SHORT TERM SCIENTIFIC MISSION (STSM) – SCIENTIFIC REPORT

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

Action number: CA16228
STSM title: *On Socially Constrained Envy-Freeness up to the Least Valued Good*
STSM start and end date: 06/04/2018 to 16/04/2018
Grantee name: Georgios Amanatidis

PURPOSE OF THE STSM

The purpose of the STSM was to study the problem of fairly allocating a set of indivisible items to a set of agents subject to an underlying social network. The social network models social acquaintances; in particular, each agent should consider the allocation "fair" with respect to her bundle and the bundles of her neighbors. Depending on the model studied, an agent may ignore completely the items allocated to her non-neighbors or require some mild fairness property. The main fairness notion we focus on is "envy-freeness up to the least valued good" (EFX), a recently introduced relaxation of envy-freeness. EFX requires that for any i, j, agent i does not envy agent j after the removal of i's least valued item from j's bundle.

There are a few very recent works studying either fairness with respect to an underlying network or EFX fairness, albeit not both. Many interesting questions remain open, so we focus on the following:

(a) For more than two agents it is not known whether EFX allocations always exist. It is a most intriguing open problem to resolve, if not in general, at least for particular families of underlying graphs.

(b) In addition to the above question, it is an open problem whether one can efficiently compute an r-approximate EFX allocation for r > 0.5 and more than two agents.

(c) From a mechanism design point of view, it is known that there exists no deterministic truthful mechanism that achieves *exact* EFX fairness for two agents and more than four items. Can we design truthful mechanisms with non-trivial approximation ratios for more items and/or agents?

DESCRIPTION OF WORK CARRIED OUT DURING THE STSMS

We focused on the case of three agents where the most fundamental questions seem to still be highly non trivial.

We applied several algorithmic ideas that have been used in other related fairness notions for discrete fair division. None of our approaches yielded a better approximation factor than 0.5 (which is the current best) without some additional restriction on the underlying topology.

With the involvement of a senior undergraduate student, Georgia Dimaki, we started conducting experiments for three agents and randomly generated values drawn from several distributions (uniform, exponential and some discrete variants). Given that for a large number of items and values drawn uniformly at random it is

known that even envy-free allocations almost always exist, we decided to focus on a small number of items (up to 9).

Finally, motivated by the lack of positive results, we tried to construct an example with three agents where no EFX allocation exists. Known impossibility constructions for other fairness notions fail. In particular, the approach of Kurokawa, Procaccia and Wang that shows that MMS allocations do not always exist for three agents does not seem to work for EFX. One further attempt towards this direction was to apply the probabilistic method to prove that such "bad" instances do exist. Apparently, the calculations for the most natural choices of the probabilistic model here fail to work towards this goal.

DESCRIPTION OF THE MAIN RESULTS OBTAINED

For three agents, the problem of computing an exact EFX allocation becomes significantly easier once we assume for any specific pair of agents that they do not care about each other, i.e., when an edge is removed in the underlying social graph. For this case we obtain a polynomial time algorithm. For the unrestricted case, however, we did not obtain any theoretical results, positive or negative.

We also conducted experiments for three agents and a small number of items. In every single instance an EFX allocation exists. Even if there is a general impossibility result here, our experimental evidence suggest that it could be possible to have a probabilistic result of the form "for three agents and m items, the values of which drawn from some family of distributions, there is an EFX allocation with probability 1-o(1), where the o(1) term degrades very rapidly with respect to m".

Also, with respect to question (3) above, by combining results of two other recent works of ours, we got as a corollary that there is the no truthful allocation mechanism for *two* agents that achieves *any constant approximation* of EFX. We show that a similar statement holds for the recently introduced notion of *pairwise maximin share fairness* (PMMS), and even for *envy-freeness up to any good* (EF1) which is a rather weaker fairness notion than EFX.