Publications of Agnes Cseh

Hasso Plattner Institut
Digital Engineering · Universität Potsdam

This document lists all peer-reviewed publications of Agnes Cseh, Chair for Algorithm Engineering, Hasso Plattner Institute, Potsdam, Germany. This listing was automatically generated on February 28, 2024. An up-to-date version is available online at hpi.de/friedrich/docs/publist/cseh.pdf.

Journal articles

[1] Sauer, P., Cseh, Á., Lenzner, P., Improving ranking quality and fairness in Swiss-system chess tournaments. In: *Journal of Quantitative Analysis in Sports*, 2024.

[2] Cseh, Á., Juhos, A., Pairwise Preferences in the Stable Marriage Problem. In: *ACM Transactions on Economics and Computation (TEAC)* 9, pp. 1–28, 2021.

We study the classical, two-sided stable marriage problem under pairwise preferences. In the most general setting, agents are allowed to express their preferences as comparisons of any two of their edges and they also have the right to declare a draw or even withdraw from such a comparison. This freedom is then gradually restricted as we specify six stages of orderedness in the preferences, ending with the classical case of strictly ordered lists. We study all cases occurring when combining the three known notions of stability—weak, strong and super-stability—under the assumption that each side of the bipartite market obtains one of the six degrees of orderedness. By designing three polynomial algorithms and two NP-completeness proofs we determine the complexity of all cases not yet known, and thus give an exact boundary in terms of preference structure between tractable and intractable cases.

[3] Cseh, Á., Kavitha, T., Popular matchings in complete graphs. In: *Algorithmica*, pp. 1–31, 2021.

Our input is a complete graph $G = (V, E)$ on n vertices where each vertex has a strict ranking of all other vertices in G. The goal is to construct a matching in G that is "globally stable" or popular. A matching M is popular if M does not lose a head-to-head election against any matching M': here each vertex casts a vote for the matching in $\{M, M'\}$ where it gets a better assignment. Popular matchings need not exist in the given instance G and the popular matching problem is to decide whether one exists or not. The popular matching problem in G is easy to solve for odd n. Surprisingly, the problem becomes NP-hard for even n, as we show here.

[4] Andersson, T., Cseh, Á., Ehlers, L., Erlanson, A., Organizing time exchanges: Lessons from matching markets. In: *American Economic Journal: Microeconomics* 13, pp. 338–73, 2021.

This paper considers time exchanges via a common platform (e.g. , markets for exchanging time units, positions at education institutions, and tuition waivers). There are several problems associated with such markets, e.g., imbalanced outcomes, coordination problems, and inefficiencies. We model time exchanges as matching markets and construct a non-manipulable mec hanism that selects an individually rational and balanced allocation that maximizes exchanges among the participating agents (and those allocations are efficient). This mechanism works on a preference domain whereby agents classify the goods provided by other participating agents as either unacceptable or acceptable, and for goods classified as acceptable, agents have specific upper quotas representing their maximum needs.

[5] Cseh, Á., Fleiner, T., The Complexity of Cake Cutting with Unequal Shares. In: *ACM Transactions on Algorithms* 16, pp. 1–21, 2020.

An unceasing problem of our prevailing society is the fair division of goods. The problem of proportional cake cutting focuses on dividing a heterogeneous and divisible resource, the cake, among n players who value pieces according to their own measure function. The goal is to assign each player a not necessarily connected part of the cake that the player evaluates at least as much as her proportional share. In this article, we investigate the problem of proportional division with unequal shares, where each player is entitled to receive a predetermined portion of the cake. Our main contribution is threefold. First we present a protocol for integer demands, which delivers a proportional solution in fewer queries than all known proto- cols. By giving a matching lower bound, we then show that our protocol is asymptotically the fastest possible. Finally, we turn to irrational demands and solve the proportional cake cutting problem by reducing it to the same problem with integer demands only. All results remain valid in a highly general cake cutting model, which can be of independent interest.

[6] Cseh, Á., Heeger, K., The stable marriage problem with ties and restricted edges. In: *Discrete Optimization* 36, pp. 100571, 2020.

In the stable marriage problem, a set of men and a set of women are given, each of whom has a strictly ordered preference list over the acceptable agents in the opposite class. A matching is called stable if it is not blocked by any pair of agents, who mutually prefer each other to their respective partner. Ties in the preferences allow for three different definitions for a stable matching: weak, strong and super-stability. Besides this, acceptable pairs in the instance can be restricted in their ability of blocking a matching or being part of it, which again generates three categories of restrictions on acceptable pairs. Forced pairs must be in a stable matching, forbidden pairs must not appear in it, and lastly, free pairs cannot block any matching. Our computational complexity study targets the existence of a stable solution for each of the three stability definitions, in the presence of each of the three types of restricted pairs. We solve all cases that were still open. As a byproduct, we also derive that the maximum size weakly stable matching problem is hard even in very dense graphs, which may be of independent interest.

[7] Cseh, Á., Irving, R. W., Manlove, D. F., The Stable Roommates Problem with Short Lists. In: *Theory of Computing Systems*, pp. 128–149, 2019.

We consider two variants of the classical Stable Roommates problem with Incomplete (but strictly ordered) preference lists (sri) that are degree constrained, i.e., preference lists are of bounded length. The first variant, egald-sri, involves finding an egalitarian stable matching in solvable instances of sri with preference lists of length at most d. We show that this problem is NP-hard even if $d = 3$. On the positive side we give a $\frac{2d+3}{7}$-approximation algorithm for $d \in \{3, 4, 5\}$ which improves on the known bound of 2 for the unbounded preference list case. In the second variant of sri, called d-srti, preference lists can include ties and are of length at most d. We show that the problem of deciding whether an instance of d-srti admits a stable matching is NP-complete even if $d = 3$. We also consider the "most stable" version of this problem and prove a strong inapproximability bound for the $d = 3$ case. However for $d = 2$ we show that the latter problem can be solved in polynomial time.

[8] Cseh, Á., Skutella, M., Paths to stable allocations. In: *International Journal of Game Theory*, pp. 835–862, 2019.

The stable allocation problem is one of the broadest extensions of the well-known stable marriage problem. In an allocation problem, edges of a bipartite graph have capacities and vertices have quotas to fill. Here we investigate the case of uncoordinated processes in stable allocation instances. In this setting, a feasible allocation is given and the aim is to reach a stable allocation by raising the value of the allocation along blocking edges and reducing it on worse edges if needed. Do such myopic changes lead to a stable solution? In our present work, we analyze both better and best response dynamics from an algorithmic point of view. With the help of two deterministic algorithms we show that random procedures reach a stable solution with probability one for all rational input data in both cases. Surprisingly, while there is a polynomial path to stability when better response strategies are played (even for irrational input data), the more intuitive best response steps may require exponential time. We also study the special case of correlated markets. There, random best response strategies lead to a stable allocation in expected polynomial time.

[9] Cechlárová, K., Cseh, Á., Manlove, D., Selected open problems in Matching Under Preferences. In: *Bulletin of the European Association for Theoretical Computer Science*, pp. 14–38, 2019.

The House Allocation problem, the Stable Marriage problem and the Stable Roommates problem are three fundamental problems in the area of matching under preferences. These problems have been studied for decades under a range of optimality criteria, but despite much progress, some challenging questions remain open. The purpose of this article is to present a range of key open questions for each of these problems, which will hopefully stimulate further research activity in this area.

[10] Cseh, Á., Matuschke, J., New and Simple Algorithms for Stable Flow Problems. In: *Algorithmica*, pp. 2557–2591, 2019.

Stable flows generalize the well-known concept of stable matchings to markets in which transactions may involve several agents, forwarding flow from one to another. An instance of the problem consists of a capacitated directed network in which vertices express their preferences over their incident edges. A network flow is stable if there is no group of vertices that all could benefit from rerouting the flow along a walk. Fleiner (Algorithms 7:1-14, 2014) established that a stable flow always exists by reducing it to the stable allocation problem. We present an augmenting path algorithm for computing a stable flow, the first algorithm that achieves polynomial running time for this problem without using stable allocations as a black-box subroutine. We further consider the problem of finding a stable flow such that the flow value on every edge is within a given interval. For this problem, we present an elegant graph transformation and based on this, we devise a simple and fast algorithm, which also can be used to find a solution to the stable marriage problem with forced and forbidden edges. Finally, we study the stable multicommodity flow model introduced by Király and Pap (Algorithms 6:161-168, 2013). The original model is highly involved and allows for commodity-dependent preference lists at the vertices and commodity-specific edge capacities. We present several graph-based reductions that show equivalence to a significantly simpler model. We further show that it is NP-complete to decide whether an integral solution exists.

[11] Arulselvan, A., Cseh, Á., Groß, M., Manlove, D. F., Matuschke, J., Matchings with Lower Quotas: Algorithms and Complexity. In: *Algorithmica*, pp. 185–208, 2018.

We study a natural generalization of the maximum weight many-to-one matching problem. We are given an undirected bipartite graph $G = (A \cup P, E)$ with weights on the edges in E, and with lower and upper quotas on the vertices in P. We seek a maximum weight many-to-one matching satisfying two sets of constraints: vertices in A are incident to at most one matching edge, while vertices in P are either unmatched or they are incident to a number of matching edges between their lower and upper quota. This problem, which we call maximum weight many-to-one matching with lower and upper quotas (WMLQ), has applications to the assignment of students to projects within university courses, where there are constraints on the minimum and maximum numbers of students that must be assigned to each project. In this paper, we provide a comprehensive analysis of the complexity of WMLQ from the viewpoints of classical polynomial time algorithms, fixed-parameter tractability, as well as approximability. We draw the line between NP-hard and polynomially tractable instances in terms of degree and quota constraints and provide efficient algorithms to solve the tractable ones. We further show that the problem can be solved in polynomial time for instances with bounded treewidth; however, the corresponding runtime is exponential in the treewidth with the maximum upper quota $u_{\max}$ as basis, and we prove that this dependence is necessary unless FPT=W[1]. The approximability of WMLQ is also discussed: we present an approximation algorithm for the general case with performance guarantee $u_{\max} + 1$, which is asymptotically best possible unless $P = NP$. Finally, we elaborate on how most of our positive results carry over to matchings in arbitrary graphs with lower quotas.

[12] Cseh, Á., Huang, C.-C., Kavitha, T., Popular Matchings with Two-Sided Preferences and One-Sided Ties. In: *SIAM Journal on Discrete Mathematics*, pp. 367–379, 2017.

We are given a bipartite graph $G = (A \cup B, E)$ where each vertex has a preference list ranking its neighbors: In particular, every $a \in A$ ranks its neighbors in a strict order of preference, whereas the preference list of any $b \in B$ may contain ties. A matching M is popular if there is no matching M' such that the number of vertices that prefer M' to M exceeds the number of vertices that prefer M to M'. We show that the problem of deciding whether G admits a popular matching or not is NP-hard. This is the case even when every $b \in B$ either has a strict preference list or puts all its neighbors into a single tie. In contrast, we show that the problem becomes polynomially solvable in the case when each $b \in B$ puts all its neighbors into a single tie. That is, all neighbors of b are tied in b's list and b desires to be matched to any of them. Our main result is an $\mathcal{O}(n^2)$ algorithm (where $n = |A \cup B|$) for the popular matching problem in this model. Note that this model is quite different from the model where vertices in B have no preferences and do not care whether they are matched or not.

[13] Cseh, Á., Dean, B. C., Improved algorithmic results for unsplittable stable allocation problems. In: *Journal of Combinatorial Optimization*, pp. 657–671, 2016.

The stable allocation problem is a many-to-many generalization of the well-known stable marriage problem, where we seek a bipartite assignment between, say, jobs (of varying sizes) and machines (of varying capacities) that is "stable" based on a set of underlying preference lists submitted by the jobs and machines. Building on the initial work of Dean et al. (The unsplittable stable marriage problem, 2006), we study a natural "unsplittable" variant of this problem, where each assigned job must be fully assigned to a single machine. Such unsplittable bipartite assignment problems generally tend to be NP-hard, including previously-proposed variants of the unsplittable stable allocation problem (McDermid and Manlove in J Comb Optim 19(3): 279–303, 2010). Our main result is to show that under an alternative model of stability, the unsplittable stable allocation problem becomes solvable in polynomial time; although this model is less likely to admit feasible solutions than the model proposed in McDermid and Manlove (J Comb Optim 19(3): 279–303, McDermid and Manlove 2010), we show that in the event there is no feasible solution, our approach computes a solution of minimal total congestion (overfilling of all machines collectively beyond their capacities). We also describe a technique for rounding the solution of a stable allocation problem to produce "relaxed" unsplit solutions that are only mildly infeasible, where each machine is overcongested by at most a single job.

[14] Cseh, Á. Marriages are made in calculation. In: *Bulletin of the European Association for Theoretical Computer Science*, pp. 180–183, 2016.

Butterflies in the stomach, racing heart and sweaty hands. . . is he really the one? The mathematician simply sits down to the computer and calculates it quickly. On the side she makes important discoveries that come handy in various other fields of life, such as job applications.

[15] Cseh, Á., Matuschke, J., Skutella, M., Stable Flows over Time. In: *Algorithms*, pp. 532–545, 2013.

In this paper, the notion of stability is extended to network flows over time. As a useful device in our proofs, we present an elegant preflow-push variant of the Gale-Shapley algorithm that operates directly on the given network and computes stable flows in pseudo-polynomial time, both in the static flow and the flow over time case. We show periodical properties of stable flows over time on networks with an infinite time horizon. Finally, we discuss the influence of storage at vertices, with different results depending on the priority of the corresponding holdover edges.

Conference papers

[16] Cseh, Á., Führlich, P., Lenzner, P., The Swiss Gambit. In: *Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2023.

In each round of a Swiss-system tournament, players of similar score are paired against each other. An intentional early loss therefore might lead to weaker opponents in later rounds and thus to a better final tournament result a phenomenon known as the Swiss Gambit. To the best of our knowledge it is an open question whether this strategy can actually work. This paper provides answers based on an empirical agent-based analysis for the most prominent application area of the Swiss-system format, namely chess tournaments. We simulate realistic tournaments by employing the official FIDE pairing system for computing the player pairings in each round. We show that even though gambits are widely possible in Swiss-system chess tournaments, profiting from them requires a high degree of predictability of match results. Moreover, even if a Swiss Gambit succeeds, the obtained improvement in the final ranking is limited. Our experiments prove that counting on a Swiss Gambit is indeed a lot more of a risky gambit than a reliable strategy to improve the final rank.

[17] Führlich, P., Cseh, Á., Lenzner, P., Improving Ranking Quality and Fairness in Swiss-System Chess Tournaments. In: *ACM Conference on Economics and Computation (EC)*, pp. 1101–1102, 2022.

The International Chess Federation (FIDE) imposes a voluminous and complex set of player pairing criteria in Swiss-system chess tournaments and endorses computer programs that are able to calculate the prescribed pairings. The purpose of these formalities is to ensure that players are paired fairly during the tournament and that the final ranking corresponds to the players' true strength order. We contest the official FIDE player pairing routine by presenting alternative pairing rules. These can be enforced by computing maximum weight matchings in a carefully designed graph. We demonstrate by extensive experiments that a tournament format using our mechanism (1) yields fairer pairings in the rounds of the tournament and (2) produces a final ranking that reflects the players' true strengths better than the state-of-the-art FIDE pairing system.

[18] Cseh, Á., Peters, J., Three-Dimensional Popular Matching with Cyclic Preferences. In: *Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 77–87, 2022.

Two actively researched problem settings in matchings under preferences are popular matchings and the three-dimensional stable matching problem with cyclic preferences. In this paper, we apply the optimality notion of the first topic to the input characteristics of the second one. We investigate the connection between stability, popularity, and their strict variants, strong stability and strong popularity in three-dimensional instances with cyclic preferences. Furthermore, we also derive results on the complexity of these problems when the preferences are derived from master lists.

[19] Cseh, Á., Friedrich, T., Peters, J., Pareto Optimal and Popular House Allocation with Lower and Upper Quotas. In: *Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 300–308, 2022.

In the house allocation problem with lower and upper quotas, we are given a set of applicants and a set of projects. Each applicant has a strictly ordered preference list over the projects, while the projects are equipped with a lower and an upper quota. A feasible matching assigns the applicants to the projects in such a way that a project is either matched to no applicant or to a number of applicants between its lower and upper quota. In this model we study two classic optimality concepts: Pareto optimality and popularity. We show that finding a popular matching is hard even if the maximum lower quota is 2 and that finding a perfect Pareto optimal matching, verifying Pareto optimality, and verifying popularity are all NP-complete even if the maximum lower quota is 3. We complement the last three negative results by showing that the problems become polynomial-time solvable when the maximum lower quota is 2, thereby answering two open questions of Cechlárová and Fleiner. Finally, we also study the parameterized complexity of all four mentioned problems.

[20] Kraiczy, S., Cseh, Á., Manlove, D., On Weakly and Strongly Popular Rankings. In: *Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 1563–1565, 2021.

Van Zuylen et al. introduced the notion of a popular ranking in a voting context, where each voter submits a strictly-ordered list of all candidates. A popular ranking pi of the candidates is at least as good as any other ranking sigma in the following sense: if we compare pi to sigma, at least half of all voters will always weakly prefer pi. Whether a voter prefers one ranking to another is calculated based on the Kendall distance. A more traditional definition of popularity—as applied to popular matchings, a well-established topic in computational social choice—is stricter, because it requires at least half of the voters who are not indifferent between pi and sigma to prefer pi. In this paper, we derive structural and algorithmic results in both settings, also improving upon the results by van Zylen et al. We also point out connections to the famous open problem of finding a Kemeny consensus with 3 voters.

[21] Aziz, H., Chan, H., Cseh, Á., Li, B., Ramezani, F., Wang, C., Multi-Robot Task Allocation—Complexity and Approximation. In: *Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 133–141, 2021.

Multi-robot task allocation is one of the most fundamental classes of problems in robotics and is crucial for various real-world robotic applications such as search, rescue and area exploration. We consider the Single-Task robots and Multi-Robot tasks Instantaneous Assignment (ST-MR-IA) setting where each task requires at least a certain number of robots and each robot can work on at most one task and incurs an operational cost for each task. Our aim is to consider a natural computational problem of allocating robots to complete the maximum number of tasks subject to budget constraints. We consider budget constraints of three different kinds: (1) total budget, (2) task budget, and (3) robot budget. We provide a detailed complexity analysis including results on approximations as well as polynomial-time algorithms for the general setting and important restricted settings.

[22] Aziz, H., Cseh, A., Dickerson, J., McElfresh, D., Optimal Kidney Exchange with Immunosuppressants. In: *Conference on Artificial Intelligence (AAAI)*, pp. 21–29, 2021.

Algorithms for exchange of kidneys is one of the key successful applications in market design, artificial intelligence, and operations research. Potent immunosuppressant drugs suppress the body's ability to reject a transplanted organ up to the point that a transplant across blood- or tissue-type incompatibility becomes possible. In contrast to the standard kidney exchange problem, we consider a setting that also involves the decision about which recipients receive from the limited supply of immunosuppressants that make them compatible with originally incompatible kidneys. We firstly present a general computational framework to model this problem. Our main contribution is a range of efficient algorithms that provide flexibility in terms of meeting meaningful objectives. Motivated by the current reality of kidney exchanges using sophisticated mathematical-programming-based clearing algorithms, we then present a general but scalable approach to optimal clearing with immunosuppression; we validate our approach on realistic data from a large fielded exchange.

[23] Cseh, Á., Juhos, A., Pairwise Preferences in the Stable Marriage Problem. In: *Symposium Theoretical Aspects of Computer Science (STACS)*, pp. 21:1–21:16, 2019.

We study the classical, two-sided stable marriage problem under pairwise preferences. In the most general setting, agents are allowed to express their preferences as comparisons of any two of their edges and they also have the right to declare a draw or even withdraw from such a comparison. This freedom is then gradually restricted as we specify six stages of orderedness in the preferences, ending with the classical case of strictly ordered lists. We study all cases occurring when combining the three known notions of stability - weak, strong and super-stability - under the assumption that each side of the bipartite market obtains one of the six degrees of orderedness. By designing three polynomial algorithms and two NP-completeness proofs we determine the complexity of all cases not yet known, and thus give an exact boundary in terms of preference structure between tractable and intractable cases.

[24] Cseh, Á., Kavitha, T., Popular Matchings in Complete Graphs. In: *Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, pp. 17:1–17:14, 2018.

Our input is a complete graph $G = (V, E)$ on n vertices where each vertex has a strict ranking of all other vertices in G. The goal is to construct a matching in G that is "globally stable" or popular. A matching M is popular if M does not lose a head-to-head election against any matching M': here each vertex casts a vote for the matching in $\{M, M'\}$ where it gets a better assignment. Popular matchings need not exist in the given instance G and the popular matching problem is to decide whether one exists or not. The popular matching problem in G is easy to solve for odd n. Surprisingly, the problem becomes NP-hard for even n, as we show here.

[25] Cseh, Á., Fleiner, T., The Complexity of Cake Cutting with Unequal Shares. In: *Symposium Algorithmic Game Theory (SAGT)*, pp. 19–30, 2018.

An unceasing problem of our prevailing society is the fair division of goods. The problem of proportional cake cutting focuses on dividing a heterogeneous and divisible resource, the cake, among n players who value pieces according to their own measure function. The goal is to assign each player a not necessarily connected part of the cake that the player evaluates at least as much as her proportional share. In this paper, we investigate the problem of proportional division with unequal shares, where each player is entitled to receive a predetermined portion of the cake. Our main contribution is threefold. First we present a protocol for integer demands that delivers a proportional solution in fewer queries than all known algorithms. Then we show that our protocol is asymptotically the fastest possible by giving a matching lower bound. Finally, we turn to irrational demands and solve the proportional cake cutting problem by reducing it to the same problem with integer demands only. All results remain valid in a highly general cake cutting model, which can be of independent interest.

[26] Cseh, Á., Matuschke, J., New and Simple Algorithms for Stable Flow Problems. In: *Workshop Graph-Theoretic Concepts in Computer Science (WG)*, pp. 206–219, 2017.

Stable flows generalize the well-known concept of stable matchings to markets in which transactions may involve several agents, forwarding flow from one to another. An instance of the problem consists of a capacitated directed network, in which vertices express their preferences over their incident edges. A network flow is stable if there is no group of vertices that all could benefit from rerouting the flow along a walk. Fleiner [13] established that a stable flow always exists by reducing it to the stable allocation problem. We present an augmenting-path algorithm for computing a stable flow, the first algorithm that achieves polynomial running time for this problem without using stable allocation as a black-box subroutine. We further consider the problem of finding a stable flow such that the flow value on every edge is within a given interval. For this problem, we present an elegant graph transformation and based on this, we devise a simple and fast algorithm, which also can be used to find a solution to the stable marriage problem with forced and forbidden edges. Finally, we study the highly complex stable multicommodity flow model by Király and Pap [24]. We present several graph-based reductions that show equivalence to a significantly simpler model. We further show that it is NP-complete to decide whether an integral solution exists.

[27] Cseh, Á., Kavitha, T., Popular Edges and Dominant Matchings. In: *International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pp. 138–151, 2016.

Given a bipartite graph $G = (A \cup B, E)$ with strict preference lists and given an edge $e^* \in E$, we ask if there exists a popular matching in G that contains e^*. We call this the popular edge problem. A matching M is popular if there is no matching M' such that the vertices that prefer M' to M outnumber those that prefer M to M'. It is known that every stable matching is popular; however G may have no stable matching with the edge e^*. In this paper we identify another natural subclass of popular matchings called "dominant matchings" and show that if there is a popular matching that contains the edge e^*, then there is either a stable matching that contains e^* or a dominant matching that contains e^*. This allows us to design a linear time algorithm for the popular edge problem. When preference lists are complete, we show an $\mathcal{O}(n^3)$ algorithm to find a popular matching containing a given set of edges or report that none exists, where $n = |A| + |B|$.

[28] Cseh, Á., Irving, R. W., Manlove, D. F., The Stable Roommates Problem with Short Lists. In: *Symposium Algorithmic Game Theory (SAGT)*, pp. 207–219, 2016.

We consider two variants of the classical Stable Roommates problem with Incomplete (but strictly ordered) preference lists (sri) that are degree constrained, i.e., preference lists are of bounded length. The first variant, egal d-sri, involves finding an egalitarian stable matching in solvable instances of sri with preference lists of length at most d. We show that this problem is NP-hard even if $d = 3$. On the positive side we give a $\frac{2d+3}{7}$-approximation algorithm for $d \in \{3, 4, 5\}$ which improves on the known bound of 2 for the unbounded preference list case. In the second variant of sri, called d-srti, preference lists can include ties and are of length at most d. We show that the problem of deciding whether an instance of d-srti admits a stable matching is NP-complete even if $d = 3$. We also consider the "most stable" version of this problem and prove a strong inapproximability bound for the $d = 3$ case. However for $d = 2$ we show that the latter problem can be solved in polynomial time.

[29] Cseh, Á., Huang, C.-C., Kavitha, T., Popular Matchings with Two-Sided Preferences and One-Sided Ties. In: *International Colloquium on Automata, Languages and Programming (ICALP)*, pp. 367–379, 2015.

We are given a bipartite graph $G = (A \cup B, E)$ where each vertex has a preference list ranking its neighbors: in particular, every $a \in A$ ranks its neighbors in a strict order of preference, whereas the preference lists of $b \in B$ may contain ties. A matching M is popular if there is no matching M' such that the number of vertices that prefer M' to M exceeds the number that prefer M to M'. We show that the problem of deciding whether G admits a popular matching or not is NP-hard. This is the case even when every $b \in B$ either has a strict preference list or puts all its neighbors into a single tie. In contrast, we show that the problem becomes polynomially solvable in the case when each $b \in B$ puts all its neighbors into a single tie. That is, all neighbors of b are tied in $b's$ list and and b desires to be matched to any of them. Our main result is an $\mathcal{O}(n^2)$ algorithm (where $n = |A \cup B|$) for the popular matching problem in this model. Note that this model is quite different from the model where vertices in B have no preferences and do not care whether they are matched or not.

[30] Arulselvan, A., Cseh, Á., Groß, M., Manlove, D. F., Matuschke, J., Many-to-one Matchings with Lower Quotas: Algorithms and Complexity. In: *International Symposium Algorithms and Computation (ISAAC)*, pp. 176–187, 2015.

We study a natural generalization of the maximum weight many-to-one matching problem. We are given an undirected bipartite graph $G = (A \cup P, E)$ with weights on the edges in E, and with lower and upper quotas on the vertices in P. We seek a maximum weight many-to-one matching satisfying two sets of constraints: vertices in A are incident to at most one matching edge, while vertices in P are either unmatched or they are incident to a number of matching edges between their lower and upper quota. This problem, which we call maximum weight many-to-one matching with lower and upper quotas (wmlq), has applications to the assignment of students to projects within university courses, where there are constraints on the minimum and maximum numbers of students that must be assigned to each project. In this paper, we provide a comprehensive analysis of the complexity of wmlq from the viewpoints of classic polynomial time algorithms, fixed-parameter tractability, as well as approximability. We draw the line between NP -hard and polynomially tractable instances in terms of degree and quota constraints and provide efficient algorithms to solve the tractable ones. We further show that the problem can be solved in polynomial time for instances with bounded treewidth; however, the corresponding runtime is exponential in the treewidth with the maximum upper quota u_max as basis, and we prove that this dependence is necessary unless FPT=W[1]. Finally, we also present an approximation algorithm for the general case with performance guarantee $u_max + 1$, which is asymptotically best possible unless $P = NP$.

[31] Cseh, Á., Manlove, D. F., Stable Marriage and Roommates Problems with Restricted Edges: Complexity and Approximability. In: *Symposium Algorithmic Game Theory (SAGT)*, pp. 15–26, 2015.

In the stable marriage and roommates problems, a set of agents is given, each of them having a strictly ordered preference list over some or all of the other agents. A matching is a set of disjoint pairs of mutually acceptable agents. If any two agents mutually prefer each other to their partner, then they block the matching, otherwise, the matching is said to be stable. We investigate the complexity of finding a solution satisfying additional constraints on restricted pairs of agents. Restricted pairs can be either forced or forbidden. A stable solution must contain all of the forced pairs, while it must contain none of the forbidden pairs. Dias et al. [5] gave a polynomial-time algorithm to decide whether such a solution exists in the presence of restricted edges. If the answer is no, one might look for a solution close to optimal. Since optimality in this context means that the matching is stable and satisfies all constraints on restricted pairs, there are two ways of relaxing the constraints by permitting a solution to: (1) be blocked by as few as possible pairs, or (2) violate as few as possible constraints on restricted pairs. Our main theorems prove that for the (bipartite) stable marriage problem, case (1) leads to NP-hardness and inapproximability results, whilst case (2) can be solved in polynomial time. For non-bipartite stable roommates instances, case (2) yields an NP-hard but (under some cardinality assumptions) 2-approximable problem. In the case of NP-hard problems, we also discuss polynomially solvable special cases, arising from restrictions on the lengths of the preference lists, or upper bounds on the numbers of restricted pairs.

[32] Cseh, Á., Skutella, M., Paths to Stable Allocations. In: *Symposium Algorithmic Game Theory (SAGT)*, pp. 61–73, 2014.

The stable allocation problem is one of the broadest extensions of the well-known stable marriage problem. In an allocation problem, edges of a bipartite graph have capacities and vertices have quotas to fill. Here we investigate the case of uncoordinated processes in stable allocation instances. In this setting, a feasible allocation is given and the aim is to reach a stable allocation by raising the value of the allocation along blocking edges and reducing it on worse edges if needed. Do such myopic changes lead to a stable solution? In our present work, we analyze both better and best response dynamics from an algorithmic point of view. With the help of two

deterministic algorithms we show that random procedures reach a stable solution with probability one for all rational input data in both cases. Surprisingly, while there is a polynomial path to stability when better response strategies are played (even for irrational input data), the more intuitive best response steps may require exponential time. We also study the special case of correlated markets. There, random best response strategies lead to a stable allocation in expected polynomial time.