Theoretical Computer Science 569 (2015) 43–57

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

The complexity of fully proportional representation for single-crossing electorates ☆

Piotr Skowron [a,*], Lan Yu [b], Piotr Faliszewski [c], Edith Elkind [d]

[a] University of Warsaw, Warsaw, Poland
[b] Google Inc., Mountain View, CA, USA
[c] AGH University of Science Technology, Krakow, Poland
[d] University of Oxford, Oxford, UK

ARTICLE INFO

Article history:
Received 22 March 2014
Received in revised form 17 October 2014
Accepted 19 December 2014
Available online 29 December 2014
Communicated by D. Peleg

Keywords:
Single-crossing
Winner determination
Chamberlin–Courant's rule
Monroe's rule
Complexity

ABSTRACT

We study the complexity of winner determination in single-crossing elections under two classic fully proportional representation rules—Chamberlin–Courant's rule and Monroe's rule. Winner determination for these rules is known to be NP-hard for unrestricted preferences. We show that for single-crossing preferences this problem admits a polynomial-time algorithm for Chamberlin–Courant's rule, but remains NP-hard for Monroe's rule. Our algorithm for Chamberlin–Courant's rule can be modified to work for elections with *bounded single-crossing width*. We then consider elections that are both single-peaked and single-crossing, and develop an efficient algorithm for the egalitarian variant of Monroe's rule for such elections. While Betzler et al. [3] have recently presented a polynomial-time algorithm for this rule under single-peaked preferences, our algorithm has considerably better worst-case running time than that of Betzler et al.

1. Introduction

Parliamentary elections, i.e., procedures for selecting a fixed-size set of candidates that, in some sense, best represent the voters, received a lot of attention in the literature. Some well-known approaches include the first-past-the-post system (FPTP), where the voters are divided into districts and in each district a plurality election is held to find this district's representative; party-list systems, where the voters vote for parties and later the parties distribute the seats among their members; SNTV (single nontransferable vote) and Bloc rules, where each voter picks t candidates to approve, and the rule picks k candidates with the highest number of approvals (here k is the target parliament size, and $t = 1$ for SNTV and $t = k$ for Bloc); and some variants of STV (single transferable vote). In this paper, we focus on two voting rules that for each voter explicitly define the candidate that will represent her in the parliament (such rules are said to provide *fully proportional representation*), namely, Chamberlin–Courant's rule [6] and Monroe's rule [18]. Winner determination algorithms for these rules can also be used for applications other than parliamentary elections, such as resource allocation [18,25] and recommender systems [16].

☆ The preliminary version of this paper was presented at SAGT-2013.

* Corresponding author.
E-mail addresses: p.skowron@mimuw.edu.pl (P. Skowron), jen.lan.yu@gmail.com (L. Yu), faliszew@agh.edu.pl (P. Faliszewski), elkind@cs.ox.ac.uk (E. Elkind).

http://dx.doi.org/10.1016/j.tcs.2014.12.012
0304-3975/

Let us consider an election where n voters aim to select a k-member parliament out of m candidates. Both Chamberlin–Courant's rule and Monroe's rule work by finding a function Φ that maps each voter v to the candidate that is to represent v in the parliament. This function is required to output at most k candidates altogether.[1] Further, under Monroe's rule each candidate is either assigned to about $\frac{n}{k}$ voters or to none, whereas under Chamberlin–Courant's rule there is no such restriction (as a consequence, parliaments elected according to Chamberlin–Courant's rule may have to use weighted voting in their proceedings).

Intuitively, under both rules, each voter should be represented by a candidate that this voter ranks as highly as possible. To specify this requirement formally, we assume that there is a global *dissatisfaction function* α, $\alpha: \mathbb{N} \to \mathbb{N}$, such that $\alpha(i)$ is voter's dissatisfaction from being represented by a candidate that she views as i-th best—for instance, Borda dissatisfaction function α_B is given by $\alpha_B(i) = i - 1$. In the utilitarian variants of Chamberlin–Courant's and Monroe's rules we seek assignments that minimize the sum of voters' dissatisfactions; in the egalitarian variants (introduced recently by Betzler et al. [3]) we seek assignments that minimize the dissatisfaction of the worst-off voter.

Chamberlin–Courant's and Monroe's rules have a number of attractive properties, which distinguish them from other multiwinner rules. Indeed, they elect parliaments that (at least in some sense) proportionally represent the voters, ensure that candidates who are not individually popular do not make it to the parliament even if they come from very popular parties, and take minority candidates into account. In contrast, FPTP can provide largely disproportionate results, party-list systems cause members of parliament to feel more responsible to the parties than to the voters, SNTV and Bloc tend to disregard minority candidates, and STV is sometimes accused of putting too much emphasis on voters' top preferences. We point readers interested in the properties of Monroe's and Chamberlin–Courant's rules, as well as some other multiwinner voting rules, to the recent work of Elkind at al. [12].

Unfortunately, Chamberlin–Courant's and Monroe's rules do have one flaw that makes them impractical: It is NP-hard to compute a winning parliament with respect to these rules [21,16,3]. Nonetheless, these rules are so attractive that there is a growing body of research on computing their outputs either exactly (e.g., through integer linear programming formulations [20], by means of fixed-parameter tractability analysis [3], or by considering restricted preference domains [3,27]) or approximately [16,25,24,23]. We continue this line of research by considering the complexity of computing the outputs of Chamberlin–Courant's and Monroe's rules for the case where voters' preferences are single-crossing. Our results complement those of Betzler et al. [3] for single-peaked electorates.

Recall that voters are said to have single-crossing preferences if it is possible to order them so that for every pair of candidates a, b the voters who prefer a to b appear together on one side of the order and the voters who prefer b to a appear together on the other side. For example, it is quite natural to assume that the voters are aligned on the standard political left-right axis. Given two candidates a and b, where a is viewed as more left-wing and b is viewed as more right-wing, the left-leaning voters would prefer a to b and the right-leaning voters would prefer b to a. While real-life elections are typically too noisy to have this property, it is plausible that they are often close to single-crossing, and it is important to understand the complexity of the idealized model before proceeding to study nearly single-crossing profiles. Indeed, in the context of single-peaked elections, Faliszewski et al. [14] have recently shown that many algorithmic results established for such elections can be extended to elections that are nearly single-peaked.

Our main results are as follows: for single-crossing elections winner determination under Chamberlin–Courant's rule is in P (for every dissatisfaction function, and both for the utilitarian and for the egalitarian variant of this rule), but under Monroe's rule it is NP-hard. Our hardness result for Monroe's rule applies to the utilitarian setting with Borda dissatisfaction function. Our algorithm for Chamberlin–Courant's rule extends to elections that have bounded *single-crossing width* (see [7,8]). The proof proceeds by showing that for single-crossing elections Chamberlin–Courant's rule admits an optimal assignment that has the *contiguous blocks property*: the set of voters assigned to an elected representative forms a contiguous block in the voters' order witnessing that the election is single-crossing. This property can be interpreted as saying that each selected candidate represents a group of voters who are fairly similar to each other, and we believe it to be desirable in the context of proportional representation.

The NP-hardness result for winner determination under Monroe's rule motivates us to search for further domain restrictions that may make this problem efficiently solvable. To this end, we focus on the egalitarian variant of Monroe's rule and consider elections that are both single-peaked and single-crossing. We provide an $O(m^2n)$ algorithm for this setting (where n is the number of voters and m is the number of candidates), thus improving over the $O(n^3m^3k^3)$ algorithm (where k is the target parliament size) of Betzler et al. [3] for single-peaked preferences. Our algorithm is based on the observation that under the egalitarian variant of Monroe's rule single-peaked single-crossing elections always admit an optimal assignment that satisfies the contiguous blocks property. The proof uses the recent characterization result of Elkind et al. [10] for this domain. We show, however, that our approach does not extend to general single-crossing elections or to the utilitarian variant of Monroe's rule: in both cases, requiring the contiguous blocks property may rule out all optimal assignments. We then ask whether an analogue of the contiguous blocks property (obtained by considering the ordering of the voters

[1] Under Monroe's rule we are required to pick exactly k winners. Some authors also impose this requirement in the case of Chamberlin–Courant's rule, but allowing for smaller parliaments appears to be more consistent with the spirit of this rule and is standard in its computational analysis (see, e.g., [16,3,25,24,27]). In any case, this distinction has no bite if there are at least k candidates such that each of them is ranked first by at least one voter, which is a typical situation in political elections.