

Typical-Case Challenges to Complexity Shields That Are Supposed to Protect Elections Against Manipulation and Control: A Survey*

Jörg Rothe

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

Lena Schend

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

Abstract

In the context of voting, manipulation and control refer to attempts to influence the outcome of elections by either setting some of the votes strategically (i.e., by reporting untruthful preferences) or by altering the structure of elections via adding, deleting, or partitioning either candidates or voters. Since by the celebrated Gibbard–Satterthwaite theorem (and other results expanding its scope) all reasonable voting systems are manipulable in principle and since many voting systems are in principle susceptible to many control types modeling natural control scenarios, much work has been done to use computational complexity as a shield to protect elections against manipulation and control. However, most of this work has yielded NP-hardness results, showing that certain voting systems resist certain types of manipulation or control only in the worst case. The typical case, where votes are given according to some natural distribution, poses a serious challenge to such worst-case complexity results and is frequently open to successful manipulation or control attempts, despite the NP-hardness of the corresponding problems.

We survey some recent results on typical-case challenges to worst-case complexity results for manipulation and control.

Introduction

In the emerging area of computational social choice, manipulation and control have been intensely studied from a computational point of view. Manipulation refers to strategic voting—voters who report untruthful preferences so as to either make their favorite candidate win or prevent their most despised candidate’s victory. By the celebrated Gibbard–Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975) and other results expanding its scope (see, e.g., the work by Duggan and Schwartz (2000)), all reasonable voting systems are manipulable in principle. This fact motivated Bartholdi et al. to study the complexity of manipulation problems in order to protect voting systems against manipulation via such complexity shields (Bartholdi III, Tovey, and Trick 1989; Bartholdi III and Orlin 1991). Recent surveys by Conitzer (2010) and Faliszewski et al. summarize the state of the art in our ongoing “war on manipulation” (Faliszewski

and Procaccia 2010; Faliszewski, Hemaspaandra, and Hemaspaandra 2010; Faliszewski et al. 2009b).

Electoral control, on the other hand, refers to attempts of an external actor, commonly called the “chair,” to influence the outcome of elections by altering their structure. The common control types, each modeling some natural control scenario, include adding, deleting, and partitioning either candidates or voters. These control types have been introduced and studied by Bartholdi III, Tovey, and Trick (1992) and Hemaspaandra, Hemaspaandra, and Rothe (2007). Although some voting systems are immune to some types of control (i.e., the chair never succeeds in making a favorite candidate win or prevent a most despised candidate’s victory), many voting systems have been shown to be susceptible (i.e., not immune) to many control types. Again, this fact motivated Bartholdi III, Tovey, and Trick (1992) to use complexity shields to protect voting systems against electoral control. Recent surveys by Faliszewski et al. (2009b) and Baumeister et al. (2010) (see also the follow-up paper by Erdélyi, Piras, and Rothe (2011)) summarize the state of the art in our ongoing “war on control.”

The overwhelming number of results on using complexity shields against manipulation or control are NP-hardness results. NP-hardness of some given (decision) problem Y is usually shown by a polynomial-time many-one reduction from a problem X already known to be NP-hard. Such a reduction is implemented by a polynomial-time computable function r transforming any (yes or no) instance x of X into an instance $y = r(x)$ of Y such that $x \in X$ if and only if $y \in Y$.

Undoubtedly, P (*deterministic polynomial time*) and NP (*nondeterministic polynomial time*) are the best known and most central complexity classes; the annoyingly intractable “P = NP?” problem has been the most important open question of theoretical computer science since decades (see also Gasarch’s “P vs. NP?” poll (Gasarch 2002), whose tenth anniversary will be celebrated by conducting a *new* poll); and the many thousands of important problems that by now are known to be NP-complete (i.e., NP-hard and a member of NP) witness the centrality of the theory of NP-completeness (Garey and Johnson 1979). Thus, the first thing to do when one encounters a seemingly hard problem is to try to prove its NP-hardness. However, NP is defined in the worst-case complexity model. All we can say about the computational hardness of NP-hard problems is that they are

*This work was supported in part by DFG grant RO-1202/15-1, SFF grant “Cooperative Normsetting” of HHU Düsseldorf, and a DAAD grant for a PPP project in the PROCOPE programme.

hard to solve on *some* instances, even if that is just, say, one instance per length. What about the other $2^n - 1$ instances for length n (assuming binary encoding)? Note that this one hard instance per length might even never occur in practice.

That is why the manipulability and controllability of voting systems has recently also been investigated, both theoretically and experimentally, with respect to typical-case instances, seeking to circumvent NP-hardness of manipulation or control problems by showing that in many cases these problems can be solved by efficient heuristics, or are even polynomial-time solvable for certain typical special cases.

In this paper, we survey some recent known results on typical-case challenges to worst-case complexity results for manipulation and control, including some previously unpublished results on analyzing control problems experimentally.

Elections and Voting Systems

An election is given by a set C of candidates and a list V of voters each having strict preferences over the candidates. As is most common, preferences are represented by linear orderings, i.e., complete rankings of the candidates. This representation will be used for the following voting systems.

Scoring rules are a very central class of voting systems. For m candidates, a scoring rule is given by a scoring vector of nonnegative integers, $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_m)$, such that $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_m$. Each candidate gets α_j points for each vote where c is ranked on the j th position, and all candidates with the most points are winners. Many important voting rules can be described by families of scoring vectors (one for each m). For example, *plurality* has scoring vectors of the form $(1, 0, \dots, 0)$; *veto* (a.k.a. *anti-plurality*) has scoring vectors of the form $(1, \dots, 1, 0)$; *Borda* has scoring vectors of the form $(m-1, m-2, \dots, 1, 0)$ for m candidates; and *k*-approval has scoring vectors $(1, \dots, 1, 0, \dots, 0)$ with k ones.

Single transferable vote (STV) proceeds in rounds, at most as many as there are candidates. Each candidate gets a point for each top position in the votes. In each round, if there is a candidate with a strict majority of points, he or she wins. Otherwise, a candidate with the smallest number of points is deleted (where ties are broken by some rule if needed), transferring his or her points to the candidate placed next, and the procedure is repeated until only one winner remains.

Plurality with runoff proceeds in two rounds. In the first round, each candidate gets a point for each top position in the votes, and the two candidates with the most points (where ties are broken by some rule if needed) move on to the second round, the “runoff,” where their points are compared after deleting all other candidates. The candidate with the most points wins (again breaking ties if needed).

Other voting rules are based on pairwise contests between the candidates. In *Condorcet voting*, the winner is a candidate who is preferred to all other candidates by a strict majority of votes (Condorcet 1785). Note that Condorcet winners don’t always exist (due to the *Condorcet paradox*), but when they exist, they are unique. A number of voting systems respect the Condorcet winner but avoid the Condorcet paradox, i.e., Condorcet winners win in these systems and there always is at least one winner. For example, in *maximin* (a.k.a. *Simpson’s rule*), let $N(c, d)$ be the number of votes

preferring c to d . The *Simpson score* of a candidate $c \in C$ is defined as $\min_{d \neq c} N(c, d)$, and all candidates with maximum Simpson score win. In *Copeland $^\alpha$ voting* with α being a rational number in $[0, 1]$, the score of a candidate $c \in C$ is defined as $\|\{d \in C - \{c\} \mid N(c, d) > N(d, c)\}\| + \alpha \|\{d \in C - \{c\} \mid N(d, c) = N(c, d)\}\|$, i.e., c gets one point for each d defeated and α points for each tie. This definition is due to Faliszewski et al. (2009a). The original definition of *Copeland voting*, which is obtained by setting $\alpha = 1/2$, is due to Copeland (1951). Nanson’s and Baldwin’s rules are based on the Borda rule, but unlike this rule, they do respect the Condorcet winner. *Nanson’s rule* proceeds in rounds, where in each round all candidates with less than the average Borda score are deleted, and this procedure is repeated until only one winner remains (Nanson 1882). *Baldwin’s rule* successively deletes a candidate with lowest Borda score in each round until only a single candidate is left, again using a tie-breaking rule if needed (Baldwin 1926).

Bucklin’s rule again proceeds in rounds (or levels). The *level i score* of a candidate $c \in C$ is defined as the number of votes from V that rank c among their top i positions. The *Bucklin score* of c is the smallest level i such that c ’s level i score strictly exceeds $\|V\|/2$. All candidates with a smallest Bucklin score, say j , and a largest level j score win.

Brams and Fishburn (1978) proposed *approval voting*, a system that unlike the voting systems above doesn’t expect linear preference orders as votes, but rather 0-1 vectors of length $\|C\|$: Every voter approves (“1”) or disapproves (“0”) each candidate, and all candidates with the most approvals win. Brams and Sanver (2009) proposed a hybrid system, called *fallback voting*, that combines Bucklin with approval voting as follows: All voters approve or disapprove of each candidate, and in addition rank their approved candidates (only those contribute to the level i scores). Every Bucklin winner in these partial rankings is also a fallback winner. However, if there exists none (due to disapprovals), then every approval winner is also a fallback winner.

Typical-Case Challenges to Complexity Results for Manipulation

Manipulation in the context of voting refers to actions of voters who seek to make their favorite candidate win (in the *constructive* case introduced by Bartholdi III, Tovey, and Trick (1989)) or to prevent their most despised candidate’s victory (in the *destructive* case introduced by Conitzer, Sandholm, and Lang (2007)), by reporting insincere preferences. Formally, for any voting system \mathcal{V} , *constructive coalitional weighted manipulation* is modeled by the following decision problem:

CONSTRUCTIVE COALITIONAL WEIGHTED MANIPULATION

Given: A set C of candidates, a list V of nonmanipulative voters over C each having a nonnegative integer weight, a list of the weights of the manipulators in S (whose votes over C are still unspecified) with $V \cap S = \emptyset$, and a distinguished candidate $c \in C$.

Question: Can the preferences of the voters in S be set such that c is a \mathcal{V} winner of $(C, V \cup S)$?

We assume that the strategic voters have complete knowledge of the sincere votes of all nonmanipulators. As special cases of the above problem, *constructive coalitional unweighted manipulation* is defined by an analogous problem, except that all weights are set to one, and the problem modeling *constructive unweighted manipulation by a single manipulator* is defined by setting the coalition size $\|S\|$ to one. The destructive variants of these problems are obtained by asking whether the preferences of the voters in S can be set such that c is *not* a \mathcal{V} winner of $(C, V \cup S)$.

Many manipulation problems—modeling different scenarios for a variety of voting systems—have been shown to be NP-complete, be it for coalitions of manipulators in the case of weighted voters (Conitzer, Sandholm, and Lang 2007; Hemaspaandra and Hemaspaandra 2007), be it for such coalitions in the case of unweighted voters (Faliszewski, Hemaspaandra, and Schnoor 2008; 2010; Betzler, Niedermeier, and Woeginger 2011; Davies et al. 2011), or be it even for single strategic voters in the case of unweighted voters (Bartholdi III, Tovey, and Trick 1989; Bartholdi III and Orlin 1991). A manipulator facing an NP-hard problem, however, does not have to despair! After all, NP-hardness merely shows that this problem is hard to solve in the worst case, and the (often technically sophisticated) reductions used to prove NP-hardness usually provide very particular instances—elections that are unlikely to appear in real-world elections. In this section we survey some of the recent approaches to and advances in tackling NP-hard manipulation problems by “typical-case” elections.

Efficient Heuristics for Junta Distributions

One of the first typical-case challenges to NP-hard manipulation problems is due to Procaccia and Rosenschein (2007b), who introduced so-called “junta distributions” that focus much weight on hard problem instances and are very light on the remaining ones. Intuitively put, they argue that when a problem is easy to solve relative to a junta distribution, it will be easy to solve relative to every typical distribution. To achieve this goal in a more formal way, they define the notion of “heuristic polynomial time” relative to a junta. Formally, these notions are defined as follows. A distribution μ is said to be a *junta* if it satisfies the following properties:

1. **Hardness:** Given an NP-hard problem X , the *restriction* $X|_{\mu}$ of X to μ , which is defined by $X|_{\mu} = \bigcup_{n \in \mathbb{N}} \{x \mid \mu_n(x) > 0 \text{ and } |x| = n\}$, is also NP-hard. Here, μ_n is the distribution of all length n instances according to μ ; in particular, for each n , μ_n sums to 1.
2. **Balance:** There are constants $k > 1$ and $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$ and for all instances x of length n , $1/k \leq \Pr_{\mu_n}[x \in X] \leq 1 - 1/k$, where $\Pr_{\mu_n}[\cdot]$ denotes the probability according to μ_n .
3. **Dichotomy:** There is a polynomial p such that for all n and for all instances x of length n , either $\mu_n(x) \geq 2^{-p(n)}$ or $\mu_n(x) = 0$.
4. **Symmetry:** If X is a manipulation problem for a voting system that expects linear preference orderings as votes,

then for each nonmanipulator $v \in V$, for any two candidates a and b distinct from the distinguished candidate c , and for each position i in the votes, v ranks a and b at position i with the same probability.

5. **Refinement:** If X is a coalitional manipulation problem, then for any length n input string x with $\mu_n(x) > 0$, if all manipulators voted identically, the distinguished candidate c would not be a winner. (Note that Procaccia and Rosenschein (2007b) restrict themselves to the case of constructive manipulation.)

A *distributional problem* (X, μ) is a (decision or search) problem X on Σ^* paired with a *distribution function* $\mu : \Sigma^* \rightarrow [0, 1]$, i.e., μ is a nondecreasing function converging to one: $\mu(0) \geq 0$, $\mu(x) \leq \mu(y)$ for each x and y with x lexicographically preceding y , and $\lim_{x \rightarrow \infty} \mu(x) = 1$. Procaccia and Rosenschein (2007b) define a *heuristic polynomial-time algorithm for* (X, μ) to be a polynomial-time algorithm \mathcal{A} for which there is a polynomial q of degree at least one and a constant $n_0 \in \mathbb{N}$ such that for each $n \geq n_0$,

$$\Pr_{\mu_n}[x \notin L \text{ if and only if } \mathcal{A} \text{ accepts } x] < 1/q(n). \quad (1)$$

Procaccia and Rosenschein (2007b) show that for each scoring rule with vector $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_m)$ satisfying $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_{m-1} > \alpha_m = 0$, there exists a junta distribution μ^* such that CONSTRUCTIVE COALITIONAL WEIGHTED MANIPULATION can be solved in heuristic polynomial time with respect to μ^* . Their heuristic polynomial-time algorithm proceeds greedily. Roughly speaking, it ranks the distinguished candidate on top of the votes of the manipulators, and in each iteration it ranks the remaining candidates by their current scores: a candidate with lowest current score is ranked highest. Their junta distribution $\mu^* = \{\mu_n\}_{n \in \mathbb{N}}$ is defined by the following sampling procedure:

1. For each manipulator $s \in S$, randomly and independently choose the weight of s to be a value in $[0, 1]$ (up to $\mathcal{O}(\log n)$ bits of precision).
2. For each candidate d distinct from the distinguished candidate c , randomly and independently choose the votes of the nonmanipulators such that the initial score of d (before the manipulators cast their votes) is in the range $[(\alpha_1 - \alpha_2)W, \alpha_1 W]$ (again, up to $\mathcal{O}(\log n)$ bits of precision), where W is the total weight of the manipulators.

In addition, c is ranked last by each nonmanipulator.

Adapting this greedy algorithm appropriately, it can also be used to show similar results for other voting systems, such as the maximin and the Copeland system with respect to a junta distribution defined similar to μ^* .

These very interesting results have been discussed by Erdélyi et al. (2009b), who consider *basic* junta distributions (defined just as junta distributions but disregarding symmetry and refinement, as these two properties are specific to manipulation problems) for general NP-hard problems, i.e., not restricted to manipulation problems. They show that very many NP-hard sets, even problems such as SAT (when suitably encoded) that are widely believed to be “really” hard problems (not only in the worst case), can be solved in heuristic polynomial time with high probability weight of

correctness with respect to basic junta distributions. They conclude that “if one were to hope to effectively use on typical NP-complete sets the notion of juntas and of heuristic polynomial time w.r.t. juntas, one would almost certainly have to go beyond the basic three conditions and add additional conditions” (Erdélyi et al. 2009b, p. 3996). Nonetheless, they stress that the approach of Procaccia and Rosenschein (2007b), which is restricted to manipulation problems only and isn’t meant to speak to general NP-hard problems, is very interesting indeed and should be further pursued.

Erdélyi et al. (2009b) also discuss the related, but different notion of *average polynomial time*, which—somewhat misleadingly—has been used in the literature for typical-case studies, such as that of Procaccia and Rosenschein (2007b), that actually concern the frequency of correctness of heuristics with respect to underlying distributions. The theory of average-case complexity was initiated by Levin (1986); see also the surveys by Goldreich (1997) and Wang (1997a; 1997b). Crucially, average polynomial time refers to taking an average of running times over the inputs according to some underlying distribution such that this average running time is low: AvgP is the class of distributional problems (X, μ) for which there is an algorithm \mathcal{A} solving X such that the running time T of \mathcal{A} is polynomial on the average with respect to distribution μ . That is, there is a constant $\varepsilon > 0$ such that $\sum_{x \in \Sigma^*} \bar{\mu}(x) \cdot (T(x)^\varepsilon / |x|) < \infty$, where $\bar{\mu}$ denotes the *density function induced by μ* , i.e., $\bar{\mu}(0) = \mu(0)$ and $\bar{\mu}(x) = \mu(x) - \mu(x-1)$ for all $x > 0$ (here, the string $x-1$ denotes the lexicographic predecessor of x).

By contrast, heuristic polynomial time with respect to a junta refers to the probability weight according to some underlying distribution for which the heuristic is correct. Note that merely a probability weight of $1 - (1/q(n))$ is required for all except a finite number of length n inputs, for some polynomial q . As the remaining inputs with probability weight $1/q(n)$ have no guarantee as to whether the heuristic solves them correctly, they need to be solved by brute force, which requires exponential time and so destroys any hope of getting a real average polynomial-time algorithm; see the more detailed discussion in (Erdélyi et al. 2008, Appendix C). A related discussion can be found in (Faliszewski, Hemaspaandra, and Hemaspaandra 2011a, Section 6), see also (Homan and Hemaspaandra 2009; Erdélyi et al. 2009a)

Approximation Algorithms

Another approach to challenge NP-hard manipulation problems in practice is due to Zuckerman, Procaccia, and Rosenschein (2009), who reformulate this as an optimization problem in the constructive, unweighted case for coalitions of manipulators and then try to approximate a solution:

CONSTRUCTIVE COALITIONAL UNWEIGHTED OPTIMIZATION

- Input:** A set C of candidates, a list V of nonmanipulative voters over C , and a distinguished candidate $c \in C$.
- Output:** The minimal n such that a coalition S of size n of (unweighted) manipulators can make c win in $(C, V \cup S)$ by casting insincere votes.
-

Zuckerman, Procaccia, and Rosenschein (2009) investigate this problem for scoring rules, Borda, and maximin, and they design greedy algorithms that approximate the corresponding problems within a certain factor, i.e., they analyze the algorithms’ windows of error. In particular, CONSTRUCTIVE COALITIONAL UNWEIGHTED OPTIMIZATION can be efficiently approximated up to an additive constant of 1 for Borda (i.e., one additional manipulator suffices) by an algorithm called REVERSE, and it can be efficiently approximated within a factor of 2 for maximin (i.e., doubling the number of manipulators suffices). On the other hand, they provide algorithms that solve the unweighted decision problem CONSTRUCTIVE COALITIONAL UNWEIGHTED MANIPULATION efficiently for plurality with runoff and for veto. Note that these algorithms also apply to the weighted case and make errors only on very few configurations of voters’ weights. In particular, the approximation algorithm for Borda improves the error analysis implicit in the above-mentioned more general result that Procaccia and Rosenschein (2007b) achieve for a certain family of scoring rules and CONSTRUCTIVE COALITIONAL WEIGHTED MANIPULATION with respect to junta distributions when tailored to Borda only (which, of course, satisfies the required condition $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_{m-1} > \alpha_m = 0$). However, as the latter result applies to a more general family of voting systems, neither result subsumes the other.

Davies et al. (2011) propose two other approximation algorithms, LARGEST FIT and AVERAGE FIT, to find optimal manipulations for the Borda rule,¹ and compare them both theoretically and experimentally (the latter approach will be treated in more detail later on) with the above-mentioned algorithm REVERSE introduced by Zuckerman, Procaccia, and Rosenschein (2009). LARGEST FIT and AVERAGE FIT use ideas from bin packing and multiprocessor scheduling. The theoretical analysis shows that both are incomparable with REVERSE in the sense that an infinite family of instances can be found where REVERSE performs better than either of them, and in turn there is an infinite family of instances where REVERSE does not find the optimal solution but the other algorithms do.

Zuckerman, Lev, and Rosenschein (2011) improve the approximation error of 2 for maximin in CONSTRUCTIVE COALITIONAL UNWEIGHTED OPTIMIZATION to a factor of $5/3$. In addition, they prove that no approximation factor for this problem can be better than $3/2$, unless $P = NP$.

This approach is very promising and should definitely be pursued and further explored, for example, for other voting systems. In particular, additional inapproximability results such as the one mentioned in the previous paragraph would be especially beneficial, as they add to the protection complexity theory can provide against manipulation.

Procaccia (2010) suggests a completely different approach to protect elections against manipulation by using approximation methods. The idea is not to analyze the ap-

¹They and, independently, Betzler, Niedermeier, and Woeginger (2011) have recently shown that CONSTRUCTIVE COALITIONAL UNWEIGHTED MANIPULATION is NP-complete for Borda, even for only two manipulators.

proximability of the manipulation problem for a given voting system but to secure the voting system itself against manipulation by approximating it with a strategyproof (i.e., nonmanipulable) randomized voting rule. Procaccia (2010) proposes such randomized voting rules providing approximations of score-based voting systems. A strategyproof randomized voting rule f is said to *approximate a given score-based voting rule within an approximation ratio of γ* if the expected score of the winner that f chooses is at least $\gamma \cdot s$, where s is the maximal score. It is shown that for m candidates, positional scoring rules can be approximated by strategyproof randomized voting rules within a factor of $\Omega(1/\sqrt{m})$, and that for plurality voting this is, asymptotically, the best approximation possible. For the Borda rule, on the other hand, it is proven that a $((1/2) + \Omega(1/m))$ -approximation can be achieved. Furthermore, Copeland $^\alpha$ and maximin are analyzed. Interestingly, maximin cannot be approximated nontrivially, which means that no strategyproof randomized voting rule exists that provides a better approximation of maximin than the trivial approximation, namely choosing a winner at random. For Copeland $^\alpha$ with $\alpha \in [1/2, 1]$, a lower bound of $1/2 + \Omega(1/m)$ can be shown, whereas for $\alpha \in [0, 1]$ the analysis provides an upper bound of $1/2 + \mathcal{O}(1/m)$.

Single-Peaked Preferences

What is a “typical” election? Well, it depends. In general, it is a nontrivial problem to say how votes in an election are “typically” distributed. However, there are certain special cases that may occur in real-world elections and that may outright be easy to solve. For example, suppose society votes on a single issue (such as taxes or health care or war on terror, etc.) that can be nicely embedded in a left-right spectrum. That is, there exists a linear (societal) ordering of candidates on this spectrum (e.g., if taxes are the issue to be voted on, a left-wing candidate would stand for high taxes and a right-wing candidate would stand for low taxes), and relative to this linear ordering, all voters’ preference utility curves raise to a single peak (representing this voter’s most preferred position on this spectrum) and then fall, or just raise, or just fall. This model of *single-peaked* preferences is a central concept in political science and has been introduced by Black (1948; 1958); see also, e.g., (Gailmard, Patty, and Penn 2009; Ballester and Haeringer 2011; Lepelley 1996) for more recent social-choice-theoretic work on single-peakedness. Formally, an election (C, V) is said to be *single-peaked* if there exists a linear ordering L on C such that for each vote $v_i \in V$ (individually represented by a linear ordering $>_i$ on C) and for each triple of candidates, c, d , and e in C , if $cLdLe$ or $eLdLc$ then $c >_i d$ implies $d >_i e$ for each i . In other words, for each triple of candidates ordered according to L , it can never happen in an individual vote that the “middle” candidate is ranked last.

Restricting an electorate to only single-peaked preferences may or may not change the computational properties and the complexity of the associated manipulation problems. Such restrictions have only recently been considered—e.g., by Escoffier, Lang, and Öztürk (2008) and Conitzer (2009)—from a computational point of view, although the concept

in political science is well-established for more than a half century now. In particular, Walsh (2007) shows that the weighted manipulation problem for STV for at least three candidates remains NP-complete, even when the given election is restricted to be single-peaked. Here, the underlying linear ordering L of candidates relative to which the votes of the nonmanipulators are single-peaked is part of the input, and the manipulators’ votes are supposed to be single-peaked relative to the same ordering L .

Faliszewski et al. (2011) prove that, depending on the voting system used, NP-hardness of CONSTRUCTIVE COALITIONAL WEIGHTED MANIPULATION can vanish or can remain in place. For example, one of their results says that CONSTRUCTIVE COALITIONAL WEIGHTED MANIPULATION for 3-candidate Borda elections is in P when restricted to single-peaked electorates (whereas it is NP-complete in the unrestricted case), while it remains NP-complete for 4-candidate Borda elections, even when restricted to single-peaked electorates (just as in the general case).

Remarkably, for m -candidate 3-veto elections, this manipulation problem is in P whenever $m \leq 4$ or $m \geq 6$, but is NP-complete for $m = 5$ (Faliszewski et al. 2011). That is, due to single-peakedness, the complexity of the problem drops down to polynomial time, although the number of candidates is incremented from five to six or more.

Faliszewski et al. (2011) also show that (for a certain artificial voting system) restricting the electorate to the single-peaked case may even increase the complexity of manipulation. In addition, they prove a dichotomy result for single-peaked electorates when a scoring rule with vector $\vec{\alpha} = (\alpha_1, \alpha_2, \alpha_3)$ is used: CONSTRUCTIVE COALITIONAL WEIGHTED MANIPULATION is NP-complete whenever $\alpha_1 - \alpha_3 > 2(\alpha_2 - \alpha_3) > 0$; otherwise, it is in P. This dichotomy result has been generalized by Brandt et al. (2010) to scoring rules with any (fixed) number of candidates. For further results, see (Faliszewski, Hemaspaandra, and Hemaspaandra 2011a).

Mattei (2011) empirically investigates huge data sets from real-world elections (drawn from the Netflix Prize data set) with respect to properties such as how likely the Condorcet paradox is to appear and how often single-peaked preference profiles are to occur.² In particular, his experiments indicate that single-peaked preferences only very rarely occur in practice, which brings us right to our next topic.

An Experimental Approach

Another recent line of research studies experimental simulation and evaluates heuristics for solving NP-hard manipulation problems empirically. These investigations were initiated by Walsh (2009; 2010), who experimentally studied NP-hard manipulation problems for veto and STV, and showed that for many instances the elections generated can be manipulated quickly.

²Peleg and Sudhölter (1999) prove that all (generalized) “median voter systems” are strategyproof and even coalition-strategyproof for single-peaked preference profiles. However, many widely used voting systems (such as plurality, Borda, etc.) are no median voter systems and thus are not strategyproof.

Generating elections for experimental analysis can be done in various ways depending on the electorates one wants to model. The possibility and frequency of successful manipulation can vary greatly for different vote distributions.

For the veto rule, Walsh (2009) investigates the problem CONSTRUCTIVE COALITIONAL WEIGHTED MANIPULATION restricted to elections with three candidates, which can be directly reduced to 2-WAY-NUMBER PARTITIONING. This reduction allows to make use of known efficient algorithms for the latter problem, such as the CKK algorithm by Korf (1995), to solve the manipulation problem.

The votes in the tested elections are generated by randomly choosing one of the three candidates to be vetoed, and the vetoes carry randomly drawn weights as well. The electorate’s distribution is varied by the generation of the voters’ weights. To generate uniform votes, the weights of the voters are drawn uniformly and independently at random from a given interval. Similarly, normally distributed votes are generated by drawing the voters’ weights independently from a normal distribution.

The weight of the manipulating coalition is crucial for the complexity of determining whether a given election is manipulable or not. If the weight is too small, the coalition is hopeless, whereas manipulation is trivially possible if the weight is too big. Instances where the manipulative coalition’s weight is between these trivial cases are conjectured to be the hard ones, so this region is of most interest in the results obtained from the conducted experiments.

Similar results were found for uniformly and for normally distributed votes: Even in the critical region, the decision whether the tested election is manipulable or not could be made with low computational costs. The probability curves for successful manipulation show in both distributions a smooth phase transition in this critical area, similarly to the phase transition observed for polynomial-time solvable decision problems.

Complementary to this setting, electorates with correlated votes have been investigated as well. For the veto system, so-called “hung” elections were generated where the manipulative coalition is twice as heavy as the nonmanipulative voters’ total weight, but all nonmanipulators veto the distinguished candidate the coalition wants to make win. This finely balanced situation is exactly what the reduction of Conitzer, Sandholm, and Lang (2007) produces from a given PARTITION instance in their proof of NP-hardness of the manipulation problem. Not surprisingly, generating such instances at random leads to higher computational costs for deciding them in the critical region. Furthermore, the probability curves resulting from these instances show a typical sharp phase transition, similar to that of other hard decision problems, namely around $(\log k)/m \approx 1$, where m is the number of manipulators and their weights are randomly chosen from $(0, k]$. Interestingly, one randomly vetoing voter in an otherwise perfectly hung election suffices to at least empirically make the problem easier.

Furthermore, the results show that in elections with uniform votes the sizes of the voters’ weights do not influence the manipulability of the tested elections, confirming empirically the theoretical conjectures by Procaccia and Rosen-

schein (2007a) and Xia and Conitzer (2008).

For preference-based voting systems such as STV and Borda, the votes are given by linear orders over all candidates and thus are permutations of the candidates. To generate uniformly distributed votes for these voting systems, each vote is drawn uniformly and independently from an urn containing all possible votes. This is the so-called *Impartial Culture (IC)* model in which each vote is equally likely to occur. To model correlated votes, Walsh (2010) uses the *Polya Eggenberger urn model (PE)*—described, e.g., by Berg (1985)—in the following sense: The first vote is drawn from an urn containing all possible votes; there are $m!$ different permutations of the m candidates. Before drawing the second vote, $m!$ votes identical to the first one are put back to the urn. Before drawing the third vote, $m!$ votes identical to the second vote are put back to the urn, and so on.³

For the conducted experiments, an improved version of an algorithm given by Conitzer, Sandholm, and Lang (2007) is used. The experiments show that independent of the underlying model a successful manipulation action for a single manipulator can easily be computed in elections where the number of candidates is not higher than 128. For coalitions of manipulators casting identical votes, the computational cost of deciding whether manipulation is possible depends on the coalition size. Increasing the number of manipulators increases the manipulability of the tested elections.

Complementary to the testing on random elections where no votes exist that are never cast, Walsh (2010) also samples real elections: An election to determine a trajectory for NASA’s Marine spacecraft and the votes cast in a faculty hiring committee at the University of Irvine (for details see (Dyer and Miles Jr. 1976; Dobra 1983)). The sampled elections show similar results as the randomly generated elections.⁴ Deciding whether a manipulator can successfully change the outcome of the election or not is easy for up to 128 candidates and voters.

Davies et al. (2011) use the experimental approach introduced by Walsh (2009; 2010) to analyze their approximation algorithms, LARGEST FIT and AVERAGE FIT, for the Borda system and compare them with the algorithm REVERSE of Zuckerman, Procaccia, and Rosenschein (2009). Random elections for their experiments are generated in the IC model and the PE model, and the optimal solution for a given instance is computed with the solver GECODE after having modeled the manipulation problem as a constraint

³This procedure generates highly correlated votes and models how homogeneity varies in society. By differently choosing the number of votes that are put back in each step, the correlation can be varied. Here we have that the second vote is the same as the first vote with probability $1/2$. This probability can be de- or increased by putting back less or more votes in each step.

⁴For comparison, the sampled elections should have the same number of candidates and voters as the randomly generated elections. To obtain this, candidate or voter sets containing too many elements are reduced by randomly choosing appropriate subsets. If the list of voters has to be extended, votes are uniformly and independently chosen from the given votes. If the candidate set is too small, the candidates are duplicated and the ranking between the clone and the original candidate is chosen randomly.

satisfaction problem.⁵ The results show that the LARGEST FIT algorithm finds an optimal solution in roughly 83% of the elections with uniform votes and in roughly 42% of the elections generated with the PE model. The AVERAGE FIT algorithm, on the other hand, finds an optimal solution in almost all (roughly 99%) of the tested elections independent of the distribution. Thus, LARGEST FIT and AVERAGE FIT behave better than REVERSE for Borda in both distribution models, as REVERSE finds an optimal manipulation in roughly 76% of the tested elections only.

Narodytska, Walsh, and Xia (2011) study unweighted and weighted manipulation of Nanson’s and Baldwin’s rules. In the unweighted case, they prove that both rules are NP-hard to manipulate, even for just one manipulator. In the weighted case, they show that coalitional manipulation is NP-hard for Nanson when there are four candidates and is in P for three candidates. Since Coleman and Teague (2007) have shown NP-hardness of this problem for Baldwin already for three candidates, Baldwin’s rule appears to be computationally more resistant to manipulation than Nanson’s rule. Narodytska, Walsh, and Xia (2011) also conduct experiments for these two rules, using the same approximation algorithms as Davies et al. (2011), LARGEST FIT and AVERAGE FIT, and two more, ELIMINATE and REVERSE ELIMINATE. Their results suggest that, at least for the algorithms studied, Nanson’s and Baldwin’s rules are harder to manipulate in practice than Borda’s rule. For Nanson and Baldwin, REVERSE works slightly better than LARGEST FIT and AVERAGE FIT, which in turn outperform ELIMINATE and REVERSE ELIMINATE, especially so when the number of candidates is large.

Typical-Case Challenges to Complexity Results for Control

Electoral control models structural changes an election’s chair (who seeks to influence the outcome of an election) can make by adding, deleting, or partitioning either candidates or voters. These control actions model real-world election issues such as campaign advertisement, get-out-the-vote drives, vote suppression, and gerrymandering. Bartholdi III, Tovey, and Trick (1992) introduced seven constructive control types and investigated them for Condorcet and plurality voting. Hemaspaandra, Hemaspaandra, and Rothe (2007) extended this study by introducing destructive control types as well, also adding two natural tie-handling rules for the partitioning cases, and specifically studied the control complexity of Condorcet, plurality, and approval voting.

A number of follow-up papers were concerned with the control complexity of further voting systems and other aspects of control—see, e.g., (Hemaspaandra, Hemaspaandra, and Rothe 2009; Faliszewski et al. 2009a; Erdélyi, Nowak, and Rothe 2009; Erdélyi and Rothe 2010; Faliszewski, Hemaspaandra, and Hemaspaandra 2011b; Erdélyi, Piras, and Rothe 2011) and the surveys by Faliszewski et al. (2009b; 2010) and Baumeister et al. (2010).

⁵The timeout is set to one hour and the variable-ordering heuristic used is the Domain Over Weight Degree.

Control actions can be formalized by decision problems whose instances always contain a distinguished candidate c and an initial election (C, V) , and the question always is whether c can be made the unique winner by modifying (C, V) according to the control action at hand.

In the adding candidates or voters cases, the instances additionally contain a set of spoiler candidates or a list of as yet unregistered voters from which the chair can choose whom to add, and a bound limiting the number of candidates or votes that may be added. In the deleting candidates or voters scenarios, only the limiting bound is given additionally.

Partitioning either the candidates or the voters changes an election’s course by transforming it into a two-stage election with one or two pre-round election(s) and one final-stage election. The two tie-handling rules mentioned above are *Ties Eliminate (TE)*, where only unique pre-round winners move on to the final round, and *Ties Promote (TP)*, where all pre-round winners participate in the final stage.

There is a total of 22 standard types of control. As an example, we state the formal decision problem corresponding to constructive control by partition of candidates by:

CONSTRUCTIVE CONTROL BY PARTITION OF CANDIDATES

Given: An election (C, V) and a distinguished candidate $c \in C$.

Question: Is it possible to partition C into C_1 and C_2 such that c is the unique winner (under the election system at hand) of election $(W_1 \cup C_2, V)$, where W_1 is the set of winners of subelection (C_1, V) surviving the tie-handling rule?

In control by runoff partition of candidates, there are two pre-rounds, (C_1, V) and (C_2, V) , and the final stage is the subelection $(W_1 \cup W_2, V)$, where W_i , $i \in \{1, 2\}$, is the set of winners of subelection (C_i, V) surviving the tie-handling rule. When the list of voters is partitioned, the pre-round of the resulting two-stage election consists of two subelections where the voters of each sublist vote over all candidates and those candidates surviving the tie-handling rule run against each other in the final round, considering all votes.

Again, we can define the destructive variants analogously, and we assume that the chair knows all the votes. Many of the arguments in the previous section on manipulation apply also to control, so an election’s chair does not have to despair either when facing an NP-hard control problem. However, some approaches—such as those of Zuckerman, Procaccia, and Rosenschein (2009; 2011) that have been proposed for manipulation—may be less suited for control problems.

Single-Peaked Preferences

Regarding the restriction of electorates to single-peaked preferences, Faliszewski et al. (2011) showed that the control problem can be solved in polynomial time in all cases in which plurality and approval voting are in general NP-hard to control by adding or deleting either candidates or voters. Brandt et al. (2010) achieved similar results for other voting systems as well, in particular those that satisfy the weak Condorcet criterion, and in addition for the case of constructive control by partition of voters.

An Experimental Approach

Among natural voting systems with efficient winner determination, the systems currently known to have the most NP-hardness results in the various control scenarios are fallback voting and Bucklin voting (Erdélyi, Piras, and Rothe 2011). For plurality, Bucklin, and fallback voting, Schend and Rothe (2011) have performed an extensive experimental study on the frequency of controllability⁶ for randomly generated elections. Inspired by the experimental setup of Walsh (2009; 2010) described in the section about experimental approaches regarding manipulation, one model used for the vote distribution is the IC model. To simulate correlated votes, however, Schend and Rothe (2011) introduce an adaption of the PE model, which they call the *Two Mainstreams (TM)* model: Two votes v_1, v_2 are drawn independently from an urn containing all possible, say t , votes. Each of the votes drawn is put back into the urn with t additional identical votes, and the list of votes is then drawn uniformly at random from this urn. Thus, each voter has with probability $1/3$ the same preference as v_1 , with probability $1/3$ the same preference as v_2 , and again with probability $1/3$ a different preference. The votes v_1 and v_2 model two mainstreams, such as liberal and conservative, a society may have. With these two distributions random elections are generated, letting the number of candidates and votes, m and n , vary in powers of 2 between 4 and 128. For each data point (i.e., for each pair of m and n), 500 elections are tested.

The algorithms implemented to solve the control problems for the different voting systems apply a heuristic approach: A successful control action for a given election is searched for by basically testing all possible control actions of the given type. This testing process is systematized by preordering the candidates or voters such that promising control actions are tested first. To ensure practicability, the algorithm aborts the computation after a fixed timeout, indicating the inconclusive result by its output. The results of Schend and Rothe (2011) allow a fine-grained analysis of those control types the three considered voting systems are resistant to (i.e., the corresponding control problem is NP-hard) except, up to now, the partition-of-candidates cases where the experiments are still running.

So far, the results for both constructive and destructive control by partition of candidates in model TE for Bucklin and fallback voting align with the conclusions arrived at by Schend and Rothe (2011) for other control scenarios, which can be summarized as follows.

Comparing the two distribution models used, for all investigated control types and in all voting systems considered, elections generated with the IC model show a higher overall number of yes-instances than elections generated in the TM model. At the same time, the number of timeouts is higher for elections generated with the TM model.

Bucklin and fallback voting show the same tendencies

⁶The experiments have been conducted for each of the three voting systems in the NP-hard control cases only, with one exception: DESTRUCTIVE CONTROL BY PARTITION OF VOTERS in model TP has nonetheless been studied experimentally for Bucklin voting, as the complexity of this control problem is still unknown.

both theoretically (regarding NP-hardness versus membership in P) and experimentally. Note that the complexity status of DESTRUCTIVE CONTROL BY PARTITION OF VOTERS in model TP is still open, whereas this problem is known to be NP-complete for fallback voting (Erdélyi, Piras, and Rothe 2011). Since Bucklin and fallback voting behave empirically very similarly with respect to this control type, we conjecture that this control problem is NP-complete for Bucklin voting as well.

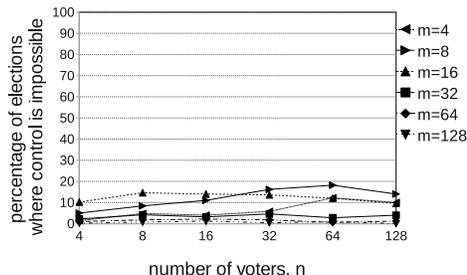


Figure 1: Experimental results for fallback voting in the TM model for CONSTRUCTIVE CONTROL BY PARTITION OF CANDIDATES in model TE, for a fixed number of candidates

Comparing Figures 1 and 2 affirms the intuition that destructive control is easier to exert than constructive control: For some election sizes, up to 100% of the tested elections are controllable by certain destructive control types, whereas for the constructive cases, especially regarding candidate control, the number of yes-instances is much smaller.

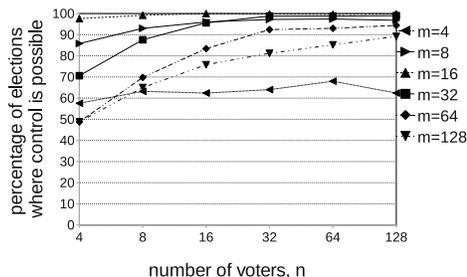


Figure 2: Experimental results for Bucklin voting in the IC model for DESTRUCTIVE CONTROL BY PARTITION OF CANDIDATES in model TE, for a fixed number of candidates

Further comparisons across the different control types show that in the constructive cases voter control seems to be easier to exert in practice than candidate control. Control by adding candidates shows particularly few yes-instances for all three voting systems, indicating that this may be the hardest control type investigated empirically so far (the experiments on partition of candidates still need to be completed). The results also show that more of the tested elections are controllable by deleting voters than by adding voters, and the same can be observed regarding candidate control.

References

- Baldwin, J. 1926. The technique the Nanson preferential majority system of election. *Transactions and Proceedings of the Royal Society of Victoria* 39:42–52.
- Ballester, M., and Haeringer, G. 2011. A characterization of the single-peaked domain. *Social Choice and Welfare* 36(2):305–322.
- 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 Modelling* 16(8/9):27–40.
- Baumeister, D.; Erdélyi, G.; Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 2010. Computational aspects of approval voting. In Laslier, J., and Sanver, R., eds., *Handbook on Approval Voting*. Springer. chapter 10, 199–251.
- Berg, S. 1985. Paradox of voting under an urn model: The effect of homogeneity. *Public Choice* 47(2):377–387.
- Betzler, N.; Niedermeier, R.; and Woeginger, G. 2011. Unweighted coalitional manipulation under the Borda rule is NP-hard. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, 55–60. IJCAI.
- Black, D. 1948. On the rationale of group decision-making. *Journal of Political Economy* 56(1):23–34.
- Black, D. 1958. *The Theory of Committees and Elections*. Cambridge University Press.
- Brams, S., and Fishburn, P. 1978. Approval voting. *American Political Science Review* 72(3):831–847.
- Brams, S., and Sanver, R. 2009. Voting systems that combine approval and preference. In Brams, S.; Gehrlein, W.; and Roberts, F., eds., *The Mathematics of Preference, Choice, and Order: Essays in Honor of Peter C. Fishburn*. Springer. 215–237.
- Brandt, F.; Brill, M.; Hemaspaandra, E.; and Hemaspaandra, L. 2010. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 715–722. AAAI Press.
- Coleman, T., and Teague, V. 2007. On the complexity of manipulating elections. In *Proceedings of Computing: the 13th Australasian Theory Symposium*, volume 65, 25–33.
- Condorcet, J. 1785. Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix. Facsimile reprint of original published in Paris, 1972, by the Imprimerie Royale. English translation appears in I. McLean and A. Urken, *Classics of Social Choice*, University of Michigan Press, 1995, pages 91–112.
- Conitzer, V.; Sandholm, T.; and Lang, J. 2007. When are elections with few candidates hard to manipulate? *Journal of the ACM* 54(3):Article 14.
- Conitzer, V. 2009. Eliciting single-peaked preferences using comparison queries. *Journal of Artificial Intelligence Research* 35:161–191.
- Conitzer, V. 2010. Making decisions based on the preferences of multiple agents. *Communications of the ACM* 53(3):84–94.
- Copeland, A. 1951. A “reasonable” social welfare function.
- Davies, J.; Katsirelos, G.; Narodytska, N.; and Walsh, T. 2011. Complexity of and algorithms for Borda manipulation. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, 657–662. AAAI Press.
- Dobra, J. 1983. An approach to empirical studies of voting paradoxes: An update and extension. *Public Choice* 41(2):241–250.
- Duggan, J., and Schwartz, T. 2000. Strategic manipulability without resoluteness or shared beliefs: Gibbard–Satterthwaite generalized. *Social Choice and Welfare* 17(1):85–93.
- Dyer, J., and Miles Jr., R. 1976. An actual application of collective choice theory to the selection of trajectories for the mariner Jupiter/Saturn 1977 project. *Operations Research* 24(2):220–244.
- Erdélyi, G., and Rothe, J. 2010. Control complexity in fallback voting. In *Proceedings of Computing: the 16th Australasian Theory Symposium*, 39–48. Australian Computer Society Conferences in Research and Practice in Information Technology Series, vol. 32, no. 8.
- Erdélyi, G.; Hemaspaandra, L.; Rothe, J.; and Spakowski, H. 2008. Frequency of correctness versus average-case polynomial time and generalized juntas. Technical Report TR-934, Department of Computer Science, University of Rochester, Rochester, NY.
- Erdélyi, G.; Hemaspaandra, L.; Rothe, J.; and Spakowski, H. 2009a. Frequency of correctness versus average polynomial time. *Information Processing Letters* 109(16):946–949.
- Erdélyi, G.; Hemaspaandra, L.; Rothe, J.; and Spakowski, H. 2009b. Generalized juntas and NP-hard sets. *Theoretical Computer Science* 410(38–40):3995–4000.
- Erdélyi, G.; Nowak, M.; and Rothe, J. 2009. Sincere-strategy preference-based approval voting fully resists constructive control and broadly resists destructive control. *Mathematical Logic Quarterly* 55(4):425–443.
- Erdélyi, G.; Piras, L.; and Rothe, J. 2011. The complexity of voter partition in Bucklin and fallback voting: Solving three open problems. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems*, 837–844. IFAA-MAS.
- Escoffier, B.; Lang, J.; and Öztürk, M. 2008. Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence*, 366–370. IOS Press.
- Faliszewski, P., and Procaccia, A. 2010. AI’s war on manipulation: Are we winning? *AI Magazine* 31(4):53–64.
- Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 2009a. Llull and Copeland voting computationally resist bribery and constructive control. *Journal of Artificial Intelligence Research* 35:275–341.
- Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 2009b. A richer understanding of the complexity of election systems. In Ravi, S., and Shukla, S., eds., *Fundamental Problems in Computing: Essays in Honor of Professor Daniel J. Rosenkrantz*. Springer. chapter 14, 375–406.
- Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 2011. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation* 209(2):89–107.
- Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. 2010. Using complexity to protect elections. *Communications of the ACM* 53(11):74–82.
- Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. 2011a. The complexity of manipulative attacks in nearly single-peaked electorates. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*, 228–237. ACM Press.
- Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. 2011b. Multimode control attacks on elections. *Journal of Artificial Intelligence Research* 40:305–351.

- Faliszewski, P.; Hemaspaandra, E.; and Schnoor, H. 2008. Copeland voting: Ties matter. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, 983–990. IFAAMAS.
- Faliszewski, P.; Hemaspaandra, E.; and Schnoor, H. 2010. Manipulation of Copeland elections. In *Proceedings of the 9th International Joint Conference on Autonomous Agents and Multiagent Systems*, 367–374. IFAAMAS.
- Gailmard, S.; Patty, J.; and Penn, E. 2009. Arrow’s theorem on single-peaked domains. In Aragonés, E.; Beviá, C.; Llavador, H.; and Schofield, N., eds., *The Political Economy of Democracy*. Fundación BBVA. 335–342.
- Garey, M., and Johnson, D. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company.
- Gasarch, W. 2002. The $P = ?$ NP poll. *SIGACT News* 33(2):34–47.
- Gibbard, A. 1973. Manipulation of voting schemes. *Econometrica* 41(4):587–601.
- Goldreich, O. 1997. Notes on Levin’s theory of average-case complexity. Technical Report TR97-058, Electronic Colloquium on Computational Complexity.
- 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. 2007. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence* 171(5–6):255–285.
- Hemaspaandra, E.; Hemaspaandra, L.; and Rothe, J. 2009. Hybrid elections broaden complexity-theoretic resistance to control. *Mathematical Logic Quarterly* 55(4):397–424.
- Homan, C., and Hemaspaandra, L. 2009. Guarantees for the success frequency of an algorithm for finding Dodgson-election winners. *Journal of Heuristics* 15(4):403–423.
- Korf, R. 1995. From approximate to optimal solutions: A case study of number partitioning. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, 266–272.
- Lepelley, D. 1996. Constant scoring rules, Condorcet criteria, and single-peaked preferences. *Economic Theory* 7(3):491–500.
- Levin, L. 1986. Average case complete problems. *SIAM Journal on Computing* 15(1):285–286.
- Mattei, N. 2011. Empirical evaluation of voting rules with strictly ordered preference data. In *Proceedings of the 2nd International Conference on Algorithmic Decision Theory*, 165–177. Springer-Verlag *Lecture Notes in Artificial Intelligence* #6992.
- Nanson, E. 1882. Methods of election. *Transactions and Proceedings of the Royal Society of Victoria* 19:197–240.
- Narodytska, N.; Walsh, T.; and Xia, L. 2011. Manipulation of nanson’s and baldwin’s rules. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, 713–718. AAAI Press.
- Peleg, B., and Sudhölter, P. 1999. Single-peakedness and coalition-proofness. *Review of Economic Design* 4(4):381–387.
- Procaccia, A., and Rosenschein, J. 2007a. Average-case tractability of manipulation in voting via the fraction of manipulators. In *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems*, volume 718–720. Short paper.
- Procaccia, A., and Rosenschein, J. 2007b. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research* 28:157–181.
- Procaccia, A. 2010. Can approximation circumvent Gibbard-Satterthwaite? In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 836–841.
- Satterthwaite, M. 1975. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10(2):187–217.
- Schend, L., and Rothe, J. 2011. Control complexity in plurality, Bucklin, and fallback voting: An experimental approach. Manuscript.
- Walsh, T. 2007. Uncertainty in preference elicitation and aggregation. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence*, 3–8. AAAI Press.
- Walsh, T. 2009. Where are the really hard manipulation problems? The phase transition in manipulating the veto rule. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, 324–329. IJCAI.
- Walsh, T. 2010. An empirical study of the manipulability of single transferable voting. In *Proceedings of the 19th European Conference on Artificial Intelligence*, 257–262. IOS Press.
- Wang, J. 1997a. Average-case computational complexity theory. In Hemaspaandra, L., and Selman, A., eds., *Complexity Theory Retrospective II*. Springer-Verlag. 295–328.
- Wang, J. 1997b. Average-case intractable NP problems. In Du, D., and Ko, K., eds., *Advances in Languages, Algorithms, and Complexity*. Kluwer Academic Publishers. 313–378.
- Xia, L., and Conitzer, V. 2008. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, 109–118. ACM Press.
- Zuckerman, M.; Lev, O.; and Rosenschein, J. 2011. An algorithm for the coalitional manipulation problem under maximin. In *Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems*, 845–852. IFAAMAS.
- Zuckerman, M.; Procaccia, A.; and Rosenschein, J. 2009. Algorithms for the coalitional manipulation problem. *Artificial Intelligence* 173(2):392–412.