k \in \mathbb{N}$.
> **Question:** Is there a subset of voters $W' \subseteq W$ of size at most k such that $p \in \mathcal{R}(C, V \cup \kappa(W'))$, where $\mathcal{R}(C, X)$ is the set of winners of the election (C, X) under the rule $\mathcal{R}$?

We note that we use here a so-called nonunique-winner model. For a control action to be successful, it suffices for p to be one of the tied winners. Throughout this work, we refer to the set W' of voters such that p wins election $(C, V \cup \kappa(W'))$ as the solution and denote k as the solution size.

$\mathcal{R}$-C-CC-AV is a generalization of the well-studied problem $\mathcal{R}$ Constructive Control by Adding Voters ($\mathcal{R}$-CC-AV) (in which κ is fixed so that for each $w \in W$ we have $\kappa(w) = \{w\}$). The non-combinatorial problem CC-AV is polynomial-time solvable for the Plurality rule [2], but is NP-complete for the Condorcet rule [20].

	m	n	k	b	d
# alternatives (m)	Non-anonymous: W[2]-h wrt. k even if $m = 2$ [Thm. 2] Anonymous: ILP-FPT wrt. m [Thm. 3]				
# unreg. voters (n)		FPT wrt. n			
solution size (k)			XP [Obs. 1] Anonymous & $b = 3$: W[1]-h wrt. k [Thm. 1]		Single-peaked & full-1 κ: W[1]-h wrt. k [Thm. 8]
max. bundle size (b)				$b = 2$: NP-h [Thm. 4] and P for full-d κ [Thm. 5] $b = 3$: NP-h even for full-d κ [Thm. 6] $b \geq 4$: NP-h even for full-1 κ [Thm. 7]	
max. swap dist. (d)					$d = 1$: W[1]-h wrt. k [Thm. 8] Single-crossing & full-d κ: P [Thm. 9]

Table 1. Computational complexity classification of Plurality-C-CC-AV (since the non-combinatorial problem CC-AV is already NP-hard for Condorcet's rule, we concentrate here on the Plurality rule). Each row and column in the table corresponds to a parameter such that each cell contains results for the two corresponding parameters combined. Due to symmetry, there is no need to consider the cells under the main diagonal, therefore they are painted in gray. ILP-FPT means FPT based on a formulation as an integer linear program.

Parameterized Complexity. An instance (I, k) of a parameterized problem consists of the actual instance I and an integer k being the *parameter* [11]. A parameterized problem is called *fixed-parameter tractable* (is in FPT) if there is an algorithm solving it in $f(k) \cdot |I|^{O(1)}$ time, for an arbitrary computable function f only depending on parameter k, whereas an algorithm with running-time $|I|^{f(k)}$ only shows membership in the class XP (clearly, FPT $\subseteq$ XP). One can show that a parameterized problem L is (presumably) not fixed-parameter tractable by devising a *parameterized reduction* from a W[1]-hard or a W[2]-hard problem to L. A parameterized reduction from a parameterized problem L to another parameterized problem L' is a function that, given an instance (I, k), computes in $f(k) \cdot |I|^{O(1)}$ time an instance (I', k'), such that $k' \leq g(k)$ and $(I, k) \in L \Leftrightarrow (I', k') \in L'$. Betzler et al. [3] survey parameterized complexity investigations in voting.

Our Contributions. As $\mathcal{R}$-C-CC-AV is generally NP-hard even for $\mathcal{R}$ being the Plurality rule, we show several fixed-parameter tractability results for some of the natural parameterizations of $\mathcal{R}$-C-CC-AV; we almost completely resolve the complexity of C-CC-AV, for the Plurality rule and the Condorcet rule, as a function of the maximum bundle size b and the maximum distance d from a voter v to the farthest element of her bundle. Further, we show that the problem remains hard even when restricting the elections to be single-peaked, but that it is polynomial-time solvable when we focus on single-crossing elections. Our results for Plurality elections are summarized in Table 1.

3 Complexity for Unrestricted Elections

In this section we provide our results for the case of unrestricted elections, where voters may have arbitrary preference orders. In the next section we will consider single-peaked and single-crossing elections that only allow "reasonable" preference orders.

Number of Voters, Number of Alternatives, and Solution Size. We start our discussion by considering parameters "the number m of alternatives", "the number n of unregistered voters", and "the solution size k". A simple brute-force algorithm, checking all possible combinations of k bundles, proves that both Plurality-C-CC-AV and Condorcet-C-CC-AV are in XP for parameter k, and in FPT for parameter n (the latter holds because $k \leq n$). Indeed, the same result holds for all voting rules that are XP/FPT-time computable for the respective parameters.

Observation 1 *Both* Plurality-C-CC-AV *and* Condorcet-C-CC-AV *are solvable in* $O(n^k \cdot n \cdot m \cdot \text{winner})$ *time, where* winner *is the complexity of determining Plurality/Condorcet winners.*

The XP result for Plurality-C-CC-AV with respect to the parameter k probably cannot be improved to fixed-parameter tractability. Indeed, for parameter k we show that the problem is W[1]-hard, even for anonymous bundling functions and for maximum bundle size three.

Theorem 1. Plurality-C-CC-AV *is* NP*-hard and* W[1]*-hard when parameterized by the solution size k, even when the maximum bundle size b is three and the bundling function is anonymous.*

Proof (Sketch). We provide a parameterized reduction from the W[1]-hard problem Clique parameterized by the parameter h [11], which asks for the existence of a complete subgraph with h vertices in an input graph G.

Let (G, h) be a Clique instance. Without loss of generality, we assume that G is connected, that $h \geq 3$, and that each vertex in G has degree at least $h-1$. We construct an election $E = (C, V)$ with $C := \{p, w, g\} \cup \{c_e \mid e \in E(G)\}$, and set p to be the preferred alternative. The registered voter set V consists of $\binom{h}{2} + h$ voters each with preference order $w \succ \langle C \setminus \{w\} \rangle$, another $\binom{h}{2}$ voters each with preference order $g \succ \langle C \setminus \{g\} \rangle$, and another h voters each with preference order $p \succ \langle C \setminus \{p\} \rangle$. For each vertex $u \in V(G)$, we define $C(u) := \{c_e \mid e \in E(G) \wedge u \in e\}$, and construct the set W of unregistered voters as follows:

(i) For each vertex $u \in V(G)$, we add an unregistered g-voter w_u with preference order $g \succ \langle C(u) \rangle \succ \langle C \setminus (\{g\} \cup C(u)) \rangle$ and we set $\kappa(w_u) = \{w_u\}$.

(ii) For each edge $e = \{u, u'\} \in E(G)$, we add an unregistered p-voter w_e with preference order $p \succ c_e \succ \langle C \setminus \{p, c_e\} \rangle$ and we set $\kappa(w_e) = \{w_u, w_{u'}, w_e\}$.

Since all the unregistered voters have different preference orders (this is so because G is connected, $h \geq 3$, and each vertex has degree at least $h-1$), every bundling function for our instance is anonymous. Finally, we set $k := \binom{h}{2}$. □

If we drop the anonymity requirement for the bundling function, then we obtain a stronger intractability result. For parameter k, the problem becomes W[2]-hard, even for two alternatives. This is quite remarkable because typically election problems with a small number of alternatives are easy (they can be solved either through brute-force attacks or through integer linear programming attacks employing the famous FPT algorithm of Lenstra [19]; see the survey of Betzler et al. [3] for examples, but note that there are also known examples of problems where a small number of alternatives does not seem to help [8]). Further, since our proof uses only two alternatives, it applies to almost all natural voting rules: For two alternatives almost all of them (including the Condorcet rule) are equivalent to the Plurality rule. The reduction is from the W[2]-complete problem SET COVER parameterized by the solution size [11].

Theorem 2. *Both* PLURALITY-C-CC-AV *and* CONDORCET-C-CC-AV *parameterized by the solution size k are* W[2]*-hard, even for two alternatives.*

If we require the bundling function to be anonymous, then C-CC-AV can be formulated as an integer linear program where the number of variables and the number of constraints are bounded by some function in the number m of alternatives. Hence, C-CC-AV is fixed-parameter tractable due to Lenstra [19].

Theorem 3. *For anonymous bundling functions, both* PLURALITY-C-CC-AV *and* CONDORCET-C-CC-AV *parameterized by the number m of alternatives are fixed-parameter tractable.*

Combinatorial Parameters. We focus now on the complexity of PLURALITY-C-CC-AV as a function of two combinatorial parameters: (a) the maximum swap distance d between the leader and his followers in one bundle, and (b) the maximum size b of each voter's bundle.

First, if $b = 1$, then C-CC-AV reduces to CC-AV and, thus, can be solved by a greedy algorithm in polynomial time [2]. However, for arbitrary bundling functions, PLURALITY-C-CC-AV becomes intractable as soon as $b = 2$.

Theorem 4. PLURALITY-C-CC-AV *is* NP*-hard even if the maximum bundle size b is two.*

Proof (Sketch). We reduce from a restricted variant of 3SAT, where each clause has either two or three literals, each variable occurs exactly four times, twice as a positive literal, and twice as a negative literal. This variant is still NP-hard (the proof is analogous to the one shown for [25, Theorem 2.1]). Given such a restricted 3SAT instance $(\mathcal{C}, \mathcal{X})$, where $\mathcal{C}$ is the set of clauses over the set of variables $\mathcal{X}$, we construct an election (C, V) with $C := \{p, w\} \cup \{c_i \mid C_i \in \mathcal{C}\}$; we call c_i the *clause* alternatives. We set $k := 4|\mathcal{X}|$. We construct the set V such that the initial score of w is $4|\mathcal{X}|$, the initial score of each clause alternative c_i is $4|\mathcal{X}| - |C_i| + 1$, and the initial score of p is zero. We construct the set W of unregistered voters as follows (we will often write ℓ_j to refer to a literal that contains variable x_j; depending on the context, ℓ_j will mean either x_j or $\neg x_j$ and the exact meaning will always be clear):

1. for each variable $x_j \in \mathcal{X}$, we construct four p-voters, denoted by $p_1^j, p_2^j, p_3^j, p_4^j$;
2. for each clause $C_i \in \mathcal{C}$ and each literal ℓ contained in C_i, we construct a c_i-voter, denoted by c_i^ℓ; we call such voter a *clause* voter.

We define the assignment function κ as follows: For each variable $x_j \in \mathcal{X}$ that occurs as a negative literal $(\neg x_j)$ in clauses C_i and C_s, and as a positive literal (x_j) in clauses C_r and C_t, we set $\kappa(p_1^j) = \{p_1^j, c_i^{\neg x_j}\}$, $\kappa(c_i^{\neg x_j}) = \{c_i^{\neg x_j}, p_2^j\}$, $\kappa(p_2^j) = \{p_2^j, c_r^{x_j}\}$, $\kappa(c_r^{x_j}) = \{c_r^{x_j}, p_3^j\}$, $\kappa(p_3^j) = \{p_3^j, c_s^{\neg x_j})\}$, $\kappa(c_s^{\neg x_j}) = \{c_s^{\neg x_j}, p_4^j\}$, $\kappa(p_4^j) = \{p_4^j, c_t^{x_j}\}$, $\kappa(c_t^{x_j}) = \{c_t^{x_j}, p_1^j\}$.

The general idea is that in order to let p win, all p-voters must be in $\kappa(W')$ and no clause alternative should gain more than $(|C_i|-1)$ points. We now show that if $(\mathcal{C}, \mathcal{X})$ has a satisfying truth assignment, then there is a size-k subset $W' \subseteq W$ such that p wins the election $(C, V \cup \kappa(W'))$ (recall that $k = 4|\mathcal{X}|$). The proof for the reverse direction is omitted.

Let $\beta : \mathcal{X} \to \{T, F\}$ be a satisfying truth assignment function for $(\mathcal{C}, \mathcal{X})$. Intuitively, β will guide us through constructing the set W' in the following way: First, for each variable x_j, we put into W' those voters $c_i^{\ell_j}$ for whom β sets ℓ_j to false (this way in $\kappa(W')$ we include $2|\mathcal{X}|$ p-voters and, for each clause c_i, at most $(|C_i|-1)$ c_i-voters). Then, for each clause voter $c_i^{\ell_j}$ already in W', we also add the voter p_a^j, $1 \leq a \leq 4$, that contains $c_i^{\ell_j}$ in his or her bundle (this way we include in $\kappa(W')$ additional $2|\mathcal{X}|$ p-voters without increasing the number of clause voters included). Formally, we define W' as follows: $W' := \{c_i^{\neg x_j}, p_a^j \mid \neg x_j \in C_i \wedge \beta(x_j) = T \wedge c_i^{\neg x_j} \in \kappa(p_a^j)\} \cup \{c_i^{x_j}, p_a^j \mid x_j \in C_i \wedge \beta(x_j) = F \wedge c_i^{\neg x_j} \in \kappa(p_a^j)\}$. As per our intuitive argument, one can verify that all p-voters are contained in $\kappa(W')$ and each clause alternative c_i gains at most $(|C_i|-1)$ points. □

The situation is different for full-d bundling functions, because we can extend the greedy algorithm by Bartholdi et al. [2] to bundles of size two.

Theorem 5. *If κ is a full-d bundling function and the maximum bundle size b is two, then* Plurality-C-CC-AV *is polynomial-time solvable.*

However, as soon as $b = 3$, we obtain NP-hardness, by modifying the reduction used in Theorem 4.

Theorem 6. *If κ is a full-d bundling function, then* Plurality-C-CC-AV *is* NP*-hard even if the maximum bundle size b is three.*

Taking also the swap distance d into account, we find out that both Plurality-C-CC-AV and Condorcet-C-CC-AV are NP-hard, even if $d = 1$. This stands in contrast to the case where $d = 0$, where $\mathcal{R}$-C-CC-AV reduces to the CC-AV problem (perhaps for the weighted voters [15]), which, for Plurality voting, is polynomial-time solvable by a simple greedy algorithm.

Theorem 7. Plurality-C-CC-AV *is* NP*-hard even for full-*1 *bundling functions and even if the maximum bundle size b is four.*

4 Single-Peaked and Single-Crossing Elections

In this section, we focus on instances with full-d bundling functions, and we do so because without this restriction the hardness results from previous sections easily translate to our restricted domains (at least for the case of the Plurality rule). We find that the results for the combinatorial variant of control by adding voters for single-peaked and single-crossing elections are quite different than those for the non-combinatorial case. Indeed, both for Plurality and for Condorcet, the voter control problems for single-peaked elections and for single-crossing elections are solvable in polynomial time for the non-combinatorial case [6, 14, 21]. For the combinatorial case, we show hardness for both Plurality-C-CC-AV and Condorcet-C-CC-AV for single-peaked elections, but give polynomial-time algorithms for single-crossing elections. We mention that the intractability results can also be seen as regarding anonymous bundling functions because all full-d bundling functions are leader-anonymous and follower-anonymous.

Theorem 8. *Both* Plurality-C-CC-AV *and* Condorcet-C-CC-AV *parameterized by the solution size k are* W[1]*-hard for single-peaked elections, even for full-*1 *bundling functions.*

Proof (Sketch). We provide a parameterized reduction from the W[1]-complete Partial Vertex Cover (PVC) parameterized by "solution size" h [17], which asks for a set of at most h vertices in a graph G, which intersects with at least ℓ edges. Given a PVC instance (G, h, ℓ), we set $k := h$, construct an election $E = (C, V)$ with $C := \{p, w\} \cup \{a_i, \overline{a}_i, b_i, \overline{b}_i \mid u_i \in V(G)\}$, and set p to be the preferred alternative, such that the initial score of w is $h + \ell$, and is zero for all other alternatives. We do so by creating $h + \ell$ registered voters who all have the same preference order $\succ$ such that it differs from the following *canonical preference order*: $p \succ w \succ a_1 \succ \overline{a}_1 \succ \ldots \succ a_{|V(G)|} \succ \overline{a}_{|V(G)|} \succ b_1 \succ \overline{b}_1 \succ \ldots \succ b_{|V(G)|} \succ \overline{b}_{|V(G)|}$ by only the first pair $\{p, w\}$, which is swapped.

For each set P of disjoint pairs of alternatives, neighboring with respect to the canonical preference order, we define the preference order diff-order(P) to be identical to the canonical preference order, except that all the pairs of alternatives in P are swapped. The unregistered voter set W consists of the following three types of voters:

(i) for each edge $e = \{u_i, u_j\} \in E(G)$ we have an *edge voter* w_e with preference order diff-order$(\{\{a_i, \overline{a}_i\}, \{a_j, \overline{a}_j\}\})$,
(ii) for each edge $e = \{u_i, u_j\} \in E(G)$ we have a *dummy voter* d_e with preference order diff-order$(\{\{p, w\}, \{a_i, \overline{a}_i\}, \{a_j, \overline{a}_j\}\})$, and
(iii) for each vertex $u_i \in V(G)$ we have a *vertex voter* w_i^u with preference order diff-order$(\{\{a_i, \overline{a}_i\}\})$.

The preference orders of the voters in $V \cup W$ are single-peaked with respect to the axis $\langle\overline{B}\rangle \succ \langle\overline{A}\rangle \succ p \succ w \succ \langle A\rangle \succ \langle B\rangle$, where $\langle\overline{B}\rangle := \overline{b}_{|V(G)|} \succ \overline{b}_{|V(G)|-1} \succ \ldots \succ \overline{b}_1$, $\langle\overline{A}\rangle := \overline{a}_{|V(G)|} \succ \overline{a}_{|V(G)|-1} \succ \ldots \succ \overline{a}_1$, $\langle A\rangle := a_1 \succ a_2 \succ \ldots \succ a_{|V(G)|}$, and $\langle B\rangle := b_1 \succ b_2 \succ \ldots \succ b_{|V(G)|}$. Finally, we define the function κ such that it is a full-1 bundling function. □

We now present some tractability results for single-crossing elections. Consider an $\mathcal{R}$-C-CC-AV instance $((C, V), W, d, \kappa, p \in C, k)$ such that $(C, V \cup W)$ is single-crossing. This has a crucial consequence for full-d bundling functions: For each unregistered voter $w \in W$, the voters in bundle $\kappa(w)$ appear consecutively along the single-crossing order restricted to only the voters in W.[1] Using the following lemmas, we can show that Plurality-C-CC-AV and Condorcet-C-CC-AV are polynomial-time solvable in some cases.

Lemma 1. *Let $I = ((C, V), W, d, \kappa, p \in C, k)$ be a* Plurality-C-CC-AV *instance such that $(C, V \cup W)$ is single-crossing and κ is a full-d bundling function. Then, the following statements hold:*

(i) The p-voters are ordered consecutively along the single-crossing order.
(ii) If I is a yes instance, then there is a subset $W' \subseteq W$ of size at most k such that all bundles of voters $w \in W'$ contain only p-voters, except at most two bundles which may contain some non-p-voters.

Lemma 2. *Let $(C, V \cup \kappa(W'))$ be a single-crossing election with single-crossing voter order $\langle x_1, x_2, \ldots, x_z\rangle$ and set $X_{\text{median}} := \{x_{\lceil z/2\rceil}\} \cup \{x_{z/2+1}$ if z is even$\}$, where $z = |V| + |\kappa(W')|$. Alternative p is a (unique) Condorcet winner in $(C, V \cup \kappa(W'))$ if and only if every voter in X_{median} is a p-voter.*

Theorem 9. *Both* Plurality-C-CC-AV *and* Condorcet-C-CC-AV *are polynomial-time solvable for the single-crossing case with full-d bundling functions.*

Proof. First, we find a (unique) single-crossing voter order for $(C, V \cup W)$ in quadratic time [12, 7]. Due to Lemma 1 and Lemma 2, we only need to store the most preferred alternative of each voter to find the solution set W'. Thus, the running-time from now on only depends on the number of voters. We start with the Plurality rule and let $\alpha := \langle w_1, w_2, \ldots, w_{|W|}\rangle$ be a single-crossing voter order.

Due to Lemma 1 (ii), the two bundles in $\kappa(W')$ which may contain non-p-voters appear at the beginning and at the end of the p-voter block, along the single-crossing order. We first guess these two bundles, and after this initial guess, all remaining bundles in the solution contain only p-voters (Lemma 1 (i)). Thus, the remaining task is to find the maximum score that p can gain by selecting k' bundles containing only p-voters. This problem is equivalent to the Maximum Interval Cover problem, which is solvable in $O(|W|^2)$ time (Golab et al. [16, Section 3.2]).

For the Condorcet rule, we propose a slightly different algorithm. The goal is to find a minimum-size subset $W' \subseteq W$ such that p is the (unique) Condorcet winner in $(C, V \cup \kappa(W'))$. Let $\beta := \langle x_1, x_2, \ldots, x_z\rangle$ be a single-crossing voter order for $(C, V \cup W)$. Considering Lemma 2, we begin by guessing at most two voters in $V \cup W$ whose bundles may contain the median p-voter (or, possibly, several p-voters) along the single-crossing order of voters restricted to the final

[1] Note that for each single-crossing election, the order of the voters possessing the single-crossing property is, in essence, unique.

election (for simplicity, we define the bundle of each registered voter to be its singleton). The voters in the union of these two bundles must be consecutively ordered. Let those voters be $x_i, x_{i+1}, \ldots, x_{i+j}$ (where $i \geq 1$ and $j \geq 0$), let $W_1 := \{x_s \in W \mid s < i\}$, and let $W_2 := \{x_s \in W \mid s > i + j\}$. We guess two integers $z_1 \leq |W_1|$ and $z_2 \leq |W_1|$ with the property that there are two subsets $B_1 \subseteq W_1$ and $B_2 \subseteq W_2$ with $|B_1| = z_1$ and $|B_2| = z_2$ such that the median voter(s) in $V \cup B_1 \cup \{x_i, x_{i+1}, \ldots, x_{i+j}\} \cup B_2$ are indeed p-voters (for now, only the sizes z_1 and z_2 matter, not the actual sets). These four guesses cost $O(|V \cup W|^2 \cdot |W|^2)$ time. The remaining task is to find two minimum-size subsets W'_1 and W'_2 such that $\kappa(W'_1) \subseteq W_1$, $\kappa(W'_2) \subseteq W_2$, $|\kappa(W'_1)| = z_1$, and $|\kappa(W'_2)| = z_2$. As already discussed, this can be done in $O(|W|^2)$ time [16]. We conclude that one can find a minimum-size subset $W' \subseteq W$ such that p is the (unique) Condorcet winner in $(C, V \cup \kappa(W'))$ in $O(|V \cup W|^2 \cdot |W|^4)$ time. □

5 Conclusion

We provide opportunities for future research. First, we did not discuss destructive control and the related problem of combinatorial deletion of voters. For Plurality, we conjecture that combinatorial addition of voters for destructive control, and combinatorial deletion of voters for either constructive or destructive control behave similarly to combinatorial addition of voters for constructive control.

Another, even wider field of future research is to study other combinatorial voting models—this may include controlling the swap distance, "probabilistic bundling", "reverse bundling", or using other distance measures than the swap distance. Naturally, it would also be interesting to consider other problems than election control (with bribery being perhaps the most natural candidate).

Finally, instead of studying a "leader-follower model" as we did, one might also be interested in an "enemy model" referring to control by adding alternatives: The alternatives of an election "hate" each other such that if one alternative is added to the election, then all of its enemies are also added to the election.

Bibliography

[1] J. J. Bartholdi, III and M. Trick. Stable matching with preferences derived from a psychological model. *Oper. Res. Lett.*, 5(4):165–169, 1986.

[2] J. J. Bartholdi, III, C. A. Tovey, and M. A. Trick. How hard is it to control an election. *Math. Comput. Model.*, 16(8–9):27–40, 1992.

[3] N. Betzler, R. Bredereck, J. Chen, and R. Niedermeier. Studies in computational aspects of voting—a parameterized complexity perspective. In *The Multivariate Algorithmic Revolution and Beyond*, volume 7370 of *LNCS*, pages 318–363. Springer, 2012.

[4] D. Black. On the rationale of group decision making. *J. Polit. Econ.*, 56(1):23–34, 1948.

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

[6] F. Brandt, M. Brill, E. Hemaspaandra, and L. A. Hemaspaandra. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. In *Proc. 24th AAAI*, pages 715–722, 2010.

[7] R. Bredereck, J. Chen, and G. Woeginger. A characterization of the single-crossing domain. *Soc. Choice Welf.*, 41(4):989–998, 2013.

[8] R. Bredereck, J. Chen, P. Faliszewski, A. Nichterlein, and R. Niedermeier. Prices matter for the parameterized complexity of shift bribery. In *Proc. 28th AAAI*, 2014. To appear.

[9] V. Conitzer. Eliciting single-peaked preferences using comparison queries. *J. Artificial Intelligence Res.*, 35:161–191, 2009.

[10] V. Conitzer, J. Lang, and L. Xia. How hard is it to control sequential elections via the agenda? In *Proc. 21st IJCAI*, pages 103–108, July 2009.

[11] R. G. Downey and M. R. Fellows. *Fundamentals of Parameterized Complexity.* Springer, 2013.

[12] E. Elkind, P. Faliszewski, and A. Slinko. Clone structures in voters' preferences. In *Proc. 13th EC*, pages 496–513, 2012.

[13] B. Escoffier, J. Lang, and M. Öztürk. Single-peaked consistency and its complexity. In *Proc. 18th ECAI*, pages 366–370, 2008.

[14] P. Faliszewski, E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Inform. and Comput.*, 209(2):89–107, 2011.

[15] P. Faliszewski, E. Hemaspaandra, and L. A. Hemaspaandra. Weighted electoral control. In *Proc. 12th AAMAS*, pages 367–374, 2013.

[16] L. Golab, H. Karloff, F. Korn, A. Saha, and D. Srivastava. Sequential dependencies. In *35th PVLDB*, volume 2(1), pages 574–585, 2009.

[17] J. Guo, R. Niedermeier, and S. Wernicke. Parameterized complexity of Vertex Cover variants. *Theory Comput. Syst.*, 41(3):501–520, 2007.

[18] E. Hemaspaandra, L. A. Hemaspaandra, and J. Rothe. Anyone but him: The complexity of precluding an alternative. *Artif. Intell.*, 171(5–6):255–285, 2007.

[19] H. W. Lenstra. Integer programming with a fixed number of variables. *Math. Oper. Res.*, 8(4):538–548, 1983.

[20] H. Liu, H. Feng, D. Zhu, and J. Luan. Parameterized computational complexity of control problems in voting systems. *Theor. Comput. Sci.*, 410:2746–2753, 2009.

[21] K. Magiera and P. Faliszewski. How hard is control in single-crossing elections? In *Proc. 21st ECAI*, 2014. To appear.

[22] K. W. Roberts. Voting over income tax schedules. *J. Public Econ.*, 8:329–340, 1977.

[23] M. H. Rothkopf, A. Pekeč, and R. M. Harstad. Computationally manageable combinational auctions. *Manage. Sci.*, 44(8):1131–1147, 1998.

[24] T. Sandholm. Optimal winner determination algorithms. In Cramton, Shoham, and Steinberg, editors, *Combinatorial Auctions*, chapter 14. MIT Press, 2006.

[25] C. A. Tovey. A simplified NP-complete satisfiability problem. *Discrete Appl. Math.*, 8(1):85—89, 1984.