Maintaining Near-Popular Matchings*

Sayan Bhattacharya[1], Martin Hoefer[2], Chien-Chung Huang[3], Telikepalli Kavitha[4], and Lisa Wagner[5]

[1] Institute of Mathematical Sciences, Chennai, India
bsayan@imsc.res.in
[2] MPI für Informatik and Saarland University, Germany
mhoefer@mpi-inf.mpg.de
[3] Dept. Computer Science and Engineering, Chalmers University, Sweden.
villars@mpi-inf.mpg.de
[4] Tata Institute of Fundamental Research, Mumbai, India
kavitha@tcs.tifr.res.in
[5] Dept. of Computer Science, RWTH Aachen University, Germany
lwagner@rwth-aachen.de

Abstract. We study dynamic matching problems in graphs among agents with preferences. Agents and/or edges of the graph arrive and depart iteratively over time. The goal is to maintain matchings that are favorable to the agent population and stable over time. More formally, we strive to keep a small unpopularity factor by making only a small amortized number of changes to the matching per round. Our main result is an algorithm to maintain matchings with unpopularity factor $(\Delta + k)$ by making an amortized number of $O(\Delta + \Delta^2/k)$ changes per round, for any $k > 0$. Here Δ denotes the maximum degree of any agent in any round. We complement this result by a variety of lower bounds indicating that matchings with smaller factor do not exist or cannot be maintained using our algorithm.

As a byproduct, we obtain several additional results that might be of independent interest. First, our algorithm implies existence of matchings with small unpopularity factors in graphs with bounded degree. Second, given any matching M and any value $\alpha \geq 1$, we provide an efficient algorithm to compute a matching M' with unpopularity factor α over M if it exists. Finally, our results show the absence of voting paths in two-sided instances, even if we restrict to sequences of matchings with larger unpopularity factors (below Δ).

1 Introduction

Matching arises as a fundamental task in many coordination, resource allocation, and network design problems. In many domains, matching and allocation problems occur among agents with preferences, e.g., in job markets, when assigning residents to hospitals, or students to dormitory rooms, or when allocating resources in distributed systems. There are a number of approaches for formal

* Supported by DFG Cluster of Excellence MMCI at Saarland University

study of allocation under preferences, the most prominent being stable and popular matchings. Usually, there is a set of agents embedded into a graph, and each agent has a preference list over his neighbors. An edge is called a blocking pair if both agents strictly prefer each other to their current partners (if any). A matching without blocking pair is a *stable matching*. In a popular matching all agents get to vote between two matchings M and M'. They vote for M if it yields a partner which is strictly preferred to the one in M', or vice versa (they don't vote if neither of them is strictly preferred). The matching that receives more votes is more popular. For a *popular matching* there exists no other matching that is more popular.

Stable and (to a lesser extent) popular matchings have been studied intensively in algorithms, economics, operations research, and game theory, but mostly under the assumption that the set of agents and the set of possible matching edges remain static. In contrast, many application areas above are inherently dynamic. For example, in a large firm new jobs open up on a repeated basis, e.g., due to expansion into new markets, retirement of workers, or the end of fixed-term contracts. Similarly, new applicants from outside arrive, or internal employees seek to get promoted or move into a different department. The firm strives to fill its positions with employees in a way that is preferable to both firm and workers. The naive approach would be to compute, e.g., a stable or popular matching from scratch every time a change happens, but then employees might get assigned differently every time. Instead, the obvious goal is to maintain a stable or popular assignment at a small rate of change. Similar problems arise also in the context of dormitory room assignment or resource allocation in distributed systems. Perhaps surprisingly, these natural problems have not been studied in the literature so far.

Maintaining graph-theoretic solution concepts like matchings or shortest paths is an active research area in algorithms. In these works, the objective is to maintain matchings of maximum cardinality while making a small number of changes. These approaches are unsuitable for systems with agent preferences, which fundamentally change the nature and the characteristics of the problem.

More fundamentally, a central theme in algorithmic game theory is to study dynamics in games such as best response or no-regret learning. However, in the overwhelming majority of these works, the games themselves (agents, strategies, payoffs) are static over time, and the interest is to characterize the evolution of strategic interaction. In contrast, there are many games in which maintaining stability concepts at a small rate of change is a natural objective, such as in routing or scheduling problems. To the best of our knowledge, our paper is the first to study algorithms for maintaining equilibria in the prominent domain of matching and network design problems.

Model and Notation. Before we state our results, let us formally introduce the model and notation. We consider a dynamic round-based matching scenario for a set V of agents. In each round t, there exists a graph $G^t = (V, E^t)$ with set E^t of possible matching edges among the agents. Initially, $E^0 = \emptyset$. For *edge* *dynamics*, in the beginning of each round $t \geq 1$ a single edge is added or deleted, i.e., E^t and E^{t+1} differ in exactly one edge. We denote this edge by e^t. Note that a particular edge e can be added and removed multiple times over time.

For *vertex dynamics*, in the beginning of each round $t \geq 1$ a single vertex arrives or departs along with all incident edges. We denote this vertex by v^t, where the same vertex can arrive and depart multiple times over time. Formally, in vertex dynamics all vertices exist throughout. We color them red and blue depending on whether they are currently present or not, respectively. Then, in the beginning of a round, if v^t arrives, it is colored red and all edges between v^t and red agents arrive. If v^t leaves, it is colored blue and all incident edges are removed. Thus, E^t and E^{t+1} differ by exactly a set of edges from v^t to red agents. Vertex dynamics also model the case when in each round the preference list of one vertex changes. Assume there is a separate vertex with the new preference list and consider two rounds in which the old vertex leaves and the new one arrives. Since our asymptotic bounds do not depend on the overall number of agents or edges, they directly extend to this case.

We consider several structures for the preferences. In the *roommates case* each agent $v \in V$ has a strict preference list $\succ_v$ over all other agents in V. In the *two-sided* case we have sets X and Y and $E^t \subseteq X \times Y$. In the *one-sided* case the elements in X do not have preferences, only agents in Y have preferences over elements in X. Each agent always prefers being matched over being unmatched.

Our goal is to maintain at small amortized cost a matching in each round that satisfies a preference criterion. Towards this end, we study several criteria in this paper. For matching M and agent v we denote by $M(v)$ the agent matched to v in M, where we let $v = M(v)$ when v is unmatched. In round t, an edge $e = (u, v) \in E^t \setminus M$ is called a *blocking pair* for matching $M \subseteq E^t$ if $u \succ_v M(v)$ and $v \succ_u M(u)$. M is a *stable matching* if it has no blocking pair.

For two matchings M and M', v is called a $(+)$-agent if $M'(v) \succ_v M(v)$. We call v a $(-)$-agent if $M(v) \succ_v M'(v)$ and (0)-agent if $M'(v) = M(v)$. We denote by V^+, V^- and V^0 the sets of $(+)$-, $(-)$- and (0)-agents, respectively. For $\alpha \geq 1$, we say M' is *α-more popular* than M if $|V^+| \geq \alpha \cdot |V^-|$. If $|V^+| = |V^-| = 0$, we say M' is *1-more popular* than M, and if $|V^+| > 0 = |V^-|$ then M' is ∞-more popular than M. In round t, the *unpopularity factor* $\rho(M) \in [1, \infty) \cup \{\infty\}$ of matching $M \subseteq E^t$ is the maximum α such that there is an α-more popular matching $M' \subseteq E^t$. M is a *c-unpopular matching* if it has unpopularity factor $\rho(M) \leq c$. A *1-unpopular matching* is called *popular matching*.

Our bounds depend on the *maximum degree* of any agent, where for one-sided instances this includes only the agents in Y. In round t, consider an agent v in G^t. We denote by $N^t(v)$ the set of current neighbors of v, by $d^t(v)$ the degree of v, by Δ^t the maximum degree of any agent. Finally, by $\Delta = \max_t \Delta^t$ we denote the maximum degree of any agent in any of the rounds. Observe that throughout the dynamics, we allow the same edge to arrive and depart multiple times. In addition, an agent v can have a much larger degree than Δ in $\bigcup_t E^t$.

Our Results. We maintain matchings when agents and/or edges of the graph arrive and depart iteratively over time. If every agent has degree at most Δ in every round, our algorithm maintains $O(\Delta)$-unpopular matchings by making an amortized number of $O(\Delta)$ changes to the matching per round. This result holds in one-sided, two-sided and roommates cases. It is almost tight with respect to the unpopularity factor, since there are instances where all matchings have unpopularity factor at least Δ. More formally, if there is one edge arriving or leaving per round, our algorithm yields a tradeoff. Given any number $k > 0$, the algorithm can maintain matchings with unpopularity factor $(\Delta + k)$ using an amortized number of $O(\Delta + \Delta^2/k)$ changes per round. If one vertex arrives or leaves per round, the algorithm needs $O(\Delta^2 + \Delta^3/k)$ changes per round.

The algorithm switches to a matching that is $\alpha > (\Delta + k)$-more popular whenever it exists, and we show that this strategy converges in every round. We can decide for a given matching M and value $\alpha \geq 1$ if there is a matching M' that yields an unpopularity factor at least α for M and compute M' if it exists. Our bounds imply the existence of matchings with small unpopularity factors in one-sided and roommates instances with bounded degree. These insights might be of independent interest.

For two-sided instances, stable and popular matchings exist, but we show that maintaining them requires an amortized number of $\Omega(n)$ changes to the matching per round, even when $\Delta = 2$. In addition, our algorithm cannot be used to maintain matchings with unpopularity factors below $\Delta - 1$. Iterative resolution of matchings with such unpopularity factors might not converge. In fact, we provide an instance and an initial matching from which every sequence of matchings with unpopularity factor greater than 1 leads into a cycle. In contrast to one-sided instances, this implies that two-sided instances might have no voting paths, even for complete and strict preferences. Furthermore, we show that cycling dynamics can evolve even when we restrict to resolution of matchings with higher unpopularity factors (up to Δ).

In summary, our results show that we can maintain a near-popular matching in a dynamic environment with relatively small changes, by pursuing a greedy improvement strategy. For the one-sided case, this achieves essentially the best unpopularity factor we can hope for. In the two-sided case, achieving a better factor with our strategy is bound to fail. Whether there are other strategies with better factors or smaller changes to maintain near-popular matchings is an interesting future direction.

Related Work. Stable matchings have been studied intensively over the last decades, and we refer to standard textbooks in the area for an overview [10,16, 19]. Perhaps closest to our paper are works on reaching stable matchings via iterative resolution of blocking pairs. Knuth [15] provided a cyclic sequence of resolutions in a two-sided instance. Hence, even though stable matchings exist, iterative resolution of blocking pairs might not always lead there. Nevertheless, Roth and Vande Vate [20] showed that there is always some sequence of polynomially many blocking-pair resolutions that leads to a stable matching.

Ackermann et al [3] constructed instances where random sequences require exponential time with high probability. Although in the roommates case (for general graphs) stable matchings might not exist, Diamantoudi et al. [7] showed that there are always sequences of resolutions leading to a stable matching if it exists. Furthermore, the problem has been studied in constrained stable matching problems [11–13]. In contrast, our aim is to maintain matchings by making a small number of changes per round. However, we also show that, perhaps surprisingly, similar sequences do not exist for popular matchings in two-sided instances.

Stable matching turns out to be a very demanding concept that cannot be maintained at small cost. We obtain more positive results for near-popular matchings. The notion of popularity was introduced by Gärdenfors [8] in the two-sided case, who showed that every stable matching is popular when all preference lists are strict. When preference lists admit ties, it was shown by Biró, Irving, and Manlove [6] that the problem of computing an arbitrary popular matching in two-sided instances is NP-hard. They also provide an algorithm to decide if a matching is popular or not in the two-sided and roommates cases.

When agents on only one side have preferences, popular matchings might not exist. Abraham et al. [1] gave a characterization of instances that admit popular matchings; when preference lists are strict, they showed a linear-time algorithm to determine if a popular matching exists and if so, to compute one. Popular matchings in the one-sided case have been well-studied; closest to our paper is Abraham and Kavitha [2] that study the *voting paths* problem. Given an initial matching M_1, the problem is to find a voting path of least length, i.e., a sequence of matchings $M_1, M_2, \ldots, M_k$ of least length such that M_k is popular. In this sequence every M_i must be more popular than M_{i-1}. If a one-sided instance admits a popular matching, then from every M_1 there is always a voting path of length at most 2, and one of least length can be determined in linear time [2].

McCutchen [17] introduced the notion of *unpopularity factor* and showed that the problem of computing a least unpopular matching in one-sided instances is NP-hard. For a roommates instance, popular matchings might not exist. Huang and Kavitha [14] show that with strict preference lists, there is always a matching with unpopularity factor at most $O(\log n)$, and there exist instances where every matching has unpopularity factor $\Omega(\log n)$.

A prominent topic in algorithms is maintaining matchings in dynamic graphs that approximate the maximum cardinality matching. In graphs with n nodes and iterative arrival and departure of edges, Onak and Rubinfeld [18] design a randomized algorithm that maintains a matching which guarantees a large constant approximation factor and requires only $O(\log^2 n)$ amortized update time. Baswana, Gupta and Sen [4] provide a randomized 2-approximation in $O(\log n)$ amortized time. For deterministic algorithms, Gupta and Peng [9] gave a $(1+\varepsilon)$-approximation in $O(\sqrt{m}/\varepsilon^2)$ worst-case update time. Very recently, Bhattacharya et al. [5] showed a deterministic $(4+\varepsilon)$-approximation in $O(m^{1/3}/\varepsilon^2)$ worst-case update time.

Algorithm 1: DeferredResolution

```
for every round t = 1, 2, ... do
    Compute for matching M an α-more popular matching M′ if it exists.
    while M′ exists do
        M ← M′
        Compute for matching M an α-more popular matching M′ if it exists.
```

2 Maintaining $(\Delta + k)$-Unpopular Matchings

In this section, we present an algorithm that, given any number $k > 0$, maintains $(\Delta + k)$-unpopular matchings. Our approach applies in one-sided, two-sided and roommates instances. In the edge-dynamic case, it makes an amortized number of $O(\Delta + \Delta^2/k)$ changes to the matching per round. In every round, our algorithm DeferredResolution iteratively replaces the current matching with an α-more popular matching until no such matching exists (see Algorithm 1). We show in Section 2.1 that such matchings can be computed efficiently. In Section 2.2 we show that when $\alpha > \Delta$ the iterative replacement converges in every round and amortized over all rounds the number of changes made to the matching is at most $O(\Delta + \Delta^2/k)$ per round.

2.1 Finding an α-More Popular Matching

Let us first show that for any given matching M and any value α, we can decide in polynomial time if the unpopularity factor is $\rho(M) \geq \alpha$ and construct an α-more popular matching if it exists. While throughout this paper we assume agents to have strict preferences, this result holds even when the preferences have ties.

Theorem 1. *Let $G = (V, E)$ be a graph, and suppose for every agent $v \in V$ there is a preference order $\succeq_v$ (possibly with ties) over $N(v) \cup \{v\}$ such that $u \succeq_v v$ for all $u \in N(v)$. Then for every matching M in G and every value $\alpha \in \mathbb{R} \cup \{\infty\}$, we can decide in polynomial time if $\rho(M) \geq \alpha$ as well as compute an α-more-popular matching M' if it exists.*

Proof. The general structure of the algorithm is shown as Algorithm 2. The main idea is to construct an adjusted graph and find a maximum-weight matching, which allows to see if an α-more popular matching exists.

We first take a closer look at α. The case $\alpha \leq 1$ is trivial. If $\alpha > |V| - 1$, any α-more popular M' has no $(-)$-agent. So we are checking if $\rho(M) = \infty$ or, equivalently, if $\rho(M) \geq \alpha = |V|$. If $\rho(M) \in (1, |V| - 1]$, it is given as a ratio of two numbers $|V^+|$ and $|V^-|$, which are both integers in $\{1, \ldots, n\}$. Let $\mathbb{Q}_n$ be the set of rational numbers that can be expressed as a fraction of two integers in $\{1, \ldots, n\}$. Thus, when $\alpha \notin \mathbb{Q}_n$, we can equivalently test for $\rho(M) \geq \alpha^*$, where α^* is the smallest number of $\mathbb{Q}_n$ larger than α (see line 3 in the algorithm). Due to

Algorithm 2: Finding an α-more popular matching for M

if $\alpha \leq 1$ **then return** M
else if $\alpha > |V| - 1$ **then** set $\alpha^* \leftarrow |V|$
else set $\alpha^* \leftarrow \min\{r \in \mathbb{Q}_n \mid r \geq \alpha\}$
Set $\alpha' \leftarrow \alpha^* - \epsilon$
Construct $\tilde{G} = (\tilde{V}, \tilde{E})$ as union of two copies (V_1, E_1), (V_2, E_2) of G and edges E_3 between copies, and assign edge weights $w_2(e)$ to every edge $e \in \tilde{E}$
Compute a maximum-weight matching M^* in $\tilde{G}$
if $w_2(M^*) > |V|(2\alpha' + 1)$ **then return** $M^* \cap E_1$
else return $\emptyset$

reasons mentioned below, we replace the test $\rho(M) \geq \alpha^*$ by testing $\rho(M) > \alpha'$, where α' is slightly smaller than α^*, but still larger than the next-smaller number of $\mathbb{Q}_n$. Formally, $\alpha' = \alpha^* - \epsilon$ with

$$\epsilon = \frac{1}{2} \cdot \min_{r, r' \in \mathbb{Q}_n} \{r - r' \mid r - r' > 0\}$$

half of the smallest strictly positive difference between any two numbers in $\mathbb{Q}_n$. Observe that $\rho(M) \geq \alpha$ if and only if $\rho(M) > \alpha'$.

For the test we construct M' via a maximum-weight matching in a graph structure $\tilde{G}$ indicating the gains and losses in popularity. $\tilde{G}$ contains two full copies of G. In addition, for each vertex v in G there is an edge connecting the two copies of v. More formally, $\tilde{G} = (\tilde{V}, \tilde{E})$, $\tilde{V} = V_1 \cup V_2$ and $\tilde{E} = E_1 \cup E_2 \cup E_3$. (V_1, E_1) and (V_2, E_2) constitute two copies of G. E_3 contains for each vertex v in G an edge (v_1, v_2) between its two copies $v_1 \in V_1$ and $v_2 \in V_2$. We define edge weights such that each maximum-weight matching M^* in $\tilde{G}$ is perfect. Then, we construct M' by restricting attention to V_1 and matching the same vertices as M^* within V_1. Vertices of V_1 matched to their copy remain unmatched in M'.

For clarity, we define the edge weights $w_2(e)$ in two steps. We first consider weights w_1 where, intuitively, $w_1(e)$ indicates whether the incident agents become $(+)$-, (0)-, or $(-)$-agents when e is added to M. The value of w_1 is used to charge the $(+)$-agents to the $(-)$-agents. Formally, let $e = (u_i, v_j) \in \tilde{E}$ and set

$$w_1(e) = \begin{cases} 2 & \text{if } v \succ_u M(u) \text{ and } u \succ_v M(v), \\ 1 & \text{if } v \succ_u M(u) \text{ and } u =_v M(v), \text{ or } v =_u M(u) \text{ and } u \succ_v M(v), \\ 0 & \text{if } v =_u M(u) \text{ and } u =_v M(v), \\ 1 - \alpha' & \text{if } v \succ_u M(u) \text{ and } M(v) \succ_v u, \text{ or } M(u) \succ_u v \text{ and } u \succ_v M(v), \\ -\alpha' & \text{if } v =_u M(u) \text{ and } M(v) \succ_v u, \text{ or } M(u) \succ_u v \text{ and } u =_v M(v), \\ -2\alpha' & \text{if } M(u) \succ_u v \text{ and } M(v) \succ_v u \end{cases}$$

We let $w_1(M) = \sum_{e \in M} w_1(e)$.

If there is an α^*-more popular matching M', there is a perfect matching $\tilde{M}$ in $\tilde{G}$ with total weight $w_1(\tilde{M}) > 0$. We simply install M' in both copies (V_1, E_1)

and (V_2, E_2) and match single vertices to their copy. Then, for every (+)-agent in V^+ we add a weight of 2 on the incident edges of $\tilde{M}$. For every (−)-agent in V^- we subtract a weight of $2\alpha'$ on the incident edges of $\tilde{M}$. The contribution of (0)-agents in V^0 to the edge weight is 0. Thus, as $2|V^+| \geq 2\alpha^*|V^-| > 2\alpha'|V^-|$, we get $w_1(\tilde{M}) > 0$. In contrast, an arbitrary matching $\tilde{M}$ with $w_1(\tilde{M}) > 0$ might not be perfect and thus impossible to be transformed into a α^*-more popular matching in G. Towards this end, we change the weights to w_2 with $w_2(e) = w_1(e) + 2\alpha' + 1$ for every $e \in \tilde{E}$. We show that there is an α^*-more popular matching M' if and only if a *maximum-weight matching* M^* for w_2 in $\tilde{G}$ has $w_2(M^*) > |V|(2\alpha'+1)$. The key difference is that $w_2(e) > 0$ for all $e \in \tilde{E}$, and therefore under w_2 every maximum-weight matching is perfect.

More formally, if there is an α^*-more popular matching M', we construct $\tilde{M}$ as above and observe that $w_1(\tilde{M}) > 0$ if and only if $w_2(\tilde{M}) > |V|(2\alpha' + 1)$. For the other direction, we first claim that every maximum-weight matching M^* for w_2 is perfect. Assume first there is some maximum matching M^* where some vertex v remains single. By $M^*(V_1)$ we denote the part of M^* which only uses vertices in V_1. Similarly, $M^*(V_2)$ is the part of M^* which only uses vertices in V_2. W.l.o.g. we assume $w_2(M^*(V_1)) \geq w_2(M^*(V_2))$, and if $w_2(M^*(V_1)) = w_2(M^*(V_2))$ we assume the number of unmatched vertices in V_1 is larger or equal to the number of unmatched vertices in V_2. If $w_2(M^*(V_1)) > w_2(M^*(V_2))$, then M^* could be improved by matching V_2 in the same manner as V_1. Thus, $w_2(M^*(V_1)) = w_2(M^*(V_2))$, and there is at least one single vertex v_1 regarding M^* in V_1. If the corresponding copy $v_2 \in V_2$ is single as well, we can improve M^* by adding (v_1, v_2). If v_2 is matched, we can rearrange the matching on V_2 to mirror the one on V_1 without loss in total weight. Then (v_1, v_2) can be added. Hence, M^* has to be a perfect matching.

Suppose $w_2(M^*) > |V|(2\alpha'+1)$, we construct an α^*-more popular matching as follows. As M^* has maximum-weight for w_2, by the observations above we can assume that $M^*(V_1)$ and $M^*(V_2)$ contain exactly the copies of the same edges of E. Since M^* is perfect, for each $v \in V$ both copies v_1, v_2 are matched. If they are matched via (v_1, v_2), we leave v single in M'. Otherwise, the non-single agents in M' are matched as their copies in $M^*(V_1)$. We claim that $w_2(M^*) > |V|(2\alpha'+1)$ implies M' is α^*-more popular. First, note that $w_2(M^*) > |V|(2\alpha' + 1)$ implies $w_1(M^*) > 0$. Especially, this implies that $|V^+| > 0$. The preference of agent v for M' corresponds to the contribution of $v_1 \in V_1$ to $w_1(M^*)$, i.e., v_1 contributes 1,0, or $-\alpha'$ when $v \in V^+$, V^0, or V^-, respectively. By symmetry of M^* and of edge weights in E_3, the total contribution of vertices in V_1 to $w_1(M^*)$ is exactly $w_1(M^*)/2$. Hence, $w_1(M^*) > 0$ implies $|V^+| > \alpha'|V^-|$ for M'. Here the choice of $\alpha' = \alpha^* - \epsilon$ becomes crucial. By the choice of ϵ we know that the smallest value of $\mathbb{Q}_n$ larger than α' is α^*. Thus, $|V^+| > \alpha'|V^-|$ also implies $|V^+| \geq \alpha^*|V^-|$ which shows $|V^+| \geq \alpha|V^-|$. Hence, $w_2(M^*) > |V|(2\alpha' + 1)$ if and only if an α-more popular matching exists.

We can use the same approach for instances with one-sided preferences by simply defining the preferences of the other side to be indifferent between all potential matching partners as well as being single. □

2.2 Convergence and Amortized Number of Changes

Given that we can decide and find α-more popular matchings efficiently, we now establish that for $\alpha > \Delta$ the iterative resolution does not lead into cycles and makes a small amortized number of changes per round.

Theorem 2. DEFERREDRESOLUTION *maintains a* $(\Delta+k)$*-unpopular matching by making an amortized number of* $O(\Delta + \Delta^2/k)$ *changes to the matching per round with edge dynamics, for any* $k > 0$.

Proof. Our proof is based on the following potential function

$$\Phi^t(M) = \sum_{v \in V} d^t(v) + 1 - rank(M(v)) \enspace ,$$

where $rank(M(v)) = i$ if in the preference list of v restricted to $N^t(v) \cup \{v\}$, partner $M(v)$ ranks at the i^{th} position. Whenever DEFERREDRESOLUTION replaces a matching M in round t with any $(\Delta+k+\epsilon)$-more popular one M' with $\epsilon > 0$, we know that $|V^+| > (\Delta + k)|V^-|$.

Consider the symmetric difference $M' \oplus M = (M \cup M') \setminus (M \cap M')$. Observe that due to strictness of preference lists, we have $v \in V^0$ if and only if $M(v) = M'(v)$. In the two-sided or roommates case this also implies $M(v) \in V^0$. This implies that the number of changes between M and M' is at most $|M \oplus M'| \leq |V^+| + |V^-|$ (or in the one-sided case $|M \oplus M'| \leq 2(|V^+| + |V^-|)$).

First, suppose $|V^-| = 0$. In these steps, the potential strictly increases by at least $|V^+|$. Thus, on the average, for every unit of increase in the potential, the number of changes from M to M' is O(1). Second, suppose $|V^-| \geq 1$. Then for every $v \in V^+$, the potential increases by at least 1. For every $v \in V^-$, it drops by at most Δ. Let $\delta = |V^+| - (\Delta + k)|V^-| > 0$. Thus,

$$\Phi^t(M') - \Phi^t(M) \geq |V^+| - \Delta|V^-| \geq \lceil \delta + k|V^-| \rceil$$

The average number of changes made per unit increase in the potential due to updates of the matching with $V^- > 0$ is at most

$$\frac{|M \oplus M'|}{\Phi^t(M') - \Phi^t(M)} = O\left(1 + \frac{\Delta}{k}\right) \enspace .$$

Finally, we bound the total increase in the potential function over time. Consider the rounds with additions and deletions of edges. If an edge is added in round t, the maximum potential value increases by at most 2 (or 1 in the one-sided case) and the current value of the potential does not decrease. If an edge is deleted, the maximum potential value decreases by at most 2 (or 1 in the one-sided case) and the current value of the potential decreases by at most 2Δ (or Δ in the one-sided case). Thus, in total we can increase the potential up to at most twice the number of edge additions. Also, each deletion creates the possibility to increase the potential by at most 2Δ in subsequent rounds. This implies an amortized potential increase of at most $O(\Delta)$ per round. Also, we get an average number of $O(1+\Delta/k)$ changes in the matching per unit of potential increase. Combining these insights yields the theorem. □

We can strengthen the latter result in case we have only edge additions.

Corollary 1. DEFERREDRESOLUTION *maintains a* $(\Delta+k)$*-unpopular matching by making an amortized number of* $O(1+\Delta/k)$ *changes to the matching per round with edge dynamics without deletions, for any* $k>0$.

Proof. In the previous proof we observed that rounds with edge additions generate an amortized potential increase of 1. Hence, we directly get the average number of $O(1+\frac{\Delta}{k})$ changes in the matching per unit of potential increase also as amortized change per round. □

The following corollary is due to the fact that we can simulate the addition or deletion of a single vertex by Δ additions or deletions of the incident edges. A similar reduction by Δ can be achieved without vertex deletions.

Corollary 2. DEFERREDRESOLUTION *maintains a* $(\Delta+k)$*-unpopular matching by making an amortized number of* $O(\Delta^2+\Delta^3/k)$ *changes to the matching per round with vertex dynamics, for any* $k>0$.

The above results apply in the roommates, two-sided, and one-sided cases. The bound on the unpopularity factor is almost tight, even in terms of existence in the one-sided case.

Proposition 1. *There exist one-sided instances with maximum degree* Δ *for every agent in* Y *such that every matching has unpopularity factor at least* Δ.

Proof. As an example establishing the lower bound consider a one-sided instance with $|X|=\Delta$ elements and $|Y|=\Delta+1$ agents. We assume there is a global ordering $x_1,\ldots,x_\Delta$ over elements and $x_i \succ_y x_{i+1}$ for all agents $y \in Y$. If a matching M leaves an element in X unmatched, we can add any single edge and thereby create a matching with $|V^+|>0$ and $|V^-|=0$. By definition this new matching is now ∞-more popular, and the unpopularity factor becomes $\rho(M)=\infty$. For any matching M that matches all of X, we w.l.o.g. denote y_i as the agent with $M(y_i)=x_i$ for $i=1,\ldots,\Delta$, and $y_{\Delta+1}$ the remaining unmatched agent. We show that M has unpopularity factor Δ by providing a matching M' that is Δ-more popular than M. Consider M' composed of edges (x_i,y_{i+1}) for $i=1,\ldots,\Delta$ and y_1 unmatched. y_1 is a $(-)$-agent, all others are $(+)$-agents. □

3 Two-Sided Matching and Lower Bounds

For the roommates case, the construction in [14] shows that there are instances in which every matching has unpopularity factor of $\Omega(\log\Delta)$. In contrast, in the two-sided case there always exists a stable matching, and every stable matching is a popular matching. However, we show that maintaining a stable or popular matching requires $\Omega(n)$ amortized changes per round, even in instances where we have only edge or vertex additions and every agent has degree at most 2.

Theorem 3. *There exist two-sided instances with $\Delta = 2$ such that maintaining a stable or popular matching requires $\Omega(n)$ amortized number of changes to the matching per round for (1) edge dynamics with only additions, (2) vertex dynamics with only additions in X and Y, (3) vertex dynamics with additions and deletions only in X.*

The case of vertex dynamics and only additions to X can be tackled using the standard DeferredAcceptance algorithm of Gale-Shapley for stable matching.

Proposition 2. DeferredAcceptance *maintains a stable matching by making an amortized number of $O(\Delta)$ changes to the matching per round with vertex dynamics and only additions to X.*

Hence, without any additional assumptions we can only expect to maintain α-unpopular matchings for $\alpha > 1$. Here we observe that our algorithm DeferredResolution cannot be used to maintain matchings with unpopularity factor significantly below Δ, even in the two-sided case. The problem is that the iterative resolution may be forced to cycle.

Theorem 4. *There is an instance with maximum degree Δ and an initial matching such that no sequence of iterative resolution of matchings with unpopularity factor $(\Delta - 1)$ leads to a α-unpopular matching, for any $\alpha < \Delta - 1$.*

It is easy to force DeferredResolution into the cycle. We first add the edges of one cycle matching, then the edges of the more popular cycle matching, and finally the edges of the third cycle matching. DeferredResolution will construct the first cycle matching and switch to the next one whenever it has arrived entirely.

The proof here uses a particular instance with degree $\Delta = 3$. Furthermore, it shows that even though two-sided instances always have popular matchings, there are instances and initial matchings such that no sequence of resolutions towards more popular matchings converges. The following corollary sharply contrasts the one-sided case, in which there always exist voting paths of length 2 whenever a popular matching exists.

Corollary 3. *There are two-sided matching instances and matchings from which there is no voting path to a popular matching.*

More generally, we can establish the following lower bound for any maximum degree $\Delta \geq 3$.

Theorem 5. *For every $\Delta \geq 3$ and $k = 3, \ldots, \Delta$ there is an instance with maximum degree Δ and an initial matching M such that any sequence of resolutions of matchings with unpopularity factor at least $k-1$ does not converge to a $(k-2)$-unpopular matching.*

We can again steer DeferredResolution into the cycle. We first let the edges (x_j, y_j) arrive that remain fixed throughout the cycle, for $j = \Delta - k + 1, \ldots, \Delta$. Then, we let the remaining incident edges arrive for these nodes. DeferredResolution will construct all edges (x_j, y_j) and keep them in the matching throughout. Then, we assume edges (x_j, y_j) arrive iteratively for $j = 1, \ldots, k-1$. DeferredResolution will include each of these edges into the matching. Subsequently, we consider the next matching from the cycle and let the edges arrive iteratively, and so on. DeferredResolution will switch to the next matching in the cycle whenever it has arrived entirely. It then infinitely runs through the cycle once all edges have arrived.

References

1. D. Abraham, R. Irving, T. Kavitha, and K. Mehlhorn. Popular matchings. *SIAM J. Comput.*, 37(4):1030–1045, 2007.
2. D. Abraham and T. Kavitha. Voting paths. *SIAM J. Disc. Math.*, 24(2):520–537, 2010.
3. H. Ackermann, P. Goldberg, V. Mirrokni, H. Röglin, and B. Vöcking. Uncoordinated two-sided matching markets. *SIAM J. Comput.*, 40(1):92–106, 2011.
4. S. Baswana, M. Gupta, and S. Sen. Fully dynamic maximal matching in $O(\log n)$ update time. In *Proc. 52nd Symp. Foundations of Computer Science (FOCS)*, pages 383–392, 2011.
5. S. Bhattacharya, M. Henzinger, and G. Italiano. Deterministic fully dynamic data structures for vertex cover and matching. In *Proc. 25th Symp. Discrete Algorithms (SODA)*, pages 785–804, 2015.
6. P. Biró, R. Irving, and D. Manlove. Popular matchings in the marriage and roommates problems. In *Proc. 17th Intl. Conf. Algorithms and Complexity (CIAC)*, pages 97–108, 2010.
7. E. Diamantoudi, E. Miyagawa, and L. Xue. Random paths to stability in the roommates problem. *Games Econom. Behav.*, 48(1):18–28, 2004.
8. P. Gärdenfors. Match making: Assignments based on bilateral preferences. *Behavioural Sciences*, 20:166–173, 1975.
9. M. Gupta and R. Peng. Fully dynamic $(1+\varepsilon)$-approximate matchings. In *Proc. 54th Symp. Foundations of Computer Science (FOCS)*, pages 548–557, 2013.
10. D. Gusfield and R. Irving. *The Stable Marriage Problem: Structure and Algorithms.* MIT Press, 1989.
11. M. Hoefer. Local matching dynamics in social networks. *Inf. Comput.*, 222:20–35, 2013.
12. M. Hoefer and L. Wagner. Locally stable marriage with strict preferences. In *Proc. 40th Intl. Coll. Automata, Languages and Programming (ICALP)*, volume 2, pages 620–631, 2013.
13. M. Hoefer and L. Wagner. Matching dynamics with constraints. In *Proc. 10th Intl. Conf. Web and Internet Economics (WINE)*, pages 161–174, 2014.
14. C. Huang and T. Kavitha. Near-popular matchings in the roommates problem. *SIAM J. Disc. Math.*, 27(1):43–62, 2013.
15. D. Knuth. *Marriages stables et leurs relations avec d'autres problemes combinatoires.* Les Presses de l'Université de Montréal, 1976.
16. D. Manlove. *Algorithmics of Matching Under Preferences.* World Scientific, 2013.

17. M. McCutchen. The least-unpopularity-factor and least-unpopularity-margin criteria for matching problems with one-sided preferences. In *Proc. 8th Latin Amer. Theoret. Informatics Conf. (LATIN)*, pages 593–604, 2008.
18. K. Onak and R. Rubinfeld. Maintaining a large matching and a small vertex cover. In *Proc. 42nd Symp. Theory of Computing (STOC)*, pages 457–464, 2010.
19. A. Roth and M. O. Sotomayor. *Two-sided Matching: A study in game-theoretic modeling and analysis*. Cambridge University Press, 1990.
20. A. Roth and J. V. Vate. Random paths to stability in two-sided matching. *Econometrica*, 58(6):1475–1480, 1990.

A Omitted Proofs

A.1 Proof of Theorem 3

Proof. Consider a two-sided instance with $|X| = |Y| = n$, where all agents in X and Y have the same preference list over Y and X, respectively. In particular, we denote $x_1, \dots, x_n$ and $y_1, \dots, y_n$ such that for all $i = 1, \dots, n$ we have $x_i \succ_y x_{i+1}$ for every $y \in Y$ and $y_i \succ_x y_{i+1}$ for every $x \in X$.

For part (1), in round $t = 1, \dots, 2n-1$, edge $e^t = (x_{n-2\lfloor (t-1)/2 \rfloor}, y_{n-2\lfloor t/2 \rfloor})$ arrives:

$$x_n \overset{e^1}{\longleftrightarrow} y_n \overset{e^2}{\longleftrightarrow} x_{n-1} \overset{e^3}{\longleftrightarrow} y_{n-1} \overset{e^4}{\longleftrightarrow} x_{n-2} \overset{e^5}{\longleftrightarrow} y_{n-2} \cdots$$

Thereby, we construct a long path, where the stable and popular matchings are unique in every round. Every edge e^t is a blocking pair upon arrival. Resolving this edge iteratively introduces a new blocking pair and requires to alternate through the entire existing path. Hence, in every round $t \in \Theta(n)$ we have to make $\Theta(n)$ changes to the matching to maintain the stability conditions. It is easy to see that the unique stable matching is also the unique popular matching here. This yields $\Theta(n)$ amortized number of changes for every round to maintain stable or popular matchings. The degree of every agent is at most 2.

For part (2), vertices of X and Y now arrive alternatingly one after the other, starting from x_n and y_n to x_1 and y_1. In round $t = 1, 3, 5, \dots$ we assume that $x_{n-2\lfloor t/2 \rfloor}$ arrives along with all edges to previously arrived vertices, and in round $t = 2, 4, 6, \dots$ vertex $y_{n-2\lfloor (t-1)/2 \rfloor}$ arrives along with all edges to previously arrived vertices. The overall edge set again forms the above described path, which is revealed iteratively as vertices arrive. It is easy to see that the same argumentation as above can be applied to show $\Theta(n)$ amortized number of changes for every round to maintain stable or popular matchings.

For part (3), we assume that in the first $O(n)$ rounds, a graph G is constructed by arrivals of the above described alternating path. In the subsequent rounds, x_n iteratively arrives, departs, arrives again, departs again, etc. Depending on x_n being present or not, the unique stable and popular matching requires to switch the edges along the entire alternating path down to x_1 in every round. Hence, for every subsequent round we need to make $\Theta(n)$ changes to the matching. $\square$

A.2 Proof of Proposition 2

Proof. It is straightfoward that when only vertices of X arrive, the standard X-proposing DeferredAcceptance algorithm can maintain a X-optimal stable matching with an overall number of $O(|E|)$ many changes to the matching. For this, we simply consecutively implement the proposal algorithm and continue the execution whenever a new vertex arrives. Each vertex has to account only for a number of changes in the order of its degree.

Let us also mention a simple adjustment, by which this algorithm maintains a 2-unpopular matching with edge dynamics and no deletions when edges arrive consecutively for their incident vertex in X. When the first edge of x_i arrives that is connected to a currently unmatched node y_j, we temporarily match x_i and y_j. This guarantees that during the consecutive arrival of x_i's incident edges, we generate at most an unpopularity factor of 2. If the last edge of x_i has arrived, we remove any edge incident to x_i and continue DeferredAcceptance by including x_i's proposals. This generates at most $O(|E|)$ many changes to the matching and, hence, amortized $O(1)$ per edge. □

A.3 Proof of Theorem 4

Proof. We consider an instance with $|X| = |Y| = \Delta = 3$ and an initial state such that it is impossible to reach any α-unpopular matching with $\alpha < \Delta - 1$ via resolution of alternating cycles with unpopularity factor of at least $\Delta - 1$. The preferences are as follows.

x_1	x_2	x_3
y_2	y_1	y_3
y_1	y_3	y_2
y_3	y_2	y_1

y_1	y_2	y_3
x_3	x_2	x_1
x_1	x_3	x_2
x_2	x_1	x_3

There exists a cycle of three perfect matchings:

$$\begin{array}{ccc} \{(x_1, y_1), (x_2, y_2), (x_3, y_3)\} & \longrightarrow & \{(x_1, y_2), (x_2, y_3), (x_3, y_1)\} \\ \nwarrow & & \swarrow \\ & \{(x_1, y_3), (x_2, y_1), (x_3, y_2)\} & \end{array}$$

Each of these matchings has unpopularity factor exactly $\Delta - 1 = 2$ over the previous one. In particular, in this cycle there is exactly one edge between two $(+)$-agents and two edges with a single $(+)$-agent and $(-)$-agent in every step. It is, thus, impossible to simply drop any of the edges and keep a positive factor. The remaining three perfect matchings are stable matchings:

$$\begin{aligned} X\text{-optimal:} \quad & \{(x_1, y_2), (x_2, y_1), (x_3, y_3)\} \\ \text{Middle:} \quad & \{(x_1, y_1), (x_2, y_3), (x_3, y_2)\} \\ Y\text{-optimal:} \quad & \{(x_1, y_3), (x_2, y_2), (x_3, y_1)\} \end{aligned}$$

Note that all popular matchings must be perfect. As we have listed all perfect matchings, the popular matchings are exactly the stable matchings here. However, from every cycle matching to any of the stable matchings the unpopularity factor is 1. For every cycle matching there is no non-perfect matching that yields an unpopularity factor greater than 1. Hence, if we initially have a cycle matching, iterative resolution is forced to follow the cycle. Thus, it is impossible to converge to a α-unpopular matching with $\alpha < \Delta - 1 = 2$. □

A.4 Proof of Theorem 5

Proof. We consider an instance as in the proof of Theorem 3 with $|X| = |Y| = k$, where every $x \in X$ has the same preference over Y, and every $y \in Y$ the same preference over X. We denote $x_1, \ldots, x_k$ such that $x_i \succ_y x_{i+1}$ for all $i = 1, \ldots, k-1$ and $y \in Y$. Similary, we denote $y_1, \ldots, y_k$ such that $y_i \succ_x y_{i+1}$ for all $i = 1, \ldots, k-1$ and $x \in X$. We assume that $E = (X \times Y) \setminus \{(x_k, y_k)\}$.

Consider the matching $M = \{(x_i, y_{k+1-i}) \mid i = 1, \ldots, k\}$. Since all agents are matched in M and all agents in X and Y have the same preference list, we create at least two $(-)$-agents when changing the matching. When we increase the preference for some agent x_i (i.e., the new partner of x_i has smaller index), we steal this better partner from some other agent $x_{i'}$. Either $x_{i'}$ becomes a $(-)$-agent, or it steals a better partner from some other agent in X. While this can generate a chain of changes, we do not change the indices of agents in Y. Hence, there must be at least one $(-)$-agent in X that ends up either single or with a partner of higher index. The same is true for the agents in Y, so we create at least two $(-)$-agents.

Hence, we require at least $2(k-1)$ $(+)$-agents to obtain the required unpopularity factor. There is a unique matching $M' = \{(x_i, y_{k-i}) \mid i = 1, \ldots, k-1\}$ that has unpopularity factor at least $k-1$ over M. Here the $(-)$-agents are x_k and y_k (unmatched in M') and all other agents are $(+)$-agents. Using a similar argumentation, M' has a unique matching M'' that verifies unpopularity factor at least $k-1$, where we cyclically shift all agents in X (Y) to the next higher agent in Y (X), and match the current partners of x_1 (y_1) to x_k (y_k), respectively. As there is no edge that simultaneously combines the $(-)$-agents, we have no opportunity to leave out any edge, and hence the new matching M'' must be perfect. Applying this argument iteratively, we see that we are forced to return to M.

To adjust the maximum degree, we can introduce additional agents x_j and y_j to X and Y, for $j = \Delta - k + 1, \ldots, \Delta$. For their preferences we assume $y_j \succ_{x_j} y_{j'}$ and $x_j \succ_{y_j} x_{j'}$ for all $j' = 1, \ldots, j-1, j+1, \ldots, \Delta$. The remaining preferences of x_j and y_j are completed arbitrarily. For every $i = 1, \ldots, k$ we have $y_{i'} \succ_{x_i} y_j$ and $x_{i'} \succ_{y_i} x_j$ for every $i' = 1, \ldots, k$ and $j = \Delta - k + 1, \ldots, \Delta$. In the initial matching we add (x_j, y_j) to M. Therefore, none of the agents x_j or y_j can be $(+)$-agents. Removing any edge (x_j, y_j) creates more $(-)$-agents and thereby makes it impossible to obtain unpopularity factor of at least $k-1$. Hence, the unique M' that has unpopularity factor at least $k-1$ over M executes the cyclic shift as above and keeps all edges (x_j, y_j) for $j = \Delta - k + 1, \ldots, \Delta$. $\square$