Journal of Computer and System Sciences 80 (2014) 697–710

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

The complexity of online manipulation of sequential elections[☆]

Edith Hemaspaandra[a,1], Lane A. Hemaspaandra[b,2], Jörg Rothe[c,*,3]

[a] *Department of Computer Science, Rochester Institute of Technology, Rochester, NY 14623, USA*
[b] *Department of Computer Science, University of Rochester, Rochester, NY 14627, USA*
[c] *Institut für Informatik, Heinrich-Heine-Universität Düsseldorf, 40225 Düsseldorf, Germany*

ARTICLE INFO

Article history:
Received 13 February 2013
Received in revised form 16 September 2013
Accepted 1 October 2013
Available online 9 October 2013

Keywords:
Computational complexity
Computational social choice
Elections
Manipulation
Online algorithms
Preferences
Sequential voting

ABSTRACT

Most work on manipulation assumes that all preferences are known to the manipulators. However, in many settings elections are open and sequential. We introduce a framework, in which manipulators can see the past votes but not the future ones, to model online coalitional manipulation of sequential elections, and we show that here manipulation can be extremely complex even for election systems with simple winner problems. We also show that for some of the most important election systems such manipulation is simple in certain settings. Among our highlights are: Depending on the size of the manipulative coalition, the online manipulation problem can be complete for each level of the polynomial hierarchy or even for PSPACE. We obtain the most dramatic contrast to date between the nonunique-winner and unique-winner models: Online weighted manipulation for plurality is in P in the nonunique-winner model, yet is coNP-hard (constructive case) and NP-hard (destructive case) in the unique-winner model. And we obtain what to the best of our knowledge are the first $\mathrm{P}^{\mathrm{NP}[1]}$-completeness and P^{NP}-completeness results in computational social choice, in particular regarding 3-candidate and 4-candidate (and unlimited-candidate) online weighted coalition manipulation of veto elections.

1. Introduction

Voting is a widely used method for preference aggregation and decision-making. In particular, *strategic* voting (or *manipulation*) has been studied intensely in social choice theory (starting with the celebrated work of Gibbard [20] and Satterthwaite [33]) and, in the rapidly emerging area of *computational* social choice, also with respect to its algorithmic properties and computational complexity (starting with the seminal work of Bartholdi, Tovey, and Trick [3]; see the surveys [15,16]). This computational aspect is particularly important in light of the many applications of voting in computer

---

[☆] A preliminary version of some of these results appears in the Fourteenth Conference on Theoretical Aspects of Rationality and Knowledge.

[*] Corresponding author.
URLs: http://www.cs.rit.edu/~eh (E. Hemaspaandra), http://www.cs.rochester.edu/u/lane (L.A. Hemaspaandra), http://ccc.cs.uni-duesseldorf.de/~rothe (J. Rothe).

[1] Supported in part by grant NSF-CCF-1101452 and a Friedrich Wilhelm Bessel Research Award. Work done in part while visiting Heinrich-Heine-Universität Düsseldorf.

[2] Supported in part by grants NSF-CCF-{0915792, 1101479}, ARC-DP110101792, and a Friedrich Wilhelm Bessel Research Award. Work done in part while visiting Heinrich-Heine-Universität Düsseldorf.

[3] Supported in part by DFG grant RO-1202/15-1, SFF grant "Cooperative Normsetting" of HHU Düsseldorf, ARC-DP110101792, and a DAAD-PPP grant in the PROCOPE program. Work done in part while visiting Stanford University.

0022-0000/$ – see front matter © 2013 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.jcss.2013.10.001

science, ranging from meta-search heuristics for the internet [14], to recommender systems [19] and multiagent systems in artificial intelligence (see the survey by Conitzer [11]).

Most of the previous work on manipulation, however, is concerned with voting where the manipulators know the nonmanipulative votes. Far less attention has been paid (see the related work below) to manipulation in the midst of elections that are modeled as dynamic processes.

We introduce a novel framework for online manipulation, where voters vote in sequence and the current manipulator, who knows the previous votes and which voters are still to come but does not know their votes, must decide—right at that moment—what the "best" vote to cast is. So, while other approaches to sequential voting are stochastic, game-theoretic (yet different from our approach, see Footnote 4), or axiomatic in nature (again, see the related work), our approach to manipulation of sequential voting is shaped by the area of "online algorithms" [8], in the technical sense of a setting in which one (for us, each manipulative voter) is being asked to make a manipulation decision just on the basis of the information one has in one's hands at the moment even though additional information/system evolution may well be happening down the line. In this area, there are different frameworks for evaluation. But the most attractive one, which pervades the area as a general theme, is the idea that one may want to "maxi-min" things—*one may want to take the action that maximizes the goodness of the set of outcomes that one can expect regardless of what happens down the line from one time-wise*. For example, if the current manipulator's preferences are Alice $>$ Ted $>$ Carol $>$ Bob and if she can cast a (perhaps insincere) vote that ensures that Alice or Ted will be a winner no matter what later voters do, and there is no vote she can cast that ensures that Alice will always be a winner, this maxi-min approach would say that that vote is a "best" vote to cast.

It will perhaps be a bit surprising to those familiar with online algorithms and competitive analysis that in our model of online manipulation we will not use a (competitive) *ratio*. The reason is that voting commonly uses an *ordinal* preference model, in which preferences are total orders of the candidates. It would be a severely improper step to jump from that to assumptions about intensity of preferences and utility, e.g., to assuming that everyone likes her nth-to-least favorite candidate exactly n times more than she likes her least favorite candidate.

Related work. Xia and Conitzer [41] (see also the related paper by Desmedt and Elkind [13]) define and study the Stackelberg voting game (also quite naturally called, in an earlier paper that mostly looked at two candidates, the roll-call voting game [34]). This basically is an election in which the voters vote in order, *and the preferences are common knowledge—everyone knows everyone else's preferences, everyone knows that everyone knows everyone else's preferences, and so on out to infinity*. Their analysis of this game is game-theoretically shaped; they compute a subgame perfect Nash equilibrium from the back end forward. Under their work's setting and assumptions, for bounded numbers of manipulators manipulation is in P, but we will show that in our model even with bounded numbers of manipulators manipulation sometimes (unless $\mathrm{P} = \mathrm{NP}$) falls beyond P.[4]

The interesting "dynamic voting" work of Tennenholtz [37] investigates sequential voting, but focuses on axioms and voting rules rather than on coalitions and manipulation. Much heavily Markovian work studies sequential decision-making and/or dynamically varying preferences; our work in contrast is nonprobabilistic and focused on the complexity of coalitional manipulation. Also somewhat related to, but quite different from, our work is the work on possible and necessary winners. The seminal paper on that is due to Konczak and Lang [29], and more recent work includes [40,7,1,5,6,10,4,30]; the biggest difference is that those are, loosely, one-quantifier settings, but the more dynamic setting of online manipulation involves numbers of quantifiers that can grow with the input size. Another related research line studies multi-issue elections [42–45]; although there the separate issues may run in sequence, each issue typically is voted on simultaneously and with preferences being common knowledge.

Organization. We first provide the needed preliminaries for (standard and sequential) elections, manipulation, and scoring rules, and give some background from complexity theory. Then, after introducing our model of online manipulation formally, we will present some general complexity results on the problems defined, and also some specific results for online manipulation in natural voting systems (i.e., for some central scoring rules). Finally, we turn to schedule-robust online manipulation, a setting in which not even the order of future voters is known to the current manipulator.

2. Preliminaries

Elections. A (*standard, i.e., simultaneous*) *election* (C, V) is specified by a set C of candidates and a list V, where we assume that each element in V is a pair (v, p) such that v is a voter name and p is v's vote. How the votes in V are

---

[4] Our work too is game-theoretically connected. Although in our model we are asking whether we can reach our goal no matter what the future nonmanipulators do, if one thinks about what the actual effect of this is, one can see that our setting is in effect well-captured by what is known as a 2-player combinatorial game (combinatorial games are a particular type of complete-information sequential game). In our setting, the goal of one player in this game will be to ensure that the winner set (which of course heavily depends on what moves have occurred already and on the election system) will have nonempty intersection with a certain subset of the candidates, and the goal of the other player will be to ensure that that does not happen. Of course, the former player is in effect the currently-under-consideration and still-to-vote members of the manipulative coalition, and the latter player is capturing the same except regarding nonmanipulators. So, the key differences between [41] and our work regard goals and coalitionality. For them, each player (and they may have many players) is in effect a completely separate agent, with a preference order, and is trying to see if a change as an individual will make a more preferred candidate win. For us, the manipulative voters function as a coalition, and one that has an all-or-nothing goal, and there are no gradations within that goal in terms of our analysis (despite the fact that we use a preference order when speaking of the coalition), and we are in effect a two-player combinatorial game.