Report from Dagstuhl Seminar 12101

Computation and Incentives in Social Choice

Edited by

Edith Elkind[1], Christian Klamler[2], Jeffrey S. Rosenschein[3], and M. Remzi Sanver[4]

1 **Nanyang TU – Singapore, SG,** eelkind@ntu.edu.sg
2 **Universität Graz, AT,** christian.klamler@uni-graz.at
3 **The Hebrew University of Jerusalem, IL,** jeff@cs.huji.ac.il
4 **Istanbul Bilgi University, TR,** sanver@bilgi.edu.tr

Abstract

Computational social choice is an active research area that combines tools and techniques of theoretical computer science and AI with those of mathematics, social sciences and economics. The aim of the Dagstuhl Seminar 12101 "Computation and Incentives in Social Choice" was to bring together the experts in these areas in order to discuss recent advances in this field and share open problems. This report collects the material presented during the course of the seminar.

Seminar 04.–09. March, 2012 – www.dagstuhl.de/12101
1998 ACM Subject Classification I.2.11. Distributed Artificial Intelligence
Keywords and phrases Computational Social Choice, Voting, Incentives, Algorithmic Game Theory, Fair Division
Digital Object Identifier 10.4230/DagRep.2.3.1
Edited in cooperation with Gábor Erdélyi

1 Executive Summary

Edith Elkind
Christian Klamler
Jeffrey S. Rosenschein
M. Remzi Sanver

The aim of classic social choice theory is to explain how groups of agents can come to a joint decision that reflects the heterogeneous preferences of individual agents. This covers a wide range of scenarios, such as, for example, voting, fair division and ranking. As such, social choice theory enhances our understanding of human societies and can be used as a theoretical foundation for the design of multiagent systems.

In recent years, the study of computational aspects of social choice received a lot of attention from AI and theoretical computer science communities. This interest was motivated by existing and potential applications of social choice ideas in AI settings, which, in turn, highlighted the importance of understanding which of the recommendations of social choice theory are computationally feasible.

The value of algorithmic analysis in the context of social choice stems from the fact that, to be practically applicable, a decision-making rule needs to be efficiently implementable. Indeed, the analysis of computational complexity of well-known voting rules, both in the general case, and in interesting special cases (such as, e.g., single-peaked preferences) is

one of the most actively studied topics in computational social choice, with a number of impressive results obtained so far.

However, computational tractability is not the only criterion for selecting a social choice procedure: an equally desirable feature is *incentive compatibility*, i.e., resilience to dishonest behavior by self-interested participants, who may want to manipulate the outcome of the procedure in their favor. There is an exciting interplay between incentive compatibility and computational tractability: in many settings of interest, computing one's optimal strategy requires solving a hard optimization problem, while acting honestly is computationally easy. Thus, one may view computational complexity as a barrier against strategic behavior, and try to design or identify social choice procedures that make strategizing difficult. This research direction was initiated more than 20 years ago and remains a major research focus of the computational social choice community.

Alternatively, one can deal with manipulative agents in the context of social choice by embracing the strategic behavior rather than trying to prevent it. This can be done either by investigating the outcomes of standard social choice procedures under the assumption that all agents act strategically, or, more ambitiously, by designing social choice procedures that result in desirable outcomes even if agents are not truthful; these two approaches are associated, respectively, with game theory and mechanism design. Both game-theoretic and mechanism design approaches are widely used by the classic social choice community; however, their computational aspects have received relatively little attention so far.

In contrast, algorithmic aspects of strategic behavior in other settings, such as, e.g., matrix games or auctions, have been studied extensively in the last few years. Indeed, computational game theory and algorithmic mechanism design are among the fastest-growing subfields of both AI and theoretical computer science. Thus, in organizing this seminar, we aimed to bring together the researchers in the areas of computational and classic social choice and those in the area of algorithmic game theory. Our goal was to foster a discussion of computational aspects of various forms of strategic behavior in social choice contexts.

Outcomes

The seminar took place on March 4–9, 2012. It was interdisciplinary in nature: among the participants, there were computer scientists, mathematicians, social choice theorists and political scientists. There were 32 regular talks, as well as an after-dinner talk by Virginia Vassilevska-Williams, who spoke about her groundbreaking work on algorithms for matrix multiplication. The seminar talks covered a broad range of topics, such as, e.g., the complexity of dishonest behavior in voting, judgement aggregation, coalitional game theory, and fair division. The program also featured a rump session consisting of short (5–8 minute) talks; these included announcements about events that were likely to be of interest to the seminar participants, short research talks, and presentations of open problems. The participants also used the seminar as an opportunity to continue ongoing research projects or start new ones. We are aware of two research papers that are largely based on discussions that happened during this Dagstuhl seminar; both of them have been recently submitted to the 4th International Workshop on Computational Social Choice. Moreover, several speakers who presented work in progress received useful feedback from other seminar participants, and, as a result, were able to improve or extend their papers significantly. To summarize, the participants of the seminar benefitted from it in a variety of ways: by being exposed to new research results and directions, by getting fresh perspectives on their work, by learning about open problems and initiating new collaborations, and by having an opportunity to work with their co-authors from all over the world on ongoing research projects.

2 Table of Contents

Executive Summary
Edith Elkind, Christian Klamler, Jeffrey S. Rosenschein, and M. Remzi Sanver . . 1

Overview of Talks

Bayesian Vote Manipulation: Optimal Strategies and Impact on Welfare
Craig Boutilier 5

When Does Approval Voting Make the "Right Choices"?
Steven J. Brams 5

On Measuring Nearly Single-Peakedness
Robert Bredereck 6

Possible and Necessary Winners of Partial Tournaments
Markus Brill 6

An incentive-compatible 2-agent kidney exchange mechanism
Ioannis Caragiannis 7

Parameterized Complexity Aspects of Optimal Lobbying
Jiehua Chen 7

Evaluating Resistance to False-Name Manipulations in Elections
Vincent Conitzer 8

Popular Spanning Trees
Andreas Darmann 8

A metatheorem for impossibility results in judgement aggregation
Daniel Eckert 9

On the geometry of voting rules with respect to the swap distance
Edith Elkind 9

Graph Aggregation
Ulle Endriss 9

Manipulation Under Voting Rule Uncertainty
Gábor Erdélyi 10

Exploring and Exploiting Clone Structures in Elections
Piotr Faliszewski 10

The Common Structure of Aggregation Paradoxes (and how to avoid them)
Umberto Grandi 11

Pareto Optimality in Coalition Formation
Paul Harrenstein 11

Search versus Decision for Election Manipulation Problems
Lane A. Hemaspaandra 11

A Simple Bargaining Mechanism That Induces Truth-Telling
Marc Kilgour 12

N-Person Cake-Cutting: There May Be No Perfect Division
Christian Klamler 12

Finding Extremal Voting Systems via Integer Linear Programming
Sascha Kurz 13

Judgment Aggregation Rules Based on Minimization
Jérôme Lang 13

Social Distance Games
Kate Larson 14

Convergence of Iterative Voting – Results & Problems
Omer Lev 14

On Worst-Case Allocations in the Presence of Indivisible Goods
Evangelos Markakis 15

Approximate Judgement Aggregation
Ilan Nehama 15

Optimal Voting Manipulation
Svetlana Obraztsova 16

Applying Social Choice Rules for the Solution of the Multi-Dimensional Knapsack Problem
Ulrich Pferschy 16

Bribery in Voting with CP-nets
Francesca Rossi 17

Bribery in Path-Disruption Games
Jörg Rothe 18

Control Complexity in Bucklin and Fallback Voting: A Theoretical and Experimental Analysis
Lena Schend 18

On the Computation of Fully Proportional Representation
Arkadii Slinko 19

The Structure, Efficacy and Manipulation of Double-Elimination Tournaments
Isabelle Stanton 19

Dividing the indivisible: elicitaton free protocols for the allocation of indivisible goods
Toby Walsh 20

Rump Session 20

Participants 22

3 Overview of Talks

3.1 Bayesian Vote Manipulation: Optimal Strategies and Impact on Welfare

Craig Boutilier (University of Toronto, CA)

Joint work of Lu, Tyler; Tang, Pingzhong; Procaccia, Ariel; Boutilier, Craig;

Most analyses of manipulation of voting schemes in computational social choice have considered the manipulation problem under two assumptions that greatly diminish their practical import. First, it is usually assumed that the manipulators have full knowledge of the votes of the nonmanipulating agents.

Second, analysis tends to focus on the probability of manipulation rather than its impact on social welfare. We relax both of these assumptions by analyzing optimal Bayesian manipulation strategies when the manipulators have only partial probabilistic information about nonmanipulator votes, and assessing the expected loss in social welfare. We present a general optimization framework for the derivation of optimal manipulation strategies given arbitrary voting rules and distributions over preferences. We theoretically and empirically analyze the optimal manipulability of some popular voting rules using distributions and real data sets that go well beyond the common, but unrealistic, impartial culture assumption. We also shed light, both theoretically and empirically, on the stark difference between the loss in social welfare and the probability of manipulation by showing that even when manipulation is likely, impact to social welfare is slight (and often negligible).

3.2 When Does Approval Voting Make the "Right Choices"?

Steven J. Brams (New York University, US)

Joint work of Brams, Steven J.; Kilgour, D. Marc
Main reference Working paper
URL http://politics.as.nyu.edu/object/stevenbrams

We assume that a voter's approval of a proposal depends on (i) the proposal's probability of being right (or good or just) and (ii) the voter's probability of making a correct judgment about its rightness (or wrongness). The state of a proposal (right or wrong), and the correctness of a voter's judgment about it, are assumed, initially, to be independent. If the average probability that voters are correct in their judgments is greater than $\frac{1}{2}$, then the proposal with the greatest probability of being right will, in expectation, receive the greatest number of approval votes. This result also holds when voters' probabilities of being correct are state dependent but not proposal dependent; when they are functionally related in a certain way; or when voters follow a leader with an above-average probability of correctly judging proposals. Sometimes, however, voters will more frequently select the right proposal by not following a leader and, instead, making their own independent judgments (as assumed by the Condorcet jury theorem). Applications of these results to different kinds of voting situations are discussed.

3.3 On Measuring Nearly Single-Peakedness

Robert Bredereck (TU Berlin, DE)

Many problems in context of voting are NP-hard in general. However, when the elections are single-peaked some voting problems become polynomial-time solvable. Often, real-world elections are not perfectly single-peaked, because some voters behave unexpectedly or few candidates do not fit into the model. In our work in progress, we investigate two distances measuring almost single-peakedness. The first distance is "the number of voters to remove to make the election single-peaked", which is also known as number of mavericks in the literature. The second distance is "the number of candidates to remove to make the election single-peaked". We show NP-hardness for the first distance as well as fixed-parameter algorithms computing both distances. Furthermore, we show that there exist effective data reduction procedures (leading to so-called polynomial-size problem kernels) useful for computing these distances (and the corresponding solution sets).

3.4 Possible and Necessary Winners of Partial Tournaments

Markus Brill (TU München, DE)

Joint work of Aziz, Haris; Brill, Markus; Fischer, Felix; Harrenstein, Paul; Lang, Jérôme; Seedig, Hans Georg
Main reference H. Aziz, M. Brill, F. Fischer, P. Harrenstein, J. Lang, H. G. Seedig, "Possible and necessary winners of partial tournaments," in V. Conitzer and M. Winikoff, (eds.), Proc. of the 11th Int'l Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS). IFAAMAS, 2012.
URL http://dss.in.tum.de/files/brandt-research/partial.pdf

We study the problem of computing possible and necessary winners for partially specified weighted and unweighted tournaments. This problem arises naturally in elections with incompletely specified votes, partially completed sports competitions, and more generally in any scenario where the outcome of some pairwise comparisons is not yet fully known. We specifically consider a number of well-known solution concepts–including the uncovered set, Borda, ranked pairs, and maximin–and show that for most of them possible and necessary winners can be identified in polynomial time. These positive algorithmic results stand in sharp contrast to earlier results concerning possible and necessary winners given partially specified preference profiles.

3.5 An incentive-compatible 2-agent kidney exchange mechanism

Ioannis Caragiannis (University of Patras, GR)

Joint work of Caragiannis, Ioannis; Filos-Ratsikas, Aris; Procaccia, Ariel D.
Main reference I. Caragiannis, A. Filos-Ratsikas, A. D. Procaccia, "An improved 2-agent kidney exchange mechanism," in Proc. of the 5th Int'l Workshop on Internet and Network Economics (WINE'11), LNCS 7090, Springer, pp. 37–48, 2011.
URL http://dx.doi.org/10.1007/978-3-642-25510-6_4

We consider a mechanism design version of matching computation in graphs that models the game played by hospitals participating in pairwise kidney exchange programs. We present a new randomized matching mechanism for two agents which is truthful in expectation and has an approximation ratio of 3/2 to the maximum cardinality matching. This is an improvement over a recent upper bound of 2 [Ashlagi et al., EC 2010] and, furthermore, our mechanism beats for the first time the lower bound on the approximation ratio of deterministic truthful mechanisms. We complement our positive result with new lower bounds. Among other statements, we show that the weaker incentive compatibility property of truthfulness in expectation in our mechanism is necessary; universally truthful mechanisms that have an inclusion-maximality property have an approximation ratio of at least 2.

3.6 Parameterized Complexity Aspects of Optimal Lobbying

Jiehua Chen (TU Berlin, DE)

Joint work of Bredereck, Robert; Chen, Jiehua; Hartung, Sepp; Kratsch, Stefan; Niedermeier, Rolf; Suchý, Ondřej

In a multi-issue election, each voter may "approve" or "disapprove" each individual issue. In the context of attacks on multi-issue elections, Optimal Lobbying asks whether a lobbyist can lobby a given number of voters such that for each issue, there is a majority of voters who vote in favor of the lobbyist.

Here, lobbying a voter means changing this voter's vote completely into the lobbyist's reference. In general, Optimal Lobbying is computational intractable (NP-complete). However, is this problem always hard for realistic scenarios? To address this, Christian et al. [Review of Economic Design 2007] studied the parameterized complexity of Optimal Lobbying and showed that this problem is W[2]-complete in the parameter "number of voters to lobby". In this talk, we look into several additional parameters which describe the structure of the input or the distance to the lobbyist's goal. We gain both tractability and intractability results. We also present an efficient greedy algorithm which solves our problem optimally if the number of issues is at most four.

3.7 Evaluating Resistance to False-Name Manipulations in Elections

Vincent Conitzer (Duke University, US)

Joint work of Waggoner, Bo; Xia, Lirong; Conitzer, Vincent
Main reference B. Waggoner, L. Xia, V. Conitzer, "Evaluating Resistance to False-Name Manipulations in Elections," in AAAI'12, *to appear*.

In many mechanisms (especially online mechanisms), a strategic agent can influence the outcome by creating multiple false identities. We consider voting settings where the mechanism designer cannot completely prevent false-name manipulation, but may use false-name-limiting methods such as CAPTCHAs to influence the amount and characteristics of such manipulation. Such a designer would prefer, first, a high probability of obtaining the "correct" outcome, and second, a statistical method for evaluating the correctness of the outcome. In this paper, we focus on settings with two alternatives. We model voters as independently drawing a number of identities from a distribution that may be influenced by the choice of the false-name-limiting method. We give a criterion for the evaluation and comparison of these distributions. Then, given the results of an election in which false-name manipulation may have occurred, we propose and justify a statistical test for evaluating the outcome.

3.8 Popular Spanning Trees

Andreas Darmann (Universität Graz, AT)

The considered problem combines Combinatorial Optimization with Social Choice Theory. In classic Combinatorial Optimization, costs are assigned to the edges of an undirected graph and one is interested in finding a spanning tree of minimum total cost. In our approach, instead of associating costs with the edges of the undirected graph, it is assumed that individuals have preferences over the single edges. A spanning tree is proposed by an external source (e.g., a central authority), and the goal is to decide on the fairness (or quality) of the proposed solution. Given the individual preferences over the edges, we evaluate quality by means of a Condorcet criterion. In particular, we perform comparisons between spanning trees that are based on scoring functions used in classic voting rules such as approval voting and Borda voting. The focus of our work is laid on the computational complexity involved in deciding on the quality of a spanning tree with respect to the different voting rules adapted.

With our results, the sharp separation line between polynomially solvable and computationally intractable instances is drawn.

3.9 A metatheorem for impossibility results in judgement aggregation

Daniel Eckert (Universität Graz, AT)

Joint work of Eckert, Daniel; Herzberg, Frederik

The close relation between the two major impossibility results in social choice theory, Arrow's "general possibility theorem" and the Gibbard- Satterthwaite theorem, has been explored in several metatheorems. In a model theoretic framework, an analogous metatheorem for impossibility results in the recent literature on judgment aggregation is provided.

3.10 On the geometry of voting rules with respect to the swap distance

Edith Elkind (Nanyang TU Singapore, SG)

Joint work of Obraztsova, Svetlana; Elkind, Edith; Faliszewski, Piotr; Slinko, Arkadii

Axioms that govern our choice of voting rules are usually defined by imposing constraints on the rule's behavior under various transformations of the preference profile. In this paper we adopt a different approach, and view a voting rule as a (multi-)coloring of the election graph—the graph whose vertices are elections over a given set of candidates, and two vertices are adjacent if they can be obtained from each other by swapping adjacent candidates in one of the votes. Given this perspective, a voting rule F is characterized by the shapes of its "monochromatic components", i.e., the sets of elections that have the same winner under F. In particular, it would be natural to expect each monochromatic component to be convex, or, at the very least, connected. We formalize the notions of connectivity and (weak) convexity for monochromatic components, and say that a voting rule is connected/(weakly) convex if each of its monochromatic components is connected/(weakly) convex. We then investigate which of the classic voting rules have these properties. It turns out that while all voting rules that we consider are connected, convexity and even weak convexity are much more demanding properties. Our study of connectivity suggests a new notion of monotonicity, which may be of independent interest.

3.11 Graph Aggregation

Ulle Endriss (University of Amsterdam, NL)

Joint work of Endriss, Ulle; Grandi, Umberto

Suppose several agents each provide us with a directed graph on the same set of vertices. Graph aggregation is the problem of computing a single collective graph that best represents the information inherent in this profile of individual graphs. A procedure to perform this kind of aggregation is called collectively rational with respect to a given property if it is the case that, whenever every individual graph satisfies the property in question, then so does the collective graph the procedure is going to return. We set up a formal framework

for analysing collective rationality in graph aggregation and discuss several possibility and impossibility results.

3.12 Manipulation Under Voting Rule Uncertainty

Gábor Erdélyi (Universität Siegen, DE)

Joint work of Elkind, Edith; Erdélyi, Gábor;
Main reference E. Elkind, G. Erdélyi, "Manipulation Under Voting Rule Uncertainty," in V. Conitzer and M. Winikoff, (eds.), Proc. of the 11th Int'l Joint Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS). IFAAMAS, 2012.

An important research topic in the field of computational social choice is the complexity of various forms of dishonest behavior, such as manipulation, control, and bribery. While much of the work on this topic assumes that the cheating party has full information about the election, recently there have been a number of attempts to gauge the complexity of non-truthful behavior under uncertainty about the voters' preferences. In this paper, we analyze the complexity of (coalitional) manipulation for the setting where there is uncertainty about the voting rule: the manipulator(s) know that the election will be conducted using a voting rule from a given list, and need to select their votes so as to succeed no matter which voting rule will eventually be chosen. We identify a large class of voting rules such that arbitrary combinations of rules from this class are easy to manipulate; in particular, we show that this is the case for single-voter manipulation and essentially all easy-to-manipulate voting rules, and for coalitional manipulation and k-approval. While a combination of a hard-to-manipulate rule with an easy-to-manipulate one is usually hard to manipulate—we prove this in the context of coalitional manipulation for several combinations of prominent voting rules—we also provide counterexamples showing that this is not always the case.

3.13 Exploring and Exploiting Clone Structures in Elections

Piotr Faliszewski (AGH University of Science and Technology Krakow, PL)

Joint work of Elkind, Edith; Faliszewski, Piotr; Slinko, Arkadii
Main reference E. Elkind, P. Faliszewski, A. Slinko, "Clone Structures in Voters' Preference," in Proc. of 13th ACM Conf. on Electronic Commerce, *to appear*.

In election, a clone set is a subset of candidates ranked consecutively by all voters. A clone structure of a given election is a family of all its clone sets. In this talk we will consider the following issues: Given an election and some beliefs as to which clone sets are a result of cloning, how to reconstruct the most likely original election? How to reconstruct the most likely election that ensures some given candidate's victory? If the election was originally single-peaked (or single-crossing), is it possible to discover this single-peakedness (single-crossingness) while "decloning" as few candidates as possible? To answer these questions, we will explore the landscape of possible clone structures in elections.

3.14 The Common Structure of Aggregation Paradoxes (and how to avoid them)

Umberto Grandi (University of Amsterdam, NL)

In this talk I will analyse some of the classical paradoxes in Social Choice Theory (namely, the Condorcet paradox, the discursive dilemma, the Ostrogorski paradox and the multiple election paradox) using a general framework for the study of aggregation problems called binary aggregation with integrity constraints. I will provide a definition of paradox that is general enough to account for the four cases mentioned, and identify a common structure in the syntactic properties of the rationality assumptions that lie behind such paradoxes. I will conclude by introducing an aggregation procedure that avoids paradoxical situations for any given rationality assumption called the average voter rule. I investigate its axiomatic properties and the computational complexity of both the problem of winner determination and strategic manipulation.

3.15 Pareto Optimality in Coalition Formation

Paul Harrenstein (TU München, DE)

Joint work of Aziz, Haris; Brandt, Felix; Harrenstein, Paul
Main reference H. Aziz, F. Brandt, P. Harrenstein, "Pareto optimality in coalition formation," in G. Persiano, (ed.), Proc. of the 4th Int'l Symp. on Algorithmic Game Theory (SAGT), LNCS, vol. 6982, pp. 93–104. Springer-Verlag, 2011.
URL http://dx.doi.org/10.1007/978-3-642-24829-0_10

A minimal requirement on allocative efficiency in the social sciences is Pareto optimality. In this paper, we identify a close structural connection between Pareto optimality and perfection that has various algorithmic consequences for coalition formation. Based on this insight, we formulate the Preference Refinement Algorithm (PRA) which computes an individually rational and Pareto optimal outcome in hedonic coalition formation games or any other discrete allocation setting. Our approach also leads to various results for specific classes of hedonic games. In particular, we show that computing and verifying Pareto optimal partitions in general hedonic games, anonymous games, three-cyclic games, room-roommate games and B-hedonic games is intractable while both problems are tractable for roommate games, W-hedonic games, and house allocation with existing tenants.

3.16 Search versus Decision for Election Manipulation Problems

Lane A. Hemaspaandra (University of Rochester, US)

Joint work of Hemaspaandra, Edith; Hemaspaandra, Lane A.; Menton, Curtis

Most theoretical definitions about the complexity of manipulating elections focus on the decision problem of recognizing which instances can be successfully manipulated, rather

than the search problem of finding the successful manipulative actions. Since the latter is a far more natural goal for manipulators, that definitional focus may be misguided if these two complexities can differ. Our main result is that they probably do differ: If integer factoring is hard, then for election manipulation, election bribery, and some types of election control, there are election systems for which recognizing which instances can be successfully manipulated is in polynomial time but producing the successful manipulations cannot be done in polynomial time.

3.17 A Simple Bargaining Mechanism That Induces Truth-Telling

Marc Kilgour (Wilfrid Laurier University, CA)

Joint work of Kilgour, Marc; Brams, Steven J.; Kaplan, Todd R.;
Main reference D. M. Kilgour, S. J. Brams, T. R. Kaplan, "Three procedures for inducing honesty in bargaining," in Proc. of 13th Conf. on Theoretical Aspects of Rationality and Knowledge (TARK'11), pp. 170–176, ACM, 2011.
URL http://dx.doi.org/10.1145/2000378.2000398

No bargaining mechanism can induce bargainers to report their reservation prices (bottom lines) truthfully. Several mechanisms that come close to achieving perfect efficiency are reviewed, including a new 2-stage mechanism that induces two bargainers to report truthfully in a 1st stage. If these prices criss-cross, the referee reports that they overlap, and the bargainers proceed to make offers in a 2nd stage. The average of the 2nd-stage offers becomes the settlement if both offers fall into the overlap interval; if only one offer falls into this interval, it is the settlement, but is implemented with probability $\frac{1}{2}$; if neither offer falls into the interval, there is no settlement. Thus, if the bargainers reach the 2nd stage, they know their reservation prices overlap even if they fail to reach a settlement, possibly motivating them to try again.

3.18 N-Person Cake-Cutting: There May Be No Perfect Division

Christian Klamler (Universität Graz, AT)

Main reference S. J. Brams, M. A. Jones, C. Klamler, "N-Person Cake-Cutting: There May Be No Perfect Division," October 2011, available at SSRN.
URL http://dx.doi.org/10.2139/ssrn.1946993

A cake is a metaphor for a heterogeneous, divisible good, such as land.

Over the past fifteen years, a substantial literature on cake-cutting has sprung up. In particular, three properties of cake-cutting algorithms have been the focus of most of cake-cutting literature. Efficiency (also called Pareto-optimality) requires that there is no other division that gives players portions that they value at least as much and gives at least one player strictly more. Envy-freeness states that each player values its portion at least as much as that of every other player and, consequently, does not envy any other player. Finally, a division is considered to be equitable if each player values its portion exactly the same as everybody else values its portion, i.e., each player thinks that its portion is the same fraction of its perceived value of the entire cake. In this paper, we consider a division of a cake that

satisfies all three properties to be a perfect division. We give an example of a cake in which it is impossible to divide it among three players such that these three properties are satisfied, however many cuts are made. It turns out that two of the three properties can be satisfied by a 3-cut and a 4-cut division, which raises the question of whether the 3-cut division, which is not efficient, or the 4-cut division, which is not envy-free, is more desirable (a 2-cut division can at best satisfy either envy-freeness or equitability but not both). We prove that no perfect division exists for more than 4 cuts and for an extension of this example to more than three players.

3.19 Finding Extremal Voting Systems via Integer Linear Programming

Sascha Kurz (Universität Bayreuth, DE)

Main reference S. Kurz, "On the inverse power index problem," *to appear* in Optimization, 23 pages, 2012.
URL http://dx.doi.org/10.1080/02331934.2011.587008

Different types of yes/no voting systems are frequently studied in the literature. Using integer linear programming we determine extremal, according to a given criterion, voting systems.

As examples we consider voting systems

- whose Shapley-Shubik vector has minimal distance to a given power distribution;
- whose Public Good Index maximally violate local monotonicity;
- which are farthest away from weighted voting games with respect to a recently introduced hierarchy of simple games.

We present the general underlying ideas and computational results for instances where exhaustive enumeration of all voting systems is infeasible.

3.20 Judgment Aggregation Rules Based on Minimization

Jérôme Lang (Université Paris-Dauphine, FR)

Joint work of Lang, Jérôme; Pigozzi, Gabriella; Slavkovik, Marija; van der Torre, Leon
Main reference J. Lang, G. Pigozzi, M. Slavkovik, L. van der Torre, "Judgment aggregation rules based on minimization," in Proc. of 13th Conf. on Theoretical Aspects of Rationality and Knowledge (TARK'11), pp. 238–246, ACM, 2011.
URL http://dx.doi.org/10.1145/2000378.2000407

Many voting rules are based on a minimization or maximization principle.

Likewise, in the field of logic-based knowledge representation and reasoning, many belief change or inconsistency handling operators make use of minimization.

Surprisingly, minimization has not played a major role in the field of judgment aggregation, in spite of its proximity to voting theory and logic-based knowledge representation and reasoning. Here we the study judgment aggregation rules based on minimization, and propose a classification of judgment aggregation rules based on some minimization or maximization principle. We distinguish four families of rules. While most of these rules are new, a few ones correspond to rules that have been defined elsewhere.

We study the inclusion relationships among these rules, and analyze them with respect to the common judgment aggregation properties.

3.21 Social Distance Games

Kate Larson (University of Waterloo, CA)

Joint work of Larson, Kate; Branzei, Simina
Main reference S. Branzei, K. Larson, "Social Distance Games," in Proc. of the 22nd Int'l Joint Conf. on Artificial Intelligence (IJCAI'11), pp. 91–96, IJCAI/AAAI, 2011.
URL http://ijcai.org/papers11/Papers/IJCAI11-027.pdf

In this paper we introduce and analyze social distance games, a family of non-transferable utility coalitional games where an agent's utility is a measure of closeness to the other members of the coalition. We study both social welfare maximisation and stability in these games using a graph theoretic perspective.

We use the stability gap to investigate the welfare of stable coalition structures, and propose two new solution concepts with improved welfare guarantees. We argue that social distance games are both interesting in themselves, as well as in the context of social networks.

3.22 Convergence of Iterative Voting – Results & Problems

Omer Lev (The Hebrew University of Jerusalem, IL)

Joint work of Lev, Omer; Rosenschein, Jeffrey S.
Main reference O. Lev, J. S. Rosenschein, "Convergence of Iterative Voting," in Proc. of the 11th Int'l Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS'12), June 2012, Valencia, Spain.
URL ftp://ftp.cs.huji.ac.il/users/jeff/aamas12lev.pdf

In multiagent systems, social choice functions can help aggregate the distinct preferences that agents have over alternatives, enabling them to settle on a single choice. Despite the basic manipulability of all reasonable voting systems, it would still be desirable to find ways to reach a stable result, i.e., a situation where no agent would wish to change its vote. One possibility is an iterative process in which, after everyone initially votes, participants may change their votes, one voter at a time. This technique, explored in previous work, converges to a Nash equilibrium when Plurality voting is used, along with a tie-breaking rule that chooses a winner according to a linear order of preferences over candidates.

In this work, we both consider limitations of the iterative voting method, as well as expanding upon it. We demonstrate the significance of tie-breaking rules, showing that when using a general tie-breaking rule, no scoring rule (nor Maximin) needs to iteratively converge. However, using a restricted tie-breaking rule (such as the linear order rule used in previous work) does not by itself ensure convergence. We demonstrate that many scoring rules (such as Borda) need not converge, regardless of the tie-breaking rule. On a more encouraging note, we prove that Iterative Veto does converge—but that voting rules "between" Plurality and Veto, k-approval rules, do not.

3.23 On Worst-Case Allocations in the Presence of Indivisible Goods

Evangelos Markakis (Athens University of Economics and Business, GR)

Joint work of Markakis, Evangelos; Psomas, Christos-Alexandros
Main reference E. Markakis, C.-A. Psomas, "On Worst-Case Allocations in the Presence of Indivisible Goods," in Proc. of the 7th Int'l Workshop on Internet and Network Economics (WINE'11), pp. 278–289, LNCS, vol. 7090, Springer, 2011.
URL http://dx.doi.org/10.1007/978-3-642-25510-6_24

We study a fair division problem, where a set of indivisible goods is to be allocated to a set of n agents. Each agent may have different preferences, represented by a valuation function that is a probability distribution on the set of goods. In the continuous case, where goods are infinitely divisible, it is well known that proportional allocations always exist, i.e., allocations where every agent receives a bundle of goods worth to him at least 1/n. In the presence of indivisible goods however, this is not the case and one would like to find worst case guarantees on the value that every agent can have. We focus on algorithmic and mechanism design aspects of this problem. In the work of [Hill, 1987], an explicit lower bound was identified, as a function of the number of agents and the maximum value of any agent for a single good, such that for any instance, there exists an allocation that provides at least this guarantee to every agent. The proof however did not imply an efficient algorithm for finding such allocations. Following upon the work of Hill, we first provide a slight strengthening of the guarantee we can make for every agent, as well as a polynomial time algorithm for computing such allocations. We then move to the design of truthful mechanisms. For deterministic mechanisms, we obtain a negative result showing that a truthful 2/3-approximation of these guarantees is impossible. We complement this by exhibiting a simple truthful algorithm that can achieve a constant approximation when the number of goods is bounded. Regarding randomized mechanisms, we also provide a negative result, showing that we cannot have truthful in expectation mechanisms under the restrictions that they are Pareto-efficient and satisfy certain symmetry requirements.

3.24 Approximate Judgement Aggregation

Ilan Nehama (The Hebrew University of Jerusalem, IL)

Joint work of Nehama, Ilan
Main reference I. Nehama, "Approximate Judgement Aggregation," Discussion Paper 574R, Center for the Study of Rationality, Hebrew University, Jerusalem, 2012.
URL http://www.ratio.huji.ac.il/dp_files/dp574R.pdf

In this paper we analyze judgement aggregation problems in which a group of agents independently votes on a set of complex propositions that has some interdependency constraint between them(e.g., transitivity when describing preferences). We consider the issue of judgement aggregation from the perspective of approximation. That is, we generalize the previous results by studying approximate judgement aggregation. We relax the main two constraints assumed in the current literature, Consistency and Independence and consider mechanisms that only approximately satisfy these constraints, that is, satisfy them up to a small portion of the inputs. The main question we raise is whether the relaxation of these notions significantly alters the class of satisfying aggregation mechanisms. The recent

works for preference aggregation of Kalai, Mossel, and Keller fit into this framework. The main result of this paper is that, as in the case of preference aggregation, in the case of a subclass of a natural class of aggregation problems termed 'truth-functional agendas', the set of satisfying aggregation mechanisms does not extend non-trivially when relaxing the constraints. Our proof techniques involve Boolean Fourier transform and analysis of voter influences for voting protocols.

The question we raise for Approximate Aggregation can be stated in terms of Property Testing. For instance, as a corollary from our result we get a generalization of the classic result for property testing of linearity of Boolean functions.

3.25 Optimal Voting Manipulation

Svetlana Obraztsova (St. Petersburg Electrotechnical University, RU)

Joint work of Obraztsova, Svetlana; Elkind, Edith;

Complexity of voting manipulation is a prominent research topic in computational social choice. In this talk the complexity of optimal manipulation, i.e., finding a manipulative vote that achieves the manipulator's goal yet deviates as little as possible from his true ranking, was discussed. This problem was studied for three natural notions of closeness, namely, swap distance, footrule distance, and maximum displacement distance, and a variety of voting rules, such as scoring rules, Bucklin, Copeland, and Maximin. For all three distances, poly-time algorithms for all scoring rules and Bucklin and hardness results for Copeland and Maximin were showed.

3.26 Applying Social Choice Rules for the Solution of the Multi-Dimensional Knapsack Problem

Ulrich Pferschy (Universität Graz, AT)

Joint work of Nussbaumer, Martin; Pferschy, Ulrich;

The multi-dimensional knapsack problem (MKP) considers a set of items, each of them with a profit and a d-dimensional weight vector, and asks for the selection of a subset of items with maximum total profit, such that the sum of weights in each dimension fulfills a capacity constraint implied by a d-dimensional capacity vector. This generalization of the standard knapsack problem (KP) with a single constraint is surprisingly difficult to solve in practice. Even relatively small benchmark instances with 500 items and 10 constraints still cannot be solved to proven optimality.

Many approaches tackling MKP make use of an ordering of items based on a generalization of the efficiency measure usually applied for KP. This measure simply calculates the profit to weight ratio for each item. However, the presence of multiple constraints requires an aggregation of the d weights for each item. Different approaches for the resulting efficiency coefficients were presented in the literature. The most successful among them is based on the optimal dual variables from the associated LP-relaxation.

In this contribution, we propose a different approach to obtain a valuation of each item. We consider each of the d weight dimensions as a voter who gives a preference relation on the set of items based only on the single constraint the voter is associated with. Then we use different voting rules to derive a complete ordering of items based on the aggregated preference profile of the d voters. This ordering is then used e.g. for greedy-type heuristics.

Computational experiments show that this new approach of applying social choice techniques for the solution of a classical combinatorial optimization problem produces reasonably good solutions and offers a highly welcome element of diversification for metaheuristic frameworks.

Squeaky wheel optimization (SWO) is a metaheuristic which was successfully applied to MKP. In particular, SWO turned out to be useful in reducing an MKP instance to a core, i.e. fixing some of the variables to 0 or 1 and leaving only a smaller instance for further treatment (e.g. by an ILP-solver). Since SWO is based on an initial ordering and the subsequent reordering of the item set, the above voting-based procedure is well suited for obtaining alternative initial orderings. Computational experiments show that the resulting cores compare favorably with those based on more involved LP-solution values.

3.27 Bribery in Voting with CP-nets

Francesca Rossi (University of Padova, IT)

Joint work of Mattei, Nicholas; Venable, Kristen Brent; Pini, Maria Silvia; Rossi, Francesca;
Main reference N. Mattei, F. Rossi, K. B. Venable, M. S. Pini, "Bribery in Voting Over Combinatorial Domains Is Easy," in Proc. of the 11th Int'l Joint Conf. on Autonomous Agents and Multiagent Systems (AAMAS'12), Extended Abstracft, June 2012, Valencia, Spain.

We investigate the computational complexity of finding optimal bribery schemes in voting domains where candidates are multi-issue decisions and agents' preferences are represented as CP-nets.

In this setting, voting can be structured as the combination of several decisions, or it can be a one-step process.

We consider both approaches, by studying voting rules such as sequential majority (SM), one-step plurality (OP), one-step veto (OV), and one-step k-approval (OK). We then consider several cost schemes for changing a vote of an agent in response to a briber's request, among which:

C_{EQUAL} (any amount of change costs the same),

C_{FLIP} (the cost is the number of flips),

C_{LEVEL} (the cost is the number of flips weighted by their position in the CP-net).

SM bribery is easy except when we use C_{EQUAL}. For OP and OV, bribery is always easy, except with C_{FLIP} or C_{LEVEL} when we can flip dependent variables. Bribery is easy also for OK when k is a power of 2.

3.28 Bribery in Path-Disruption Games

Jörg Rothe (Universität Düsseldorf, DE)

Joint work of Rey, Anja; Rothe, Jörg
Main reference A. Rey, J. Rothe, "Bribery in Path-Disruption Games," in Proc. of the 2nd Int'l Conf. on Algorithmic Decision Theory (ADT'11), LNCS, vol. 6992, pp. 247–261, Springer-Verlag, 2011.
URL http://dx.doi.org/10.1007/978-3-642-24873-3_19

Bachrach and Porat (AAMAS 2010) introduced path-disruption games. In these coalitional games, agents are placed on the vertices of a graph, and one or more adversaries want to travel from a source vertex to a target vertex. In order to prevent them from doing so, the agents can form coalitions, and a coalition wins if it succeeds in blocking all paths for the adversaries. In this paper, we introduce the notion of bribery for path-disruption games. We analyze the question of how hard it is to decide whether the adversaries can bribe some of the agents such that no coalition can be formed that blocks all paths for the adversaries. We show that this problem is NP-complete, even for a single adversary. For the case of multiple adversaries, we provide an upper bound by showing that the corresponding problem is in Σ_2^p, the second level of the polynomial hierarchy, and we suspect it is complete for this class.

3.29 Control Complexity in Bucklin and Fallback Voting: A Theoretical and Experimental Analysis

Lena Schend (Universität Düsseldorf, DE)

Joint work of Erdélyi, Gábor; Rothe, Jörg; Schend, Lena;
Main reference J. Rothe, L. Schend, "Control Complexity in Bucklin, Fallback, and Plurality Voting: An Experimental Approach," in Proc. of the 11th Int'l Symp. on Experimental Algorithms (SEA'12), LNCS, vol. 7276, pp. 356–368, Springer-Verlag, 2012.
URL http://dx.doi.org/10.1007/978-3-642-30850-5_31

In [1] we complete the study of control complexity in fallback voting (FV) initiated by Erdelyi and Rothe. FV displays the broadest resistance to control currently known to hold among natural voting systems with a P-time winner problem. We also prove that Bucklin voting (BV) behaves almost as good in terms of control resistance.

Complementary to these worst-case results, an experimental analysis for FV and BV has been made inspired by Walsh's empirical investigation of manipulation complexity. Our findings indicate that NP-hard control problems can often be solved effectively in practice. Moreover, our experiments allow a more fine-grained analysis and comparison across various control scenarios, vote distribution models, and voting systems.

References

1 G. Erdélyi, L. Piras, J. Rothe. *The Complexity of Voter Partition in Bucklin and Fallback Voting: Solving Three Open Problems.* Proc. AAMAS'11. IFAAMAS, pages 837–844, 2011.

3.30 On the Computation of Fully Proportional Representation

Arkadii Slinko (University of Auckland, NZ)

My talk will consist of two parts. Firstly, I will outline challenges that Computational Social Choice faces in case of multi-winner elections. There is no such thing as an ideal voting system and we must sacrifice something but what are the trade-offs?

In the second part I will dwell on some parameterized complexity results in relation to fully proportional representation methods of Chamberlin-Courant and Monroe (joint work with Nadja Betzler and Johannes Ulhmann).

3.31 The Structure, Efficacy and Manipulation of Double-Elimination Tournaments

Isabelle Stanton (University of California, Berkeley, US)

Joint work of Stanton, Isabelle; Vassilevska Williams, Virginia

A double-elimination (DE) tournament is a competition where no participant is eliminated until they have lost two matches. It is structured as two single-elimination tournaments: the winner bracket and the loser bracket. Players who lose once in the winner bracket are mapped to positions in the loser bracket, according to a mapping called the link function. Surprisingly, although the same structure of the winner and loser brackets is used universally, there is no standard definition of the link function. By investigating several design goals, we show that the functions used in practice are not optimal and propose a similar function that is optimal with respect to avoiding repeated match-ups. We empirically show that use of the new link function does not impact the ability of a DE tournament to select a strong winner. Given our definitions, we address the manipulability of DE tournaments. We show that they are vulnerable to manipulation by a coalition of players who can improve their chance of winning by throwing matches. We also discuss the computational complexity of manipulation by a tournament organizer (agenda control) in two settings: by changing the player seeding in the winner bracket, or by picking the mapping of losers to the loser bracket. We provide algorithms, hardness proofs, and we formulate open problems for future research. Finally, we empirically compare single and double-elimination tournaments in terms of the probability that the strongest player wins the tournament and show that this probability can be drastically higher in DE tournaments, confirming the intuition that DE tournaments are more robust than SE tournaments.

3.32 Dividing the indivisible: elicitaton free protocols for the allocation of indivisible goods

Toby Walsh (NICTA and University of New South Wales, Sydney, AU)

Joint work of Kalinowski, Thomas; Narodytska, Nina; Walsh, Toby; Xia, Lirong

We study in detail a simple sequential procedure for allocating a set of indivisible goods to multiple agents. Agents take turns to pick items according to a policy. For example, in the alternating policy, agents simply alternate who picks the next item. A similar procedure has been used by Harvard Business School to allocate courses to students. When agents behave truthfully, this sequential allocation procedure returns precisely the Pareto optimal allocations. Supposing agents behave truthfully is a strong assumption. Indeed, strategic behavior has been observed in students selecting courses at the Harvard Business School. We study therefore the impact of strategic behavior on the complete information extensive-form game of such sequential allocation procedures. We show that computing the subgame-perfect Nash equilibrium is PSPACE-hard in general, but takes only linear time with two agents. Finally we compute the optimal policies for two agents in different settings, including when agents behave strategically and when agents can give away items.

4 Rump Session

The first presentation in the rump session was by Marcel Ackermann, who talked about recent developments concerning the DBLP database and asked the seminar participants to share their opinions and experiences using DBLP. Dr. Ackermann was available for discussions in the evening of the same day, and many seminar participants used this opportunity to talk to him.

The talk of Dr. Ackermann was followed by two conference announcements: Vincent Merlin reminded the participants about the 11th Meeting of the Society for Social Choice and Welfare, to take place in New Delhi in August 2012, and Felix Brandt and Piotr Faliszewski gave a brief presentation about the Fourth International Workshop on Computational Social Choice, to be held in Krakow in September 2012.

The announcements were followed by 9 short research talks.

- Vangelis Markakis talked about approximation algorithms for maxsum and minmax procedures in the election of committees; he mentioned several conjectures about the lower and upper bounds for this problem.
- Toby Walsh introduced various prices (of manipulation, information or computation) in voting and suggested them as a means to compare voting rules.
- Andreas Darmann considered the problem of deciding whether a given spanning tree is popular (in the sense of a weak Condorcet criterion) given that agents have preferences over the edges of an undirected graph. Whereas it is possible to draw a sharp separation between polynomially solvable and computationally intractable instances, the computational complexity of the existence of a popular spanning tree is still open.
- Ioannis Cariagiannis talked about complexity issues in bribery problems under scoring rules with scoring vectors of the form $(s, t, 0, \ldots, 0)$. Results for plurality, 2-approval and 3-approval are known, but the general case remains open.

- Ulle Endriss asked whether opinion polls provide relevant information in elections. He argued that the answer is positive for the plurality rule; however, it is unclear whether this is also the case for other rules.
- Jérôme Lang suggested that the problem of selecting the social activity for the free afternoon in Dagstuhl can be viewed as a social choice problem, proposed a formal model for it, and mentioned several research questions that can be stated within this model. He invited the participants of the seminar to contact him if they are interested in working on this problem. As a result, 5 seminar participants (including Jerome himself) and one external co-author wrote a paper about this problem that was submitted to COMSOC'12.
- Edith Elkind talked about open problems regarding the complexity of finding a safe strategic vote, as defined in the COMSOC'08 paper of Slinko and White. While the complexity of this problem has been resolved for a large class of scoring rules (including Borda and k-approval) and the Bucklin rule, for Maximin and Copeland this question remains open. She has also mentioned the problem of finding the Condorcet dimension of a given profile (this notion was introduced in the IJCAI'11 paper by Elkind, Lang and Saffidine).
- Vincent Conitzer introduced a new measure of manipulability of a voting rule, which is based on comparing the benefits from submitting a non-truthful vote and those from being able to submit multiple truthful votes. He showed how to compute this measure for some simple voting rules; for others, the associated algorithmic question is open.
- Craig Boutilier talked about matching models for preference-sensitive group purchasing: How should buyers be assigned to vendors?

Participants

- Craig Boutilier
University of Toronto, CA
- Steven J. Brams
New York University, US
- Felix Brandt
TU München, DE
- Robert Bredereck
TU Berlin, DE
- Markus Brill
TU München, DE
- Ioannis Caragiannis
CTI & University of Patras, GR
- Jiehua Chen
TU Berlin, DE
- Yann Chevaleyre
Université Paris-Dauphine, FR
- Vincent Conitzer
Duke University, US
- Andreas Darmann
Universität Graz, AT
- Britta Dorn
Universität Tübingen, DE
- Daniel Eckert
Universität Graz, AT
- Edith Elkind
Nanyang TU – Singapore, SG
- Ulle Endriss
University of Amsterdam, NL
- Gábor Erdélyi
Universität Siegen, DE
- Piotr Faliszewski
AGH University of Science & Technology – Krakow, PL
- Felix Fischer
University of Cambridge, GB
- Umberto Grandi
University of Amsterdam, NL
- Paul Harrenstein
TU München, DE
- Edith Hemaspaandra
Rochester Institute of Technology, US
- Lane Hemaspaandra
University of Rochester, US
- Marc Kilgour
Wilfrid Laurier University, CA
- Christian Klamler
Universität Graz, AT
- Sascha Kurz
Universität Bayreuth, DE
- Jérôme Lang
Université Paris-Dauphine, FR
- Kate Larson
University of Waterloo, CA
- Omer Lev
The Hebrew University of Jerusalem, IL
- Vangelis Markakis
Athens University of Economics and Business, GR
- Nicolas Maudet
UPMC – Paris, FR
- Vincent Merlin
Caen University, FR
- Ilan Nehama
The Hebrew University of Jerusalem, IL
- Svetlana Obraztsova
St. Petersburg Electrotechnical University, RU
- Ulrich Pferschy
Universität Graz, AT
- Francesca Rossi
University of Padova, IT
- Jörg Rothe
Universität Düsseldorf, DE
- Lena Schend
Universität Düsseldorf, DE
- Arkadii Slinko
University of Auckland, NZ
- Isabelle Stanton
University of California – Berkeley, US
- Toby Walsh
NICTA – Kensington, AU
- Virginia Vassilevska Williams
Stanford University, US
- Gerhard Woeginger
TU Eindhoven, NL
- Michael Zuckerman
The Hebrew University of Jerusalem, IL