

Morphing between Stable Matching Problems

Ciaran McCreesh*, Patrick Prosser, and James Trimble

University of Glasgow, Glasgow, Scotland

Abstract. In the stable roommates (SR) problem we have n agents, where each agent ranks all other agents in strict order of preference. The problem is then to match agents into pairs such that no two agents prefer each other to their matched partners, and this is a stable matching. The stable marriage (SM) problem is a special case of SR, where we have two equal sized sets of agents, men and women, where men rank only women and women rank only men. Every instance of SM admits at least one stable matching, whereas for SR as the number of agents increases the number of instances with stable matchings decreases. So, what will happen if in SM we allow men to rank men and women to rank women, i.e. we relax gender separation? Will stability abruptly disappear? And what happens in a stable roommates scenario if agents do not rank all other agents? Again, is stability uncommon? And finally, what happens if there are an odd number of agents? We present empirical evidence to answer these questions.

1 Introduction

In the **Stable Roommates** (SR) problem [7, 6, 9] we have n agents, where each agent ranks all $n - 1$ other agents in strict order of preference. The problem is then to match pairs of agents in a one-one correspondence (bijection) such that the matching is *stable*. A matching is stable if there does not exist a *blocking pair* of agents ($agent_i$ and $agent_j$) in the matching such that $agent_i$ and $agent_j$ find each other *acceptable* (i.e. they rank each other) and (a) $agent_i$ is unmatched or prefers $agent_j$ to his matched partner and (b) $agent_j$ is unmatched or prefers $agent_i$ to his matched partner (see [9], 1.4.2).

The **Stable Marriage** problem (SM) [3, 16, 4, 6, 17, 9] is a specialized instance of SR where agents have gender, such that we have two sets of agents, namely men and women, both the same cardinality. A stable matching is a one-one correspondence between the men and women such that the matching admits no blocking pair (as defined above, also see [9] 1.3.4). Figure 1 shows (on the left) an instance of SR with 6 agents (sr6), and (in the middle) an instance of SM with three men and three women (sm3).

SM instance sm3 can be represented as an instance of the Stable Roommates Problem with Incomplete lists (SRI), i.e. we have incomplete lists where some agents find each other unacceptable (see [9], 1.4.2). This is shown on the right

* Ciaran McCreesh was supported by the Engineering and Physical Sciences Research Council [grant number EP/K503058/1]

of Figure 1 (sri6), where agents 1 to 3 represent the men in sm3, and agents 4 to 6 represent the women in sm3. Therefore the matching $\{(1, 1), (2, 2), (3, 3)\}$ in sm3 corresponds to the matching $\{(1, 4), (2, 5), (3, 6)\}$ in sri6.

1: 6 3 5 2 4			1: 4 5 6
2: 3 5 1 6 4	1: 1 2 3	1: 1 2 3	2: 5 6 4
3: 2 6 1 5 4	2: 2 3 1	2: 1 3 2	3: 4 6 5
4: 5 1 2 3 6	3: 1 3 2	3: 3 1 2	4: 1 2 3
5: 6 1 2 3 4			5: 1 3 2
6: 4 2 5 1 3			6: 3 1 2

Fig. 1. A Stable Roommates (SR) instance on the left sr6, with 6 agents. In the middle an instance of Stable Marriage (SM) with three men and three women, sm3. On the right, sm3 is represented as an instance of Stable Roommates with Incomplete lists (SRI), instance sri6. Instance sr6 has two stable matchings, namely $\{(1, 5), (2, 3), (4, 6)\}$ and $\{(1, 4), (2, 3), (5, 6)\}$. Instance sm3 has a single stable matching $\{(1, 1), (2, 2), (3, 3)\}$. Instance sri6 has only one stable matching $\{(1, 4), (2, 5), (3, 6)\}$ and this corresponds to the stable matching for sm3.

The underlying structure of SR is a simple graph [10], where an edge corresponds to a pair of agents who find each other acceptable. The SR graph is therefore a clique K_n and preference lists of agents are permutation of their adjacency lists. The SM graph is a complete biclique, with two sets of vertices (men and women), each of size $n/2$ with $n^2/4$ edges. In SRI the underlying graph is again a simple graph, but not complete. Every instance of SM admits at least one stable matching [3], but this is not the case for SR [7], where some instances admit no stable matching (the same is true of SRI). As the number of agents increase the proportion of SR instances with stable matchings falls [7, 14, 12, 11, 15].

Therefore, we have a spectrum of stable matching problems. At one extreme we have SM, highly structured (a biclique), always with stable matchings and at the other extreme we have SR, unstructured, with a falling number of stable matchings as the number of agents increases [3, 7, 6, 9]. Therefore, what happens as we replace some of the structure in SM with randomness from SR, i.e. what happens when we blend SM and SR? Will there be an abrupt change in behaviour, where the average number of instances with stable matchings falls, or will it be smooth and gradual with stability declining slowly?

There is also a spectrum of stable roommates problems. At one extreme we have SRI instances with empty preference lists, where every agent finds every other agent unacceptable, and corresponds to the edgeless graph. This instance has one stable matching, where every agent is happy to be unmatched¹. At the other extreme we have SR, where every agent finds every other agent acceptable,

¹ We assume that every agent ranks himself in last position and can potentially be self-matched.

not always admitting a stable matching, and this corresponds to the complete graph. Therefore we can move across this spectrum, gradually increasing the number of acceptable pairs (edges in the graph). As we do so, will the proportion of instances with stable matchings fall gradually, or will it fall abruptly?

And finally, it is tacitly assumed that the number of agents is even. Why should that be so? Imagine we had a conference where delegates share rooms in the students' halls of residence, two to a room. Would it be possible to limit attendance to only an even number of delegates? And if delegates rank one another in order of preference, is it more likely to be an unstable scenario when the number of delegates is odd? We investigate this also.

In the following section we describe how we can mix SM and SRI in a controlled manner (problem generation). We then present empirical results for mixing SM and SRI followed by experiments on gradually moving from SRI with no acceptable pairs to SR, where the number of agents is even and when the number of agents is odd.

2 Problem Generation

Given a graph $G = (V, E)$, where V is the set of vertices and E the set of edges, we can create a stable matching problem from G as follows. The set of vertices correspond to the set of agents and an edge (i, j) is in E if and only if $agent_i$ and $agent_j$ are an acceptable pair (i.e. they rank each other). Assume vertex neighbourhood is represented as an adjacency list. Given an edge $(u, v) \in E$, add u to the list $adjacent[v]$ and add v to the list $adjacent[u]$. Once this has been done for all edges, perform a Knuth shuffle [2] on each adjacency list and treat these as preference lists. This is essentially the technique used in Section 2 of [10].

We now describe two techniques for creating a blended graph, with m edges, from two input graphs. Consider simple graphs G_1 and G_2 , both of the same order n (number of vertices), where G_1 is the complete bipartite graph $K_{n/2, n/2}$ and G_2 is the clique K_n . Assume we have a mixing proportion p , where $0 \leq p \leq 1$. We can mix these two graphs to produce a new graph G_3 with $(1 - p)m$ edges taken from G_1 and $p.m$ edges taken from G_2 . This can be done in two ways. The first is similar to the rewiring technique of Watts and Strogatz [8]. This is presented in Algorithm 1 and will be referred to as *model A*. The algorithm returns a set of edges E , i.e. acceptable pairs, where that set is of size $m = n^2/4$, the same size as the biclique.

Algorithm 2 corresponds to the type-B morph described in [5], where again $m = n^2/4$ edges are to be produced (line 3). The set E_1 contains m randomly selected edges from K_n , the set E_2 contains all edges in $K_{n/2, n/2}$ and set E is the intersection of these two sets (lines 4, 5 and 6). The remaining number of edges to be added to E is $\delta = m - |E|$, where $p.\delta$ edges are randomly selected from $E_1 \setminus E$ and $(1 - p).\delta$ edges are randomly selected from $E_2 \setminus E$. Therefore when $p = 0$ both models deliver an instance of SM, and when $p = 1$ both deliver an instance of SRI with $p.n^2/4$ edges drawn at random from K_n .

Algorithm 1: modelA: select $m.(1-p)$ from biclique and $m.p$ from clique

```

1 Set⟨Edge⟩ modelA(int  $n$ , real  $p$ )
2 begin
3   int  $m \leftarrow n^2/4$ 
4    $E_1 \leftarrow \{(i, j) \mid 1 \leq i < j \leq n\}$ 
5    $E_2 \leftarrow \{(i, j) \mid 1 \leq i \leq n/2, n/2 < j \leq n\}$ 
6    $E \leftarrow \text{select}((1-p).m, E_2)$ 
7    $E \leftarrow E \cup \text{select}(p.m, E_1 \setminus E)$ 
8   return  $E$ 

```

Algorithm 2: modelB: a type B morph

```

1 Set⟨Edge⟩ modelB(int  $n$ , real  $p$ )
2 begin
3   int  $m \leftarrow n^2/4$ 
4    $E_1 \leftarrow \text{select}(m, \{(i, j) \mid 1 \leq i < j \leq n\})$ 
5    $E_2 \leftarrow \{(i, j) \mid 1 \leq i \leq n/2, n/2 < j \leq n\}$ 
6    $E \leftarrow E_1 \cap E_2$ 
7    $\delta \leftarrow m - |E|$ 
8    $E \leftarrow E \cup \text{select}(p.\delta, E_1 \setminus E) \cup \text{select}((1-p).\delta, E_2 \setminus E)$ 
9   return  $E$ 

```

3 The Empirical Study

The majority of the study used a constraint programming formulation of the stability constraint proposed in [15]. In all our models each agent ranks himself in last position. Consequently an agent can self-match, if and only if this results in stability. The experiments were run on an Intel Xeon CPU E5-2660 processor at 2.2 GHz with 20Mb of cache and 128 Gb of RAM. In many of the studies the control parameter is p (as an edge probability or mixing proportion), and is varied in steps of 0.001 with a sample size of 1,000.

3.1 Morphing from SM to SRI

The first experiment investigates what happens as we morph from SM to SRI using models A and B. This is shown in Figure 2. On the x axis we have p , the mixing parameter, and when $p = 0$ all instances are bipartite and when $p > 0$ instances are non-bipartite. On the y axis we have the average percentage of instances that admit a stable matching, i.e. the percentage that were satisfiable. On the left we have two contours, both for $n = 100$, one for model A the other for model B. The models produce similar behaviour, with model B preserving the SM properties (bipartite) longer than model A.

On the right are contours for n varying from 50 to 400 using model A. At $p = 0$ all instances are satisfiable, as expected, but what is surprising is

how rapidly behaviour changes with a small degree of mixing and that this becomes more abrupt as the number of agents increases. Although not shown, as p increases, the number of stable matchings per instance falls rapidly. In summary, a small degree of within-gender acceptability results in a rapid loss of stability.²




Fig. 2. Morphing from SM to SRI. The x-axis is p , the mixing proportion, and the y-axis is the percentage of instances admitting a stable matching. On the left, model A and model B with $n = 100$. On the right, contours for a variety of n , 50 to 400.

3.2 Morphing between SR and SRI

In [10] Mertens empirically investigated SR (measuring P_n , the probability of a stable matching existing in K_n), SRI where agents exist on a grid (and rank only their Moore neighbourhood) and SRI on graphs with a given average degree (where random graphs were generated with increasing order but with an average degree of only 35, 45 or 60). We now investigate what happens in the stable roommates problem as we vary the amount of acceptability between agents, i.e. we vary average degree. Again, viewing the problem as a simple graph, we vary the edge probability p and with that the degree of the graph, and this corresponds to the average length of preference lists. When edge probability is one every agent finds every other agent acceptable and problem instances are SR, when p is zero every agent finds every other agent unacceptable and agents are happy to be alone (i.e. self-matched), and when $0 < p < 1$ preference lists are incomplete and we have SRI instances. Figure 3 shows, on the left, contours for $n = 100$ and $n = 101$, with probability of acceptability on the x axis and percentage of stable instances on the y axis. On the right is the average size of a stable matching when one exists³, and it should be noted that for a given instance all its matchings are the same size [16].

² We leave any social interpretation of these observations to others.

³ The x axis is cut short due to small sample size for odd n and large p .




Fig. 3. Increasing the proportion p of agents found to be acceptable in stable roommates. On the x axis p , probability that two agents find each other acceptable. On the left, the y axis is percentage of instances that admit a stable matching. On the right, y axis is average size of stable matching when one exists. Note that when an instance admits a stable matching, all its stable matchings are the same size [6]. Contours are shown for even ($n = 100$) and odd ($n = 101$) number of agents.

The shape of these curves are surprising. For even n all instances are stable when $p = 0$, and this then falls away. Tabulated results for SRI predict that this will happen [7, 13, 14, 12, 11, 15]. But what was not predictable was the shape of the $n = 100$ contour: falling sharply, climbing abruptly and then tapering off to its final value when SRI becomes SR. Note that the “knee” in the contour on the left of Figure 3 comes some time after the point where random graphs $G(n, p)$ become a single component. Therefore the “knee” is not due to the emergence of a giant component.

When n is odd there are indeed stable matchings in our model (an odd number of agents must self-match), but these stable matchings are typically either small or rare. However, the shape of the odd contour on the left is perhaps not so surprising. For $n = 101$, suppose all preference lists are complete, and without loss of generality suppose that agent 101 is unmatched. Then for agent 101 not to block the matching, every other agent must have a better partner than him, which is unlikely to be satisfied in general. Obviously the longer the preference lists in general, the more chances agent 101 has to block the matching, hence the shape of the $n = 101$ contour (on the left)⁴.

Figure 4 separates even (contours on the left) from odd (contours on the right) for various number of agents. This time the x axis is $n.p$ and that is the average degree of the underlying graph, and is then the average size of preference lists. For compactness, the x-axis is shown on a log scale.

The experiments were then repeated for larger values of even n and odd $n+1$. Graphs were generated using the algorithm of Batagelj and Brandes [1], which allows much faster generation of large, sparse graphs than simple quadratic-

⁴ We thank David Manlove for this explanation.




Fig. 4. Percentage of SRI instances with stable matchings as we vary probability that agents find each other acceptable. On the left n is even and on the right n is odd. The x-axis is a logscale of $n.p$, i.e. average degree.

time methods. Each instance was solved using both Irving’s algorithm [7] and a simple SAT encoding using MiniSat, to verify correctness of our implementation. Sample size varied from 2,000 to 1,000,000. We ran our program for values of np up to 896; these results are shown in Figure 5. In addition, for large even values of n the figure shows results for $p = 1$ from [12]; these appear as the rightmost points. It appears that when n is even and n increases, its odd partner *hugs* onto it for longer and longer, leading us to expect that in the limit both curves will be indistinguishable. That is, when there are a large number of agents it will not matter if that number is odd or even.

4 Conclusion

Stable marriage problems always admit a stable matching but this is not true for stable roommates. In our experiments, a certain amount of disorder was added to SM such that we permit within-gender matching. This brought about an abrupt change in the behaviour of the problem, with an abrupt fall in the proportion of instances with stable matchings.

In the roommates problem, with a small number of agents, it appears to matter if the number of agents is odd or even. If the number of agents is odd, stable matchings tend to be small or scarce. When the number of agents is even, the proportion of instances with stable matchings falls as acceptability increases then abruptly climbs, and this climb then gradually tapers off as the instances tend to SR with complete preference lists. However, when n is large, we conjecture that behaviour will tend to be that observed when n is odd.

Acknowledgements

We would like to thank David Manlove, Augustine Kwanashie, Rob Irving, Ian Gent and Craig Reilly.

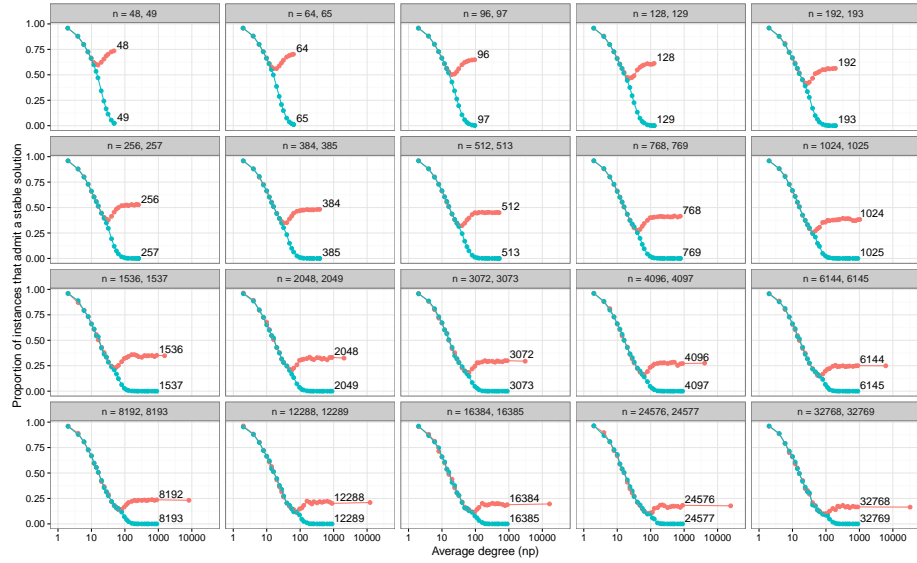


Fig. 5. Increasing the proportion p of agents found to be acceptable. On the x axis the log of average degree ($p.n$), the y axis is percentage of instances that admit a stable matching. The red contour is for an even number of agents n and the blue contour for odd number of agents $n + 1$.

References

1. Vladimir Batagelj and Ulrik Brandes. Efficient generation of large random networks. *Phys. Rev. E*, 71:036113, Mar 2005.
2. Richard Durstenfeld. Algorithm 235: Random permutation. *CACM*, 7:420, 1964.
3. D. Gale and L.S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
4. D. Gale and M. Sotomayor. Some remarks on the stable matching problem. *Discrete Applied Mathematics*, 11:223–232, 1985.
5. Ian P. Gent, Holger H. Hoos, Patrick Prosser, and Toby Walsh. Morphing: Combining structure and randomness. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence*, pages 654–660, 1999.
6. D. Gusfield and R. W. Irving. *The Stable Marriage Problem: Structure and Algorithms*. The MIT Press, 1989.
7. Robert W. Irving. An Efficient Algorithm for the Stable Roommates Problem. *J. Algorithms*, 6(4):577–595, 1985.
8. Duncan J. Watts and Steven H. Strogatz. Collective dynamics of small-world networks. *Nature*, 393:440–442, 1998.
9. David Manlove. *Algorithmics of Matching under Preferences*, volume 2 of *Theoretical Computer Science*. World Scientific, 2013.
10. Stepan Mertens. Random stable matchings. *Journal of Statistical Mechanics: Theory and Experiments*, (10):P10008, 2005.
11. Stephan Mertens. Small random instances of the stable roommates problem. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(6):P06034, 2015.

12. Stephan Mertens. Stable roommates problem with random preferences. *Journal of Statistical Mechanics: Theory and Experiment*, 2015(1):P01020, 2015.
13. C. Ng and D.S. Hirschberg. Lower bounds for the stable marriage problem and its variants. *SIAM Journal on Computing*, 19:71–77, 1990.
14. Boris Pittel and Robert W. Irving. An upper bound for the solvability of a random stable roommates instance. *Random Struct. Algorithms*, 5(3):465–487, 1994.
15. Patrick Prosser. Stable roommates and constraint programming. In *Integration of AI and OR Techniques in Constraint Programming - 11th International Conference, CPAIOR*, pages 15–28, 2014.
16. A.E. Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of Political Economy*, 92(6):991–1016, 1984.
17. A.E. Roth and M.A.O. Sotomayor. *Two-sided matching: a study in game-theoretic modeling and analysis*, volume 18 of *Econometric Society Monographs*. Cambridge University Press, 1990.