

Random Popular Matchings

[Extended Abstract]

Mohammad Mahdian
Microsoft Research
Redmond, WA, USA
mahdian@microsoft.com

ABSTRACT

We consider matching markets where a centralized authority must find a matching between the *agents* on one side of the market, and the *items* on the other side. Such settings occur, for example, in mail-based DVD rental services such as NetFlix or in some job markets. The objective is to find a *popular matching*, or a matching that is preferred by a majority of the agents to any other matching. This concept was first defined and studied by Abraham et al. The main drawback of this concept is that popular matchings sometimes do not exist. We partially address this issue in this paper, by proving that in a probabilistic setting where preference lists are drawn at random and the number of items is more than the number of agents by a small multiplicative factor, popular matchings almost surely exist. More precisely, we prove that there is a threshold $\alpha \approx 1.42$ such that if the number of items divided by the number of agents exceeds this threshold, then a solution almost always exists. Our proof uses a characterization result by Abraham et al., and a number of tools from the theory of random graphs and phase transitions.

Categories and Subject Descriptors

G.2 [Discrete Mathematics]: Applications

General Terms

Theory, Algorithms, Economics

Keywords

matching, stable matching, popular matching

1. INTRODUCTION

In many centralized real-world markets, the task is to form a matching between two sides of the market. Examples of such markets include matchmaking markets (where boys are matched to girls), job markets such as medical residency

matching program [13] (where applicants are matched to positions), mail-based DVD rental systems such as NetFlix [1] (where available DVDs are matched to the subscribers), and kidney exchange markets [11] (where patients are matched to kidneys available for transplant). In such markets, agents on one or both sides of the market have preferences over their possible matches. These preferences are often *ordinal*, i.e., they only express the relative ranking of the options. The objective is to find a matching that is “optimal” with respect to these preferences. This is, however, not a well-defined goal, since preferences are only ordinal, and there is no way to compare the utility of one agent with the utility of another. This has motivated the search for reasonable *solution concepts* in this setting.

In cases where both sides of the market have preferences over the other side, a satisfactory answer to this question, called *stable matching*, was proposed by Gale and Shapley in 1962 [6]. Since then, this concept was used in many applications such as matching medical students to hospitals or matching students to high schools in New York City (See [12] for a list of such applications). One of the most important properties of stable matchings, proved by Gale and Shapley [6], is that they always exist. This is a vital property for a solution concept, as in most practical applications having no solution is simply not an option.

For markets where only one side has preference over the other, the situation is less satisfactory. Such markets correspond to situations where one side of the market consists of agents with preferences, and the other side consists of items that can be allocated to the agents. In cases where there is an initial assignment of items to the agents (e.g., in some models of the housing market), there is an algorithm commonly known as the *top trading cycle* algorithm that always finds a solution with nice properties [14]. This, however, does not apply to many settings, most notably the NetFlix setting where DVDs are to be allocated to the subscribers who don’t initially own any DVDs.

In [2], Abraham et al. proposed a solution for this problem by defining the notion of *popular matchings*, and giving an efficient algorithm that finds a popular matching, if one exists. Roughly speaking, a popular matching is a matching that is preferred by a majority of the agents to any other matching. The main drawback of this notion is that it does not necessarily exist. In fact, it is easy to see that even a simple instance with three agents and three items, all having the same preference list, does not contain any popular matching.

In this paper, we partially answer this question by showing

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC’06, June 11–15, 2006, Ann Arbor, Michigan, USA.
Copyright 2006 ACM 1-59593-236-4/06/0006 ...\$5.00.

that in a certain probabilistic model, if the number of items is larger than the number of agents by a small multiplicative factor, then almost surely popular matchings exist. More precisely, we show that there is a threshold $\alpha \approx 1.42$, such that if the number of items is at least α times the number of agents, then the instance admits a popular matching with high probability. This is particularly applicable in a setting like Netflix, where the number of available DVDs tend to be considerably larger than the number of subscribers.

Related Work. The notion of popular matchings (a.k.a. *majority assignment*) was first introduced by Gardenfors [7] in the context of stable matchings. Abraham, Irving, Kavitha, and Mehlhorn [2] introduced this notion for markets with preferences on one side, and proved several interesting characterization results and gave fast algorithms for finding a popular matching, even if the preference lists are not strict. Their characterization result (presented in the next section) provides the basis for our proof. They also give some experimental results with a probabilistic model similar to ours, except they only present experiments for the case that the number of agents is equal to the number of items, and therefore their experimental results are mostly negative. Mestre [10] and Abraham and Kavitha [3] show several generalizations of the setting and the results of [2]. Abraham, Chen, Kumar, and Mirrokni [1] study and compare several solution concepts (including popular matching) for the DVD rental market both in a static and in a dynamic setting.

The rest of this paper is organized as follows. In Section 2, we present the definitions, and a characterization result from [2] that will be used in our proof. In Section 3, we prove the main result of our paper. Section 4 contains a discussion of various extensions of our results, and some open questions and future directions.

2. DEFINITIONS AND PRELIMINARIES

Consider a situation where a number of items (e.g., Netflix DVDs, or jobs) must be allocated among n agents. We denote the set of agents by A and the set of items by B , and let $m = |B|$. Throughout this paper we assume $m \geq n$. Each agent can receive at most one item, and each item can be allocated to at most one agent. In other words, an allocation of the items to the agents is a *matching* between A and B .

The preference of each agent $a \in A$ is given by an ordered list p_a of a subset of items in B . In general, this ordering does not have to be *strict*, and can express *ties* between items. An item i that is not listed on p_a is considered unacceptable to a , i.e., a prefers to receive no item than to receive i . Throughout this paper, we assume that preference lists are strict and complete, i.e. each list p_a is a permutation of B . We will see in Section 4 that this assumption is only a simplifying assumption for our positive result.

Consider two matchings M and M' between the sets A and B . If the number of agents in A that prefer their match in M to their match in M' is more than the number of agents that prefer their match in M' to their match in M , we say that M is *more popular* than M' (denoted by $M \succ M'$). A matching M is called *popular* if there is no other matching that is more popular than M . Notice that the relation \succ is not necessarily acyclic; therefore, popular matchings need not exist in general.

Abraham et al. [2] gave a nice characterization of popular matchings that will be used in this paper. Here we present their result in the case that preference lists are all strict and complete. This characterization is in terms of two functions f and s defined below. For every agent a , we define $f(a)$ as the first item on p_a . An item $i \in B$ is called an f -item if $i = f(a)$ for some agent a . For every agent a , $s(a)$ is defined as the first item on a 's list that is not an f -item (notice that since preference lists are complete and $m \geq n$, such an item always exists). We define a bipartite graph G with agents in A on one side and items in B on the other side. An agent a has two edges connecting it to the items $f(a)$ and $s(a)$. A matching M between A and B is A -perfect if every agent $a \in A$ is matched under M . The following result is proved in [2].

THEOREM A. [2] *A matching M is popular if and only if*

- (i) *every f -item is matched in M , and*
- (ii) *M is an A -perfect matching in G .*

COROLLARY A. [2] *An instance of the popular matching problem admits a popular matching if and only if the graph G defined above has an A -perfect matching.*

3. EXISTENCE OF POPULAR MATCHINGS

As we mentioned in the previous section, there are simple instances that do not admit any popular matching. This is an unsatisfactory feature of the concept of popular matching (as opposed to, for example, stable matchings which always exist). A natural question is whether such instances are rare, i.e., whether we can expect a random instance of the popular matching problem to have a solution. In order to answer this question, we first need to clarify what we mean by a *random instance*.

Following previous work on random instances of the stable marriage problem [9], we consider the following distribution: Each agent in A picks her preference list p_a independently and uniformly at random from the set of all permutation of B . In this section, we show that a random instance drawn from the above distribution almost surely admits a popular matching if the number of items m is greater than a certain linear threshold.

THEOREM 1. *Assume $\alpha = \alpha^* + \epsilon$, where $\epsilon > 0$ and $\alpha^* \approx 1.42$ is the solution of the equation $x^2 e^{-1/x} = 1$. Then the probability that a random instance of the popular matching problem with n agents and $m = \alpha n$ items admits a popular matching tends to 1 as n tends to infinity.*

PROOF. By Corollary A, we only need to show that the probability that the graph G defined in the previous section has an A -perfect matching tends to 1. Notice that the graph G for a random instance can be constructed using the following procedure: each node $a \in A$ independently and uniformly picks a vertex in B as their first neighbor (this corresponds to $f(a)$). Let F denote the set of all vertices in B that are chosen as the first neighbor of some $a \in A$. Next, every node in $a \in A$ independently picks a uniformly random vertex in $S := B \setminus F$ as their second neighbor (this corresponds to $s(a)$). It is easy to see that the distribution of G constructed using the above procedure is exactly the same as the distribution of G for a random instance of the popular matching problem.

We start by proving the following lemma, which provides an estimate for the size of S . Notice that the random variable X defined in this lemma has exactly the same distribution as $|S|$.

LEMMA 1. *Let $m = \alpha n$, and define the random variable X by picking n elements of the set $\{1, \dots, m\}$ independently and uniformly at random and letting X be the number of elements in this set that are not picked. Then, $E[X] = e^{-1/\alpha}m - \Theta(1)$ and $\text{Var}[X] < E[X]$.*

PROOF. For $j \in \{1, \dots, m\}$, let X_j be 1 if and only if j is not picked in our experiment. Clearly, $X = \sum_j X_j$. By linearity of expectations we have

$$E[X] = \sum_j E[X_j] = m \cdot (1 - \frac{1}{m})^n = (e^{-1/\alpha} - \Theta(\frac{1}{m}))m.$$

Furthermore, we can estimate the variance of X using the following inequality from [4]:

$$\begin{aligned} \text{Var}[X] &\leq E[X] + \sum_{i \neq j} \text{Cov}[X_i, X_j] \\ &= E[X] + \binom{m}{2} \left((1 - \frac{2}{m})^n - (1 - \frac{1}{m})^{2n} \right) \\ &< E[X], \end{aligned}$$

where $\text{Cov}[X_i, X_j] := E[X_i X_j] - E[X_i]E[X_j]$ is the covariance of the variables X_i and X_j . This completes the proof of the lemma. \square

We now give a necessary and sufficient condition for the existence of an A -perfect matching in G in terms of unicyclicity of the components of another random graph G' defined below. Graph G' is bipartite with parts F' and S' , with $F' = B$ and $S' = S$. Corresponding to every vertex $a \in A$ in G , we put one edge e_a in G' between the vertex corresponding to $f(a)$ in F' , and the vertex corresponding to $s(a)$ in S' . We call the parts F' and S' the f -side and the s -side of the graph, and $f(a)$ and $s(a)$ the f -endpoint and the s -endpoint of e_a , respectively. We claim the following.

LEMMA 2. *There is an A -perfect matching in G if and only if every connected component of G' contains at most one cycle.*

PROOF. Assume we are given an A -perfect matching M in G . For every $a \in A$, we can orient the edge e_a of G' toward the endpoint corresponding to $M(a)$. In this way, we obtain an orientation of G' in which every vertex has at most one incoming edge. Conversely, it is easy to see that every orientation of G' that satisfies this property corresponds to an A -perfect matching in G . Therefore, we only need to prove that G' admits such an orientation if and only if every component of G' has at most one cycle.

If every connected component of G' has at most one cycle, we can orient the edges of G' as follows: for every component C of G' , if C contains a cycle, then orient this cycle in one of the two directions, and orient every other edge of C away from this cycle. For components C that do not contain any cycle, we pick an arbitrary vertex and orient all edges away from this vertex. It is easy to see that in this orientation, every vertex has at most one incoming edge.

On the other hand, assume G' has an orientation in which every vertex has at most one incoming edge, and consider a

connected component C of G' . Since the in-degree of each vertex in C is at most one, the number of edges in C is at most the number of vertices in C . This means that C is a tree plus at most one edge, and therefore contains at most one cycle. \square

We now need to bound the probability that G' contains a connected component with more than one cycle. Such connected components are often called *complex components*. It is well-known [5, 4] that the Erdős-Renyi random graph almost surely does not contain any complex component before the appearance of the giant component. Unfortunately, the graph G' is not exactly an Erdős-Renyi random graph. However, we can use similar techniques with a few additional tricks to prove our result.

The main difficulty with G' is that the number of vertices on its s -side is not fixed, and depends on the outcome of the random choices. The following lemma gets around this difficulty, by showing that fixing the number of vertices on the s -side does not increase the probability of a rare event defined on the graph G' significantly. More precisely, we define the random graph $G(a, b, M)$ as follows: $G(a, b, M)$ is a bipartite graph, with a vertices in the first part and b vertices in the second part, and M edges selected independently and uniformly at random from the set of all possible ab edges. Notice that in this definition edges are picked with replacement, so the graph might contain parallel edges.

LEMMA 3. *Assume $m = \alpha n$, and E is an arbitrary event defined on graphs. If for every fixed integer $h \in [e^{-1/\alpha}m - m^{2/3}, e^{-1/\alpha}m + m^{2/3}]$, the probability of E on the random graph $G(m, h, n)$ is at most $O(1/n)$, then the probability of E on G' is at most $O(n^{-1/3})$.*

PROOF. Let us denote the probability of E on the random graphs G' and $G(a, b, M)$ by $\Pr_{G'}[E]$ and $\Pr_{G(a,b,M)}[E]$, respectively. We sometimes drop the subscript if it is clear from the context which random graph the probability is taken over. By definition, for every fixed value of h , the distribution of G' , conditioned on $|S'| = h$ is exactly the same as the distribution of $G(m, h, n)$ conditioned on having exactly h isolated vertices (i.e., zero-degree vertices) in the first part of this graph (i.e., the part with m vertices). Let us denote the number of isolated vertices in the first part of $G(m, h, n)$ by $s(G(m, h, n))$. Therefore, by the above statement, we have

$$\begin{aligned} \Pr_{G'}[E] &= \sum_h \Pr_{G'}[E \mid |S'| = h] \cdot \Pr_{G'}[|S'| = h] \\ &= \sum_h \Pr_{G(m,h,n)}[E \mid s(G(m, h, n)) = h] \cdot \Pr[X = h], \end{aligned} \tag{1}$$

where X is the random variable defined in Lemma 1. By the Bayes rule and the fact that $s(m, h, n)$ has the same distribution as X , we have

$$\begin{aligned} \Pr[E \mid s(G(m, h, n)) = h] &= \\ &= \frac{\Pr[s(G(m, h, n)) = h \mid E] \cdot \Pr_{G(m,h,n)}[E]}{\Pr[X = h]}. \end{aligned} \tag{2}$$

Let I denote the interval $[E[X] - \delta, E[X] + \delta]$ for a value of

δ that will be fixed later. Equations 1 and 2 together imply

$$\Pr_{G'}[E] \leq \Pr[|X - \mathbb{E}[X]| > \delta] + \sum_{h \in I} \Pr[s(G(m, h, n)) = h|E] \cdot \Pr_{G(m, h, n)}[E] \quad (3)$$

By the Chebyshev's inequality, the first term in the right-hand side of the above inequality is at most $\text{Var}[X]\delta^{-2}$. In the second term, we bound $\Pr[s(G(m, h, n)) = h|E]$ by 1. Therefore, we obtain

$$\begin{aligned} \Pr_{G'}[E] &\leq \frac{\text{Var}[X]}{\delta^2} + \sum_{h \in I} \Pr_{G(m, h, n)}[E] \\ &\leq \frac{\mathbb{E}[X]}{\delta^2} + 2\delta \max_{h \in I} \Pr_{G(m, h, n)}[E] \end{aligned} \quad (4)$$

Now, we set $\delta = \frac{1}{2}m^{2/3}$. By Lemma 1, the interval I is contained in $[e^{-1/\alpha}m - m^{2/3}, e^{-1/\alpha}m + m^{2/3}]$. Therefore, $\max_{h \in I} \Pr_{G(m, h, n)}[E] = O(1/n)$. Thus,

$$\Pr_{G'}[E] \leq \frac{O(m)}{m^{4/3}} + m^{2/3} \cdot O\left(\frac{1}{n}\right) = O(n^{-1/3}),$$

as desired. \square

Given the above lemma, we only need to show that the graph $G(m, h, n)$ almost surely does not contain any complex components. This graph is essentially a bipartite variant of the Erdős-Renyi graph $G(n, M)$ [5, 4], and we can use similar techniques to prove our result.

LEMMA 4. *Let $m = \alpha n$, and h be an arbitrary number in the interval $[e^{-1/\alpha}m - m^{2/3}, e^{-1/\alpha}m + m^{2/3}]$. Then the probability that the random graph $G(m, h, n)$ contains a complex component (i.e., a component with more than one cycle) is at most $O(1/n)$.*

PROOF. Let X and Y be subsets of the vertices of the first and the second part of the graph $G(m, h, n)$, respectively. We define an event $BAD_{X,Y}$ that indicates that the subgraph of $G(m, h, n)$ induced by the vertices $X \cup Y$ contains one of the following graphs as a *spanning* subgraph: (i) two vertices joined by three disjoint paths (ii) two disjoint cycles joined by a path disjoint from the two cycles. We call these subgraphs *bad* subgraphs. Notice that every graph that contains a complex component must contain a bad subgraph.

We prove that with high probability none of the events $BAD_{X,Y}$ happen. Fix X and Y , and denote $k_1 = |X|$, $k_2 = |Y|$, and $k = k_1 + k_2$. Notice that if $|k_1 - k_2| > 1$, then the graph induced by $X \cup Y$ cannot contain any of the bad subgraphs as a spanning subgraph. Thus, we may assume that $k_1, k_2 \geq (k-1)/2$.

It is easy to see that the number of non-isomorphic bad graphs with k_1 vertices in the first part and k_2 vertices in the second part is at most $2k^2$, and for each such graph there are $k_1! \times k_2!$ ways to "put" this graph on the vertex set $X \cup Y$. The probability that all $k+1$ edges of this subgraph are picked in our procedure is at most

$$(k+1)! \binom{n}{k+1} \left(\frac{1}{mh}\right)^{k+1}$$

Therefore, the probability of $BAD_{X,Y}$ can be bounded using

the union bound by

$$2k^2 k_1! k_2! (k+1)! \binom{n}{k+1} \left(\frac{1}{mh}\right)^{k+1} \leq 2k^2 k_1! k_2! \left(\frac{n}{mh}\right)^{k+1}.$$

By the union bound, the probability that at least one of the events $BAD_{X,Y}$ happens is at most

$$\begin{aligned} \Pr\left[\bigvee_{X,Y} BAD_{X,Y}\right] &\leq \sum_{k_1, k_2} \binom{m}{k_1} \binom{h}{k_2} 2k^2 k_1! k_2! \left(\frac{n}{mh}\right)^{k+1} \\ &\leq \sum_{k_1, k_2} \frac{m^{k_1}}{k_1!} \times \frac{h^{k_2}}{k_2!} \times 2k^2 k_1! k_2! \left(\frac{1}{\alpha h}\right)^{k+1} \\ &= \sum_{k_1, k_2} \frac{2k^2}{h} \cdot \alpha^{-(k+1)} \left(\frac{m}{h}\right)^{k_1} \\ &\leq \sum_{k=1}^{\infty} \frac{O(k^2)}{n} \cdot \alpha^{-k} \left(e^{-1/\alpha} - m^{-1/3}\right)^{-k/2} \\ &= \frac{O(1)}{n} \sum_{k=1}^{\infty} k^2 (\alpha^2 (e^{-1/\alpha} - m^{-1/3}))^{-k/2}. \end{aligned}$$

By the assumption $\alpha > \alpha^*$, if n is large enough, we have $\alpha^2 (e^{-1/\alpha} - m^{-1/3}) > 1$, and therefore the above sum converges. Thus, the probability that at least one of the events $BAD_{X,Y}$ happens is at most $O(1/n)$. \square

The above lemma together with Lemmas 2 and 3 show that the probability that a random instance of the popular matching problem admits a popular matching is at least $1 - O(n^{-1/3}) = 1 - o(1)$. \square

4. DISCUSSION

In this paper we showed that in a probabilistic model, if the number of items is more than the number of agents by a small factor, then popular matchings almost surely exist. In our probabilistic model, every agent has a complete and strict preference list, drawn independently and uniformly at random. This raises several questions, some of which addressed below, and some posed as open questions.

Tightness. Using the tools developed in the previous section and standard techniques used to prove the emergence of a giant component in random graph theory [5, 4], we can prove that if the ratio m/n is bounded by a constant smaller than 1.42, then almost surely the instance does not contain a popular matching. In other words, there is a *phase transition* at $\alpha = \alpha^*$. The proof is long, and is not included in this paper, but the sketch of the proof is as follows: we use Lemmas 2 and 3 to reduce the problem to analyzing the probability that $G(m, h, n)$ does not have any complex component (notice that this time we need to define the event E in Lemma 3 as *not* having a complex component). Then, we show that for a small enough δ , with high probability, the graph $G(m, h, (1-\delta)n)$ contains a *giant component*. This is done by approximating the breath-first search tree starting from an arbitrary vertex of $G(m, h, (1-\delta)n)$ with a *Poisson branching process*, similar to the argument used to show the emergence of the giant component in the Erdős-Renyi graph. Finally, we observe that with high probability, at least two of the remaining δn edges will land in the giant component, thereby creating a complex component.

Dealing with ties. In our model we assumed that there is no tie in the preference list of the agents. The main reason behind this assumption was that there is no *clean* probabilistic model with ties. Our model is an accurate model in some settings (e.g., in NetFlix the system does not allow the subscribers to express ties in their preferences), but not in others. However, intuitively ties make the task of finding a popular matching easier. More precisely, if we resolve the ties in all preference lists in an arbitrary way, we obtain an instance without any ties, and every popular matching for this instance is also a popular matching for the original instance. Therefore, our result extends to any model with ties in which there is a way to break the ties so that the resulting distribution is close to our distribution. In particular, this includes the probabilistic model introduced by Abraham et al. [2]. On the other hand, it might be (and intuitively it should be) possible to improve the threshold α^* for models with ties. This is confirmed by experimental result in [2], where it is suggested that if there are enough ties, even an instance with $m = n$ has a reasonable probability of containing a popular matching.

Short lists. We assumed in our model that the agents have complete preference lists. This is far from being realistic in most applications. However, it is not hard to see that restricting the lengths of the preference lists in any instance only increases the likelihood that the instance contains a popular matching (this statement is easy to prove using the characterization result of [2]). Again, one might ask if restricting the lengths of the preference lists can improve the threshold α^* . Our conjecture is that restricting the length of all preference lists to k in our model can only improve the threshold by an amount exponentially small in k . The intuitive reason for this is that for each agent a , the probability that a cannot find an $s(a)$ on its list is exponentially small in k .

Incentive compatibility. One important issue in real-world markets is the issue of incentive compatibility. For example, do users of NetFlix have an incentive to reveal their true preferences? Some argue that in many such systems it is beneficial not to disclose some of the options that are lower on one's preference list, since if such options are disclosed, there is a chance that the system selects one of the inferior options and this might decrease the chance of getting the top options. Such situations are not desirable from a game theoretic point of view. Therefore, a natural question is whether a mechanism that always computes a popular matching is incentive compatible. This is an interesting open question. A similar question for stable matchings was solved by Immorlica and Mahdian [8]. We don't know if similar techniques can solve the incentive compatibility question for popular matchings.

Other solution concepts. Popular matching is not the only solution concept proposed for matching markets with one-sided preferences. For example, another solution concept is *rank-maximal* matching, which is a matching that maximizes the number of agents who receive their first choice, and then (subject to the first maximization) the number of agents who receive their second choice, and so on. Analyzing various other solution concepts from the point of view of performance and incentive compatibility is an interesting open direction. A first step in this direction is the recent paper by Abraham et al. [1] which studies the performance of

several solution concepts both in static (i.e., one-shot) and dynamic (i.e., repeated) scenarios.

Acknowledgments. I would like to thank David Wilson and Joel Spencer for fruitful discussions on Lemma 3, and Vahab Mirrokni and Nicole Immorlica for discussions about the concept of popular matchings.

5. REFERENCES

- [1] D. Abraham, N. Chen, V. Kumar, and V.S. Mirrokni. Assignment problems in rental markets. 2006.
- [2] D. J. Abraham, R. W. Irving, T. Kavitha, and K. Mehlhorn. Popular matchings. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 424–432, 2005.
- [3] D. J. Abraham and T. Kavitha. Voting paths. manuscript, 2005.
- [4] N. Alon and J. Spencer. *The probabilistic method*. John Wiley & Sons, 2000.
- [5] B. Bollobás. *Random Graphs*. Cambridge University Press, 2001.
- [6] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
- [7] P. Gardenfors. Match making: assignment based on bilateral preferences. *Behavioural Sciences*, 20:166–173, 1975.
- [8] N. Immorlica and M. Mahdian. Marriage, honesty, and stability. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 53–62, 2005.
- [9] D. E. Knuth, R. Motwani, and B. Pittel. Stable husbands. *Random Structures and Algorithms*, 1:1–14, 1990.
- [10] J. Mestre. Weighted popular matchings. manuscript, 2005.
- [11] A. Roth, T. Sonmez, and M. U. Unver. Kidney exchange. *Quarterly Journal of Economics*, 119(2):457–488, May 2004.
- [12] A. E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *forthcoming in Econometrica*.
- [13] A.E. Roth and E. Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American Economic Review*, 89:748–780, 1999.
- [14] L. Shapley and H. Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1:23–28, 1974.