

Matching Markets: Design and Analysis

Thesis Proposal

David J. Abraham

1 Introduction

A market consists of buyers and sellers of some commodity, say a DVD movie. In this thesis, we assume the role of market operator. Our goal is to ensure that the market has certain desirable properties, which may include fairness, revenue maximization, and so on. We achieve these properties by designing rules for how buyers and sellers can perform transactions. Our main focus will be on matching markets.

An example of a matching market is Netflix [31], the mail-based DVD rental business. Each day, Netflix allocates DVDs from its warehouses to customers entitled to rent a DVD on this day. Typically, there are not enough DVDs to give all customers their first choice, so Netflix must decide on an allocation that somehow takes into account the preferences of customers over DVDs. This is an example of a bipartite market, since there are two disjoint sets of participants, namely the customers and the DVDs/Netflix, and each set of participants only interacts with the other set. Other bipartite matching markets include the allocation of doctors to internships, children to public schools, and midshipman to naval jobs.

More generally, a matching market can be non-bipartite, in which case participants can interact with any other participant. One type of non-bipartite matching market is a barter-exchange market. An example of a such a market is Swaptree [36], the mail-based DVD swapping business. (Swaptree also allows swapping of CDs, games, books, and so on.) Customers give Swaptree a list of DVDs they want to own, and a list of DVDs they want to trade. Whenever Swaptree finds an exchange (possibly involving three or more customers), it organizes for the customers to swap their DVDs by mail. In this market, the buyers and sellers coincide, and no money is exchanged.

There are many other examples of barter-exchange markets. For example, there is a currently a move in the United States to establish a national kidney-exchange market [38]. Many kidney-disease patients have a willing but blood-type incompatible donor amongst their family. In the past, these potential donors have been sent home, leaving the patients to wait for cadaver kidneys, which are in very limited supply. However, more and more, doctors are performing exchanges in which patients swap their incompatible donors in order

to obtain a compatible donor. We have worked with doctors and economists on solving computational and incentive problems for this market. In particular, our recent work on an algorithm for finding optimal exchanges [2] has been used by the Alliance for Paired Donation, which covers many hospitals networks in the US, Korea and India.

Matching markets have long held a central place in the game theory, mechanism design and computer science literature. In this thesis, as well as making contributions on the design and analysis of traditional matching markets, we explore a range of new problems that have arisen from more modern markets. These problems include special constraints on which buyers and sellers can interact, and buyers and sellers arriving and leaving the market over time. Perhaps the most significant new aspect of modern markets though is their size, and as such, the problem of finding efficient algorithms for matching buyers and sellers is the central theme in this thesis.

2 Matching Agents with Items

A bipartite matching market with one-sided preferences consists of a set of *agents* and a set of *items*. Each agent has a *preference list*, ranking the subset of items it finds *acceptable* in order of preference (first choice, second choice, and so on). The *outcome* we seek is a *matching* of agents to acceptable items in which no agent is matched with more than one item, and no item is matched with more than one agent.

Preferences over items extend naturally to preferences over outcomes / matchings: an agent prefers all matchings in which it is allocated an acceptable item. Given two such matchings, an agent prefers the one in which it obtains the better-ranked item. If an agent does not prefer one matching over another, it is *indifferent* between the two.

A *matching algorithm* selects a matching based on the preferences of agents over outcomes. One area of research involves trying to find computationally-efficient algorithms that produce optimal matchings according to some social welfare objective.

2.1 Static Objective Functions

There are many possible social welfare objectives. For example, we might assign weights to possible allocations, with better-ranked allocations having more weight. The aim then is to find a *maximum-weight* matching, where the weight of a matching is the weight sum of all the allocations it makes. This setting is itself quite general, as many weight functions are possible.

With a sharply decreasing weight function, we can find *rank-maximal* matchings, which match the maximum number of agents to their first choice, and subject to this, the maximum number to their second choice, and so on. Alternatively, we can also find *egalitarian* matchings, which match the minimum

number of agents to say their 10th choice, and subject to this, match the minimum number to their 9th choice, and so on. There has been some nice recent work on finding faster algorithms for these objectives [18, 29]

2.1.1 Pareto Optimality

The weakest reasonable objective is *Pareto optimality*. A matching M is *Pareto optimal* if there is no matching M' such that some agent prefers M' to M , and no agent prefers M to M' . Such a matching has the property that no coalition of agents can all improve their allocations (say by exchanging items with one another), without requiring some other agent to be worse off.

In [4] (not included in this thesis), we give several structural and algorithmic results on Pareto optimality. For example, we introduce the concept of a *matching signature* to show that a natural class of greedy algorithms characterizes the set of Pareto-optimal matchings. We also give an algorithmic characterization of Pareto optimality, which leads to a faster algorithm for finding a maximum-cardinality Pareto-optimal matching.

2.1.2 Popularity

For any social welfare objective, we need a way to compare two matchings (for example, by weight, or Pareto dominance). This comparison can be made democratic by allowing the agents to vote between the matchings. A matching M' is *more popular than* another M if the number of agents preferring M' to M exceeds the number of agents preferring M to M' . We say that M is *popular* if there is no matching M' more popular than M . This definition of optimality extends the well-known Condorcet voting rule for presidential elections - a candidate wins only if she loses no pairwise election with any other candidate.

The notion of popular matching was originally introduced by Gardenfors [12] in the context of the full stable marriage problem. It is well known that every stable marriage instance admits a weakly stable matching (one for which there is no pair who strictly prefer each other to their partners in the matching). In fact, there can be an exponential number of weakly stable matchings, and so Gardenfors considered the problem of finding one with additional desirable properties, such as popularity. Gardenfors showed that when preference lists are strictly ordered, every stable matching is popular. He also showed that when preference lists contain ties, there may be no popular matching.

Figure 1 contains an example matching market in which $\{a_1, a_2, a_3\}$ is the set of agents, $\{i_1, i_2, i_3\}$ is the set of items, and each agent prefers i_1 to i_2 , and i_2 to i_3 . Consider the matchings $M_1 = \{(a_1, i_1), (a_2, i_2), (a_3, i_3)\}$, $M_2 = \{(a_1, i_3), (a_2, i_1), (a_3, i_2)\}$ and $M_3 = \{(a_1, i_2), (a_2, i_3), (a_3, i_1)\}$. It is easy to verify that none of these matchings is popular, since M_2 is more popular than M_1 , M_3 is more popular than M_2 , and M_1 is more popular than M_3 . In fact, this instance admits no popular matching, the problem being, of course, that the *more popular than* relation is not acyclic.

a_1	:	i_1	i_2	i_3
a_2	:	i_1	i_2	i_3
a_3	:	i_1	i_2	i_3

Figure 1: An instance for which there is no popular matching.

Given a matching market, the *popular matching problem* is to find a popular matching, or prove that no such matching exists. Unfortunately, the number of matchings may be exponential in the number of agents, and so it is unclear if we can efficiently solve this problem through a sequence of pairwise comparisons of matchings.

As such, the time complexity of popular matching was unknown, until, in a recent paper [5], we developed a novel characterization of popularity that led to the following surprising results: the popular matching problem can be solved in *linear* time when preference lists contain no ties. More generally, when preference lists are unrestricted, popular matching has the same time complexity as the classical maximum matching problem, and hence can be solved in polynomial time.

Since the publication of these results, several papers have helped build the theory of popular matchings. For example, [22] shows how to use a faster randomized maximum matching algorithm to give a faster randomized algorithm for popular matching. [28] extends the popular matching algorithm to handle agents with different voting weights. [26] also extends the algorithm to efficiently handle items with multiple copies. [23] gives an efficient combinatorial algorithm for finding various *optimal* popular matchings. [25] addresses the existence problem by showing that a popular matching exists with high probability when the number of items is a small factor larger than the number of agents, and agent preference lists are drawn at random. If there is no popular matching, [27] shows that the problem of finding a least-unpopular matching is NP-hard. However, [23] shows that it is always possible to find an approximately popular matching in polynomial time. Finally, [35] makes the first attempt to study the time complexity of popular matchings in the stable marriage problem by giving a polynomial-time algorithm to check if a given matching is popular in a restricted setting.

2.2 Dynamic and Decentralized Markets

In a *dynamic* matching market, agents and items can enter and leave the market, and agents can change their preferences arbitrarily. Whenever such a change occurs, the matching in place may no longer be optimal, in which case we are required to recompute a new optimal matching. In a *decentralized* matching market, there is no central authority to find an optimal matching. Instead, the market itself has some mechanism to make local improvements to the current matching. This process of local improvements may eventually lead to a locally optimal matching, which, hopefully, is not too far from being globally optimal.

We studied these two types of markets when the notion of optimality is a popular matching [6]. In both cases, we assume that moving from one matching to another requires a consensus of votes - i.e. the new matching must be strictly more popular. However, even if a popular matching exists, there may be no popular matching that is pairwise strictly more popular than the current matching. Hence, we look for *voting paths*, which are *sequences* of matchings, each strictly more popular than the last, ending in a popular matching.

There is no a priori reason to expect that such paths always exist. After all, the *more popular than* relation is not acyclic, and so perhaps there are some matchings from which we cannot avoid cycling. And even if such a path does exist, it may have exponential length, since there can be an exponential number of matchings.

Nonetheless, we show the following surprising result: Given a market with a popular matching and an existing matching M , there is always a voting path of length at most 2 from M to some popular matching. Moreover, we can find a shortest-length voting path in only linear time.

This result is analogous to a series of papers [19, 33, 37, 9] on decentralized mechanisms in the stable matching literature. The well-known mechanisms for stable matching, due to Gale/Shapley [11] and Irving [17], require a central body to collect preferences and dictate the final matching. Alternatively, in a decentralized setting, a blocking pair (i.e. a man and woman who prefer each other to their current partners) will *act* locally by divorcing their current partners and marrying each other. Knuth [19] showed that if the divorced partners also marry each other, it is possible for this process to cycle. However, when divorced partners are not required to marry each other, and every blocking pair has some probability of acting next, Roth and Vande Vate [33] show by way of a potential argument that there is always a path to a stable matching.

In our setting, the analogue of a blocking pair is a coalition of applicants who (i) prefer some matching M' to the current matching M , and (ii) have sufficient numbers to win a vote between M' and M . It is not too difficult to give a potential argument to prove the existence of voting paths (at least for the restriction to strictly-ordered preference lists). However, in [6], we used more powerful techniques from matching theory, which in addition to proving existence, also give the length-2 bound. As with the result in [33], this means that as long as every matching more popular than the current one has some probability of an up-or-down vote, then in the limit, a decentralized mechanism will lead to a popular matching.

2.3 Keyword Auctions

Web search engines use keyword auctions to sell the advertising space alongside their algorithmic search results. When a user searches for a keyword (e.g. “digital camera”), merchants wanting to target this keyword (e.g. Canon, Nikon and Amazon.com) bid to have their advertisements displayed on the search results page. Based on these bids, the search engine performs a matching of merchants (agents) to advertisement slots (items). Since users are more likely to click on

slots at the top of the page, merchants value these slots more highly than slots lower down the page.

Small changes in the auction mechanism can lead to big changes in both the bidding behavior of merchants, and the revenue collected by the search engine. As such, there has been significant recent interest in the design and analysis of keyword auction mechanisms (see [24] for a recent summary).

In this work [1], we introduce the class of *layerable* mechanisms for the keyword auctions. Several existing mechanisms are layerable, including the Generalized Second Price mechanism [10], which is used by the major search engines, the ladder auction [13], and the VCG mechanism [39, 8, 15]. Layerable mechanisms can be decomposed into a collection of layers, where each layer consists of an auction in which merchants are indifferent between the slots. This decomposition leads to a simple technique for designing and analyzing keyword auction mechanisms: instead of working in the general setting, we can focus on the restricted layer setting.

The main focus of this work is to characterize the class of layerable mechanisms by giving necessary and sufficient conditions for when the decomposition technique is possible. We demonstrate our new technique by simplifying the correctness proofs of some existing mechanisms, and showing how they can be extended to more general settings. We also introduce a new mechanism that allows a search engine to truthfully bid in its own keyword auction. This is becoming increasingly common as Google, Microsoft and Yahoo compete to sell their other services, such as operating systems, video-game consoles, photo-sharing websites and so on. With existing mechanisms, this leads to a conflict of interest as the search engine may mis-represent its bids in order to make other merchants pay the search engine more money. Building on the ladder auction, our new mechanism ensures that the merchants can not be exploited in this way, whilst also guaranteeing that the search engine gains at least as much utility as in the standard ladder auction. Microsoft has submitted two patent applications based on this work.

Our work on layerable mechanisms is most closely related to Aggarwal, Goel, and Motwani’s paper [13] on ladder auctions, and Roughgarden and Sundararajan’s paper [34] on the trade-off between efficient and optimal auctions. In a ladder auction, if merchant wins the i th-best slot, then for one part of its allocation, it pays the same price as the merchant who wins the worst slot, and for a second part of its allocation it pays the same price as the merchant who wins the second-worst slot, and so on. Our definition of a layerable mechanism is inspired by this property of a ladder auction. Roughgarden and Sundararajan [34] independently noted that efficient auctions and optimal auctions can both be thought of as superposition of multi-item auctions. It turns out this is similar to our definition of a layerable mechanism. In contrast, we work in a significantly more general setting, and our focus is on defining the class of layerable mechanisms, exploring its properties, and using the decomposition as a way to design new mechanisms.

3 Matching Agents with other Agents

The setting in the previous section involved matching agents and items, where agents had preferences over items, whilst items had no preferences over agents. In this section, we consider a more general model in which the aim is to match agents with other agents. Both the stable marriage and stable roommate problems are modelled in this setting. The stable marriage problem is a bipartite market in which we match, for example men with women. The stable roommate problem is a non-bipartite market in which we match, for example, college students to be dormitory roommates. As well as considering matching markets, we will also consider the following generalization.

A *barter-exchange market* consists of a set of agents, each of whom owns an item. Each agent has a preference list over the set of items it finds acceptable. These preferences extend naturally to the set of agents who own these items. An *exchange* is a disjoint collection of cycles of agents, where each agent wants the item owned by the next agent in its cycle.

Our main focus has been on kidney-exchange markets, where patients with kidney disease trade their incompatible donors. In the United States, more than 60,000 patients are on the waiting list for the transplant of a cadaver kidney. Because demand for cadaver kidneys far outstrips supply, thousands of patients die each year waiting for their transplant opportunity. Since it is illegal to buy and sell human organs, barter-exchange markets, which are almost universally regarded as ethical and legal, present the only real options for patients unlikely to survive the wait for a cadaver kidney.

3.1 Finding a Maximum Exchange

The problem of finding an exchange involving the maximum number of agents, also known as the directed cycle cover problem, is polynomial-time solvable. However, if a fixed limit (more than 2) is placed on the maximum allowable cycle size, we show the problem becomes NP-hard [2]. The kidney-exchange market gives one example of why such a limit may be imposed. In a 3-way kidney exchange, all 6 operations have to be done *simultaneously* for incentive/legal reasons. Cycles larger than 2 or 3 are not allowed then because transplant centers do not have the capacity to perform them.

Natural approaches to dealing with NP-hard problems include looking for good heuristics or approximation algorithms. However, such approaches are unacceptable in this context - any loss in optimality may lead to patients dying unnecessarily. As such, several attempts have been made at devising an exact algorithm, however all have had trouble in scaling up to solve markets of the size expected by the upcoming national kidney-exchange market.

In recent work, we have written a new algorithm that successfully surpasses what is required by this market. The algorithm uses column generation techniques from linear programming theory, as well as novel upper and lower bounding techniques. We have used this algorithm, in conjunction with the Alliance for Paired Donation, to organize optimal exchanges of donors from hospital net-

works throughout the US. The United Network for Organ Sharing is currently reviewing our algorithm for use in the upcoming national exchange program in the US.

3.2 Stable Roommates with Globally-Ranked Pairs

The stable roommates problem [11, 17] involves pairing up a set of agents, each of whom ranks the others in order of preference. In this work [7], we introduce a restriction of the roommate problem in which preferences are derived from a global ranking function on the agent pairs. This restriction is motivated by the following application:

When two (patient,donor) pairs are matched with each other in order to swap donors, we are not certain if the swap can occur until expensive last-minute compatibility tests are performed on the donors and patients. If either potential transplant in the swap is incompatible, the swap is cancelled and the two patients must wait for a future match run. Doctors can rank potential swaps by their chance of success. This ranking induces patient preferences, since patients prefer to be involved in swaps with better chances of success.

In [7], we study the globally-ranked pairs restriction and give a polynomial-time algorithm to find a rank-maximal stable matching (amongst all stable matchings, it maximizes the number of rank-1 pairs, and subject to this, maximizes the number of rank-2 pairs and so on). This is the first generalization of the rank-maximal matching algorithm due to Irving et al. [18] to a non-bipartite setting. Also, we prove several hardness results in an even more restricted setting, including for example the problem of finding a stable matching with the minimum number of weakly blocking pairs.

3.3 Fair Exchanges

Even in a maximum exchange, it may not be possible to include all the agents. This exposes the market operator to charges of bias - why was one maximum exchange selected over another? Can we be fair to the agents, even though some will not end up in the final exchange?

The fair exchange problem is to find a distribution over maximum exchanges, so that, as much as possible, the *probability* of each agent being in final exchange is equalized. Recently, it was shown that in pairwise exchange markets (where the maximum cycle size is 2), there always exists a distribution with very strong equality and incentive properties [32].

In recent work, we derived an easier proof of the existence of this distribution. We then used the associated characterization, together with ideas from [21], as the basis for a polynomial-time algorithm to find an egalitarian distribution.

We are currently trying to extend this work to settings in which the agents have more complicated preferences. In particular, our aim is to find efficient algorithms for producing a distribution over “optimal” matchings, such that the utility of the agent with the worst expected allocation is maximized, and subject to this, the utility of the agent with the second-worst expected utility

is maximized, and so on. Various definitions of optimality can be considered, including Pareto optimality, cardinality, rank-maximality and popularity.

3.4 Online Markets

In a kidney-exchange market, patients and donors enter and leave the market over time. Because we cannot see the future, any exchanges we make now may turn out to be suboptimal later. However, we cannot delay making exchanges, since some patients may die in the meantime, and anyway, patients will have an incentive to leave the market and organize their own exchanges.

Incentives in offline markets have been well-studied. To avoid coalitions of agents from leaving the market and organizing their own exchanges, we need to find a *core* exchange in which these coalitions are at least as well off. We have generalized this solution concept to the online setting. In the strongest generalization, whenever a change occurs in the market, we must find and immediately commit to a maximum exchange. In a weaker generalization, we must find a maximum exchange, but instead of committing to it immediately, we only need to guarantee that the patients involved will be in some exchange in the future. This gives us more control - when a new patient arrives, we may be able to include them by augmenting the existing exchange when otherwise there would be no one left for them to exchange with.

We have several competitive ratio results for algorithms in these core settings. For example, in the weaker core setting for pairwise exchange markets, we derived an algorithm that, if allowed to see a small way into the future, exchanges at least two-thirds as many patients as is optimal. Combined with a statistical model of patients and donors entering and leaving the market, this lookahead algorithm could be used to help make better decisions about which pairs to exchange in the current market.

3.4.1 Query-Commit Model

There is another separate online aspect to the kidney-exchange market. To determine if a patient and potential donor are compatible, there is a sequence of increasingly accurate and expensive medical tests. The spectrum goes from basic health checks for the donor, to blood-type checks, to mixing samples of the patient and donor's blood. Unfortunately, due to financial and logistical reasons, it is not feasible to perform all the pairwise patient-donor tests upfront. Instead, when finding a exchange, we must *query* the proposed transplants before we *commit* to the exchange. In the weaker core setting for pairwise exchange markets, this is not a problem - if we discover that a pair is incompatible, they simply rejoin the market, and we can augment our possibly sub-optimal exchange. However, in the stronger core setting, we may be forced to commit to cycles that are not in any maximum exchange. We call this the *query-commit* model, and we are currently working on several different ways to approach and solve the problem, in particular for pairwise exchange markets.

In one approach, we are trying to generalize the ranking algorithm [20, 14] for online bipartite matching to the non-bipartite setting. The online aspect of the bipartite matching problem involves vertices from one side of the market arriving one by one. However, this same algorithm can be reinterpreted in the query-commit model, and so this seems like a promising avenue of research. In a separate approach, we are looking at a relaxation of the problem in which we can make several simultaneous queries per patient, and only have to commit to one of these, if any queries are successful. Finally, in another approach, we consider the problem when information is given on the probability that each donor and patient are compatible. In principle, there is an optimal plan/ordering to query the edges in the graph. Similar to our work on finding large exchanges [2], we would like to explore the feasibility of efficiently producing an optimal plan for querying the edges, even though, in the worst-case, this is not polynomial-time solvable.

4 Proposed Work

Listed below are several problems, in order of priority, that we may investigate for this thesis.

1. Query-Commit Model: Generalize the ranking algorithm for online bipartite matching to the setting in which we can query the existence of an edge, but if it is there, we are committed to including it in the matching.
2. Online Cycle Covers in the Core: Improve competitive ratio bounds for the kidney-exchange setting in which vertices arrive to the market over time and have a deadline before leaving the market.
3. Optimal Query-Commit Algorithm: In the query-commit model with probabilities on the edges, derive and implement an algorithm that finds an optimal plan for querying edges and can solve large real-world instances.
4. Popular matching in the stable marriage setting: Following [35], attempt to extend the theory of popular matching from matching markets with one-sided preferences to markets with two-sided preferences (i.e. the stable marriage setting).
5. Fair matching: Find efficient algorithms for producing a distribution over "optimal" matchings, such that the utility of the agent with the worst expected allocation is maximized, and subject to this, the utility of the agent with the second-worst expected utility is maximized, and so on. Various definitions of optimality can be considered, including Pareto optimality, cardinality, rank-maximality and popularity.
6. Auctioneer-as-Merchant Mechanism: Improve mechanism that allows a search engine to bid in its own keyword auction so that it applies in more general settings.

7. Optimal Keyword Auction: Using our decomposition technique, derive a generalization of Myerson's optimal revenue mechanism for more general settings.

References

- [1] D. J. Abraham, A. Asadpour and K. Jain. *Layerable Mechanisms for Keyword Auctions*. In submission.
- [2] D. J. Abraham, A. Blum, T. Sandholm. *Clearing Algorithms for Barter-Exchange Markets: Enabling Nationwide Kidney Exchanges*. In Proceedings of ACM-EC 2007: the Eighth ACM Conference on Electronic Commerce, 2007.
- [3] D. J. Abraham, N. Chen, V. Kumar and V. Mirrokni. *Assignment problems in rental markets*. In Proceedings of WINE '06: the 2nd Workshop on Internet and Network Economics, 2006.
- [4] D. J. Abraham, K. Cechlárová, D. F. Manlove, K. Mehlhorn. *Pareto-optimality in house allocation problems*. In Proceedings of ISAAC '04: the 15th Annual International Symposium on Algorithms and Computation, volume 3341 of Lecture Notes in Computer Science, pages 3-15, Springer-Verlag, 2004.
- [5] D. J. Abraham, R. W. Irving, T. Kavitha, and K. Mehlhorn. *Popular Matchings*. In Proceedings of SODA '05: the 16th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 424-432, 2005.
- [6] D. J. Abraham and T. Kavitha. *Dynamic matching markets and voting paths*. In Proceedings of SWAT '06: the 10th Scandinavian Workshop on Algorithm Theory, 2006.
- [7] D. J. Abraham, A. Levavi, D. F. Manlove, and G. O'Malley. *The stable roommates problem with globally-ranked pairs*. In Proceedings of WINE '07: the 3rd International Workshop on Internet and Network Economics, 2007.
- [8] E. H. Clarke *Multipart pricing of public goods*. Public Choice **11**(1) (1971) 17-33
- [9] E. Diamantoudi, E. Miyagawa, and L. Xue. *Random paths to stability in the roommate problem*. Games and Economic Behavior, 48(1):18-28, 2004.
- [10] B. Edelman, M. Ostrovsky and M. Schwarz. *Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords*. American Economic Review **97**(1) (2007) 242-259
- [11] D. Gale and L.S. Shapley. *College admissions and the stability of marriage*. American Mathematical Monthly, 69:9-15, 1962.

- [12] P. Gardenfors. *Match Making: assignments based on bilateral preferences*. Behavioural Sciences, 20:166–173, 1975.
- [13] G. Aggarwal, A. Goel, and R. Motwani. *Truthful auctions for pricing search keywords*. In Proceedings of the 7th ACM conference on Electronic commerce. (2006) 1–7
- [14] B. Birnbaum and C. Mathieu. *On-line bipartite matching made simple*. SIGACT News, Vol. 39, No. 1. (2008), pp. 80-87.
- [15] T. Groves. *Incentives in teams*. Econometrica **41**(4) (1973) 617–631
- [16] C. Huang, T. Kavitha, D. Michail, and M. Nasre. *Bounded Unpopularity Matchings* In Proceedings of SWAT '08: the 11th Scandinavian Workshop on Algorithm Theory
- [17] R. W. Irving. *An efficient algorithm for the "stable roommates" problem*. Journal of Algorithms, 6:577-596, 1985
- [18] R. W. Irving, T. Kavitha, K. Mehlhorn, D. Michail, and K. Paluch. *Rank-maximal matchings*. In Proceedings of SODA '04: the 15th ACM-SIAM Symposium on Discrete Algorithms, pages 68–75. ACM-SIAM, 2004.
- [19] D. E. Knuth. *Stable Marriage and its Relation to Other Combinatorial Problems*. CRM Proceedings and Lecture Notes, volume 10, American Mathematical Society. English translation of *Mariages Stables*, Les Presses de L'Université de Montréal, 1976
- [20] R. M. Karp, U. V. Vazirani and V. V. Vazirani. *An optimal algorithm for on-line bipartite matching*. In Proceedings of STOC '90: the twenty-second annual ACM symposium on Theory of Computing, 352 – 358, 1990
- [21] A. Katta and J. Sethuraman. *A solution to the random assignment problem on the full preference domain*. Journal of Economic Theory, 131(1):231-250, 2006.
- [22] T. Kavitha and C. Shah. *Efficient Algorithms for Weighted Rank-Maximal Matchings and Related Problems*. In Proceedings of ISAAC '06: the 17th Annual International Symposium on Algorithms and Computation, pages 153-162, 2006.
- [23] T. Kavitha and M. Nasre. *Optimal Popular Matchings* In MATCH-UP: Matching Under Preferences - Algorithms and Complexity
- [24] S. Lahaie, D. Pennock, A. Saberi and R. Vohra. *Chapter 28: Sponsored search auctions*. In Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V., eds.: Algorithmic Game Theory. Cambridge University Press (2007)
- [25] M. Mahdian. *Random popular matchings*. In Proceedings of the 7th ACM Conference on Electronic-Commerce, pages 238–242, 2006.

- [26] D. F. Manlove and C. Sng. *Popular matchings in the capacitated house allocation problem*. In Proceedings of ESA 2006, the 14th Annual European Symposium on Algorithms, pages 492-503, 2006.
- [27] M. McCutchen *Least-Unpopularity-Factor Matching*. In Proceedings of LATIN 2008, the 8th Latin American Theoretical Informatics Symposium, pages 593-604, 2008.
- [28] J. Mestre. *Weighted popular matchings*. In Proceedings of the 33rd International Colloquium on Automata, Languages and Programming, (LNCS 4051), pages 715–726, 2006.
- [29] D. Michail. *Reducing rank-maximal to maximum weight matching*. Theoretical Computer Science, volume 389, issue 1-2, pages 125-132 , 2007
- [30] R. B. Myerson *Optimal auction design*. Mathematics of Operations Research **6**(1) (1981) 58–73
- [31] Netflix DVD Rental: see <http://www.netflix.com>.
- [32] A. E. Roth, T. Snmez and U. Unver. *Pairwise Kidney Exchange*. Journal of Economic Theory, (December, 2005) 125: 151-188.
- [33] A. E. Roth and J. H. Vande Vate. *Random Paths to Stability in Two-Sided Matching*. Econometrica, 58:1475-1480, 1990.
- [34] T. Roughgarden and M. Sundararajan *Is Efficiency Expensive?*. In the Third Workshop on Sponsored Search, 2007.
- [35] C. Sng. *Efficient algorithms for bipartite matching problems with preferences*. PhD thesis, University of Glasgow.
- [36] Swaptree barter-exchange market: see <http://www.swaptree.com>.
- [37] A. Tamura. *Transformation from arbitrary matchings to stable matchings*. Journal of Combinatorial Theory, Series A, 62(2):310-323, 1993.
- [38] UNOS pilot national kidney exchange market, press release: <http://www.unos.org/news/newsDetail.asp?id=1098>
- [39] W. Vickery *Counterspeculation, auctions, and competitive sealed tenders*. Journal of Finance **16** (1961) 8–37