The Complexity of Computing Minimal Unidirectional Covering Sets

Dorothea Baumeister[1], Felix Brandt[2], Felix Fischer[3], Jan Hoffmann[2], and Jörg Rothe[1]

[1] Institut für Informatik, Heinrich-Heine-Universität Düsseldorf, 40225 Düsseldorf, Germany
[2] Institut für Informatik, Ludwig-Maximilians-Universität München, 80538 München, Germany
[3] Harvard School of Engineering and Applied Sciences, Cambridge, MA 02138, USA

Abstract. Given a binary dominance relation on a set of alternatives, a common thread in the social sciences is to identify subsets of alternatives that satisfy certain notions of stability. Examples can be found in areas as diverse as voting theory, game theory, and argumentation theory. Brandt and Fischer [4] proved that it is NP-hard to decide whether an alternative is contained in some inclusion-minimal unidirectional (i.e., either upward or downward) covering set. For both problems, we raise this lower bound to the Θ_2^p level of the polynomial hierarchy and provide a Σ_2^p upper bound. Relatedly, we show that a variety of other natural problems regarding minimal or minimum-size unidirectional covering sets are hard or complete for either of NP, coNP, and Θ_2^p. An important consequence of our results is that neither minimal upward nor minimal downward covering sets (even when guaranteed to exist) can be computed in polynomial time unless P = NP. This sharply contrasts with Brandt and Fischer's result that minimal bidirectional covering sets are polynomial-time computable.

1 Introduction

A common thread in the social sciences is to identify sets of alternatives that satisfy certain notions of stability according to some binary dominance relation. Applications range from cooperative to noncooperative game theory, from social choice theory to argumentation theory, and from multi-criteria decision analysis to sports tournaments (see, e.g., [15, 4] and the references therein).

In social choice settings, the most common dominance relation is the pairwise majority relation, where an alternative x is said to dominate another alternative y if the number of individuals preferring x to y exceeds the number of individuals preferring y to x. McGarvey [16] proved that *every* asymmetric dominance relation can be realized via a particular preference profile, even if the individual preferences are linear. For example, Condorcet's well-known paradox says that the majority relation may contain cycles and thus does not always have maximal elements, even if all of the underlying individual preferences do. This means that the concept of maximality is rendered useless in many cases, which is why various so-called *solution concepts* have been proposed. Solution concepts can be used in place of maximality for nontransitive relations (see, e.g., [15]). In particular, concepts based on so-called *covering relations*—transitive subrelations of the dominance relation at hand—have turned out to be very attractive [11, 17, 9].

Computational social choice, an emerging new field at the interface of social choice theory, economics, and computer science, focuses on the computational properties of social-choice-related concepts and problems [7]. In this paper, we study the computational complexity of problems related to the notions of upward and downward covering sets in dominance graphs. An alternative x is said to *upward cover* another alternative y if x dominates y and every alternative dominating x also dominates y. The intuition is that x "strongly" dominates y in the sense that there is no alternative that dominates x but not y. Similarly, an alternative x is said to *downward cover* another alternative y if x dominates y and every alternative dominated by y is also dominated by x. The intuition here is that x "strongly" dominates y in the sense that there is no alternative dominated by y but not by x. A *minimal upward* or *minimal downward covering set* is defined as an inclusion-minimal set of alternatives that satisfies certain notions of internal and external stability with respect to the upward or downward covering relation [9, 4].

Recent work in computational social choice has addressed the computational complexity of most solution concepts proposed in the context of binary dominance (see, e.g., [21, 1, 8, 5, 4, 6]). Brandt and Fischer [4] have shown NP-hardness of deciding whether an alternative is contained in some minimal upward (respectively, downward) covering set. For both problems, we raise their NP-hardness lower bounds to the Θ_2^p level of the polynomial hierarchy, and we provide an upper bound of Σ_2^p. We also analyze the complexity of a variety of other problems related to minimal and minimum-size upward and downward covering sets that have not been studied before. In particular, we provide hardness and completeness results for the complexity classes NP, coNP, and Θ_2^p. Remarkably, these new results imply that neither minimal upward covering sets nor minimal downward covering sets (even when guaranteed to exist) can be found in polynomial time unless P = NP. This sharply contrasts with Brandt and Fischer's result that minimal *bidirectional* covering sets are polynomial-time computable [4]. Note that, notwithstanding the hardness of computing minimal upward covering sets, the decision version of this search problem is trivially in P: Every dominance graph always contains a minimal upward covering set.

Our Θ_2^p-hardness results apply Wagner's method [20]. To the best of our knowledge, our constructions for the first time apply his method to problems defined in terms of minimality rather than minimum size of a solution.

2 Definitions and Notation

We now define the necessary concepts from social choice and complexity theory.

Let A be a finite set of alternatives, let $B \subseteq A$, and let $\succ \subseteq A \times A$ be a dominance relation on A, i.e., $\succ$ is asymmetric and irreflexive (in general, $\succ$ need not be transitive or complete).[4] A dominance relation $\succ$ on a set A of alternatives can be conveniently represented as a *dominance graph*, denoted by $(A, \succ)$, whose vertices are the alternatives from A, and for each $x, y \in A$ there is a directed edge from x to y if and only if $x \succ y$. For any two alternatives x and y in B, define the following covering relations (see, e.g., [11, 17, 3]): *x upward covers y in B*, denoted by $x\, C_u^B\, y$, if $x \succ y$ and for all $z \in B$, $z \succ x$ implies $z \succ y$, and *x downward covers y in B*, denoted by $x\, C_d^B\, y$, if $x \succ y$ and for all $z \in B$, $y \succ z$ implies $x \succ z$. When clear from the context, we omit mentioning "in B" explicitly.

Definition 1. *Let A be a set of alternatives, let $B \subseteq A$ be any subset, and let $\succ$ be a dominance relation on A. The* upward uncovered set *of B is defined as* $\mathrm{UC}_u(B) = \{x \in B \mid y\, C_u^B\, x \text{ for no } y \in B\}$*, whereas the* downward uncovered set *of B is defined as* $\mathrm{UC}_d(B) = \{x \in B \mid y\, C_d^B\, x \text{ for no } y \in B\}$.

In the dominance graph $(A, \succ)$ in Figure 1, b upward covers c in A, and a downward covers b in A (i.e., $b\, C_u^A\, c$ and $a\, C_d^A\, b$), so $UC_u(A) = \{a, b, d\}$ is the upward uncovered set and $UC_d(A) = \{a, c, d\}$ is the downward uncovered set of A. For both the upward and the downward covering relation (henceforth unidirectional covering relations), transitivity of the relation implies nonemptiness of the corresponding uncovered set for each nonempty set of alternatives. The intuition underlying covering sets is that there should be no reason to restrict the selection by excluding some alternative from it (internal stability) and there should be an argument against each proposal to include an outside alternative into the selection (external stability).

Definition 2. *Let A be a set of alternatives and $\succ$ be a dominance relation on A. A subset $B \subseteq A$ is an* upward covering set *of A if* $\mathrm{UC}_u(B) = B$ *(internal stability) and for all $x \in A - B$, $x \notin \mathrm{UC}_u(B \cup \{x\})$ (external stability).* Downward covering sets *are defined analogously using* UC_d*. An upward (respectively, a downward) covering set M is said to be* (inclusion-)minimal *if no $M' \subset M$ is an upward (respectively, a downward) covering set for A.*

Every upward uncovered set contains one or more minimal upward covering sets, whereas minimal downward covering sets do not always exist [4]. Dutta [9] proposed minimal covering sets in the context of tournaments, i.e., complete dominance relations, where both notions of covering coincide. Minimal unidirectional covering sets are one of several possible generalizations to incomplete dominance relations (see [4]). Occasionally, it might be helpful to specify the dominance relation explicitly to avoid ambiguity. In such cases we refer to the dominance graph and write, e.g., "M is an upward covering set for $(A, \succ)$." The unique minimal upward covering set for the dominance graph shown in Figure 1 is $\{b, d\}$, and the unique minimal downward covering set is $\{a, c, d\}$.

[4] For alternatives x and y, $x \succ y$ (alternatively, $(x, y) \in \succ$) is interpreted as x being strictly preferred to y (we say "x dominates y"), e.g., due to a strict majority of voters preferring x to y.

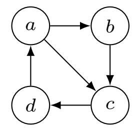

Fig. 1. Dominance graph $(A, \succ)$ as an example for upward and downward covering relations.

One computational problem of central interest in this paper is Minimal Upward Covering Set Member ($\mathrm{MC_u}$-Member, for short): Given a set of alternatives A, a dominance relation $\succ$ on A, and a distinguished element $d \in A$, is d contained in some minimal upward covering set for A? Another important problem is the search problem $\mathrm{MC_u}$-Find: Given a set of alternatives A and a dominance relation $\succ$ on A, find a minimal upward covering set for A. The problems $\mathrm{MC_d}$-Member and $\mathrm{MC_d}$-Find are defined analogously for minimal downward covering sets.

We assume that the reader is familiar with the basic notions of complexity theory, such as polynomial-time many-one reducibility and the related notions of hardness and completeness, and also with standard complexity classes such as P, NP, coNP, and the polynomial hierarchy (see, e.g., [18]). In particular, coNP is the class of sets whose complements are in NP. $\Sigma_2^p = \mathrm{NP}^{\mathrm{NP}}$, the second level of the polynomial hierarchy, consists of all sets that can be solved by an NP oracle machine that has access (in the sense of a Turing reduction) to an NP oracle set such as SAT. SAT denotes the satisfiability problem of propositional logic, which is one of the standard NP-complete problems (see, e.g., Garey and Johnson [12]) and is defined as follows: Given a boolean formula in conjunctive normal form, does there exist a truth assignment to its variables that satisfies the formula? Papadimitriou and Zachos [19] introduced the class of problems solvable in polynomial time via asking $\mathcal{O}(\log n)$ sequential Turing queries to NP. This class is also known as the Θ_2^p level of the polynomial hierarchy, and has been shown to coincide with the class of problems that can be decided by a P machine that accesses its NP oracle in a parallel manner. Equivalently, Θ_2^p is the closure of NP under polynomial-time truth-table reductions. It follows immediately from the definitions that $\mathrm{P} \subseteq \mathrm{NP} \cap \mathrm{coNP} \subseteq \mathrm{NP} \cup \mathrm{coNP} \subseteq \Theta_2^p \subseteq \Sigma_2^p$. The class Θ_2^p captures the complexity of various optimization problems [20]. In the field of computational social choice, the winner problems for Dodgson, Young, and Kemeny elections have been shown to be Θ_2^p-complete (see [10] and the references cited therein).

3 Results and Discussion

Results. Brandt and Fischer [4] proved that it is NP-hard to decide whether a given alternative is contained in some minimal unidirectional covering set. Using the notation of this paper, their results state that the problems $\mathrm{MC_u}$-Member and $\mathrm{MC_d}$-Member are NP-hard. The question of whether these problems are

NP-complete or of higher complexity was left open. Our contribution is *(i)* to raise Brandt and Fischer's NP-hardness lower bounds for MC_u-MEMBER and MC_d-MEMBER to Θ_2^p-hardness and to provide (simple) Σ_2^p upper bounds for these problems, and *(ii)* to extend the techniques we developed to also apply to various other covering set problems that will be defined in Section 6 and in particular to the search problems. Due to space constraints we focus here on our results for MC_u-MEMBER and MC_u-FIND but we mention that we obtained many more results on upward and downward covering set problems (see Theorem 9 and Table 1 in Section 6) the proofs of which are provided in the full version of this paper [2].

Discussion. We consider the problem of *finding* a minimal unidirectional covering set (MC_u-FIND and MC_d-FIND) to be particularly important and natural.

Regarding upward covering sets, we stress that our result that, assuming $P \neq NP$, MC_u-FIND is hard to compute (Theorem 8) does not follow directly from the NP-hardness of MC_u-MEMBER in an obvious way.[5] Our reduction that raises the lower bound of MC_u-MEMBER from NP-hardness to Θ_2^p-hardness, however, allows us to prove that MC_u-FIND is not polynomial-time solvable unless $P = NP$.

Regarding downward covering sets, the hardness of MC_d-FIND (assuming $P \neq NP$) is an immediate consequence of Brandt and Fischer's result that it is NP-complete to decide whether there exists a minimal downward covering set [4, Thm. 9]. We provide an alternative proof based on our reduction showing that MC_d-MEMBER is Θ_2^p-hard [2, Thm. 5.13]. In contrast to Brandt and Fischer's proof, our proof shows the hardness of MC_d-FIND even when the existence of a (minimal) downward covering set is guaranteed.

As mentioned above, the problem MC_u-MEMBER was already known to be NP-hard [4] and is here shown to be even Θ_2^p-hard. One may naturally wonder whether raising its (or any problem's) lower bound from NP-hardness to Θ_2^p-hardness gives us any more insight into the problem's inherent computational complexity. After all, $P = NP$ if and only if $P = \Theta_2^p$. However, this question is a bit more subtle than that and has been discussed carefully by Hemaspaandra et al. [14]. They make the case that the answer to this question crucially depends on what one considers to be the most natural computational model. In particular, they argue that raising NP-hardness to Θ_2^p-hardness potentially (i.e., unless longstanding open problems regarding the separation of the corresponding complexity classes could be solved) is an improvement in terms of randomized polynomial time and in terms of unambiguous polynomial time [14].

---

[5] The decision version of MC_u-FIND is: Given a dominance graph, does it contain a minimal upward covering set? However, this question has always an affirmative answer, so this problem is trivially in P. Note also that MC_u-FIND is no harder (with respect to "polynomial-time disjunctive truth-table" reductions) than the search version of MC_u-MEMBER. The converse, however, is not at all obvious. Brandt and Fischer's results only imply the hardness of finding an alternative that is contained in *all* minimal upward covering sets [4].

4 Upward Covering Constructions and Their Key Properties

In this section, we provide the constructions and their key properties to be used in Sections 5 and 6 to prove lower bounds for problems such as $\mathrm{MC_u}$-Member.

Construction 1 (to be used for showing coNP-hardness of $\mathrm{MC_u}$-Member). *Given a boolean formula in conjunctive normal form, $\varphi(w_1, w_2, \ldots, w_k) = f_1 \wedge f_2 \wedge \cdots \wedge f_\ell$, over the set $W = \{w_1, w_2, \ldots, w_k\}$ of variables, we construct a set of alternatives $A = \{u_i, \overline{u}_i, u'_i, \overline{u}'_i \mid w_i \in W\} \cup \{e_j, e'_j \mid f_j \text{ is a clause in } \varphi\} \cup \{a_1, a_2, a_3\}$, and a dominance relation $\succ$ on A that is defined by:* (i) *for each i, $1 \leq i \leq k$, there is a cycle $u_i \succ \overline{u}_i \succ u'_i \succ \overline{u}'_i \succ u_i$;* (ii) *if variable w_i occurs in clause f_j as a positive literal, then $u_i \succ e_j$, $u_i \succ e'_j$, $e_j \succ \overline{u}_i$, and $e'_j \succ \overline{u}_i$;* (iii) *if variable w_i occurs in clause f_j as a negative literal, then $\overline{u}_i \succ e_j$, $\overline{u}_i \succ e'_j$, $e_j \succ u_i$, and $e'_j \succ u_i$;* (iv) *if variable w_i does not occur in clause f_j, then $e_j \succ u'_i$ and $e'_j \succ \overline{u}'_i$;* (v) *for each j, $1 \leq j \leq \ell$, we have $a_1 \succ e_j$ and $a_1 \succ e'_j$; and* (vi) *there is a cycle $a_1 \succ a_2 \succ a_3 \succ a_1$.*

Figures 2(a)–2(c) show some parts of the dominance graph that results from the given formula φ. As a more complete example, Figure 2(d) shows the entire dominance graph that corresponds to the concrete formula $(\neg w_1 \vee w_2) \wedge (w_1 \vee \neg w_3)$, which can be satisfied by setting, for example, each of w_1, w_2, and w_3 to true. A minimal upward covering set for A corresponding to this assignment is $M = \{u_1, u'_1, u_2, u'_2, u_3, u'_3, a_1, a_2, a_3\}$. Note that neither e_1 nor e_2 occurs in M, and none of them occurs in any other minimal upward covering set for A either. For alternative e_1 in the example shown in Figure 2(d) this can be seen as follows. If there were a minimal upward covering set M' for A containing e_1 (and thus also e'_1, since they both are dominated by the same alternatives) then neither $\overline{u}_1$ nor u_2 (which dominate e_1) must upward cover e_1 in M', so all alternatives corresponding to the variables w_1 and w_2 (i.e., $\{u_i, \overline{u}_i, u'_i, \overline{u}'_i \mid i \in \{1, 2\}\}$) would also have to be contained in M'. Due to $e_1 \succ u'_3$ and $e'_1 \succ \overline{u}'_3$, all alternatives corresponding to w_3 (i.e., $\{u_3, \overline{u}_3, u'_3, \overline{u}'_3\}$) are in M' as well. Consequently, e_2 and e'_2 are no longer upward covered and must also be in M'. The alternatives a_1, a_2, and a_3 are contained in every minimal upward covering set for A. But then M' is not minimal because the upward covering set M, which corresponds to the satisfying assignment stated above, is a strict subset of M'. Hence, e_1 cannot be contained in any minimal upward covering set for A.

We now list some properties of the dominance graph created by Construction 1 in general. The first property, stated in Claim 2(1), has already been seen in the example above. All proofs omitted due to space constraints are given in the full version of this paper [2].

Claim 2. *Let $(A, \succ)$ be the dominance graph created from φ by Construction 1.*

1. *Fix any j, $1 \leq j \leq \ell$. For each minimal upward covering set M for A, if the alternative e_j is in M then all alternatives are in M (i.e., $A = M$).*

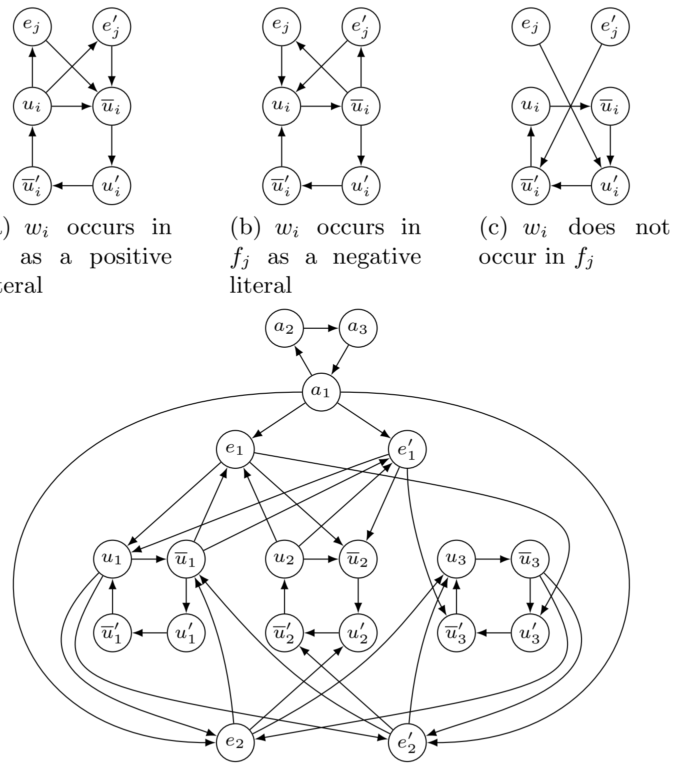

(a) w_i occurs in f_j as a positive literal

(b) w_i occurs in f_j as a negative literal

(c) w_i does not occur in f_j

(d) Example for the formula $(\neg w_1 \vee w_2) \wedge (w_1 \vee \neg w_3)$

Fig. 2. Dominance graph from Construction 1

2. φ is satisfiable if and only if there is no minimal upward covering set for A that contains any of the e_j, $1 \leq j \leq \ell$.

Construction 1 and Claim 2(2) already prove $\mathrm{MC_u}$-Member coNP-hard, via a reduction from the complement of SAT. Building on this reduction and that of Brandt and Fischer [4] to show $\mathrm{MC_u}$-Member NP-hard, we raise the lower bound to Θ_2^p-hardness. Wagner provided a sufficient condition for proving Θ_2^p-hardness that was useful also in other contexts (e.g., [13]):

Lemma 3 (Wagner [20]). *Let S be some* NP*-complete problem and let T be any set. If there is a polynomial-time computable function f such that, for all $m \geq 1$ and all strings $x_1, x_2, \ldots, x_{2m}$ satisfying that if $x_j \in S$ then $x_{j-1} \in S$ $(1 < j \leq 2m)$, $\|\{i \mid x_i \in S\}\|$ is odd if and only if $f(x_1, x_2, \ldots, x_{2m}) \in T$, then T is Θ_2^p-hard.*

One subtlety in our construction is due to the fact that we consider not only minimum-size but also (inclusion-)minimal covering sets. To the best of our knowledge, our constructions for the first time apply Wagner's technique [20] to problems defined in terms of minimality/maximality rather than

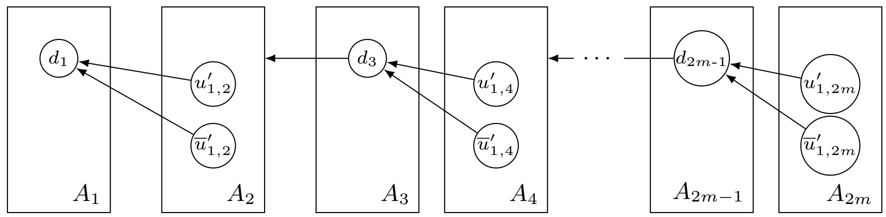

Fig. 3. Dominance graph from Construction 4. Most alternatives, and all edges between pairs of alternatives, in A_j for $1 \leq j \leq 2m$ have been omitted. All edges between alternatives in A_i and alternatives in A_j for $i \neq j$ are shown. An edge incident to a set of alternatives represents an edge incident to *each* alternative in the set.

minimum/maximum size of a solution: In Construction 4 below, we define a dominance graph based on Construction 1 and the construction from Brandt and Fischer [4] (which is also presented in the proof sketch of Thm. 4.1 [2]) such that Lemma 3 can be applied to prove Θ_2^p-hardness of MC_u-Member (see Theorem 6).

Construction 4 (for applying Lemma 3 to MC_u-Member). *We apply Wagner's Lemma with the* NP*-complete problem $S = \text{SAT}$ and construct a dominance graph. Fix an arbitrary $m \geq 1$ and let $\varphi_1, \varphi_2, \ldots, \varphi_{2m}$ be $2m$ boolean formulas in conjunctive normal form such that if φ_j is satisfiable then so is φ_{j-1}, for each j, $1 < j \leq 2m$. Without loss of generality, we assume that for each j, $1 \leq j \leq 2m$, the first variable of φ_j does not occur in all clauses of φ_j. It is easy to see that if φ_j does not have this property, it can be transformed into a formula that does have it, without affecting the satisfiability of the formula.*

We will now define a polynomial-time computable function f, which maps the given $2m$ boolean formulas to a dominance graph $(A, \succ)$ with useful properties for upward covering sets. Define $A = \bigcup_{j=1}^{2m} A_j$ and the dominance relation $\succ$ on A by $\left(\bigcup_{j=1}^{2m} \succ_j\right) \cup \left(\bigcup_{i=1}^{m} \left\{(u'_{1,2i}, d_{2i-1}), (\overline{u}'_{1,2i}, d_{2i-1})\right\}\right) \cup \left(\bigcup_{i=2}^{m} \{(d_{2i-1}, z) \mid z \in A_{2i-2}\}\right)$, where we use the following notation:

1. *For each i, $1 \leq i \leq m$, let $(A_{2i-1}, \succ_{2i-1})$ be the dominance graph that results from the formula φ_{2i-1} according to Brandt and Fischer's construction [4]. We use the same names for the alternatives in A_{2i-1} as in the proof sketch of [2, Thm. 4.1] that presents their construction, except that we attach the subscript $2i-1$. For example, alternative d from the proof sketch of [2, Thm. 4.1] now becomes d_{2i-1}, x_1 becomes $x_{1,2i-1}$, y_1 becomes $y_{1,2i-1}$, etc.*
2. *For each i, $1 \leq i \leq m$, let $(A_{2i}, \succ_{2i})$ be the dominance graph that results from the formula φ_{2i} according to Construction 1. We use the same names for the alternatives in A_{2i} as in that construction, except that we attach the subscript $2i$. For example, alternative a_1 from Construction 1 now becomes $a_{1,2i}$, e_1 becomes $e_{1,2i}$, u_1 becomes $u_{1,2i}$, and so on.*
3. *For each i, $1 \leq i \leq m$, connect the dominance graphs $(A_{2i-1}, \succ_{2i-1})$ and $(A_{2i}, \succ_{2i})$ as follows. Let $u_{1,2i}, \overline{u}_{1,2i}, u'_{1,2i}, \overline{u}'_{1,2i} \in A_{2i}$ be the four alternatives*

in the cycle corresponding to the first variable of φ_{2i}*. Then both* $u'_{1,2i}$ *and* $\overline{u}'_{1,2i}$ *dominate* d_{2i-1}*. The resulting dominance graph is denoted by* $(B_i, \succ_i^B)$*.*

4. *Connect the* m *dominance graphs* $(B_i, \succ_i^B)$*,* $1 \leq i \leq m$*, as follows: For each* i*,* $2 \leq i \leq m$*,* d_{2i-1} *dominates all alternatives in* A_{2i-2}*.*

Figure 3 sketches the dominance graph $(A, \succ)$ created by Construction 4. Clearly, $(A, \succ)$ is computable in polynomial time. Before we use this construction to prove $\mathrm{MC_u}$-MEMBER Θ_2^p-hard, we again list the key properties of this construction. Note that the first item of Claim 5 considers, for any fixed i with $1 \leq i \leq m$, the dominance graph $(B_i, \succ_i^B)$ resulting from the formulas φ_{2i-1} and φ_{2i} in Step 3 of Construction 4. Doing so will simplify our arguments for the whole dominance graph $(A, \succ)$ in the second and third item of Claim 5.

Claim 5. *Consider Construction 4.*

1. *For each* i*,* $1 \leq i \leq m$*, alternative* d_{2i-1} *is contained in some minimal upward covering set for* $(B_i, \succ_i^B)$ *if and only if* φ_{2i-1} *is satisfiable and* φ_{2i} *is not.*
2. *For each* i*,* $1 \leq i \leq m$*, let* M_i *be the minimal upward covering set for* $(B_i, \succ_i^B)$ *according to the cases in the proof of the first item. Then each of the sets* M_i *must be contained in every minimal upward covering set for* $(A, \succ)$*.*
3. *It holds that* $\|\{i \mid \varphi_i \in \mathrm{SAT}\}\|$ *is odd if and only if* d_1 *is contained in some minimal upward covering set* M *for* A*.*

5 Complexity of $\mathrm{MC_u}$-MEMBER and $\mathrm{MC_u}$-FIND

We now apply the constructions from Section 4 to show that $\mathrm{MC_u}$-MEMBER is Θ_2^p-hard and that $\mathrm{MC_u}$-FIND cannot be solved in polynomial time unless $\mathrm{P} = \mathrm{NP}$.

Theorem 6. $\mathrm{MC_u}$-MEMBER *is hard for* Θ_2^p *and in* Σ_2^p*.*

Our main goal is to determine the complexity of *finding* minimal unidirectional covering sets. As mentioned in the discussion in Section 3, the hardness of $\mathrm{MC_u}$-FIND does not follow directly from the NP-hardness of $\mathrm{MC_u}$-MEMBER, and neither from its Θ_2^p-hardness (Theorem 6). However, Construction 1 has another important property, stated in Claim 7, that can be applied to show coNP-hardness of the problem $\mathrm{MC_u}$-UNIQUE: Given a set A of alternatives and a dominance relation $\succ$ on A, does there exist a unique minimal upward covering set for A? And *this* property can be used to establish the hardness of the search problem $\mathrm{MC_u}$-FIND.

Claim 7. *Consider Construction 1. The boolean formula* φ *is not satisfiable if and only if there is a unique minimal upward covering set for* A*.*

Theorem 8. *If* $\mathrm{P} \neq \mathrm{NP}$ *then minimal upward covering sets cannot be found in polynomial time, i.e.,* $\mathrm{MC_u}$-FIND *is not polynomial-time solvable unless* $\mathrm{P} = \mathrm{NP}$*.*

Table 1. Overview of complexity results for various covering set problems. As indicated, previously known results are due to Brandt and Fischer [4]; all other results are new to this paper.

Problem Type	MC_u	MSC_u	MC_d	MSC_d
SIZE	NP-complete	NP-complete	NP-complete	NP-complete
MEMBER	Θ_2^p-hard, in Σ_2^p	Θ_2^p-complete	Θ_2^p-hard, in Σ_2^p	coNP-hard, in Θ_2^p
MEMBER-ALL	coNP-complete [4]	Θ_2^p-complete	coNP-complete [4]	coNP-hard, in Θ_2^p
UNIQUE	coNP-hard, in Σ_2^p	coNP-hard, in Θ_2^p	coNP-hard, in Σ_2^p	coNP-hard, in Θ_2^p
TEST	coNP-complete	coNP-complete	coNP-complete	coNP-complete
FIND	not in polynomial time unless P = NP	not in polynomial time unless P = NP	not in polynomial time unless P = NP (follows from [4])	not in polynomial time unless P = NP

Proof. Consider the problem of deciding whether there exists a *nontrivial* minimal upward covering set, i.e., one that does *not* contain all alternatives. By Construction 1 and Claim 7, there exists a trivial minimal upward covering set for A (i.e., one containing all alternatives in A) if and only if this set is the only minimal upward covering set for A. Thus, the coNP-hardness proof for MC_u-UNIQUE that is based on Claim 7 immediately implies that the problem of deciding whether there is a nontrivial minimal upward covering set for A is NP-hard. However, since the latter problem can easily be reduced to the search problem (because the search problem, when used as a function oracle, will yield the set of all alternatives if and only if this set is the only minimal upward covering set for A), it follows that the search problem cannot be solved in polynomial time unless P = NP. □

6 Generalizations

In addition to the (inclusion-)minimal unidirectional covering sets considered by Brandt and Fischer [4], we will also consider *minimum-size* covering sets, i.e., unidirectional covering sets of smallest cardinality. For some of the computational problems we study, different complexities can be shown for the minimal and minimum-size versions of the problem (see Theorem 9 and Table 1). Specifically, we will consider six types of computational problems, for both upward and downward covering sets, and for each both their "minimal" and "minimum-size" versions. In addition to MC_u-MEMBER and MC_u-FIND (which were defined in Section 2) and to MC_u-UNIQUE (which was defined in Section 5) we now define three more problem types:

1. MC_u-SIZE: Given a set A of alternatives, a dominance relation $\succ$ on A, and a positive integer k, does there exist some minimal upward covering set for A containing at most k alternatives?
2. MC_u-MEMBER-ALL: Given a set A of alternatives, a dominance relation $\succ$ on A, and a distinguished element $d \in A$, is d contained in all minimal upward covering sets for A?

3. $\mathrm{MC_u}$-Test: Given a set A of alternatives, a dominance relation $\succ$ on A, and a subset $M \subseteq A$, is M a minimal upward covering set for A?

If we replace "upward" by "downward" in the six problem types, we obtain the six corresponding "downward covering" versions, denoted by $\mathrm{MC_d}$-Size, $\mathrm{MC_d}$-Member, $\mathrm{MC_d}$-Member-All, $\mathrm{MC_d}$-Unique, $\mathrm{MC_d}$-Test, and $\mathrm{MC_d}$-Find. And if we replace "minimal" by "minimum-size" in the twelve problems already defined, we obtain the corresponding "minimum-size" versions: $\mathrm{MSC_u}$-Size, $\mathrm{MSC_u}$-Member, $\mathrm{MSC_u}$-Member-All, $\mathrm{MSC_u}$-Unique, $\mathrm{MSC_u}$-Test, $\mathrm{MSC_u}$-Find, $\mathrm{MSC_d}$-Size, $\mathrm{MSC_d}$-Member, $\mathrm{MSC_d}$-Member-All, $\mathrm{MSC_d}$-Unique, $\mathrm{MSC_d}$-Test, and $\mathrm{MSC_d}$-Find. Note that the four problems $\mathrm{MC_u}$-Find, $\mathrm{MC_d}$-Find, $\mathrm{MSC_u}$-Find, and $\mathrm{MSC_d}$-Find are search problems, whereas the other twenty problems are decision problems. Our results are stated in the following theorem.

Theorem 9. *The complexity of the covering set problems defined in this paper is as shown in Table 1.*

Acknowledgements This work was supported in part by DFG grants BR-2312/3-2, BR-2312/6-1, RO-1202/11-1, RO-1202/12-1, the ESF EUROCORES program LogICCC, and the Humboldt Foundation's TransCoop program.

References

1. N. Alon. Ranking tournaments. *SIAM Journal on Discrete Mathematics*, 20(1):137–142, 2006.
2. D. Baumeister, F. Brandt, F. Fischer, J. Hoffmann, and J. Rothe. The complexity of computing minimal unidirectional covering sets. Technical Report arXiv:0901.3692v3 [cs.CC], ACM Computing Research Repository (CoRR), 2009.
3. G. Bordes. On the possibility of reasonable consistent majoritarian choice: Some positive results. *Journal of Economic Theory*, 31:122–132, 1983.
4. F. Brandt and F. Fischer. Computing the minimal covering set. *Mathematical Social Sciences*, 56(2):254–268, 2008.
5. F. Brandt, F. Fischer, and P. Harrenstein. The computational complexity of choice sets. *Mathematical Logic Quarterly*, 55(4):444–459, 2009.
6. F. Brandt, F. Fischer, P. Harrenstein, and M. Mair. A computational analysis of the tournament equilibrium set. *Social Choice and Welfare*. Forthcoming.
7. Y. Chevaleyre, U. Endriss, J. Lang, and N. Maudet. A short introduction to computational social choice. In *Proceedings of the 33rd Conference on Current Trends in Theory and Practice of Computer Science*, volume 4362 of *Lecture Notes in Computer Science (LNCS)*, pages 51–69. Springer-Verlag, 2007.
8. V. Conitzer. Computing Slater rankings using similarities among candidates. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, pages 613–619. AAAI Press, 2006.
9. B. Dutta. Covering sets and a new Condorcet choice correspondence. *Journal of Economic Theory*, 44:63–80, 1988.

10. P. Faliszewski, E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. A richer understanding of the complexity of election systems. In S. Ravi and S. Shukla, editors, *Fundamental Problems in Computing: Essays in Honor of Professor Daniel J. Rosenkrantz*, chapter 14, pages 375–406. Springer, 2009.
11. P. Fishburn. Condorcet social choice functions. *SIAM Journal on Applied Mathematics*, 33(3):469–489, 1977.
12. M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, 1979.
13. E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Exact analysis of Dodgson elections: Lewis Carroll's 1876 voting system is complete for parallel access to NP. *Journal of the ACM*, 44(6):806–825, 1997.
14. E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Raising NP lower bounds to parallel NP lower bounds. *SIGACT News*, 28(2):2–13, 1997.
15. J. Laslier. *Tournament Solutions and Majority Voting*. Springer-Verlag, 1997.
16. D. McGarvey. A theorem on the construction of voting paradoxes. *Econometrica*, 21(4):608–610, 1953.
17. N. Miller. A new solution set for tournaments and majority voting: Further graph-theoretical approaches to the theory of voting. *American Journal of Political Science*, 24(1):68–96, 1980.
18. C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
19. C. Papadimitriou and S. Zachos. Two remarks on the power of counting. In *Proceedings of the 6th GI-Conference on Theoretical Computer Science*, volume 145 of *Lecture Notes in Computer Science (LNCS)*, pages 269–276. Springer-Verlag, 1983.
20. K. Wagner. More complicated questions about maxima and minima, and some closures of NP. *Theoretical Computer Science*, 51:53–80, 1987.
21. G. Woeginger. Banks winners in tournaments are difficult to recognize. *Social Choice and Welfare*, 20:523–528, 2003.