


Fast, Fair and Truthful Distributed Stable Matching for Common Preferences

Juho Hirvonen ✉ 

Helsinki Institute for Information Technology HIIT, Finland
Department of Computer Science, Aalto University, Espoo, Finland

Sara Ranjbaran ✉ 

Department of Computer Science, Aalto University, Espoo, Finland

Abstract

Stable matching is a fundamental problem studied both in economics and computer science. The task is to find a matching between two sides of agents that have preferences over who they want to be matched with. A matching is stable if no pair of agents prefer each other over their current matches. The deferred acceptance algorithm of Gale and Shapley solves this problem in polynomial time. Further, it is a mechanism: the proposing side in the algorithm is always incentivised to report their preferences truthfully.

The deferred acceptance algorithm has a natural interpretation as a distributed algorithm (and thus a distributed mechanism). However, the algorithm is slow in the worst case and it is known that the stable matching problem cannot be solved efficiently in the distributed setting. In this work we study a natural special case of the stable matching problem where all agents on one of the two sides share common preferences. We show that in this case the deferred acceptance algorithm does yield a fast and truthful distributed mechanism for finding a stable matching. We show how algorithms for sampling random colorings can be used to break ties fairly and extend the results to fractional stable matching.

2012 ACM Subject Classification Theory of computation → Algorithmic mechanism design; Theory of computation → Distributed algorithms

Keywords and phrases stable matching, deferred acceptance, local algorithm, mechanism design

Digital Object Identifier 10.4230/LIPIcs.OPODIS.2024.30

Funding *Sara Ranjbaran*: Supported by Helsinki Institute for Information Technology HIIT.

1 Introduction

In mechanism design the goal is to design optimisation algorithms that work when the input is private information held by strategic agents. For example, in the case of auctions, each agent has a private valuation for the item being sold and the goal is to find an auction mechanism that incentivises the agents to reveal their true valuations while assigning the item to the agent that values it the most.

Stable matching is one of the seminal problems studied in mechanism design [28]. There is a bipartite matching graph where each node is an agent and each edge indicates a potential match. The agents have a linear order over their neighbours that is initially only known to them. Each agent reports *some* linear order and based on these an algorithm computes a matching. The *deferred acceptance* algorithm of Gale and Shapley [18] finds a stable matching: there is no pair of agents that would prefer each other over their assigned matches. See Figure 1 for an illustration.

The algorithm of Gale and Shapley consists of rounds of proposals. One side of the bipartition is the proposing side and one is the receiving side. In each round each unmatched proposing-side agent proposes to its next most preferred match. Agents that receive proposals accept the best proposal if they prefer it over their current match (deferred acceptance).



© Juho Hirvonen and Sara Ranjbaran;

licensed under Creative Commons License CC-BY 4.0


28th International Conference on Principles of Distributed Systems (OPODIS 2024).

Editors: Silvia Bonomi, Letterio Galletta, Etienne Rivière, and Valerio Schiavoni; Article No. 30; pp. 30:1–30:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** A matching instance and a stable matching in bold. The preferences are given as a list for the neighbours from left to right, with 1 indicating the highest preference. The matching is *stable*: there is no pair of unmatched agents connected by an edge such that they would prefer each other over their current match. Since there are more agents on one side, some agents must remain unmatched.

The deferred acceptance algorithm has the important property that, for the proposing side, revealing their true preferences is a dominant strategy [11, 32]. This makes the deferred acceptance algorithm one of the few examples of a mechanism without money. It is also known that there is no mechanism that finds a stable matching and is truthful for all agents [32].

The deferred acceptance algorithm can be applied to any setting where two sides need to be matched based on preference criteria held by the individual agents. The original motivation of Gale and Shapley was college admission: how to match students to school seats [18]. The algorithm naturally generalises to this kind of one-to-many matching, where a single school can be matched to many students. Due to its simplicity and strategic properties, the deferred acceptance algorithm has also been widely deployed in practice [28, 34], for example in the National Resident Matching Program (NRMP) in the US [33] and New York City and Boston public high schools [35].

The deferred acceptance algorithm has a natural interpretation as a distributed proposal algorithm. However, its worst-case running time is $O(n^2)$ rounds [30] and computing a stable matching requires $\Omega(n)$ rounds in the LOCAL model [17]. We study a special case of the general stable matching problem where one of the two sides of agents has common preferences, i.e. the preferences of the agents are restrictions of the same partial order [21]. This special case corresponds to the case where the preferences of one side come from some objective ranking, such as standardised tests in school choice. This is a common assumption when studying matching between autonomous agents, for example in biology [3], economics [7], game theory [12], and resource assignment in computer networks [19].

1.1 Our contribution

We present several extensions of the deferred acceptance algorithm in the setting where one of the sides has common preferences. Our algorithms implement the deferred acceptance algorithm as an efficient distributed mechanism [14], and thus can be applied in settings where a single centralised entity running the mechanism is infeasible. They are simple and retain the potential of the centralised deferred acceptance algorithm for practical applications.

We will next present informal versions of our main results. We show that for one-sided common preferences, the deferred acceptance algorithm can be implemented efficiently as a distributed algorithm. Preferences are represented as a score function with S different preference classes, and let Δ_C and Δ_P denote the maximum degrees of the two sides in the matching graph. Let n denote the number of agents on the proposing side.

► **Theorem 1** (Informal, Theorem 8). *There is an incentive-compatible distributed implementation of the deferred acceptance algorithm in the CONGEST model for one-sided common preferences, with running time $\min\{O(\Delta_C \Delta_P S + \log^* n), O(n)\}$.*

The formal version of this Theorem is given in Theorem 8 in Section 3.

Our algorithms are incentive-compatible, i.e. the agents on the proposing side are never worse off by reporting their true preferences. The most interesting application for our algorithms is in sparse graphs: We are not aware of any other incentive-compatible mechanisms with efficient distributed implementations, even in this setting. The classic deferred acceptance algorithm does run in $O(S)$ rounds, but requires that there are no ties. Providing this fast and fair tie-breaking procedure is one of our main contributions.

It should be noted that the distributed execution of our algorithm is not claimed to be strategy-proof: we assume the standard mechanism design setting where the agents are strategic in revealing their private information, but the mechanism itself is run faithfully.

The dependency on both Δ and S is necessary, as we discuss in Section 2.3. First, stable matching is at least as hard as maximal matching, which is known to require $\Omega(\min\{\Delta, \log n / \log \log n\})$ rounds [5]. Second, we show that even for $\Delta = 2$, finding a stable matching with common preferences requires $\Omega(S)$ rounds.

We also show how to extend the algorithm to the fractional case that we call *capacitated fractional stable matching*. The matching itself can take fractional values and both sides have a maximum capacity. This can be seen as a generalisation of stable matching to a setting where the clients have divisible jobs that are assigned to the providers for completion.

For general preferences, this setting is not well defined, as one must first define how agents preferences weight quality of match against the total amount of a match [8]. We show that in the case of one-sided common preferences, this is not an issue and the natural greedy extension of the deferred acceptance algorithm is incentive-compatible as the truth maximises both the quality and the quantity of the matching.

► **Theorem 2** (Informal, Theorem 13). *There is an incentive-compatible distributed implementation of a generalised deferred acceptance algorithm for capacitated fractional stable matching in the CONGEST model for one-sided common preferences, with running time $\min\{O(\Delta_C \Delta_P S + \log^* n), O(n)\}$.*

We assume that the preferences may contain ties. This means that an algorithm must break these ties in some way, and the procedure has implications on the properties of the mechanism. For example, if the tie-breaking does not retain the commonness of the preferences, then it may cause the deferred acceptance algorithm to become slow to converge. We solve this issue by using graph coloring as a tie-breaking procedure. While any coloring is incentive-compatible, we also want the coloring to be fair. We propose a simple and fast algorithm for sampling a random coloring (Section 5.2) and also show how existing algorithms for sampling random colorings can be used as a subroutine in our algorithm (Section 5.1).

The first algorithm samples an almost uniformly random coloring using an existing subroutine [16]. Let $\delta > 0$ denote the *total variation distance* from uniform random coloring.

► **Theorem 3** (Informal, Theorem 17). *There is an incentive-compatible, randomised distributed implementation of the deferred acceptance algorithm in the CONGEST model for one-sided common preferences with a tie-breaking routine that has total variation distance $\delta > 0$ from random and running time $O(\Delta_C \Delta_P S + \log(n/\delta))$.*

A formal version of this Theorem is given in Theorem 17, Section 5.1.

The second algorithm uses a simpler color sampling that only guarantees that some subset of agents succeed in being properly colored. These agents have a tie-breaking rule that is conditionally uniform, but the remaining agents have an arbitrary tie-breaking.

► **Theorem 4** (Informal, Theorem 19). *For any $\delta > 0$ there is an incentive-compatible, randomised distributed implementation of the deferred acceptance algorithm in the CONGEST model for one-sided common preferences with a tie-breaking routine that is uniformly random for a subset of $(1 - \delta)n$ proposing-side agents and runs in time $O(\delta^{-1}\Delta_C\Delta_P S + \log^* n)$.*

The formal version is given in Theorem 19, Section 5.2. Both of these fair tie-breaking procedures can also be applied to capacitated fractional stable matching, yielding algorithms with the same running times and guarantees.

Since finding a stable matching with general preferences is hard in the distributed setting [17], our results suggest a trade-off for market design: in certain scenarios, it might be preferable to force one of the sides to use common preferences in order to make the matching procedure more efficient.

1.2 Related work

The *deferred acceptance* or *Gale-Shapley algorithm* computes a stable matching: a matching (i.e. a subset of edges such that each agent belongs to at most one edge) such that no unmatched potential pair prefers each other over their assigned matches [18]. As a mechanism it is incentive-compatible for the proposing side of the algorithm [11, 32]. There does not exist a mechanism (without payments) that is incentive-compatible for both sides at the same time [32]. We refer to surveys by Knuth [25] and by Roth [34] for more details.

Irving, Manlove, and Scott studied the stable matching problem with common preferences (called master preference lists there) [21]. This setting was also studied by Scott [37] and O'Malley [29]. Generalisations of this setting were studied by Kamiyama [22, 23].

In the distributed setting, Khanchandani and Wattenhofer studied a variant of stable matching where preference lists are similar on one side [24]. The running time of their algorithm is $O((\Delta + 1)n)$, where Δ is the similarity parameter. When $\Delta = 0$ their setting corresponds to ours, except that we consider incomplete matching graphs. They show that in the CONGEST model computing a stable matching with $\Delta = 0$ requires $\Omega(n/\log n)$ rounds. This is further support for restricting our attention to the sparse case. Amira, Giladi, and Lotker studied a variant where the preferences are given as edge weights and gave an algorithm that finds a stable matching in $O(\sqrt{n})$ rounds [4].

Floréen, Kaski, Polishchuk, and Suomela [17] showed that finding a stable matching requires $\Omega(n)$ rounds in the LOCAL model. Kipnis and Patt-Shamir showed that solving the stable matching problem, even in graphs of diameter $\Theta(n)$, requires $\Omega(\sqrt{n}/\log n)$ rounds in the CONGEST model. Most variants of stable matching are at least as hard as finding a maximal matching, and this is known to require $\Omega(\Delta)$ rounds, where Δ is the maximum degree of the graph [5].

Finding a stable matching is hard in the distributed setting, and there has been a lot of interest in finding *almost stable matchings*. An edge is unstable if it is not in the matching but both of its endpoints prefer each other to their current matches, and a matching M is ε -stable if the number of unstable edges is at most $\varepsilon|M|$. Floréen, Kaski, Polishchuk, and Suomela [17] showed that truncating the deferred acceptance algorithm after $O(\Delta^2/\varepsilon)$ rounds results in an ε -stable matching. Ostrovsky and Rosenbaum present an $O(\log^5 n)$ -round deterministic algorithm and an $O(\log^2 n)$ -round randomised algorithm for computing ε -stable

matchings [30]. Hassidim, Mansour, and Vardi [20] study the problem in the centralised *local computation algorithm (LCA)* model. None of these algorithms have the strategic guarantees of the deferred acceptance algorithm, as we discuss in Section 2.3.

We also study fractional stable matchings, which have been studied in the literature due to their connection to integral stable matchings [38, 36]. Caragiannis, Filos-Ratsikas, Kanellopoulos, and Vaish [8] resolve the issue ill-defined utility by studying cardinal preferences. See their work for further references on the subject.

Our algorithms can be seen as examples of *distributed algorithmic mechanism design (DAMD)* [14]. In DAMD the goal is to design mechanisms that can be run by the participating agents. Work in this field has typically dealt with distributed mechanisms for global problems in distributed networks, such as multi-cast cost sharing [13], leader election [1], consensus, and renaming [2]. We are not aware of any distributed mechanisms implemented by a sublinear-time algorithm (even for sparse graphs) before our work.

2 Background and definitions

2.1 Distributed algorithms

We study distributed algorithms in the standard LOCAL and CONGEST models of distributed computing [31, 26]. The system is represented as a graph $G = (V, E)$ where the set of nodes V are the computational agents and the edges E are the communication links: two nodes u and v can communicate directly if and only if $\{u, v\} \in E$. Nodes connected to u are its neighbors and are denoted by $N(u)$. The degree (number of neighbours) of node v is denoted by $\deg(v)$. The maximum degree Δ of a graph G is $\max_{v \in V} \deg(v)$. We assume that n and Δ are known to all nodes. Each node has a unique identifier from $[\text{poly}(n)]$.

Computation proceeds in synchronous rounds. In each round, each node can send a message to each of its neighbors, receive the messages sent by its neighbors, and update its state. In the LOCAL model these messages can be unbounded in size, and the CONGEST model only differs from the LOCAL in that the size of the messages is $O(\log n)$ bits. Eventually each node must stop and announce its own output. The running time of an algorithm is the number of communication rounds until all nodes have stopped.

Unless otherwise stated, all algorithms are deterministic. In randomised algorithms, nodes additionally have a uniformly random and sufficiently long string of bits as input.

2.2 Mechanism design

A mechanism M is a game: It consists of a set V of *agents* and each agent $v \in V$ has a private *type* $t(v)$ representing its utility function. In the case of stable matching, the type of the agent is its preference relation over its neighbours. For each agent there is a set of possible *actions* $A(v)$ representing the possible ways of reporting its type (preferences) to the mechanism. The *strategy* $a(v)$ of agent v is its chosen action. The mechanism has an *allocation function* F that maps the strategies of the agents to a set O of possible *outcomes*.

A mechanism is *incentive-compatible* or *truthful* if reporting the truth (i.e. choosing $a(v) = t(v)$) is a dominant strategy: agent is never worse off by reporting truthfully. For randomised mechanisms, we say that a mechanism is (universally) incentive-compatible if truth is a dominant strategy for all choices of randomness.

In this work we consider *distributed mechanisms* where the allocation function can be implemented as a distributed algorithm. Each agent corresponds to a node in graph G . The strategy of each agent is part of the input of the corresponding node. After the strategies

are revealed in the input, the agents can no longer affect the mechanism and we assume that the algorithm implementing the allocation rule is run faithfully. Designing distributed algorithms that are strategy-proof *to run* is a distinct hard problem, and even very simple greedy algorithms can be difficult to analyse [10].

2.3 Stable matching and deferred acceptance

We study the graph-theoretic version of the stable matching problem. The agents form a bipartite matching graph $G = (C \cup P, E)$. We call the side C *clients* and the side P *providers*: the clients correspond to the proposing side in the deferred acceptance algorithm. The edges represent the possible matches of each agent, and these are all assumed to be preferable to being unmatched. The goal is to compute a (weakly) stable matching: a matching where no pair of agents strictly prefer each other to their current matches.

We assume that the communication network $N = (C \cup P, E')$ is G with additional edges between any pair of clients that are connected to the same provider. The modelling choice is somewhat arbitrary, and for most parameter values it allows us to simplify the description of the CONGEST algorithms. For some large values of Δ , the maximum degree of the network, there is a difference between the running time of the best LOCAL and CONGEST algorithms if the matching graph is the communication network. We discuss this in Section A.3.

When we refer to degrees, we mean the degrees of the matching graph unless otherwise specified. We denote by Δ_C the maximum degree among the clients and by Δ_P the maximum degree among the providers, and assume that these are known to the agents.

The type of each client and each provider is a linear order on its incident providers, determining in which order it prefers to be matched with them. The basic deferred acceptance algorithm can be implemented as a distributed algorithm as follows. Let $>_v$ denote the announced preferences of agent v . Each client v initialises a list $L(v)$ of its $\deg(v)$ neighbours in the preference order, starting from the most preferred. Each client and provider v maintains a match variable $m(v)$ initially set to unmatched.

On odd rounds $2i - 1$ for $i = 1, \dots$:

1. Each unmatched client v sends message *proposal* to its most preferred match in $L(v)$ and removes it from $L(v)$.
2. Each matched client waits.
3. Each provider p that receives one or more proposals picks the one that is from the most preferred client v , and checks if $v >_p m(p)$. If yes, set $m(p) \leftarrow v$.


On even rounds $2i$ for $i = 1, \dots$:

1. Each provider p sends $m(p)$ to its neighbours (as acknowledgement of proposal acceptance or rejection).
2. Each client v that receives $m(p) = v$ from a neighbour sets $m(v) \leftarrow p$. Each client u with $m(u) = p$ that receives $m(p) \neq u$ from p sets $m(u) = \perp$ (unmatch).

The algorithm runs in $O(n^2)$ rounds. For common preferences, this can be improved to $O(n)$: by using the unique identifiers to break ties consistently, we obtain common preferences with n preference classes. In each round the client in the highest preference class becomes matched.

It is known that computing a stable matching requires $\Omega(n)$ rounds [17].

► **Theorem 5** (Floreen et al. [17]). *Computing a stable matching in a system where each client and provider has maximum degree 2 requires $\Omega(n)$ rounds in the LOCAL model.*



■ **Figure 2** Two instances of stable matching that only differ in the preferences of agent v on the left. This causes the unique stable matching to change, and therefore agent u on the right has to learn the preferences of v in order to compute a stable matching.

The theorem follows from a simple argument: for any path, there exist two assignments of preferences to the agents such that 1) they only differ at one agent at one end of the path and 2) the resulting instances both have a unique stable matching and these matchings are edge-disjoint. See Figure 2 for an illustration.

Stable matching does not become easy when preference lists of the providers are *very similar*. Khanchandani and Wattenhofer [24] studied the setting where all clients and providers can be matched (matching graph is the complete bipartite graph). Their results cannot be directly compared with the incomplete case, as providers might not share more than one client. We can show that even if all providers that share adjacent clients share at least δ adjacent clients computing a stable matching requires $\Omega(n/\delta)$ rounds in the LOCAL model. The details are given in Appendix A.2.

We study a special case of the stable matching problem where the providers have *common preferences* (also known as *master lists*): they always agree when comparing two clients. We assume that the preferences of the providers can be expressed as a single *score function* $s: C \rightarrow \{1, \dots, S\}$. We call the value $s(v)$ of agent v its *preference class*. This score function assumption is necessary, as in the example of Theorem 5 no providers share a pair of adjacent clients and thus trivially agree in their preferences.

We must bound S to have any hope of having a fast distributed algorithm, as any distributed algorithm for finding a stable matching requires $\Omega(S)$ rounds, even when $\Delta = 2$.

► **Theorem 6.** *Any distributed algorithm in the LOCAL model for computing a stable matching when providers have common preferences with S preference classes requires $\Omega(S)$ rounds.*

The theorem follows from a similar construction as Theorem 5. It is possible to construct two paths such that they only differ in the preferences of one client in one end, but the output of another client at the other end depends on this difference. See Appendix A.1 for details, and Figure 3 there for an illustration.

It is easy to see that *if there are no ties among the neighbours of any provider* then the deferred acceptance algorithm also terminates in $O(S)$ rounds: in each round of proposals, the agents in the highest remaining preference class become matched and will stay matched.

If the preferences are represented in a more fine-grained manner, for example as real numbers from an interval [7], then by discretising them into S intervals it is possible to achieve a trade-off between solution quality and running time. While explicitly knowing the possible range of values of preferences s is not important, an upper bound on the number of distinct preference classes needs to be known.

2.3.1 Tie-breaking with common preferences

Several natural procedures for breaking ties fail, in particular when used in the distributed setting. A first come, first served rule that prefers clients that are already matched fails because it loses the truthfulness property: clients may be incentivised to propose in a different order to affect the tie-breaker. Using the unique identifiers to break the ties fails because it introduces too many preference classes: as shown in Theorem 6, finding a stable matching requires $\Omega(S)$ rounds. Similarly, using the port numbering to break symmetry makes the problem as hard as general case.

We use colorings to break ties in a consistent way. This approach has the disadvantage that closing the gap to the lower bound of $\Omega(\max\{\Delta, S\})$ appears difficult, as the tie-breaking graph has maximum degree $O(\Delta_C \Delta_P)$. It is an interesting open question whether the tie-breaking can be computed in a more efficient manner while retaining the truthfulness and the fast running time of the deferred acceptance algorithm.

2.3.2 Variants of deferred acceptance

It is typically assumed that the preferences of the agents are strict. However, the deferred acceptance algorithm generalises naturally to the case where there are ties by using a suitable tie-breaking mechanism. It is important to note that, in order to maintain incentive-compatibility, the procedure to break ties should not depend on the preferences of the agents. For example, a tie-breaking rule preferring the existing match in case of a tie could incentive agents to lie in order to obtain a match early.

The deferred acceptance algorithm also generalises to one-sided one-to-many matching. For example, in the case of school choice, schools have multiple seats and can thus accept multiple applicants. A simple reduction represents a provider with k seats as k identical copies of a provider with unit capacity. It is also possible to modify the algorithm itself so that the providers accept k best matches. This has the advantage of not increasing the maximum degree of the graph.

However, the deferred acceptance algorithm does not generalise to the case of many-to-many matching, for example clients having multiple jobs that need to be assigned to providers for completion. In this case it is not clear what the utility function of a client is: for example, how does it prefer match quality compared to match size? We show in Section 4 that in the case of common preferences, there is a natural generalisation that can be solved efficiently.

2.3.3 Almost stable matching

Several existing algorithms solve the almost stable matching problem [17, 30] in the distributed setting. An edge is unstable if it is not in the matching but both of its endpoints prefer each other to their current matches. The goal is to find a matching M with at most $\varepsilon|M|$ unstable edges.

It is important to note that previous efficient algorithms for almost stable matching are not incentive-compatible. For example, truncating the deferred acceptance algorithm leads to an almost stable matching [17]. This algorithm, however, is not incentive-compatible. As an example, it could be that in the last round, agent v becomes unmatched from u when it accepts a better proposal. Since this was the last round, v will remain unmatched. Now assume it had a neighbor u' that also is unmatched: by lying that $u' >_v u$ agent v could have ensured that it is matched with u' .

3 Fast and fair deferred acceptance algorithm

3.1 Our basic algorithm

Since we assume that the clients are scored with a bounded number of values, we may have ties among the clients connected to any provider. Since the original deferred acceptance algorithm assumes that the preferences are distinct, these ties must be broken in a way that preserves the truthfulness of the mechanism and does not create long dependencies that would slow down the algorithm. We do this by computing a coloring of the virtual conflict network formed by clients that share a preference class and a provider. These colors are used to break ties: smaller color is preferred over a larger color if the preference classes are equal. Since the tie breaking rule is precomputed, clients cannot affect it by lying about their preferences, and the mechanism remains truthful. We note that tie-breaking can affect the quality of the solution, in particular when the preference lists are not complete. We do not attempt to optimise this.

The algorithm consists of three phases:

- (1) Given the matching graph $G = (C \cup P, E)$, construct the virtual conflict network $H = (C, E')$, with the set of clients C as nodes, where two clients v and u are connected by an edge $\{v, u\} \in E'$ if and only if $s(v) = s(u)$ and there exists $w \in P$ such that $\{v, w\}, \{u, w\} \in E$.
- (2) Compute a c -coloring ϕ of H for some c . Create a new score function $s \circ \phi$ defined as $(s \circ \phi)(v) = c(s(v) - 1) + \phi(v)$.
- (3) Run the distributed deferred acceptance algorithm with the provider-side common preferences given by the score function $s \circ \phi$.

In step 2, any coloring that does not depend on the preferences of the clients preserves truthfulness. In Section 5 we discuss how the coloring can be done in a *fair* way. Without specifying the coloring algorithm, we get the following Lemma.

► **Lemma 7.** *There exists a distributed mechanism in the CONGEST model for stable matching when the providers have common preferences. The mechanism is incentive-compatible for the clients and the algorithm implementing it runs in $O(T_c(n, \Delta_H) + S \cdot c)$ rounds, where $T_c(n, \Delta)$ is the running time of a distributed c -coloring algorithm on networks of size n and maximum degree Δ and S is the number of common preferences.*

The mechanism is deterministic unless the coloring algorithm is randomised. By using the best deterministic $(\Delta + 1)$ -coloring algorithm we get the following.

► **Theorem 8.** *There exists a deterministic distributed mechanism in the CONGEST model for stable matching when the providers have common preferences. The mechanism is incentive-compatible for the clients and the algorithm implementing it runs in $\min\{O(S\Delta_H + \log^* n), O(n)\}$ rounds.*

Here, the maximum degree of the conflict graph is bounded by $\Delta_H \leq \Delta_C \Delta_P$.

Proof. Using the $(\Delta + 1)$ -coloring algorithm of Maus and Tonoyan [27] that runs in $O(\sqrt{\Delta \log \Delta} + \log^* n)$ rounds we can color H with $\Delta_H + 1$ colors in time $O(\sqrt{\Delta_H \log \Delta_H} + \log^* n)$. Using this in Lemma 7, the algorithm then runs in $S(\Delta_H + 1)$ rounds, giving a total running time $O(S\Delta_H + \log^* n)$.

If running time of this algorithm is more than $2n$, we can revert to another simple algorithm: use the unique identifiers to break ties. There are at most n preference classes among the clients and thus the running time will be $O(n)$. ◀

We will now proceed to prove Lemma 7.

3.2 Proof of Lemma 7

If the providers have common preferences with S preference classes, then the deferred acceptance algorithm runs in $O(S)$ rounds in the CONGEST model.

► **Lemma 9.** *If providers have common preferences s represented with at most S preference classes and the preferences of each provider form a strict linear order over its adjacent clients, then the distributed deferred acceptance algorithm runs in $2S - 1$ rounds in the CONGEST model of distributed computing.*

Proof. This follows from the fact that in each proposal-answer iteration, consisting of 2 communication rounds, all clients with the highest preference class remaining become matched and stop.

In the first round, the proposals of clients in preference class 1 are always accepted by the providers, as the linear order is strict. They will therefore stop proposing and will not participate in the algorithm further. In the third round, clients in the second preference class either had their proposal accepted in the first round, or have their second proposal accepted in the third round. The latter follows from the fact that all clients of (higher) preference one have already stopped. By induction, by round $2i - 1$, for $i = 1, 2, \dots, S$, all agents with preference class at most $i - 1$ have stopped. In the final round no acknowledgement is required, as by the strictness of the preferences the clients that are active before round $2S - 1$ are the only such neighbours of their active adjacent providers. ◀

The score function $s \circ \phi$ constructed in step 2 of the algorithm induces a strict linear order for each provider over its adjacent clients, as required in Lemma 9. The total number of preference classes is Sc .

► **Lemma 10.** *Given a score function s and a c -coloring ϕ , the score function $s \circ \phi$, restricted to the neighborhood of any single provider, forms a strict linear order. The number of preference classes is Sc .*

Proof. By construction, the new score function assigns values between 1 and Sc . If two clients v and u have $s(v) \neq s(u)$, then by construction they also have $(s \circ \phi)(v) \neq (s \circ \phi)(u)$. Finally, if they have $s(v) = s(u)$ and share a neighbour, then $\{u, v\}$ is in the conflict graph and $\phi(v) \neq \phi(u)$ and therefore $(s \circ \phi)(v) \neq (s \circ \phi)(u)$. ◀

We also want the algorithm to be a mechanism, i.e. revealing the true preferences should be a dominant strategy. To prove this in the special case of common preferences for providers, we use the following technical lemma that states that the match of a client is its highest preference after the matches of clients with a higher preference have been removed.

► **Lemma 11.** *Assume that the providers have common preferences s such that the preferences of each provider form a strict linear order over its adjacent clients. Let $C(i)$ denote the set of clients v with $s(v) = i$, and let $M(i)$ denote the matching obtained by running the deferred acceptance algorithm in the subgraph induced by $\cup_{j=1}^i C(j) \cup P$. For all i and each $v \in C(i)$ it holds that the match of v in $G = (C \cup P, E)$ under the deferred acceptance algorithm is its most preferred match in the subgraph of G induced by $(C \cup P) \setminus M(i - 1)$.*

Proof. We note three facts about the deferred acceptance algorithm. First, each client can only become matched with its k th preferred match if each of the $k - 1$ more preferred providers 1) have already been proposed to by v , and those proposals were rejected, 2) have

unmatched v due to a later more preferable proposal, or 3) have announced a match with a client with a better priority (P1). Second, each provider only accepts proposals from more preferred clients and therefore its match can only improve during the algorithm (P2). Third, as in the proof of Lemma 9, all clients in $C(1), \dots, C(j)$ have stopped after round $2j$ (P3).

The proof is by induction on i . For the base case $i = 1$, the first proposal of each $v \in C(1)$ is accepted. Therefore $M(1)$ consists of the most preferred match of each client. Now assume the claim holds up to some j . Consider $v \in C(j+1)$ and round $2j+1$. If v is matched to some provider p , all higher priority providers are matched to a higher priority client by P1 and P2. By P3, v will remain matched to p until the end of the execution. If v is unmatched, let p' be the highest priority provider adjacent to v that has not announced that it is matched to a client in $\cup_{i=1}^j C(i)$. Since v is unmatched, it has not yet proposed to p' (P2). It will now propose to it, and, since higher priority clients have stopped, become matched. If no p' exists, then by (P1) and (P2) all neighbours of v are matched to a more preferred client. ◀

Proof of Lemma 7. By Lemma 10, the score function $s \circ \phi$ gives a partial linear order for the providers with S_c classes such that the order is strict for any provider. By Lemma 9 the distributed deferred acceptance algorithm then runs in $2S_c - 1$ rounds. Constructing the virtual conflict graph requires 1 round of communication (recall the modelling assumptions from Section 2.3). The coloring algorithm, by definition, takes $O(T_c(n, \Delta_H))$ rounds. In total the running time is $O(T_c(n, \Delta) + S_c)$ rounds.

The algorithm computes a stable matching by the combination of Lemmas 9 and 10. It remains to argue that the algorithm is incentive-compatible for the clients. The first step, where the conflict graph H is formed, depends on the preferences of the providers, but not on the preferences of the clients. Therefore clients cannot affect the structure of H by lying. In the second step, the coloring algorithm only depends on the structure of H and (possibly) the unique identifiers of the clients. Since we assume that these are not part of the private input, the clients cannot affect the coloring by lying. It is known that the final step, the original deferred acceptance algorithm is incentive compatible for the clients [11, 32]. In the case of common preferences we have a simpler argument for this. By Lemma 11, the match of each client is its most preferred provider once the matches of higher priority clients have been removed. Conversely, since the match of the higher priority clients is not affected by the preferences of lower priority clients, the agents should always report their true preferences so that they get matched to their most preferred remaining provider. ◀

4 Generalisation to fractional loads and capacities

In this section we show that our algorithm can be generalised to the case where each client v has some load $\ell(v)$ and each provider u has some capacity $\ell(u)$, these have fractional values, and the goal is to find a stable matching such that for each provider, the total load assigned to it is at most its capacity. Note that this also captures the setting where clients have discrete load that is matched to the providers.

More formally, we call the following problem the *fractional capacitated stable matching* problem. The output is a fractional matching $m: E \rightarrow \mathbb{Q}$ that assigns a rational number to each edge. For agent v , let $m(v) = \sum_{e:v \in e} m(e)$. We have a capacity constraint $m(v) \leq \ell(v)$ for all agents. We say that agents $v \in C$, $u \in P$: $\{v, u\} \in E$ form a *blocking pair* if they would prefer to assign more load to the matching between them. More formally, if 1) there exist $v' \in N(u)$ such that $m(\{v', u\}) > 0$ and $v \succ_u v'$, or $m(u) < \ell(u)$ and 2) there exists $u' \in N(v)$ such that $m(\{v, u'\}) > 0$ and $u \succ_v u'$ or $m(v) < \ell(v)$.

With general preferences, it is not clear how utility should be defined: how does an agent prefer x units of load matched at preference y against $2x$ units of load matched at preference $y + 1$. A typical solution is to define a new utility function that fixes some relation between these two. Here we will show that the natural greedy modification of the deferred acceptance algorithm optimises both when one side has common preferences: the clients will maximise both the quality and quantity of their matching by being truthful.

4.1 Modified deferred acceptance algorithm

To account for capacities (and fractional loads), we modify the original deferred acceptance algorithm. We call this simple modification *batch proposals*. First, in each acknowledgement round, all providers send their current loads with the associated preference scores to all of their neighbours. Second, in each proposal round, each client proposes greedily based on the available capacity of its neighbours.

We use the basic algorithm template given in Section 3.1, but use the following, modified subroutine for deferred acceptance. For simplicity, we present the non-adaptive version of this algorithm. It has the same worst-case running time as the adaptive version presented in Section 3.1, but has slightly better message complexity.

Each client v holds a vector $r(v) = (r_1(v), r_2(v), \dots, r_d(v))$, consisting of the remaining capacities of its neighboring providers and where $d = \deg(v)$. On round 0, these are set to the initial capacities of the neighbours. For a given vector $r(v)$, the greedy batch proposal of v to the i th preferred provider is calculated as $p(v, i) = \min\{r_i(v), l(v) - \sum_{j=1}^{i-1} p(v, j)\}$. That is, go over the providers one by one in preference order and assign the maximum amount of load that is feasible.

We assume the clients v know their own preference class $s(v)$. If not, the providers can use one round to send this value to the clients. On odd rounds $2t - 1$ for $t = 1, \dots, S$:

1. For a client, if $s(v) = t$, for each neighbor u_i send them the batch proposal $p(v, i)$ and stop with $m(\{v, u_i\}) = p(v, i)$.
2. If a provider receives a non-zero proposal, accept it.

On even rounds $2t$ for $t = 1, \dots, S$:

1. Each provider, if it accepted a non-zero proposal in the previous round, sends its new remaining capacity to all neighbours.
2. Each client v updates $r(v)$ based on received messages.

► **Lemma 12.** *There exists a distributed mechanism in the CONGEST model for fractional capacitated stable matching when the providers have common preferences. The mechanism is incentive-compatible for the clients and the algorithm implementing it runs in $O(T_c(n, \Delta_H) + S_c)$ rounds.*

Analogous to Theorem 8, we have the following.

► **Theorem 13.** *There exists a deterministic distributed mechanism in the CONGEST model for fractional capacitated stable matching when the providers have common preferences. The mechanism is incentive-compatible for the clients and the algorithm implementing it runs in $\min\{O(S\Delta_H + \log^* n), O(n)\}$ rounds.*

While all fractional capacitated matchings are not comparable based on which an agent would prefer, we show that in the case of common preferences, we can define a natural partial order over fractional matchings the algorithm can compute such that agents will always prefer the matching produced when the agent is truthful. This implies that in the case of common preferences the deferred acceptance algorithm is incentive-compatible for the clients.

For a fixed client v , let $m(v, i)$ denote the load assigned by fractional matching m on the edge between v and its i th most preferred provider. We define that m is preferable to m' for client v if the following holds: for all $i \in \{1, \dots, \deg(v)\}$, we have that $\sum_{j=1}^i m(v, j) \geq \sum_{j=1}^i m'(v, j)$. That is, assignment m always assigns at least as much load at at least as high preference as m' .

To prove the truthfulness of the modified algorithm, we show that under the modified deferred acceptance algorithm, the solution computed under true preferences is preferable to the solution computed under any other reported preferences.

► **Lemma 14.** *For each client v and each i in $\{1, \dots, \deg(v)\}$, the modified deferred acceptance algorithm computes an assignment such that reporting the true preferences maximises $\sum_{j=1}^i m(v, j)$.*

This lemma means that by truthful reporting of preferences a client maximises the total matching assigned to any first i most preferred providers. In particular, it means that it also maximises the total weight of its matching. The proof is given in Appendix A.4

Proof of Lemma 12. The score function s and a c -coloring ϕ work the same way for the fractional version of the algorithm. Therefore, by Lemma 10, the score function $s \circ \phi$ provides a strict order for the providers with Sc preference classes. At each round $2t - 1$ for $t \in \{1, \dots, Sc\}$, the proposals of clients in preference class t make their proposals and stop. Given the one round of additional preprocessing, the modified deferred acceptance algorithm runs in $2Sc$ rounds. The coloring takes, by definition, $O(T_c(n, \Delta))$ rounds and the total running time is $O(T_c(n, \Delta) + S \cdot c)$.

Since the $s \circ \phi$ are strict, each provider receives at most one proposal in each round and these are always accepted. The resulting matching is stable, as the clients compute the best available matching with respect to existing matching. By Lemma 14 the algorithm is also incentive-compatible, as the matching computed given true preferences is always preferred by a client. ◀

Note that here we assumed that initially all capacities are values that can be sent in one round in the CONGEST model. The operations performed are addition and subtraction, and we assume that the values obtained this way can also be sent in the CONGEST model.

5 Fair tie-breaking

In this section we consider two *fair* ways of breaking ties in the deferred acceptance algorithm. This does not affect the truthfulness of the algorithm, but the goal is to provide an additional guarantee: despite having ties, from the perspective of each client, each way of breaking ties is (almost) equally probable. This implies that the algorithm not biased against any of the clients.

We define a tie-breaking rule as a function s' that assigns a new preference class to each client from some set S' such that each pair of clients that are assigned the same preference class by the original score function s and share an adjacent provider, are assigned a different class. Further, it must respect the original order, i.e. if $s(v) < s(u)$, then $s'(v) < s'(u)$. The tie-breaking rule given by a proper coloring of the conflict graphs, as described in Section 3.1 satisfies these properties.

Tie-breaking is uniformly random, if each legal tie-breaking is equally likely. A uniformly random coloring ϕ of the conflict graph H , where all clients are colored from a palette of the same size c , would give a uniformly random tie-breaking rule $s \circ \phi$, as described in Section 3.1. Computing a uniformly random coloring in the distributed setting is hard [15], so we will consider two different ways of computing almost random colorings.

5.1 Sampling an almost random coloring

The first approach is to use an existing sampling algorithm to find an almost random coloring. The state of the art distributed sampling algorithm is due to Carlson et al. [9].

► **Theorem 15** ([9], Theorem 1). *For all $\varepsilon > 0$, all $\Delta \geq 2$, all $\delta > 0$, and any $k > (11/6 + \varepsilon)\Delta$, for any graph $G = (V, E)$ of maximum degree Δ , a random k -coloring within total variation distance $\leq \delta$ from uniform can be generated in $O(\log(n/\delta))$ rounds.*

Given two distributions μ, ν over the same space Ω , the *total variation distance* is defined as $d(\mu, \nu) = \sum_{\sigma \in \Omega} |\mu(\sigma) - \nu(\sigma)|$. It is essentially the total absolute distance between the two distributions.

Fischer and Ghaffari present a simpler algorithm that finds a coloring with slightly more colors, but can also be used (in our context) in the CONGEST model.

► **Theorem 16** ([16], Theorem 1). *A uniform proper q -coloring of an n -node graph with maximum degree Δ can be sampled within total variation distance $\delta > 0$ in $O(\log(\frac{n}{\delta}))$ rounds, where $q = \alpha\Delta$ for any $\alpha > 2$.*

The disadvantage of using these algorithms compared to the algorithm we will present in Section 5.2 is that almost uniform sampling is slower when maximum degree is low: finding a coloring from a distribution with a total variation distance $\leq \delta$ to a random coloring requires $\Omega(\log(\frac{n}{\delta}))$ rounds [15].

Using the algorithm of Fischer and Ghaffari we get the following theorem.

► **Theorem 17**. *There exists a randomised distributed mechanism in the CONGEST model that computes a stable matching with high probability when the providers have common preferences. The mechanism is incentive-compatible for the clients, for any $\delta > 0$, the tie-breaking is within total variation distance δ from the uniformly random tie-breaking, and the algorithm implementing it runs in $O(\Delta_H S + O(\log(n/\delta)))$ rounds.*

We note that this theorem has an analogous variant for the fractional stable matching as well, as the computation of the fractional matching is independent of how the tie-breaking rule was obtained.

Proof. The running time and the incentive-compatibility follow from Lemma 7 and Theorem 16 by using $c = 2\Delta + 1$. Similarly, the bound on the total variation distance of the coloring implies directly a corresponding bound on the total variation distance of the tie-breaking rule, as each order is produced by the same number of colorings.

The algorithm runs in the CONGEST model, even if the matching graph is the communication graph. Initially each node v starts with some color $\phi(v)$. In each round, each node samples a new color $\phi'(v) \in [c]$ uniformly at random. Then it sends its current color $\phi(v)$ and new color $\phi'(v)$ to each of its neighbours. If for all neighbors u it holds that $\phi'(v) \notin \{\phi(u), \phi'(u)\}$, then v sets $\phi'(v)$ as its new color. Otherwise there is a potential conflict, and it keeps its old color. To implement this in the CONGEST model, the providers function as intermediaries: for each neighbor u , each provider v checks whether there is a conflict with any of u 's neighbours in H through v and reports this back to u . If no provider of u reports a conflict, u assumes new color $\phi'(u)$. ◀

5.2 Fast random coloring with failures

The second approach is to sample a coloring once and to fix any improperly colored parts arbitrarily. The sampling takes constant time, but the bias for the agents that fail to be properly colored is arbitrarily bad. Additionally the error probability is inversely related to the number of colors, and thus requiring high fairness slows down the matching step. We prove the following lemma.

► **Lemma 18.** *For any $0 < \delta < 1$ there exists a randomised distributed coloring algorithm with the following properties:*

- (P1) *The running time is $O(\sqrt{\Delta \log \Delta} + \log^* n)$ rounds.*
- (P2) *There is a subset $F \subseteq V$ such that each node is in F with probability at least $1 - \delta$ (and thus the expected size of F is $(1 - \delta)n$) and F is colored u.a.r. with $\lceil \delta^{-1} \rceil \Delta$ colors.*
- (P3) *Remaining nodes are colored with $\Delta + 1$ colors from $\{\lceil \delta^{-1} \rceil \Delta + 1, \dots, (\lceil \delta^{-1} \rceil + 1)\Delta + 2\}$.*

Recall that in Section 3 the algorithm requires that the conflict graph of clients is colored. In what follows, let $H = (V, E)$ denote the conflict graph and let Δ denote the maximum degree of the conflict graph. As input, the algorithm takes parameter δ : this is the target probability for failure at a single node. Our first coloring algorithm consists of two steps:

1. Each node v picks a color $\phi(v)$ uniformly at random from $\{1, \dots, \lceil \delta^{-1} \rceil \Delta\}$. Here assume that $\lceil \delta^{-1} \rceil \Delta \geq \Delta + 1$.
2. Let X denote the set of nodes with coloring conflicts: if $v \in X$ then there exists u such that $\phi(v) = \phi(u)$. Color nodes in X with $\Delta + 1$ colors from the palette $\{\lceil \delta^{-1} \rceil \Delta + 1, \dots, (\lceil \delta^{-1} \rceil + 1)\Delta + 2\}$.

The idea behind this simple algorithm is that all nodes that are either properly colored in step 1 are colored completely at random. The fraction of nodes colored properly in step 1 is at least $(1 - \delta)$. The remaining nodes are colored in a way that does not affect the matching computed by nodes colored in step 1. The proof of Lemma 18 is given in Appendix A.5.

Using this algorithm for the template given in Lemma 7 we get the following. The tie-breaking rule guarantees that there is a large set of agents that have fair tie-breaking in the sense that 1) the tie-breaking is completely random among them, and 2) the remaining agents have worse tie-breaking. The trade-off is that the size of the fair set affects the number of colors and thus the running time of the deferred acceptance algorithm.

► **Theorem 19.** *For any $\delta > 0$, there exists a randomised distributed mechanism in the LOCAL model for stable matching when the providers have common preferences. The mechanism is incentive-compatible for the clients, there exists a subset F of clients of expected size $(1 - \delta)n$ such that the tie-breaking is uniformly random when restricted to F , and $C \setminus F$ has arbitrary tie-breaking at lower preference than F . The running time of the algorithm is $O(\delta^{-1} S \Delta_H + \log^* n)$.*

Proof. The properties of the algorithm follow from Lemma 7 and Lemma 18. ◀

Again, this theorem has an analogous variant for the fractional stable matching as the modification to the deferred acceptance algorithm is unaffected.

References

- 1 Ittai Abraham, Danny Dolev, and Joseph Y. Halpern. Distributed protocols for leader election: A game-theoretic perspective. *ACM Trans. Economics and Comput.*, 7(1):4:1–4:26, 2019. doi:10.1145/3303712.
- 2 Yehuda Afek, Yehonatan Ginzberg, Shir Landau Feibish, and Moshe Sulamy. Distributed computing building blocks for rational agents. In Magnús M. Halldórsson and Shlomi Dolev, editors, *Proc. 2014 ACM Symposium on Principles of Distributed Computing (PODC 2014)*, pages 406–415. ACM, 2014. doi:10.1145/2611462.2611481.
- 3 Steve Alpern and Diane Reyniers. Strategic mating with common preferences. *Journal of Theoretical Biology*, 237(4):337–354, 2005. doi:10.1016/j.jtbi.2003.09.021.
- 4 Nir Amira, Ran Giladi, and Zvi Lotker. Distributed weighted stable marriage problem. In *Proc. 17th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2010)*, pages 29–40, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg. doi:10.1007/978-3-642-13284-1_4.
- 5 Alkida Balliu, Sebastian Brandt, Juho Hirvonen, Dennis Olivetti, Mikaël Rabie, and Jukka Suomela. Lower bounds for maximal matchings and maximal independent sets. *J. ACM*, 68(5):39:1–39:30, 2021. doi:10.1145/3461458.
- 6 Leonid Barenboim, Michael Elkin, and Uri Goldenberg. Locally-iterative distributed ($\Delta + 1$)-coloring and applications. *J. ACM*, 69(1):5:1–5:26, 2022. doi:10.1145/3486625.
- 7 Ken Burdett and Melvyn G. Coles. Marriage and class. *The Quarterly Journal of Economics*, 112(1):141–168, 1997. URL: <http://www.jstor.org/stable/2951279>.
- 8 Ioannis Caragiannis, Aris Filos-Ratsikas, Panagiotis Kanellopoulos, and Rohit Vaish. Stable fractional matchings. In *Proc. 2019 ACM Conference on Economics and Computation (EC 2019)*, pages 21–39. ACM, 2019. doi:10.1145/3328526.3329637.
- 9 Charlie Carlson, Daniel Frishberg, and Eric Vigoda. Improved Distributed Algorithms for Random Colorings. In *Proc. 27th International Conference on Principles of Distributed Systems (OPODIS 2023)*, volume 286 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 13:1–13:18, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.OPODIS.2023.13.
- 10 Simon Collet, Pierre Fraigniaud, and Paolo Penna. Equilibria of games in networks for local tasks. In *Proc. 22nd International Conference on Principles of Distributed Systems (OPODIS 2018)*, volume 125 of *LIPIcs*, pages 6:1–6:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.OPODIS.2018.6.
- 11 Lester E. Dubins and David A. Freedman. Machiavelli and the gale-shapley algorithm. *The American Mathematical Monthly*, 88(7):485–494, 1981. doi:10.2307/2321753.
- 12 Kimmo Eriksson and Olle Häggström. Instability of matchings in decentralized markets with various preference structures. *Int. J. Game Theory*, 36(3-4):409–420, 2008. doi:10.1007/S00182-007-0081-6.
- 13 Joan Feigenbaum, Christos H. Papadimitriou, and Scott Shenker. Sharing the cost of multicast transmissions. *J. Comput. Syst. Sci.*, 63(1):21–41, 2001. doi:10.1006/JCSS.2001.1754.
- 14 Joan Feigenbaum, Michael Schapira, and Scott Shenker. *Distributed Algorithmic Mechanism Design*, pages 363–384. Cambridge University Press, 2007.
- 15 Weiming Feng, Yuxin Sun, and Yitong Yin. What can be sampled locally? In *Proc. 2017 ACM Symposium on Principles of Distributed Computing (PODC 2017)*, pages 121–130, 2017. doi:10.1145/3087801.3087815.
- 16 Manuela Fischer and Mohsen Ghaffari. A simple parallel and distributed sampling technique: Local glauber dynamics. In *Proc. 32nd International Symposium on Distributed Computing (DISC 2018)*, volume 121 of *LIPIcs*, pages 26:1–26:11. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPIcs.DISC.2018.26.
- 17 Patrik Floréen, Petteri Kaski, Valentin Polishchuk, and Jukka Suomela. Almost stable matchings by truncating the gale-shapley algorithm. *Algorithmica*, 58(1):102–118, 2010. doi:10.1007/s00453-009-9353-9.

- 18 David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962. doi:10.2307/2312726.
- 19 Zihan Gao, Wanming Hao, and Shouyi Yang. Joint offloading and resource allocation for multi-user multi-edge collaborative computing system. *IEEE Trans. Veh. Technol.*, 71(3):3383–3388, 2022. doi:10.1109/TVT.2021.3139843.
- 20 Avinatan Hassidim, Yishay Mansour, and Shai Vardi. Local computation mechanism design. *ACM Trans. Economics and Comput.*, 4(4):21:1–21:24, 2016. doi:10.1145/2956584.
- 21 Robert W. Irving, David F. Manlove, and Sandy Scott. The stable marriage problem with master preference lists. *Discrete Applied Mathematics*, 156(15):2959–2977, 2008. doi:j.dam.2008.01.002.
- 22 Naoyuki Kamiyama. Stable matchings with ties, master preference lists, and matroid constraints. In Martin Hoefer, editor, *Algorithmic Game Theory (SAGT 2015)*, pages 3–14, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg. doi:10.1007/978-3-662-48433-3_1.
- 23 Naoyuki Kamiyama. Many-to-many stable matchings with ties, master preference lists, and matroid constraints. In *Proc. 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019)*, pages 583–591, 2019. URL: <http://dl.acm.org/citation.cfm?id=3331743>, doi:10.5555/3306127.3331743.
- 24 Pankaj Khanchandani and Roger Wattenhofer. Distributed stable matching with similar preference lists. In *Proc. 20th International Conference on Principles of Distributed Systems, (OPODIS 2016)*, volume 70 of *LIPICs*, pages 12:1–12:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2016. doi:10.4230/LIPICs.OPODIS.2016.12.
- 25 Donald Ervin Knuth. *Stable marriage and its relation to other combinatorial problems: An introduction to the mathematical analysis of algorithms*, volume 10. American Mathematical Soc., 1997.
- 26 Nathan Linial. Locality in Distributed Graph Algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992. doi:10.1137/0221015.
- 27 Yannic Maus and Tigran Tonoyan. Local conflict coloring revisited: Linial for lists. In Hagit Attiya, editor, *Proc. 34th International Symposium on Distributed Computing (DISC 2020)*, volume 179 of *LIPICs*, pages 16:1–16:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPICs.DISC.2020.16.
- 28 Economic Sciences Prize Committee of the Royal Swedish Academy of Sciences. Scientific background on the sveriges riksbank prize in economic sciences in memory of alfred nobel 2012: Stable allocations and the practice of market design. <https://www.kva.se/app/uploads/2012/10/globalassets-priser-ekonomi-2012-scibackeken12.pdf>, 2012. Accessed:15-02-2024.
- 29 Gregg O’Malley. *Algorithmic aspects of stable matching*. PhD thesis, University of Glasgow, 2007.
- 30 Rafail Ostrovsky and Will Rosenbaum. Fast distributed almost stable matchings. In *Proc. 2015 ACM Symposium on Principles of Distributed Computing (PODC 2015)*, pages 101–108. ACM, 2015. doi:10.1145/2767386.2767424.
- 31 David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. Society for Industrial and Applied Mathematics, 2000. doi:10.1137/1.9780898719772.
- 32 Alvin E. Roth. The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7(4):617–628, 1982. doi:10.1287/moor.7.4.617.
- 33 Alvin 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. doi:10.1086/261272.
- 34 Alvin E. Roth. Deferred acceptance algorithms: history, theory, practice, and open questions. *Int. J. Game Theory*, 36(3-4):537–569, 2008. doi:10.1007/S00182-008-0117-6.
- 35 Alvin E. Roth. What have we learned from market design? *The Economic Journal*, 118(527):285–310, 2008. doi:10.1111/j.1468-0297.2007.02121.x.
- 36 Alvin E. Roth, Uriel G. Rothblum, and John H. Vande Vate. Stable matchings, optimal assignments, and linear programming. *Mathematics of Operations Research*, 18(4):803–828, 1993. URL: <http://www.jstor.org/stable/3690124>, doi:10.1287/MOOR.18.4.803.

- 37 Sandy Scott. *A Study Of Stable Marriage Problems With Ties*. PhD thesis, University of Glasgow, 2005.
- 38 John H. Vande Vate. Linear programming brings marital bliss. *Operations Research Letters*, 8(3):147–153, 1989. doi:10.1016/0167-6377(89)90041-2.

A Additional details

A.1 Proof of Theorem 6




Figure 3 Two instances of stable matching that only differ in the preferences of agent v on the left. For each client (bottom, red) their preferences are indicated below them in the order of the adjacent providers and their preference class on their right. The change causes the unique stable matching to change, and therefore agent u on the right has to learn the preferences of v in order to compute a stable matching.

We show that even when $\Delta = 2$, providers have common preferences, and there are no ties among the neighbours of any provider, finding a stable matching requires $\Omega(S)$ rounds in the LOCAL model. The construction is illustrated in Figure 3.

Fix any S . We can construct a pair of matching graphs on $n = 2S$ nodes such that a node must see distance $\Omega(S)$ to decide its output. Denote by $C = \{c_1, c_2, \dots, c_S\}$ the set of clients and by $P = \{p_1, p_2, \dots, p_S\}$ the set of providers. Each client c_i is connected to providers p_i and p_{i+1} (except c_S is connected only to p_S). The preference class of each c_i is i . In the first matching graph each client c_i has the preference relation $p_i > p_{i+1}$. In the second graph we have the same preference relation, except for c_1 these are flipped: $p_2 > p_1$.

Clearly client c_S has to see receive information from distance $\Omega(S)$ to see the difference between these two instances. In the first instance the most preferable client c_1 prefers p_1 and therefore the only stable matching consists of edges $\{c_i, p_i\}$ for $i \in \{1, \dots, S\}$. In the second instance c_1 prefers p_2 , and therefore $\{c_1, p_2\}$ is in any stable matching. This means that c_2 is matched to p_3 , and so on: the unique stable matching consists of edges $\{c_i, p_{i+1}\}$ for $i \in \{1, \dots, S - 1\}$.

A.2 Similar preferences

In this section we show that the similarity of preference lists does not help in computing a stable matching.

► **Theorem 20.** *Assume that all pairs of providers that share a neighbor share at least δ neighbours and agree on their relative preference order. Computing a stable matching requires $\Omega(n/\delta)$ rounds in the LOCAL model.*

Proof. The construction is a simple modification of the lower bound from Theorem 5. Fix any odd $k = 2l + 1$. We construct a pair of matching graphs G and G' as follows. Denote by $C = \{c_1, c_2, \dots, c_{l+1}\}$ the set of clients and by $P = \{p_1, p_2, \dots, p_l\}$ the set of providers. Each provider p_i is connected to clients c_i and c_{i+1} . Additionally, for each pair p_i, p_{i+1} create a set C_i of $\delta - 1$ additional clients connected to both p_i and p_{i+1} . In both instances the preferences of the clients in C are such that client c_i prefers p_i over p_{i+1} . In the first instance provider p_1 prefers c_1 over c_2 , and in the second instance these are flipped. All other providers p_i prefer c_i over c_{i+1} . Providers always prefer clients in C over the clients in the sets C_i .

It is easy to see that a similar argument as in the proof of Theorem 5 applies. The providers ignore clients in sets C_i when computing the stable matching. The unique stable matching in the first instance consists of edges $\{c_i, p_i\}$ and in the second instance of edges $\{p_i, c_{i+1}\}$. To decide whether it is matched, client c_{k+1} must see $\Omega(k)$ hops away. Since for each provider there are δ clients, we have that $k = \Omega(n/\delta)$. ◀

A.3 Coloring in the CONGEST model

In Section 2 we assumed that the communication network has edges between clients that are connected to the same provider (specifically we need edges between clients who are also in the same preference class under s). Removing this assumption has two effects: first, simulating the coloring algorithm in the conflict graph slows down by a factor of 2, as the messages have to be relayed through the providers, and second, providers form bottlenecks since, in the worst case, all connected clients want to send a message to each other connected client. This causes a potential slowdown by a factor of Δ .

First, it should be noted that in this case it is impossible to form the explicit conflict graph in constant time. However, since we need $\Omega(\Delta_H)$ time to run the deferred acceptance algorithm, we can spend $O(\Delta_P) = O(\Delta_H)$ rounds to do this.

The most suitable coloring algorithm depends on the maximum degree of the conflict graph. When $\Delta_H = \Omega(\log n)$, we can use a simple variant of Luby's coloring algorithm to color with 2Δ colors. In each round, each uncolored node tries to pick a free color uniformly at random. Since there are 2Δ possible colors, the probability to choose a color that is not picked by others is at least $1/2$. Providers can be used to verify whether a trial was successful: each client sends its color to each provider and each provider indicates whether a neighbor in the conflict graph chose the same color. If no provider indicates a conflict, the color can be chosen safely. This algorithm runs in $O(\log n)$ rounds with high probability. We note that greedy color reduction does not work in this setting: simulating one round requires sending $O(\Delta_P)$ colors.

For fast $(\Delta + 1)$ -coloring, we can use the additive-group coloring algorithm of Barenboim and Elkin [6] combined with the $O(\Delta^2)$ -coloring algorithm of Linial [26]. For sufficiently small values of $\Delta_H = O(\log / \log \log n)$ there is no overhead (as again, we must spend $\Omega(\Delta_H)$ rounds to compute the stable matching). For Δ_H larger than this, there is a slowdown in the algorithm, at least without further optimisation.

A.4 Proof of Lemma 14

Proof of Lemma 14. By Lemma 10, the score function $s \circ \phi$ provides a strict linear order for each provider. Therefore when client v makes its proposal, the adjacent providers receive no other proposals. At the beginning of each round $2t - 1$ for $t = 1, \dots, S$, the clients (in preference class t) know the available capacities $r_t(v)$ of connected providers. Denote the capacity of the i th preferred provider u by $r_t(v, i)$.

By round $2t - 1$, all clients in preference classes $i < t$ have finalised their matches. It is important to note that a client cannot influence its preference class $(s \circ \phi)(v)$, determined by the coloring algorithm, nor can it affect the resources assigned to higher priority classes. Further, since the lower priority clients do not replace any of the load assigned in round t , the best v can do is to optimise its matching given the current remaining capacities $r_t(v)$.

Now let $>_v$ denote the true preferences of v and let $>'_v$ denote some different preference order. Let m and m' denote the assignments computed under $>_v$ and $>'_v$, respectively (with other preferences remaining constant). For each provider in $N(v)$, there are exactly two possibilities: either the provider becomes saturated (v assigns $m(\{v, u_i\}) = r_t(v, i)$), or v assigns its remaining load (and therefore assigns 0 load to remaining providers).

Now consider the matchings m and m' . Again, there are two cases: either all providers become saturated, or not. In the first case, the assignments m and m' agree, and the claim holds trivially. In the second case, m saturates some first k most preferred providers, and then assigns some load to the $k + 1$ th most preferred provider. This directly implies that for all $i \in \{1, \dots, k\}$, we have that $\sum_{j=1}^i m(v, i) \geq \sum_{j=1}^i m'(v, i)$ (as m saturates these edges, but m' may assign partial or no load on some). However, since $\sum_{j=1}^{k+1} m(v, i) = \ell(v)$ (it assigns all of its load), we also must have $\sum_{j=1}^{k+1} m(v, i) \geq \sum_{j=1}^{k+1} m'(v, i)$. ◀

A.5 Proof of Lemma 18

Proof of Lemma 18. The running time of the algorithm is dominated by step 2, as step 1 requires no communication (we assume that the value Δ or an upper bound for it is common knowledge). Computing a $(\Delta + 1)$ -coloring of the nodes that fail in step 1 can be computed in $O(\sqrt{\Delta \log \Delta} + \log^* n)$ rounds [27]. This establishes P1.

Next, we establish property 2. In step 1, each node v picks a random color from $\{1, \dots, \lceil \delta^{-1} \rceil \Delta\}$. For each neighbor, the probability of a conflict is $1/(\lceil \delta^{-1} \rceil \Delta)$. By a union bound, the total probability of a conflict over at most Δ neighbors is bounded by $1/\lceil \delta^{-1} \rceil \leq \delta$.

The coloring is uniformly random for nodes in $F = V \setminus X$, as conditioned on having no conflicts, each color is equally likely. This establishes P2.

Finally, P3 follows from the fact that subgraph induced by X has maximum degree Δ and can therefore be colored with $\Delta + 1$ colors. ◀