

Single-Peakedness and Total Unimodularity: New Polynomial-Time Algorithms for Multi-Winner Elections

Dominik Peters

Department of Computer Science
University of Oxford, UK
dominik.peters@cs.ox.ac.uk

Abstract

The winner determination problems of many attractive multi-winner voting rules are NP-complete. However, they often admit polynomial-time algorithms when restricting inputs to be single-peaked. Commonly, such algorithms employ dynamic programming along the underlying axis. We introduce a new technique: carefully chosen integer linear programming (IP) formulations for certain voting problems admit an LP relaxation which is totally unimodular if preferences are single-peaked, and which thus admits an integral optimal solution. This technique gives efficient algorithms for finding optimal committees under Proportional Approval Voting (PAV) and the Chamberlin–Courant rule with single-peaked preferences, as well as for certain OWA-based rules. For PAV, this is the first technique able to efficiently find an optimal committee when preferences are single-peaked. An advantage of our approach is that no special-purpose algorithm needs to be used to exploit structure in the input preferences: any standard IP solver will terminate in the first iteration if the input is single-peaked, and will continue to work otherwise.

1 Introduction

In a departure from classical voting theory, a growing literature from computational social choice has recently studied *multi-winner* voting rules (Faliszewski et al. 2017b). Given diverse preferences of a collection of agents, instead of identifying a single best alternative, we are aiming for a (fixed-size) *set* of alternatives that jointly are able to represent the preferences of the agents best. Such procedures are useful in a wide variety of circumstances: obvious examples include the election of a parliament, or of a committee representing the interests of members of an organisation. Other applications can be found in group recommendation systems, or for making decisions about which products or services to offer: Which courses should be offered at a university? Which movies should be presented on an airline entertainment system?

Several attractive rules for such tasks have been designed by researchers in political science (e.g., Chamberlin and Courant 1983, Monroe 1995) and more recently by computer scientists (Faliszewski et al. 2016a; 2016b; Skowron, Faliszewski, and Lang 2015). Many of these rules are defined in terms of some objective function: a winning committee is a

set of k candidates that maximises this objective. Unsurprisingly, then, the winner determination problems of such rules are typically NP-hard (Lu and Boutilier 2011). To tackle the complexity of these problems, approximation algorithms (Skowron, Faliszewski, and Slinko 2015) and fixed-parameter tractability approaches (Betzler, Slinko, and Uhlmann 2013; Bredereck et al. 2015) have been developed for these problems, and integer programming formulations have also been designed for them (Pothhoff and Brams 1998).

Another approach seeks to exploit underlying structure in the preferences reported by the agents (Elkind, Lackner, and Peters 2017). A particularly influential preference restriction is the notion of *single-peaked preferences*, due to Black (1948) and Arrow (1950). In this model, the alternative space has a one-dimensional structure: alternatives are ordered on a left-to-right *axis*; and agents’ preferences are monotonically decreasing as we move further away from their *peak* (most-preferred alternative). In particular, we can expect preferences to be structured this way when voting over the value of a numerical quantity (such as a tax rate). While single-peaked preferences were first employed to escape impossibility results in social choice theory (Moulin 1991), it also yields positive algorithmic results: Notably, Betzler et al. (2013) showed that Chamberlin–Courant’s (1983) multi-winner rule is efficiently computable for single-peaked preferences. This result can be extended to some other multi-winner voting rules (Elkind and Ismaili 2015), and to preferences that are single-crossing, or single-peaked on trees (Skowron et al. 2013; Elkind, Faliszewski, and Skowron 2014; Peters and Elkind 2016).

The Chamberlin–Courant rule usually takes as input preferences specified by linear orders, i.e., rankings of the candidates. An alternative way of specifying preferences is by *approval ballots*, where voters submit a *set* of approved candidates. Recent work has uncovered a rich variety of multi-winner voting rules in this framework (Kilgour 2010). A particularly well-studied rule is *Proportional Approval Voting* (PAV) due to Thiele (1895). PAV has particularly nice axiomatic properties in terms of proportional representation (Aziz et al. 2017), and has recently been axiomatically characterised as the only consistent extension of d’Hondt’s proportionality criterion for party list elections to the general setting of approval votes (Lackner and Skowron 2017, see also Brill, Laslier, and Skowron 2017). Like for Chamberlin–Courant,

the winner determination problem of PAV is NP-complete (Aziz et al. 2015). For PAV, however, tractability results for cases of structured preferences have proven somewhat elusive. Elkind and Lackner (2015) have studied this question in depth, but for single-peaked preferences¹ they were only able to establish that the problem lies in FPT for some natural parameters, and that it becomes polynomial-time solvable for a very restrictive subclass of single-peaked approval preferences. The difficulty is that standard approaches based on dynamic programming seem to not be powerful enough to capture the *multirepresentation* nature of PAV: Under the PAV rule, voters derive (implicit) utility from potentially *many* committee members, whereas in the case of Chamberlin–Courant, they are only represented by one member of the committee. This difficulty led Elkind and Lackner (2015) to conjecture that evaluating PAV is NP-complete even for single-peaked preferences.

In this paper, we introduce a new technique that allows evaluating a large class of multi-winner rules in polynomial time if the input preferences are single-peaked. This class includes PAV; thus we disconfirm Elkind and Lackner’s conjecture and show that PAV does become tractable with single-peaked preferences. Other rules in this class include Chamberlin–Courant, the *t*-Borda rules (Faliszewski et al. 2017a), and most OWA-based multi-winner rules (Skowron, Faliszewski, and Lang 2015).

Our technique is based on designing integer linear programming formulations for these rules, and proving that these formulations are *totally unimodular* in case the input preferences are single-peaked. Total unimodularity is a condition on the determinants of the matrix of coefficients appearing in the constraints of the integer program. Since totally unimodular integer programs are optimally solved by their linear programming relaxations, these rules are polynomial-time computable with such inputs. In fact, since all standard IP solvers first solve the LP relaxation, they will terminate with the correct answer in their first iteration. If the instance is not single-peaked, the IP solver might enter further iterations while solving: our formulations are correct whether or not the input is single-peaked. This makes it easy to implement our algorithms, and stands in contrast to algorithms based on dynamic programming, which are specialised to work only under the assumption that the input is single-peaked.

Another difference to specialised algorithms (such as the one due to Betzler et al. 2013) is that we do not need to know an underlying single-peaked axis of the input profile in order to solve the integer program in polynomial time. Dynamic programs that work along the single-peaked axis need a preprocessing stage, where we run a *recognition algorithm* which constructs a single-peaked axis if one exists. Such algorithms have been designed by Escoffier, Lang, and Öztürk (2008), Bartholdi III, Tovey, and Trick (1989), and Doignon and Falmagne (1994), among others. Our algorithms skip this step, again making implementation easier.

¹For approval preferences, single-peakedness requires that there is a left-to-right ordering of the alternatives so that each approval set forms an interval of this ordering (Faliszewski et al. 2011; Elkind and Lackner 2015).

2 Preliminaries

We write $[n] = \{1, \dots, n\}$.

Total Unimodularity A matrix $A = (a_{ij})_{ij} \in \mathbb{Z}^{m \times n}$ with $a_{ij} \in \{-1, 0, 1\}$ is called *totally unimodular* if every square submatrix B of A has $\det B \in \{-1, 0, 1\}$. The following results are well-known. Proofs and much more about their theory can be found in the textbook by Schrijver (1998).

Theorem 1. *Suppose $A \in \mathbb{Z}^{m \times n}$ is a totally unimodular matrix, $b \in \mathbb{Z}^m$ is an integral vector of right-hand sides, and $c \in \mathbb{Q}^n$ is an objective vector. Then the linear program*

$$\max c^T x \text{ subject to } Ax \leq b \quad (P)$$

has an integral optimum solution, which is a vertex of the polyhedron $\{x : Ax \leq b\}$. Thus, the integer linear program

$$\max c^T x \text{ subject to } Ax \leq b, x \in \mathbb{Z}^n \quad (IP)$$

is solved optimally by its linear programming relaxation (P).

An optimum solution to (IP) can be found in polynomial time. We will now state some elementary results about totally unimodular matrices.

Proposition 2. *If A is totally unimodular, then so is*

- (1) *its transpose A^T ,*
- (2) *the matrix $[A \mid -A]$ obtained from A by appending the negated columns of A ,*
- (3) *the matrix $[A \mid I]$ where I is the identity matrix,*
- (4) *any matrix obtained from A through permuting or deleting rows or columns.*

In particular, from (3) and (4) it follows that appending a unit column $(0, \dots, 1, \dots, 0)^T$ will not destroy total unimodularity. Further, using these transformations, we can see that Theorem 1 remains true even if we add to (P) constraints giving lower and upper bounds to some variables, if we replace some of the inequality constraints by equality constraints, or change the direction of an inequality.

A binary matrix $A = (a_{ij}) \in \{0, 1\}^{m \times n}$ has the *strong consecutive ones property* if the 1-entries of each row form a contiguous block, as in the example on the right. A binary matrix has the *consecutive ones property* if its columns can be permuted so that the resulting matrix has the strong consecutive ones property. The key result that will allow us to connect single-peaked preferences to total unimodularity is as follows:

Proposition 3. *Every binary matrix with the consecutive ones property is totally unimodular.*

We remark that by a celebrated result of Seymour (1980), it is possible to decide in polynomial time whether a given matrix is totally unimodular, though we do not use this fact.

Single-Peaked Preferences Let A be a finite set of *alternatives*, or *candidates*, and let $m = |A|$. A *weak order*, or *preference relation*, is a binary relation \succsim over A that is complete and transitive. We write \succ and \sim for the strict and indifference parts of \succsim . A *linear order* is a weak order that, in addition, is antisymmetric, so that $x \sim y$ only if $x = y$.

Every preference relation \succsim induces a partition of A into indifference classes A_1, \dots, A_r so that $A_1 \succ A_2 \succ \dots \succ A_r$ and $x \sim y$ for all $x, y \in A_t$. We will say that an alternative $a \in A_t$ has *rank* t in the ordering \succsim and write $\text{rank}(a) = t$; thus the alternatives of rank 1 are the most-preferred alternatives under \succsim . Finally, we say that any set of the form $\{x \in A : \text{rank}(x) \geq t\}$ is a *top-initial segment* of \succsim .

Let \triangleleft be (the strict part of) a linear order over A ; we call \triangleleft an *axis*. A linear order \succsim_i with most-preferred alternative p (the *peak*) is *single-peaked with respect to* \triangleleft if for every pair of candidates $a, b \in A$ with $p \triangleleft b \triangleleft a$ or $a \triangleleft b \triangleleft p$ it holds that $b \succsim_i a$. For example, if the alternatives in A correspond to different proposed levels of a tax, and the numbers in A are ordered by \triangleleft in increasing order, then it is sensible to expect voters' preferences over A to be single-peaked with respect to \triangleleft .

We need an equivalent definition of single-peakedness:

Proposition 4. *A linear order \succ is single-peaked with respect to \triangleleft if and only if all top-initial segments of \succ form an interval of \triangleleft .*

The proof is straightforward, see Figure 1 for an illustration. An advantage of this alternative definition is that it allows us to generalise single-peakedness to *weak* orders. Thus, we will define a weak order \succsim to be *single-peaked* with respect to \triangleleft exactly if all top-initial segments of \succsim form an interval of \triangleleft . This concept is often known as ‘possibly single-peaked’ (Lackner 2014) because it is equivalent to requiring that the ties in the weak order can be broken in such a way that the resulting linear order is single-peaked.

A *profile* $P = (\succsim_1, \dots, \succsim_n)$ over a set of alternatives A is a list of weak orders over A . Each of the orders represents the preferences of a *voter*; we write $N = [n]$ for the set of voters. The profile will be called *single-peaked* if there exists some axis \triangleleft over A so that each order \succsim_i in P is single-peaked with respect to \triangleleft .

Based on Proposition 4, we can characterise single-peakedness of a profile in terms of the consecutive ones property of a certain matrix. Indeed, a profile is single-peaked if and only if the following matrix M_{SP}^P has the consecutive ones property: take one column for each alternative, and one row for each top-initial segment S of each voter's preference relation; the row is just the incidence vector of S . Figure 1 shows an example. Clearly, M_{SP}^P has the consecutive ones property if and only if the condition of Proposition 4 is satisfied. This construction is due to Bartholdi III and Trick (1986), see also Fitzsimmons (2015). Note that, because it is possible to decide whether a matrix has the consecutive ones property in linear time (Booth and Lueker 1976), this yields an $O(m^2n)$ algorithm for checking whether a given profile is single-peaked; however, there are faster, more direct algorithms for this task (Escoffier, Lang, and Öztürk 2008; Doignon and Falmagne 1994).

Dichotomous Preferences A weak order \succsim is *dichotomous* if it partitions A into at most two indifference classes $A_1 \succ A_2$. The alternatives in A_1 are said to be *approved*




Figure 1: Translation of single-peakedness of the profile P into the consecutive ones property of the matrix M_{SP}^P : each row corresponds to a top-initial segment.

by the voter \succsim . On dichotomous preferences, the notion of single-peakedness essentially coincides with the consecutive ones property (Faliszewski et al. 2011): there needs to be an ordering \triangleleft of the alternative so that each approval set A_1 is an interval of \triangleleft . Thus, Elkind and Lackner (2015) use the name *Candidate Interval (CI)* for single-peakedness in this context.

3 Proportional Approval Voting

In this section, we will consider Proportional Approval Voting (PAV), a multi-winner voting rule defined for dichotomous (approval) preferences. A naïve way to form a committee would be to select the k alternatives with highest approval score, but this method tends to ignore minority candidates, and so is not representative (Aziz et al. 2017). PAV attempts to fix this issue: it is based on maximising a sum of voters' utilities, where a voter i 's utility is a concave function of the number of candidates in the committee that i approves of. The rule was first proposed by Thiele (1895). In the general case, a winning committee under PAV is NP-hard to compute (Aziz et al. 2015), even if each voter approves only 2 candidates and each candidate is approved by only 3 voters.

Let us define PAV formally. Each voter i submits a set $v_i \subseteq C$ of *approved* candidates (or, equivalently, a dichotomous weak order with $v_i \succ_i C \setminus v_i$). We aim to find a good committee $W \subseteq C$ of size $|W| = k$. The intuition behind Proportional Approval Voting (PAV) is that voters are happier with committees that contain more of their approved candidates, but that there are decreasing marginal returns to extra approved candidates in the committee. Concretely, each voter obtains a ‘utility’ of 1 for the first approved candidate in W , of $\frac{1}{2}$ for the second, of $\frac{1}{3}$ for the third, and so on. The objective value of a committee $W \subseteq C$ is thus

$$\sum_{i \in N} 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{|W \cap v_i|}.$$

The choice of harmonic numbers might seem arbitrary, and one can more generally define a rule α -PAV where $\alpha \in \mathbb{R}_+^k$ is a non-increasing scoring vector (so $\alpha_i \geq \alpha_j$ when $i \geq j$). This rule gives W the objective value $\sum_{i \in N} \sum_{\ell=1}^{|W \cap v_i|} \alpha_\ell$. Then PAV is just $(1, \frac{1}{2}, \frac{1}{3}, \dots, \frac{1}{k})$ -PAV. However, the choice

of harmonic numbers is the only vector α that lets α -PAV satisfy an axiom called ‘extended justified representation’ (Aziz et al. 2017), making this a natural choice after all.

For single-peaked preferences, Elkind and Lackner (2015) presented algorithms for finding an optimal PAV committee running in time $O(2^s nm)$ and $\text{poly}(d, m, n, k^d)$, where s is the maximum cardinality of the approval sets, and d is the maximum number of voters that approve a given candidate. These algorithms, based on dynamic programming, are efficient if these parameters s or d are small. They also showed, by extending an algorithm of Betzler, Slinko, and Uhlmann (2013), that α -PAV is easy for scoring vectors $\alpha = (\alpha_1, \dots, \alpha_r, 0, 0, \dots)$ that are ‘truncated’ in that they only have a constant number of non-zero values at the beginning. Finally, if one imposes further (restrictive) assumptions on the structure of the input preferences, they also found polynomial-time results. However, none of these algorithms could be extended to cover the general case for solving PAV and α -PAV for non-truncated scoring vectors. Our method, being very different from dynamic programming, can solve PAV in polynomial time if preferences are single-peaked.

Let us now give and analyse our IP formulation for PAV. This formulation has one binary variable y_c for each candidate $c \in C$, indicating whether candidate c is part of the committee. Constraint (2) requires that the committee contains exactly k candidates. The binary variables $x_{i,\ell}$ indicate whether voter $i \in N$ approves of at least ℓ candidates in the committee; this interpretation is implemented by the constraints (3).

$$\text{maximise } \sum_{i \in N} \sum_{\ell \in [k]} \alpha_\ell \cdot x_{i,\ell} \quad (\text{PAV-IP})$$

$$\text{subject to } \sum_{c \in C} y_c = k \quad (2)$$

$$\sum_{\ell \in [k]} x_{i,\ell} \leq \sum_{i \text{ approves } c} y_c \quad \text{for } i \in N \quad (3)$$

$$x_{i,\ell} \in \{0, 1\} \quad \text{for } i \in N, \ell \in [k] \quad (4)$$

$$y_c \in \{0, 1\} \quad \text{for } c \in C$$

To help intuition, let us explicitly write down an example for PAV with harmonic weights, with 4 candidates a, b, c, d , target committee size $k = 2$, and two voters, where voter 1 approves $\{a, b, c\}$, and voter 2 approves $\{c, d\}$. This profile is single-peaked on the axis $a \triangleleft b \triangleleft c \triangleleft d$.

$$\text{maximise } (x_{1,1} + \frac{1}{2}x_{1,2}) + (x_{2,1} + \frac{1}{2}x_{2,2}) \quad (\text{PAV-IP}')$$

$$\text{subject to } y_a + y_b + y_c + y_d = 2 \quad (2)$$

$$x_{1,1} + x_{1,2} \leq y_a + y_b + y_c \quad (3)$$

$$x_{2,1} + x_{2,2} \leq y_c + y_d \quad (3)$$

all variables binary

An optimal committee for this profile is, for example, $W = \{c, d\}$. An optimum feasible solution would then set $(x_{1,1}, x_{1,2}, x_{2,1}, x_{2,2}) = (1, 0, 1, 1)$. We must set $x_{1,2} = 0$, because voter 1 only approves of a single candidate in the

committee. The resulting objective value is $1 + 0 + 1 + \frac{1}{2} = 2.5$, which is precisely the PAV value of the committee $\{c, d\}$.

We now argue formally that our formulation (PAV-ILP) captures the winner determination problem of α -PAV.

Proposition 5. *Program (PAV-IP) correctly computes an optimal committee according to α -PAV.*

Proof. In any feasible solution of (PAV-IP), the variables y_c encode a committee of size k . Fix such a committee $W = \{c \in C : y_c = 1\}$. We show that the optimum objective value of a feasible solution with these choices for the y_c -variables is the α -PAV-value of this committee.

Since $\alpha \geq 0$, in optimum, as many $x_{i,\ell}$ will be set to 1 as constraint (3) allows. Thus, for each i , exactly $|W \cap v_i|$ many variables $x_{i,\ell}$ will be set to 1. Since w is non-increasing, wlog, in optimum, these will be variables $x_{i,1}, \dots, x_{i,|W \cap v_i|}$. Then the objective value equals the α -PAV-value of W . \square

The constraint matrix of our example (PAV-ILP') is

$$A_{\text{PAV-IP}'} = \begin{bmatrix} x_{1,1} & x_{1,2} & x_{2,1} & x_{2,2} & y_a & y_b & y_c & y_d \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ -1 & -1 & 0 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & -1 & -1 & 0 & 0 & 1 & 1 \end{bmatrix}$$

Here we have rewritten constraint (3) as $-x_{1,1} - x_{1,2} + y_a + y_b + y_c \geq 0$. Note that, because the input profile is single-peaked, the submatrix corresponding to the y_c -variables has the consecutive ones property: it is just M_{SP}^P with an all-1s row added. Thus, this submatrix is totally unimodular. The whole constraint matrix is the obtained by adding negative unit columns corresponding to the $x_{i,\ell}$ -variables. By Proposition 2, adding these columns yields a matrix that is still totally unimodular. Generalising this argument to arbitrary profiles, we get the following.

Proposition 6. *The constraint matrix of (PAV-IP) is totally unimodular when the input preferences are single-peaked.*

Proof. We will use the manipulations allowed by Proposition 2 liberally. In particular, it allows us to ignore the constraints $0 \leq x_{i,\ell}, y_c \leq 1$, and to ignore the difference between equality and inequality constraints. Thus, after permuting columns corresponding to variables $x_{i,\ell}$ so that they are sorted by i , the constraint matrix of (PAV-IP) is

$$A_{\text{PAV-IP}} = \begin{bmatrix} \overbrace{-I_n \ \dots \ -I_n}^{k \text{ times}} & M_{\text{SP}}^P \\ \mathbf{0}_n & \mathbf{1}_m \end{bmatrix},$$

where I_n is the $n \times n$ identity matrix, and $\mathbf{1}_m$ is an all-1s row vector with m entries.

If preferences P are single-peaked, then M_{SP}^P has the consecutive ones property, and this is also true after appending a row with all-1s. Thus, $\begin{bmatrix} M_{\text{SP}}^P \\ \mathbf{1}_m \end{bmatrix}$ is totally unimodular. Applying Proposition 2 repeatedly to append negations of unit columns, we obtain $A_{\text{PAV-IP}}$, which is thus totally unimodular. \square

Using Theorem 1, that integer programming is easy for totally unimodular programs, we obtain our desired result.

Theorem 7. *α -PAV can be computed in polynomial time for single-peaked approval preferences.*

An interesting feature of (PAV-IP) is that the integrality constraints (4) on the variables $x_{i,\ell}$ can be relaxed to just $0 \leq x_{i,\ell} \leq 1$; this does not change the objective value of the optimum solution. This is because in optimum, the quantity $\sum_{\ell} x_{i,\ell}$ is integral by (3), and it never pays to have one of the $x_{i,\ell}$ to be fractional, because this fractional amount can be shifted to $x_{i,\ell'}$ with $\ell' < \ell$ to get a (weakly) higher value.² This observation might tell us that solving (PAV-IP) is relatively easy as it is “close” to being an LP, and this also seems to be true in practice. On the other hand, it is not necessarily beneficial to relax the integrality constraints (4) when passing (PAV-IP) to an IP solver: the presence of integrality constraints might nudge the solver to keep numerical integrality gaps smaller. Of course, the point of this paper is to give another reason why (PAV-IP) is “close” to being an LP, namely when preferences are single-peaked.

4 Chamberlin–Courant’s Rule

Now we leave the domain of dichotomous preferences, and consider the full generality of profiles of weak orders. The definition of Chamberlin–Courant’s rule (1983) is based on the notion of having a *representative* in the elected committee: each voter is represented by their favourite candidate in the committee, and voters are happier with more preferred representatives. Let $\mathbf{w} \in \mathbb{N}^m$ be a (non-increasing) scoring vector; the standard choice for \mathbf{w} are Borda scores: $\mathbf{w} = (m, m-1, \dots, 2, 1)$. Let $P = (\succ_1, \dots, \succ_n)$ be a profile. Then the objective value of a committee $W \subseteq C$ according to Chamberlin–Courant’s rule is

$$\sum_{i \in N} \max\{w_{\text{rank}_i(c)} : c \in W\}.$$

Chamberlin–Courant now returns any committee $W \subseteq C$ with $|W| = k$ that maximises this objective.

Betzler et al. (2013) gave a dynamic programming algorithm solving this problem if preferences are single-peaked. We now give an alternative algorithm that has the advantage of being extensible to OWA-based rules, as we will see in the next section.

Chamberlin–Courant can be seen as a (non-metric) facility location problem: each candidate $c \in C$ is a potential facility location, we are allowed to open exactly k facilities, and the distance between customers and facilities are determined through \mathbf{w} . There is a standard integer programming formulation for this problem using binary variables y_c , denoting whether c will be opened or not, and variables $x_{i,c}$, denoting whether facility c will service voter i . However, this formulation is not totally unimodular, since it encodes the preferences in its objective function rather than in its constraints.

To construct our alternative IP formulation, we need another definition of the Chamberlin–Courant objective function based on maximising a number of *points*. For expositional simplicity, let us take \mathbf{w} to be Borda scores; other scoring rules can be obtained by weighting the points. Here is another way of thinking about the objective value as defined above: each voter i can earn a point for each *rank*

in i ’s preference order: for every rank $r \in [m]$, i earns the point $x_{i,r}$ if there is a committee member $c \in W$ with $\text{rank}_i(c) \geq r$. Then the number of points obtained in total equals the objective value: if i ’s favourite committee member is in rank r , then i will earn precisely $w_r = m - r + 1$ points, namely the points $x_{i,r}, x_{i,r+1}, \dots, x_{i,m}$. This view suggests the following integer programming formulation, where we put $w'_r = w_r - w_{r-1}$ and $w'_1 = w_1$.

$$\text{maximise } \sum_{i \in N} \sum_{r \in [m]} w'_r \cdot x_{i,r} \quad (\text{CC-IP})$$

$$\text{subject to } \sum_{c \in C} y_c = k \quad (2)$$

$$x_{i,r} \leq \sum_{c : \text{rank}(c) \geq r} y_c \quad \text{for } i \in N, r \in [m] \quad (3)$$

$$x_{i,r} \in \{0, 1\} \quad \text{for } i \in N, r \in [m]$$

$$y_c \in \{0, 1\} \quad \text{for } c \in C$$

Again let us consider an illustrative instantiation of this program, based on Borda scores, committee size $k = 2$, and the profile from Figure 1, with $b \succ_1 c \succ_1 a \succ_1 d$ and $c \succ_2 d \succ_2 b \succ_2 a$. The resulting program is as follows.

$$\text{maximise } (x_{1,1} + x_{1,2} + x_{1,3} + x_{1,4}) \\ + (x_{2,1} + x_{2,2} + x_{2,3} + x_{2,4})$$

$$\text{subject to } y_a + y_b + y_c + y_d = 2 \quad (2)$$

$$x_{1,1} \leq y_b \quad (3)$$

$$x_{1,2} \leq y_b + y_c \quad (3)$$

$$x_{1,3} \leq y_a + y_b + y_c \quad (3)$$

$$x_{1,4} \leq y_a + y_b + y_c + y_d \quad (3)$$

$$x_{2,1} \leq y_c \quad (3)$$

$$x_{2,2} \leq y_c + y_d \quad (3)$$

$$x_{2,3} \leq y_b + y_c + y_d \quad (3)$$

$$x_{2,4} \leq y_a + y_b + y_c + y_d \quad (3)$$

all variables binary

To understand the program, it is convenient to rephrase it in logical language: we can interpret “ \leq ” as “only if”, and “+” as “or” in constraints (3). For example, the second constraint of type (3) can be read as $x_{1,2}$ only if the committee contains either b or c .

Clearly, for this profile, the committee $\{b, d\}$ is optimal, which allows us to set all the $x_{i,\ell}$ -variables to 1, yielding objective value 8. If we were to take the committee size $k = 1$ (in which case Chamberlin–Courant is the same as the Borda count), then $\{c\}$ is the optimum committee. That committee forces us to set $x_{1,1} = 0$, but allows us to set all the others to 1, yielding objective value 7, which is c ’s Borda score. For general profiles, we can see correctness of the formulation as follows.

Proposition 8. *Program (CC-IP) correctly computes an optimal committee according to \mathbf{w} -Chamberlin–Courant.*

Proof. In any feasible solution of (CC-IP), the variables y_c encode a committee of size k . Fix such a committee $W =$

²A similar property is used by Bredereck et al. (2015, Thm 1).

$\{c \in C : y_c = 1\}$. We show that the optimum objective value of a feasible solution with these choices for the y_c -variables is the objective value of this committee according to **w-CC**.

Since $\mathbf{w}' \geq 0$, in optimum, every $x_{i,r}$ will be set to 1 if constraint (3) allows this. This is the case iff there is a committee member $c \in W$ with $\text{rank}_i(c) \geq r$, i.e., iff the ‘point’ $x_{i,r}$ is earned as described above. Thus, the objectives of (CC-IP) and Chamberlin–Courant coincide. \square

Similarly to the case of PAV, the constraint matrix of our example instance can be obtained by taking the matrix in Figure 1, appending a row of all-1s for constraint (2), and then adding negative unit columns. By Proposition 2, the resulting matrix is totally unimodular, because the matrix M_{SP}^P that we started out with has the consecutive ones property. More generally, the proof is as follows.

Proposition 9. *The constraint matrix of (CC-IP) is totally unimodular when the input preferences are single-peaked.*

Proof. After similar simplification as in Proposition 6 we see that the constraint matrix of (CC-IP) is

$$A_{\text{CC-IP}} = \begin{bmatrix} -I_{nm} & M_{\text{SP}}^P \\ 0 & \mathbf{1}_m \end{bmatrix}.$$

Again, if preferences P are single-peaked, then M_{SP}^P has the consecutive ones property, and this is also true after appending a row with all-1s. Thus, $\begin{bmatrix} M_{\text{SP}}^P \\ \mathbf{1}_m \end{bmatrix}$ is totally unimodular. Applying Proposition 2 repeatedly to append unit columns, we obtain $A_{\text{CC-IP}}$, which is thus totally unimodular. \square

Theorem 10. *Chamberlin–Courant with score vector \mathbf{w} can be solved in polynomial time for single-peaked preferences.*

5 OWA-based Rules

When using Chamberlin–Courant, each voter is represented by exactly one committee member. In many applications, we may instead seek *multirepresentation* (Skowron, Faliszewski, and Lang 2015): e.g., you might watch several of the movies offered by an inflight-entertainment system, and thus should be represented by several committee members. In such scenarios, Chamberlin–Courant might design a suboptimal committee; Skowron et al. (2015) introduce *OWA-based* multi-winner rules as a more flexible alternative (see also Faliszewski et al. 2016b). OWA-based rules generalise both Chamberlin–Courant and PAV.

Another way of motivating OWA-based rules, as defined below, is by examining the contrast between the k -Borda rule and Chamberlin–Courant, a line of thought explored by Faliszewski et al. (2017a). The k -Borda rule returns the committee formed of the k -candidates whose Borda scores are highest. This rule is suitable if we aim to identify a collection of ‘excellent candidates’, such as when shortlisting candidates for job interviews, or when finding finalists of various competitions, based on evaluations provided by judges. On the other hand, Chamberlin–Courants identifies committees that are *diverse*, in that they represent as many of the voters as possible. In many real-life applications, we desire a compromise between these two extremes; for example, we might

want our shortlist of excellent candidates to also be somewhat diverse. Faliszewski et al. (2017a) propose the t -Borda rules, $t = 1, \dots, k$, as a family of rules that interpolate between k -Borda and Chamberlin–Courant. Under the t -Borda rule, a voter i ’s utility is measured by the sum of the Borda scores of i ’s top t most-preferred members of the committee. Thus, Chamberlin–Courant is identical to 1-Borda, and k -Borda is identical to, well, k -Borda.

Let us now formally define OWA-based rules. Given a vector $\mathbf{x} \in \mathbb{R}^k$, a weight vector $\alpha \in \mathbb{R}^k$ defines an *ordered weighted average* (OWA) operator as follows: first, sort the entries of \mathbf{x} into non-increasing order, so that $x_{\sigma(1)} \geq \dots \geq x_{\sigma(k)}$; second, apply the weights: the ordered weighted average of \mathbf{x} with weights α is given by $\alpha(\mathbf{x}) := \sum_{i=1}^k \alpha_i x_{\sigma(i)}$. For example, $\alpha = (1, 0, \dots, 0)$ gives the maximum, and $\alpha = (1, 1, \dots, 1)$ gives the sum of the numbers in \mathbf{x} .

Now, a scoring vector $\mathbf{w} \in \mathbb{N}^m$ and an OWA α define an OWA-based multi-winner rule as follows: Given a profile P , the rule outputs a committee $W = \{c_1, \dots, c_k\}$ that maximises the objective value

$$\sum_{i \in N} \alpha(w_{c_1}, \dots, w_{c_k}).$$

Thus, choosing $\alpha = (1, 0, \dots, 0)$ gives us **w-Chamberlin–Courant** as a special case. Choosing $\alpha = (1, 1, 0, \dots, 0)$ gives us 2-Borda, the analogue of Chamberlin–Courant where voters are represented by their favourite *two* members of the committee. The OWA-based rules with $\alpha = (1, \frac{1}{2}, \dots, \frac{1}{k})$ and $\mathbf{w} = (1, 0, \dots, 0)$ gives us PAV, when given dichotomous preferences as input. Thus, OWA-based rules generalise both Chamberlin–Courant and PAV, and it turns out that we can apply our method to these rules by merging the ideas of (PAV-IP) and (CC-IP). However, our formulation is only valid for *non-increasing* OWA vectors with $\alpha_i \geq \alpha_j$ whenever $i \geq j$. For example, this excludes the rule where voters are represented by their *least*-favourite committee member.³

In the IP, we use variables $x_{i,\ell,r}$ indicating whether voter $i \in N$ ranks at least ℓ candidates in the committee in rank r or higher. We again put $w'_r = w_r - w_{r-1}$ and $w'_1 = w_1$.

$$\text{maximise } \sum_{i \in N} \sum_{\ell \in [k]} \sum_{r \in [m]} \alpha_\ell \cdot w'_r \cdot x_{i,\ell,r} \quad (\text{OWA-IP})$$

$$\text{subject to } \sum_{c \in C} y_c = k \quad (2)$$

$$\sum_{\ell \in [k]} x_{i,\ell,r} \leq \sum_{c : \text{rank}(c) \geq r} y_c \quad \text{for } i \in N, r \in [m] \quad (3)$$

$$x_{i,\ell,r} \in \{0, 1\} \quad \text{for } i \in N, \ell, r \quad (4)$$

$$y_c \in \{0, 1\} \quad \text{for } c \in C$$

³Still, this case is also efficiently solvable in the single-peaked case: note that a voter’s least-favourite committee members will be either the left-most or the right-most member of the committee; thus it suffices to consider committees of size 2. This idea can be extended to OWA operators $\alpha = (0, \dots, 0, \alpha_{k-c}, \dots, \alpha_k)$ that are zero except for constantly many values at the end.

Proposition 11. *If α and \mathbf{w} are non-increasing, (OWA-IP) correctly computes an optimal committee according to the OWA-based rule based on α and \mathbf{w} .*

Proof sketch. Similarly to previous arguments, in optimum, we will have $x_{i,\ell,r} = 1$ if and only if the committee $W = \{c \in C : y_c = 1\}$ contains at least ℓ candidates that voter i ranks in rank r or above. Thus, the objective value of (OWA-IP) agrees with the defined objective of the OWA-based rule. \square

The following property is proved very similarly to before: the constraint matrix is obtained from M_{SP}^P by appending unit columns.

Proposition 12. *The constraint matrix of (OWA-IP) is totally unimodular when input preferences are single-peaked.*

Theorem 13. *OWA-based rules with non-increasing OWA operator can be solved in polynomial time for single-peaked preferences.*

Since Chamberlin–Courant and PAV are special cases of OWA-based rules, Theorem 13 implies our previous results.

6 Some Extensions

More than Single-Peakedness. Our polynomial-time results apply to a slightly larger class than just single-peaked profiles: they also apply when M_{SP}^P (with an all-1s row appended) is totally unimodular but does not necessarily have the consecutive ones property. It can be shown that this is the case whenever P contains only two distinct voters, or, more generally, when the set of all top-initial segments of P can also be induced by a two-voter profile. Together with single-peaked profiles, we conjecture that these classes of profiles are precisely the profiles for which the relevant constraint matrices are totally unimodular.

Single-Peaked on a Circle. Peters and Lackner (2017) introduce the notion of preferences single-peaked on a circle, a generalisation of single-peakedness. They show, by adapting our total unimodularity technique, that all the multiwinner rules we consider remain tractable even for their larger preference restriction. Notably, their argument requires running a recognition algorithm first, and then rearranging our ILP formulations in order to obtain total unimodularity.

Egalitarian versions. We can obtain egalitarian versions of the multi-winner rules that we have discussed by replacing the sum over N by a minimum in their objective values (Betzler, Slinko, and Uhlmann 2013). For PAV and Chamberlin–Courant, our IP formulations can easily be adapted to answer the question “is there a committee with egalitarian objective value $\geq L$?” while preserving total unimodularity in the case of single-peaked preferences. An optimum committee can then be found by a binary search on L . However, it is unclear how this can be achieved for OWA-based rules. It is also unclear how to handle other utility aggregation operators such as leximin (see Elkind and Ismaili 2015).

PAV and Voter Intervals. Elkind and Lackner (2015) define an analogue of single-crossingness for dichotomous preferences called *voter interval* (VI), which requires the *transpose*

of M_{SP}^P to have the consecutive ones property. As for CI, they conjectured that PAV remains hard on VI preferences. We could not solve this problem using our method: the constraint $\sum_{c \in C} y_c = k$ of (PAV-IP) destroys total unimodularity.

Monroe’s rule. Monroe (1995) proposed a version of Chamberlin–Courant in which each committee member should represent (roughly) the same number of voters. Since this rule (in a certain version) remains NP-complete to evaluate for single-peaked preferences (Betzler et al. 2013), our approach will likely not extend to Monroe’s rule.

Dodgson’s rule. An alternative is a *Dodgson winner* if it can be made a Condorcet winner using a minimum number of swaps of adjacent alternatives. This number of swaps is the *Dodgson score* of an alternative. Bartholdi III, Tovey, and Trick (1989) give an IP formulation for this problem, which is also used in the treatment of Caragiannis et al. (2009). Sadly, while ‘most’ of the constraint matrix is again identical to M_{SP}^P , some extra constraints (saying that the swaps in each vote should only count once) destroy total unimodularity, so our method cannot be employed for this formulation. Note that while Brandt et al. (2015) give an efficient algorithm for finding a *Dodgson winner* in the case of single-peaked preferences, the problem of efficiently calculating *scores* appears to be open and non-trivial. Maybe there is an alternative IP formulation that can be made to work using our approach.

Kemeny’s rule. Conitzer, Davenport, and Kalagnanam (2006) present several IP formulations for Kemeny’s rule. The poly-size formulation they give involves constraints enforcing transitivity of the Kemeny ranking; these constraints are not totally unimodular. In any case, most strategies for calculating Kemeny’s rule first calculate all pairwise majority margins; we might as well check for transitivity at this stage – trying to use fancy total unimodularity is unnecessary.

7 Conclusions

We presented a new algorithmic technique that can evaluate some popular multi-winner voting rules in polynomial time when preferences are single-peaked. Interestingly, this approach works even if we have not checked in advance whether single-peakedness applies. In future work, it will be interesting to see how these formulations perform in practice, and to see whether other problems admit formulations of this type.

Acknowledgements. I thank Martin Lackner and Piotr Skowron for helpful discussions and insights. I am supported by EPSRC and through ERC grant 639945 (ACCORD).

References

- Arrow, K. J. 1950. A difficulty in the concept of social welfare. *The Journal of Political Economy* 328–346.
- Aziz, H.; Gaspers, S.; Gudmundsson, J.; Mackenzie, S.; Mattei, N.; and Walsh, T. 2015. Computational aspects of multi-winner approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 107–115.
- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare* 48(2):461–485.

- Bartholdi III, J., and Trick, M. A. 1986. Stable matching with preferences derived from a psychological model. *Operation Research Letters* 5(4):165–169.
- Bartholdi III, J.; Tovey, C. A.; and Trick, M. A. 1989. Voting schemes for which it can be difficult to tell who won the election. *Social Choice and Welfare* 6(2):157–165.
- Betzler, N.; Slinko, A.; and Uhlmann, J. 2013. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research* 47(1):475–519.
- Black, D. 1948. On the rationale of group decision-making. *Journal of Political Economy* 23–34.
- Booth, K. S., and Lueker, G. S. 1976. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. *Journal of Computer and System Sciences* 13(3):335–379.
- Brandt, F.; Brill, M.; Hemaspaandra, E.; and Hemaspaandra, L. A. 2015. Bypassing combinatorial protections: Polynomial-time algorithms for single-peaked electorates. *Journal of Artificial Intelligence Research* 53:439–496.
- Bredereck, R.; Faliszewski, P.; Niedermeier, R.; Skowron, P.; and Talmon, N. 2015. Elections with few candidates: Prices, weights, and covering problems. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT)*, 414–431. Springer.
- Brill, M.; Laslier, J.; and Skowron, P. 2017. Multiwinner approval rules as apportionment methods. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 414–420.
- Caragiannis, I.; Covey, J. A.; Feldman, M.; Homan, C. M.; Kakkamanis, C.; Karanikolas, N.; Procaccia, A. D.; and Rosenschein, J. S. 2009. On the approximability of Dodgson and Young elections. In *Proceedings of the 20th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1058–1067.
- Chamberlin, B., and Courant, P. 1983. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review* 77(3):718–733.
- Conitzer, V.; Davenport, A.; and Kalagnanam, J. 2006. Improved bounds for computing Kemeny rankings. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, 620–626.
- Doignon, J.-P., and Falmagne, J.-C. 1994. A polynomial time algorithm for unidimensional unfolding representations. *Journal of Algorithms* 16(2):218–233.
- Elkind, E., and Ismaili, A. 2015. OWA-based extensions of the Chamberlin–Courant rule. In *Proceedings of the 4th International Conference on Algorithmic Decision Theory (ADT)*, 486–502. Springer.
- Elkind, E., and Lackner, M. 2015. Structure in dichotomous preferences. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2019–2025.
- Elkind, E.; Faliszewski, P.; and Skowron, P. 2014. A characterization of the single-peaked single-crossing domain. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 654–660.
- Elkind, E.; Lackner, M.; and Peters, D. 2017. Structured preferences. In Endriss, U., ed., *Trends in Computational Social Choice*. AI Access.
- Escoffier, B.; Lang, J.; and Öztürk, M. 2008. Single-peaked consistency and its complexity. In *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI)*, 366–370.
- Faliszewski, P.; Hemaspaandra, E.; Hemaspaandra, L. A.; and Rothe, J. 2011. The shield that never was: Societies with single-peaked preferences are more open to manipulation and control. *Information and Computation* 209(2):89–107.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2016a. Multiwinner analogues of the plurality rule: Axiomatic and algorithmic perspectives. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 482–488.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2016b. Committee scoring rules: Axiomatic classification and hierarchy. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 482–488.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017a. Multiwinner rules on paths from k -Borda to Chamberlin–Courant. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 192–198.
- Faliszewski, P.; Skowron, P.; Slinko, A.; and Talmon, N. 2017b. Multiwinner voting: A new challenge for social choice theory. In Endriss, U., ed., *Trends in Computational Social Choice*. AI Access.
- Fitzsimmons, Z. 2015. Single-peaked consistency for weak orders is easy. In *Proceedings of the 15th Conference on Theoretical Aspects of Rationality and Knowledge (TARK)*, 103–110.
- Kilgour, D. M. 2010. Approval balloting for multi-winner elections. *Handbook on Approval Voting* 105–124.
- Lackner, M., and Skowron, P. 2017. Consistent approval-based multi-winner rules. *arXiv preprint arXiv:1704.02453*.
- Lackner, M. 2014. Incomplete preferences in single-peaked electorates. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 742–748.
- Lu, T., and Boutilier, C. 2011. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI)*, 280–286.
- Monroe, B. L. 1995. Fully proportional representation. *American Political Science Review* 89(04):925–940.
- Moulin, H. 1991. *Axioms of Cooperative Decision Making*. Cambridge University Press.
- Peters, D., and Elkind, E. 2016. Preferences single-peaked on nice trees. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 594–600.
- Peters, D., and Lackner, M. 2017. Preferences single-peaked on a circle. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 649–655.
- Potthoff, R. F., and Brams, S. J. 1998. Proportional representation: Broadening the options. *Journal of Theoretical Politics* 10(2):147–178.
- Schrijver, A. 1998. *Theory of Linear and Integer Programming*. John Wiley & Sons.
- Seymour, P. 1980. Decomposition of regular matroids. *Journal of Combinatorial Theory, Series B* 28(3):305–359.
- Skowron, P.; Yu, L.; Faliszewski, P.; and Elkind, E. 2013. The complexity of fully proportional representation for single-crossing electorates. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, 1–12.
- Skowron, P.; Faliszewski, P.; and Lang, J. 2015. Finding a collective set of items: From proportional multirepresentation to group recommendation. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 2131–2137.
- Skowron, P.; Faliszewski, P.; and Slinko, A. 2015. Achieving fully proportional representation: Approximability results. *Artificial Intelligence* 222:67–103.
- Thiele, T. N. 1895. Om flerfoldevalg. In *Oversigt over det Kongelige Danske Videnskaberne Selskabs Forhandling*. 415–441.