RWTH Aachen

Department of Computer Science
Technical Report

Uncoordinated Two-Sided Markets

Heiner Ackermann, Paul W. Goldberg, Vahab S. Mirrokni,
Heiko Röglin, and Berthold Vöcking

ISSN 0935–3232 · Aachener Informatik Berichte · AIB-2007-22

RWTH Aachen · Department of Computer Science · December 2007

Uncoordinated Two-Sided Markets*

Heiner Ackermann[1], Paul W. Goldberg[2],
Vahab S. Mirrokni[3], Heiko Röglin[1], and Berthold Vöcking[1]

[1] Department of Computer Science, RWTH Aachen, Germany
{ackermann,roeglin,voecking}@cs.rwth-aachen.de
[2] Department of Computer Science, University of Liverpool, U. K.
P.W.Goldberg@liverpool.ac.uk
[3] Microsoft Research, Redmond, WA, USA
mirrokni@microsoft.com

Abstract. Various economic interactions can be modeled as two-sided markets. A central solution concept to these markets are *stable matchings*, introduced by Gale and Shapley. It is well known that stable matchings can be computed in polynomial time, but many real-life markets lack a central authority to match agents. In those markets, matchings are formed by actions of self-interested agents. Knuth introduced uncoordinated two-sided markets and showed that the uncoordinated better response dynamics may cycle. However, Roth and Vande Vate showed that the random better response dynamics converges to a stable matching with probability one, but did not address the question of *convergence time*.

In this paper, we give an *exponential lower bound* for the convergence time of the random better response dynamics in two-sided markets. We also extend these results to the *best response* dynamics, i. e., we present a cycle of best responses, and prove that the random best response dynamics converges to a stable matching with probability one, but its convergence time is exponential. Additionally, we identify the special class of *correlated two-sided markets* with real-life applications for which we prove that the random best response dynamics converges in expected polynomial time.

1 Introduction

One main function of many markets is to match agents of different kinds to one another, for example men and women, students and colleges [GS62], interns and hospitals [Rot84,Rot96], and firms and workers. Gale and Shapley [GS62] introduced *two-sided markets* to model these problems. A two-sided market consists of two disjoint groups of agents. Each agent has some preferences about the agents on the other side and can be matched to one of them. A matching is *stable* if it does not contain a *blocking pair*, i. e., a pair of agents from different sides who can deviate from this matching and both benefit. Gale and Shapley [GS62] showed that stable matchings always exist and can be found in polynomial time. Besides their theoretical appeal, two-sided matching models have proved useful in the empirical study of many labor markets such as the National Resident Matching Program (NRMP). Since the seminal work of Gale and Shapley, there has been a significant amount of work in studying two-sided markets, especially on extensions to many-to-one matchings and preference lists with ties [KC82,RS90,Fle03,EO06]. See for example, the book by Knuth [Knu76], the book by Gusfield and Irving [GI89], or the book by Roth and Sotomayor [RS90].

* This work was supported by DFG grant VO 889/2, EPSRC Grant GR/T07343/02, and by the EU within the 6th Framework Programme under contract 001907 (DELIS).

In many real-life markets, there is no central authority to match agents, and agents are self-interested entities. This motivates the study of *uncoordinated two-sided markets*, first proposed by Knuth [Knu76]. Uncoordinated two-sided markets can be modeled as a game among agents of one side, which we call the *active* side. The strategy of each active agent is to choose one agent from the *passive* side. Stable matchings correspond to Nash equilibria of the corresponding games. In these uncoordinated markets, it is important to analyze better response dynamics among agents, and bound the number of steps for agents to converge to a stable matching. In this regard, Knuth showed that a sequence of better responses of agents can cycle, and posed a question concerning the convergence of this dynamics. Consider the following *random better response dynamics*: at each step, pick a blocking pair of agents at random and let the agents in this pair match to each other. Roth and Vande Vate [RV90] proved that the random better response dynamics converges to a stable matching with probability one. However, they do not address the question of *convergence time*.

Our first result in this paper is an *exponential lower bound* for the convergence time of this better response dynamics in uncoordinated two-sided markets (Theorem 1). Both Knuth's cycle [Knu76], and Roth and Vande Vate's proof [RV90] hold only for the better response dynamics, and not for the *best response dynamics*. We strengthen the results in [Knu76,RV90] to best responses. That is, we illustrate a cycle of best responses of agents (Theorem 2), and then, using a potential function argument, we show that starting from any matching, there exists a short sequence of best responses of agents to a stable matching (Theorem 3). Moreover, we study the *random best response dynamics* and show an exponential lower bound for its convergence time to stable matchings (Theorem 4).

The above lower bounds show that the decentralized game theoretic approach for stable matchings does not converge in polynomial time. This motivates studying special cases of two-sided markets for which the convergence time is polynomial. In this regard, we consider a natural class of *correlated two-sided markets*, which are inspired from real-life one-sided market games in which players have preferences about a set of markets, and the preferences of markets are correlated with the preferences of players. This special class of two-sided markets is shown to be a potential game in [AGM$^+$07] and complexity related questions are studied in [ALMO07]. Two illustrative examples of these markets are market sharing games [GLMT04], and distributed caching games [FGMS06,Mir05]. These markets have been also studied for finding stable geometric configurations with applications in VLSI design [HHP06]. In a correlated two-sided market, there is a payoff associated with every possible pair of active and passive agent. Both active and passive agents are interested in maximizing their payoff, that is, an agent i prefers an agent j to an agent j' if the payoff associated with pair (i, j) is larger than the payoff associated with pair (i, j'). In contrast to general two-sided markets, we show that the random best response dynamics converges in polynomial time to a stable matching in correlated two-sided markets (Theorem 5).

2 Preliminaries and Notations

In this section, we define the problems and notations that will be used throughout the paper.

Two-sided Markets. A two-sided market consists of two disjoint groups of agents $\mathcal{X}$ and $\mathcal{Y}$, e. g., women and men. Each agent has a preference list over the agents of the other side. An agent $i \in \mathcal{X} \cup \mathcal{Y}$ can be assigned to one agent j in the other side. Then she gets *payoff* $p_i(j)$. If the preference list of agent i is $(a_1, a_2, \ldots, a_n)$, we say that agent i has payoff $k \in \{0, \ldots, n-1\}$ if she is matched to agent a_{n-k}. Also, we say that an agent has payoff -1 if she is unmatched. Given a matching M, we denote the payoff of an agent i in matching M by $p_i(M)$. Throughout the paper, we use women or players as active agents, and men, or resources, or markets as passive agents in the corresponding market game.

Given a matching M, an agent $x \in \mathcal{X}$ and an agent $y \in \mathcal{Y}$ form a *blocking pair* if $\{x, y\} \notin M$ and $p_x(y) > p_x(M)$ and $p_y(x) > p_y(M)$. Given a matching M and a blocking pair (x, y) in M, we say that a matching M' is obtained from M by *resolving* the blocking pair (x, y) if the following holds: $\{x, y\} \in M'$, any partners with whom x and y are matched in M are unmatched in M', and all other edges in M and M' coincide. A matching is *stable* if it does not contain a blocking pair.

Uncoordinated Two-sided Markets. We model the uncoordinated two-sided market $(\mathcal{X}, \mathcal{Y})$ as a game $G(\mathcal{X}, \mathcal{Y})$ among agents of the *active* side $\mathcal{X}$. The strategy of each active agent $x \in \mathcal{X}$ is to choose one agent y from the *passive* side $\mathcal{Y}$. The goal of each active agent $x \in \mathcal{X}$ is to maximize her payoff $p_x(y)$. Given a strategy vector of active players, an active agent x obtains payoff $p_x(y)$ if she proposes to y, and if she is the *winner* of y. Agent x is the winner of y if y ranks x highest among all active agents who currently propose to her. Additionally, passive agent y obtains $p_y(x)$ if x is the winner of y.

Remark 1. Stable matchings in an uncoordinated two-sided market $(\mathcal{X}, \mathcal{Y})$ correspond to *pure Nash equilibria* of the corresponding game $G(\mathcal{X}, \mathcal{Y})$ and vice versa.

Consider two agents $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. If a blocking pair (x, y) is resolved, we say that x plays a *better response*. If there does not exist a blocking pair (x, y') with $p_x(y') > p_x(y)$, then we say that x plays a *best response* when the blocking pair (x, y) is resolved. In the *random better response dynamics* at each step a blocking pair is chosen uniformly at random and resolved. In the *random best response dynamics* at each step an active agent from $\mathcal{X}$ is chosen uniformly at random and allowed to play a best response.

Correlated Two-sided Markets. In general, there are no dependencies between the preference lists of agents. Correlated two-sided markets are examples in which the preference lists are correlated. Assume that there is a payoff $p_{x,y} \in \mathbb{N}$ associated with every pair (x, y) of agents $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ such that $p_x(y) = p_y(x) = p_{x,y}$. The preference lists of both active and passive agents are then defined according to these payoffs, e. g., a passive agent y prefers an active agent x to an active agent x' if $p_{x,y} > p_{x',y}$. We assume that for every agent i, the payoffs associated to all pairs including agent i are pairwise distinct. Then the preference lists are uniquely determined by the ordering of the payoffs.

3 Better Response Dynamics

In this section, we consider the random better response dynamics and present instances for which with high probability the better response dynamics takes exponential time. We present our instances using an edge-weighted bipartite graph with an edge for each pair of woman and man. A woman w prefers a man m to a man m' if the weight of the edge $\{w, m\}$ is *smaller* than the weight of $\{w, m'\}$. On the other hand, a man m prefers a woman w to a woman w' if the weight of the edge $\{m, w\}$ is *larger* than the weight of the edge $\{m, w'\}$. The bipartite graph is depicted in Figure 1. Before we analyze the number of better

	m_1	m_2	m_3	$\dots$	m_{n-2}	m_{n-1}	m_n
w_1	1	2	3	$\dots$	$n-2$	$n-1$	n
w_2	n	1	2	$\dots$	$n-3$	$n-2$	$n-1$
w_3	$n-1$	n	1	$\dots$	$n-4$	$n-3$	$n-2$
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
w_{n-1}	3	4	5	$\dots$	n	1	2
w_n	2	3	4	$\dots$	$n-1$	n	1

Fig. 1: The weights of the edges in our construction.

responses needed to reach a stable matching, we prove a structural property of the instances we construct.

Lemma 1. *For the family of instances of the two-sided market problem that is depicted in Figure 1, a matching M is stable if and only if it is perfect and every woman has the same payoff in M.*

Proof. First we show that every perfect matching M in which every woman has the same payoff is stable. One crucial property of our construction is that whenever a woman w and a man m are married, the sum $p_w(m) + p_m(w)$ of their payoffs is $n-1$. In order to see this, assume that the edge between w and m has weight $l+1$. Then there are l men whom woman w prefers to m, i. e., $p_w(m) = n-1-l$. Furthermore, there are $n-1-l$ women whom man m prefers to w, i. e., $p_m(w) = l$. This implies $p_w(m)+p_m(w) = n-1$. We consider the case that every woman has payoff k and hence every man has a payoff of $n-1-k$ in M. Assume that there exists a blocking pair (w, m). Currently w has payoff k, m has payoff $n-1-k$, and w and m are not married to each other. Since (w, m) is a blocking pair, $p_w(m) > k$ and hence $p_m(w) = n-1-p_w(m) < n-1-k = p_m(M)$, contradicting the assumption that (w, m) is a blocking pair.

Now we have to show that a state M in which not every woman has the same payoff cannot be a stable matching. We can assume that M is a perfect matching as otherwise it obviously cannot be stable. Let M be a perfect matching and define $l(M)$ to be the lowest payoff that one of the women receives, i. e., $l(M) = \min\{p_w(M) \mid w \in \mathcal{X}\}$. Furthermore, by $L(M)$ we denote the set of women receiving payoff $l(M)$, i. e., $L(M) = \{w \in \mathcal{X} \mid p_w(M) = l(M)\}$. We claim that there exists at least one woman in $L(M)$ who forms a blocking pair with one of the men.

First we consider the case that the lowest payoff is unique, i. e., $L(M) = \{w\}$. Let m be the man with $p_w(m) = l(M)+1$. We claim that (w, m) is a blocking pair.

To see this, let M' denote the matching obtained from M by resolving (w, m). We have to show that the payoff $p_m(M)$ of man m in matching M is smaller than his payoff $p_m(M')$ in M'. Due to our construction $p_m(M') = n - 1 - p_w(m)$ and $p_m(M) = n - 1 - p_{w'}(m)$, where w' denotes m's partner in M. Due to our assumption, w is the unique woman with the lowest payoff in M. Hence, $p_{w'}(m) = p_{w'}(M) > p_w(M) = p_w(m) - 1$. This implies $p_m(M') \geq p_m(M)$, and hence (w, m) is a blocking pair.

It remains to consider the case that the woman with the lowest payoff is not unique. We claim that also in this case we can identify one woman in $L(M)$ who forms a blocking pair. Let $w^{(1)} \in L(M)$ be chosen arbitrarily and let $m^{(1)}$ denote her partner in M. Let $m^{(2)}$ denote the man with $p_{w^{(1)}}(m^{(2)}) = p_{w^{(1)}}(m^{(1)}) + 1$ and let $w^{(2)}$ denote the woman married to $m^{(2)}$ in M. If the payoff of $w^{(2)}$ in M is larger than the payoff of $w^{(1)}$ in M, then by the same arguments as for the case $|L(M)| = 1$, it follows that $(w^{(1)}, m^{(2)})$ is a blocking pair. Otherwise, if $p_{w^{(1)}}(M) = p_{w^{(2)}}(M)$, we continue our construction with $w^{(2)}$. To be more precise, we choose the man $m^{(3)}$ with $p_{w^{(2)}}(m^{(3)}) = p_{w^{(2)}}(m^{(2)}) + 1$ and denote by $w^{(3)}$ his partner in M. Again either $w^{(3)} \in L(M)$ or $(w^{(2)}, m^{(3)})$ is a blocking pair. In the former case, we continue the process analogously, yielding a sequence $m^{(1)}, m^{(2)}, m^{(3)}, \ldots$ of men. If the sequence is finite, a blocking pair exists. Now we consider the case that the sequence is not finite. Let $j \in \{1, \ldots, n\}$ be chosen such that $m^{(1)} = m_j$. Due to our construction, it holds $m^{(i)} = m_{(j-i \bmod n)+1}$ for $i \in \mathbb{N}$. Hence, in this case, every man appears in the sequence, and hence every woman has the same payoff $l(M)$.

Now we can prove that with high probability the number of better responses needed to reach a stable matching is exponential.

Theorem 1. *There exists an infinite family of two-sided market instances $I_1, I_2, I_3, \ldots$ and corresponding matchings $M_1, M_2, M_3, \ldots$ such that, for $n \in \mathbb{N}$, I_n consists of n women and n men and a sequence of random better responses starting in M_n needs $2^{\Omega(n)}$ steps to reach a stable matching with probability $1 - 2^{-\Omega(n)}$.*

Proof. We consider the instances shown in Figure 1. In Lemma 1, we have shown that in any stable matching all women have the same payoff. For a given matching M, we are interested in the most common payoff among the women and denote by $\chi(M)$ the number of women having this payoff, i. e.,

$$\chi(M) = \max_{i \in \{0,\ldots,n-1\}} |\{w \in \mathcal{X} \mid p_w(M) = i\}| \quad .$$

In the following, we show that whenever $\chi(M)$ is at least $15n/16$, then $\chi(M)$ is more likely to decrease than to increase. This yields a biased random walk which takes with high probability exponentially many steps to reach $\chi(M) = n$. If the most common payoff is unique, which is always the case if $\chi(M) > n/2$, then we denote by $\mathcal{X}'(M)$ the set of women having this payoff and by $\mathcal{Y}'(M)$ the set of men married to women from $\mathcal{X}'(M)$.

Let $\delta = 15/16$ and assume that $\chi(M) \geq \delta n$. First, we consider the case that the current matching M is not perfect, i. e., there exists at least one unmatched woman w and at least one unmatched man m. We call a blocking pair *good* if for the matching M' obtained from resolving it, $\chi(M') = \chi(M) - 1$. On the other hand, we call a blocking pair *bad* if $\chi(M') = \chi(M) + 1$ or if M' is a perfect matching. We count now the number of good and of bad blocking pairs. Let k denote the most common payoff. Both the unmarried woman w and the unmarried man m form a blocking pair which each person who prefers her/him to his/her current partner. Since the current payoff of the women in $\mathcal{X}'(M)$ is k, at most k of these women do not improve their marriage by marrying the unmarried man m. Analogously, since the payoff of the men in $\mathcal{Y}'(M)$ is $n-1-k$, at most $n-1-k$ of these men do not improve their marriage by marrying the unmarried woman w. This implies that the number of good blocking pairs is at least $\max\{\delta n - k, \delta n - n + 1 + k\} \geq (\delta - 1/2)n$. On the other hand, there can be at most $(1-\delta)n + 1$ bad blocking pairs. This follows easily because only women from $\mathcal{X} \setminus \mathcal{X}'(M)$ can form bad blocking pairs and each of these women forms at most one bad blocking pair as there is only one man who is at position $n-k$ in her preference list. Furthermore, there exists at most one blocking pair which makes the matching perfect.

For a matching M with $\chi(M) \geq \delta n$, the ratio of good blocking pairs to bad blocking pairs is bounded from below by

$$\frac{(\delta - 1/2)n}{(1-\delta)n + 1} \geq \frac{7}{2} \ .$$

This implies that the conditional probability of choosing a good blocking pair under the condition that either a good or a bad blocking pair is chosen is bounded from below by 7/9.

If a good blocking pair is chosen, χ decreases by 1. If a bad blocking pair is chosen χ increases by 1 or the matching obtained is perfect. In the latter case, after the next step again a matching M'' is obtained that is not perfect. For this matching M'', we have $\chi(M'') \leq \chi(M) + 2$. Since we are interested in proving a lower bound, we can pessimistically assume that the current matching is not perfect and that whenever a bad blocking pair is chosen, χ increases by 2. Hence, we can obtain a lower bound on the number of better responses needed to reach a stable state, i. e., a state M with $\chi(M) = n$, by considering a random walk on the set $\{\lceil \delta n \rceil, \lceil \delta n \rceil + 1, \ldots, n\}$ that starts at $\lceil \delta n \rceil$, terminates when it reaches n, and has the transition probabilities as shown in Figure 2. This is a biased

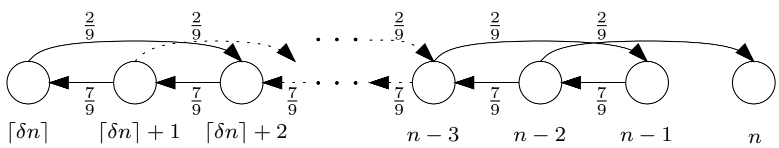

Fig. 2: Transition probabilities of the random walk.

random walk. If we start with an arbitrary matching M satisfying $\chi(M) \leq \delta n$, then one can show by applying a Chernoff bound that the biased random walk takes $2^{\Omega(n)}$ steps with probability $1 - 2^{-\Omega(n)}$ to reach state n.

4 Best Response Dynamics

In this section, we study the best response dynamics in two-sided markets.

Theorem 2. *There exists an instance of the two-sided market problem with three women and three men in which the best response dynamics can cycle.*

Proof. Let w_1, w_2, w_3 denote the women and let m_1, m_2, m_3 denote the men. Let the preference of w_1, w_2, w_3, m_1, m_2, and m_3 be (m_2, m_3, m_1), (m_1, m_2, m_3), (m_3, m_1, m_2), (w_1, w_3, w_2), (w_2, w_1, w_3), and (w_1, w_2, w_3) respectively. We describe a state by a triple (x, y, z), meaning that the first woman is married to the man m_x, the second woman to man m_y, and the third woman to man m_z. A value of -1 indicates that the corresponding woman is unmarried. The following sequence of states constitutes a cycle in the best response dynamics:

$$
\begin{aligned}
&(-1,2,3) \to (3,2,-1) \to (3,1,-1) \to (3,-1,1) \\
&\to (2,-1,1) \to (-1,2,1) \to (-1,2,3) \enspace .
\end{aligned}
$$

Theorem 3. *For every two-sided market instance with n women and n men and every matching M, there exists a sequence of $2n^2$ best responses starting in M and leading to a stable matching.*

Proof. We divide the sequence of best responses into two phases. In the first phase only married women are allowed to change their marriages. If no married woman can improve her marriage anymore, then the second phase starts. In the second phase, all women are allowed to play best responses in an arbitrary order. In the first phase, we use the potential function

$$
\Phi(M) = \sum_{x \in X} (n - p_x(M)) \enspace ,
$$

where X denotes the set of married women. This potential function decreases with every best response of a married woman by at least 1 because this woman increases her payoff and the set X can only become smaller. Since Φ is bounded from above by n^2, the first phase terminates after at most n^2 best responses in a state in which no married woman can improve her marriage.

Now consider the second phase. We claim that if we start in a state M' in which no married woman can improve her marriage, then every sequence of best responses terminates after at most n^2 steps in a stable matching. Assume that we start in a state M' in which no married woman can improve her marriage and that an unmarried woman w plays a best response and marries a man m, leading to state M''. Then the payoff of m can only increase. Hence, man m does not accept proposals in state M'' which he did not accept in M'. This implies that also in M'' no married woman can improve her marriage. Since no married woman becomes unhappy with her marriage, men are never left and therefore they can only improve their payoffs. With every best response one man increases his payoff by at least 1. This concludes the proof of the theorem as each of the n men can increase his payoff at most n times.

Theorem 4. *There exists an infinite family of two-sided market instances $I_1, I_2, I_3, \ldots$ and corresponding matchings $M_1, M_2, M_3, \ldots$ such that, for $n \in \mathbb{N}$, I_n consists of n women and n men and a sequence of random best responses starting in M_n needs $2^{\Omega(n)}$ steps to reach a stable matching with probability $1-2^{-\Omega(n)}$.*

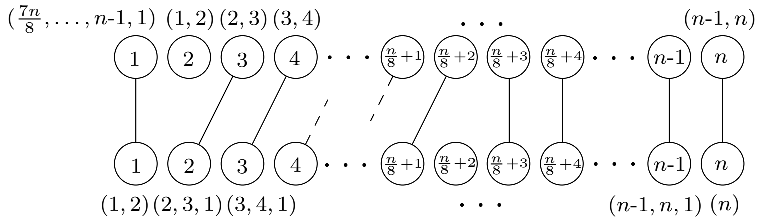

Fig. 3: Nodes in the upper and lower row correspond to women and men respectively. The figure also shows the initial state and the preference lists. The lists are only partially defined, but they can be completed arbitrarily.

Proof. For every large enough $n \in \mathbb{N}$, we construct an instance I_n with n women and n men in which the preference lists and the initial state M_n are chosen as shown in Figure 3.

Let $\mathcal{M}$ denote the set of matchings that contain the edges

$$(w_1, m_1), \ldots, (w_{j-2}, m_{j-2}), (w_j, m_{j-1}), \ldots, (w_k, m_{k-1}),$$
$$(w_{k+1}, m_{k+1}), \ldots, (w_l, m_l), (w_{l+2}, m_{l+1}), \ldots, (w_n, m_{n-1})$$

for some $j < k < l$ with $n/16 \leq k - j \leq n/4$, $k < n/4$, and $l \geq 5n/8$ (cf. Figure 4a). We claim that if one starts in a matching that belongs to $\mathcal{M}$, then with probability $1 - 2^{-cn}$, for an appropriate constant $c > 0$, another matching from $\mathcal{M}$ is reached after $\Theta(n)$ many steps. Since no matching from $\mathcal{M}$ is stable, this implies the theorem.

(a) Matching from $\mathcal{M}$. (b) w_1 proposes to m_k if $\frac{7n}{8} \leq k < n$.

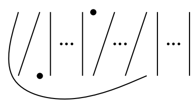

(c) A new diagonal is introduced.

Fig. 4: One phase of the best response dynamics.

If the current matching belongs to $\mathcal{M}$, then there are at most three women who have an incentive to change their marriage. Woman w_{j-1} can propose to man m_{j-1}, woman w_{k+1} can propose to man m_k, and, if $l < n$, woman w_{l+1} can propose to man m_{l+1}. Intuitively, as long as we are in a state that belongs to $\mathcal{M}$, there exists one block of diagonal marriages in the first half, and possibly a second block at the right end of the construction. In every step the left end of the first block, the right end of the first block, and the left end of the second block move with the same probability one position to the right. Since the length of the first block is $\Omega(n)$, one can show by a standard application of a Chernoff bound that the probability that the first block vanishes, i. e., its left end catches up with its right end, before its right end reaches man m_n is exponentially small.

Furthermore, since the distance between the first and the second block is $\Omega(n)$, the probability that the right end of the first block catches up with the left end of the second block before the second block has vanished is also exponentially small.

When the right end of the first block has reached man $m_{7n/8}$, i. e., $m_{7n/8}$ is unmarried, then with probability exponentially close to 1, the second block has already vanished (see Figure 4b) because the initial distance between the two blocks is at least $3n/8$ and only with probability $2^{-\Omega(n)}$ it decreases to $n/8$ before the second block vanishes. As long as the right end of the second block lies in the interval $\{7n/8, \dots, n-1\}$, woman w_1 has an incentive to change her marriage since she prefers m_k with $k \in \{7n/8, \dots, n-1\}$ to m_1. Once she has changed her strategy, a new block of diagonals can be created on the left end of the construction (see Figure 4c). In particular, woman w_1 will only return to m_1 if no man m_k with $k \in \{7n/8, \dots, n-1\}$ is unmarried, that is, she will only return to m_1 if the right end of the first block has reached man m_n. Since it is as likely that a new diagonal at the beginning is inserted as it is that the right end of the block moves one position further to the right, the expected length of the newly created block is $n/8 - 2$. By Lemma 2 it follows that the length of the new block lies with high probability in the interval $[n/16, n/4]$. Only with exponentially small probability the left end of the block has not passed man $m_{5n/8}$ when the right end has reached man m_n because this would imply that the length of the block has increased from at most $n/4$ to $3n/8$. If none of these exponentially unlikely failures events occurs, we are again in a matching from $\mathcal{M}$.

Lemma 2. *Let X be the sum of $n/8$ geometric random variables with parameter $p = 1/2$. There exists a constant $c > 0$ such that*

$$\mathbf{Pr}\left[X \notin [n/16, n/4]\right] \leq 2e^{-cn} \ .$$

Proof. The random variable X is negative binomially distributed with parameters $n/8$ and $1/2$. For a series of independent Bernoulli trials with success probability $1/2$, the random variable X describes the number of failures before the $(n/8)$th success is obtained. For $a \in \mathbb{N}$, let Y_a be a binomially distributed random variable with parameters a and $1/2$. Then

$$\mathbf{Pr}\left[X > n/4\right] = \mathbf{Pr}\left[Y_{3n/8} < n/8\right] = \mathbf{Pr}\left[Y_{3n/8} < \frac{2}{3}\mathbf{E}\left[Y_{3n/8}\right]\right] \leq e^{-cn} \ ,$$

where the last inequality follows, for an appropriate constant $c > 0$, from a Chernoff bound. Furthermore

$$\mathbf{Pr}\left[X < n/16\right] = \mathbf{Pr}\left[Y_{3n/16} > n/8\right] = \mathbf{Pr}\left[Y_{3n/16} > \frac{4}{3}\mathbf{E}\left[Y_{3n/16}\right]\right] \leq e^{-cn} \ .$$

5 Correlated Two-Sided Markets

In this section, we show that, in contrast to general two-sided markets, the convergence time of the random best response dynamics in correlated two-sided markets is polynomial.

Theorem 5. *In every correlated two-sided market the random best response dynamics converges to a stable matching in polynomial time with high probability.*

Proof. We denote by a *round* a consecutive sequence of best responses such that every player is activated at least once. Due to the coupon collector's problem, each round has length $\Theta(n \log n)$ with high probability. Let p denote the highest possible payoff that can be achieved. After the first round there will be a pair $(x, y) \in \mathcal{X} \times \mathcal{Y}$ contained in the matching such that $p_y(x) = p$ because players play best responses. After the value p occurs in the potential function Φ, player x will never leave market y again. Furthermore, x cannot be displaced from y since no player is strictly preferred to x by resource y. Hence, the assignment of x to y can be fixed and we can remove x and y from the game. Now we can inductively apply the same argument to the remaining game. This implies that after at most n rounds a stable state is reached. Hence, the best response dynamics terminates after $O(n^2 \log n)$ steps in expectation and with high probability.

Acknowledgements. We thank Fuhito Kojima and and Yuval Peres for pointing out related work.

References

[AGM+07] Heiner Ackermann, Paul W. Goldberg, Vahab S. Mirrokni, Heiko Röglin, and Berthold Vöcking. A unified approach to congestion games and two-sided markets. In *Proc. of the 3rd Int. Workshop on Internet and Network Economics (WINE)*, pages 30–41, 2007.

[ALMO07] David Abraham, Ariel Levavi, David Manlove, and Gregg O'Malley. The stable roommates problem with globally-ranked pairs. In *Proc. of the 3rd Int. Workshop on Internet and Network Economics (WINE)*, pages 431–444, 2007.

[EO06] F. Echenique and J. Oviedo. A theory of stability in many-to-many matching markets. *Theoretical Economics, Society for Economic Theory*, 1(2):233–273, 2006.

[FGMS06] L. Fleischer, M. Goemans, V. S. Mirrokni, and M. Sviridenko. Tight approximation algorithms for maximum general assignment problems. In *Proceedings of the 16th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA)*, pages 611–620, 2006.

[Fle03] T. Fleiner. A fixed-point approach to stable matchings and some applications. *Mathematics of Operations Research*, 28(1):103–126, 2003.

[GI89] D. Gusfield and R. Irving. *The stable Marriage Problem: Structure and Algorithms.* MIT Press, 1989.

[GLMT04] M. Goemans, L. Li, V. S. Mirrokni, and M. Thottan. Market sharing games applied to content distribution in ad-hoc networks. In *Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 1020–1033, 2004.

[GS62] D. Gale and L. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.

[HHP06] C. Hoffman, A. Holroyd, and Y. Peres. A stable marriage of poisson and lebesgue. *Annals of Probability*, 34(4):1241–1272, 2006.

[KC82] A. Kelso and V. Crawford. Job matchings, coalition formation, and gross substitute. *Econometrica*, 50:1483–1504, 1982.

[Knu76] D. Knuth. *Marriage Stables et leurs relations avec d'autres problèmes Combinatories.* Les Presses de l'Université de Montréal, 1976.

[Mir05] V. S. Mirrokni. *Approximation Algorithms for Distributed and Selfish Agents.* Massachusetts Institute of Technology, 2005.

[Rot84] A. E. Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, 92:991–1016, 1984.

[Rot96] A. E. Roth. The national residency matching program as a labor market. *Journal of American Medical Association*, 275(13):1054–1056, 1996.

[RS90] A. E. Roth and M. A. O. Sotomayor. *Two-sided Matching: A study in game-theoretic modeling and analysis*. Cambridge University Press, 1990.

[RV90] A. E. Roth and J. H. Vande Vate. Random paths to stability in two-sided matching. *Econometrica*, 58(6):1475–1480, 1990.

Aachener Informatik-Berichte

This list contains all technical reports published during the past five years. A complete list of reports dating back to 1987 is available from http://aib.informatik.rwth-aachen.de/. To obtain copies consult the above URL or send your request to: Informatik-Bibliothek, RWTH Aachen, Ahornstr. 55, 52056 Aachen, Email: `biblio@informatik.rwth-aachen.de`

2001-01 *	Jahresbericht 2000
2001-02	Benedikt Bollig, Martin Leucker: Deciding LTL over Mazurkiewicz Traces
2001-03	Thierry Cachat: The power of one-letter rational languages
2001-04	Benedikt Bollig, Martin Leucker, Michael Weber: Local Parallel Model Checking for the Alternation Free mu-Calculus
2001-05	Benedikt Bollig, Martin Leucker, Thomas Noll: Regular MSC Languages
2001-06	Achim Blumensath: Prefix-Recognisable Graphs and Monadic Second-Order Logic
2001-07	Martin Grohe, Stefan Wöhrle: An Existential Locality Theorem
2001-08	Mareike Schoop, James Taylor (eds.): Proceedings of the Sixth International Workshop on the Language-Action Perspective on Communication Modelling
2001-09	Thomas Arts, Jürgen Giesl: A collection of examples for termination of term rewriting using dependency pairs
2001-10	Achim Blumensath: Axiomatising Tree-interpretable Structures
2001-11	Klaus Indermark, Thomas Noll (eds.): Kolloquium Programmiersprachen und Grundlagen der Programmierung
2002-01 *	Jahresbericht 2001
2002-02	Jürgen Giesl, Aart Middeldorp: Transformation Techniques for Context-Sensitive Rewrite Systems
2002-03	Benedikt Bollig, Martin Leucker, Thomas Noll: Generalised Regular MSC Languages
2002-04	Jürgen Giesl, Aart Middeldorp: Innermost Termination of Context-Sensitive Rewriting
2002-05	Horst Lichter, Thomas von der Maßen, Thomas Weiler: Modelling Requirements and Architectures for Software Product Lines
2002-06	Henry N. Adorna: 3-Party Message Complexity is Better than 2-Party Ones for Proving Lower Bounds on the Size of Minimal Nondeterministic Finite Automata
2002-07	Jörg Dahmen: Invariant Image Object Recognition using Gaussian Mixture Densities
2002-08	Markus Mohnen: An Open Framework for Data-Flow Analysis in Java
2002-09	Markus Mohnen: Interfaces with Default Implementations in Java
2002-10	Martin Leucker: Logics for Mazurkiewicz traces
2002-11	Jürgen Giesl, Hans Zantema: Liveness in Rewriting
2003-01 *	Jahresbericht 2002
2003-02	Jürgen Giesl, René Thiemann: Size-Change Termination for Term Rewriting
2003-03	Jürgen Giesl, Deepak Kapur: Deciding Inductive Validity of Equations

2003-04 Jürgen Giesl, René Thiemann, Peter Schneider-Kamp, Stephan Falke: Improving Dependency Pairs
2003-05 Christof Löding, Philipp Rohde: Solving the Sabotage Game is PSPACE-hard
2003-06 Franz Josef Och: Statistical Machine Translation: From Single-Word Models to Alignment Templates
2003-07 Horst Lichter, Thomas von der Maßen, Alexander Nyßen, Thomas Weiler: Vergleich von Ansätzen zur Feature Modellierung bei der Softwareproduktlinienentwicklung
2003-08 Jürgen Giesl, René Thiemann, Peter Schneider-Kamp, Stephan Falke: Mechanizing Dependency Pairs
2004-01 * Fachgruppe Informatik: Jahresbericht 2003
2004-02 Benedikt Bollig, Martin Leucker: Message-Passing Automata are expressively equivalent to EMSO logic
2004-03 Delia Kesner, Femke van Raamsdonk, Joe Wells (eds.): HOR 2004 – 2nd International Workshop on Higher-Order Rewriting
2004-04 Slim Abdennadher, Christophe Ringeissen (eds.): RULE 04 – Fifth International Workshop on Rule-Based Programming
2004-05 Herbert Kuchen (ed.): WFLP 04 – 13th International Workshop on Functional and (Constraint) Logic Programming
2004-06 Sergio Antoy, Yoshihito Toyama (eds.): WRS 04 – 4th International Workshop on Reduction Strategies in Rewriting and Programming
2004-07 Michael Codish, Aart Middeldorp (eds.): WST 04 – 7th International Workshop on Termination
2004-08 Klaus Indermark, Thomas Noll: Algebraic Correctness Proofs for Compiling Recursive Function Definitions with Strictness Information
2004-09 Joachim Kneis, Daniel Mölle, Stefan Richter, Peter Rossmanith: Parameterized Power Domination Complexity
2004-10 Zinaida Benenson, Felix C. Gärtner, Dogan Kesdogan: Secure Multi-Party Computation with Security Modules
2005-01 * Fachgruppe Informatik: Jahresbericht 2004
2005-02 Maximillian Dornseif, Felix C. Gärtner, Thorsten Holz, Martin Mink: An Offensive Approach to Teaching Information Security: "Aachen Summer School Applied IT Security"
2005-03 Jürgen Giesl, René Thiemann, Peter Schneider-Kamp: Proving and Disproving Termination of Higher-Order Functions
2005-04 Daniel Mölle, Stefan Richter, Peter Rossmanith: A Faster Algorithm for the Steiner Tree Problem
2005-05 Fabien Pouget, Thorsten Holz: A Pointillist Approach for Comparing Honeypots
2005-06 Simon Fischer, Berthold Vöcking: Adaptive Routing with Stale Information
2005-07 Felix C. Freiling, Thorsten Holz, Georg Wicherski: Botnet Tracking: Exploring a Root-Cause Methodology to Prevent Distributed Denial-of-Service Attacks
2005-08 Joachim Kneis, Peter Rossmanith: A New Satisfiability Algorithm With Applications To Max-Cut

2005-09 Klaus Kursawe, Felix C. Freiling: Byzantine Fault Tolerance on General Hybrid Adversary Structures

2005-10 Benedikt Bollig: Automata and Logics for Message Sequence Charts

2005-11 Simon Fischer, Berthold Vöcking: A Counterexample to the Fully Mixed Nash Equilibrium Conjecture

2005-12 Neeraj Mittal, Felix Freiling, S. Venkatesan, Lucia Draque Penso: Efficient Reductions for Wait-Free Termination Detection in Faulty Distributed Systems

2005-13 Carole Delporte-Gallet, Hugues Fauconnier, Felix C. Freiling: Revisiting Failure Detection and Consensus in Omission Failure Environments

2005-14 Felix C. Freiling, Sukumar Ghosh: Code Stabilization

2005-15 Uwe Naumann: The Complexity of Derivative Computation

2005-16 Uwe Naumann: Syntax-Directed Derivative Code (Part I: Tangent-Linear Code)

2005-17 Uwe Naumann: Syntax-directed Derivative Code (Part II: Intraprocedural Adjoint Code)

2005-18 Thomas von der Maßen, Klaus Müller, John MacGregor, Eva Geisberger, Jörg Dörr, Frank Houdek, Harbhajan Singh, Holger Wußmann, Hans-Veit Bacher, Barbara Paech: Einsatz von Features im Software-Entwicklungsprozess - Abschlußbericht des GI-Arbeitskreises "Features"

2005-19 Uwe Naumann, Andre Vehreschild: Tangent-Linear Code by Augmented LL-Parsers

2005-20 Felix C. Freiling, Martin Mink: Bericht über den Workshop zur Ausbildung im Bereich IT-Sicherheit Hochschulausbildung, berufliche Weiterbildung, Zertifizierung von Ausbildungsangeboten am 11. und 12. August 2005 in Köln organisiert von RWTH Aachen in Kooperation mit BITKOM, BSI, DLR und Gesellschaft fuer Informatik (GI) e.V.

2005-21 Thomas Noll, Stefan Rieger: Optimization of Straight-Line Code Revisited

2005-22 Felix Freiling, Maurice Herlihy, Lucia Draque Penso: Optimal Randomized Fair Exchange with Secret Shared Coins

2005-23 Heiner Ackermann, Alantha Newman, Heiko Röglin, Berthold Vöcking: Decision Making Based on Approximate and Smoothed Pareto Curves

2005-24 Alexander Becher, Zinaida Benenson, Maximillian Dornseif: Tampering with Motes: Real-World Physical Attacks on Wireless Sensor Networks

2006-01 * Fachgruppe Informatik: Jahresbericht 2005

2006-02 Michael Weber: Parallel Algorithms for Verification of Large Systems

2006-03 Michael Maier, Uwe Naumann: Intraprocedural Adjoint Code Generated by the Differentiation-Enabled NAGWare Fortran Compiler

2006-04 Ebadollah Varnik, Uwe Naumann, Andrew Lyons: Toward Low Static Memory Jacobian Accumulation

2006-05 Uwe Naumann, Jean Utke, Patrick Heimbach, Chris Hill, Derya Ozyurt, Carl Wunsch, Mike Fagan, Nathan Tallent, Michelle Strout: Adjoint Code by Source Transformation with OpenAD/F

2006-06 Joachim Kneis, Daniel Mölle, Stefan Richter, Peter Rossmanith: Divide-and-Color

2006-07 Thomas Colcombet, Christof Löding: Transforming structures by set interpretations

2006-08 Uwe Naumann, Yuxiao Hu: Optimal Vertex Elimination in Single-Expression-Use Graphs

2006-09 Tingting Han, Joost-Pieter Katoen: Counterexamples in Probabilistic Model Checking

2006-10 Mesut Günes, Alexander Zimmermann, Martin Wenig, Jan Ritzerfeld, Ulrich Meis: From Simulations to Testbeds - Architecture of the Hybrid MCG-Mesh Testbed

2006-11 Bastian Schlich, Michael Rohrbach, Michael Weber, Stefan Kowalewski: Model Checking Software for Microcontrollers

2006-12 Benedikt Bollig, Joost-Pieter Katoen, Carsten Kern, Martin Leucker: Replaying Play in and Play out: Synthesis of Design Models from Scenarios by Learning

2006-13 Wong Karianto, Christof Löding: Unranked Tree Automata with Sibling Equalities and Disequalities

2006-14 Danilo Beuche, Andreas Birk, Heinrich Dreier, Andreas Fleischmann, Heidi Galle, Gerald Heller, Dirk Janzen, Isabel John, Ramin Tavakoli Kolagari, Thomas von der Maßen, Andreas Wolfram: Report of the GI Work Group “Requirements Management Tools for Product Line Engineering”

2006-15 Sebastian Ullrich, Jakob T. Valvoda, Torsten Kuhlen: Utilizing optical sensors from mice for new input devices

2006-16 Rafael Ballagas, Jan Borchers: Selexels: a Conceptual Framework for Pointing Devices with Low Expressiveness

2006-17 Eric Lee, Henning Kiel, Jan Borchers: Scrolling Through Time: Improving Interfaces for Searching and Navigating Continuous Audio Timelines

2007-01 * Fachgruppe Informatik: Jahresbericht 2006

2007-02 Carsten Fuhs, Jürgen Giesl, Aart Middeldorp, Peter Schneider-Kamp, René Thiemann, and Harald Zankl: SAT Solving for Termination Analysis with Polynomial Interpretations

2007-03 Jürgen Giesl, René Thiemann, Stephan Swiderski, and Peter Schneider-Kamp: Proving Termination by Bounded Increase

2007-04 Jan Buchholz, Eric Lee, Jonathan Klein, and Jan Borchers: coJIVE: A System to Support Collaborative Jazz Improvisation

2007-05 Uwe Naumann: On Optimal DAG Reversal

2007-06 Joost-Pieter Katoen, Thomas Noll, and Stefan Rieger: Verifying Concurrent List-Manipulating Programs by LTL Model Checking

2007-07 Alexander Nyßen, Horst Lichter: MeDUSA - MethoD for UML2-based Design of Embedded Software Applications

2007-08 Falk Salewski and Stefan Kowalewski: Achieving Highly Reliable Embedded Software: An empirical evaluation of different approaches

2007-09 Tina Kraußer, Heiko Mantel, and Henning Sudbrock: A Probabilistic Justification of the Combining Calculus under the Uniform Scheduler Assumption

2007-10 Martin Neuhäußer, Joost-Pieter Katoen: Bisimulation and Logical Preservation for Continuous-Time Markov Decision Processes

2007-11 Klaus Wehrle: 6. Fachgespräch Sensornetzwerke

2007-12 Uwe Naumann: An L-Attributed Grammar for Adjoint Code

2007-13 Uwe Naumann, Michael Maier, Jan Riehme, and Bruce Christianson: Second-Order Adjoints by Source Code Manipulation of Numerical Programs

2007-14 Jean Utke, Uwe Naumann, Mike Fagan, Nathan Tallent, Michelle Strout, Patrick Heimbach, Chris Hill, and Carl Wunsch: OpenAD/F: A Modular, Open-Source Tool for Automatic Differentiation of Fortran Codes

2007-15 Volker Stolz: Temporal assertions for sequential and current programs

2007-16 Sadeq Ali Makram, Mesut Güneç, Martin Wenig, Alexander Zimmermann: Adaptive Channel Assignment to Support QoS and Load Balancing for Wireless Mesh Networks

2007-17 René Thiemann: The DP Framework for Proving Termination of Term Rewriting

2007-20 Joost-Pieter Katoen, Daniel Klink, Martin Leucker, and Verena Wolf: Three-Valued Abstraction for Probabilistic Systems

* These reports are only available as a printed version.

Please contact `biblio@informatik.rwth-aachen.de` to obtain copies.