COST Action IC1205 on Computational Social Choice: STSM Report
Applicant: Aris Filos-Ratsikas
Home institution: University of Oxford
Home country: United Kingdom
Host: Ioannis Caragiannis
Host institution: University of Patras
Host country: Greece
Dates: 03/06/2016 to 17/06/2016

During my visit to the University of Patras, I worked together with Ioannis Caragiannis and Panagiotis Kanellopoulos on several different problems regarding matching with metric preferences, fair resource allocation and stable marriage.

On the problem of matching with metric preferences in the resource augmentation framework, Ioannis and worked on proving several results, related to our recent work on the topic. Most of our attempts were focused on proving a stronger upper bound on the approximation ratio of random priority as well as lower bound for the mechanism under resource augmentation. These problems seem to be quite involved however and although we put in significant effort, we did not manage to improve upon our already proven bounds.

Together with Ioannis and Panagiotis, we investigated the problem of fair division of divisible resources, with the objective of egalitarian welfare maximization. Specifically, given a set of divisible resources, the question is "what is the best approximation to the maximum egalitarian welfare that we can achieve, when constrained to output envy-free allocations?". This is equivalent to the "Price of Envy-freeness" for the egalitarian welfare objective. We have a Theta(n^(1/3)) lower bound, which improves on current results in the literature, and we worked on proving some upper bound as well. We have not managed to do that yet, but I think we have some intuition on how to proceed.

Together with Ioannis and Panagiotis, we worked on a different problem, with more success. The problem is a variation of the well-known stable marriage problem, in which men and women have cardinal utilities over the members of the other side, which express their preferences. We look at fractional matchings, which are convex combinations over deterministic matchings. We say that a fractional matching is stable, if the weighted average (the expected) utility of all agents from their matched partners is such that no man-woman pair would prefer to deviate from the fractional matching to the single edge connecting them in the bipartite graph.

We look at the social welfare objective of the best fractional stable matching. (The question of the welfare-best deterministic stable matching has been studied since as early as 1987). First, we provide examples where the best fractional stable matching can be arbitrarily better (in terms of welfare) than the best deterministic stable matching. Secondly, we provide a polynomial-time algorithm for finding a fractional stable matching which is better than the best deterministic stable matching, if one exists. This algorithm can be used to characterize the instances on which a fractional stable matching with higher welfare than all deterministic ones exists.

We then tackle the computational complexity of finding the best fractional stable matching. For now, we have proved that we can find the best fractional stable matching in polynomial-time when the number of possible matchings (the support) is constant. At this point, we believe that the general problem will be computationally hard, but if the opposite turns out to be true, we think that the result will be quite interesting algorithmically.

We plan to continue working on this problem in the following months, hoping to answer the question above, as well as explore what happens when we consider other notions of stability that exist in literature.