

Trends in Computational Social Choice

6

Cite as: Ágnes Cseh. Popular Matchings. In Ulle Endriss (editor), *Trends in Computational Social Choice*, chapter 6, pages 105–122. AI Access, 2017.

<http://www.illc.uva.nl/COST-IC1205/Book/>

CHAPTER 6

Popular Matchings

Ágnes Cseh

6.1 Introduction

Matching problems lie at the heart of discrete mathematics. Their rich history reaches back over 100 years (König, 1916), including some milestones of complexity and algorithms, such as perfect, maximum weight and minimum cost matchings, together with their connection to network flow and vertex cover problems.

In this chapter we focus on matching markets under preferences, where each market participant expresses their preferences as an ordered list of possible scenarios. Our task is to find a matching that is optimal with respect to these preferences. If the agents express their preferences in a cardinal manner, then the most common aim is to maximize the total utility of the agents. This yields the concept of *maximum weight* or *minimum cost matchings*. If preferences are ordinal, one might want to guarantee that no two agents are inclined to form a coalition in order to deviate from the given solution. This concept corresponds to the well-known notion of *stable matchings*. In coordinated allocation mechanisms, the central authority of control usually aims at a solution that matches a large number of agents. Thus, negotiating the size and the optimality of the matching with respect to agents' preferences is a problem that occurs naturally.

Popularity is a concept that offers an attractive trade-off between these two notions. In short, a popular matching M guarantees that no matter what alternative matching is offered on the market, the majority of the agents will opt for M . Moreover, $|M|$ is relatively close to the size of the maximum matching in the market. The notion was first defined by Gärdenfors (1975) and surprisingly, it comes from cognitive science, where such a majority decision is a well-motivated potential focus of investigation. After Gärdenfors' paper, decades passed without any achievement in the topic. Recently, an impressive amount of top-tier publications have demonstrated the importance of popular matchings.

6.1.1 Definition of Popular Matchings

Popular matchings can be defined in various market settings. For the sake of generality we assume that we are given a not necessarily complete and not necessarily bipartite graph with n vertices and m edges, where each vertex represents

an agent and each edge stands for an acceptable agent-agent pair. Each agent expresses her preferences over her adjacent agents in the form of a strictly ordered list—every time ties are allowed, we explicitly say so. Figure 6.2 depicts such an instance. The ranking of each vertex can be seen on the edges incident to it. Lower numbers mean a better rank: for instance, agent a_1 's top choice is b_2 , her second choice is b_1 , her third choice is b_3 , while she finds b_4 unacceptable. To the left of the figure, the preference lists are shown. A matching is a set of edges in the graph so that no agent is matched to more than one other agent.

We compare matchings M_1 and M_2 in the following manner. Each vertex casts a vote for M_1 or M_2 or abstains from voting. If vertex v is matched to a better partner in M_1 than in M_2 or it is matched in M_1 and unmatched in M_2 , then v votes for M_1 . Analogous rules specify when v votes for M_2 . Finally, v abstains from voting if its situation is the same in both matchings: either because it is matched to a vertex with the same rank or because it is unmatched in both matchings. If the number of votes for M_1 is at least as large as the number of votes for M_2 , then we say that M_1 is *at least as popular as* M_2 . If M_1 receives strictly more votes than M_2 , then M_1 *defeats* M_2 , in other words, M_1 is more popular than M_2 . Note that the notion of defeat here is not transitive. Figure 6.2 shows an instance in which four matchings defeat each other in a circular manner.

Matching M is *popular* if it is at least as popular as any other matching in the instance. In other words, M does not get defeated by any matching in a comparison.

Besides the aforementioned roots in cognitive science, an approach to motivate the notion of popularity comes from voting theory. If we consider all possible matchings in an instance as the set of alternatives and let the agents vote, then it turns out that popular matchings form a well-defined subset of alternatives, namely the set of weak Condorcet winners (Condorcet, 1785). This set consists of the alternatives that beat or tie with every other alternative in a pairwise comparison.

In a similar spirit, popular matchings can be viewed as a special case of *maximal lotteries*, defined by Kreweras (1965) and Fishburn (1984) and rediscovered by several other researchers since then (Felsenthal and Machover, 1992; Laffond et al., 1993; Rivest and Shen, 2010). Chapter 1 of this book elaborates on maximal lotteries and points out the connection to popular mixed matchings, which we will discuss in Section 6.3.2.

6.1.2 Models and Chapter Structure

The most general setting involves an arbitrary graph representing a set of agents and the possible connections between them. We will refer to this setting as the non-bipartite model. Bipartite graphs play a distinguished role in matching markets. In bipartite graphs, the popular matching problem has been studied in the following two models.

- *One-sided model.* One side of the graph consists of agents who have preferences and votes, while the other side is formed by objects with no preferences or votes. This setting is analogous to the house allocation market model.

- *Two-sided model.* Vertices on both sides are agents, so they all have preferences and cast votes. This setting is analogous to the stable marriage model and it is a subcase of the non-bipartite model.

The remainder of this chapter is structured as follows. We start in Section 6.2 with a literature review of optimality concepts that can be seen as alternatives of popularity. After this, the two main building blocks follow in Sections 6.3 and 6.4, centered around the two above described bipartite models. Finally, in Section 6.5 we discuss results in the non-bipartite setting. Our approach is mainly algorithmic, but at the end of each section we also elaborate on more applied studies in the literature. The aim of this chapter is to give a structured overview of the rapidly growing field of the theory of popular matchings.

6.2 Related Literature

Defining optimality on markets with ordinal preferences is far from straightforward. In this subsection we sketch a number of alternative optimality concepts to popularity. These concepts are grouped based on the model they are most common to be used in.

6.2.1 One-sided Model

A number of optimality concepts for one-side markets have been studied in the literature. The most prevalent concept is *Pareto-optimality*. Informally, a matching is Pareto-optimal if there is no other matching in which at least one agent is better off, whilst no agent is worse off. Pareto-optimal matchings always exist in the one-sided model and at least one can be found using the strategyproof *Random Serial Dictatorship* mechanism, as shown by Abdulkadiroğlu and Sönmez (1998). The shortcomings of Pareto-optimal matchings are that even the largest one of them can be as small as half the size of a maximum matching. Moreover, the definition allows all but one agents to receive poor choices in order to avoid a single agent to be allocated to a slightly worse object than she has.

Other optimality concepts are defined based on the *profile* of the matching. This is a array of numbers, where the i th element is the number of agents who are matched to their i th choice object. Matchings that maximize the profile in a lexicographic sense are called *rank-maximal* matchings, defined by Irving (2003). Similarly to Pareto-optimal matchings, rank-maximal matchings always exist and can be found in polynomial time (Irving et al., 2006). On the other hand, even the largest rank-maximal matchings can be as small as half the size of a maximum matching asymptotically. To overcome this disadvantage, *greedy maximum matchings* (Michail, 2007) and *generous maximum matchings* (Abraham et al., 2006) were also defined, both of them are based on the profile of the matching.

6.2.2 Two-sided Model

The literature on two-sided markets is clearly dominated by *stable matchings* first discussed by Gale and Shapley (1962). A matching is called stable if it is not

blocked by any pair of agents. A *blocking pair* comprises two agents not matched to each other who are either single or prefer to be matched to one another than to their respective partners in the matching. By their well-known deferred acceptance algorithm, Gale and Shapley showed that a stable matching always exists and can be found in linear time. A characteristic feature of stable matchings is the so called Rural Hospitals Theorem (Roth, 1984), part of which states that the set of matched agents is identical in all stable matchings. In particular, all stable solutions have the same cardinality.

Pareto-optimal matchings can be defined in two-sided markets analogously to one-sided markets. Clearly every stable matching in a market with strict preferences is Pareto optimal, but Pareto optimal matchings can be twice as large as stable matchings. Sng (2008) showed that a maximum Pareto optimal matching can be found in polynomial time. Profile-based optimality concepts were studied in the paper of Huang and Kavitha (2012).

6.2.3 Non-bipartite Model

The non-bipartite version of the stable matching problem is usually referred to as the *stable roommates problem*, which is quite different from its classical variant from an algorithmic point of view. First of all, a stable solution is not guaranteed to exist, which was pointed out by Gale and Shapley (1962) already, but there is a polynomial algorithm to find one, or a proof for its nonexistence (Irving, 1985).

The definition of Pareto-optimal matchings carries over to this setting. Just as in the simpler models, a Pareto-optimal matching always exists and a largest Pareto-optimal matching can be found in polynomial time (Abraham and Manlove, 2004). Profile-based optimality concepts were studied by Abraham et al. (2008).

The detailed study of existing literature on popular matchings is spread thorough the upcoming sections. Nevertheless, we would like to point out earlier surveys on the topic, such as those of Kavitha (2008), Mestre (2008) and Chapter 7 in the book of Manlove (2013).

6.3 One-sided Model

We start the study of popular matchings in one-sided models, where the two sides of the bipartite graph $G = (A \cup B, E)$ represent agents (A) and objects (B), respectively. The defining property of this setting is that only vertices in A cast votes, the objects have neither preferences nor a right to vote. Such one-sided markets are particularly suitable for modeling object allocation, such as in the well-known house allocation problem.

This section starts with the existence of popular matchings and the problem of finding a maximum size popular matching. Then we turn to the most important extensions of the problem from a theoretical point of view. Finally, we discuss some more applied approaches, such as computational studies and fairness concepts.

6.3.1 Finding a Max Size Popular Matching

As we have already mentioned, the initial paper of Gärdenfors (1975) was followed by decades of silence in the matching community. The notion of popular matchings in bipartite graphs reappeared in 2005, in the conference version of a paper by Abraham et al. (2007), who worked on the one-sided model. The main result of their paper is a polynomial algorithm for deciding whether a popular matching exists.




Figure 6.1: No popular matching exists in this instance. The dotted gray matching $\{a_2b_1, a_3b_2\}$ is more popular than the dashed gray matching $\{a_1b_1, a_2b_2\}$, because both a_2 and a_3 prefer it. Similarly, the black matching $\{a_1b_2, a_3b_1\}$ defeats the dotted gray, and the dashed gray defeats the black.

In the one-sided model, the existence of a popular matching is not guaranteed. Figure 6.1 depicts an instance equivalent to the famous voting paradox of Condorcet (1785), where none of the matchings is popular. In this context, the following result answers the most striking algorithmic question of the topic.

Theorem 6.1 (Abraham et al., 2007). *There is an $O(n+m)$ algorithm that outputs either a largest cardinality popular matching or a proof for its nonexistence.*

Note that this result not only answers the question on the existence of a popular matchings but it also guarantees a maximum cardinality solution, if any exists. Maximizing the cardinality of the outputted matching is particularly important, since the main motivation behind popular matchings is that the concept unites preference-optimality and large size.

The notion of first and second choice objects plays a crucial role in the algorithm of Abraham et al. (2007). The object ranked highest by agent a_i is called a_i 's *first choice object*. The *second choice object* of agent a_i is the object that was not marked as a first object by anyone and it is ranked the highest among such objects in a_i 's preference list. The following lemma sheds light to the importance of these two definitions.

Lemma 6.2. *M is popular in an instance of the one-sided model if and only if*

1. *every first choice object is assigned in M and*
2. *each agent is matched to either their first or second choice object.*

From this lemma, it is easy to see how to search for a popular matching. First we need to construct the graph, where each agent is adjacent to their first and second choice objects only and then check for a matching that matches all agents and all first choice objects. To reach a maximum cardinality popular matching, one needs to ensure that as few agents are matched to their dummy last resort object as possible. This can be done by a simple augmenting path algorithm, for example.

A slightly modified version of the above described algorithm serves to solve the general case in which preference lists may contain ties. This is also presented by Abraham et al. (2007), who gave an $O(\sqrt{nm})$ time algorithm for the maximum cardinality one-sided popular matching problem with ties.

6.3.2 Theoretical Results

This subsection is built up by three parts, each of them centered around capacitated instances, the relaxation of popular matchings and weighted instances, respectively.

Capacitated Extension

The many-to-one matching case clearly belongs to the most intuitive generalizations of the popular matching problem. In this setting, each object is assigned a positive capacity, which is the upper bound on the number of agents who can get this object allocated to them. On the other hand, each agent is allowed to receive one object at most. Due to this latter point, the notion of comparing two matchings does not need to be modified at all.

A characterization analogous to the one in Lemma 6.2 was given by Sng and Manlove (2010). They also presented the following results on the complexity of finding a popular matching.

Theorem 6.3 (Sng and Manlove, 2010). *There is an $O(\sqrt{C}n + m)$ algorithm to determine if an instance of the capacitated popular matching problem admits a popular matching, and if so, to find a largest such matching, where C is the total capacity of the objects. If ties are allowed, the time complexity of the algorithm changes to $O((\sqrt{C} + n)m)$.*

Defining a voting rule in the many-to-many setting is complex, and as a matter of fact, there are several legitimate options to study. Lexicographic order was studied by Paluch (2014). She provided a characterization of popular matchings and showed that finding a popular matching or a proof for its nonexistence is NP-hard.

Relaxing Popular Matchings

Having established in Theorem 6.1 that we can distinguish instances with and without popular matchings in polynomial time, the relaxation of popularity is the next intuitive move. We will now sketch the two most common relaxations, namely least unpopular matchings and popular mixed matchings.

Least unpopular matchings. First, the notion of *least unpopular* matchings was proposed to deal with instances that had no popular matchings (McCutchen, 2008). Assume that M_1 and M_2 are two matchings in the same instance. We say that M_2 dominates M_1 by a factor of $\frac{u}{v}$, if u is the number of agents who strictly prefer M_2 to M_1 and v is the number of agents who strictly prefer M_1 to M_2 . For instance, matching $\{a_2b_1, a_3b_2\}$ in Figure 6.1 dominates matching $\{a_1b_1, a_2b_2\}$ by a factor of 2. The *unpopularity factor* of a matching M is the maximum factor by which it is dominated by any other matching, ignoring matchings that give $u = v = 0$. According to this definition, a matching is popular if and only if its unpopularity factor is exactly 1.

McCutchen (2008) also defined an alternative concept to measure the degree of popularity, called the *unpopularity margin*. This is defined in the same manner as the unpopularity factor, except that one subtracts the numbers of votes instead of dividing them. More precisely, M_2 dominates M_1 by a margin of $u - v$, if u is the number of agents who strictly prefer M_2 to M_1 and v is the number of agents who strictly prefer M_1 to M_2 . Returning to the same example in Figure 6.1, we can state that $\{a_2b_1, a_3b_2\}$ dominates $\{a_1b_1, a_2b_2\}$ by a margin of 1. The *unpopularity margin* of M is the maximum margin by which M is dominated by any other matching. According to this definition, a matching is popular if and only if its unpopularity margin is exactly 0.

Theorem 6.4 (McCutchen, 2008, Manlove, 2013). *There is an $O(m\sqrt{n})$ time algorithm to find the unpopularity factor of a matching and there is an $O(m\sqrt{n} \cdot \log n)$ time algorithm to find the unpopularity margin of a matching. These algorithms work even in the presence of ties.*

Theorem 6.5 (McCutchen, 2008). *The problems of finding a least unpopularity factor matching and a least unpopularity margin matching are NP-hard.*

McCutchen (2008) also showed that the unpopularity factor of any matching is always an integer. In particular, if G does not admit a popular matching, then the unpopularity factor is at least 2 for all matchings in G .

Note that matchings with least unpopularity factor are exactly the matchings with least unpopularity margin. The least unpopularity margin is equivalent to the Simpson-Kramer voting rule (Kramer, 1977; Simpson, 1969), which selects as the winner the candidate whose greatest pairwise defeat is smaller than the greatest pairwise defeat of any other candidate.

Popular mixed matchings. The second optimality concept proposed for instances without popular matchings is *popular mixed matchings* (Kavitha et al., 2011). The notion of popularity is kept intact here, while the matching condition is relaxed. A mixed matching is a probability distribution over matchings in the input graph. The vote of an agent can be adjusted in a straightforward manner to this setting. For instance, taking each of the three matchings of cardinality 2 in Figure 6.1 with probability $\frac{1}{2}$ is a mixed matching that defeats matching $\{a_1b_1, a_2b_2\}$ by exactly one vote, because a_1 casts half a vote for $\{a_1b_1, a_2b_2\}$, a_1 casts half a vote for the mixed matching and finally, a_3 fully votes for the mixed matching.

Theorem 6.6 (Kavitha et al., 2011). *Popular mixed matchings exist even in the presence of ties in preference lists, and they can be found in polynomial time.*

Kavitha et al. (2011) presented two algorithms for the problem. Interestingly, one of them relies on the algorithm of McCutchen (2008) to determine the unpopularity margin of a matching, while the other one uses linear programming techniques.

Optimizing over Weights

A natural extension of the popular matching problem is to consider graphs with edge or vertex weights and search for the weight-optimal popular solution.

Edge weights. McDermid and Irving (2011) gave a structural characterization of popular matchings, and efficient algorithms to enumerate them. This led to the following result.

Theorem 6.7 (McDermid and Irving, 2011). *In the presence of edge weights, a maximum weight maximum cardinality popular matching or a proof for its nonexistence can be found in $O(n + m)$ time.*

Presenting a reduction to the minimum cost assignment problem Matsui and Hamaguchi (2016) proposed a polynomial time algorithm for finding a maximum weight popular matching, irrespective of its cardinality.

Vertex weights. Another intuitive extension of the problem is to assign an arbitrary positive weight to each agent. The vote of that agent then counts with the multiplicity given by this weight. Mestre (2014) considered this extension and showed the following.

Theorem 6.8 (Mestre, 2014). *In the presence of vertex weights, a maximum weight maximum cardinality popular matching or a proof for its nonexistence can be found in polynomial time even in the presence of ties.*

6.3.3 Applied Approaches

Upon establishing the characterization of popular matchings in the one-sided model, Abraham et al. (2007) ran experiments to test the probability of the existence of a popular matching in randomly generated instances with $|A| = |B|$. Their results show that the ratio of solvable instances drops radically as the length of preference lists increase. Obviously, if every list is of length 1, a popular matching is guaranteed to exist. Out of 1000 instances with 100 agents and lists of length 10 only 2 were solvable, while the same setting with preference list of length 20 or more did not allow a single instance to admit a popular matching. The intuition behind this phenomenon is that dummy posts as second choice objects increase the probability of a matching assigning all agents. Due to Lemma 6.2, this latter is a necessary condition for the existence of a popular matching. To complement these slightly discouraging results, Mahdian (2006)

showed that a popular matching exists with high probability, if $|B|$ is a small multiplicative factor larger than $|A|$.

Popular mixed matchings were studied from the view of fairness concepts by Aziz et al. (2013). They showed that in some instances, popularity and envy-freeness are incompatible if $n \geq 3$. On the other hand, if a popular and envy-free assignment exists, it can be computed in polynomial time. They also proved that there is no strategyproof popular random assignment rule if $n \geq 3$. Weaker notions of envy-freeness and strategyproofness were also showed to be incompatible with popularity by Brandt et al. (2017), for $n \geq 5$ and $n \geq 7$, respectively.

Nasre (2013) studied strategyproofness in the classical integral matching case. She assumed that a_1 is the sole manipulative agent who is aware of the true preference lists of all other agents and that a central authority chooses an arbitrary popular matching. Thus, the goal of a_1 is to falsify her preference list to weakly improve the post she gets matched to in the falsified instance with any chosen popular outcome. She showed that the optimal cheating strategy for a single agent to get better always can be computed in $O(n + m)$ time when preference lists are all strict and in $O(\sqrt{nm})$ time when preference lists are allowed to contain ties.

6.4 Two-sided Model

In this section we turn to bipartite graphs with preferences on both sides. Such instances model situations where vertices on both sides represent agents and thus are given the right to vote. Initially, Gärdenfors (1975) defined the notion of popularity for these two-sided markets with preferences on both sides. He also showed that if all preference lists are strict, then any stable matching is popular; thus a popular matching always exists and can be found in linear time using the well-known deferred acceptance algorithm of Gale and Shapley (1962). Huang and Kavitha (2013) later gave a characterization of popular matchings based on augmenting paths. They also came up with an $O(m)$ algorithm to test whether a given matching is popular.

This section is structured similarly to Section 6.3. It starts with the problem of finding a maximum size popular matching, then we elaborate on extensions of the problem, such as the case of ties or maximum weight popular matchings. Finally, we discuss some more applied approaches.

6.4.1 Finding a Max Size Popular Matching

Popular matchings of the same instance can differ in size, as illustrated by a sample instance from Kavitha (2015) in Figure 6.2. Besides the two stable matchings $M_1 = \{a_1b_1, a_2b_2\}$ and $M_2 = \{a_1b_2, a_2b_1\}$ the perfect matching $M_3 = \{a_1b_3, a_2b_4, a_3b_2, a_4b_1\}$ is also popular. This gives us popular matchings of size 2 and 4. None of the four matchings of size 3 is popular, because they defeat each other in a circular manner. Note one more nicety of this instance: no popular matching defeats any of these size 3 matchings strictly in a comparison.




Figure 6.2: Sample instance with popular matchings of size 2 and 4.

As demonstrated by this instance, a strikingly important feature of popular matchings is that they beat stable matchings in size. As a matter of fact, any stable matching is a minimum size popular matching (Huang and Kavitha, 2013). The size of a stable matching in G can be as small as $|M_{\max}|/2$, where M_{\max} is a maximum matching in G . Relaxing stability to popularity yields larger matchings and it is easy to show that a largest popular matching has size at least $\frac{2}{3} \cdot |M_{\max}|$. This result begs for the question about finding a maximum size popular matching.

Efficient algorithms for computing a maximum size popular matching were given by Huang and Kavitha (2013) and Kavitha (2014). Here we present the latter one.

The algorithm can be seen as a 2-round Gale-Shapley algorithm. Each man in the instance can have two states: unpromoted or promoted. At start, every man is unpromoted and the deferred acceptance rounds of the Gale-Shapley algorithm begin. According to the rules of that, each man proposed to his most preferred woman. As a response, each woman temporarily accepts the offer she ranks highest and rejects the rest of the proposing men. Rejected men now proceed to their second-choice woman and compete for her by submitting a proposal. Later proposals can result in the rejection of the earlier temporarily accepted man. The Gale-Shapley algorithm terminates with a stable matching.

At this stage, all men in the instance are unpromoted. The second round of the algorithm starts with the promotion of all men who remained unmatched at the end of the Gale-Shapley algorithm. These men now get the chance to walk through their original preference list one more time, from the top to the bottom. Women find promoted men more attractive than unpromoted men, irrespective of their original preferences. Two men of the same state will always be compared according to the original list of the woman. It is easy to see that the proposals of promoted men can result in some other men becoming single. Every time a man reaches the end of his preference list for the first time, he gets promoted. If a man reaches the end of his preference list for a second time as well, he is deactivated.

This algorithm outputs a maximum size popular matching, moreover, its time-complexity is the same as of the Gale-Shapley algorithm.

Theorem 6.9 (Huang and Kavitha, 2013, Kavitha, 2014). *In the two-sided model with strictly ordered lists there is an $O(m)$ algorithm that outputs a largest cardinality popular matching.*

It is easy to see that once a woman got a proposal in this algorithm, she will never become single. In particular, women matched in the output of the Gale-Shapley algorithm will be matched in the computed popular matching. As Hirakawa et al. (2015) have shown, more is true: every maximum cardinality popular matching assigns the same set of agents, which is a superset of the agents matched in any stable matching.

Naturally, one could allow men to be promoted after the second round as well. Kavitha (2014) showed that more Gale-Shapley rounds yield an even larger matching, but this increment in size comes at a price of an increased unpopularity factor. This latter can be defined in the two-sided model analogously as in the one-sided model.

Theorem 6.10 (Kavitha, 2014). *For every k where $2 \leq k \leq n$, there is a matching M_k such that $|M_k| \geq \frac{k}{k+1}|M_{\max}|$ and $u(M_k) \leq k - 1$, where M_{\max} is a maximum matching in G and $u(M_k)$ is the unpopularity factor of M_k . This matching can be computed in $O(km)$ time via a k -round Gale-Shapley procedure.*

6.4.2 Theoretical Results

Popularity among Maximum Matchings

Motivated by the search for a matching that is of largest cardinality among popular matchings, Kavitha (2014) investigated the question of finding a maximum cardinality matching that is never defeated by any other maximum cardinality matching.

Theorem 6.11 (Kavitha, 2014). *A matching that is popular among maximum cardinality matchings always exists and can be found in $O(nm)$ time.*

Ties in Preference Lists

It turns out that ties have a massive effect on the complexity of popular matching problems in the two-sided model. When ties are allowed in preference lists on both sides, Biró et al. (2010) showed that deciding whether a popular matching exists is NP-complete. This result was further strengthened by Cseh et al. (2015) who also studied an intermediate variant between the 1- and 2-sided models with strict lists, namely if only agents in A have ordered preference lists ranking their neighbors, however agents on both sides cast votes—in this case, agents in B only care about being matched. Their results can be summarized as follows.

Theorem 6.12 (Cseh et al., 2015). *If one side of the bipartite graph has strict preference lists while on the other side each agent either puts its neighbors into a single tie or into a strict list, then deciding whether a popular matching exists is NP-complete.*

If one side of the bipartite graph has strict preference lists while on the other side each agent puts its neighbors into a single tie, then a popular matching or a proof for its nonexistence can be found in $O(n^2)$ time.

Optimizing over Weights

Currently there is no known method to find a maximum weight popular matching in a graph equipped with edge weights. Several results point in this direction, which justifies that the problem is clearly among the most riveting open questions in the area.

Cseh and Kavitha (2016) investigated the case of a forced edge in the graph. This refers to the problem in which there is a given forced edge e and we seek popular matchings that contain e . The problem is equivalent to searching for a maximum weight matching with weight function 1 on e and 0 elsewhere.

Theorem 6.13 (Cseh and Kavitha, 2016). *A popular matching containing a given forced edge e or a proof for its nonexistence can be found in $O(m)$ time.*

The same authors investigated the maximum-weight popular matching problem with complete lists.

Theorem 6.14 (Cseh and Kavitha, 2016). *If all preference lists are complete, then a maximum weight popular matching can be found in polynomial time.*

Besides considering special weight functions or preference lists, another approach is to relax the matching condition by permitting mixed matchings. A special case of those is half-integral matchings, in which edges are allowed to occur with value 0, $\frac{1}{2}$ or 1.

Theorem 6.15 (Kavitha, 2016). *The maximum weight popular half-integral matching problem can be solved in polynomial time.*

Most recently, Huang and Kavitha (2017) achieved remarkable structural results using LP techniques. Alongside other results they showed that there is always a half-integral popular matching among the maximum weight fractional popular matchings.

Theorem 6.16 (Huang and Kavitha, 2017). *The popular fractional matching polytope is half-integral and in the special case where a stable matching in the graph is a perfect matching, it is integral.*

6.4.3 Applied Approaches

Bhattacharya et al. (2015) studied a dynamic matching scenario, when agents and edges of the graph arrive and depart iteratively over time. The question is whether one can maintain a popular matching after each timeslot by modifying the given matching only in a few edges. They showed that maintaining popularity requires an amortized number of $\Omega(n)$ changes to the matching per round. Their result also answers an algorithmic question of independent interest. No algorithm is known for finding a popular matching by gradually building it up from a given matching, stepping from one matching to a more popular matching in each round. The negative result about maintaining popularity implies that two-sided instances might have no such paths to a popular matching, even for complete and strict preferences.

Chisca et al. (2016) propose the first constraint programming formulation of the popular matching problem. They encode preferences using the global cardinality constraint (Régín, 1996).

Popular matchings were proposed as a solution concept for task allocation in multi-camera networks by Cui and Jia (2013). According to users priority and task nature, different tasks are prioritized. For example, routine patrolling may have the lowest rank, while tasks that are triggered by motion detection are ranked highest. The authors run extensive simulations and demonstrate that popular matchings offer an attractive and efficient alternative to baseline approaches based on various greedy matching procedures.

6.5 Non-bipartite Model

The notion of popularity can be defined in not necessarily bipartite instances by a straightforward adjustment of the definition introduced in Section 6.1.1. We assume that all vertices represent agents and cast votes.

This section also follows the outline of the previous two sections. Due to the smaller volume, we do not separate the parts on existence, theoretical and applied approaches.

Chung (2000) was the first to observe that stable matchings are popular even in the non-bipartite case. Thus, if an instance with strict lists admits a stable matching, then the existence of a popular matching is also guaranteed. Some instances of the stable roommates problem do not admit a stable solution, yet they admit a popular matching, as demonstrated by Figure 6.3, first presented by Biró et al. (2010). Surprisingly, the complexity of deciding whether a non-bipartite instance admits a popular matching is unknown. Biró et al. (2010) proved that validating whether a given matching is popular can be done in polynomial time, even if ties are present in the preference lists.




Figure 6.3: The dotted gray edges mark the unique popular matching $M = \{a_1a_2, a_3a_4\}$. It is blocked by the edge a_2a_3 . The instance admits no stable matching.

Even though the main complexity question about popular matchings in non-bipartite instances has not been answered yet, there is a number of results marking a promising path towards it.

Huang and Kavitha (2017) showed that the polytope of popular fractional matchings is half-integral in the non-bipartite case, analogously to Theorem 6.16. This means that one can compute a maximum weight popular half-integral matching in polynomial time. They also showed that the problem of computing an integral maximum weight popular matching in a non-bipartite instance is NP-hard. Note that this still does not answer the open question on finding a largest cardinality popular matching, if any exists in the instance.

Some studies about extensions of the problem are also present in the literature. Huang and Kavitha (2013) have proved that the problem of computing a least unpopularity factor matching is NP-hard and presented instances where every matching has unpopularity factor $\Omega(\log n)$. On the positive side, they proved that every instance admits a matching whose unpopularity factor is $O(\log n)$, and such a matching can be computed in linear time.

6.6 Conclusion

In this chapter we have discussed the popular matching problem from an algorithmic point of view. We discussed existence, maximum size popular matchings, various theoretical and applied contributions in the cases of a bipartite market with one-sided and two-sided preferences, and finally in non-bipartite instances.

We have posed three open questions.

1. What is the complexity of finding a maximum weight popular matching in the two-sided model?
2. What is the complexity of finding a popular matching in the non-bipartite model?
3. What is the complexity of finding a largest cardinality popular matching in the non-bipartite model, if any exists in the instance?

Acknowledgment

This work was supported by the Hungarian Academy of Sciences under its Momentum Programme (LP2016-3/2016) and OTKA grant K108383.

Bibliography

- A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.
- D. J. Abraham and D. F. Manlove. Pareto optimality in the Roommates problem. Technical Report TR-2004-182, University of Glasgow, Department of Computing Science, December 2004.

- D. J. Abraham, N. Chen, V. Kumar, and V. S. Mirrokni. Assignment problems in rental markets. In *Proceedings of the 2nd International Workshop on Internet and Network Economics (WINE)*, volume 4286 of *Lecture Notes in Computer Science*, pages 198–213. Springer, 2006.
- D. J. Abraham, R. W. Irving, T. Kavitha, and K. Mehlhorn. Popular matchings. *SIAM Journal on Computing*, 37:1030–1045, 2007.
- D. J. Abraham, A. Levavi, D. F. Manlove, and G. O’Malley. The stable roommates problem with globally-ranked pairs. *Internet Mathematics*, 5:493–515, 2008.
- H. Aziz, F. Brandt, and P. Stursberg. On popular random assignments. In *Proceedings of the 6th International Symposium on Algorithmic Game Theory (SAGT)*, pages 183–194. Springer, 2013.
- S. Bhattacharya, M. Hoefer, C.-C. Huang, T. Kavitha, and L. Wagner. Maintaining near-popular matchings. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 504–515, 2015.
- P. Biró, R. W. Irving, and D. F. Manlove. Popular matchings in the marriage and roommates problems. In *Proceedings of the 7th International Conference on Algorithms and Complexity (CIAC)*, volume 6078 of *Lecture Notes in Computer Science*, pages 97–108. Springer, 2010.
- F. Brandt, J. Hofbauer, and M. Suderland. Majority graphs of assignment problems and properties of popular random assignments. *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2017.
- D. S. Chisca, M. Siala, G. Simonin, and B. O’Sullivan. A CP-based approach for popular matching. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*. AAAI Press, 2016.
- K. S. Chung. On the existence of stable roommate matchings. *Games and Economic Behavior*, 33(2):206–230, 2000.
- M. Condorcet. *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix*. L’Imprimerie Royale, 1785.
- Á. Cseh and T. Kavitha. Popular edges and dominant matchings. In *Proceedings of the 18th International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, *Lecture Notes in Computer Science*, pages 138–151. Springer, 2016.
- Á. Cseh, C.-C. Huang, and T. Kavitha. Popular matchings with two-sided preferences and one-sided ties. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 9134 of *Lecture Notes in Computer Science*, pages 367–379. Springer, 2015.
- L. Cui and W. Jia. Weighted capacitated popular matching for task assignment in multi-camera networks. In *Proceedings of the 25th International Teletraffic Congress (ITC)*, pages 1–4. IEEE, 2013.

- D. S. Felsenthal and M. Machover. After two centuries, should Condorcet's voting procedure be implemented? *Behavioral Science*, 37(4):250–274, 1992.
- P. C. Fishburn. Probabilistic social choice based on simple voting comparisons. *The Review of Economic Studies*, 51(4):683–692, 1984.
- D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
- P. Gärdenfors. Match making: assignments based on bilateral preferences. *Behavioural Science*, 20:166–173, 1975.
- M. Hirakawa, Y. Yamauchi, S. Kijima, and M. Yamashita. On the structure of popular matchings in the stable marriage problem - Who can join a popular matching? In *Proceedings of the 3rd International Workshop on Matching Under Preferences (MATCH-UP)*, 2015.
- C.-C. Huang and T. Kavitha. Weight-maximal matchings. In *Proceedings of the 2nd International Workshop on Matching Under Preferences (MATCH-UP)*, pages 87–98, 2012.
- C.-C. Huang and T. Kavitha. Popular matchings in the stable marriage problem. *Information and Computation*, 222:180–194, 2013.
- C.-C. Huang and T. Kavitha. Popularity, mixed matchings, and self-duality. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2294–2310. SIAM, 2017.
- R. W. Irving. An efficient algorithm for the “stable roommates” problem. *Journal of Algorithms*, 6(4):577–595, 1985.
- R. W. Irving. Greedy matchings. Technical report, TR-2003-136. University of Glasgow, Glasgow, UK, 2003.
- R. W. Irving, T. Kavitha, K. Mehlhorn, D. Michail, and K. Paluch. Rank-maximal matchings. *ACM Transactions on Algorithms*, 2(4):602–610, 2006.
- T. Kavitha. Ranked Matching. In *Encyclopedia of Algorithms*, pages 744–748. Springer, 2008.
- T. Kavitha. A size-popularity tradeoff in the stable marriage problem. *SIAM Journal on Computing*, 43:52–71, 2014.
- T. Kavitha. Popular matchings. Talk given at COST Action IC1205 meeting on Matching and Fair Division, 2015. URL <http://www.optimalmatching.com/COST2015/slides/KavithaTelikepalli.pdf>.
- T. Kavitha. Popular half-integral matchings. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 55, pages 22:1–22:13, Dagstuhl, Germany, 2016.
- T. Kavitha, J. Mestre, and M. Nasre. Popular mixed matchings. *Theoretical Computer Science*, 412:2679–2690, 2011.

- D. König. Über Graphen und ihre Anwendung auf Determinantentheorie und Mengenlehre. *Mathematische Annalen*, 77(4):453–465, 1916.
- G. H. Kramer. A dynamical model of political equilibrium. *Journal of Economic Theory*, 16(2):310–334, 1977.
- G. Kreweras. Aggregation of preference orderings. In *Mathematics and Social Sciences I: Proceedings of the seminars of Menthon-Saint-Bernard, France (1–27 July 1960) and of Gösing, Austria (3–27 July 1962)*, pages 73–79, 1965.
- G. Laffond, J.-F. Laslier, and M. Le Breton. The bipartisan set of a tournament game. *Games and Economic Behavior*, 5(1):182–201, 1993.
- M. Mahdian. Random popular matchings. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC)*, pages 238–242. ACM, 2006.
- D. F. Manlove. *Algorithmics of Matching Under Preferences*. World Scientific, 2013.
- T. Matsui and T. Hamaguchi. Characterizing a set of popular matchings defined by preference lists with ties. *arXiv preprint arXiv:1601.03458*, 2016.
- R. M. McCutchen. The least-unpopularity-factor and least-unpopularity-margin criteria for matching problems with one-sided preferences. In *Proceedings of the 8th Latin-American Theoretical Informatics Symposium (LATIN)*, volume 4957 of *Lecture Notes in Computer Science*, pages 593–604. Springer, 2008.
- E. McDermid and R. W. Irving. Popular matchings: Structure and algorithms. *Journal of Combinatorial Optimization*, 22(3):339–358, 2011.
- J. Mestre. Weighted popular matchings. In *Encyclopedia of Algorithms*, pages 1023–1024. Springer, 2008.
- J. Mestre. Weighted popular matchings. *ACM Transactions on Algorithms*, 10(1):2:1–2:16, 2014.
- D. Michail. Reducing rank-maximal to maximum weight matching. *Theoretical Computer Science*, 389(1-2):125–132, 2007.
- M. Nasre. Popular matchings: Structure and cheating strategies. In *Proceedings of the 30th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 412–423. Citeseer, 2013.
- K. Paluch. Popular and clan-popular b -matchings. *Theoretical Computer Science*, 544:3–13, 2014.
- J.-C. Régin. Generalized arc consistency for global cardinality constraint. In *Proceedings of the 13th AAAI Conference on Artificial Intelligence*, pages 209–215. AAAI Press, 1996.
- R. L. Rivest and E. Shen. An optimal single-winner preferential voting system based on game theory. In *Proceedings of the 3rd International Workshop on Computational Social Choice (COMSOC)*, pages 399–410. University of Düsseldorf, 2010.

-
- 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.
- P. B. Simpson. On defining areas of voter choice: Professor Tullock on stable voting. *The Quarterly Journal of Economics*, pages 478–490, 1969.
- C. T. S. Sng. *Efficient Algorithms for Bipartite Matching Problems with Preferences*. PhD thesis, University of Glasgow, Department of Computing Science, 2008.
- C. T. S. Sng and D. F. Manlove. Popular matchings in the weighted capacitated house allocation problem. *Journal of Discrete Algorithms*, 8(2):102–116, 2010.