COST
EUROPEAN COOPERATION
IN SCIENCE & TECHNOLOGY

SHORT TERM SCIENTIFIC MISSION (STSM) – SCIENTIFIC REPORT

The STSM applicant submits this report for approval to the STSM coordinator

Action number: CA16228
STSM title: Computational Complexity of Automated Mechanism Design
STSM start and end date: 2018-05-01 to 2018-06-10
Grantee name: Dominik Peters

PURPOSE OF THE STSM/

Certain problems in the theories of mechanism design, social choice, and fair allocation have proven analytically intractable. Examples include the design of revenue-optimal auctions, and of finding strategyproof voting and allocation rules satisfying fairness properties. One way around this problem is to use computers to automatically design suitable mechanisms, an approach proposed as Automated Mechanism Design by Conitzer and Sandholm in 2002. One problem with this approach is that it doesn't scale well: typically it can only design mechanism in settings that involve a small number of agents and a very small number of objects/states/alternatives. However, many real-world applications concern problem sizes which are currently out of reach. In this STSM, we want to study whether there are methods and algorithms for these problems that scale better, or whether there are fundamental obstacles to this (in form of NP-hardness results). We hope to find general results that apply to many different types of settings, and different axioms and design objectives.

DESCRIPTION OF WORK CARRIED OUT DURING THE STSMS

We performed some initial work in formalizing the problem described in the research plan for the STSM, but had trouble gaining traction on their computational complexity. Since it seemed difficult to make further progress, we switched gears and worked on a new and different project, the problem of *portioning*.

In portioning, our task is to split a commonly owned budget (of money or some other resource such as time) among different tasks. A popular example is *participatory budgeting*, where a local government asks citizens how (part of) the government budget should be divided among different tasks, where citizens report their preferences among projects. My host, Jerome Lang, with coauthors, has recently introduced a family of portioning rules that can be used to arrive at such splits. The most promising member of this family is a rule translating voters' ordinal preferences into utilities using, say, Borda scores, and then selects the split of the budget which maximises Nash welfare. With Jerome, and with Justin Kruger, we started a more detailed axiomatic analysis of this rule, studying fairness and monotonicity properties. We also worked on developing connections of these rules to the problem of electing a parliament, and to the problem of assigning weights to criteria in multi-criteria decision making.

In another project, we have worked on the communication complexity of the popular voting STV (aka IRV or alternative vote). The question is whether we can figure out the election winner without having to ask all voters about all of their preferences. A classical result shows that an STV winner can be found in only O((log m)^2) queries. We considered ways in which this can be sped up even further in practice, but found that a natural technique, while usually requiring significantly fewer queries, requires the central authority to solve an NP-complete problem.

DESCRIPTION OF THE MAIN RESULTS OBTAINED

For the rule based on maximising Nash welfare described above, we found an example showing that it violates a natural monotonicity property (that, if we shift the position of one of the alternatives up in a vote, then the share of the budget spent on that alternative should not decrease). Based on extensive computer searches, we conjecture that monotonicity can be retained if we do not use Borda scores, but k-Approval scores, but a proof seems difficult. We also conjecture the Nash rule to be monotonic in case the input preferences over projects are dichotomous (0-1).

Another interesting question that we investigated is the formalisation of fairness notions. By a "fairness notion", we mean a concept which says that the rule spends enough on projects so that each voter, and each coalition of voters, is reasonably happy. A minimal such notion is "positive share", which requires that the rule does not allocate all of the budget to some voter's least-preferred project. It is easy to see that the Nash rule satisfies this notion, and that many other rules also do. However, the Nash welfare maximum is known to satisfy much stronger notions in many settings. In particular, when voters have cardinal utilities, the Nash outcome lies in the core of a natural associated cooperative game. We defined a notion of the core (the SD-core) which is appropriate to the present ordinal input format.

FUTURE COLLABORATIONS (if applicable)

We aim to continue the collaborations begun during my visit in Paris. Several results are currently only conjectural, and we are working on formal proofs. Some of us are planning to meet again later this summer to continue this project.