Committee Scoring Rules, Banzhaf Values, and Approximation Algorithms

Edith Elkind
University of Oxford
Oxford, UK

Piotr Faliszewski
AGH University
Krakow, Poland

Martin Lackner
University of Oxford
Oxford, UK

Dominik Peters
University of Oxford
Oxford, UK

Nimrod Talmon
Weizmann Institute
Rehovot, Israel

ABSTRACT

We consider committee scoring rules (a family of multiwinner voting rules) and define a class of cooperative games based on elections held according to these rules. We show that there is a polynomial-time algorithm for computing the Banzhaf value for a large subclass of these games and we show, using this Banzhaf value, an appealing heuristic algorithm for computing winning committees. We evaluate this algorithm experimentally for the case of the Chamberlin-Courant voting rule.

CCS Concepts

•Computing methodologies → Multi-agent systems; Cooperation and coordination;

Keywords

committee scoring rules, approximation algorithms, Banzhaf value, greedy algorithms, winner determination

1. INTRODUCTION

The goal of a multiwinner election is to choose a subset (a committee) of presented items (candidates) based on the preferences of a group of agents (the voters). Committee elections are a natural model for various tasks [12], ranging from shortlisting [4], through numerous business applications [13, 19, 20, 27], to tasks involving proportional representation, such as parliamentary elections [1, 6].

In consequence, there is a great number of very diverse multiwinner voting rules, based on many different principles. For example, Kilgour [18] discusses various approval-based rules (see also the work of Aziz et al. [2] for a more computational persective), Gehrlein [16] and Ratliff [25] discuss elections in the ordinal model that are based on the Condorcet principle[1], and many rules can be seen as extensions of single-winner scoring rules (e.g., the Bloc rule, the k-Borda rule [9], or the Chamberlin–Courant [8] and Monroe [22] rules). Elkind et al. [12] recently provided the formalism of committee scoring rules that captures many of the rules from this last group.

In this paper, we focus on the family of committee scoring rules (with a particular focus on decomposable committee scoring rules rules [13] in the theoretical part of the paper, and on the Chamberlin–Courant rule[2], in the experimental part). We show that committee scoring rules can be used to define a class of cooperative games, consider the Banzhaf value of this game, and show that it leads to improved approximation algorithms for our rules.

Committee Scoring Rules. We consider elections where the voters express preferences over the candidates by ranking them from the most to the least desired ones (we refer to such preference orders as votes), and where the goal is to select a group of candidates, i.e., a committee, of a given size k. A single-winner scoring function associates each position in a preference order with a score. For example, the Borda scoring function for m candidates, β_m, associates value $m-i$ with position i (so the top ranked candidate has Borda score $m-1$, the next one $m-2$, and so on). A committee scoring function associates a position of a committee in a preference order (i.e., a sequence of the positions of the committee members, sorted in the increasing order) with a numerical score, and the score of a committee in an election is the sum of the scores it receives from all the voters. For example, the k-Borda rule uses the committee scoring function that sums up the Borda scores of the committee members (within a vote), whereas the Chamberlin–Courant rule uses a scoring function under which the score of a committee (in a given vote) is the Borda score of the top-ranked committee member (referred to as the representative of this voter in the committee). In consequence, the k-Borda rule tends to select very similar candidates, whereas the Chamberlin–Courant rule tries to select a very diverse committee [11].

Cooperative Games for Committee Scoring Rules. Multiwinner elections and committee scoring rules define a certain natural class of cooperative games. Let us consider an election E (with candidate set C and voter collection V) and some committee scoring rule $\mathcal{R}_f$ with underlying committee

[1]In single-winner voting, a candidate is a Condorcet winner if he or she is preferred to every other candidate by a majority of the voters (albeit, possibly a different majority in each case).

[2]Naturally, we intend to extend these results, but for the current—preliminary report of our work—in the experimental evaluation we limit ourselves to this rule only.

Appears at: 4th Workshop on Exploring Beyond the Worst Case in Computational Social Choice (EXPLORE 2017). Held as part of the Workshops at the 16th International Conference on Autonomous Agents and Multiagent Systems. May 8th-9th, 2017. Sao Paulo, Brazil.

scoring function f (see the preliminaries for exact formal definitions).

We define a cooperative game $G(E, \mathcal{R}_f)$, where the candidates are the players and the value of a coalition S is the score assigned by f to committee S. Given such a game, we can use the full set of tools developed within cooperative game theory to analyze properties of the underlying election; here, in particular, we consider the Banzhaf values of the candidates, where the Banzhaf value of a player i in a cooperative game is its average marginal contribution to the value of a coalition, where the coalitions are chosen uniformly at random from the set of all possible coalitions that do not include i. Intuitively, the Banzhaf value of a candidate (a player) in our game measures the importance of this candidate. There are, however, two issues to resolve. First, the definition of Banzhaf value requires us to consider all coalitions, while in committee elections the size of the committee is usually constrained. Second, we may already know some committee members (see below) and the definition of the Banzhaf value does not account for this. Thus, we consider a variant of the Banzhaf value where we measure the average marginal contribution of a player to a randomly chosen coalition that (a) contains some given preselected players, and (b) is of a required size.

> We show that our generalized variant of the Banzhaf value is polynomial-time computable for a large class of committee scoring rules (including the Chamberlin–Courant rule, but also several more involved rules).

Application of the Banzhaf Value. Unfortunately, many committee scoring rules are NP-hard to compute [5, 19, 24, 27]. One way to deal with this problem is to use approximation algorithms, the first of which was proposed by Lu and Boutilier [19] for the case of the Chamberlin–Courant rule. This algorithm, which is a natural incarnation of the celebrated greedy approximation algorithm for maximizing submodular functions of Nemhauser et al. [23], achieves approximation ratio of $(1 - \frac{1}{e})$. Indeed, the algorithm was shown to be applicable to a very large subclass of committee scoring rules [13] (and to date it is the only general algorithm that applies to a wide class of committee scoring rules and provides guarantees regarding the output quality).

The greedy algorithm proceeds as follows. We start with an empty committee and in a sequence of k iterations (where k is the number of candidates that we want to select) we keep adding to the committee those candidates that increase its score by the largest value. While this algorithm seems to be achieving very good results in practice (see the works of Lu and Boutilier [19] and Skowron et al. [28] for some experimental results), it also has some drawbacks. For example, if it is used to select an approximate Chamberlin–Courant committee, then in the first iteration it always selects the candidate with the highest individual Borda score (i.e., the so-called Borda winner), even though in many cases this candidate does not belong to the optimal committee. While it is clear that an approximation algorithm does not always choose the optimal committee, always selecting the Borda winner creates a huge systematic bias against some candidates that might otherwise be selected (this effect is illustrated in the full version of the work of Elkind et al. [11]). This bias is a strong argument against using the algorithm.

> We modify the greedy algorithm so that instead of choosing the candidate that increases the score of the committee the most, it chooses the candidate with the highest Banzhaf value (for a given committee size and the candidates already selected in previous iterations).

We evaluate our algorithm experimentally for the case of the Chamberlin–Courant rule. In our experiments, the Banzhaf-based algorithm essentially always outperforms the original greedy algorithm, and often outperforms the algorithm of Skowron et al. [28] that underlies their polynomial-time approximation scheme (PTAS) for the Chamberlin–Courant rule (we use a variant of the algorithm improved by Elkind et al. [11]; see the full version of their paper).

Our approach is partially inspired by the line of work regarding the use of the Shapley value to extend centrality notions for networks (see the work of Michalak et al. [21] as a representative paper). We also model our problem (finding a good, approximate committee) using the language of cooperative game theory and use its tools to obtain improved results.

Organization of the Paper. In Section 2, we provide necessary background regarding elections, committee scoring rules, and cooperative games. Then, in Section 3, we discuss the Banzhaf value for the games we define as well as algorithms for computing it. In Section 4, we present the greedy algorithm based on the Banzhaf value, and in Section 5 we evaluate it experimentally for the Chamberlin–Courant rule. We discuss further research directions in Section 6.

2. PRELIMINARIES

In this section we present necessary background regarding elections and cooperative games. For a positive integer t, we write $[t]$ to denote the set $\{1, \dots, t\}$. For a logical expression Λ, by $[\Lambda]$ we mean 1 if Λ is true and 0 otherwise.

2.1 Multiwinner Elections

An election is a pair $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ is the set of candidates and $V = (v_1, \dots, v_n)$ is a collection of voters. Each voter v is associated with a preference order $\succ_v$, i.e., with a ranking of the candidates from C (from best to worst). A single-winner voting rule is a function that, given an election, outputs a set of tied winners.[3] In multiwinner elections we are interested in choosing whole committees of candidates of a given size k. A multiwinner voting rule $\mathcal{R}$ is a function that, given an election $E = (C, V)$ and a positive integer k, $k \le |C|$, returns a family of size-k committees that tie as winners. Before we discuss committee scoring rules, used for multiwinner elections, we discuss single-winner scoring rules, on which they are based.

Single-Winner Scoring Rules. Let E be an election with m candidates. For a candidate c and a voter v, we write $\mathrm{pos}_v(c)$ to denote the position of c in v's preference order. A single-winner scoring function γ_m for m candidates, $\gamma_m \colon [m] \to \mathbb{R}$, is a non-increasing function that associates each possible position in a vote with a score. For example, $\beta_m(i) = m - i$ is the Borda scoring function; for each positive integer t, $\alpha_t(i) = [i \le t]$ is the t-Approval scoring function. Typically,

[3] In practice, it is necessary to have some tie-breaking scheme. We disregard this issue for convenience.

we consider families of scoring functions, e.g., $\gamma = (\gamma_m)_{m\in\mathbb{N}}$, with one function for each number of candidates. For such a family γ, the γ-score of candidate c in election $E = (C, V)$ is $\gamma\text{-score}_E(c) = \sum_{v\in V} \gamma_{|C|}(\mathrm{pos}_v(c))$. For each scoring function γ, we have a single-winner scoring rule, denoted by $\mathcal{R}_\gamma$, which is defined as the voting rule which outputs the candidates with the highest γ-score.

Committee Scoring Rules. Elkind et al. [12] generalize the idea of single-winner scoring functions to the committee setting. Consider an election E with m candidates and a given committee size k. We define the *position* of a committee S (i.e., of a size-k set of candidates) in some vote v to be the increasing sequence resulting from sorting the set $\{\mathrm{pos}_v(c) \mid c \in S\}$. We write $[m]_k$ to denote the set of all length-k increasing sequences of numbers from $[m]$. For two committee positions $I = (i_1, \dots, i_k)$ and $J = (j_1, \dots, j_k)$, $I, J \in [m]_k$, we say that I *dominates* J if $i_1 \leq j_1, \dots, i_k \leq j_k$.

A *committee scoring function* $f_{m,k}$ for m candidates and committee size k, $f_{m,k}\colon [m]_k \to \mathbb{R}$, is a function that associates each committee position with a score in such a way that if some committee position I dominates some committee position J, then it holds that $f_{m,k}(I) \geq f_{m,k}(J)$.

Definition 1 (Elkind et al. [12]). *Let $f = (f_{m,k})_{k\leq m}$ be a family of committee scoring functions. The* committee scoring rule *$\mathcal{R}_f$ is a multiwinner rule that, given an election $E = (C, V)$ and a committee size k, outputs those committees S that maximize the value:*

$$f\text{-score}_E(S) = \sum_{v\in V} f_{|C|,k}\big(\mathrm{pos}_v(S)\big).$$

To distinguish single-winner scoring functions and committee scoring functions, we use Greek letters to denote the former and Latin letters to denote the latter.

Examples of Committee Scoring Rules. It turns out that the family of committee scoring rules is quite rich [13]; specifically, a number of well-known multiwinner voting rules are, in fact, committee scoring rules. For example, the k-Borda rule (which outputs committees of k candidates with the highest Borda scores) is the committee scoring rule defined via the following scoring function:

$$f^{k\text{-Borda}}_{m,k}(i_1, \dots, i_k) = \beta_m(i_1) + \dots + \beta_m(i_k).$$

The SNTV and the Bloc rules are defined analogously, where, instead of the Borda scoring function, the former uses 1-Approval (known as the Plurality scoring function) while the latter uses the k-Approval scoring function. Committee scoring rules of this form are known as *weakly separable.* Identifying a winning committee under a weakly separable rule can be done in polynomial time (provided the underlying single-winner scoring functions are polynomial-time computable), since one can compute the corresponding score of each candidate independently from the others.

The Chamberlin–Courant rule (β-CC) is the committee scoring rule defined by scoring functions of the following form:

$$f^{\beta\text{-CC}}_{m,k}(i_1, \dots, i_k) = \beta_m(i_1).$$

Intuitively, under the Chamberlin–Courant rule, each voter is assigned a representative (the committee member that he or she ranks the highest, among all the selected committee members) and each voter increases the score of the committee by exactly the Borda score of his or her representative. The committees with the highest scores tie as co-winners. Elkind et al. [12] argue that this rule is useful where we aim to find a diverse committee in which each voter is well represented (but where each committee member can represent different numbers of voters). The Chamberlin–Courant rule is computationally hard [19, 24] but there are good parameterized algorithms for it [5] as well as approximation algorithms and heuristics [15, 28].

Our final example of a committee scoring rule is the α_k-PAV rule, which is a variant of the *proportional approval voting (PAV)* rule, adapted to the format of committee scoring rules. It uses scoring functions of the following form:

$$f^{\alpha_k\text{-PAV}}_{m,k}(i_1, \dots, i_k) = \sum_{t=1}^{k} {}^1\!/\!_t \cdot \alpha_k(i_t).$$

Aziz et al. [1], Elkind et al. [11], Sánchez-Fernández et al. [26] and Brill et al. [6] provide strong evidence that the α_j-PAV rule is a very good choice if the goal is to elect a committee that represents voters proportionally. Winner determination for this rule is computationally hard, but there exists an approximation algorithm [27] and parameterized algorithms [14].

Decomposable Committee Scoring Rules. Throughout this paper we are mostly interested in *decomposable committee scoring rules* [13]. A committee scoring rule is decomposable if it can be defined through committee scoring functions of the following form:

$$f_{m,k}(i_1, \dots, i_k) = \gamma^{(1)}_{m,k}(i_1) + \dots + \gamma^{(k)}_{m,k}(i_k),$$

where $\gamma = (\gamma^{(t)}_{m,k})_{1\leq t\leq k\leq m}$ is a vector of single-winner scoring functions. All the committee scoring rules mentioned above are decomposable committee scoring rules; indeed, the subclass of decomposable committee scoring rules is quite a profound subclass of all committee scoring rules [13].

2.2 Cooperative Games

A cooperative game $G = (N, \nu)$ consists of a set of players $N = \{1, \dots, n\}$ and a characteristic function $\nu\colon 2^N \to \mathbb{R}$ such that $\nu(\emptyset) = 0$. Intuitively, for each coalition N' of players ($N' \subseteq N$), $\nu(N')$ is the joint payoff that the players in N' receive for working together. We refer to subsets of players as coalitions. Throughout this paper we consider monotone games only, i.e., games where for each two coalitions N' and N'' such that $N' \subseteq N''$, it holds that $\nu(N') \leq \nu(N'')$.

There are many solution concepts in cooperative game theory that describe which coalitions may form and/or how to distribute the coalitions' payoffs among their members (see, e.g., the overview of Chalkiadakis et al. [7] for a computer science perspective on the theory of cooperative games). Among the solution concepts, we focus on the Banzhaf value [3, 10], as the basic notion which turns out to be useful in our context.

Definition 2. *Let $G = (N, \nu)$ be a cooperative game. The Banzhaf value of player $i \in N$ is defined as follows:*

$$B_G(i) = \frac{1}{2^{|N-\{i\}|}} \sum_{S\subseteq N-\{i\}} \big(\nu(S \cup \{i\}) - \nu(S)\big).$$

In other words, the Banzhaf value of player i is its marginal contribution to a randomly selected coalition. Intuitively, we can view it as the player's importance: the higher the Banzhaf value, the more useful the player is for a (random) coalition.

3. COMMITTEE SCORING RULES AND BANZHAF VALUES

One of the contributions of this paper is in the following connection between multiwinner elections (committee scoring rules in particular) and cooperative games (Banzhaf values in particular). Next, we define a class of cooperative games associated with multiwinner elections and committee scoring rules.

DEFINITION 3. *Let $E = (C, V)$ be an election with $C = (c_1, \dots, c_m)$ and let $\mathcal{R}_f$ be a committee scoring rule defined through scoring functions $f = (f_{m,k})_{k \leq m}$. We define the game $G(E, \mathcal{R}_f) = (C, \nu)$, associated with an election E and an election rule $\mathcal{R}$ so that for each coalition S of candidates (players) we have:*

$$\nu(S) = \begin{cases} f\text{-score}_E(S) & , \text{ if } S \neq \emptyset, \\ 0 & , \text{ otherwise.} \end{cases}$$

The above definition requires some explanations and comments. First, the most important aspect of this definition is that the candidates are the players. The payoff of a coalition S is simply the score that this coalition—interpreted as a committee—would obtain in the underlying election. Thus, we use the terms *committee* and *coalition* interchangeably, depending on the context. Second, the characteristic function encompasses the committee scoring functions for all committee sizes between 1 and m. While it may seem somewhat strange at first, we view it as a natural approach (for the cases where we want to focus on particular election sizes only, we simply limit ourselves to coalitions of this size).

3.1 Computing Banzhaf Values

Let us consider the Banzhaf value of the game G associated with election $E = (C, V)$, where $C = \{c_1, \dots, c_m\}$ and $V = (v_1, \dots, v_n)$, and with the committee scoring rule $\mathcal{R}_f$, where $f = (f_{m,k})_{k \leq m}$. Intuitively, the Banzhaf value of a candidate measures the importance of this candidate in the given election.

Our first goal is to show that, for decomposable committee scoring rules, computing the Banzhaf values of all candidates can be done in polynomial-time. Our polynomial-time algorithm builds on the following two observations: (1) the Banzhaf value of a candidate can be computed for each voter separately (this follows since committee scoring rules treat each voter separately), and (2) the Banzhaf value of a candidate with respect to a given voter depends only on the position of this candidate (and on the positions of committee members already fixed to be present; see below); this follows since committee scoring rules depend only on the positions of the committee members within each voter's preference order.

Before we formally describe our polynomial time algorithm, we need some definitions and observations. Later we will need a variant of the Banzhaf value that considers committees of a given size only, such that some committee members are fixed, so we define the following variant of the Banzhaf value. We let k be a committee size, c_i be the candidate we are interested in, and W be a coalition of size smaller than k (which does not include c_i):

$$B_G(c_i, k, W) = \sum_{S \subseteq C : W \subseteq S, |S| = k-1} \nu(S \cup \{c_i\}) - \nu(S).$$

Note that:

$$B_G(c_i) = \frac{1}{2^{m-1}} \sum_{k=1}^{m} B_G(c_i, k, \emptyset),$$

so it suffices to focus on computing the values $B(c_i, k, W)$.

Next we show observation (1), which says that instead of considering the whole election E, it suffices to focus on each vote separately. We write $G(v_j)$ to denote the game G where the voter set is restricted to v_j only. Specifically, for each candidate c_i it holds that:

$$B_G(c_i, k, W) = \sum_{j=1}^{n} B_{G(v_j)}(c_i, k, W).$$

Corresponding to observation (2), now we prove the following technical lemma.

LEMMA 1. *Let $f = (f_{m,k})_{k \leq m}$ be a family of decomposable committee scoring rules (defined through polynomial-time computable single-winner scoring functions), let $E = (C, V)$ be an election, let k be the committee size, and let $G = (\mathcal{R}_f, E)$ be the game associated with $\mathcal{R}_f$ and E. Then, for each voter v in V, each candidate $c \in C$, and each set W such that $W \subseteq C - \{c\}$ and $|W| < k$, the value $B_{G(v)}(c, k, W)$ can be computed in polynomial time.*

PROOF. We set $m = |C|$. Let $\gamma_{m,k}^{(1)}, \dots, \gamma_{m,k}^{(t)}$ be the polynomial-time computable single-winner scoring functions such that:

$$f_{m,k}(i_1, \dots, i_k) = \gamma_{m,k}^{(1)}(i_1) + \dots + \gamma_{m,k}^{(k)}(i_k).$$

Let us rename the candidate set to $C = \{a_1, \dots, a_x, c, b_1, \dots, b_y\}$, such that voter v has the following preference order:

$$v \colon a_1 \succ \dots \succ a_x \succ c \succ b_1 \succ \dots \succ b_y.$$

Then, we partition W into two sets, W_A and W_B, such that v ranks all the candidates in W_A before c and all the candidates in W_B after c. Let ν be the characteristic function associated with our game $G(v)$. Our goal is to compute the following quantity:

$$B_{G(v)}(c, k, W) = \sum_{S \subseteq C \colon W \subseteq S, |S| = k-1} \nu(S \cup \{c\}) - \nu(S). \quad (1)$$

To this end, for each candidate $d \in C - \{c\}$ and each integer $t \in [k]$, we write $\mathcal{C}(d, t)$ to denote the set of coalitions S such that: (a) $W \subseteq S$; (b) $|S| = k - 1$; (c) $d \in S$; and (d) voter v ranks d as his or her t'th most desirable member of S. We define $r(d)$ to be 0 if v ranks d ahead of c and define it to be 1 otherwise. Further, we define:

$$\begin{aligned} \Delta(d) &= \sum_{t=1}^{k} \sum_{S \in \mathcal{C}(d,t)} \left(\gamma_{m,k}^{(t)}(\mathrm{pos}_v(d)) - \gamma_{m,k-1}^{(t+r(d))}(\mathrm{pos}_v(d)) \right) \\ &= \sum_{t=1}^{k} |\mathcal{C}(d, t)| \cdot \left(\gamma_{m,k}^{(t)}(\mathrm{pos}_v(d)) - \gamma_{m,k-1}^{(t+r(d))}(\mathrm{pos}_v(d)) \right). \end{aligned}$$

Intuitively, $\Delta(d)$ is the contribution of candidate d to the sum in Equation (1).

We define $\Delta(c)$ in a similar (but not identical) way. That is, for each $t \in [k]$, we let $\mathcal{C}(c,t)$ be the set of coalitions S such that $W \subset S$, $|S| = k-1$ and voter v ranks c as his or her t'th-best among the candidates in $S \cup \{c\}$. Further, we set:

$$\Delta(c) = \sum_{t=1}^{k} \sum_{S \in \mathcal{C}(c,t)} \gamma^{(t)}_{m,k}(\mathrm{pos}_v(c))$$
$$= \sum_{t=1}^{k} |\mathcal{C}(c,t)| \cdot \gamma^{(t)}_{m,k}(\mathrm{pos}_v(c)).$$

As in the case of $\Delta(d)$, $\Delta(c)$ is the contribution of c to the sum in Equation (1). We conclude the following:

$$B_{G(v)}(c,k,W) = \Delta(c) + \sum_{d \in C - \{c\}} \Delta(d).$$

To complete the proof, it suffices to note that, for each candidate $e \in C$ and each $t \in [k]$, the value $|\mathcal{C}(e,t)|$ can be computed in polynomial time. For example, for $t > |W_A|$ we have the following:

$$|\mathcal{C}(c,t)| = \binom{\mathrm{pos}_v(c) - 1 - |W_A|}{t - 1 - |W_A|} \cdot \binom{m - \mathrm{pos}_v(c) - |W_B|}{t - k - |W_B|}.$$

The idea behind the formula above is as follows. For c to be ranked on the t'th position among the candidates in $S \cup \{c\}$, S has to contain exactly $t-1$ candidates that v ranks ahead of c. S has to contain all members of W so it contains the $|W_A|$ members of W ranked ahead of t, and it suffices to add the missing $t-1-|W_A|$ in an arbitrary way (altogether there are $\mathrm{pos}_v(c) - 1 - |W_A|$ candidates that do not belong to W and that v ranks ahead of c). We calculate the number of ways in which we can choose members of S that v ranks after c analogously. $\square$

By our preceding reasoning, Lemma 1 immediately implies the following.

THEOREM 2. *For each decomposable committee scoring rule $\mathcal{R}_f$ defined through polynomial-time computable single-winner scoring functions, there is a polynomial-time algorithm that computes the Banzhaf value for each candidate in a given election.*

We conclude this section with the following important remark. Let $\mathcal{R}_f$ be some committee scoring rule and let $E = (C, V)$ be an election. Notice that, for each candidate $c \in C$ and each voter $v \in V$, the Banzhaf value of c in the game $G(\mathcal{R}_f, v)$ depends only on the position of c in v. This means that we can define a single-winner scoring rule γ (for $|C|$ candidates) so that $\gamma(i)$ is the Banzhaf value of the candidate ranked on the i'th position among the $|C|$ candidates in an election with a single vote. Then, the Banzhaf value of candidate c in the game $G(\mathcal{R}_f, E)$ is simply the γ-score of c in election E.

In particular, the above remark implies that forming a committee by choosing k candidates in the order of their decreasing Banzhaf values (with respect to some initial committee scoring rule) means simply using a weakly separable committee scoring rule (albeit, based on a fairly complicated single-winner scoring function).

Another consequence of the above remark is that for every committee scoring rule $\mathcal{R}_f$ based on a family of committee scoring functions with values bounded by functions exponential in the number of candidates, the problem of computing $B_{G(\mathcal{R}_f,E)}(c)$ (for some election E and a candidate c) belongs to the complexity class P/poly (see, e.g., the book of Hemaspaandra and Ogihara [17] for an extensive catalog of complexity classes). Intuitively, the class P/poly contains those problems for which, given an instance I and a value $h(|I|)$ (where $|I|$ is the length of the encoding of I and h is some, not necessarily computable, function whose output is polynomially bounded in $|I|$) it is possible to solve the problem in polynomial time. In our case, the value of $h(I)$ would consist of the description of functions γ from the preceding two paragraphs. The consequence is that, under standard complexity-theoretic assumptions, the problem of computing $B_{G(\mathcal{R}_f,E)}(c)$ cannot be NP-hard (however this does not apply to the more general problem from Lemma 1).

4. BANZHAF VALUES AND APPROXIMATION ALGORITHMS

In this section we show how to use the ideas concerning Banzhaf values, as discussed above, in order to design good heuristic algorithms for computing winning committees under decomposable committee scoring rules (indeed, winner determination under such rules is typically NP-hard).

Lu and Boutilier [19] introduced a greedy algorithm for computing committees of a given size with score close to the optimal one (they did it for the β-CC rule, and later other authors applied the algorithm to further rules, with Faliszewski et al. [13] providing the most general application). Let $E = (C, V)$ be the input election, let k be the committee size, and let $\mathcal{R}_f$ be the committee scoring rule to use (defined by committee scoring functions $f = (f_{m,k})_{k \le m}$). The algorithm starts with an empty committee S and then executes k iterations. In the i'th iteration, it selects a candidate $c \notin S$ that maximizes the value $f\text{-score}_E(S \cup \{c\})$.

Faliszewski et al. [13] invoke the classic result of Nemhauser et al. [23] to argue that this greedy algorithm outputs a committee whose score is at least a $(1 - 1/e)$-fraction of the optimal one, provided that the underlying scoring rule is decomposable and based on single-winner scoring functions $\gamma = (\gamma^{(t)}_{m,k})_{1 \le t \le k \le m}$ such that $\gamma^{(t)}_{m,k}(i) \ge \gamma^{(t+1)}_{m,k}(i)$ for each $m \in \mathbb{N}$ ($m > 0$), $k \in [m]$, $t \in [k-1]$, and $i \in [m]$. In particular, this applies both to β-CC and α_k-PAV.

Unfortunately, while this greedy algorithm has reasonably good approximation guarantee and it usually performs quite well in practice (in terms of its approximability [19, 28]), it has one serious drawback. Namely, it selects committees that are in some sense quite biased. For example, for β-CC, the greedy algorithm always starts by selecting the Borda winner, whereas for α_k-PAV it always starts by selecting a candidate with the highest k-Approval score.

We propose to rectify this issue by using a "non-myopic" variant of the algorithm. The only difference is that instead of selecting in the i'th iteration a candidate that maximizes the marginal increase of the given (partial) committee, we select a candidate with the highest Banzhaf value; specifically, in the i'th iteration, we compute Banzhaf values focusing on committees of size k, but under the assumption that the $i-1$ committee members from the previous iterations

Notation: $f = (f_{m,k})_{k \leq m} \leftarrow$ committee scoring functions.
$E = (C, V) \leftarrow$ input election.
$k \leftarrow$ committee size.
$G \leftarrow G(\mathcal{R}_f, E)$
$S \leftarrow \emptyset$
$A \leftarrow C$
for $i \leftarrow 1$ **to** K **do**
  $c \leftarrow \mathrm{argmax}_{a \in A} B_G(a, k, S)$
  $S \leftarrow S \cup \{c\}$
  $A \leftarrow S \setminus \{c\}$
return S

Algorithm 1: Banzhaf-based approximation algorithm.

are already belong to the coalition. The exact pseudocode is given as Algorithm 1.

As hinted above, and formally stated below, for many decomposable committee scoring rules (including all rules described in this paper), Algorithm 1 runs in polynomial time.

THEOREM 3. *For a decomposable committee scoring rule $\mathcal{R}_f$ defined by a family of polynomial-time computable single-winner scoring functions $\gamma = (\gamma^{(t)}_{m,k})_{1 \leq t \leq k \leq m}$, Algorithm 1 is polynomial-time computable.*

PROOF. Follows by Lemma 1 and a simple analysis of the algorithm. $\square$

5. EXPERIMENTAL ANALYSIS

In this section we report on experiments we performed in order to assess the approximation quality of Algorithm 1. We tested our Banzhaf-based approximation algorithms for β-CC (we intend to extend the experiments to other rules; the current manuscript is a preliminary presentation of our ideas). We performed two experiments; in the first we investigated biases of the algorithm, whereas in the second one we compared the quality of the generated committees.

Histograms. In the first experiment, building on the work of Elkind et al. [11], we computed histograms which visually demonstrate the behavior of our Banzhaf-based approximation algorithm, and compared it to the greedy algorithm of Lu and Boutilier [19] (we note that the histogram for this algorithm was already presented in the full version of the paper of Elkind et al. [11]).

Specifically, we generated elections from the two-dimensional Euclidean domain, where both the candidates and the voters are drawn uniformly at random from a uniform square (each candidate and each voter is a point; a voter ranks the candidates by sorting them in increasing order of the Euclidean distances). We generated 100 candidates and 100 voters for each election, and computed winning committees with 10 committee members each. We generated one histogram for the actual β-CC rule (using an ILP solver to find the optimal solutions), another histogram for the greedy approximation algorithm of Lu and Boutilier [19], and another histogram for our Banzhaf-based approximation algorithm. Each histogram was created by aggregating 10000 elections.

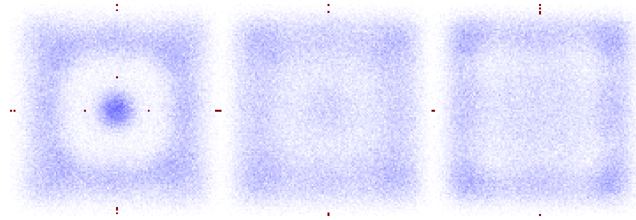

(a) greedy (b) Banzhaf (c) optimal

Figure 1: Comparison of the histograms for β-CC using the greedy algorithm (on the left), the Banzhaf-based one (in the center), and the optimal ILP-based one (on the right). The histograms were computed for 10000 elections, with 100 candidates and 100 voters each, for committee size $k = 10$.

Figure 1 shows the results of the first experiment. It is quite visible that, while the greedy algorithm of Lu and Boutilier [19] has a bias towards the center of the histogram (as explained above), the Banzhaf-based algorithm does not suffer from this problem. Further, the histogram of the Banzhaf-based algorithm looks very similarly to the histogram showing the optimal ILP-based algorithm for β-CC (the rangingCC algorithm, described below, also gets a histogram very similar to that for β-CC [11]).

Positions of the Representatives. In the second experiment, we checked the average position (in voters' preference orders) of the representatives chosen by various approximation rules. This is a more direct way of assessing the approximation quality of the algorithms, since β-CC minimizes this value. Indeed, β-CC maximizes the sum of the Borda scores of the representatives; we believe, however, that it is more useful to consider this "reversed" measure. There are two reason for this. First, measuring the direct approximation ratio suggests that all the algorithms have near-perfect performance (e.g., we have to choose between 0.98 and 0.99 approximation ratios; this is clearly possible but inconvenient). Second, the average position of the representative is a very intuitive measure from the point of view of the voters: A voter can more easily interpret information that "on the average he or she will be represented by someone he or she ranks as third best" than that "he or she will be represented by someone he or she prefers to 97 candidates, on average" (in particular, the former does not require the voter to know how many candidates there were in the election).

We considered two distributions of voters' preference orders. The first distribution is the one used in the first experiment (we generated candidates and voters as points on a square, uniformly at random, and the preference order of each voter is formed by sorting the candidates with respect to the distance from the voter). The second one followed the impartial culture assumption (each voter chose his or her preference order uniformly at random). For the first distribution we created 1000 elections, each with 100 voters and 100 candidates, and we varied the committee size k to be any integer between 2 and 30; we checked the average position of each voter's representative for β-CC (computed using an ILP solver), rangingCC,[4] the greedy approximation

[4]RangingCC is a variant of Algorithm P [28], improved by Elkind et al. [11]. The algorithm proceeds as follows: Given an election $E = (C, V)$ and a committee size k, it considers

algorithm, and our Banzhaf-based approximation algorithm. For the second distribution we generated only 250 elections for each committee size between 2 and 30 (with step 2), but otherwise the experiment was analogous (the reason for this restriction was that we ran out of time for our computations; we plan to have 1000 elections per data point for the final version of this paper).

Figure 2 shows the results of the second experiment. The results are normalized to β-CC, i.e., the figure shows the ratio between the average position of a representative under a given algorithm and the optimal average position (thus, the values are always greater or equal to 1). While rangingCC performs very well for small committee sizes, both the greedy algorithm and our Banzhaf-based approximation algorithm perform much better as the committee size increases. Further, our Banzhaf-based approximation algorithm consistently outperforms the greedy algorithm.

Running Times. In the analysis above we have disregarded the running times of our algorithms. Indeed, both rangingCC and greedyCC can be significantly faster than the Banzhaf-based algorithm (by an order of magnitude in our experiments)[5]. Thus, one might say that the Banzhaf-based heuristic has unfair advantage over the greedy algorithm and, in particular, one might consider a greedy algorithm that picks two candidates in each iteration instead of one. Our very preliminary experiments suggest that this does not improve the performance of the algorithm, but—in our setting—increases the running time by two orders of magnitude as compared to the classic greedy algorithm (resulting in, altogether, a 10 times slower algorithm than the Banzhaf-based one, with worse quality of results).

6. OUTLOOK

We considered multiwinner elections (held using committee scoring rules) as cooperative games and, building on this idea, were able to design improved approximation algorithms for winner determination for a rich class of multiwinner voting rules. We provided some preliminary experiments to assess the quality of our algorithms. While the experiments showed that in some cases the quality of approximation of these new approximation algorithms is quite good, more experiments are needed to fully understand when these algorithms are most useful. Specifically, one might consider further election distributions as well as real-world elections.

In this paper we concentrated on the Banzhaf value, as a fundamental solution concept in cooperative game theory. It might be interesting to consider other solution concepts, such as the Shapley value; guiding greedy algorithms by solution concepts other than the Banzhaf value might lead to efficient algorithms with a better quality of approximation.

Finally, in this paper we concentrated on committee scoring rules, especially decomposable committee scoring rules. While the subclass of decomposable committee scoring rules is quite rich, it is natural to wonder whether the ideas presented here can be useful for other multiwinner voting rules.

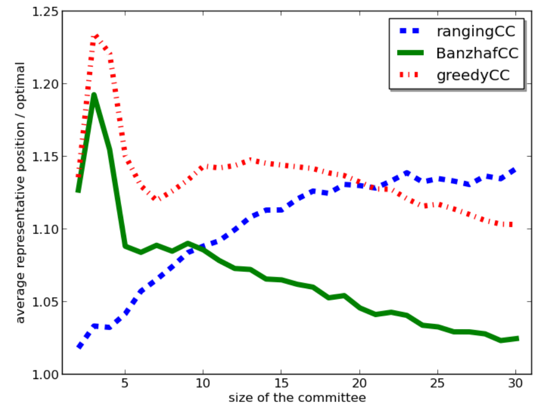

(a) 2D model: uniform distribution on a square

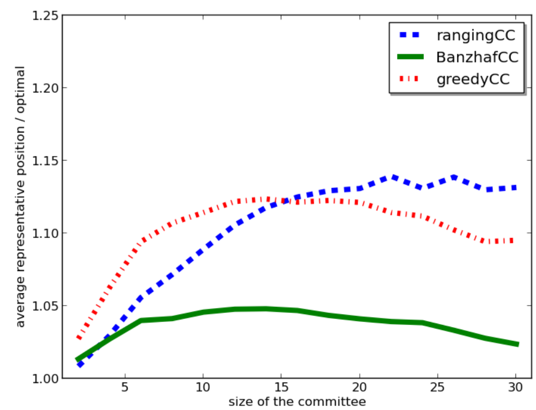

(b) Impartial culture

Figure 2: Comparison of the average representative position for β-CC for two distributions: uniform square Euclidean domain (on the top) and impartial culture (on the bottom); for three algorithms: rangingCC, greedyCC, and BanzhafCC (our Banzhaf-based algorithm). The values are computed for elections, with 100 candidates and 100 voters each (1000 elections per data point). For each committee size we present the ratio between the average position of a voter's representative under the given algorithm and under β-CC (computed using an ILP solver).

Acknowledgments Edith Elkind, Martin Lackner, and Dominik Peters were supported by the ERC grant 639945 (ACCORD). Piotr Faliszewski was supported by the National Science Centre, Poland, under project 2016/21/B/ST6/01509.

each threshold value $t \in [m]$; for a given threshold value t, it greedily finds a committee such that as many voters as possible rank some committee member among the top t positions (i.e., for a given value t, it first adds to the committee the candidate ranked among the top t positions by most voters, removes these voters from consideration, and repeats the process until k candidates are selected). Then, for each computed committee (one computed committee for each value of $t \in [m]$) it computes its β-CC score; then, it outputs the committee with the highest score. This algorithm is the basis of a PTAS for β-CC and has the highest theoretically-established approximation guarantee for this rule.

[5]However, we should mention that we used a highly-optimized variant of our algorithm. In particular, our algorithm never recomputed already-used values of binomial coefficients and used formulas from Lemma 1 optimized for β-CC, to not compute values that have to add up to zero.

REFERENCES

[1] H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48(2):461–485, 2017.

[2] H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. Computational aspects of multi-winner approval voting. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, pages 107–115, 2015.

[3] J. Banzhaf. Weighted voting doesn't work: A mathematical analysis. *Rutgers Law Review*, 19:317–343, 1965.

[4] S. Barberà and D. Coelho. How to choose a non-controversial list with k names. *Social Choice and Welfare*, 31(1):79–96, 2008.

[5] N. Betzler, A. Slinko, and J. Uhlmann. On the computation of fully proportional representation. *Journal of Artificial Intelligence Research*, 47:475–519, 2013.

[6] M. Brill, J. Laslier, and P. Skowron. Multiwinner approval rules as apportionment methods. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 414–420, 2017.

[7] G. Chalkiadakis, E. Elkind, and M. Wooldridge. *Computational Aspects of Cooperative Game Theory.* Morgan & Claypool, 2012.

[8] B. Chamberlin and P. Courant. Representative deliberations and representative decisions: Proportional representation and the Borda rule. *American Political Science Review*, 77(3):718–733, 1983.

[9] B. Debord. An axiomatic characterization of Borda's k-choice function. *Social Choice and Welfare*, 9(4):337–343, 1992.

[10] P. Dubey and L. Shapley. Mathematical properties of the Banzhaf power index. *Mathematics of Operations Research*, 4(2):99–131, May 1979.

[11] E. Elkind, P. Faliszewski, J. Laslier, P. Skowron, A. Slinko, and N. Talmon. What do multiwinner voting rules do? An experiment over the two-dimensional euclidean domain. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 494–501, 2017. Full version available at `http://home.agh.edu.pl/~faliszew/2d.pdf`.

[12] E. Elkind, P. Faliszewski, P. Skowron, and A. Slinko. Properties of multiwinner voting rules. *Social Choice and Welfare*, 48(3):599–632, 2017.

[13] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Committee scoring rules: Axiomatic classification and hierarchy. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, pages 250–256, 2016.

[14] P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. Multiwinner analogues of the plurality rule: Axiomatic and algorithmic views. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pages 482–488, 2016.

[15] P. Faliszewski, A. Slinko, K. Stahl, and N. Talmon. Achieving fully proportional representation by clustering voters. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, pages 296–304, 2016.

[16] W. Gehrlein. The Condorcet criterion and committee selection. *Mathematical Social Sciences*, 10(3):199–209, 1985.

[17] L. Hemaspaandra and M. Ogihara. *The Complexity Theory Companion.* Springer-Verlag, 2002.

[18] M. Kilgour. Approval balloting for multi-winner elections. In *Handbook on Approval Voting.* Springer, 2010. Chapter 6.

[19] T. Lu and C. Boutilier. Budgeted social choice: From consensus to personalized decision making. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, pages 280–286, 2011.

[20] T. Lu and C. Boutilier. Value-directed compression of large-scale assignment problems. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pages 1182–1190, 2015.

[21] T. Michalak, K. Aadithya, P. Szczepanski, B. Ravindran, and N. Jennings. Efficient computation of the shapley value for game-theoretic network centrality. *Journal of Artificial Intelligence Research*, 46:607–650, 2013.

[22] B. Monroe. Fully proportional representation. *American Political Science Review*, 89(4):925–940, 1995.

[23] G. Nemhauser, L. Wolsey, and M. Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, 14(1):265–294, December 1978.

[24] A. Procaccia, J. Rosenschein, and A. Zohar. On the complexity of achieving proportional representation. *Social Choice and Welfare*, 30(3):353–362, 2008.

[25] T. Ratliff. Some startling inconsistencies when electing committees. *Social Choice and Welfare*, 21(3):433–454, 2003.

[26] L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. Proportional justified representation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 670–676, 2017.

[27] P. Skowron, P. Faliszewski, and J. Lang. Finding a collective set of items: From proportional multirepresentation to group recommendation. *Artificial Intelligence*, 241:191–216, 2016.

[28] P. Skowron, P. Faliszewski, and A. Slinko. Achieving fully proportional representation: Approximability result. *Artificial Intelligence*, 222:67–103, 2015.