High-Multiplicity Election Problems[1]

Zack Fitzsimmons and Edith Hemaspaandra

Abstract

The computational study of elections generally assumes that the preferences of the electorate come in as a list of votes. Depending on the context, it may be much more natural to represent the list succinctly, as the distinct votes of the electorate and their counts, i.e., high-multiplicity representation. We consider how this representation affects the complexity of election problems. High-multiplicity representation may be exponentially smaller than standard representation, and so many polynomial-time algorithms for election problems in standard representation become exponential-time. Surprisingly, for polynomial-time election problems, we are often able to either adapt the same approach or provide new algorithms to show that these problems remain polynomial-time in the high-multiplicity case; this is in sharp contrast to the case where each voter has a weight, where the complexity usually increases. In the process we explore the relationship between high-multiplicity scheduling and manipulation of high-multiplicity elections. And we show that for any fixed set of job lengths, high-multiplicity scheduling on uniform parallel machines is in P, which was previously known for only two job lengths. We did not find any natural case where a polynomial-time election problem does not remain in P when moving to high-multiplicity representation. However, we found one natural NP-hard election problem where the complexity does increase, namely winner determination for Kemeny elections.

1 Introduction

Elections are an important and widely used tool for determining an outcome given the preferences of several agents (see, e.g., [8]). In most of the computational studies on elections, the preferences of the voters are represented as a list of votes. Though this may be a reasonable representation for paper ballots in political elections, in artificial intelligence applications a more succinct representation where the preferences of the electorate are represented as a list of distinct votes and their counts may be more natural. For example, this representation is used by the online preference repository PREFLIB for election data [38]. Following terminology from scheduling introduced by Hochbaum and Shamir [32] we refer to this as high-multiplicity representation.

We consider how high-multiplicity representation can affect the complexity of different election problems, and contrast this with the case of weighted voters.

High-multiplicity representation may be exponentially smaller than standard representation, and so many polynomial-time algorithms for election problems in standard representation become exponential-time. Surprisingly, we find that the complexity of polynomial-time election problems generally remains in P. We explain this phenomenon by showing that many common proof techniques that show that election problems are in P can be adapted to the high-multiplicity case (see Section 3).

We also explore the relationship between the well-studied problem of high-multiplicity scheduling and high-multiplicity election manipulation (see Section 4). In the process, we show that for any fixed set of job lengths, high-multiplicity scheduling on uniform parallel machines (this is the case where different machines can have different speeds) is in P

(Theorem 5). Previous work showed that for two job lengths, high-multiplicity scheduling on uniform parallel machines is in P [39], and very recent work shows that for any fixed number of job lengths, high-multiplicity scheduling on *identical* parallel machines (basically bin packing) is in P [21].

In general, we find that the complexity of polynomial-time election problems in standard representation does not increase when moving to the high-multiplicity case. This is in line with related work that considers high-multiplicity representation of elections (there simply called succinct elections), which did not find a case where the complexity increases [16, 14, 25, 17]. And this general behavior is also found in the literature on high-multiplicity scheduling. For example, with respect to scheduling, Clifford and Posner [9] states: "...we do not know of a problem that can be solved in polynomial time, under standard encoding, yet is NP-complete under HM encoding."

However, this does not mean that it is not possible for the complexity to increase. For example, consider the following election system ExactlyHalfApproval. Every voter votes by an approval vector (i.e., states an approval or disapproval for each candidate), and a candidate is a winner if they receive exactly half of all approvals. Though simple, this election system is not particularly natural. It has been chosen to show that an increase in complexity is possible. Consider the problem of constructive control by adding candidates (CCAC) for election system ExactlyHalfApproval. Given an election in high-multiplicity representation, a set of unregistered candidates, and a preferred candidate p, it is NP-complete to determine whether p can be made a winner of the election using system ExactlyHalfApproval by adding unregistered candidates to the election. This is because this problem is basically subset sum, which is in P for unary numbers by dynamic programming and is NP-complete for binary numbers [20]. (See Theorem 10 for a straightforward proof of this example. Also, we mention in passing that ExactlyHalfApproval-Weighted-CCAC is also NP-complete.)

We did find a natural case where the complexity increases, but for an NP-hard election problem. We find that determining the winner of a Kemeny election, which is well-known to be Θ_2^p-complete for standard representation [31], is Δ_2^p-complete for high-multiplicity representation. This solves an open problem from the previous work by Hemaspaandra, Spakowski, and Vogel [31]. (See Section 6.)

2 Preliminaries

An *election* (in *standard representation*) is defined as a finite set of candidates C, and a list of voters V, where each voter $v \in V$ has a preference order over the candidates. In high-multiplicity representation, V is not a list of voters, but instead a list of distinct votes v (preference orders) and their positive integer counts $\kappa(v)$. To see how these two definitions compare, consider the election that in standard representation consists of $C = \{a, b, c\}$ and the following four voters: $(a > b > c)$, $(a > b > c)$, $(a > b > c)$, and $(b > c > a)$. This list of four voters in high-multiplicity representation is: $3 \times (a > b > c)$ and $1 \times (b > c > a)$.

We will compare complexity results between high-multiplicity and weighted election problems. In a *weighted election*, V is still a list of voters, but now each $v \in V$ has an associated positive integer weight $\omega(v)$, and can be thought of as a coalition of $\omega(v)$ voters all voting the same. It is important to note that weights are indivisible, while the counts in high-multiplicity representation are not.[2] We assume that the preference order of each voter is a total order, i.e., he or she strictly ranks each candidate from most to least preferred.

An election system $\mathcal{E}$ maps an election to a set of winners, where the winners can be any

---

[2]Xia, Conitzer, and Procaccia [44] have a notion of "divisible" weights, but they allow noninteger divisions and so this is very different from high-multiplicity. It is interesting that their paper reduces manipulation for this notion to a scheduling problem, though a different scheduling problem than what we use.

subset of the candidate set. (This is known as the nonunique winner model.) The problem $\mathcal{E}$-Winner is defined below.

Name: $\mathcal{E}$-Winner

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

Question: Is p a winner of (C, V) using election system $\mathcal{E}$?

In the high-multiplicity case, $\mathcal{E}$-High-Multiplicity-Winner, we use high-multiplicity representation for V, and in the weighted case, $\mathcal{E}$-Weighted-Winner, each voter has a corresponding positive integer weight.

A scoring vector $\langle \alpha_1, \alpha_2, \ldots, \alpha_m \rangle$, $\alpha_i \geq \alpha_{i+1}$, defines an election system over m candidates. Each candidate receives α_i points for each vote where they are ranked ith, and the candidate(s) with the highest score win.

A *pure scoring rule* defines a family of scoring vectors where the m-candidate scoring vector can be computed in polynomial time in m, and the $(m+1)$-candidate scoring vector can be obtained from the m-candidate scoring vector by adding a single coefficient [4].

2.1 Manipulative Actions

We examine the complexity of the following manipulative actions on elections. For each problem, we present the definition for the standard representation, and describe how the high-multiplicity and weighted cases differ.

Manipulation is the most widely studied manipulative action on elections. The computational study of manipulation was introduced by Bartholdi, Tovey, and Trick [2], and was later extended for the coalitional, destructive, and weighted cases by Conitzer, Sandholm, and Lang [10]. We present the definition of constructive unweighted coalitional manipulation (CUCM) below.

Name: $\mathcal{E}$-CUCM

Given: A set of candidates C, a list of nonmanipulative voters V_1, a list of k manipulators V_2,[3] and a preferred candidate $p \in C$.

Question: Is there a way to set the preferences of the manipulators so that p is a winner of the election $(C, V_1 \cup V_2)$ using election system $\mathcal{E}$?

The weighted case (CWCM) is essentially the same as above, except each of the voters (both the nonmanipulators and the manipulators) have an associated positive integer weight. In the corresponding high-multiplicity case, we use high-multiplicity representation for the nonmanipulators as well as for the manipulators, which are represented as k encoded in binary.

Electoral control denotes the family of manipulative actions that consider an agent with control over the structure of the election, called the election chair, who wants to ensure his or her preferred outcome [3]. Below we define constructive control by adding voters (CCAV), which is a very natural case of control and can be thought of as modeling get-out-the-vote drives.

Name: $\mathcal{E}$-CCAV

Given: A set of candidates C, a list of registered voters V, a list of unregistered voters U, an add limit $k \in \mathbb{N}$, and a preferred candidate $p \in C$.

---

[3]This can also be thought of as specifying the number of manipulators, k, in unary.

Question: Does there exist a list U' of unregistered voters $U' \subseteq U$ such that U' consists of at most k unregistered voters and p is a winner of $(C, V \cup U')$ using election system $\mathcal{E}$?

In the standard model of weighted voter control the parameter k denotes the number of weighted voters the chair can add/delete [15]. In the high-multiplicity case, the only change from the $\mathcal{E}$-CCAV definition above is that we use high-multiplicity representation for both lists of voters.

2.2 Computational Complexity

We assume that the reader is familiar with the complexity classes P and NP. We also have results concerning the complexity classes Θ_2^p and Δ_2^p. Δ_2^p denotes the class of problems solvable in P with access to an NP oracle. Θ_2^p is a subset of Δ_2^p, where the P-machine can ask a log number of queries to an NP oracle (which is equivalent to the class of problems solvable by a P-machine that can ask one round of parallel queries to an NP oracle [22]).

Most of our polynomial-time results for high-multiplicity problems modify the proof of the problem in standard representation, but some results will use the well-known result due to Lenstra, which shows that even though solving an integer linear program is NP-complete in the general case [33, 7], it is in P when the number of variables is fixed [35]. Informally, an integer linear program is a system of linear inequalities with integer variables and coefficients.

3 Adapting Approaches

The general theme of this paper is that even though we would expect the complexity of election problems to increase (with respect to the length of the input) when using high-multiplicity representation, we find that in general such increases do not occur. In this section we will discuss several common approaches used to show that election problems are in P (in standard representation), and describe how they can be adapted for high-multiplicity representation, sometimes straightforwardly and sometimes in a more complicated way.

In particular, we show how greedy approaches, limited brute-forcing, network flow, and edge matching/cover techniques can be adapted. To showcase these adaptations, we will show that the dichotomy result for constructive control by adding voters (CCAV) for pure scoring rules [27] holds for the high-multiplicity case.

Assuming P $\neq$ NP, the following pure scoring rules are asymptotically (i.e., for a large enough number of candidates) the only cases where CCAV is in P, both for standard and high-multiplicity representation.

- $\langle \alpha, \beta, 0, \ldots, 0 \rangle$, where $\alpha > \beta$.

- t-Approval: $\langle 1^t, 0, \ldots, 0 \rangle$, where $t \leq 3$.

- t-Veto: $\langle 1, \ldots, 1, 0^t \rangle$, where $t \leq 2$.

- $\langle 2, 1, \ldots, 1, 0 \rangle$.

In contrast, all weighted cases are NP-complete, except triviality (i.e., a scoring rule where every candidate scores the same), 1-approval (plurality), 2-approval, and 1-veto [15, 37].

The following is a case where we need to adapt a greedy approach as well as to handle limited brute-forcing.

Theorem 1 *For $\alpha > \beta \geq 0$, $\langle \alpha, \beta, 0, \ldots, 0 \rangle$-CCAV is in* P *for elections in high-multiplicity representation.*

The proof of the above theorem can be found in Appendix A.1, but we will briefly describe the approach here. The greedy part of the algorithm is adapted by instead of adding one voter at a time, adding groups of voters at a time. The limited brute-forcing is more interesting. Essentially, we are able to replace each case of limited brute-forcing used in the proof of the standard case with an integer linear program (ILP) with a fixed number of variables for the high-multiplicity case, which can be solved in polynomial time by the result from Lenstra [35].

We now show how edge matching/cover techniques can be adapted.

Theorem 2 *CCAV is in* P *for t-approval when $t \leq 3$ and for t-veto when $t \leq 2$ for high-multiplicity elections.*

Proof. The results for 1-approval, 2-approval, and 1-veto follow via simple greedy algorithms that can easily be adapted to work in the high-multiplicity case.

3-approval-CCAV was shown to be in P by Lin [37] using a reduction to Simple b-Edge Matching for Multigraphs (see [43]). The essence of the reduction is that every unregistered voter of the form $(\{p, x, y\} > \cdots)$ (by which we mean a voter who gives a point to p, x, and y) corresponds to an edge (x, y) in the constructed multigraph. For every candidate $c \neq p$, c is a vertex in the graph and $b(c)$ is the final score of p minus the initial score of c (i.e., the number of points that we can add to c while keeping c's score less than or equal to p's score). In the high-multiplicity case, we would have too many edges in the graph. However, Capacitated b-Edge Matching (for graphs where edges have integer capacities) is also in P (see [43]). So, we simply set the capacity of edge (x, y) equal to the number of unregistered voters of the form $(\{p, x, y\} > \cdots)$, and we let $b(c)$ be the final score of p minus the initial score of c as previously.

Similarly, 2-veto-CCAV was shown to be in P by reduction to Simple b-Edge Cover for Multigraphs. In that reduction, every unregistered voter of the form $(\cdots > \{x, y\})$ corresponds to an edge (x, y) in the constructed multigraph. Again, we can modify this construction to a reduction to Capacitated b-Edge Cover, which is also in P, by letting the capacity of edge (x, y) be the number of unregistered voters of the form $(\cdots > \{x, y\})$. ❑

Another common technique to show that election problems are in P is network flow.

Theorem 3 *$\langle 2, 1, \ldots, 1, 0 \rangle$-CCAV is in* P *for high-multiplicity elections.*

Proof. The essence of the standard representation proof is to (after some easy preprocessing) build a min-cost flow network. The capacities of the edges are the scores of the candidates from the registered voters and the multiplicities of the votes of the unregistered voters. In the high-multiplicity case, we can use the exact same network. The capacities are now binary integers, but min-cost network flow is still in P in that case. ❑

Dynamic Programming There is one other common technique that is used to prove that election problems are in P. This is dynamic programming. As alluded to in the introduction, dynamic programming approaches do not generalize to the high-multiplicity case. Let's for example look at the result by Hemaspaandra and Schnoor [30] that shows that for all pure scoring rules with a constant number of different coefficients, manipulation is in P. This is shown by dynamic programming. And this algorithm is not in P for the high-multiplicity case, not even for very restricted versions of this problem. In the next section we consider a different approach to this problem using scheduling.

4 Manipulation as Scheduling

We now turn to the problem of manipulation for high-multiplicity elections for pure scoring rules. What does our election manipulation problem have to do with scheduling? As an example, consider scoring rule $\langle 4, 3, 3, 2, 0, \ldots, 0 \rangle$ and let $\{p, c_1, \ldots, c_m\}$ be the set of candidates. Let s_c be the score of candidate c from the nonmanipulators, and let k be the number of manipulators. Note that we can assume that all manipulators rank p first, so that p's final score is $s_p + 4k$. We need to see if we can set the manipulators such that every candidate $c \neq p$ scores at most $s_p + 4k$. As also pointed out in Bachrach et al. [1] for the case without nonmanipulators, we can view this as the following scheduling problem: We have $2k$ jobs of length 3 and k jobs of length 2, we have m machines, machine i corresponds to candidate c_i, and machine i has deadline $s_p + 4k - s_{c_i}$. It is easy to see that if p can be made a winner, then we can schedule the jobs such that every machine meets its deadline. However, we note that the converse does not hold: We need to make sure that every machine has at most k jobs scheduled on it. We also need to be careful that every voter ranks each candidate exactly once. The construction from Hemaspaandra and Schnoor [29] shows that there is a successful manipulation if and only if there is a successful schedule such that each machine has at most k jobs scheduled on it. Note that this also gives the obvious equivalence for the corresponding high-multiplicity versions of these problems.

More generally, we can state the following equivalence, which is implicit in the work by Hemaspaandra and Schnoor [29], which shows that manipulation is in P for all pure scoring rules with a constant number of different coefficients in standard representation.

Theorem 4 *Let $\langle \alpha_0, \alpha_1^{m_1}, \alpha_2^{m_2}, \ldots, \alpha_C^{m_C}, 0^{m_{C+1}} \rangle$ be a scoring rule such that $\alpha_0 \geq \alpha_1 > \alpha_2 > \cdots > \alpha_C > 0$ and $m_1, \ldots, m_{C+1}$ are positive integers such that $m_1 + \cdots + m_{C+1} = m$. Let the set of candidates be $\{p, c_1, \ldots, c_m\}$ and let there be k manipulators.*

*Then p can be made a winner if and only if for $1 \leq j \leq C$, we can schedule km_j jobs of length α_j on m machines such that the ith machine has deadline $s_p + \alpha_0 k - s_{c_i}$ **with the additional restriction that we schedule at most k jobs per machine**.*

Since our problem is closely related to scheduling and since high-multiplicity scheduling has been well-studied, it makes sense to see what is known there, so that we can see how easy or hard these problems are and so that we can try to adapt the results.

First some terminology. It is easy to see that scheduling with machine-dependent deadlines is equivalent to the well-known scheduling problem called scheduling on uniform parallel machines (where different machines can have different speeds) [39]. In the high-multiplicity version of the scheduling with machine-dependent deadlines problem, the input is given as $\ell_1, \ldots, \ell_C, n_1, \ldots, n_C, D_1, \ldots, D_m$, meaning that there are n_i jobs of length ℓ_i, and m machines with deadlines $D_1, \ldots, D_m$. If all deadlines are the same, the problem is equivalent to scheduling on identical parallel machines (this problem is also known as bin packing and as the cutting-stock problem). In the high-multiplicity version of that problem, the input is given as $\ell_1, \ldots, \ell_C, n_1, \ldots, n_C, m, D$ (note that this makes the input even more succinct, and so a polynomial-time algorithm for scheduling with machine-dependent deadlines does not necessarily imply a polynomial-time algorithm for the same problem where all machines have the same deadline).

In many contexts, the number of job lengths or item types is small and so high-multiplicity scheduling for a constant number of item types is an important problem. Surprisingly, this problem was open for a long time. Only in 2001 was this determined to be in P for two job lengths, both in the case where all machines have the same deadline as well as the case with machine-dependent deadlines [39]. The case for any constant number of job lengths was only very recently shown to be in P for the case where all machines have

the same deadline [21]. However, the algorithm from [21] does not give a polynomial-time result for the case of machine-dependent deadlines.

Given all this, do we have any chance of solving our high-multiplicity manipulation problems? There are some glimmers of hope. The first one is that the restriction to having at most k jobs per machine could make the problem easier (it could also make it harder of course). Secondly, we know that the number of jobs of each length is a multiple of k (this could make the problem simpler). Finally, we can take $\ell_1, \ldots, \ell_C$ to be constant, since for a pure scoring rule with a constant number of different coefficients, there are only a fixed number of coefficients that occur for any number of candidates, and by Theorem 4, the occurring coefficients are exactly the occurring job lengths. We will first show that for every fixed $C, \ell_1, \ldots, \ell_C$, high-multiplicity scheduling with machine-dependent deadlines (and thus also high-multiplicity scheduling on uniform parallel machines) is in P. Note that this is a very natural problem, since in many contexts, the set of job lengths or item types is fixed. So this is a very interesting result in its own right.

Theorem 5 *For any fixed set of job lengths, high-multiplicity scheduling on uniform parallel machines is in* P.

Proof. Let C be the number of job lengths and let $\ell_1, \ldots, \ell_C$ be these job lengths. As mentioned previously, high-multiplicity scheduling on uniform parallel machines is equivalent to high-multiplicity scheduling on machines with machine-dependent deadlines [39], and this is the problem that we will show to be in P.

Given $n_1, \ldots, n_C$, where n_i is the number of jobs of length ℓ_i, and $D_1, \ldots, D_m$, where D_i is the deadline for machine i, we can solve our scheduling problem in polynomial time, and pretty easily at that.

Example Case. We will first explain the core argument by looking at $C = 2$, $\ell_1 = 2$, and $\ell_2 = 3$. So, we have n_1 jobs of length 2, n_2 jobs of length 3, and m machines with deadlines $D_1, \ldots, D_m$. And the question is whether there exist nonnegative integers x_i (the number of jobs of length 2 scheduled on machine i) and y_i (the number of jobs of length 3 scheduled on machine i) such that $2x_i + 3y_i \leq D_i$, $\sum_{i=1}^m x_i = n_1$, and $\sum_{i=1}^m y_i = n_2$.

Consider the ith machine. We can group the jobs as "sixes," by which we mean 3 jobs of length 2 or 2 jobs of length 3, and "the leftovers," which consist of 0, 1, or 2 jobs of length 2 and 0 or 1 jobs of length 3. Note that the sum of the leftovers is at most 7.

Let D_i' be the largest integer that is $\leq D_i - 7$ and divisible by 6. Then D_i' time can be scheduled with just sixes and no leftovers. And so we can schedule successfully if and only if there exist a_1 jobs of length 2 and a_2 jobs of length 3 that can be scheduled on m machines with deadlines $D_1 - D_1', D_2 - D_2', \ldots, D_m - D_m'$ such that the remaining $(n_1 - a_1)$ jobs of length 2 and the remaining $(n_2 - a_2)$ jobs of length 3 can be scheduled on m machines with deadlines $D_1', \ldots, D_m'$ in sixes, i.e., as groups of 3 jobs of length 2 and groups of 2 jobs of length 3.

Why is this in P? Note that $D_i - D_i' \leq 12$, since D_i' is the largest integer that is $\leq D_i - 7$ and divisible by 6, and so $a_1 \leq 6m$ and $a_2 \leq 4m$, and we can use dynamic programming to compute all values of a_1 and a_2 such that a_1 jobs of length 2 and a_2 jobs of length 3 fit into $D_1 - D_1', D_2 - D_2', \ldots, D_m - D_m'$. Finally, check that $n_1 - a_1$ is divisible by 3, that $n_2 - a_2$ is divisible by 2, and that $(n_1 - a_1)/3 + (n_2 - a_2)/2 \leq (D_1' + \cdots + D_m')/6$.

This takes time $O(m^3 \log n)$, where n is the number of jobs. We can bring this down to $O(m^2 \log m \log n)$, by using Leung's approach and Leung's $O(n^2 \log m)$ dynamic programming algorithm [36].

It is easy to see how the approach for the example case generalizes to the case for each fixed set of job lengths $\ell_1, \ldots, \ell_C$. Let $\ell = lcm(\ell_1, \ldots, \ell_C)$ and instead of "sixes," look at "ℓs." Details can be found in Appendix A.2. ❏

Note that the above does not give a polynomial-time algorithm if only C is fixed and $\ell_1, \ldots, \ell_C$ are given as part of the input. And so the above does not imply the result that for two job lengths, high-multiplicity scheduling with machine-dependent deadlines is in P from McCormick, Smallwood, and Spieksma [39].

However, as mentioned previously, the case where $\ell_1, \ldots, \ell_C$ are fixed is very natural, and so this is an interesting result for scheduling. But recall from Theorem 4 that to solve manipulation, we in addition have to ensure that there are at most k jobs scheduled on each machine. As mentioned earlier, such a restriction can potentially make the complexity harder (as well as easier).

What are the problems with adding the restriction that we have at most k jobs per machine in the argument of the proof of Theorem 5? Let's look at the $C = 2, \ell_1 = 2, \ell_2 = 3$ case. In the case without the restriction that we have at most k jobs per machine, all deadlines with the same value mod 6 are treated the same, and every "six" is treated the same (no matter what machine it's on). This is no longer the case if we also require that every machine handles at most k jobs, since it then matters if you schedule three jobs of length 2 or two jobs of length 3 in a six.

However, we can show a number of high-multiplicity manipulation cases to be in P using scheduling, which is in contrast to the weighted case where only triviality and plurality are in P. Note that the theorem below does not cover all of the polynomial-time cases of manipulation for pure scoring rules in standard representation, since for example, we do not cover $\langle 2^{m/3}, 1^{m/3}, 0^{m/3} \rangle$ (though cases with at most two different coefficients can be handled easily). Also note that there is a limit to how far this can be generalized, since even in standard representation manipulation for Borda is NP-complete [5, 11].

Theorem 6 *Let $\alpha_0 \geq \alpha_1 > \alpha_2 > \cdots > \alpha_C > 0$. High-multiplicity manipulation for pure scoring rule $\langle \alpha_0, \alpha_1^{m_1}, \alpha_2^{m_2}, \ldots, \alpha_C^{m_C}, 0^{m-(m_1+\cdots+m_C)} \rangle$ is in P.*

Proof. By Theorem 4, it suffices to show that the following problem is in P: Given $k, D_1, \ldots, D_m$, can we schedule km_j jobs of length α_j on m machines with deadlines $D_1, \ldots, D_m$ such that every machine meets its deadline and such that there are at most k jobs scheduled on each machine.

Example Case. We first explain the core argument for $C = 2$, $\ell_1 = 2$, $\ell_2 = 3$, $m_1 = 1$, $m_2 = 1$. So, we are given $k, D_1, \ldots, D_m$, and we are asking whether we can we schedule k jobs of length 2 and k jobs of length 3 on m machines with deadlines $D_1, \ldots, D_m$ such that every machine meets its deadline and such that there are at most k jobs scheduled on each machine. That is, we are asking whether there exist nonnegative integers x_i (the number of jobs of length 2 scheduled on machine i) and y_i (the number of jobs of length 3 scheduled on machine i) such that $2x_i + 3y_i \leq D_i$, $\sum_{i=1}^m x_i = k$, $\sum_{i=1}^m y_i = k$, and $x_i + y_i \leq k$.

Note that we can assume that $D_i \leq 3k$. Also note that if $D_i \leq 2k$, our new requirement that machine i has at most k jobs, i.e., $x_i + y_i \leq k$, follows from the old requirement that $2x_i + 3y_i \leq D_i$.

If $D_i \leq 2k$ for all i, then we can solve our problem by using the algorithm from Theorem 5. If there are more than two machines with deadline $> 2k$, we can schedule k jobs of length 2 on machine 1, $\lceil k/2 \rceil$ jobs of length 3 on machine 2, and $\lfloor k/2 \rfloor$ jobs of length 3 on machine 3.

It remains therefore to look at the case where there are one or two machines with deadline $> 2k$. Without loss of generality, assume that the deadlines are in nonincreasing order and that we have at least two machines. Then our case is solvable if and only if there exist x_1, y_1, x_2, y_2 such that $2x_1 + 3y_1 \leq D_1$, $2x_2 + 3y_2 \leq D_2$, $x_1 + y_1 \leq k$, and $x_2 + y_2 \leq k$, such that $k - x_1 - x_2$ jobs of length 2 and $k - y_1 - y_2$ jobs of length 3 can be scheduled on $m - 2$ machines with deadlines $D_3, \ldots, D_m$ (since all these deadlines are $\leq 2k$, the requirement

that we schedule at most k jobs on each machine will be automatically satisfied). We would like to use the algorithm from Theorem 5, but we would need to run that algorithm for all values of x_1, y_1, x_2, y_2, which is not polynomial. The solution is to rephrase the algorithm from Theorem 5 as an integer linear program with a fixed number of variables (see Lemma 18 in Appendix A.3), which we then combine with the equations involving x_1, y_1, x_2, y_2 above. This ILP can be solved in polynomial time by Lenstra [35].

It is easy to see how to generalize this proof. For details, see Appendix A.4. ❑

5 Fixed Numbers of Candidates

It is reasonable to assume that the number of candidates in an election may be fixed. For weighted voters many problems are hard, even when the number of candidates is fixed.

Theorem 7   *1. m-candidate α-CWCM is NP-complete for every scoring rule α that is not plurality or triviality [23].*

2. m-candidate α-CCAV and m-candidate α-CCDV are each NP-complete for every scoring rule $\alpha = \langle \alpha_1, \ldots, \alpha_m \rangle$, when $\|\{\alpha_1, \ldots, \alpha_m\}\| \geq 3$, for weighted elections [15].

Faliszewski, Hemaspaandra, and Hemaspaandra [14] showed that high-multiplicity manipulation for scoring rules is in P for any fixed number of candidates by describing an ILP with a fixed number of variables.

Theorem 8 *[14] m-candidate α-CUCM is in P for every scoring rule α, for elections in high-multiplicity representation.*

A similar approach can be used to show that high-multiplicity constructive control by adding/deleting voters is in P for every scoring rule for a fixed number of candidates. Details can be found in Appendix A.5.

Theorem 9 *m-candidate α-CCAV and m-candidate α-CCDV are each in P for every scoring rule $\alpha = \langle \alpha_1, \ldots, \alpha_m \rangle$, for high-multiplicity elections.*

6 Natural Increase in Complexity

In general, we find that the complexity of polynomial-time election problems in standard representation does not increase when moving to the high-multiplicity case. However, this does not mean that such an increase is not possible. Recall the election system Exactly-HalfApproval defined in the introduction.

Theorem 10 *ExactlyHalfApproval-CCAC   is   in   P   and   ExactlyHalfApproval-High-Multiplicity-CCAC is NP-complete.*

Proof.   This problem is basically subset sum, which is in P for unary numbers and NP-complete for binary numbers [20]. To show that ExactlyHalfApproval-CCAC is in P, let $s_1, \ldots, s_m$ be the approval scores of candidates $c_1, \ldots, c_m$, let c_1 be the preferred candidate and let $c_1, \ldots, c_\ell$ be the registered candidates. We want to know if there exists a subset I of $\{\ell+1, \ldots, m\}$ of size at most k such that

$$2s_1 = \sum_{1 \leq i \leq \ell} s_i + \sum_{i \in I} s_i.$$

This can easily be computed in polynomial time, using dynamic programming. Simply let, for all $\ell + 1 \leq i \leq m$, $0 \leq k' \leq k$, $0 \leq t \leq \sum_{1 \leq i \leq m} s_i$, $S[i, t, k'] = 1$ if and only if there exists a size-k' subset I of $\{\ell + 1, \ldots, m\}$ such that $\sum_{1 \leq i \leq \ell} s_i + \sum_{i \in I} s_i = t$. This can be done in polynomial time, since there are only polynomially many i, k', and t to consider.

When the input is given in high-multiplicity representation, there are exponentially many values for t possible. In this case, we can show that it is NP-complete. We reduce from Partition: Given $\{k_1, \ldots, k_m\}$ such that $\sum k_i = 2K$, we need to determine if a subset of these integers sums to K. We have $m + 1$ candidates, preferred candidate p and candidates $c_1, \ldots, c_m$. We have K voters that approve of only p, and k_i voters that approve of only c_i. p is the only registered voter and the addition limit is m. It is immediate that there is a subset of $k_1, \ldots, k_m$ that sums to K if and only if p can be made a winner by adding candidates. ❑

As mentioned in the introduction, this system is not particularly natural. We will now present a natural case where the complexity increases, namely winner determination for Kemeny elections.

Kemeny elections were introduced in [34]. It is the unique election system that is neutral, consistent, and Condorcet consistent [46]. A candidate p is a Kemeny winner if p is ranked first in a Kemeny consensus, i.e., a linear order $>$ over the candidates that minimizes the Kendall's tau distance to V, i.e., that minimizes $\sum_{a,b \in C, a > b} \|\{v \in V \mid b >_v a\}\|$, where $>_v$ is the preference order of voter v.

Observation 11 *Kemeny-Weighted-Winner is equivalent to Kemeny-High-Multiplicity-Winner.*

Notice that the above observation holds since we do not modify the votes to score the election. This is not the case for Dodgson and Young elections, which we discuss in Appendix A.10. Kemeny-Winner was shown NP-hard in [2] and Θ_2^p-complete in [31]. Θ_2^p-hardness was shown by a chain of three reductions, ultimately showing that the Θ_2^p-complete problem Min-Card-Vertex-Cover-Compare reduces to Kemeny-Winner.

Footnote 3 in [31] points out that Kemeny-High-Multiplicity-Winner is in Δ_2^p and explicitly leaves the exact complexity of this problem as an open question. We will now show that Kemeny-High-Multiplicity-Winner (and Kemeny-Weighted-Winner) are in fact Δ_2^p-complete. We first define the following weighted version of Min-Card-Vertex-Cover-Compare.

Name: Min-Weight-Vertex-Cover-Compare

Given: Vertex-weighted graphs G and H such that $\omega(G) = \omega(H)$, where $\omega(G)$ denotes the weight of the graph.[4]

Question: Is the weight of G's minimum-weight vertex cover $\leq$ the weight of H's minimum-weight vertex cover?

Lemma 12 *Min-Weight-Vertex-Cover-Compare polynomial-time many-one reduces to Kemeny-High-Multiplicity-Winner.*

Proof. This follows by careful inspection and modification of the proof from [31]. In particular, since that proof consists of a chain of three reductions between Θ_2^p-complete problems, we need to define suitable Δ_2^p-complete weighted versions of the two intermediate problems and show that the weighted versions of the reductions still hold. In essence, we

_____
[4]Note that the weight of a vertex-weighted graph is the sum of the weights of each of its vertices.

"lift" the constructions and proofs from Θ_2^p to Δ_2^p. This works surprisingly smoothly. For details, see Appendix A.6. ❑

It remains to show the following lemma.

Lemma 13 *Min-Weight-Vertex-Cover-Compare is Δ_2^p-hard.*

Proof. For a formula $\phi(x_1, \ldots, x_n)$, we view an assignment for ϕ as the n-bit string $\alpha_1 \cdots \alpha_n$ such that α_i gives the assignment for variable x_i. We identify α with the binary number (between 0 and $2^n - 1$) that it represents. [26] shows that it is Δ_2^p-hard to compare the maximal satisfying assignments of two cnf formulas. By negating the variables in the formulas, we obtain the following Δ_2^p-hard promise problem, which we will reduce to Min-Weight-Vertex-Cover-Compare to prove Lemma 13.

Name: MINSATASG$_\leq$

Given: Two satisfiable 3cnf formulas $\phi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$.

Question: Is the minimal satisfying assignment of $\phi \leq$ the minimal satisfying assignment of ψ?

We explain the main idea behind the reduction here. Details can be found in Appendix A.7. Let $\phi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$ be two satisfiable 3cnf formulas. Without loss of generality, assume that ϕ and ψ have the same number m of clauses (simply pad).

Let $f(\phi)$ be the graph computed by the standard reduction from 3SAT to Vertex-Cover from [33]. We will not give the definition of the graph here, but list only the properties that we need, which follow immediately from the proof of [33]. Graph $f(\phi)$ consists of $2n + 3m$ vertices: a set L of $2n$ vertices corresponding to the literals in ϕ, i.e., $L = \{x_i, \overline{x_i} \mid 1 \leq i \leq n\}$ and a set of $3m$ clause vertices. Then the size of a minimum vertex cover of $f(\phi)$ is $n + 2m$ and satisfying assignments of ϕ correspond to minimum-size vertex covers in the following way:

- If W is a vertex cover of size $n + 2m$, then the set of literals $W \cap L$ corresponds to a satisfying assignment for ϕ.

- If α is a satisfying assignment for ϕ, then there is a vertex cover W of size $n + 2m$ such that $W \cap L$ corresponds to α.

Now we set the weights of the vertices so that the minimal satisfying assignment corresponds to the minimum-weight vertex cover. For each of the n vertices corresponding to variables, set $\omega(x_i) = 2^n + 2^{n-i}$ and set the weight of all other vertices to 2^n. It is not too hard to see that a minimum-weight vertex cover has size $n + 2m$ and that the weight of a size $n + 2m$ vertex cover corresponding to satisfying assignment α is $(n + 2m)2^n + \alpha$ and so α is the smallest satisfying assignment of ϕ if and only if $f(\phi)$ has a minimum-weight vertex cover of weight $(n + 2m)2^n + \alpha$.

Since ψ is also a satisfiable 3cnf formula over $x_1, \ldots, x_n$ with m clauses, it also holds α is the smallest satisfying assignment of ψ if and only if $f(\psi)$ has a minimum-weight vertex cover of weight $(n + 2m)2^n + \alpha$.

This then implies that ϕ's minimal satisfying assignment is $\leq \psi$'s minimal satisfying assignment if and only if the weight of $f(\phi)$'s minimum-weight vertex cover is less than or equal to the weight of $f(\psi)$'s minimum-weight vertex cover. Note that $\omega(f(\phi)) = \omega(f(\psi))$. This completes the reduction from MINSATASG$_\leq$ to Min-Weight-Vertex-Cover-Compare. ❑

This completes the proof of the main theorem of this section.

Theorem 14 *Kemeny-Weighted-Winner and Kemeny-High-Multiplicity-Winner are Δ_2^p-complete.*

Dwork et al. [13] show that Kemeny-Winner is already NP-hard for four voters. In fact, one can easily combine the techniques from Hemaspaandra, Spakowski, and Vogel [31] and Dwork et al. [13] to obtain the following theorem. For details, see Appendix A.8.

Theorem 15 *Kemeny-Winner for four voters is Θ_2^p-complete.*[5]

Do we get the same complexity jump if we look at the high-multiplicity and weighted case for four votes? Note that for the high-multiplicity case we allow many voters with the same vote, and so the Kemeny scores can be large. However, even though the scores can be large, it is easy to see that for each instance there are only a polynomial number of possible scores for each candidate, namely, if the multiplicities of the four votes are k_1, k_2, k_3, and k_4, the only possible scores are $\ell_1 k_1 + \ell_2 k_2 + \ell_3 k_3 + \ell_4 k_4$, where $0 \le \ell_i \le \|C\|(\|C\| - 1)/2$. For each candidate, simply query the Kemeny-Weighted-Score oracle for all these possible scores in parallel, and then determine whether p has the highest score.

Theorem 17 *Kemeny-Weighted-Winner and Kemeny-High-Multiplicity-Winner for four votes are Θ_2^p-complete.*

7 Conclusions and Future Work

High-multiplicity representation is a very natural way to represent an election. This representation may be exponentially smaller than the standard representation and so we would expect to see an increase in complexity (with respect to the length of the input). However, we were able to either adapt the approaches used for standard representation or provide new algorithms to show that polynomial-time election problems generally remain in P. We also explored the relationship between high-multiplicity scheduling and manipulation of high-multiplicity elections, which led to a new result in scheduling.

There are several interesting directions for future work. Is it possible for a high-multiplicity election problem to be harder than the corresponding weighted problem? Can we find natural cases, for election problems as well as for scheduling problems, where the complexity increases from P to NP-hard when going to high-multiplicity representation?

References

[1] Y. Bachrach, O. Lev, Y. Lewenberg, and Y. Zick. Misrepresentation in district voting. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 81–87. IJCAI/AAAI Press, July 2016.

---

[5]The definition of Kemeny winner can naturally be extended to voters voting any binary preference relation, and we note that the same argument as for Theorem 15 also shows the following. Details can be found in Appendix A.9.

Theorem 16 *Kemeny-Winner for two voters who vote a partial order and Kemeny-Winner for one voter who votes a binary relation are Θ_2^p-complete.*

[2] J. Bartholdi, III, C. Tovey, and M. Trick. The computational difficulty of manipulating an election. *Social Choice and Welfare*, 6(3):227–241, 1989.

[3] J. Bartholdi, III, C. Tovey, and M. Trick. How hard is it to control an election? *Mathematical and Computer Modeling*, 16(8/9):27–40, 1992.

[4] N. Betzler and B. Dorn. Towards a dichotomy of finding possible winners in elections based on scoring rules. In *Proceedings of the 34th International Symposium on Mathematical Foundations of Computer Science*, pages 124–136. Springer-Verlag *Lecture Notes in Computer Science #5734*, Aug. 2009.

[5] N. Betzler, R. Niedermeier, and G. Woeginger. Unweighted coalitional manipulation under the Borda rule is NP-hard. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, pages 55–60. IJCAI/AAAI Press, July 2011.

[6] D. Black. *The Theory of Committees and Elections*. Cambridge University Press, 1958.

[7] I. Borosh and L. Treybig. Bounds on positive integral solutions of linear Diophantine equations. *Proceedings of the American Mathematical Society*, 55(2):299–304, 1976.

[8] F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia. *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

[9] J. Clifford and M. Posner. Parallel machine scheduling with high multiplicity. *Mathematical Programming*, 89(3):359–383, 2001.

[10] V. Conitzer, T. Sandholm, and J. Lang. When are elections with few candidates hard to manipulate? *Journal of the ACM*, 54(3):Article 14, 2007.

[11] J. Davies, G. Katsirelos, N. Narodytska, and T. Walsh. Complexity and algorithms for Borda manipulation. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, pages 657–662. AAAI Press, Aug. 2011.

[12] C. Dodgson. A method of taking votes on more than two issues. Pamphlet printed by the Clarendon Press, Oxford, and headed "not yet published" (see the discussions in [41, 6], both of which reprint this paper), 1876.

[13] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th International World Wide Web Conference*, pages 613–622. ACM Press, Mar. 2001.

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

[15] P. Faliszewski, E. Hemaspaandra, and L. Hemaspaandra. Weighted electoral control. *Journal of Artificial Intelligence Research*, 52:507–542, 2015.

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

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

[18] Z. Fitzsimmons and E. Hemaspaandra. The complexity of succinct elections. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (Student Abstract)*, pages 4921–4922. AAAI Press, Feb. 2017.

[19] Z. Fitzsimmons and E. Hemaspaandra. High-multiplicity election problems. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems*, July 2018. To appear.

[20] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.

[21] M. Goemans and T. Rothvoß. Polynomiality for bin packing with a constant number of item types. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 830–839. SIAM, Jan. 2014.

[22] L. Hemachandra. The strong exponential hierarchy collapses. *Journal of Computer and System Sciences*, 39(3):299–322, 1989.

[23] E. Hemaspaandra and L. Hemaspaandra. Dichotomy for voting systems. *Journal of Computer and System Sciences*, 73(1):73–83, 2007.

[24] 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.

[25] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. Hybrid elections broaden complexity-theoretic resistance to control. *Mathematical Logic Quarterly*, 55(4):397–424, 2009.

[26] E. Hemaspaandra, L. Hemaspaandra, and J. Rothe. The complexity of online manipulation of sequential elections. *Journal of Computer and System Sciences*, 80(4):697–710, 2014.

[27] E. Hemaspaandra, L. Hemaspaandra, and H. Schnoor. A control dichotomy for pure scoring rules. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 712–720. AAAI Press, July 2014.

[28] E. Hemaspaandra, L. Hemaspaandra, and H. Schnoor. A control dichotomy for pure scoring rules. Technical Report arXiv:1404.4560 [cs.GT], arXiv.org, Apr. 2014.

[29] E. Hemaspaandra and H. Schnoor. Complexity dichotomies for unweighted scoring rules. Technical Report arXiv:1604.05264 [cs.CC], arXiv.org, Apr. 2016.

[30] E. Hemaspaandra and H. Schnoor. Dichotomy for pure scoring rules under manipulative electoral actions. In *Proceedings of the 22nd European Conference on Artificial Intelligence*, pages 1071–1079. IOS Press, Aug/Sep 2016.

[31] E. Hemaspaandra, H. Spakowski, and J. Vogel. The complexity of Kemeny elections. *Theoretical Computer Science*, 349(3):382–391, 2005.

[32] D. Hochbaum and R. Shamir. Strongly polynomial algorithms for the high multiplicity scheduling problem. *Operations Research*, 39(4):648–653, 1991.

[33] R. Karp. Reducibilities among combinatorial problems. In R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103, 1972.

[34] J. Kemeny. Mathematics without numbers. *Daedalus*, 88:577–591, 1959.

[35] H. Lenstra, Jr. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983.

[36] J. Leung. On scheduling independent tasks with restricted execution times. *Operations Research*, 30(1):163–171, 1982.

[37] A. Lin. *Solving Hard Problems in Election Systems*. PhD thesis, Rochester Institute of Technology, Rochester, NY, 2012.

[38] N. Mattei and T. Walsh. PREFLIB: A library for preferences. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory*, pages 259–270, 2013.

[39] T. McCormick, S. Smallwood, and F. Spieksma. A polynomial algorithm for multiprocessor scheduling with two job lengths. *Mathematics of Operations Research*, 26(1):31–49, 2001.

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

[41] I. McLean and A. Urken. *Classics of Social Choice*. University of Michigan Press, 1995.

[42] J. Rothe, H. Spakowski, and J. Vogel. Exact complexity of the winner problem for Young elections. *Theory of Computing Systems*, 36(4):375–386, 2003.

[43] A. Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency*. Springer-Verlag, 2003.

[44] L. Xia, V. Conitzer, and A. Procaccia. A scheduling approach to coalitional manipulation. In *Proceedings of the 11th ACM Conference on Electronic Commerce*, pages 275–284. ACM Press, June 2010.

[45] H. Young. Extending Condorcet's rule. *Journal of Economic Theory*, 16(2):335–353, 1977.

[46] H. Young and A. Levenglick. A consistent extension of Condorcet's election principle. *SIAM Journal on Applied Mathematics*, 35(2):285–300, 1978.

A   Appendix

A.1   Proof of Theorem 1: For $\alpha > \beta \geq 0$, $\langle\alpha, \beta, 0, \ldots, 0\rangle$-CCAV is in P for elections in high-multiplicity representation.

Proof.   We follow the proof of the claim from [27], which is found in its corresponding technical report [28], and show how to adapt this to the high-multiplicity case.

The proof shows that there is a constant ℓ such that it is never better to add ℓ votes that put p second than it is to add ℓ different votes that put p first (assuming that votes that rank the first two candidates identically also rank all other candidates identically). Let V_1 be the set of unregistered voters that put p first and let V_2 be the set of unregistered voters that put p second.

So, if p can be made a winner, p can be made a winner by adding at most ℓ voters from V_2 or in a way that does not leave ℓ different votes in V_1 unused.

In the first case, we can brute-force over all sets of V_2 voters of size $\leq \ell$. Add these voters. We are then left with the problem of checking whether we can add at most k' V_1 voters to an election to make p a winner, where k' is k minus the number of V_2 voters added.

In Hemaspaandra, Hemaspaandra, and Schnoor [27], this is done by brute-forcing over every j in $\{0, \ldots, k'\}$ and checking whether we can add j voters from V_1 to make p a winner. If we know j, we know what the score of p after control will be, and so the last part can then be done greedily, by for each candidate adding as many voters that put a second as possible one voter at the time (until we have added j voters total).

We have two problems in the high-multiplicity case. The first is that we cannot in polynomial time brute-force over all possible values of j, since k' is not polynomially bounded. The second problem is that we don't have the time to add voters one at the time. The second problem can of course easily be fixed by, for each candidate a, adding as many voters with a in second place as possible all at once: If s_a is the initial score of a and fs_p is the final score of p (note that $fs_p = s_p + j\alpha$), we can add at most $\lfloor (fs_p - s_a)/\beta \rfloor$ voters with a second.

To handle the first problem, note that we can formulate the problem as an integer linear program with one variable (j), namely, there is a j, $0 \leq j \leq k'$ such that

$$\sum_{a \in C - \{p\}} \min(\lfloor (s_p + j\alpha - s_a)/\beta \rfloor, v_a) \geq j,$$

where v_a is the number of voters in V_1 that rank a second.[6] This integer linear program can be solved in polynomial time by Lenstra [35].

It remains to show that the second case, where p can be made a winner in a way that does not leave ℓ different votes in V_1 unused can be determined in polynomial time in the high-multiplicity case.

For the standard case, in Hemaspaandra, Hemaspaandra, and Schnoor [27] this case is handled by brute-forcing over all sets S of at most $\ell - 1$ candidates who are ranked second by unused V_1 voters, and then brute-forcing over all possible sets of unused V_1 voters consistent with S. We can describe a set of voters consistent with S as a function $u : S \rightarrow \{1, \ldots, \|V_1\|\}$ such that $u(c)$ is the number of unused V_1 voters that rank c second. We can brute-force over all such functions u in polynomial time. This way, we loop over all possible sets of unused V_1 voters, and so we also loop over all possible sets of added V_1 voters. After adding the V_1 voters, we need to see if p can be made a winner by adding V_2 voters. This can be done in a similar way as adding the V_1 voters in the first case above.

In the high-multiplicity case, we can still loop over all sets S of at most $\ell - 1$ candidates. We can not brute-force all functions u, but this will turn again into an integer linear program, this time with $\ell - 1$ variables (for the u values) and one more variable to handle the V_2 voters similarly to how we handled the V_1 voters in the first case above. This integer linear program can be solved in polynomial time by Lenstra [35]. ❑

A.2 General Case of Theorem 5: For any fixed set of job lengths, high-multiplicity scheduling on uniform parallel machines is in P.

Proof.   It is easy to see how the approach for the example case (see Theorem 5 in Section 4) generalizes to the case for each fixed set of job lengths $\ell_1, \ldots, \ell_C$. Let $\ell = lcm(\ell_1, \ldots, \ell_C)$. Our input is $n_1, \ldots, n_C, D_1, \ldots, D_m$. Consider the ith machine. We can group the jobs as "ℓs," by which we mean ℓ/ℓ_j jobs of length ℓ_j for some j, and "the

---

[6]Notice that the ILP above checks that if we add the maximum number of voters such that the score of each of the non-p candidates is at most $s_p + j\alpha$, then we have added at least j voters. This implies that we can add j voters in such a way that p is still a winner.

leftovers," which consist of $< \ell/\ell_j$ jobs of length ℓ_j for each j. Note that the sum of the leftovers is at most $\sum_{j=1}^{C}(\ell/\ell_j - 1)$.

For each i, let D_i' be the largest integer that is $\leq \sum_{j=1}^{C}(\ell/\ell_j - 1)$ and divisible by ℓ. Then D_i' time can be scheduled with just ℓs and no leftovers. And so we can schedule successfully if and only if there exist a_j jobs of length ℓ_j, for $1 \leq j \leq C$, that can be scheduled on m machines with deadlines $D_1 - D_1', D_2 - D_2', \ldots, D_m - D_m'$ such that all remaining $(n_j - a_j)$ jobs of length ℓ_j, for $1 \leq j \leq C$, can be scheduled on m machines with deadlines $D_1', \ldots, D_m'$ in ℓs, i.e., as groups of ℓ/ℓ_j jobs of length ℓ_j.

Why is this in P? Note that $D_i - D_i' \leq \sum_{j=1}^{C}(\ell/\ell_j - 1) + \ell - 1$, since D_i' is the largest integer that is $\leq \sum_{j=1}^{C}(\ell/\ell_j - 1)$ and divisible by ℓ. Since C and $\ell_1, \ldots, \ell_C$ are fixed, $a_j = O(m)$ and we can use dynamic programming to compute all values of $a_1, \ldots, a_C$ such that a_j jobs of length ℓ_j fit into $D_1 - D_1', D_2 - D_2', \ldots, D_m - D_m'$. Finally, check that $n_j - a_j$ is divisible by ℓ_j and that $\sum_{j=1}^{C}(n_j - a_j)/\ell_j \leq (D_1' + \cdots + D_m')/\ell$.

Note that since the ℓ_i's and C are constant the running time provided for the example case (in the proof of Theorem 5 in Section 4) also holds for the general case. ❏

A.3   ILP Formulation of Theorem 5.

Lemma 18 *For any fixed set of job lengths, there is an ILP with a fixed number of variables that solves high-multiplicity scheduling on parallel machines with machine-dependent deadlines.*

Proof.   Let C be the number of job lengths and let $\ell_1, \ldots, \ell_C$ be these job lengths. Given $n_1, \ldots, n_C$, where n_i is the number of jobs of length ℓ_i, and $D_1, \ldots, D_m$, where D_i is the deadline for machine i. In the proof of Theorem 5, we proved that this scheduling problem is in P using dynamic programming. However, for the proof of Theorem 6 we need a different approach. Below we present an ILP formulation with a fixed number of variables that computes if we can schedule successfully that replaces the dynamic programming. Again, to show the essence of the argument, we consider the case of $C = 2$, $\ell_1 = 2$, and $\ell_2 = 3$.

Look at a proof of Theorem 5 in Section 4 (the example case). We are already pretty close: There exist a_1 and a_2 such that a_1 jobs of length 2 and a_2 jobs of length 3 fit into $D_1 - D_1', D_2 - D_2', \ldots, D_m - D_m'$, $n_1 - a_1$ is divisible by 3, $n_2 - a_2$ is divisible by 2, and $(n_1 - a_1)/3 + (n_2 - a_2)/2 \leq (D_1' + \cdots + D_m')/6$.

So all that is left to do is to write "a_1 jobs of length 2 and a_2 jobs of length 3 fit into $D_1 - D_1', D_2 - D_2', \ldots, D_m - D_m'$" as an integer linear program with a fixed number of variables. Recall that $D_i - D_i' \leq 12$. For $0 \leq j \leq 12$, let d_j be the number of machines i such that $D_i' = d_j$.

Our integer linear program formulation look as follows. We introduce variables $a_{j,(r,t)}$ for $0 \leq j \leq 12$ and $2r + 3t \leq d_j$. $a_{j,(r,t)}$ stands for the number of machines i such that $D_i' = d_j$ and such that D_i' is scheduled with r 2s and t 3s. We add the following obvious equations:

$$\sum_{0 \leq j \leq 12, 2r+3t \leq d_j} a_{j,(r,t)} r = a_1$$

$$\sum_{0 \leq j \leq 12, 2r+3t \leq d_j} a_{j,(r,t)} t = a_2$$

It is easy to see that the example case generalizes to the general case, using the general case of the proof of Theorem 5 ❏

A.4 General Case of Theorem 6: Let $\alpha_0 \geq \alpha_1 > \alpha_2 > \cdots > \alpha_C > 0$. High-multiplicity manipulation for pure scoring rule $\langle \alpha_0, \alpha_1^{m_1}, \alpha_2^{m_2}, \ldots, \alpha_C^{m_C}, 0^{m-(m_1+\cdots+m_C)} \rangle$ is in P.

Proof. By Theorem 4, it suffices to show that the following problem is in P: Given $k, D_1, \ldots, D_m$, can we schedule km_j jobs of length α_j on m machines with deadlines $D_1, \ldots, D_m$ such that every machine meets its deadline and such that there are at most k jobs scheduled on each machine.

Note that we can assume that $D_i \leq \alpha_1 k$. Also note that if $D_i \leq \alpha_C k$, our new requirement that machine i has at most k jobs follows from the old requirement that D_i meets its deadline.

If there are at least $\alpha_1(m_1 + \cdots + m_C)$ machines with deadline $> \alpha_C k \geq k$, then we can always schedule $\lceil k/\alpha_1 \rceil$ jobs on each such machine.

Let $\widehat{m} = \alpha_1(m_1 + \cdots + m_C) - 1$. Note that $\widehat{m}$ is a constant. It remains therefore to look at the case where there are at most $\widehat{m}$ machines with deadline $> k$. Without loss of generality, assume that the deadlines are in nonincreasing order and that we have at least $\widehat{m}$ machines. Our case is solvable if and only if there exist $x_{i,j}$ for $1 \leq i \leq \widehat{m}$ and $1 \leq j \leq C$, the number of jobs scheduled on machine i of length ℓ_j, such that $\sum_{j=1}^C \alpha_i x_{i,j} \leq D_i$, $\sum_{i=1}^m x_{i,j} = km_j$, $\sum_{j=1}^C x_{i,j} \leq k$, and such that $k - \sum_{i=1}^{\widehat{m}} x_{i,j}$ jobs of length ℓ_j can be scheduled on $m - \widehat{m}$ machines with deadlines $D_{\widehat{m}+1}, \ldots, D_m$ (since all these deadlines are $\leq k$, the requirement that we schedule at most k jobs on each machine will be automatically satisfied). Use the ILP formulation of the algorithm from Theorem 5 to solve this last requirement and combine this with our equations for $x_{i,j}$ for $1 \leq i \leq \widehat{m}$ and $1 \leq j \leq C$, to obtain an ILP with a fixed number of variables, which can be solved in polynomial time by Lenstra [35]. ❑

A.5 Proof of Theorem 9: m-candidate α-CCAV and m-candidate α-CCDV are each in P for every scoring rule $\alpha = \langle \alpha_1, \ldots, \alpha_m \rangle$, for high-multiplicity elections.

Proof. Given a set of candidates C, a list of registered voters V in high-multiplicity representation, a list of unregistered voters U in high-multiplicity representation, an addition limit $k \in \mathbb{N}$, and a preferred candidate $p \in C$ we can determine if the chair can ensure that p wins in polynomial time. The argument closely follows the approach used by Faliszewski, Hemaspaandra, and Hemaspaandra [14] to show that m-candidate α-bribery and m-candidate α-CUCM, for every scoring rule α are each in P for high-multiplicity (there called succinct) elections.

We describe an integer linear program with a fixed number of variables, which due to the result from Lenstra [35], can be solved in polynomial time.

Since the number of candidates m is fixed, we know that only $m!$ different preference orders are possible. We list these preference orders in order in the following way: $o_1, \ldots, o_{m!}$. For $1 \leq i \leq m!$, let k_i be the number of voters in V with preference order o_i, and let k_i' be the number of voters in U with preference order o_i. Let $x_1, \ldots, x_{m!}$ be variables. The constants used as input to our linear program are the coefficients of the scoring vector $\alpha_1, \ldots, \alpha_m$, the counts for the registered voters $k_1, \ldots, k_{m!}$, the counts for the unregistered voters $k_1', \ldots, k_{m!}'$, and the addition limit k. We have the following constraints.

We first need to ensure that all of our variables have a nonnegative value. So, $\forall i, 1 \leq i \leq m!$,

$$x_i \geq 0.$$

We also need to ensure that no more than the number of unregistered votes of each type are added. So, $\forall i, 1 \le i \le m!$,

$$x_i \le k_i'.$$

The following constraint ensures that no more than k total votes are added.

$$\sum_{i=1}^{m!} x_i \le k.$$

Our final constraint ensures that no other candidate has a score greater than p. (Note that $\mathrm{pos}(c, j)$ below denotes the position of candidate c in preference order o_j.) So, $\forall i, 1 \le i \le m$,

$$\sum_{j=1}^{m!} \alpha_{\mathrm{pos}(p,j)}(k_j + x_j) \ge \sum_{j=1}^{m!} \alpha_{\mathrm{pos}(c_i,j)}(k_j + x_j).$$

It is easy to see that the above constraints are satisfied if and only if it is possible to add at most k unregistered voters from U such that p wins.

The above integer linear program can be easily modified for the case of deleting voters. ❏

A.6 Proof of Lemma 12: Min-Weight-Vertex-Cover-Compare polynomial-time many-one reduces to Kemeny-High-Multiplicity-Winner.

As mentioned in the core text, this follows by careful inspection and modification of the proof from [31] that Min-Card-Vertex-Cover-Compare polynomial-time many-one reduces to Kemeny-Winner. That proof consists of a chain of three reductions:

1. Min-Card-Vertex-Cover-Compare reduces to Vertex-Cover-Member.

2. Vertex-Cover-Member reduces to Feedback-Arc-Set-Member.

3. Feedback-Arc-Set-Member reduces to Kemeny-Winner.

We define the following weighted versions of the two intermediate problems.

Name: Weighted-Vertex-Cover-Member

Given: Vertex-weighted graph G and vertex v in G.

Question: Is v a member of a minimum-weight vertex cover of G?

Name: Weighted-Feedback-Arc-Set-Member

Given: Irreflexive and antisymmetric edge-weighted digraph G and vertex v in G.

Question: Is there a minimum-weight feedback arc set of G that contains all arcs entering v?

And we show the following.

1. Min-Weight-Vertex-Cover-Compare reduces to Weighted-Vertex-Cover-Member (Lemma 19).

2. Weighted-Vertex-Cover-Member reduces to Weighted-Feedback-Arc-Set-Member (Lemma 20).

3. Weighted-Feedback-Arc-Set-Member reduces to Kemeny-High-Multiplicity-Winner (Lemma 21).

Lemma 19 *Min-Weight-Vertex-Cover-Compare polynomial-time many-one reduces to Weighted-Vertex-Cover-Member.*

Proof. We extend the construction from Lemma 4.12 from [31] to vertex-weighted graphs. Given two vertex-weighted graphs G and H with $\omega(G) = \omega(H)$, construct graph F as follows. $F = (G + (\{v\}, \emptyset)) \times (H + (\{w\}, \emptyset))$, keeping the weights of the vertices in G and H, and letting $\omega(v) = \omega(w) = 1$. Map (G, H) to (F, w). Write mvc for the minimum vertex cover weight of a vertex-weighted graph. We extend the argument from Lemma 4.12 from [31] to show that this mapping reduces Min-Weight-Vertex-Cover-Compare to Weighted-Vertex-Cover-Member.

Note that V is a vertex cover of F if and only if

1. V contains all vertices in $G + (\{v\}, \emptyset)$ and a vertex cover of H. The lightest such vertex covers have weight $\omega(G) + 1 + mvc(H)$ and do not contain w.

2. V contains all vertices in $H + (\{w\}, \emptyset)$ and a vertex cover of G. The lightest such vertex covers have weight $mvc(G) + \omega(H) + 1$ and contain w.

Since $\omega(G) = \omega(H)$ it is immediate that $mvc(G) \leq mvc(H)$ if and only if there is a minimum-weight vertex cover of F that contains w. ❏

Lemma 20 *Weighted-Vertex-Cover-Member polynomial-time many-one reduces to Weighted-Feedback-Arc-Set-Member.*

Proof. We extend the construction from the proof of Lemma 4.8 from [31], which itself is an extension of the standard reduction from Vertex-Cover to Feedback-Arc-Set from [33], to graphs with weights. Given a vertex-weighted graph G, define edge-weighted digraph $H = (W, A)$ as follows.

1. $W = \{v, v' \mid v \in V(G)\}$.

2. A consists of the following edges.

   (a) For $v \in V(G)$, edge (v, v') with weight $\omega(v)$.
   (b) For $\{v, w\} \in E(G)$, edges (v', w) and (w', v) with weight $\omega(G)$.

The reduction maps $(G, \hat{v})$ to $(H, \hat{v}')$.

If V' is a vertex cover of G that contains $\hat{v}$, then, by the proof of Lemma 4.9 from [31], $\{(v, v') \mid v \in V'\}$ is a feedback arc set for H that contains all arcs entering $\hat{v}'$. Note that the weight of this feedback arc set is $\omega(V')$.

Now suppose that A' is a feedback arc set of H that contains $(\hat{v}, \hat{v}')$. Again by the proof of Lemma 4.9 from [31], $V' = \{v \in V \mid \exists w \in W : (v, w) \in A' \text{ or } (v', w) \in A'\}$ forms a vertex cover of G that contains $\hat{v}$. Note that $\omega(V') \leq \omega(A')$. ❏

Lemma 21 *Weighted-Feedback-Arc-Set-Member polynomial-time many-one reduces to Kemeny-High-Multiplicity-Winner.*

Proof. Given an irreflexive and antisymmetric edge-weighted digraph G and vertex v in G, we first multiply all edge weights by 2 (this does not change the membership of (G, v) in Weighted-Feedback-Arc-Set-Member). We then use McGarvey's construction [40] to construct in polynomial time an election in high-multiplicity representation that has G as its weighted majority graph. The remainder of the proof is identical to the proof of Lemma 4.2 of [31]. ❑

A.7 Proof of Lemma 13: Min-Weight-Vertex-Cover-Compare is Δ_2^p-hard.

Proof. For $\phi(x_1, \ldots, x_n)$ a formula, we view an assignment for ϕ as the n-bit string $\alpha_1 \cdots \alpha_n$ such that α_i gives the assignment for variable x_i. We identify α with the binary number (between 0 and $2^n - 1$) that it represents. [26] shows that it is Δ_2^p-hard to compare the maximal satisfying assignments of two cnf formulas. By negating the variables in the formulas, we obtain the following Δ_2^p-hard promise problem, which we will reduce to Min-Weight-Vertex-Cover-Compare to prove Lemma 13.

Name: MINSATASG$_\leq$

Given: Two satisfiable 3cnf formulas $\phi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$.

Question: Is the minimal satisfying assignment of $\phi \leq$ the minimal satisfying assignment of ψ?

Let $\phi(x_1, \ldots, x_n)$ and $\psi(x_1, \ldots, x_n)$ be two satisfiable 3cnf formulas. Without loss of generality, assume that ϕ and ψ have the same number m of clauses (simply pad).

Let $f(\phi)$ be the graph computed by the standard reduction from 3SAT to Vertex-Cover from [33]. Then $f(\phi)$ consists of $2n+3m$ vertices: $\{x_i, \overline{x_i} \mid 1 \leq i \leq n\} \cup \{a_i, b_i, c_i \mid 1 \leq i \leq m\}$, and the following edges: $\{x_i, \overline{x_i}\}$ for $1 \leq i \leq n$, $\{a_j, b_j\}$, $\{a_j, c_j\}$, and $\{b_j, c_j\}$ for $1 \leq j \leq m$, and if the jth clause in ϕ is $\ell_1 \vee \ell_2 \vee \ell_3$, where $\ell_r \in \{x_i, \overline{x_i} \mid 1 \leq i \leq n\}$, then we have edges $\{a_j, \ell_1\}$, $\{b_j, \ell_2\}$, $\{c_j, \ell_3\}$.

The properties of f that we need here, which follow immediately from the proof of [33], are as follows.

1. $f(\phi)$ does not have a vertex cover of size less than $n + 2m$.

2. If W is a vertex cover of size $n + 2m$, then $W \cap \{x_i, \overline{x_i} \mid 1 \leq i \leq n\}$ corresponds to a satisfying assignment of ϕ, in the sense that $\|W \cap \{x_i, \overline{x_i}\}\| = 1$ for $1 \leq i \leq n$, and $\alpha_1 \cdots \alpha_n$ defined as $\alpha_i = 1$ if and only if $x_i \in W$ is a satisfying assignment for ϕ.

3. If α is a satisfying assignment for ϕ, then there is a vertex cover of size $n + 2m$ such that $W \cap \{x_i, \overline{x_i} \mid 1 \leq i \leq n\}$ corresponds to this assignment, i.e., the set $\{x_i \mid \alpha_i = 1\} \cup \{\overline{x_i} \mid \alpha_i = 0\}$ can be extended to a vertex cover of size $n + 2m$ by adding $2m$ vertices from $\{a_i, b_i, c_i \mid 1 \leq i \leq m\}$.

Now set the weights of the vertices as follows: $\omega(a_i) = \omega(b_i) = \omega(c_i) = 2^n$, $\omega(x_i) = 2^n + 2^{n-i}$, and $\omega(\overline{x_i}) = 2^n$. Note that for W a set of vertices, $\|W\| = \lfloor \omega(W)/2^n \rfloor$. In particular, a minimum-weight vertex cover will also have minimum size.

We will now show that n-bit string α is the smallest satisfying assignment of ϕ if and only if $f(\phi)$ has a minimum-weight vertex cover of weight $(n + 2m)2^n + \alpha$. (Recall that we interpret n-bit string α as a binary number between 0 and $2^n - 1$.)

If α is a satisfying assignment for ϕ, then by property 3 above the set $\{x_i \mid \alpha_i = 1\} \cup \{\overline{x_i} \mid \alpha_i = 0\}$ can be extended to a vertex cover of size $n + 2m$ by adding $2m$ vertices

from $\{a_i, b_i, c_i \mid 1 \leq i \leq m\}$. The weight of this vertex cover is $(n+2m)2^n + \alpha$. Since the minimum size of a vertex cover is $n + 2m$, the weight of a minimum-weight vertex cover is $(n+2m)2^n + \beta$, for some β such that $0 \leq \beta < 2^n$. Such a β corresponds to a satisfying assignment by property 2 above.

Since ψ is also a satisfiable 3cnf formula over $x_1, \ldots, x_n$ with m clauses, it also holds that n-bit string α is the smallest satisfying assignment of ψ if and only if $f(\psi)$ has a minimum weight vertex cover of weight $(n+2m)2^n + \alpha$.

This then implies that ϕ's minimal satisfying assignment is $\leq \psi$'s minimal satisfying assignment if and only if the weight of $f(\phi)$'s minimum-weight vertex cover is less than or equal to the weight of $f(\psi)$'s minimum-weight vertex cover. Note that $\omega(f(\phi)) = \omega(f(\psi))$. This completes the reduction from MINSATASG$_\leq$ to Min-Weight-Vertex-Cover-Compare. ❑

A.8 Proof of Theorem 15: Kemeny-Winner for four voters is Θ_2^p-complete.

We combine the constructions from [31] and [13], to reduce Feedback-Arc-Set-Member to Kemeny-Winner for four voters.

Given an irreflexive and antisymmetric digraph $G = (V, A)$ and vertex $v \in V$, we first compute graph G' as done in [13]. $G' = (V', A')$ such that $V' = V \cup A$ and $A' = \{(v, (v, w)), ((v, w), w) \mid (v, w) \in A\}$. Note that G has a feedback arc set of size k if and only if G' has a feedback arc set of size k. In addition, and important for our proof, (G, v) is in Feedback-Arc-Set-Member if and only if (G', v) is in Feedback-Arc-Set-Member. Also note that G' is irreflexive and antisymmetric. Now apply the reduction from Feedback-Arc-Set-Member to Kemeny-Winner from [31] on (G', v). This reduction outputs an election $g(G')$ that corresponds to G' and (G', v) is in Feedback-Arc-Set-Member if and only if $(g(G'), v)$ is in Kemeny-Winner. The construction from [13] shows that we need only four voters in $g(G')$.

A.9 Proof of Theorem 16: Kemeny-Winner for two voters who vote a partial order and Kemeny-Winner for one voter who votes a binary relation are Θ_2^p-complete.

The definition of Kemeny winner can naturally be extended to voters voting any binary relation, in the following way.

A candidate p is a Kemeny winner if p is ranked first in a Kemeny consensus. A Kemeny consensus is a linear order $>$ over the candidates that minimizes the Kendall's tau distance to V, i.e., that minimizes

$$\sum_{a,b \in C, a > b} \|\{v \in V \mid b >_v a\}\| + \|\{v \in V \mid b \sim_v a\}\|/2,$$

where $>_v$ is the preference relation of voter v.

The proof of Θ_2^p-completeness proceeds as in the proof of Theorem 15. All that is needed is to observe that

1. $g(G')$ needs only two voters v and w if voters can vote with a partial order: $x >_v (x, y)$ for all $(x, y) \in A$, $(x, y) >_w y$ for all $(x, y) \in A$, and v and w are indifferent for all other pairs of candidates.

2. $g(G')$ needs only one voter v if voters can voter any binary relation: Let $x >_v y$ if and only if $(x, y) \in A'$.

A.10 Dodgson and Young Elections

In Section 6 we consider the complexity of winner determination for Kemeny elections. In this section we consider the complexity of winner determination for Dodgson [12] and for Young [45], which as with Kemeny, is Θ_2^p-complete for standard representation [24, 42]. The Dodgson score of a candidate c is the number of swaps between adjacent candidates needed in the preferences of the voters in the election to make c a Condorcet winner (a candidate that beats each other candidate pairwise). The Dodgson winner(s) of an election are the candidate(s) with minimum score. The Young score of a candidate c is the size of a largest subcollection of voters for which c is a weak Condorcet winner (a candidate that beats-or-ties each other candidate pairwise). The Young winner(s) of an election are the candidate(s) with maximum score.

For Young and Dodgson for weighted elections we can consider both the case where we charge by the voter (e.g., for Dodgson the score is based on the number of swaps regardless of the weight of the voters where candidates are swapped) and the case where we charge by the voter-weight (e.g., for Dodgson a swap in the vote of a weight ω voter counts as ω swaps). Clearly, by brute-force, we have the following.[7]

Theorem 22 *For a fixed number of voters, Young-Weighted-Winner and Dodgson-Weighted-Winner are in* P, *both in the model where we charge by the voter and in the model where we charge by the voter-weight.*

The high-multiplicity case for Young is also in P, but this seems to need much more powerful machinery.

Theorem 23 *For a fixed number of votes, Young-High-Multiplicity-Winner is in* P.

Proof. Given an election (C, V) in high-multiplicity representation and a candidate p, we can determine if p is a Young winner in the following way.

First, consider the following ILP that determines if the Young score of a candidate c is $\geq$ a given k. Let $v_1, \ldots, v_\ell$ be the distinct votes in V. The ILP has variables $k_1, \ldots, k_\ell$ that correspond to the number of voters "kept" for each vote and the following constraints.

$$\forall i, 1 \leq i \leq \ell, \ \ k_i \geq 0 \tag{1}$$

$$\forall i, 1 \leq i \leq \ell, \ \ k_i \leq \kappa(v_i) \tag{2}$$

$$\sum_{i=1}^{\ell} k_i \geq k \tag{3}$$

$$\forall a \neq c, \sum_{1 \leq i \leq \ell, c >_i a} k_i \geq \sum_{1 \leq i \leq \ell, a >_i c} k_i \tag{4}$$

Constraints 1, 2, and 3 ensure that the number of voters kept for each vote is consistent with the given election and that the total number of voters kept is at least k. Constraint 4 ensures that c is a weak Condorcet winner.

We can then determine if a candidate p is a Young winner by separately computing the Young score of each candidate using the ILP described above (which can be solved in polynomial time by Lenstra [35]) and binary search, and p is a winner if and only if p has maximum Young score. ❑

---

[7]Note that to compute the Dodgson score for a given candidate c it suffices to consider only swaps that move c up. So for each voter we only have to look at a linear number of possibilities (depending only on the position of c).

This case of a fixed number of votes is a natural case to consider, since this models scenarios where all voters aligning with the same political party or group may vote the same. This is reminiscent to related work where the preferences of the voters are single-peaked (and so only a restricted number of the total possible votes occur) (see, e.g., [17]), but instead we restrict the number of votes in an election without assuming structure among them.

For the general weighted cases where we charge by the voter, it is easy to see that the scores are polynomially bounded, and so we can state the following result.

Theorem 24 *Young-Weighted-Winner and Dodgson-Weighted-Winner in the model where we charge by the voter are Θ_2^p-complete.*

It is easy to see that Young-Weighted-Winner and Dodgson-Weighted-Winner in the model where we charge by voter-weight as well as Young-High-Multiplicity-Winner and Dodgson-High-Multiplicity-Winner are in Δ_2^p, by using binary search to find the scores of all candidates.

We conjecture that, as in the case for Kemeny, we can "lift" the Θ_2^p-hardness proofs for Dodgson and Young to Δ_2^p-hardness proofs for the weighted cases. We do not think that this "lifting" will work for the high-multiplicity cases and leave the exact complexity of these problems, which we know to be Θ_2^p-hard and in Δ_2^p, as an open question.

Zack Fitzsimmons
Department of Mathematics and Computer Science
College of the Holy Cross
Worcester, MA, 01610 USA
Email: zfitzsim@holycross.edu

Edith Hemaspaandra
Department of Computer Science
Rochester Institute of Technology
Rochester, NY, 14623 USA
Email: eh@cs.rit.edu