

Complexity of Unweighted Coalitional Manipulation Under Some Common Voting Rules

Lirong Xia* Michael Zuckerman† Ariel D. Procaccia‡
Vincent Conitzer* Jeffrey S. Rosenschein†

* Department of Computer Science, Duke University, Durham, NC 27708, USA. {lxia,conitzer}@cs.duke.edu

† School of Computer Science and Engineering, The Hebrew Univ. of Jerusalem, Israel. {michez,jeff}@cs.huji.ac.il

‡ Microsoft Israel R&D Center, 13 Shenkar Street, Herzeliya 46725, Israel. arielpro@gmail.com

Abstract

Understanding the computational complexity of manipulation in elections is arguably the most central agenda in Computational Social Choice. One of the influential variations of the the problem involves a coalition of manipulators trying to make a favorite candidate win the election. Although the complexity of the problem is well-studied under the assumption that the voters are weighted, there were very few successful attempts to abandon this strong assumption.

In this paper, we study the complexity of the *unweighted coalitional manipulation problem (UCM)* under several prominent voting rules. Our main result is that UCM is NP-complete under the *maximin* rule; this resolves an enigmatic open question. We then show that UCM is NP-complete under the *ranked pairs* rule, even with respect to a single manipulator. Furthermore, we provide an extreme hardness-of-approximation result for an optimization version of UCM under ranked pairs. Finally, we show that UCM under the *Bucklin* rule is in P.

1 Introduction

Voting is a methodology that enables a group of agents (or voters) to make a joint choice from a set of candidates. Each agent reports his or her preferences over the candidates; then, a *voting rule* is applied to aggregate the preferences of the agents—that is, to select a winning candidate. However, sometimes a subset of the agents can report their preferences insincerely to make the outcome more favorable to them. This phenomenon is known as *manipulation*. A rule for which no group of agents can ever beneficially manipulate is said to be *group strategy-proof*; if no single agent can ever beneficially manipulate, the rule is said to be *strategy-proof* (a weaker requirement).

Unfortunately, any strategy-proof voting rule will fail to satisfy some natural property. The celebrated Gibbard-Satterthwaite theorem [Gibbard, 1973; Satterthwaite, 1975] states that when there are three or more candidates, there is no strategy-proof voting rule that satisfies non-imposition (for every candidate, there exist votes that would make that candidate win) and non-dictatorship (the rule does not simply

always choose the most-preferred candidate of a single fixed voter). However, the mere existence of beneficial manipulations does not imply that voters will use them: in order to do so, voters must also be able to *discover* the manipulation, and this may be computationally hard. Recently, the approach of using computational complexity to prevent manipulation has attracted more and more attention.

In early work [Bartholdi *et al.*, 1989; Bartholdi and Orlin, 1991], it was shown that when the number of candidates is not bounded, the second-order Copeland and STV rules are hard to manipulate, even by a single voter. More recent research has studied how to modify other existing rules to make them hard to manipulate [Conitzer and Sandholm, 2003; Elkind and Lipmaa, 2005].

One of the most prominent problems considered in the context of manipulation in elections is known as *weighted coalitional manipulation (WCM)*. In this setting, there is a coalition of manipulative voters trying to coordinate their actions in a way that makes a specific candidate win the election. In addition, the voters are weighted; a voter with weight k counts as k voters voting identically. Previous work has established that this problem is computationally hard under a variety of prominent voting rules, even when the number of candidates is constant (see, e.g., [Conitzer *et al.*, 2007; Hemaspaandra and Hemaspaandra, 2007]).

However, the current literature contains very few results regarding the *unweighted* version of the coalitional manipulation problem (UCM), which is in fact more natural in most settings. This is not completely surprising since, unlike WCM, the UCM problem is quite unwieldy under many voting rules, that is, it has proven mathematically difficult to resolve its complexity, despite some effort by various researchers over the past few years. Notice that, as UCM is a special case of WCM, any tractability results from the weighted setting carry over to the unweighted setting. However, hardness results do not carry over. We argue that, if one wishes to claim that a given voting rule is resistant to coalitional manipulation, in most cases hardness of UCM would be more relevant than hardness of WCM.

A few very recent papers have directly dealt with UCM. Faliszewski *et al.* [2008] have shown that UCM is NP-complete under a family of voting rules derived from the Copeland rule, even when there are only two manipulators. In addition, Zuckerman *et al.* [2008] have established, as corollary

aries of their main theorems, that unweighted coalitional manipulation is tractable under the Veto and Plurality with Runoff voting rules. However, Zuckerman et al. also conjectured that UCM is intractable under the prominent Borda and Maximin voting rules.

Zuckerman et al. further observed that the unweighted coalitional manipulation setting admits an optimization problem that they called *unweighted coalitional optimization (UCO)*. The goal is to find the minimum number of manipulators required to make a given candidate win the election. They gave a 2-approximation algorithm for this problem under maximin (even though this problem was not previously known to be NP-hard), and an algorithm for Borda that finds an optimal solution up to an additive term of one.

Our results. In this paper, we study the computational complexity of the unweighted coalitional manipulation problem under several voting rules. Our main result is that the UCM problem under maximin is NP-complete for any fixed number of manipulators (≥ 2); thus we resolve the abovementioned conjecture of Zuckerman et al. in the positive. We next show that the UCM problem under ranked pairs is NP-complete, even when there is only one manipulator. This means that ranked pairs is a member of a very exclusive club of “natural” voting rules that have this property (which previously included only second-order Copeland and STV). We strengthen this result by providing a surprising, extreme hardness-of-approximation result for UCO under ranked pairs: it is hard to approximate the problem within a factor of $N^{1-\epsilon}$ where N is the input size and $\epsilon > 0$ is an arbitrary constant. Finally, we present a polynomial-time algorithm for the UCM problem under Bucklin.

2 Preliminaries

Let \mathcal{C} be the set of *candidates*. A linear order on \mathcal{C} is a transitive, antisymmetric, and total relation on \mathcal{C} . The set of all linear orders on \mathcal{C} is denoted by $L(\mathcal{C})$. An n -voter profile P on \mathcal{C} consists of n linear orders on \mathcal{C} . That is, $P = (R_1, \dots, R_n)$, where for every $i \leq n$, $R_i \in L(\mathcal{C})$. The set of all profiles on \mathcal{C} is denoted by $\mathcal{P}(\mathcal{C})$. In the remainder of the paper, we let m denote the number of candidates (that is, $|\mathcal{C}|$).

A *voting rule* r is a function from the set of all profiles on \mathcal{C} to \mathcal{C} , that is, $r : \mathcal{P}(\mathcal{C}) \rightarrow \mathcal{C}$. Below we formally define the three prominent voting rules that we study in this paper. For a definition of other voting rules that we mention, the reader is referred to the book by Tideman [2006], or to the preliminaries of one of the many papers on the subject (e.g., [Conitzer et al., 2007]).

- *Maximin*: Let the *advantage* of c_i over c_j with respect to P , denoted $\text{adv}_P(c_i, c_j)$, be the number of votes in P that rank c_i ahead of c_j . The winner is the candidate c that maximizes $\min\{\text{adv}_P(c, c') : c' \in \mathcal{C} \setminus \{c\}\}$.
- *Bucklin* (see, e.g., [Tideman, 2006]): A candidate c 's Bucklin score is the smallest number k such that more than half of the votes rank c among the top k candidates. The winner is the candidate that has the smallest Bucklin score.¹

¹Sometimes, ties are broken by the number of votes that rank a

- *Ranked pairs* [Tideman, 1987]: This rule first creates an entire ranking of all the candidates, as follows. Define $\text{adv}_P(c_i, c_j)$ as for the maximin rule. In each step, we consider a pair of candidates c_i, c_j that we have not previously considered (as a pair): specifically, we choose the remaining pair with the highest $\text{adv}_P(c_i, c_j)$. We then fix the order $c_i \succ c_j$, unless this contradicts previous orders that we fixed (that is, it violates transitivity). We continue until we have considered all pairs of candidates (hence, in the end, we have a full ranking). The candidate at the top of the ranking wins.

All of these rules allow for the possibility that multiple candidates end up tied for the win. Technically, therefore, they are really *voting correspondences*; a correspondence can select more than one winner. In the remainder of this paper, we will sometimes somewhat inaccurately refer to the above correspondences as rules.

Let us now turn to the definition of the problems that we are interested in investigating. We study the *constructive* manipulation variations, in which the goal is to make a given candidate win.

Definition 2.1. The *Unweighted Coalitional Manipulation (UCM)* problem is defined as follows. An instance is a tuple (r, P^{NM}, c, M) , where r is a voting rule, P^{NM} is the non-manipulators' profile, c is the candidate preferred by the manipulators, and M is the set of manipulators. We are asked whether there exists a profile P^M for the manipulators such that $r(P^{NM} \cup P^M) = \{c\}$.

The above definition uses the *unique winner* formulation, which is the common one in the literature. It is also possible to consider a *co-winner* formulation which is similar, only we require that $c \in r(P^{NM} \cup P^M)$, that is, instead of being the unique winner, c should be included among the set of winners. Our results hold for the co-winner formulation as well.

For any $k \in \mathbb{N}$, we let UCM_k be the subproblem of UCM in which the number of manipulators is k . That is, a UCM_k instance is a tuple (r, P^{NM}, c, M) where $|M| = k$.

Zuckerman et al. [2008] suggested that the unweighted manipulation setting allows for a natural optimization problem: the *unweighted coalitional optimization* problem. Given, essentially, an unweighted coalitional manipulation instance, we ask *how many* manipulators are needed in order to make c win. Formally:

Definition 2.2. The *Unweighted Coalitional Optimization (UCO)* problem is defined as follows. An instance is a tuple (r, P^{NM}, c) , where r is a voting rule, P^{NM} is the non-manipulators' profile, and c is the candidate preferred by the manipulators. We must find the minimum k such that there exists a set of manipulators M with $|M| = k$, and a profile P^M , that satisfies $r(P^{NM} \cup P^M) = \{c\}$.

3 Maximin

In this section, we prove our main result: the UCM problem under maximin is NP-complete. We thus resolve an an

candidate among the top k , but for simplicity we will not consider this tie-breaking rule here.

open question that has proved quite enigmatic over the past few years (see, e.g., [Zuckerman *et al.*, 2008]). The proof uses a reduction from the *two vertex disjoint paths in directed antisymmetric graph* problem, which is known to be NP-complete [Fortune *et al.*, 1980].

Definition 3.1. The *two vertex disjoint paths in directed graph* problem is defined as follows. We are given a directed graph G and two disjoint pairs of vertices (u, u') and (v, v') , where u, u', v, v' are all different from each other. We are asked whether there exist two directed paths $u \rightarrow u_1 \rightarrow \dots \rightarrow u_{k_1} \rightarrow u'$ and $v \rightarrow v_1 \rightarrow \dots \rightarrow v_{k_2} \rightarrow v'$ such that $u, u', u_1, \dots, u_{k_1}, v, v', v_1, \dots, v_{k_2}$ are all different from each other.

For any profile P and any pair of candidates c_1, c_2 , let $D_P(c_1, c_2)$ denote the number of times that c_1 is ranked higher than c_2 in P minus the number of times that c_2 is ranked higher than c_1 in P , that is,

$$D_P(c_1, c_2) = |\{R \in P : c_1 \succ_R c_2\}| - |\{R \in P : c_2 \succ_R c_1\}|.$$

We shall require the following previously known lemma.

Lemma 3.2. [McGarvey, 1953] *Given a function $F : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{Z}$ such that*

1. *for all $c_1, c_2 \in \mathcal{C}$, $c_1 \neq c_2$, $F(c_1, c_2) = -F(c_2, c_1)$, and*
2. *either for all pairs of candidates $c_1, c_2 \in \mathcal{C}$ (with $c_1 \neq c_2$), $F(c_1, c_2)$ is even, or for all pairs of candidates $c_1, c_2 \in \mathcal{C}$ (with $c_1 \neq c_2$), $F(c_1, c_2)$ is odd,*

there exists a profile P such that for all $c_1, c_2 \in \mathcal{C}$, $c_1 \neq c_2$, $D_P(c_1, c_2) = F(c_1, c_2)$ and

$$|P| \leq \frac{1}{2} \sum_{c_1, c_2: c_1 \neq c_2} |F(c_1, c_2) - F(c_2, c_1)|.$$

Theorem 3.3. *The UCM_k problem under maximin is NP-complete for any number of manipulators $k \geq 2$.*

Proof of Theorem 3.3: It is easy to verify that the UCM problem under maximin is in NP. We now show that UCM is NP-hard, by giving a reduction from the two vertex disjoint paths in directed graph problem.

Let the instance of the two vertex disjoint paths in directed graph problem be denoted by $G = (V, E)$, (u, u') and (v, v') where $V = \{u, u', v, v', c_1, \dots, c_{m-5}\}$. Without loss of generality, we assume that every vertex is reachable from u or v (otherwise, we can remove the vertex from the instance). We also assume that $(u, v') \notin E$ and $(v, u') \notin E$ (since such edges cannot be used in a solution). Let $G' = (V, E \cup \{(v', u), (u', v)\})$, that is, G' is the graph obtained from G by adding (v', u) and (u', v) .

We construct a UCM instance as follows.

Set of candidates: $\mathcal{C} = \{c, u, u', v, v', c_1, \dots, c_{m-5}\}$.

Candidate preferred by the manipulators: c .

Number of unweighted manipulators: $|M|$ (for some $|M| \geq 2$).

Non-manipulators' profile: P^{NM} satisfying the following conditions:

1. For any $c' \neq c$, $D_{P^{NM}}(c, c') = -4|M|$.
2. $D_{P^{NM}}(u, v') = D_{P^{NM}}(v, u') = -4|M|$.

3. For any $(s, t) \in E$ such that $D_{P^{NM}}(t, s)$ is not defined above, we let $D_{P^{NM}}(t, s) = -2|M| - 2$.
4. For any $s, t \in \mathcal{C}$ such that $D_{P^{NM}}(t, s)$ is not defined above, we let $|D_{P^{NM}}(t, s)| = 0$.

The existence of such a P^{NM} , whose size is polynomial in m , is guaranteed by Lemma 3.2.

We can assume without loss of generality that each manipulator ranks c first. Therefore, for any $c' \neq c$,

$$D_{P^{NM} \cup P^M}(c, c') = -3|M|. \quad (1)$$

We are now ready to show that there exists P^M such that $Maximin(P^{NM} \cup P^M) = \{c\}$ if and only if there exist two vertex disjoint paths from u to u' and from v to v' in G . First, we prove that if there exist such paths in G , then there exists a profile P^M for the manipulators such that $Maximin(P^{NM} \cup P^M) = \{c\}$.

Let $u \rightarrow u_1 \rightarrow \dots \rightarrow u_{k_1} \rightarrow u'$ and $v \rightarrow v_1 \rightarrow \dots \rightarrow v_{k_2} \rightarrow v'$ be two vertex disjoint paths. Further, let

$$V' = \{u, u', v, v', u_1, \dots, u_{k_1}, v_1, \dots, v_{k_2}\}.$$

Then, because any vertex is reachable from u or v in G , there exists a connected subgraph G^* of G' (which still includes all the vertices) in which $u \rightarrow u_1 \rightarrow \dots \rightarrow u_{k_1} \rightarrow u' \rightarrow v \rightarrow v_1 \rightarrow \dots \rightarrow v_{k_2} \rightarrow v' \rightarrow u$ is the only cycle. In other words, such a subgraph G^* can be obtained by possibly removing some of the edges of G' . Therefore, by arranging the vertices of $V \setminus V'$ according to the direction of the edges of G^* , we can obtain a linear order O over $V \setminus V'$ with the following property: for any $t \in V \setminus V'$, it holds that either

1. there exists $s \in V \setminus V'$ such that $s \succ_O t$ and $(s, t) \in E$, or
2. there exists $s \in V'$ such that $(s, t) \in E$.

We define P^M by letting $|M| - 1$ manipulators vote

$$c \succ u \succ u_1 \succ \dots \succ u_{k_1} \succ u' \succ v \succ v_1 \succ \dots \succ v_{k_2} \succ v' \succ O$$

and letting one manipulator vote

$$c \succ v \succ v_1 \succ \dots \succ v_{k_2} \succ v' \succ u \succ u_1 \succ \dots \succ u_{k_1} \succ u' \succ O.$$

Then, we have the following calculations:

$$\begin{aligned} D_{P^{NM} \cup P^M}(u, v') &= -4|M| + (|M| - 1) - 1 \\ &= -3|M| - 2 < -3|M|, \end{aligned}$$

$$\begin{aligned} \text{and } D_{P^{NM} \cup P^M}(v, u') &= -4|M| + 1 - (|M| - 1) \\ &= -5|M| + 2 < -3|M|. \end{aligned}$$

Moreover, for any $t \in \mathcal{C} \setminus \{c, u, v\}$, there exists $s \in \mathcal{C} \setminus \{c\}$ such that $(s, t) \in E$ and $D_{P^M}(t, s) = -|M|$, which means that

$$\begin{aligned} D_{P^{NM} \cup P^M}(t, s) &= -2|M| - 2 - |M| = -3|M| - 2 \\ &< -3|M|. \end{aligned}$$

It now follows from Equation (1) that

$$Maximin(P^{NM} \cup P^M) = \{c\}.$$

Next, we prove that if there exists a profile P^M for the manipulators such that $\text{Maximin}(P^{NM} \cup P^M) = \{c\}$, then there exist two vertex disjoint paths from u to u' and from v to v' .

We define a function $f : V \rightarrow V$ such that

$$D_{P^{NM} \cup P^M}(t, f(t)) < -3|M| .$$

Indeed, such a function exists since

$$\text{Maximin}(P^{NM} \cup P^M) = \{c\} ,$$

hence for any $t \neq c$ there must exist s such that

$$D_{P^{NM} \cup P^M}(t, s) < -3|M| .$$

Moreover, s must be a parent of t in G' . If there exists more than one such s , define $f(t)$ to be any one of them.

It follows that if $(t, f(t))$ is neither (u, v') or (v, u') , then $(f(t), t) \in E$ and $D_{P^M}(t, f(t)) = -|M|$, which means that $f(t) \succ t$ in each vote of P^M ; otherwise, if $(t, f(t))$ is (u, v') or (v, u') , then $D_{P^M}(t, f(t)) \leq |M| - 2$, which means that $f(t) \succ t$ in at least one vote of P^M .

Now, since $|V| = m$ is finite, there must exist $l_1 < l_2 \leq m$ such that $f^{l_1}(u) = f^{l_2}(u)$. That is,

$$f^{l_1}(u) \rightarrow f^{l_1+1}(u) \rightarrow \dots \rightarrow f^{l_2-1}(u) \rightarrow f^{l_2}(u)$$

is a cycle in G' . We assume that for any $l_1 \leq l'_1 < l'_2 < l_2$, $f^{l'_1}(u) \neq f^{l'_2}(u)$. Now we claim that (v', u) and (u', v) must be both in the cycle, because

1. if neither of them is in the cycle, then in each vote of P^M , we must have

$$f^{l_2}(u) \succ f^{l_2-1}(u) \succ f^{l_1}(u) = f^{l_2}(u) ,$$

which contradicts the assumption that each vote is a linear order;

2. if exactly one of them is in the cycle—without loss of generality, $f^{l_1}(u) = v$, $f^{l_1+1}(u) = u'$ —then in at least one of the votes of P^M , we must have

$$f^{l_2}(u) \succ f^{l_2-1}(u) \succ \dots \succ f^{l_1}(u) = f^{l_2}(u) ,$$

which contradicts the assumption that each vote is a linear order.

Without loss of generality, let us assume that $f^{l_1}(u) = u$, $f^{l_1+1}(u) = v'$, $f^{l_3}(u) = v$, $f^{l_3+1}(u) = u'$, where $l_3 \leq l_2 - 2$. We immediately obtain two vertex disjoint paths:

$$u = f^{l_1}(u) = f^{l_2}(u) \rightarrow f^{l_2-1}(u) \rightarrow \dots \rightarrow f^{l_3+1}(u) = u' ,$$

$$\text{and } v = f^{l_3}(u) \rightarrow f^{l_3-1}(u) \rightarrow \dots \rightarrow f^{l_1+1}(u) = v' .$$

Therefore, UCM under maximin is NP-complete. \square

Notice that the NP-completeness of UCM implies the NP-hardness of UCO under maximin. Zuckerman et al. [2008] have designed a 2-approximation algorithm for UCO under maximin, even though it was unclear at the time that the problem was indeed NP-hard. It still remains open whether the approximation ratio can be improved, or whether a hardness-of-approximation result precludes this.

4 Ranked pairs

We now turn to investigating the ranked pairs voting rule. This interesting voting rule satisfies some important social choice desiderata [Tideman, 1987]. Moreover, we assert below that ranked pairs is hard to manipulate even by a single manipulator, making it one of very few “natural” voting rules with this property. The hardness easily extends to multiple manipulators as well.

Theorem 4.1. *The UCM_k problem under ranked pairs is NP-complete for any number of manipulators.*

Crucially, we can prove an even stronger result regarding ranked pairs: the UCO problem under this rule is extremely hard to approximate. More accurately, we have the following theorem.

Theorem 4.2. *Let N be the size of the non-manipulators' profile in the UCO problem. For every constant $\epsilon > 0$, it is NP-hard to approximate UCO under ranked pairs within a factor of $N^{1-\epsilon}$.*

In particular, we show that it is NP-hard to distinguish between the following two extreme cases: there is a successful manipulation via a single manipulator, or any successful manipulation requires more than $N^{1-\epsilon}$ manipulators. It can be argued that this result makes ranked pairs especially appealing in terms of its resistance to manipulation. In fact, this is the strongest hardness of manipulation result currently known for any voting rule.

The proofs of Theorems 4.1 and 4.2 are the most involved in this paper (considerably more so than the proof of Theorem 3.3). The proof of Theorem 4.1 requires an elaborate reduction from 3SAT. In order to obtain Theorem 4.2, we extend the proof of Theorem 4.1 using an interesting property of ranked pairs: if there is a successful manipulation for a set of manipulators M , there is a successful manipulation where all the manipulators vote identically. This allows us to design instances where manipulations using one manipulator and many manipulators are equivalent. Unfortunately, we must omit the details of the proofs due to lack of space.

5 Bucklin

In this section, we present a polynomial-time algorithm for the UCM problem under Bucklin.

For any candidate $x \in \mathcal{C}$, any natural number $d \in \mathbb{N}$, and any profile P , let $B(x, d, P)$ denote the number of times that x is ranked among the top d candidates in P . The idea behind the algorithm is as follows. Let d_{min} be the minimal depth so that the favorite candidate c is ranked among the top d_{min} candidates in more than half of the votes (when all of the manipulators rank c first). Then, we simply check if there is a way to assign the manipulators' votes so that none of the other candidates is ranked among the top d_{min} candidates in more than half of the votes. In other words, the order of the candidates is not crucial, only their membership in the set of d_{min} top-ranked candidates is relevant.

Algorithm 1.

Input. A UCM instance $(\text{Bucklin}, P^{NM}, c, M)$,

$$\mathcal{C} = \{c, c_1, \dots, c_{m-1}\} .$$

Stage 0.

0.1 Calculate the minimal depth d_{min} such that

$$B(c, d_{min}, P^{NM}) + |M| > \frac{1}{2}(|NM| + |M|) .$$

0.2 If there exists $c' \in C$, $c' \neq c$ such that

$$B(c', d_{min}, P^{NM}) > \frac{1}{2}(|NM| + |M|) , \quad (2)$$

then output that there is no successful manipulation.

Aside. Notice that d_{min} is defined under the assumption that all the manipulators rank c first. Consider a candidate $c' \neq c$ that satisfies the condition in Equation (2). Such a candidate is ranked in the top d_{min} positions of half the votes $P^{NM} \cup P^M$, regardless of P^M . Hence, c cannot be a unique winner.

Stage 1.

1.1 For every $c' \in C \setminus \{c\}$, let

$$d_{c'} = \left\lfloor \frac{1}{2}(|NM| + |M|) \right\rfloor - B(c', d_{min}, P^{NM}) ,$$

and let $k_{c'} = \min\{d_{c'}, |M|\}$.

1.2 If

$$\sum_{c' \neq c} k_{c'} < (d_{min} - 1)|M| , \quad (3)$$

then output that there is no successful manipulation.

Aside. $k_{c'}$ is the number of times that we can place c' in the first d_{min} positions of the votes of P^M , without compromising the victory of c . In particular, $k_{c'}$ cannot be greater than $|M|$.

Notice that there are exactly $(d_{min} - 1)|M|$ problematic positions to fill, since c is ranked first by all the manipulators. Now, if the condition in Equation (3) is satisfied, for any P^M there must be a candidate c' that appears too many times in the first d_{min} positions, that is,

$$k_{c'} < B(c', d_{min}, P^M) .$$

Since $B(c', d_{min}, P^M) \leq |M|$, we have in particular that $k_{c'} < |M|$, hence it must hold that $k_{c'} = d_{c'}$. It follows that

$$\begin{aligned} & B(c', d_{min}, P^{NM} \cup P^M) \\ &= B(c', d_{min}, P^{NM}) + B(c', d_{min}, P^M) \\ &> B(c', d_{min}, P^{NM}) + d_{c'} \\ &= \left\lfloor \frac{1}{2}(|NM| + |M|) \right\rfloor . \end{aligned}$$

Therefore, c cannot be a unique winner.

Stage 2. Construct P^M by assigning the candidates to the first d_{min} positions of the votes in a way that for every $t = 1, \dots, m - 1$,

$$B(c_t, d_{min}, P^M) \leq k_{c_t} . \quad (4)$$

Complete the rest of the votes arbitrarily. Return P^M as a successful manipulation.

Aside. Given that (3) does not hold, it is clearly possible to construct P^M such that (4) holds for every $c' \neq c$. Moreover,

this can be done in polynomial time, e.g., by enumerating the candidates and placing each candidate in the next position in $k_{c'}$ of the votes of the manipulators, until the crucial positions are filled.

Now, for every $t = 1, \dots, m - 1$ it holds that

$$\begin{aligned} B(c_t, d_{min}, P^{NM} \cup P^M) &\leq B(c_t, d_{min}, P^{NM}) + k_{c_t} \\ &\leq \frac{1}{2}(|NM| + |M|) , \end{aligned}$$

which implies that $Bucklin(P^{NM} \cup P^M) = \{c\}$.

We have obtained the following result.

Theorem 5.1. *Algorithm 1 correctly decides the UCM problem in polynomial time.*

It is easy to see that the tractability of UCM under Bucklin implies that UCO can be solved in polynomial time as well.

6 Discussion

We have studied the computational complexity of the UCM and UCO problems under the maximin, ranked pairs, and Bucklin rules. The UCM problem is NP-complete under the maximin rule for any fixed number (at least two) of manipulators. The UCM problem is also NP-complete under the ranked pairs rule; in this case, the hardness holds even if there is only a single manipulator, similarly to the second-order Copeland [Bartholdi *et al.*, 1989] and STV [Bartholdi and Orlin, 1991] rules. Furthermore, we have shown that UCO under ranked pairs is NP-hard to approximate to a factor of $N^{1-\epsilon}$, where N is the size of the input and ϵ is an arbitrary constant. Finally, we have given a polynomial-time algorithm for the UCM problem under the Bucklin rule. Table 1 summarizes our results, and puts them in the context of previous results on the UCM problem.

These results may seem to be at odds with the results of Conitzer *et al.* [2007] about weighted coalitional manipulation (WCM). In particular, they show that WCM under maximin is in P if the number of candidates is two or three. Even though UCM is easier than WCM, there is no conflict since our results rely on the number of candidates being a parameter. In fact, if the number of candidates is constant, UCM is tractable under any voting rule (that can be executed in polynomial time), via simple enumeration [Conitzer *et al.*, 2007].

It should be noted that all of our hardness results, as well as the ones mentioned in the introduction, are *worst-case* results. Hence, there may still be an efficient algorithm that can find a beneficial manipulation for *most* instances. Indeed, nearly a dozen recent papers suggest that finding manipulations is easy with respect to some typical distributions on preference profiles (see, e.g., [Procaccia and Rosenschein, 2007b; 2007a; Conitzer and Sandholm, 2006; Friedgut *et al.*, 2008; Dobzinski and Procaccia, 2008; Xia and Conitzer, 2008a; 2008b] and the references therein). However, the fascinating question of the frequency of manipulation in elections is far from being resolved, and it is even not completely clear how this question should be approached. Hence, worst-case hardness still remains the primary tool in the study of complexity of manipulation.

Number of manipulators	One		At least two	
Copeland (specific tie-breaking)	P	[Bartholdi <i>et al.</i> , 1989]	NP-c	[Faliszewski <i>et al.</i> , 2008]
STV	NP-c	[Bartholdi and Orlin, 1991]	NP-c	[Bartholdi and Orlin, 1991]
Veto	P	[Bartholdi <i>et al.</i> , 1989]	P	[Zuckerman <i>et al.</i> , 2008]
Plurality with Runoff	P	[Zuckerman <i>et al.</i> , 2008]	P	[Zuckerman <i>et al.</i> , 2008]
Cup	P	[Conitzer <i>et al.</i> , 2007]	P	[Conitzer <i>et al.</i> , 2007]
Maximin	P	[Bartholdi <i>et al.</i> , 1989]	NP-c	(Thm 3.3)
Ranked pairs	NP-c	(Thm 4.1)	NP-c	(Thm 4.1)
Bucklin	P	(Thm 5.1)	P	(Thm 5.1)
Borda	P	[Bartholdi <i>et al.</i> , 1989]	?	

Table 1: Complexity of UCM under prominent voting rules. Boldface results appear in this paper.

There are many interesting problems left for future research. For example, settling the complexity of UCM under positional scoring rules such as Borda is a challenging problem that remains open despite several attacks. An especially intriguing open problem concerns the hardness of approximating UCO. We have seen that UCO under ranked pairs is extremely hard to approximate. We conjecture that hardness of approximation results can also be obtained with respect to other voting rules where manipulation is hard even for a single manipulator (e.g., STV [Bartholdi and Orlin, 1991]), although the inapproximability ratio may not be as extreme as it is for ranked pairs.

Acknowledgments

We thank anonymous COMSOC and IJCAI reviewers for helpful comments and suggestions. This work is supported in part by the United States-Israel Binational Science Foundation under grant 2006-216. Lirong Xia is supported by a James B. Duke Fellowship, Vincent Conitzer is supported by an Alfred P. Sloan Research Fellowship, and Ariel Procaccia was supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities. Xia and Conitzer are also supported by the NSF under award number IIS-0812113.

References

- [Bartholdi and Orlin, 1991] J. Bartholdi and J. Orlin. Single Transferable Vote resists strategic voting. *Social Choice and Welfare*, 8:341–354, 1991.
- [Bartholdi *et al.*, 1989] J. Bartholdi, C. A. Tovey, and M. A. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6:227–241, 1989.
- [Conitzer and Sandholm, 2003] V. Conitzer and T. Sandholm. Universal voting protocol tweaks to make manipulation hard. In *Proc. of 18th IJCAI*, pages 781–788, 2003.
- [Conitzer and Sandholm, 2006] V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In *Proc. of 21st AAAI*, pages 627–634, 2006.
- [Conitzer *et al.*, 2007] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):1–33, 2007.
- [Dobzinski and Procaccia, 2008] S. Dobzinski and A. D. Procaccia. Frequent manipulability of elections: The case of two voters. In *Proc. of 4th WINE*, 2008. To appear.
- [Elkind and Lipmaa, 2005] E. Elkind and H. Lipmaa. Hybrid voting protocols and hardness of manipulation. In *Algorithms and Computation*, volume 3827 of *LNCS*, pages 206–215. Springer-Verlag, 2005.
- [Faliszewski *et al.*, 2008] P. Faliszewski, E. Hemaspaandra, and H. Schnoor. Copeland voting: Ties matter. In *Proc. of 7th AAMAS*, pages 983–990, 2008.
- [Fortune *et al.*, 1980] S. Fortune, J. E. Hopcroft, and J. Wyllie. The directed subgraph homeomorphism problem. *Theoretical Computer Science*, 10:111–121, 1980.
- [Friedgut *et al.*, 2008] E. Friedgut, G. Kalai, and N. Nisan. Elections can be manipulated often. In *Proc. of 49th FOCS*, pages 243–249, 2008.
- [Gibbard, 1973] A. Gibbard. Manipulation of voting schemes. *Econometrica*, 41:587–602, 1973.
- [Hemaspaandra and Hemaspaandra, 2007] E. Hemaspaandra and L. A. Hemaspaandra. Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1):73–83, 2007.
- [McGarvey, 1953] D. C. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21:608–610, 1953.
- [Procaccia and Rosenschein, 2007a] A. D. Procaccia and J. S. Rosenschein. Average-case tractability of manipulation in elections via the fraction of manipulators. In *Proc. of 6th AAMAS*, pages 718–720, 2007.
- [Procaccia and Rosenschein, 2007b] A. D. Procaccia and J. S. Rosenschein. Junta distributions and the average-case complexity of manipulating elections. *Journal of Artificial Intelligence Research*, 28:157–181, 2007.
- [Satterthwaite, 1975] M. Satterthwaite. Strategy-proofness and Arrow’s conditions: Existence and correspondence theorems for voting procedures and social welfare functions. *Journal of Economic Theory*, 10:187–217, 1975.
- [Tideman, 1987] N. Tideman. Independence of clones as a criterion for voting rules. *Social Choice and Welfare*, 4(3):185–206, 1987.
- [Tideman, 2006] N. Tideman. *Collective Decisions and Voting*. Ashgate, 2006.
- [Xia and Conitzer, 2008a] L. Xia and V. Conitzer. Generalized Scoring Rules and the frequency of coalitional manipulability. In *Proc. of 9th ACM-EC*, pages 109–118, 2008.
- [Xia and Conitzer, 2008b] L. Xia and V. Conitzer. A sufficient condition for voting rules to be frequently manipulable. In *Proc. of 9th ACM-EC*, pages 99–108, 2008.
- [Zuckerman *et al.*, 2008] M. Zuckerman, A. D. Procaccia, and J. S. Rosenschein. Algorithms for the coalitional manipulation problem. In *Proc. of 19th SODA*, pages 277–286, 2008.