

Procedural Fairness in Stable Marriage Problems

(Extended Abstract)

Mirco Gelain
University of Padova, Italy
mgelain@math.unipd.it

Maria Silvia Pini
University of Padova, Italy
mpini@math.unipd.it

Francesca Rossi
University of Padova, Italy
frossi@math.unipd.it

Kristen Brent Venable
Universita' di Padova
Padova, Italy
kvenable@math.unipd.it

Toby Walsh
NICTA and UNSW
Sydney, Australia
toby.walsh@nicta.com.au

ABSTRACT

The stable marriage problem is a well-known problem of matching men to women so that no man and woman, who are not married to each other, both prefer each other. It has a wide variety of practical applications, ranging from matching resident doctors to hospitals, to matching students to schools, or more generally to any two-sided market. Given a stable marriage problem, it is possible to find a male-optimal (resp., female-optimal) stable marriage in polynomial time. However, it is sometimes desirable to find stable marriages without favoring one group at the expenses of the other one. To achieve this goal, we consider a local search approach to find stable marriages with the aim of exploiting the non-determinism of local search to give a fair procedure. We test our algorithm on classes of stable marriage problems, showing both its efficiency and its sampling capability over the set of all stable marriages, and we compare it to a Markov chain approach.

Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem solving, Control Methods, and Search

General Terms

Algorithms

Keywords

Stable marriage problem, local search, fairness

1. STABLE MARRIAGE PROBLEMS

The stable marriage problem (SMP) [5] is a well-known problem of matching n men to n women to achieve a certain type of ‘stability’. Given n men and n women, where each person expresses a strict preference ordering over the members of the opposite sex, the problem is to match the men to the women such that no two people of the opposite sex, who

Cite as: Procedural Fairness in Stable Marriage Problems (Extended Abstract), Mirco Gelain, Maria Silvia Pini, Francesca Rossi, Kristen Brent Venable, and Toby Walsh, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. XXX-XXX.

Copyright © 2011, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

are not married to each other, both prefer each other to their current partners. If there are no such pairs, called blocking pairs, every marriage is stable. In [2] Gale and Shapley provided an $O(n^2)$ time algorithm for finding two specific stable marriages, called male-optimal and female-optimal, that favour one gender over the other one. It is known that the set of the stable marriages forms a distributive lattice where the male-optimal is the top and the female-optimal is the bottom [5]. Male-optimality (and also female-optimality) may be considered too unfair between the two genders: although stability is assured, only one of the genders is as happy as possible. For this reason, other kinds of fairer stable marriages have been considered, such as the minimum regret stable marriage [4]. Besides the fairness of the generated stable marriage, it is also interesting to consider the fairness of how a stable marriage is generated. We now describe how the fairness of stable marriage procedures can be achieved by considering a local search approach.

2. LOCAL SEARCH FOR SMPs

In [3] we presented a local search algorithm to find stable marriages. Given an SMP instance P , we start from a randomly generated marriage M . Then, at each search step, we compute the set BP of blocking pairs in M and the neighborhood, which is the set of all marriages obtained by removing one of the blocking pairs in BP from M . To select the neighbor M' of M to move to, we use an evaluation function that counts the number of blocking pairs in all neighboring marriages, and we move to the one with the smallest number of blocking pairs. To avoid stagnation in a local minimum of the evaluation function, at each search step we perform a random walk with a certain probability which removes a randomly chosen blocking pair in BP from the current marriage M . The algorithm terminates if a stable marriage is found or when a maximal number of search steps, or a timeout, is reached. The number of such blocking pairs may be very large. Also, some of them may be useless, since their removal would surely lead to new marriages that will not be chosen by the evaluation function. This is the case for the so-called *dominated* blocking pairs. In our procedure we consider only undominated blocking pairs. Let m be a man and (m, w) and (m, w') two blocking pairs. Then (m, w) dominates (from the men’s point of view) (m, w') if m prefers w to w' . Since dominance between blocking pairs is defined from one gender’s point of view, to ensure gender

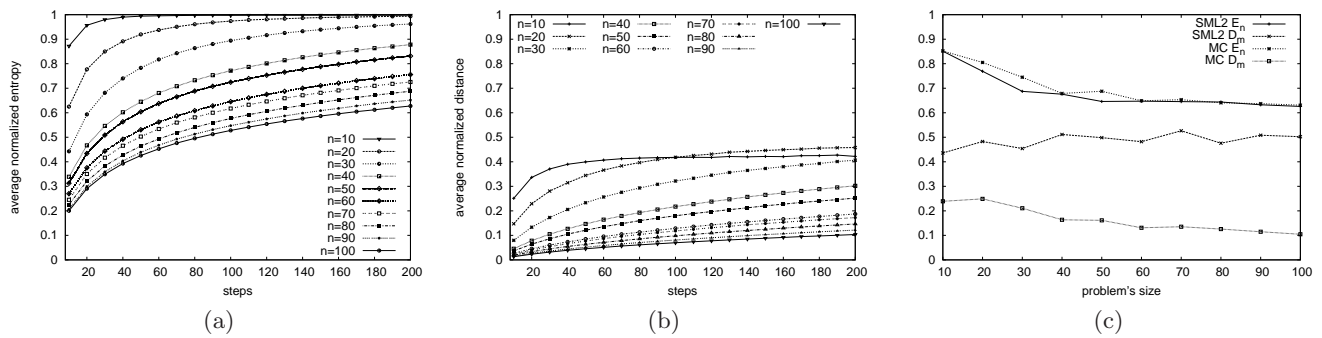



Figure 1: Average runtime entropy of MC (a), average runtime distance from the male-optimal of MC (b), Local Search vs. MC in terms of entropy and distance from the male-optimal (c).

neutrality, at each search step we swap the role of the two genders.

3. MEASURING PROCEDURAL FAIRNESS

We ran experiments on randomly generated SMPs of different size, up to 500 men and 500 women, with random walk probability 0.2. Our algorithm always found a stable marriage. Also, its runtime behavior suggests that the number of steps grows as little as $O(n \log n)$ [3]. Here we show that the algorithm is able to find a stable marriage for all the problems in the test set within 10000 steps, and for each set of problems of the same size, the probability to find a stable marriage grows very fast within a small interval of steps (see the figure below). This means that it is possible to predict the number of steps needed by our algorithm to find a stable marriage with a reasonable precision.



In [3] we evaluated the sampling capability over the lattice of stable marriages of a given SMP. To do this, we randomly generated 100 SMP instances for each size between 10 and 100, with step 10. We then measured the distance D_m of the found stable marriages (on average) from the male-optimal marriage. If D_m is equal to 0 (resp., 1), it means that all the stable marriages returned coincide with the male-optimal (resp., female-optimal) marriage. The average distance from the male-optimal is around 0.5 as shown in Fig. 1(c), where our algorithm is called SML2. This is encouraging but not very informative, since also an algorithm which always returns the same stable marriage, with distance 0.5 from the male-optimal, would have $D_m = 0.5$. To have more informative results on the sampling capabilities, we considered the entropy of our algorithm, say E_n , that is, the uncertainty to find a specific stable marriage. More precisely, E_n is the average normalized frequency of the stable marriages returned by our algorithm over the whole lattice (see [3] for the formal

definition). Experimental results showed that this entropy is in general very high (about 70% of the maximum and even higher as shown in Fig. 1(c)) and thus we are not far from the ideal behavior.

To better evaluate the sampling capability of our approach, here we compare it to a *Markov chain* approach (MC) [1], defined by using rotations exposed in each stable marriage. This approach converges in exponential time to the uniform distribution over the stable marriages. We consider the entropy and distance from the male-optimal of MC computed on executions where we vary the number of steps from 10 to 200. While the entropy of MC increases quite rapidly, the distance from the top of the lattice (i.e., from the male-optimal) increases more slowly (see Fig. 1(a) and Fig. 1(b)). For each problem instance in the test set, we start MC from the male-optimal marriage and take the stable marriage returned by MC after exactly the same number of steps needed by our algorithm to find a stable marriage for that instance. Then we measure and compare the entropy and the distance from the male-optimal for MC to those of our algorithm (SML2). While the entropy of MC is roughly the same as that of our algorithm, the distance from the male-optimal achieved by our approach (about 0.5) is on average higher than that achieved by MC (about 0.2) (see Fig. 1(c)).

Summarizing, our approach is efficient and it has sampling capabilities comparable with a Markov chain approach considering the same number of steps, and may even perform slightly better considering the distance measured from the top or the bottom of the lattice.

4. REFERENCES

- [1] N. Bhatnagar, S. Greenberg, and D. Randall. Sampling stable marriages: why spouse-swapping won't work. In *SODA*, pages 1223–1232, 2008.
- [2] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [3] M. Gelain, M. S. Pini, F. Rossi, K. B. Venable, and T. Walsh. Local search algorithms on the stable marriage problem: Experimental studies. In *ECAI'10*, pages 1085–1086, 2010.
- [4] D. Gusfield. Three fast algorithms for four problems in stable marriage. *SIAM J. Comput.*, 16(1), 1987.
- [5] D. Gusfield and R. W. Irving. *The Stable Marriage Problem: Structure and Algorithms*. MIT Press, Boston, MA, 1989.