DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD

A Complexity-of-Strategic-Behavior Comparison Between Schulze's Rule and Ranked Pairs

Citation	Parkes, David C., and Lirong Xia. 2012. A complexity-of-strategic-behavior comparison between Schulze's rule and ranked pairs. Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI '12), July 22–26, 2012, Ontario, Canada, 1429-1435. AAAI Press.
Published Version	http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/view/4976
Accessed	February 19, 2015 1:11:13 PM EST
Citable Link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:11956917
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP

(Article begins on next page)

A Complexity-of-Strategic-Behavior Comparison between Schulze's Rule and Ranked Pairs

David C. Parkes
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
parkes@eecs.harvard.edu

Lirong Xia
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
lxia@seas.harvard.edu

Abstract

Schulze's rule and ranked pairs are two Condorcet methods that both satisfy many natural axiomatic properties. Schulze's rule is used in the elections of many organizations, including the Wikimedia Foundation, the Pirate Party of Sweden and Germany, the Debian project, and the Gento Project. Both rules are immune to control by cloning alternatives, but little is otherwise known about their strategic robustness, including resistance to manipulation by one or more voters, control by adding or deleting alternatives, adding or deleting votes, and bribery. Considering computational barriers, we show that these types of strategic behavior are NP-hard for ranked pairs (both constructive, in making an alternative a winner, and destructive, in precluding an alternative from being a winner). Schulze's rule, in comparison, remains vulnerable at least to constructive manipulation by a single voter and destructive manipulation by a coalition. As the first such polynomial-time rule known to resist all such manipulations, and considering also the broad axiomatic support, ranked pairs seems worthwhile to consider for practical applications.

Introduction

In multi-agent systems, voting is a popular method used to aggregate agents' preferences to make a joint decision. Yet social choice does not give a clear prescription for the best choice for a voting rule. A common approach is to look for rules that satisfy particular axiomatic properties. The *Condorcet criterion* is a prominent example of such a property. A voting rule satisfies the Condorcet criterion (and is called a *Condorcet method*), if it always selects the *Condorcet winner* whenever one exists. An alternative c is the Condorcet winner if it beats all other alternatives in pairwise elections, that is, if for each other alternative d, a majority of voters prefer c to d. Unfortunately, some natural axioms are not compatible with others, as famously shown in Arrow's impossibility theorem (Arrow 1950). For this reason, it is plausible to adopt a voting rule that satisfies as many natural axiomatic properties as possible.

Both Schulze's rule (*Schulze* for short) (Schulze 2011) and the *ranked pairs* rule (Tideman 1987), satisfy many natural and well-studied axiomatic properties. For example, among all voting rules we are aware of (including all

common voting rules), they are the only two that satisfy anonymity, Condorcet criterion, resolvability, Pareto optimality, reversal symmetry, monotonicity, and independence of clones (Schulze 2011). Moreover, the winner(s) can be computed in polynomial time in both rules. As a comparison to some well-known voting rules, for example, Borda does not satisfy the Condorcet criterion, Copeland does not satisfy resolvability, STV does not satisfy reversal symmetry, monotonicity, or the Condorcet criterion, and the maximin rule does not satisfy reversal symmetry. The Kemeny winner cannot be computed in polynomial-time unless P=NP. Only STV amongst these rules satisfies independence of clones.[1]

Schulze is currently being used in many real-world elections, including Wikimedia Foundation, the Pirate Party of Sweden, the Pirate Party of Germany, the Debian project, and the Gento Project (Schulze 2011), and is recognized as currently "the most widespread Condorcet method" on Wikipedia (http://en.wikipedia.org/wiki/Schulze_method#Use_of_the_Schulze_method). In sharp contrast, and despite sharing all these axiomatic properties, ranked pairs is not widely used. Schulze himself has remarked that ranked pairs "comes closest" amongst other election methods to his rule in terms of its axiomatic properties, differing only in that Schulze selects an alternative from the "min max" set: when the Schulze winner is unique, then the winner minimizes the maximum defeat in pairwise elections against other alternatives (Schulze 2011).

Our paper sets out to look for additional distinctions between Schulze and ranked pairs that can come from a computational viewpoint. A voting rule is *immune* to a type of strategic behavior if, for any profile, it is impossible to make the winner more preferable to a strategic individual. Otherwise the rule is *susceptible*. A voting rule is *resistant* (respectively, *vulnerable*) to a type of strategic behavior, if deciding whether the strategic individual can make the winner more preferable by using only the particular type of strategic behavior is NP-hard (respectively, is in P).

It is common to consider the following types of strategic behavior:

- *Manipulation*: a voter or a coalition of voters cast false vote(s) to make the winner more preferable. Having

---

[1]The definition of these rules and axiomatic properties can be found in the references. We omit them due to the space constraint.

rules that are susceptible to manipulation is inevitable due to the Gibbard-Satterthwaite theorem (Gibbard 1973; Satterthwaite 1975), but resistance can sometimes be achieved.

- *Control*: a chairman changes the procedure, for example by adding or deleting alternatives, adding or deleting votes, or introducing clones of alternatives, in order to make the winner more preferable (Bartholdi, Tovey, and Trick 1992; Tideman 1987). Control by adding votes is equivalent to *false-name manipulation* (Conitzer and Yokoo 2010).

- *Bribery*: a briber can change some votes by bribing the voters, in order to make the winner preferable (Faliszewski, Hemaspaandra, and Hemaspaandra 2009). The bribery problem is closely related to the problem of determining the *margin of victory* (Cary 2011; Magrino et al. 2011; Xia 2012).

All of these strategic behaviors have both *constructive* and *destructive* variants. In the constructive (respectively, destructive) variant, the strategic individual wants to make a favored alternative win (respectively, a disfavored alternative lose). These types of strategic behavior are generally not known to be comparable, in the sense that establishing resistance or vulnerability of one kind does not imply resistance or vulnerability of another kind.[2]

Many common voting rules are susceptible to control or bribery and it is common to look for computational resistance. But prior to our work we know of no voting rule that has been shown to be resistant (or immune) to constructive and destructive control in regard to the strategic behaviors considered here, or even in regard to control via changing the set of alternatives or adding or deleting votes. For example, Copeland and Llull are vulnerable to destructive control by adding (deleting) alternatives (Faliszewski et al. 2009), plurality is vulnerable to constructive and destructive control by adding (deleting) votes, approval voting and sincere-strategy preference-based approval voting are vulnerable to destructive control by adding (deleting) votes (Hemaspaandra, Hemaspaandra, and Rothe 2007; Erdélyi, Nowak, and Rothe 2008).

Our contributions. For Schulze and ranked pairs, it was only known that both are immune to control by cloning alternatives (Tideman 1987; Schulze 2011), and that ranked pairs is resistant to constructive manipulation for any fixed number of manipulators (Xia et al. 2009).[3] We show that ranked pairs is resistant to all types of strategic behavior under consideration. To the best of our knowledge, this is the first time that a voting rule is known to resist all these types of strategic behavior. Schulze also provides good re-

sistance, but remains vulnerable to constructive manipulation by a single manipulator and destructive manipulation for any fixed number of manipulators, while the computational resistance of Schulze in regard to destructive control by adding alternatives or deleting votes remains open. Due to the space constraint, some proofs are omitted. They can be found on the second author's homepage.

Admittedly, NP-hardness is merely a worst-case concept. For example, manipulation has been shown to be computationally easy for any "reasonable" voting rule, including ranked pairs and Schulze, in the sense that with a high probability, the manipulators can decide whether they can make their favored alternative win by just looking at the number of manipulators, and this leads to a polynomial-time algorithm they can use to compute a successful manipulation (Xia and Conitzer 2008). See (Faliszewski and Procaccia 2010; Faliszewski, Hemaspaandra, and Hemaspaandra 2010; Rothe and Schend 2012) for recent surveys. Still, having a worst-case barrier is better than no barrier, and no average-case analysis is available for the other kinds of strategic behavior considered here. Our results provide a computational differentiator between Schulze and ranked pairs, and also distinguish ranked pairs in satisfying resistance properties that are not known for other rules. Given the broad axiomatic support for both Schulze and ranked pairs and given that Schulze is in wide use, there seems to be good support to adopt ranked pairs in practical applications.

Preliminaries

Let $\mathcal{C}$ denote the set of *alternatives*, $|\mathcal{C}| = m$. We assume strict preference orders. That is, a vote is a linear order over $\mathcal{C}$. The set of all linear orders over $\mathcal{C}$ is denoted by $L(\mathcal{C})$. A *preference-profile* P is a collection of n votes for some $n \in \mathbb{N}$, that is, $P \in L(\mathcal{C})^n$. A *voting rule* r is a mapping that assigns to each preference-profile a set of non-empty winning alternatives. That is, $r : L(\mathcal{C})^n \to (2^{\mathcal{C}} \setminus \emptyset)$. Throughout the paper, we let n denote the number of votes and m denote the number of alternatives.

For any profile P and any pair of alternatives $\{c, d\}$, let $D_P(c, d)$ denote the number of times that $c \succ d$ in P minus the number of times that $d \succ c$ in P. The *weighted majority graph (WMG)* is a directed graph whose vertices are the alternatives, and there is an edge between every pair of vertices, where the weight on $c \to d$ is $D_P(c, d)$. We note that in the WMG of any profile, all weights on the edges have the same parity (and whether this is odd or even depends on the number of votes), and $D_P(c, d) = -D_P(d, c)$.

Ranked pairs: This rule selects the winner by first creating a ranking of all the alternatives and then selecting as the winner the alternative at the top. In each step, consider a pair of alternatives c_i, c_j that have not previously been considered; specifically, choose a remaining pair with the highest $D_P(c_i, c_j)$. When there are multiple pairs with the highest weight, break ties according to a fixed tie-breaking method on pairs of alternatives. Fix the order $c_i \succ c_j$, unless this contradicts previous orders that we fixed (that is, it violates transitivity), in which case this pair is ignored going forward. Continue until all pairs of alternatives are considered,

---

[2]An exception is that the vulnerability to the constructive variant of some type of strategic behavior implies the vulnerability to the destructive variant. (For every alternative d other than the disfavored alternative c, we just check if it is possible to make d win.) However, if we know that a voting rule is resistant to the destructive variant, we only know that computing the constructive variant is not in P (unless P= NP); we do not immediately know whether computing the constructive variant is NP-hard. Some other connections can be found in Section 4 of (Faliszewski, Hemaspaandra, and Hemaspaandra 2009).

[3]The proof in (Xia et al. 2009) works for ranked pairs with fixed tie-breaking and ranked pairs with parallel-universe tie-breaking.

ending with a ranking of all alternatives.[4]

Schulze: In the WMG of a profile P, the *weight* of a path $d_1 \rightarrow d_2 \rightarrow \cdots \rightarrow d_l$ is $\min_{i \leq l-1} D_P(d_i, d_{i+1})$.[5] That is, the weight of a path is the weight of the weakest edge along the path. We note that the weight can be negative. For any pair of alternatives (c, d), let $S_P(c, d)$ denote the maximum weight path from c to d, giving the Schulze score a "max min" semantics. Sometimes the subscript P is omitted when there is no chance of confusion. A Schulze winner is an alternative c such that for any other alternative d, $S_P(c, d) \geq S_P(d, c)$. Such an alternative always exists (Schulze 2011).

Example 1 *The alternatives are $\{a, b, c, d\}$. Let P be a profile whose WMG is illustrated in Figure 1.*
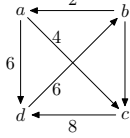

Figure 1: The weighted majority graph of the profile in Example 1. For simplicity, this figure only shows edges with positive weights, ignoreing edges with negative weights.

For ranked pairs, the following edges are fixed in order: $b \rightarrow c$, $c \rightarrow d$, $a \rightarrow d$, $b \rightarrow a$, and the winner is b. For Schulze, we have $S_P(a, b) = S_P(a, c) = S_P(a, d) = 6 > 2 = S(b, a) = S(c, a) = S(d, a)$, and the winner is a.

Manipulation, Control, and Bribery

In the context of strategic behavior, we will say that an alternative "wins" for Schulze, if it is the unique winner. Our results also extend to the case when being in the winning set is sufficient. If not mentioned specifically, all our results about ranked pairs hold for any fixed tie-breaking mechanism.

Definition 1 *In a constructive (respectively, destructive)* MANIPULATION *problem, we are given a voting rule r, M manipulators, a profile P^{NM} of the non-manipulators, and a (dis)favored alternative $c \in \mathcal{C}$. We are asked whether the manipulators can cast their votes P^* such that $\{c\} = r(P^{NM} \cup P^*)$ (respectively, $\{c\} \neq r(P^{NM} \cup P^*)$).*

This problem is also known as the *unweighted coalitional manipulation* problem in the literature.

Definition 2 *In a constructive (respectively, destructive)* CONTROL-ADD-ALT *problem, we are given a set of alternatives $\mathcal{C}$, a set of new alternatives $\mathcal{C}'$, a profile over $\mathcal{C} \cup \mathcal{C}'$, a quota $k \leq |\mathcal{C}'|$, and a (dis)favored alternative $c \in \mathcal{C}$. We are asked whether the chairman can add no more than k alternatives in $\mathcal{C}'$ such that c is the unique winner (repsectively, c is not the unique winner).*

Definition 3 *In a constructive (respectively, destructive)* CONTROL-DEL-ALT *problem, we are given a profile over*

$\mathcal{C}$, *a quota $k < |\mathcal{C}|$ and a (dis)favored alternative $c \in \mathcal{C}$. We are asked whether the chairman can delete no more than k alternatives such that c is the unique winner (respectively, c is not the unique winner, and without deleting c).*

Definition 4 *In a constructive (respectively, destructive)* CONTROL-ADD-VOTE *problem, we are given n votes and a set N' of additional votes, a quota $k \leq |N'|$ and a (dis)favored alternative $c \in \mathcal{C}$. We are asked whether the chairman can add no more than k votes from N' such that c is the unique winner (respectively, c is not the unique winner).*

Definition 5 *In a constructive (respectively, destructive)* CONTROL-DEL-VOTE *problem, we are given n votes, a quota $k < n$, and a (dis)favored alternative $c \in \mathcal{C}$. We are asked whether the chairman can delete no more than k votes such that c is the unique winner (respectively, c is not the unique winner).*

Definition 6 *In a constructive (respectively, destructive)* BRIBERY *problem, we are given a set of n votes, a quota $k < n$, and a (dis)favored alternative $c \in \mathcal{C}$. We are asked whether the briber can change no more than k votes such that c is the unique winner (respectively, c is not the unique winner).*

Many proofs in this paper use the *McGarvey's trick* (McGarvey 1953), which constructs a profile whose WMG is the same as some targeted weighted directed graph. This will be helpful because when we present the proof, we only need to specify the WMG instead of the whole profile, and then by using the McGarvey's trick, a profile can be constructed in polynomial time. The trick works as follows. For any pair of alternatives, (c, d), if we add two votes $[c \succ d \succ c_3 \succ \cdots \succ c_m], [c_m \succ c_{m-1} \succ \cdots \succ c_3 \succ c \succ d]$ to a profile P, then in the WMG, the weight on the edge $c \rightarrow d$ is increased by 2 and the weight on the edge $d \rightarrow c$ is decreased by 2, while the weights on other edges are unchanged. Moreover, for any given alternative $c' \in \mathcal{C} \setminus \{c, d\}$, we can switch c' and $c_{\lceil m/2 \rceil}$ in the two votes, such that c' is always ranked within the top $\lceil m/2 \rceil + 1$ positions in the two votes. If we require c or d be among top $\lceil m/2 \rceil + 1$ positions, then we add the following two votes. $[c_3 \succ \cdots \succ c_{\lceil m/2 \rceil + 1} \succ c \succ d \succ c_{\lceil m/2 \rceil + 2} \succ \cdots \succ c_m], [c_m \succ \cdots \succ c_{\lceil m/2 \rceil + 2} \succ c \succ d \succ c_{\lceil m/2 \rceil + 1} \succ \cdots \succ c_3]$. All of these variants on vote pairs provide the same effect on the weights of $c \rightarrow d$ and $d \rightarrow c$ while leaving other weights unchanged.

Manipulation

We first investigate constructive MANIPULATION for Schulze when there is one manipulator. Fix alternative c. Let $S(\cdot, \cdot)$ denote $S_{P^{NM}}(\cdot, \cdot)$. We note that the manipulator's vote can only change the weight of any edge in the WMG by 1. Hence, if there exists another alternative $d \neq c$ such that $S(d, c) \geq S(c, d) + 2$, then c cannot be made the (unique) winner by a single manipulator. In other words, if c can be made the winner by a single manipulator, then c must be a co-winner of P^{NM} under Schulze. It has been proved that for any Schulze co-winner c, there exists a vote V such that after V is added to the profile, c becomes the unique

---

[4]In this paper, ranked pairs is a *resolute* rule, that is, it only selects a single winner. For another version of ranked pair that selects multiple winners by using the *parallel-universe tie-breaking mechanism* (Conitzer, Rognlie, and Xia 2009), it is NP-complete to compute the winners (Brill and Fischer 2012).

[5]The original definition of the Schulze rule (Schulze 2011) is slightly different, but the voting rule is the same when the voters' preferences are linear orders.

winner (the Claim in Section 4.2.2 in (Schulze 2011)). This property is called "second version of the resolvability property". Moreover, Schulze's proof for the satisfiability of the property is constructive, where V can be computed in polynomial time. Therefore, given a MANIPULATION instance, if c is not a co-winner of P^{NM} then c cannot be made the winner by a single manipulator, otherwise c can be made the winner by a single manipulator. We immediately have the following corollary.[6]

Corollary 1 *There is a polynomial time algorithm for Constructive MANIPULATION with one manipulator in Schulze.*

The algorithm does not have a straightforward extension to two or more manipulators. For example, suppose there are four alternatives $\{a, b, c, d\}$ and two manipulators. The WMG of P^{NM} is illustrated in the following figure.
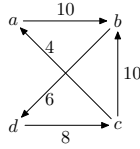

Figure 2: The weighted majority graph of P^{NM}.

Note that it is optimal for the manipulator to rank c in the top position. It is possible for the manipulators to make c beat a (because $S_{PNM}(c, a) + 4 = 8 > 6 = S_{PNM}(a, c)$) or make c beat d (because $S_{PNM}(c, d) + 4 = 10 > 8 = S_{PNM}(d, c)$). However, the weight on $b \to d$ must be increased by two for c to beat d, and it must be decreased by two for c to beat a. These two objectives cannot be achieved at the same time, which means that c cannot be made the winner.

Now we turn to destructive MANIPULATION. The following lemma gives a necessary and sufficient condition for c to be the unique Schulze winner.

Lemma 1 *Alternative c is the unique Schulze winner of a profile P if and only if for every other alternative d, $S_P(c, d) > S_P(d, c)$.*

We next present a polynomial-time algorithm for destructive MANIPULATION for Schulze. Let $S(\cdot, \cdot)$ denote $S_{PNM}(\cdot, \cdot)$. We will show that the M manipulators can make c not be the unique winner if and only if there exists an alternative $d \neq c$ such that $S(d, c) \geq S(c, d) - 2M$. It is easy to check that if $S(d, c) < S(c, d) - 2M$ holds for every alternative d that is different from c, then c must remain the unique winner. Otherwise, the algorithm will find a successful destructive manipulation in the following way. Let d denote an alternative such that $S(d, c) \geq S(c, d) - 2M$ and $S(d, c)$ is maximized. Let $k = S(d, c)$. Then, we choose an arbitrary simple path (that is, a path that does not contain a cycle) $d(= f_0) \to f_1 \to \cdots \to f_p \to c$ with weight k. The algorithm will find an alternative f_T along the path to beat c. That is, the algorithm will compute a profile P^* for the manipulators such that $S_{PNM \cup P^*}(f_T, c) = k + M \geq$

---
[6]We thank Markus Schulze for pointing out the relationship between constructive MANIPULATION with one manipulator and the second version of resolvability.

Algorithm 1: DestructiveManipulation

Input: P^{NM}, c, and the number of manipulators M.

1   Compute $S_{PNM}(\cdot, \cdot)$. Denote this as $S(\cdot, \cdot)$
2   **if** *for all $d \neq c$, $S(d, c) < S(c, d) - 2M$* **then**
3 | Output: No destructive manipulation.
4   **end**
5   Let $d = \arg\max_{d': S(d', c) \geq S(c, d') - 2M} \{S(d', c)\}$ and $k = S(d, c)$. Let $[d(= f_0) \to f_1 \to \cdots \to f_p \to c]$ denote an arbitrary simple path from d to c with weight k. Let $R = \{d \succ f_1, f_1 \succ f_2, \ldots, f_p \succ c\}$.
6   **if** $S(d, c) \geq S(c, d)$ **then**
7 | **return** a profile where all manipulators cast the same vote that extends R.
8   **end**
9   **else**
10 | Let $0 \leq T \leq p$ denote the maximum number such that $S(f_T, c) = k$ and $S(c, f_T) \leq k + 2M$.
11 | Let $R_1 = \{f_T \succ f_{T+1}, \ldots, f_p \succ c\}$ and $R_2 = \emptyset$.
12 | Compute all pairs of alternatives (a, b) such that $S(b, f_T) \geq S(c, f_T) + 2$ and $S(a, f_T) \leq S(c, f_T)$. Add $b \succ a$ to R_2.
13 | **return** a profile where all manipulators cast the same vote that extends $R = R_1 \cup R_2$.
14   **end**

$S(c, f_T) - M = S_{PNM \cup P^*}(c, f_T)$, which means that c is not the unique Schulze winner in $P^{NM} \cup P^*$ (Lemma 1).

The algorithm also illustrates that, if the manipulators can cast votes such that c is not the unique winner, then they can also do so by casting the same vote.

Theorem 1 *Algorithm 1 runs in polynomial time and computes destructive MANIPULATION for Schulze.*

Proof sketch: It is not hard to show that if all manipulators' votes extend R, then c is not the unique winner because at least d beats c. We now prove that there is no no cycle in $R_1 \cup R_2$.

We first prove that for every i such that $T < i \leq p$, $S(c, f_i) > k + 2M$. For contradiction, suppose there exists an i, with $T < i \leq p$, for which $S(c, f_i) \leq k + 2M$. Because the weight of the path from d to c is k, then $S(f_i, c) \geq k$. If $S(f_i, c) = k$, then i violates the maximality of T, which is a contradiction. So $S(f_i, c) > k$. Now, since k is the largest $S(d, c)$ such that $S(d, c) \geq S(c, d) - 2M$, we must have $S(f_i, c) < S(c, f_i) - 2M$. Hence, we have $S(c, f_i) > S(f_i, c) + 2M > k + 2M$ and a contradiction.

We are now ready to prove that there is no cycle in $R_1 \cup R_2$. For contradiction, suppose there is a cycle in $R_1 \cup R_2$. Because R_1 is composed of edges of a simple path from f_T to c, one or more edges from R_2 must be involved, and we need to consider the following two cases.

Case 1: The cycle contains an edge in R_1. Then, there exist $(a \succ b) \in R_1$ and $(b \succ e) \in R_2$ in the cycle. This case is illustrated in Figure 3. By construction, since $(a \succ b) \in R_1$ then $b = f_i$ for some $T < i \leq p$, or $b = c$. Because $(b \succ e) \in R_2$, $S(b, f_T) \geq S(c, f_T) + 2$ (Step 12). We see that $b \neq c$. Now, we have shown above that $S(c, b) > k + 2M \geq S(c, f_T)$. Therefore, $S(c, f_T) \geq \min\{S(c, b), S(b, f_T)\} > S(c, f_T)$, and a contradiction.

	MANIPULATION	Cons.ADDALT	Des.ADDALT	DELALT	{ADD,DEL}VOTE	BRIBERY
Ranked pairs	**R** (Xia et al. 2009) (Thm. 2)	**R** (Thm. 3)	**R** (Thm. 3)	**R** (Thm. 4)	**R** (Thm. 5, Thm. 6)	**R** (Xia 2012) (Thm. 7)
Schulze	**V** (Coro. 1, Thm. 1)		?	?		

Table 1: Summary of results. If not mentioned specifically, the results hold for both constructive and destructive cases. Question marks represent open questions.
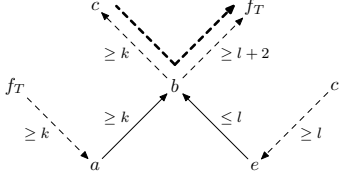

Figure 3: A contradiction for the case $(a \succ b) \in R_1$ and $(b \succ e) \in R_2$, where $l = S(c, f_T)$.

Case 2: The cycle is only composed of edges in R_2. Then, there exist $(e \succ b) \in R_2$ and $(b \succ a) \in R_2$ in the cycle. Because $(b \succ a)$ and $(e \succ b)$ are in R_2 then by applying the first and second conditions in Step 12, we have $S(b, f_T) \geq S(c, f_T) + 2$ and $S(b, f_T) \leq S(c, f_T)$, respectively, and a contradiction.

It follows that Algorithm 1 runs in polynomial time and computes destructive MANIPULATION. $\square$

Theorem 2 *It is* NP-*complete to compute destructive* MA-NIPULATION *for ranked pairs, for any fixed number of manipulators and any fixed tie-breaking mechanism.*

Proof sketch: The proof is by tweaking the NP-completeness proof of constructive manipulation for ranked pairs (Xia et al. 2009), which was a reduction from 3SAT. In the proof in (Xia et al. 2009), a 3SAT instance has a solution if and only if c can be made to win. We add one alternative d, and let d beat all other alternatives except c in their pairwise elections, and the weights on these edges in the WMG are high, so that these edges will be fixed first. We also let $d \to c$ with weight $2M - 2$, so that if all manipulators rank c in the top places, then c barely beats d in their pairwise election. We are asked whether d can be made to lose. We note that if d is not the winner for ranked pairs, then c must be the winner, which happens if and only if the 3SAT instance has a solution. This proves the NP-hardness. $\square$

Control and Bribery

It is easy to come up with examples where Schulze and ranked pairs are not immune to constructive and destructive variants of control by adding or deleting alternatives, control by adding or deleting votes, and bribery. In this section, we study whether Schulze and ranked pairs are able to *resist* these types of control and bribery.

Theorem 3 *Schulze and ranked pairs resist constructive* CONTROL-ADD-ALT. *Ranked pairs resists destructive* CONTROL-ADD-ALT.

Proof sketch: We first prove that Schulze and ranked pairs resist constructive CONTROL-ADD-ALT through a reduction from EXACT COVER BY 3-SETS (X3C) (Garey and

Johnson 1979). In an X3C instance, we are given two sets $\mathcal{A} = \{a_1, \ldots, a_q\}$ (where q is a multiple of 3) and $\mathcal{E} = \{E_1, \cdots, E_t\}$, where for each $E \in \mathcal{E}$, $E \subseteq \mathcal{A}$ and $|E| = 3$. We are asked whether there exist $q/3$ elements $\mathcal{E}' = \{E_{j_1}, \ldots, E_{j_{q/3}}\}$ in $\mathcal{E}$ such that each element in $\mathcal{A}$ appears in one and exactly one element in $\mathcal{E}'$. Given an X3C instance, we construct an election, where the old alternatives are $\{c\} \cup \mathcal{A}$ and the new alternatives are $\mathcal{E}$. If no new alternative is added, then c loses to all other alternatives in pairwise elections. Each alternative E in $\mathcal{E}$ is used to introduce a strong path from c to the three alternatives in E, and we are asked whether the chair can introduce no more than $q/3$ alternatives to make c win. Formally, the election is defined as follows.

• **Alternatives:** the old alternatives are $\{c\} \cup \mathcal{A}$ and the new alternatives are $\mathcal{E}$. We are asked whether c can be made win by adding no more than $q/3$ alternatives in $\mathcal{E}$.

• **Profile:** Because of the McGarvey's trick, it suffices to specify the WMG of the profile, where we have the following edges.

– For each $i \leq q$, there is an edge $a_i \to c$ with weight 4.

– For each $E \in \mathcal{E}$, there is an edge $c \to E$ with weight 6, and for every $c' \in E$, there is an edge $E \to c'$ with weight 6.

– All edges not defined above have zero weight.

It is easy to check that c is the (unique) Schulze (respectively, ranked pairs) winner if and only if the $q/3$ new alternatives correspond to an exact cover of $\mathcal{A}$.

For the destructive variant, we add d to the set of old alternatives, and in the WMG, let $c \to d$ with weight 2, and there are edges from d to all other alternatives with weight 6. It follows that if d is not the ranked pairs winner if and only if c is the ranked pairs winner. $\square$

Theorem 4 *Ranked pairs resists constructive and destructive* CONTROL-DEL-ALT.

Proof sketch: We first prove that ranked pairs resist constructive CONTROL-DEL-ALT through a reduction from X3C. Given an X3C instance $\mathcal{E} = \{E_1, \cdots, E_t\}$ over $\mathcal{A} = \{a_1, \ldots, a_q\}$, we construct an election, where the alternatives are $\{c\} \cup \mathcal{A} \cup \mathcal{E} \cup \{c_i^j : \forall c_i \in E_j\}$. We are asked whether the chair can delete no more than $q/3$ alternatives to make c win. If no alternative is deleted, then all alternatives in $\mathcal{A}$ are ranked above c. Removing each alternative E in $\mathcal{E}$ will introduce a strong path from c to the three alternatives in E. Therefore, the only way to make c win is to remove the $q/3$ alternatives in $\mathcal{E}$ that correspond to a 3-cover. Formally, the election is defined as follows.

• **Alternatives:** $\{c\} \cup \mathcal{A} \cup \mathcal{E} \cup \{c_i^j : \forall c_i \in E_j\}$. We are asked whether c can be made win by deleting no more than

$q/3$ alternatives.

- **Profile:** Again, it suffices to specify the WMG of the profile. In the WMG, we have the following edges.

    − For each $i \leq q$, there is an edge $a_i \to c$ with weight 4; for all j such that $a_i \in E_j$, there is an edge $c \to c_i^j$ with weight 8 and an edge $c_i^j \to a_i$ with weight 6.

    − For each $E_j \in \mathcal{E}$ and each $a_i \in E_j$, there is an edge $a_i \to E_j$ with weight 8, and an edge $E_j \to c_i^j$ with weight 8.

    − All edges $c \to c'$ not defined above have weight 2. Other edges have weight 0.

It is easy to check that c is the ranked pairs winner if and only if the $q/3$ deleted alternatives are in $\mathcal{E}$, and they correspond to an exact cover of $\mathcal{A}$.

For the destructive variant, we add d to the set of alternatives, and in the WMG, let $c \to d$ with weight 2, and there are edges from d to all other alternatives with weight 8. It follows that if d is not the ranked pairs winner if and only if c is the ranked pairs winner. □

Theorem 5 *Schulze and ranked pairs resist constructive and destructive* CONTROL-ADD-VOTE.

Proof sketch: We first prove that Schulze and ranked pairs resist constructive CONTROL-ADD-VOTE through a reduction from X3C. Given an X3C instance $\mathcal{E} = \{E_1, \cdots, E_t\}$ over $\mathcal{A} = \{a_1, \ldots, a_q\}$, we construct an election as follows.

- **Alternatives:** $\mathcal{C} = \{c, d\} \cup \mathcal{A}$. We are asked whether c can be made to win by adding no more than $q/3$ votes. For ranked pairs, the proof is for tie-breaking order where $c \to d$ is ranked in the bottom. Our proof can be extended to any other tie-breaking orders. (For the co-winner variant of the problem for Schulze, we add a new alternative to introduce strong paths from c to $\mathcal{A}$.)

- **Profile:** For each $j \leq t$, there is a new vote $V_j = [E_j \succ c \succ d \succ (\mathcal{A} \setminus E_j)]$. The WMG of the old votes is defined as follows.

    − There is an edge $c \to d$ with weight $2q/3 + 2$.

    − For each $i \leq q$, there are an edge $d \to a_i$ with weight $2q$, and an edge $a_i \to c$ with weight $4q/3 - 2$.

    − All other edges have weight 0.

We note that adding V_j will increase the weight on $c \to d$ and $c \to c'$ (for all $c' \in (\mathcal{A} \setminus E_j)$) by 1, and decrease $c \to c'$ for all $c' \in E_j$ by 1. The only way for c to win for Schulze or ranked pairs is that after $q/3$ new votes are added, in the WMG $c \to d$ with weight $q + 2$, and for all $i \leq q$, $a_i \to c$ with weight q. This happens if and only if the votes added by the chairman correspond to an exact cover of $\mathcal{A}$. Therefore, constructive control by adding votes is NP-hard for Schulze and ranked pairs.

For the destructive variant, we are asked whether the chairman can add no more than $q/3$ votes such that d does not win. We note that if c is not the ranked pairs winner, then d is the ranked pairs winner. Therefore, destructive control for ranked pairs is NP-hard. For Schulze, we make the following changes to the above reduction: the weight on the edge $c \to d$ is $2q/3$. □

Using similar reductions, we can prove the following two theorems.

Theorem 6 *Schulze and ranked pairs resist constructive and destructive* CONTROL-DEL-VOTE.

Proof sketch: The proof is similar to the proof of Theorem 5. The differences are: we specify the WMG for the old votes plus the new votes, and for each $j \leq t$, there is a voter whose vote is $V_j = [d \succ (\mathcal{A} \setminus E_j) \succ c \succ E_j]$. Therefore, removing V_j will increase the weight on $c \to d$ and $c \to c'$ (for all $c' \in (\mathcal{A} \setminus E_j)$) by 1, and decrease $c \to c'$ for all $c' \in E_j$ by 1. When using McGarvey's trick, we always rank c among top $\lfloor q/2 \rfloor + 2$ positions. We need this constraint to make sure that the chairman can only make c to win by removing votes in $\{V_j\}$.

- **Alternatives:** $\mathcal{C} = \{c, d, c_1, \ldots, c_q\}$. We are asked whether c can be made win by deleting no more than $q/3$ votes. For ranked pairs, the proof is for tie-breaking order where $c \to d$ is ranked in the bottom. Our proof can be extended to other tie-breaking orders. (For co-winner variant of the problem for Schulze, we add a new alternative to introduce strong paths from c to $\{c_1, \ldots, c_q\}$.)

- **Profile:** For each $j \leq t$, there is a voter whose vote is $V_j = [d \succ (\mathcal{C} \setminus S_j) \succ c \succ S_j]$. The WMG of the profile is defined as follows. When using McGarvey's trick, we always rank c among top $\lfloor q/2 \rfloor + 2$ positions. We need this constraint to make sure that the chairman can only make c to win by removing votes in $\{V_j\}$.

    − There is an edge $c \to d$ with weight $2q/3 + 2$.

    − For each $i \leq q$, there are an edge $d \to c_i$ with weight $2q$, and an edge $c_i \to c$ with weight $4q/3 - 2$.

    − All other edges not defined above have weight 0.

The only way for c to win for Schulze or ranked pairs is that after $q/3$ votes are deleted, in the WMG $c \to d$ with weight $q + 2$, and for all $i \leq q$, $c_i \to c$ with weight q. This happens if and only if the chair only delete votes in $\{V_j\}$, and these votes correspond to an exact cover of C. Therefore, constructive control by deleting votes is NP-hard for Schulze and ranked pairs.

For destructive control, we are asked whether the chairman can add no more than $q/3$ votes such that d does not win. We note that if c is not the ranked pairs winner, then d is the ranked pairs winner. Therefore, destructive control for ranked pairs is NP-hard. For Schulze, we make the following changes to the above reduction: the weight on the edge $c \to d$ is $2q/3$. □

It was shown (Xia 2012) that for some fixed-order tie-breaking mechanisms, ranked pairs resists constructive and destructive BRIBERY, even when the briber's quota is 1. We note that if the briber's quota is bounded above by a constant, then destructive BRIBERY for Schulze is easy (by enumerating all possible combinations of votes to remove, and then the problem becomes a standard destructive MANIPULATION problem, which is in P due to Theorem 1). We show that if the quota is not bounded, then Schulze also resists constructive and destructive BRIBERY.

Theorem 7 *Schulze resists constructive and destructive* BRIBERY.

Proof sketch: The proof is similar to the proof of Theorem 6. The differences are: For each $i \leq q$, there is an edge

$a_i \rightarrow c$ with weight $2q - 2$. Again, when using McGarvey's trick, we always rank c among top $\lceil q/2 \rceil + 2$ positions. $\square$

Summary and Future Work

In this paper, we investigated the computational complexity of strategic behavior in Schulze and ranked pairs rules. Our results and open questions are summarized in Table 1. We also conjecture some other voting rules that resist constructive MANIPULATION for even one manipulator, for example, STV (Bartholdi and Orlin 1991), Nanson's rule and Baldwin's rule (Davies et al. 2012), also broadly resist these types of strategic behavior. There remain other types of bribery and control to consider, for example control by partitioning alternatives or voters, and run-off partitioning of alternatives (Bartholdi, Tovey, and Trick 1992), as well as multi-mode control and bribery (Faliszewski, Hemaspaandra, and Hemaspaandra 2011). Most importantly, the analysis here adopts worst-case intractability and it is important to consider average case analysis for these richer kinds of strategic behavior that reach beyond manipulation.

Acknowledgements

References

Arrow, K. J. 1950. A difficulty in the concept of social welfare. *Journal of Political Economy* 58(4):328–346.

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. 1992. How hard is it to control an election? *Math. Comput. Modelling* 16(8-9):27–40. Formal theories of politics, II.

Brill, M., and Fischer, F. 2012. The Price of Neutrality for the Ranked Pairs Method. In *Proceedings of AAAI*.

Cary, D. 2011. Estimating the margin of victory for instant-runoff voting. In *Proceedings of 2011 EVT/WOTE Conference*.

Conitzer, V., and Yokoo, M. 2010. Using mechanism design to prevent false-name manipulations. *AI Magazine* 31(4):65–77.

Conitzer, V.; Rognlie, M.; and Xia, L. 2009. Preference functions that score rankings and maximum likelihood estimation. In *Proceedings of IJCAI*, 109–115.

Davies, J.; Katsirelos, G.; Narodytska, N.; Walsh, T.; and Xia, L. 2012. Complexity of and Algorithms for the Manipulation of Borda, Nanson and Baldwin's Voting Rules. *Artificial Intelligence, to appear*.

Erdélyi, G.; Nowak, M.; and Rothe, J. 2008. Sincere-strategy preference-based approval voting broadly resists control. In *Proceedings of MFCS*, 311–322.

Faliszewski, P., and Procaccia, A. D. 2010. AI's war on manipulation: Are we winning? *AI Magazine* 31(4):53–64.

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

Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2009. How hard is bribery in elections? *Journal of Artificial Intelligence Research* 35:485–532.

Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2010. Using complexity to protect elections. *Commun. ACM* 53:74–82.

Faliszewski, P.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2011. Multimode control attacks on elections. *Journal of Artificial Intelligence Research* 40:305–351.

Garey, M., and Johnson, D. 1979. *Computers and Intractability*. W. H. Freeman and Company.

Gibbard, A. 1973. Manipulation of voting schemes: A general result. *Econometrica* 41:587–601.

Hemaspaandra, E.; Hemaspaandra, L. A.; and Rothe, J. 2007. Anyone but him: The complexity of precluding an alternative. *Artificial Intelligence* 171(5–6):255–285.

Magrino, T. R.; Rivest, R. L.; Shen, E.; and Wagner, D. 2011. Computing the Margin of Victory in IRV Elections. In *Proceedings of 2011 EVT/WOTE Conference*.

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

Rothe, J., and Schend, L. 2012. Typical-Case Challenges to Complexity Shields That Are Supposed to Protect Elections Against Manipulation and Control: A Survey. In *International Symposium on Artificial Intelligence and Mathematics*.

Satterthwaite, M. 1975. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory* 10:187–217.

Schulze, M. 2011. A new monotonic, clone-independent, reversal symmetric, and condorcet-consistent single-winner election method. *Social Choice and Welfare* 36(2):267—303.

Tideman, T. N. 1987. Independence of clones as a criterion for voting rules. *Social Choice and Welfare* 4(3):185–206.

Xia, L., and Conitzer, V. 2008. Generalized scoring rules and the frequency of coalitional manipulability. In *Proceedings of EC*, 109–118.

Xia, L.; Zuckerman, M.; Procaccia, A. D.; Conitzer, V.; and Rosenschein, J. 2009. Complexity of unweighted coalitional manipulation under some common voting rules. In *Proceedings of IJCAI*, 348–353.

Xia, L. 2012. Computing the margin of victory for various voting rules. In *Proceedings of EC*.