Brief Announcement: Distributed Almost Stable Marriage*

Patrik Floréen, Petteri Kaski, Valentin Polishchuk, and Jukka Suomela
Helsinki Institute for Information Technology HIIT, University of Helsinki
P.O. Box 68, FI-00014 University of Helsinki, Finland
patrik.floreen@cs.helsinki.fi, petteri.kaski@cs.helsinki.fi,
valentin.polishchuk@cs.helsinki.fi, jukka.suomela@cs.helsinki.fi

ABSTRACT

We study the stable marriage problem in a distributed setting. The communication network is a bipartite graph, with men on one side and women on the other. Acceptable partners are connected by edges, and each participant has chosen a linear order on the adjacent nodes, indicating the matching preferences.

The classical Gale–Shapley algorithm could be simulated in such a network to find a stable matching. However, the stable matching problem is inherently global: the worst-case running time of any distributed algorithm is linear in the diameter of the network.

Our work shows that if we tolerate a tiny fraction of unstable edges, then a solution can be found by a fast local algorithm: simply truncating a distributed simulation of the Gale–Shapley algorithm is sufficient. Among others, this shows that an almost stable matching can be maintained efficiently in a very large network that undergoes frequent changes.

Categories and Subject Descriptors

C.2.4 [**Computer-Communication Networks**]: Distributed Systems; F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*computations on discrete structures*

General Terms

Algorithms, Theory

1. PROBLEM FORMULATION

In the stable marriage problem, the input is a bipartite graph $\mathcal{G}$ whose parts represent men and women. Each individual ranks adjacent nodes in $\mathcal{G}$, by imposing a linear order on them. The orders define the matching preferences; in particular, two nodes that are not connected are unacceptable partners. Let Δ denote the maximum degree of $\mathcal{G}$; assume $\Delta = O(1)$. In terms of marriages, Δ is the maximum number of acceptable partners for an individual.

A *matching* M in $\mathcal{G}$ is a subset of edges such that every node is incident to at most one edge from M. An edge $\{u, v\} \notin M$ of $\mathcal{G}$ is *unstable* if both (a) node u is unmatched or prefers v over its partner in M, and (b) node v is unmatched or prefers u over its partner in M. A matching M is *stable* if $\mathcal{G}$ has no unstable edges.

The example in Figure 1 shows that, in general, a stable matching cannot be computed by using only local information in the network: if every participant initially knows only his/her neighbours and personal preferences, the number of communication rounds required to arrive at a stable matching can be linear in the network size.

Fortunately, we find that local information suffices to arrive at a matching where the fraction of unstable edges can be made arbitrarily small in the sense of the following definition:

Definition 1. For a constant $\epsilon > 0$, a matching M in $\mathcal{G}$ is *ϵ-stable* if the number of unstable edges in $\mathcal{G}$ is at most $\epsilon|M|$.

2. RESULTS

The following theorem summarises our main contribution.

THEOREM 1. *There exists a deterministic distributed algorithm that finds an ϵ-stable matching in $\mathcal{G}$ in $4 + 2\Delta^2/\epsilon$ synchronous communication rounds.*

We prove Theorem 1 by analysing the "transient phase" of the classical Gale–Shapley algorithm [2] for the stable marriage problem. The algorithm consists of a sequence of propose–accept rounds; it is straightforward to implement the rounds in parallel. We show that if the algorithm's execution is stopped after a constant number of rounds, one obtains a matching with only a constant fraction of unstable edges.

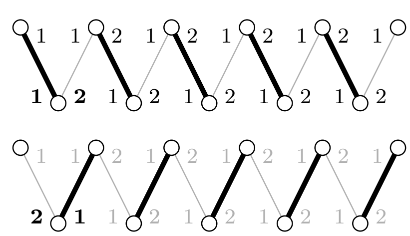

Figure 1: The numbered edge ends indicate preference rankings; the most preferred match has rank 1. In these examples, a stable matching is unique; the thick lines show the unique solution. After transposing the preferences of a single participant, the unique stable matching changes completely.

*The full version of this work will appear in *Algorithmica* [1].

PODC'10, July 25–28, 2010, Zurich, Switzerland.
ACM 978-1-60558-888-9/10/07.

In more precise terms, assume that we implement the "man-optimal" variant of the Gale–Shapley algorithm: men propose and women accept/reject the proposals. The key idea is to define the total *potential* of men as the number of unmatched men with proposals left; the potential reflects how much the men could gain altogether if the Gale–Shapley algorithm were not truncated. Initially, the potential is high; if the Gale–Shapley is run to the termination, the potential becomes 0. We prove that the potential is monotonically decreasing in the number of the propose–accept rounds: a man loses his partner only due to another man gaining a partner.

Now, if after round $k > 1$ the potential is α, then α men received a negative message (either a 'no' to a proposal, or a 'break' from a partner) in round k. The key property of the Gale–Shapley algorithm is that the edge along which the negative message was sent is of no use in the future: the woman at the other end of the edge is happier with the current partner than with the rejected man (and women can only improve in the subsequent rounds). Hence, α edges can be removed from $\mathcal{G}$ in round k. Because the potential is non-decreasing, in rounds $2, 3, \ldots, k$ at least $(k-1)\alpha$ edges were removed. But the total number of removed edges is $O(\Delta|M|)$, hence the potential after round k is $O(\Delta|M|/k)$, implying that there are $O(\Delta^2|M|/k)$ unstable edges.

2.1 Maximum-Weight Matchings

Suppose now that each edge of $\mathcal{G}$ has a weight. We can define the preferences by the edge weights: each participant prefers his or her neighbours in order of the weights of the incident edges (the heavier the edge, the more preferable the neighbour is). In general, an almost stable matching need not be a good approximation to the maximum-*weight* matching in $\mathcal{G}$ (due to, say, a very heavy edge not present in the almost stable matching). An interesting property of the almost stable matching computed by the truncated Gale–Shapley algorithm is that its weight is quite close to the maximum possible weight of *any* matching in $\mathcal{G}$ (stable or not):

THEOREM 2. *There exists a deterministic distributed algorithm that finds a $(2+\epsilon)$-approximation of a maximum-weight matching in $\mathcal{G}$ in $4 + 2\Delta/\epsilon$ communication rounds.*

2.2 Estimating the Size of a Stable Matching

A classical result in the theory of stable marriages is that all stable matchings in $\mathcal{G}$ have the same size [3, Section 1.4.2]; denote the size by m. The truncated Gale–Shapley algorithm can be applied in a centralised setting as well. We only need to have access to a *preference oracle* that, when queried with a node, returns the neighbours of the node in order of preference. We show that a constant number of queries suffice to estimate m. Put otherwise, our algorithm gives a constant-time centralised randomised approximation scheme for computing the size of a stable matching in $\mathcal{G}$:

THEOREM 3. *For any $0 < \delta \leq 1/2$, $0 < \epsilon \leq 1$, and $\Delta \geq 3$, there exists a randomised algorithm that, given access to a preference oracle for $\mathcal{G}$, makes at most*

$$25000\epsilon^{-2}(\Delta-1)^{3+4\Delta/\epsilon}\ln\delta^{-1}$$

queries to the oracle and outputs with probability at least $1-\delta$ an estimate $\hat{m}$ such that $\left|\hat{m}-m\right| \leq \epsilon m$.

We refer to the full version of this work [1] for detailed proofs.

3. ACKNOWLEDGEMENTS

This work was supported in part by the Academy of Finland, Grants 116547, 117499, 132380, and 118653 (ALGODAN), by Helsinki Graduate School in Computer Science and Engineering (Hecse), and by the Foundation of Nokia Corporation.

4. REFERENCES

[1] Patrik Floréen, Petteri Kaski, Valentin Polishchuk, and Jukka Suomela. Almost stable matchings by truncating the Gale–Shapley algorithm. *Algorithmica*, 2009. To appear.

[2] David Gale and Lloyd S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

[3] Dan Gusfield and Robert W. Irving. *The Stable Marriage Problem: Structure and Algorithms.* The MIT Press, Cambridge, MA, USA, 1989.