An Experimental Comparison of Single-Sided Preference Matching Algorithms

Dimitrios Michail
Dept. of Informatics and Telematics
Harokopio University of Athens, Greece
michail@hua.gr

Consider the scenario where a set of applicants $\mathcal{A}$ has an interest in obtaining a set of posts $\mathcal{P}$ and suppose that associated with each member of $\mathcal{A}$ is a preference list (possibly including ties) comprising a subset of elements of $\mathcal{P}$. A matching of $\mathcal{A}$ to $\mathcal{P}$ is an allocation of each applicant to at most one post such that each post is filled by at most one applicant. Stated differently, it is a matching in the bipartite graph $G = (\mathcal{A} \cup \mathcal{P}, E)$ where E consists of all pairs (a, p) where p belongs in the ordered preference list of a.

The main focus of this work is to experimentally study matchings computed by various one-sided preference matching algorithms with respect to their unpopularity. On the other hand, since it would be unfair to judge algorithms based solely on the unpopularity, we include additional quality measurements such as cardinality, total rank, maximum rank and running time. We compare several different algorithms for the computation of rank-maximal matchings [3, 4], the algorithm of [1] for the computation of popular matchings, and the algorithm of [2]. While popular matchings seem to be unrelated to rank-maximal matchings, the algorithmic techniques required in order to efficiently compute both types are very much related. Thus, all algorithms are implemented using similar heuristics and graph representations.

The experimental comparison of the aforementioned algorithms is performed on instances created by three random structured instance generators. All generated problem instances try to mimic different real life situations, while maintaining as few parameters as possible. Moreover, in addition to synthetic datasets, we experiment with two real-world datasets.

References

[1] David J. Abraham, Robert W. Irving, Telikepalli Kavitha, and Kurt Mehlhorn. Popular matchings. *SIAM Journal on Computing*, 37(4):1030–1045, 2007.

[2] Chien-Chung Huang, Telikepalli Kavitha, Dimitrios Michail, and Meghana Nasre. Bounded unpopularity matchings. *Algorithmica*, pages 1–20, 2010.

[3] Robert W. Irving, Telikepalli Kavitha, Kurt Mehlhorn, Dimitrios Michail, and Katarzyna E. Paluch. Rank-maximal matchings. *ACM Transactions on Algorithms*, 2(4):602–610, 2006.

[4] Dimitrios Michail. Reducing rank-maximal to maximum weight matching. *Theoretical Computer Science*, 389(1-2):125 – 132, 2007.