

# Fast Distributed Almost Stable Matchings

Rafail Ostrovsky  
University of California, Los Angeles  
Departments of Computer Science and  
Mathematics  
4732 Boelter Hall  
Los Angeles, CA 90095  
rafail@cs.ucla.edu

Will Rosenbaum  
University of California, Los Angeles  
Department of Mathematics  
520 Portola Plaza  
Los Angeles, CA 90095-1555  
wrosenbaum@math.ucla.edu

## ABSTRACT

In their seminal work on the Stable Marriage Problem, Gale and Shapley [4] describe an algorithm which finds a stable matching in  $O(n^2)$  communication rounds. Their algorithm has a natural interpretation as a distributed algorithm where each player is represented by a single processor. In this distributed model, Floréen, Kaski, Polishchuk, and Suomela [3] recently showed that for bounded preference lists, terminating the Gale-Shapley algorithm after a constant number of rounds results in an *almost* stable matching. In this paper, we describe a new deterministic distributed algorithm which finds an almost stable matching in  $O(\log^5 n)$  communication rounds for *arbitrary* preferences. We also present a faster randomized variant which requires  $O(\log^2 n)$  rounds. This run-time can be improved to  $O(1)$  rounds for “almost regular” (and in particular complete) preferences. To our knowledge, these are the first sub-polynomial round distributed algorithms for any variant of the stable marriage problem with unbounded preferences.

## Categories and Subject Descriptors

C.2.4 [Computer-Communication Networks]: Distributed Systems—*Distributed applications*; F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*Computations on discrete structures*; G.2.1 [Discrete Mathematics]: Combinatorics—*Combinatorial algorithms*; J.4 [Social and Behavioral Sciences]: Economics

## Keywords

stable marriage problem; stable matchings; distributed algorithms; approximation algorithms

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

PODC'15, July 21–23, 2015, Donostia-San Sebastián, Spain.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3617-8 /15/07 ...\$15.00.

<http://dx.doi.org/10.1145/2767386.2767424>.

## 1. INTRODUCTION

### 1.1 Historical Background

In their seminal work, Gale and Shapley [4] consider the following problem. Members of disjoint sets of  $n$  *men* and  $n$  *women* each rank all members of the opposite sex. The men and women (which we collectively call *players*) wish to form a *matching*—a one-to-one correspondence between the men and women—which is *stable* in the sense that it contains no *blocking pairs*: pairs of players who mutually prefer each other to their assigned partners in the matching. Gale and Shapley showed that a stable matching always exists by giving an explicit algorithm for finding one. The centralized Gale-Shapley algorithm runs in time  $\tilde{O}(n^2)$ , and this run-time is asymptotically optimal for centralized algorithms [5]. The Gale-Shapley algorithm easily generalizes to the case of *incomplete preferences* where each player ranks only a subset of the members of the opposite sex [5].

The Gale-Shapley algorithm has a natural interpretation as a distributed algorithm, where each player is represented by a separate processor which privately holds that player’s preferences. The communication links between players are formed by pairs of players who appear on each other’s preference lists. This model is natural in, for example, social networks where players may be constrained to be matched with acquaintances and do not communicate with strangers. In this model, the input to each processor has size  $\tilde{O}(n)$ , yet there is still no known distributed algorithm which improves upon the Gale-Shapley algorithm’s  $\tilde{O}(n^2)$  run-time for arbitrary preferences.<sup>1</sup>

Recently, there has been interest in approximate versions of the stable marriage problem [1, 2, 3, 7, 9], where the goal is to find a matching which is “almost stable.” There is no consensus in the literature on precisely how to measure almost stability, but typically almost stability requires that a matching induces relatively few blocking pairs. Eriksson and Häggström [2] argue that, “the proportion of blocking pairs among all possible pairs is usually the best measure of instability.” Using a finer notion of almost stability, Floréen, Kaski, Polishchuk, and Suomela show [3] that for *bounded* preference lists, truncating the Gale-Shapley algorithm after

<sup>1</sup>In the distributed computational model with complete preferences, each player can broadcast their preferences to all other players in  $O(n)$  rounds, after which each player runs a centralized version of the Gale-Shapley algorithm. While this process requires only  $O(n)$  communication rounds, the synchronous distributed run-time is still  $\tilde{O}(n^2)$  in the worst case.

boundedly many communication rounds yields a matching that induces at most  $\varepsilon|M|$  blocking pairs. Here  $|M|$  is the size of the matching produced. More recently, Hassidim, Mansour and Vardi [7] show a similar result in a more restrictive “local” computational model, so long as the men’s preferences are chosen uniformly at random.

Kipnis and Patt-Shamir [9] give an algorithm which finds an almost stable matching using  $O(n)$  communication rounds in the worst case, using a finer notion of approximate stability than we consider. Specifically, in their notion of almost-stability, a matching is almost stable if no pair of players can both improve their match by more than an  $\varepsilon$ -fraction of their preference list by deviating from their assigned partners. They also prove an  $\Omega(\sqrt{n}/\log n)$  communication round lower bound for finding an almost stable matching for this notion of approximation.

## 1.2 Overview of Results

We consider an approximate version of the stable marriage problem where an almost stable matching is allowed to have  $\varepsilon|E|$  blocking pairs. Here  $E$  is the set of pairs of men and women who rank one another (that is, the set of edges in the communication graph). Our notion of approximation, which generalizes almost stability as described in [2], is strictly coarser than those used in [3] and [9]. However, for bounded preferences (the context of [3]), our notion of instability agrees with that of [3] up to a constant factor.

Using the notion of almost stability given above, we describe a deterministic distributed algorithm, **ASM**, which produces an almost stable matching in  $O(\log^5(n))$  rounds. We note that in order to obtain this sub-polynomial runtime, we cannot use a finer notion of approximation than [9], who prove an  $\Omega(\sqrt{n}/\log n)$  round lower bound for their model. We remark that after removing an arbitrarily small fraction of “bad” players, the output of **ASM** is almost stable in the sense of [9] as well. We further describe a faster randomized variant of **ASM** which runs in  $O(\log^2(n))$  rounds. For preferences which are “almost regular,” (and in particular for complete or bounded preferences) this run-time can be improved to  $O(1)$ .

**THEOREM 1.** *There exists a deterministic distributed algorithm **ASM** which produces a  $(1 - \varepsilon)$ -stable matching in  $O(\log^5(n))$  communication rounds. A randomized variant of the algorithm, **RandASM** runs in  $O(\log^2(n))$  rounds for general preferences, and can be improved to  $O(1)$  rounds for almost regular (and in particular complete or bounded) preferences.*

**ASM** can be viewed as a generalization of the classical Gale-Shapley algorithm [4] which allows for multiple simultaneous proposals by the men and acceptances by the women. In **ASM**, the players quantize their preferences into  $O(\varepsilon^{-1})$  quantiles of equal size. In each step of the algorithm, the men propose to all women in their best nonempty quantile. Each woman accepts proposals only from her best quantile receiving proposals. A maximal matching is then found among the accepted proposals, and matched women reject men they do not prefer to their matches. This procedure is iterated until a large fraction of men are either matched or have been rejected by all women.

The analysis of our algorithm follows in two steps. We first show that by quantizing preferences, the matching found by **ASM** cannot contain a large fraction of blocking pairs

among the matched (or rejected) players. We bound the number of blocking pairs from the remaining “bad” players by showing there are few such players, and that only a small fraction can participate in many blocking pairs.

The remainder of the paper is organized as follows. In Section 2 we formalize our notion of almost stable matchings and our computational model. We also overview methods of computing maximal matchings which our algorithm will require as subroutines. Section 3 describes **ASM** and its subroutines in detail and states basic guarantees for the subroutines. Section 4 proves the performance guarantees for **ASM**. Finally, in Section 5 we describe the randomized variants of **ASM**.

## 2. PRELIMINARIES

### 2.1 Stable and almost stable matchings

We consider the stable marriage problem as originally described by Gale and Shapley [4] with incomplete preferences or, equivalently, unacceptable partners (cf. [5, 10]). Let  $X$  and  $Y$  be sets of women and men, respectively. For simplicity, we assume  $|X| = |Y| = n$ . Each player  $v \in X \cup Y$  holds a *preference list* or *ranking*  $P^v$ —a linear order on a subset of the members of the opposite sex. We denote the set of all player’s preferences by  $\mathcal{P} = \{P^v | v \in X \cup Y\}$ . We refer to the players  $u$  that appear on  $v$ ’s preference list  $P^v$  as  $v$ ’s *acceptable partners*. We call  $\mathcal{P}$  *complete* if each player ranks all players of the opposite sex. If a man  $m$  precedes  $m'$  on woman  $w$ ’s preference list, we write  $m \succ_w m'$ , and we say that  $w$  *prefers*  $m$  to  $m'$ . For simplicity, we assume that preferences are symmetric in the sense that if  $m$  appears in  $P^w$ , then  $w$  appears in  $P^m$ . Given players  $v$  and  $u$  of opposite genders, we let  $P^v(u)$  denote  $v$ ’s *rank* of  $u$ . For example,  $P^v(u) = 1$  means that  $u$  is  $v$ ’s most favored partner, et cetera.

We define the *communication graph*  $G = (V, E)$  for a set of preferences  $\mathcal{P}$  to be

$$V = Y \times X, \quad E = \{(m, w) | m \in P^w, w \in P^m\}.$$

For a communication graph  $G = (V, E)$ , we denote the degree of  $v \in V$  by  $\deg v$ , which is the number of players that appear on  $v$ ’s preference list.

A *matching*  $M \subseteq E$  is a set of edges in  $E$  such that no two edges share a vertex. Given a matching  $M$  and  $(m, w) \in M$ , we call  $m$  and  $w$  *partners* and write  $p(w) = m$  and  $p(m) = w$ . Given preferences and a matching  $M$ , we say that an edge  $(m, w) \in E$  is a *blocking pair* if  $(m, w) \notin M$ , but  $m$  and  $w$  mutually prefer each other to their partners in  $M$ ; that is,

$$m \succ_w p(w) \quad \text{and} \quad w \succ_m p(m).$$

By convention, we assume each unmatched player ( $p(v) = \emptyset$ ) prefers all acceptable partners to being without a partner. A *stable matching* is a matching which contains no blocking pairs.

We are primarily concerned with finding matchings which are “almost stable” in the sense that they induce relatively few blocking pairs. We use a definition of almost stability given by Eriksson and Häggström [2], modified to allow for incomplete preference lists.

*Definition 1.* Given  $\varepsilon \geq 0$  and preferences  $\mathcal{P}$ , we say that a matching  $M$  is  $(1 - \varepsilon)$ -*stable* with respect to  $\mathcal{P}$  if  $M$  induces at most  $\varepsilon|E|$  blocking pairs with respect to  $\mathcal{P}$ .

We refer to the problem of finding a matching which is  $(1 - \varepsilon)$ -stable for fixed  $\varepsilon > 0$  as the *almost stable matching problem*. Note that a 1-stable matching corresponds precisely to the classical stable matching definition.

*Remark 1.* We reiterate that there is no consensus in the literature on the precise definition of almost stability. For example, the authors of [3] compare the number of blocking pairs to  $|M|$ , the size of the matching rather than  $|E|$ , as we do. Since in [3], they only consider bounded preference lists, their notion of almost stability agrees with Definition 1 up to a constant factor.

The following definition is due to Kipnis and Patt-Shamir [9].

*Definition 2.* Given  $\varepsilon \geq 0$ , preferences  $\mathcal{P}$ , and a matching  $M$  we call an edge  $(m, w) \in E$   $\varepsilon$ -**blocking** if  $m$  and  $w$  appear an  $\varepsilon$ -fraction higher on each other's preferences than their assigned partners. Specifically,  $(m, w)$  is  $\varepsilon$ -blocking if

$$P^m(p(m)) - P^m(w) \geq \varepsilon \deg m$$

and

$$P^w(p(w)) - P^w(m) \geq \varepsilon \deg w.$$

We say that  $M$  is  $\varepsilon$ -**blocking-stable** if it contains no  $\varepsilon$ -blocking pairs.

*Remark 2.* Kipnis and Patt-Shamir [9] prove an  $\Omega(\sqrt{n}/\log n)$  round lower bound for finding an  $\varepsilon$ -blocking-stable matching. That we are able to achieve a polylogarithmic round algorithm for the almost stable matching problem using Definition 1 bolsters the use of Definition 1 for almost stability, at least for practical applications. Further, **ASM** produces a matching which is nearly  $\varepsilon$ -blocking-stable in the sense that after the removal of an arbitrarily small fraction of “bad” men, the resulting matching is  $\varepsilon$ -blocking-stable with respect to the remaining players.

## 2.2 Computational model

We describe **ASM** in terms of the CONGEST model formalized by Peleg [11]. In this distributed computational model, each player  $v \in X \cup Y$  represents a processor. Given preferences  $\mathcal{P}$ , the communication links between the players are given by the set of edges  $E$  in the communication graph  $G$ . Communication is performed in synchronous rounds. Each communication round occurs in three stages: first, each processor receives messages (if any) sent from its neighbors in  $G$  during the previous round. Next, each processor performs local calculations based on its internal state and any received messages. We make no restrictions on the complexity of local computations. Finally, each processor sends short ( $O(\log n)$  bit) messages to its neighbors in  $G$ —the processor may send distinct messages to distinct neighbors. In the CONGEST model, complexity is measured by the number of communication rounds needed to solve a problem.

*Remark 3.* Although the CONGEST model allows for unbounded local computation during each round, the computations required by **ASM** can be implemented in linear or near-linear time in each processor's input.

## 2.3 Maximal matchings

As a subroutine, **ASM** requires a method for computing maximal matchings in a graph.

*Definition 3.* A matching  $M$  is a *maximal matching* if it is not properly contained in any larger matching. Equivalently,  $M$  is maximal if and only if every  $v \in V$  satisfies precisely one of the following conditions:

1. there exists a unique  $u \in V$  with  $(v, u) \in M$ ;
2. for all  $u \in N(v)$  there exists  $v' \in V$  with  $v' \neq v$  such that  $(v', u) \in M$ .

For the deterministic version of our algorithm, we invoke the work of Hańćkowiak, Karoński, and Panconesi [6] who give a deterministic distributed algorithm which finds a maximal matching in a polylogarithmic number of rounds.

**THEOREM 2** (HAŃĆKOWIAK, ET AL.[6]). *There exists a deterministic distributed algorithm, **MaximalMatching**, which finds a maximal matching in a communication graph  $G = (V, E)$  in  $O(\log^4(n))$  rounds, where  $n = |V|$ .*

We remark that while the authors of [6] do not explicitly use the CONGEST model of computation, their algorithm can easily be implemented in this model.

The randomized variants of **ASM** require faster (randomized) subroutines for computing maximal and “almost maximal” matchings in a communication graph. In Appendix A, we describe how to modify an algorithm of Israeli and Itai [8] to give the necessary results.

## 3. DETERMINISTIC ALGORITHM DESCRIPTION

In this section, we describe in detail the almost stable matching algorithm, **ASM**. The main algorithm invokes the subroutine **QuantileMatch** which in turn calls **ProposalRound**. In Section 3.1 we introduce notation, and describe the internal state of each processor during the execution of **ASM**. Section 3.2 contains a description of the **ProposalRound** subroutine, while Section 3.3 describes the **QuantileMatch** subroutine. Finally, Section 3.4 describes **ASM**.

### 3.1 The state of a processor

In our algorithm, we assume that each player is represented by an independent processor. Each processor has a unique id and a gender (male or female) both of which are known to that processor. The only global information known to each processor is (an upper bound on)  $n$ , the total number of processors. At each step of the algorithm, we specify the state of each processor as well as any messages the processor might send or receive. Throughout,  $k$  is a parameter (the number of quantiles) to be chosen later. The state of a player  $v$  consists of:

- Quantized preferences  $Q_1, Q_2, \dots, Q_k$  where we denote  $Q = \bigcup Q_i$ . For each  $u$  in  $v$ 's preference list, denote  $q(u) = \lceil P(u)/k \rceil$ . Initially we set  $Q_i = \{u | q(u) = i\}$ , so that  $Q_1$  is the set of  $v$ 's  $\deg(v)/k$  favorite partners,  $Q_2$  is her next favorite  $\deg(v)/k$ , and so on. We call  $Q_i$   $v$ 's  *$i$ th quantile*. If we wish to make explicit the player to whom the preferences belong, we may adorn these symbols with a superscript. For example,  $Q_i^v$  is  $v$ 's  $i$ th quantile. Throughout the execution of the algorithm, elements may be removed from  $Q$  and the  $Q_i$ 's, but elements will never be added to any of these sets.

- A partner  $p$  (possibly empty). The partner  $p$  is  $v$ 's current partner in the matching  $M$  our algorithm constructs. To emphasize that  $p$  is player  $v$ 's partner, we will write  $p(v)$ . The (partial) matching  $M$  produced by the algorithm at any step is given by  $M = \{(p(w), w) | w \in X, p(w) \neq \emptyset\}$ .

Additionally, subroutines of our algorithm will require each processor to store the following variables:

- A set  $G_0$  of “neighbors” of the opposite sex which correspond to accepted proposals.
- A partner  $p_0$  in a matching found in the graph determined by  $G_0$ .

Thus each player knows their preferences, partners (if any) as well as any of their accepted proposals from the current round (stored in  $G_0$ ). The men  $m \in Y$  hold the following additional information:

- A set  $A$  of “active” potential mates, initially set to  $Q_1$ .

### 3.2 The ProposalRound subroutine

At the heart of our algorithm is the **ProposalRound** subroutine (Algorithm 1). **ProposalRound** works in 5 steps which are described in Algorithm 1.

---

#### Algorithm 1 ProposalRound( $Q, k, A$ )

---

**Step 1:** Each man  $m$  proposes to all women in  $A^m$  by sending each  $w \in A$  the message PROPOSE.

**Step 2:** Each woman  $w$  receiving proposals responds with the message ACCEPT to all proposals from her most preferred quantile  $Q_i^w$  from which at least one man proposed in Step 1.

**Step 3:** Let  $G_0$  denote the bipartite graph  $G_0$  of accepted proposals from Step 2. The players compute a maximal matching  $M_0$  in  $G_0$ , using **MaximalMatching**( $G_0$ ), storing their match in  $G_0$  as  $p_0$ .

**Step 4:** Each woman  $w$  matched in  $M_0$  sends REJECT to all men  $m \in Q^w$  in a lesser or equal quantile to her partner  $p_0(w)$  in  $M_0$  other than  $p_0(w)$ . She then removes all of these men from  $Q^w$  and the corresponding  $Q_i^w$ . The matched women then set  $p \leftarrow p_0$ , so the partial matching  $M$  now contains the edge  $(p_0(w), w)$ . Any man  $m$  matched in  $M_0$  sets  $p \leftarrow p_0$  and sets  $A \leftarrow \emptyset$ .

**Step 5:** The men remove all  $w$  from whom they received the message REJECT from their preferences  $Q$ , the various  $Q_i$ , and  $A$ . If a man  $m$  receives a rejection from his match  $p(m)$  from a previous round, he sets  $p \leftarrow \emptyset$ .

---

We observe that if each player  $v$  takes  $k = \deg v$ , then **ProposalRound** mimics the classical (extended) Gale-Shapley algorithm [4, 5]. In this case, each man proposes to his most favored woman that has not yet rejected him, and each woman rejects all but her most favored suitor. Computing a maximal matching is trivial, as the accepted proposals

already form a matching. The general case has one crucial feature in common with the Gale-Shapley algorithm, which follows immediately from the description of **ProposalRound**.

LEMMA 1 (MONOTONICITY). *Once a woman  $w$  has  $p(w) \neq \emptyset$  in some execution of **ProposalRound**, she is guaranteed to always have  $p(w) \neq \emptyset$  after each subsequent execution of **ProposalRound**. Further, once matched, she will only accept proposals from men in a strictly higher quantile than  $p(w)$ .*

### 3.3 The QuantileMatch subroutine

Here we describe the **QuantileMatch** subroutine (Algorithm 2), which simply iterates **ProposalRound** until each man  $m$  has either been rejected by all women in  $A^m$  or is matched with some woman in  $A^m$ . In either case,  $A^m = \emptyset$  when **QuantileMatch** terminates. We will argue that  $k$  (the number of quantiles) iterations suffice.

---

#### Algorithm 2 QuantileMatch( $Q, k$ )

---

```

 $i \leftarrow \min \{i | Q_i \neq \emptyset\} \cup \{k\}$  (male only)
if  $p = \emptyset$  then
   $A \leftarrow Q_i$  (male only)
end if
for  $i \leftarrow 1$  to  $k$  do
  ProposalRound( $Q, k, A$ )
end for

```

---

LEMMA 2 (**QuantileMatch** GUARANTEE). *At the termination of **QuantileMatch**( $Q, k$ ) every man  $m$  satisfies  $A^m = \emptyset$ . In particular, each man who had  $A^m \neq \emptyset$  before the first iteration of the loop in **QuantileMatch** has either been rejected by all women in  $A^m$  or is matched with some woman in  $A^m$ .*

PROOF. Suppose a woman  $w$  receives proposals in the first iteration of the loop in **QuantileMatch**. If she is matched with one of these suitors when **ProposalRound** terminates, she rejects all other men and receives no further proposals during the current **QuantileMatch**. On the other hand, if she is not matched with one of these suitors after the first round, then by the maximality of the matching found in Step 3 of **ProposalRound**, all of the suitors in her best quantile receiving proposals are matched with other women. Thus, in the next iteration, she only receives proposals from men in strictly worse quantiles than she accepted in the first. Similarly, in each iteration of the loop, her best quantile receiving proposals (if any) is strictly worse than the previous iteration. Therefore, after  $k$  iterations, no woman will receive proposals, hence each man  $m$  must have  $A^m = \emptyset$ .  $\square$

### 3.4 The ASM algorithm

In this section, we describe the main algorithm **ASM** (Algorithm 3). The idea of **ASM** is to iterate **QuantileMatch** until a large fraction men with high degree are either matched or have been rejected by all acceptable partners. We call such men *good*. By iterating **QuantileMatch** a constant number of times, we can ensure that the fraction of good men is close to 1. In order to bound the number of blocking pairs from men which are *bad* (not good), we must ensure that bad men comprise only a small fraction of players

with relatively high degree. To this end, we only allow men who are potentially involved in many blocking pairs (that is, with  $|Q|$  relatively large) to participate in later calls to **QuantileMatch**.

---

**Algorithm 3**  $\text{ASM}(P, \varepsilon, n)$

---

```

 $k \leftarrow \lceil 8\varepsilon^{-1} \rceil, \delta \leftarrow \varepsilon/8$ 
for all  $i \leq k$  do
   $Q_i \leftarrow \{v | q(v) = i\}$ 
end for
 $Q \leftarrow \bigcup_i Q_i, p \leftarrow \emptyset$ 
for  $i \leftarrow 0$  to  $\log n$  do
  if  $|Q| \geq 2^i$  then
    for  $j \leftarrow 1$  to  $2\delta^{-1}k$  do
      QuantileMatch $(Q, k)$ 
    end for
  end if
end for

```

---

## 4. PERFORMANCE GUARANTEES

Here we analyze the performance of **ASM** and its sub-routines. The run-time guarantee (Theorem 4) is a simple consequence of the description of **ASM** and its subroutines. To prove the approximation guarantee (Theorem 3), we consider blocking edges from two sets of men separately. We call a man  $m$  *good* if when **ASM** terminates, he is either matched or has been rejected by all of his acceptable partners. A man who is not good is *bad*. We denote the sets of good and bad men by  $G$  and  $B$ , respectively.

**THEOREM 3** (APPROXIMATION GUARANTEE). *The matching  $M$  output by **ASM** induces at most  $\varepsilon|E|$  blocking pairs with respect to  $\mathcal{P}$ . Thus  $M$  is  $(1 - \varepsilon)$ -stable.*

### 4.1 Bounding blocking pairs from good players

We bound the number of blocking pairs from good men in two steps. First we show that the good men are not involved in any  $(2/k)$ -blocking pairs (see Definition 2). Next, we show that as a result, the good men can only be incident with a small fraction of blocking pairs.

**LEMMA 3** ( $(2/k)$ -BLOCKING-STABILITY OF GOOD MEN). *Let  $m \in G$  be good. Then  $m$  is not incident with any  $(2/k)$ -blocking pairs.*

**PROOF.** Suppose  $m \in G$  and that  $(m, w)$  is  $(2/k)$ -blocking. First consider the case where  $m$  is matched,  $p(m) \neq \emptyset$ . Since  $m$ 's preferences are divided into  $k$  quantiles,  $w$  must be in a strictly better quantile than  $p(m)$ . Thus,  $m$  must have proposed to  $w$  in a strictly earlier call to **QuantileMatch** than the call in which he was matched with  $p(m)$ . Thus, by Lemma 2,  $m$  must have been rejected by  $w$ , implying that  $w$  was matched with a man  $m'$  in the same or better quantile than  $m$  in this round. By Lemma 1,  $w$ 's partner when **ASM** terminates is at least as desirable as  $m'$ . This contradicts that  $(m, w)$  is  $\varepsilon$ -blocking.

On the other hand, if  $p(m) = \emptyset$ , then since  $m$  is good, he must have been rejected by all of his acceptable partners, and in particular, by  $w$ . Thus, as in the previous paragraph,  $w$  must be matched with a man in the same or better quantile than  $m$ .  $\square$

**LEMMA 4** (FEW NON- $(2/k)$ -BLOCKING PAIRS). *There are at most  $4|E|/k$  blocking pairs which are not  $(2/k)$ -blocking.*

**PROOF.** Suppose  $(m, w)$  is a blocking pair which is not  $(2/k)$ -blocking. Thus, at least one of the following holds:

$$P^m(w) - P^m(p(m)) \leq 2 \deg(m)/k \quad (1)$$

$$P^w(m) - P^w(p(w)) \leq 2 \deg(w)/k, \quad (2)$$

where by convention we take  $P^m(\emptyset) = \deg(m) + 1$ . Let  $E_N$  denote the set of blocking pairs which are not  $(2/k)$ -blocking. For each  $m$ , the number of edges satisfying (1) is at most  $2 \deg(m)/k$ , and similarly for the women. Thus

$$|E_N| \leq \sum_{m \in Y} 2 \deg(m)/k + \sum_{w \in X} 2 \deg(w)/k = 4|E|/k,$$

as desired.  $\square$

Lemma 3 shows that no good player is involved in any  $(2/k)$ -blocking pairs. Combining Lemmas 3 and 4, we can bound the number of blocking pairs incident with good men. All that remains is to bound the number of  $(2/k)$ -blocking pairs incident with bad men. In the next section, we show that the proportion of bad men is small (at most  $\delta n$ ), and bound the number of  $(2/k)$ -blocking pairs they contribute. We remark that by Lemma 4 and the lower bound of Kipnis and Patt-Shamir [9], we cannot hope to have all men be good in  $o(\sqrt{n}/\log n)$  rounds.

### 4.2 Bounding blocking pairs from bad players

In this section, we prove the following bound on the number of blocking pairs contributed by the bad men at the termination of **ASM**. Throughout the section, for simplicity of notation, we assume that  $\log n$  is an integer.

**LEMMA 5** (BAD MEN GUARANTEE). *At the termination of **ASM**, for any  $\delta \leq \frac{1}{2}$  the bad men contribute at most  $4\delta|E|$   $(2/k)$ -blocking pairs.*

The proof of Lemma 5 is in two parts corresponding to guarantees for each of the two nested loops in **ASM**. We refer to men  $m$  with  $|Q^m| \geq 2^i$  as *active in the  $i$ th iteration* of the outer loop; the remaining men are *inactive* in the  $i$ th iteration.

**LEMMA 6** (FEW BAD MEN). *When the inner loop in **ASM** terminates, at most a  $\delta$ -fraction of active men are bad.*

**PROOF.** Let  $\mathcal{A}$  denote the set of active men before executing the inner loop in **ASM**. Suppose that after  $\ell$  iterations of the inner loop, there are  $b$  bad men in  $\mathcal{A}$ . We claim that there must have been at least  $b$  bad players in every iteration of the inner loop. To see this, first note that by Lemma 1, the number of matched players (and hence matched men) can only increase with each call to **ProposalRound**. Second, if a man is rejected by all women on his preference list, he will never become bad. Therefore, the number of good players can only increase with each iteration of the inner loop. Thus there must have been at least  $b$  bad men after each of the  $\ell$  iterations of the inner loop.

Suppose  $m$  was bad before some call to **QuantileMatch**, so that  $A^m \neq \emptyset$ . By Lemma 2, after **QuantileMatch**  $m$  is either matched, or has been rejected by all women  $w \in A^m$ . In the former case,  $p(m)$  rejected all men in her quantile

containing  $m$ . In either case,  $m$  witnessed the rejection of a quantile of men—either by precipitating the rejection of  $p(m)$ 's quantile, or by being rejected by all women in  $A$ . Notice that the number of women who are matched with new partners during an iteration of the outer loop cannot exceed  $|\mathcal{A}|$ , as if  $|\mathcal{A}|$  women did receive new partners, all active men would be matched. Therefore, the women can send at most  $k|\mathcal{A}|$  quantile rejections (after which all active men will be rejected by all women). Similarly, the men can receive at most  $k|\mathcal{A}|$  quantile rejections. Thus, in total the active men can witness at most  $2k|\mathcal{A}|$  quantile rejections. Therefore, if there are  $b$  bad men after  $\ell$  calls to **QuantileMatch**, we must have  $b\ell \leq 2k|\mathcal{A}|$ . Choosing  $\ell = 2\delta^{-1}k$  gives the desired result.  $\square$

We say that a man  $m$  is *bad in the  $i$ th iteration* of the outer loop in **ASM** if  $m$  became bad during the  $i$ th iteration and  $|Q^m| < 2^i$ . We denote the set bad men in the  $i$  iteration by  $B_i$ , so that  $B = B_1 \cup B_2 \cup \dots \cup B_{\log n}$ . Thus,  $m \in B_i$  is bad and will not participate in any further calls to **QuantileMatch**, so he will be bad when **ASM** terminates.

**LEMMA 7 (FEW  $(2/k)$ -BLOCKING PAIRS).** *Each  $m \in B_i$  participates in fewer than  $2^i$   $(2/k)$ -blocking pairs at the termination of **ASM**.*

**PROOF.** We will show that each bad  $m \in B$  participates in at most  $|Q^m|$   $(2/k)$ -blocking pairs, whence the lemma follows. To this end, notice that if  $w \notin Q^m$ , then  $w$  must have rejected  $m$  in some call to **QuantileMatch**. Therefore,  $w$  must have been matched with some  $m'$  that is in the same or better quantile as  $m$ . By Lemma 1, when **ASM** terminates,  $w$  is still matched with someone in at least as desirable quantile as  $m$ , implying that  $(m, w)$  is not  $(2/k)$ -blocking. Thus, every  $(2/k)$ -blocking pair  $(m, w)$  must have  $w \in Q^m$ .  $\square$

**PROOF PROOF OF LEMMA 5.** Let  $G_i \subseteq G$  be the set of men which are good at the termination of **ASM** and active after the  $i$ th iteration of the outer loop in **ASM**. Then we have  $G = G_1 \cup G_2 \cup \dots \cup G_{\log n}$ . Since the number of bad men cannot increase after a call to **QuantileMatch**. By Lemma 6, if there were  $b$  men which became bad in some iteration of the outer loop of **ASM**, there were  $\frac{1-\delta}{\delta}b$  good men still active during the  $i$ th iteration. Since the number of good men can only increase in subsequent iterations, we have

$$\begin{aligned} |B_i \cup B_{i+1} \cup \dots \cup B_{\log n}| &\leq b \\ &\leq \frac{\delta}{1-\delta} |G_i \cup G_{i+1} \cup \dots \cup G_{\log n}|. \end{aligned} \quad (3)$$

Applying (3), we can greedily form disjoint sets

$$H_{\log n} \subseteq G_{\log n}, H_{\log n-1} \subseteq G_{\log n-1} \cup G_{\log n}, \dots, H_1 \subseteq G$$

such that for all  $i$ ,  $H_i$  is active in the  $i$ th iteration and  $|H_i| = \frac{1-\delta}{\delta} |B_i|$ . Then we compute

$$\begin{aligned} \sum_{m \in B} |Q^m| &= \sum_{i=1}^{\log n} \sum_{m \in B_i} |Q^m| \\ &\leq \sum_{i=1}^{\log n} |B_i| 2^i \\ &\leq \sum_{i=1}^{\log n} \frac{2\delta}{1-\delta} |H_i| 2^i \\ &\leq \frac{2\delta}{1-\delta} \sum_{m \in G} |Q^m| \\ &\leq \frac{2\delta}{1-\delta} |E|. \end{aligned}$$

The first inequality holds by Lemma 7, while the second holds by the choice of the  $H_i$  and the definition of the  $G_i$ .  $\square$

### 4.3 Approximation guarantee

**PROOF PROOF OF THEOREM 3.** By Lemma 4, there are at most  $4|E|/k$  blocking pairs which are not  $(2/k)$ -blocking. By Lemma 3, all  $(2/k)$ -blocking pairs are incident with  $B$ . Finally, by Lemma 5, the bad men contribute at most  $4\delta|E|$  blocking pairs for  $\delta \leq 1/2$ . Therefore, the total number of blocking pairs is at most  $4(\delta + 1/k)|E|$ . Choosing  $\delta = \varepsilon/8$  and  $k = \lceil 8/\varepsilon \rceil$  gives the desired result.  $\square$

### 4.4 Run-time Guarantee

**THEOREM 4.** **ASM** $(P, \varepsilon, n)$  runs in  $O(\varepsilon^{-3} \log^5(n))$  communication rounds.

**PROOF.** Notice that the only communication between processors occurs in **ProposalRound**. **ASM** $(P, \varepsilon, n)$  iterates **QuantileMatch** $(P, k)$  a total of  $O(\varepsilon^{-2} \log n)$  times, while quantile match invokes **ProposalRound** $(Q, k, A)$   $O(\varepsilon^{-1})$  times. Finally, each step of **ProposalRound** can be performed in  $O(1)$  communication rounds, except for Step 3, which calls **MaximalMatching**. By [6], **MaximalMatching** runs in  $O(\log^4 n)$  communication rounds. Thus, **ASM** requires  $O(\varepsilon^{-3} \log^5(n))$  communication rounds, as claimed.  $\square$

*Remark 4.* While the CONGEST model allows for unbounded local computation in each round, the local computations required by **ASM** are quite simple. In fact, each communication round can easily be implemented in nearly-linear time in  $n$ . Thus the synchronous run-time of **ASM** is  $\tilde{O}(n)$ . To our knowledge, this gives the first distributed algorithm whose synchronous run-time is sub-quadratic in  $n$ , even for unbounded preferences.

## 5. RANDOMIZED ALGORITHMS

The main source of complexity in **ASM** comes from finding a maximal matching. While Hańćkowiak, Karoński, and Panconesi's algorithm [6] is the most efficient known deterministic algorithm, faster randomized algorithms are known. Specifically, we consider the algorithm of Israeli and Itai [8]. They describe a simple randomized distributed algorithm which finds a maximal matching in expected  $O(\log n)$  rounds. By simply replacing **MaximalMatching** with a truncated version Israeli and Itai's algorithm, we obtain a

faster randomized algorithm for finding almost stable matchings. We refer the reader to Appendix A for details on the guarantees for Israeli and Itai’s algorithm.

## 5.1 General preferences

**THEOREM 5.** *There exists a randomized distributed algorithm, **RandASM**( $P, \varepsilon, n, \delta$ ), which for any  $\delta, \varepsilon > 0$  finds a  $(1 - \varepsilon)$ -stable matching with probability at least  $1 - \delta$  in  $O(\varepsilon^{-3} \log^2(n/\delta\varepsilon^3))$  rounds.*

**PROOF SKETCH.** We take **RandASM** to be exactly the same as **ASM**, except that we use Israeli and Itai’s algorithm [8] for the **MaximalMatching** subroutine. Specifically, for **MaximalMatching**, we iterate **MatchingRound** (see Appendix A)  $O(\log(n/\delta\varepsilon^3))$  times. By Corollary 1, each call to **MaximalMatching** will succeed in finding a maximal matching with probability at least  $1 - O(\delta\varepsilon^3/\log n)$ . Since **RandASM** calls **MaximalMatching**  $O(\varepsilon^{-3}/\log n)$  times, by the union bound, every call to **MaximalMatching** succeeds with probability at least  $1 - \delta$ . The remaining analysis of **RandASM** is identical to that of **ASM**.  $\square$

## 5.2 Almost-regular preferences

For  $\alpha \geq 1$ , we call preferences  $\mathcal{P}$   $\alpha$ -almost-regular if  $\max_{m \in Y} \deg m \leq \alpha \min_{m \in Y} \deg m$ . For example, complete preferences (where all men rank all women) are 1-almost-regular, while uniformly bounded preferences are  $\alpha$ -almost-regular for  $\alpha = \max_{m \in Y} \deg m$ . From an algorithmic standpoint,  $\alpha$ -almost-regular preferences are advantageous because in order to bound the proportion of blocking edges from bad men, it suffices only to bound the number of bad men. By Lemma 6, to obtain such a guarantee, one need only iterate **QuantileMatch**  $O(1)$  rounds (instead of  $O(\log n)$  times as required by **ASM**).

Further, for  $\alpha$ -almost-regular preferences, we can relax our requirement that **MaximalMatching** actually find a maximal matching. We say that a player  $v$  is *unmatched* in  $G_0$  if  $v$  does not satisfy property 1 or 2 in Definition 3. We call a subroutine **AMM**( $\eta, \delta$ ) which finds a matching in which only an  $\eta$ -fraction of players are left unmatched with probability at least  $1 - \delta$  (see Appendix A for details). These unmatched players are immediately removed from play. With these simplifications, we obtain the following result.

**THEOREM 6.** *There exists a randomized distributed algorithm **AlmostRegularASM**( $P, \varepsilon, \delta, \alpha$ ) which for  $\alpha$ -almost-regular preferences  $P$  finds a  $(1 - \varepsilon)$ -stable matching with probability at least  $1 - \delta$ . The run-time of **AlmostRegularASM**( $P, \varepsilon, \delta, \alpha$ ) is  $O(\alpha\varepsilon^3 \log(\alpha/\delta\varepsilon))$  rounds.*

**PROOF SKETCH.** **AlmostRegularASM**( $P, \varepsilon, \delta, \alpha$ ) works by iterating **QuantileMatch**  $O(\alpha\varepsilon^{-2})$  times, which by Lemma 6 implies that only  $\varepsilon/4\alpha$  fraction of men are bad.

We modify **ProposalRound** to call **AMM**( $\eta, \delta'$ ) instead of **MaximalMatching**. **AMM** runs in  $O(\log((\eta\delta')^{-1}))$  and finds a  $(1 - \eta)$ -maximal matching with probability  $(1 - \delta')$ . Since **AMM** is called  $O(\alpha\varepsilon^{-3})$  times, choosing  $\eta = O(\varepsilon^4/\alpha)$  and  $\delta' = O(\delta\varepsilon^3/\alpha)$ , **AMM** will leave at most an  $\varepsilon/4\alpha$  fraction of men unmatched in any call **AMM** with probability at least  $1 - \delta$ , by the union bound. Such unmatched men are immediately removed from play.

By the preceding two paragraphs, **AlmostRegularASM** produces a matching in which at most an  $\varepsilon/2\alpha$  fraction of men are either bad or unmatched. By  $\alpha$ -almost-regularity,

these men can contribute at most  $\frac{\varepsilon}{2}|E|$  blocking pairs. The remaining men are good, and therefore by Lemmas 3 and 4<sup>2</sup> contribute at most  $\frac{\varepsilon}{2}|E|$  blocking pairs.  $\square$

## 6. ACKNOWLEDGMENTS

The first author is supported in part by NSF grants 09165174, 1065276, 1118126 and 1136174, US-Israel BSF grant 2008411, OKAWA Foundation Research Award, IBM Faculty Research Award, Xerox Faculty Research Award, B. John Garrick Foundation Award, Teradata Research Award, and Lockheed-Martin Corporation Research Award. This material is based upon work supported by the Defense Advanced Research Projects Agency through the U.S. Office of Naval Research under Contract N00014-11-1-0392. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

The second author thanks the anonymous referees for their helpful comments.

## 7. REFERENCES

- [1] David J Abraham, Péter Biró, and David F Manlove. Almost stable matchings in the roommates problem. In *Approximation and online algorithms*, pages 1–14. Springer, 2006.
- [2] Kimmo Eriksson and Olle Häggström. Instability of matchings in decentralized markets with various preference structures. *International Journal of Game Theory*, 36(3):409–420, March 2008.
- [3] Patrik Floréen, Petteri Kaski, Valentin Polishchuk, and Jukka Suomela. Almost Stable Matchings by Truncating the Gale–Shapley Algorithm. *Algorithmica*, 58(1):102–118, 2010.
- [4] D Gale and L S Shapley. College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, 69(1):pp. 9–15, 1962.
- [5] Dan Gusfield and Robert W Irving. *The stable marriage problem: structure and algorithms*, volume 54. MIT press Cambridge, 1989.
- [6] Michal Hańkowiak, Michal Karoński, and Alessandro Panconesi. On the distributed complexity of computing maximal matchings. *SIAM Journal on Discrete Mathematics*, 15(1):41–57, 2001.
- [7] Avinatan Hassidim, Yishay Mansour, and Shai Vardi. Local computation mechanism design. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 601–616. ACM, 2014.
- [8] Amos Israeli and Alon Itai. A fast and simple randomized parallel algorithm for maximal matching. *Information Processing Letters*, 22(2):77–80, 1986.
- [9] Alex Kipnis and Boaz Patt-Shamir. A note on distributed stable matching. In *Proceedings of the 28th ACM symposium on Principles of distributed computing*, pages 282–283, New York, NY, USA, 2009. ACM.
- [10] David Manlove. *Algorithmics of matching under preferences*. World Scientific Publishing, 2013.

<sup>2</sup>Although these lemmas were proven assuming that **MaximalMatching** found a maximal matching (not an almost maximal matching) the proofs remain valid as long as the small fraction of unmatched players immediately remove themselves from play.

[11] David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. Society for Industrial and Applied Mathematics, 2000.

## APPENDIX

### A. RANDOMIZED MAXIMAL AND ALMOST MAXIMAL MATCHINGS

Israeli and Itai’s [8] algorithm for finding a maximal matching works by identifying a sparse subgraph of  $G$ , then finding a large matching  $M_1$  in the sparse subgraph. The edges and incident vertices of  $M_1$ , as well as remaining isolated vertices, are removed from  $G$  resulting in a subgraph  $G_1$ . The process is iterated, giving a sequence of subgraphs  $G_1, G_2, \dots$  and matchings  $M_1, M_2, \dots$ , until  $G_k = \emptyset$ . At this point,  $M = \bigcup_{i=1}^k M_i$  is a maximal matching. We give pseudocode for Israeli and Itai’s main subroutine, which we call **MatchingRound**, in Algorithm 4. In [8], Israeli and Itai prove the following performance guarantee for **MatchingRound**.

---

**Algorithm 4 MatchingRound( $G$ ):** Finds a large matching in a graph

---

- 1: Each  $v \in V$  picks a neighbor  $w$  uniformly at random, forms oriented edge  $(v, w)$ .
  - 2: Each  $v \in V$  with  $\deg_{in}(v) > 0$  picks one in-coming edge  $(w, v)$  uniformly at random, deletes remaining in-edges. Let  $G'$  be the (undirected) graph formed by the chosen edges with orientation ignored.
  - 3: Each  $v \in V$  with  $\deg_{G'}(v) > 0$  chooses one incident edge  $(v, w)$  uniformly at random.
  - 4: The matching  $M_1$  consists of edges  $(v, w) \in G'$  which were chosen by both  $v$  and  $w$  in the previous round.  $G_1 = (V_1, E_1)$  is the induced subgraph of  $G$  formed by removing all vertices contained in  $M_1$  and any remaining isolated vertices from  $G$ .
  - 5: Output  $(G_1, M_1)$ .
- 

LEMMA 8. (Israeli and Itai [8]) *There exists an absolute constant  $c < 1$  such that on input  $G = G_0 = (V_0, E_0)$ , the resulting graph  $G_1 = (V_1, E_1)$  found by **MatchingRound** satisfies  $\mathbf{E}(|V_1|) \leq c|V_0|$ .*

As a consequence of Lemma 8, we obtain the following useful result.

COROLLARY 1. *Let  $\eta > 0$  be a parameter. Then  $s = O(\log(n/\eta))$  iterations of **MatchingRound** suffice to produce a maximal matching in  $G$  with probability at least  $1 - \eta$ .*

PROOF. By Lemma 8, we have  $\mathbf{E}(|V_s|) \leq c^s n$ . Therefore, applying Markov’s inequality gives

$$\Pr(|V_s| \geq 1) \leq \frac{\mathbf{E}(|V_s|)}{1} \leq c^s n.$$

The result follows by taking  $s \geq \log(n/\eta)/\log(c^{-1})$ .  $\square$

The almost regular variant of **ASM** only requires a subroutine that we finds matchings which are almost maximal.

Definition 4. Let  $G = (V, E)$  be a communication graph and  $M \subset E$  a matching in  $G$ . For  $0 < \eta \leq 1$ , we say that  $M$  is  $(1 - \eta)$ -**maximal** if the set  $V'$  of vertices not satisfying conditions 1 or 2 in Definition 3 satisfies  $|V'| \leq \eta|V|$ .

We can apply Lemma 8 to give a constant round algorithm which finds almost maximal matchings.

COROLLARY 2. *There exists a randomized distributed algorithm **AMM**( $G, \eta, \delta$ ) which finds a  $(1 - \eta)$ -maximal matching with probability at least  $(1 - \delta)$ . **AMM**( $G, \eta, \delta$ ) runs in  $O(\log(\eta^{-1}\delta^{-1}))$  rounds.*

PROOF. Consider the algorithm which iterates **MatchingRound**  $s$  times. We apply Lemma 8 and Markov’s inequality to obtain

$$\Pr(|V_s| \geq \eta n) \leq \frac{c^s n}{\eta n} = \eta^{-1} c^s.$$

Choosing  $s = O(\log(\delta^{-1}\eta^{-1}))$ , we have  $\eta^{-1} c^s \leq \delta$ , which gives the desired result.  $\square$