Decentralized Matching Schemes for the Navy Detailing Process

Wei Yang Katia Sycara Joseph Giampapa

CMU-RI-TR-03-51

December 2003

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213-3890

© Carnegie Mellon University

Abstract

In addition to the centralized matching algorithms developed in [19], we present two decentralized matching schemes for the Navy detailing process, *decentralized linear programming matching* and *decentralized random matching*, which accommodate the requirements for matching married couples, achieve high fill rate, and fill both priority and undesirable billets. In the decentralized linear programming (DLP) matching scheme Sailors submit applications to a detailer who determines the tentative matching in each stage of a multiple-stage process until the process converges. An exact winner determination model and a heuristic are provided. A near-optimal matching is shown to be achievable. The second scheme, a decentralized random matching (DRM), allows random preference and bilateral negotiation between the Sailor and the Command. We show that it produces a stable matching, just as a centralized deferred acceptance algorithm would. A performance comparison of both centralized and decentralized matching algorithms is presented to provide guidelines on when an algorithm should be preferred and how effective each algorithm is for the Navy detailing process.

Contents

1 Introduction

The US Navy has an active-duty personnel force of 371,800 to manage [7]. A significant part of managing the Navy's manpower assets is assigning Sailors on the job market to fill the available billets, which is called the Navy detailing process. This process involves the Sailors to be assigned, the Commands seeking Sailors and the detailers who make the primary assignment decisions by advising both Sailors and Commands. The Navy detailing process occurs on a biweekly cycle. During this cycle, upcoming vacancies are advertised online through the Job Advertising Selection System (JASS). After reviewing JASS Sailors submit preferences, or applications, to detailers during the biweekly cycle. The Navy currently uses a hierarchical planning method for matching Sailors with billets. This method relies on detailers striking a balance between the Command's needs and the Sailor's preference, which is inherently difficult to achieve [6, 9]. One major drawback of the current detailing practice is that the matching is manual and subject to human errors, making it nearly impossible to optimally match Sailors to billets, as detailers must rely on a "rule of thumb" method to handle the multitude of variables and trade-offs involved. This centralized labor-intensive detailing method leaves many stake holders (e.g. Sailors, detailers and Commands) discontent and frustrated. In some instances, Sailors' preferences are often overridden by the detailer to best fit Commands' needs. This leads to unhappiness with the process, thereby reducing intrinsic motivation and increasing the Sailors' propensity to leave the Navy. By the same token, some Commands have been forced to accept less qualified Sailors to avoid vacancies in key positions, reducing mission effectiveness. The other drawback of the current process lies in its requirement for a clear and deterministic preference list from the market participants. Because of the biweekly frequency of the detailing process and the magnitude of Sailors and billets involved, it is difficult and time-consuming for market participants to determine their matching preferences and behaviors before the assignment begins. Therefore, it is desirable to have a matching scheme that facilitates the dynamic revelation and update of Sailor and Command preferences that would enable the two parties to resolve disagreements automatically and without intermediaries.

There are several special considerations required in the Navy detailing process which impose additional complications to the matching process. A high *fill rate* for the Sailors is expected, which means a matching algorithm should assign as many Sailors as possible to available positions. In addition to regular billets, there are certain *priority billets* requested directly by the Navy Commands that must be filled. Occasionally, Sailors need to be assigned to billets that they do not desire. Finally, the matching process should consider the special case in which married Sailor couples look for *co-locations*.

This paper studies the design of decentralized matching schemes for the Navy detailing process. Our contributions are three-fold: For the first time two decentralized matching algorithms are presented for the detailing process, which is consistent with the proposal by the Department of Navy (DoN) to build a web-based electronic labor market [4]. Second, we complement the work of [19] by presenting alternative matching schemes allowing random preferences and bilateral negotiation, and satisfying most or all of the performance requirements for the detailing process. Finally, the algorithms we developed are analyzed and compared given different performance dimensions, providing guidelines for algorithm preference and for the effectiveness that each algorithm has for the Navy detailing process.

The rest of the paper is structured as follows. In Section 2 we introduce the existing centralized matching methods used for the labor market of medical students and describe their limitations in fulfilling the requirements of the Navy detailing system. Section 3 presents two decentralized matching algorithms for the Navy detailing process and shows how they address the Navy detailing requirements. Finally, in Section 4, we summarize and provide an overview that compares the four algorithms.

2 Existing Centralized Matching Methods

Two-sided matching algorithms have been applied to the medical labor market assigning medical students to hospital positions. We review the deferred acceptance algorithm implemented in

America and the linear programming method achieved field success in Britain.

2.1 Matching by Deferred Acceptance

American hospitals need to seek new medical graduates to fill the internship positions. This entry-level labor market for new American physicians is organized via a centralized clearinghouse called the National Resident Matching Program (NRMP) [10, 12, 13]. Graduating physicians and other applicants first interview at residency programs throughout the country. They then compose and submit Rank Order Lists (ROLs) to the NRMP, each indicating an applicant's preference ordering among the positions for which he has interviewed. Similarly, the residency programs submit ROLs of the applicants they have interviewed, along with the number of positions they wish to fill. Using some appropriate two-sided matching algorithm, the NRMP processes these ROLs and capacities to produce a matching of applicants to residency programs. One commonly implemented two-sided matching method is called the deferred acceptance (DA) algorithm first proposed in [3].

The DA algorithm commonly used in the NRMP cannot satisfy the special requirements for the detailing process. In 2000, upon the completion of the match in the NRMP, 72.3% of the programs were filled and 74.7% of active applicants were matched to a position. The rest of the unmatched applicants must apply for positions through personal channels. Hospitals with unfilled positions even search for international candidates to fill their positions. However, in the Navy assignment situation, it is unacceptable to leave 25% of Sailors without a match and expect them to spend time appealing to various Commands in search of employment. It is equally burdensome to have Commands negotiate with each other to try to fill critical manning needs. Therefore, an alternative matching algorithm may be needed to ensure the high fill rate requirement. In the DA algorithm each agent is treated on the same priority level as every other agent. Using DA in the detailing process cannot guarantee that priority billets will always be matched. One of the key assumptions in the DA algorithm is that agents cannot be forced into a match that they do not find acceptable. Agents have the option to remain unemployed or have their positions unfilled. Forcing Sailors to billets contradicts this vital assumption, which may result in unstable matches and even failure of the process. Although an enhanced DA algorithm is presented in [19] to address the market complications, it still requires a deterministic preference list and the intervention of the detailers. Therefore, an alternative matching method is necessary to accommodate these matching requirements.

2.2 Matching by Linear Programming

As distinguished from the NRMP, the matching schemes developed in two regional hospital systems in Britain, the London Hospital and Medical College and the Cambridge School of Clinical Medicine, involve the linear programming (LP) assignment algorithm [11]. Both schemes take the rank order lists of students and job consultants as inputs and assign numerical weights to their choices. The differences lie in how the weights are assigned. The London scheme sums up the weights of each student-consultant pair and uses it as the basis of the linear programming algorithm so as to maximize the total weights of the match. In the Cambridge scheme, weights are in lexicographic form and ranked by consultant's preferences, so that a consultant-student pair (A,A) has the highest weight, followed by (A,B), (A,C) and so forth. Both matching schemes are still in use. In fact, economical experiments in existing literature show that LP outperforms deferred acceptance in the area of fill rate for both sides of the market, which is a desirable performance in consideration of the high fill rate requirement for the detailing process [17, 5]. In addition, LP achieves an almost equivalent performance in terms of market utility compared to DA. These appealing features indicate that LP may be a better matching mechanism for the Navy detailing process. We adopt the London scheme in our study.

To explore the LP method in detail we use the following notations:

$s_i =$	Sailor indexed for $i = 1, ..., K,\ s_i \in S$.
$n_j =$	billet indexed by $j = 1, ..., M,\ n_j \in N$.
$a_{ij} =$	the weight of matching Sailor i with billet j.
$x_{ij} =$	decision variable. $x_{ij} = 1$ if Sailor i is matched with billet j; $x_{ij} = 0$, otherwise.
$\mu =$	matching between Sailors and billets. $\mu(i) = j$ means Sailor i is matched with billet j.

When Sailor i lists billet j in kth place in his ROL and the same billet ranks the Sailor in lth place, the resulting (i, j) match is called a $[k, l]$ match. Choices $1, 2, 3, 4$ are given weights of 20, 14, 9 and 5 respectively in the London scheme [11]. Thus, a $[1, 1]$ match receives weight $a_{ij} = 40$, $[1, 2]$ and $[2, 1]$ matches each receives weight $a_{ij} = 34$, and so forth. An unlisted choice is given a large negative weight. These matching weights are used as the basis to create the optimal match by the following LP formulation.

$$\begin{aligned} Max \quad & \sum_i \sum_j a_{ij} x_{ij} \\ s.t. \quad & \sum_i x_{ij} \leq 1 \quad \forall\, j && (1) \\ & \sum_j x_{ij} \leq 1 \quad \forall\, i && (2) \\ & x_{ij} \geq 0 \quad \forall\, i, j. && (3) \end{aligned}$$

Proposition 1 *The LP method may produce unstable matchings. Furthermore, it may fail to make* $[1, 1]$ *matches.*

<u>Proof</u> We use a simple matching example to demonstrate the above points. Assume the ROLs of the Sailors and billets are as follows:

$$\begin{aligned} n_1 &: s_1, s_2, s_3 \qquad & s_1 &: n_1, n_2, n_3 \\ n_2 &: s_1, s_3, s_2 \qquad & s_2 &: n_1, n_3, n_2 \\ n_3 &: s_3, s_1, s_2 \qquad & s_3 &: n_3, n_2, n_1. \end{aligned}$$

The unique stable matching is μ such that $\mu(n_1) = s_1, \mu(n_2) = s_2$, and $\mu(n_3) = s_3$ equal a total weight of 98. However, applying the LP method yields the highest total weight of 108 with an unstable matching μ', such that $\mu'(n_1) = s_2, \mu'(n_2) = s_1$, and $\mu'(n_3) = s_3$. The $[1, 1]$ matches are not guaranteed for the Sailor-billet pairs (s_1, n_1) and (s_3, n_3). $\square$

Since the LP is an unstable matching mechanism, it is worthwhile to investigate the reasons for its field success in Britain. Some hypotheses are presented in [11]. One is that the environments where the markets are conducted are so different that there may exist social or other kinds of pressures that make it difficult to circumvent the formal matching scheme. Sailors on the Navy market clearly have less leverage and freedom than the applicants on other free job markets. Any behavior that is not aligned with the formal personnel assignment scheme is discouraged. Another hypothesis to account for the high reported percentage of $[1, 1]$ matches is that agents manage to adapt to the system by coordinating among themselves before the formal match or by modifying the ROLs they submit. Therefore, the match created by the linear programming method can even be stable. In the Navy labor market the detailers function as the advisors to Sailors and Commands. They are the natural source for the coordination before the formal matching. In addition, the Navy labor market is significantly different from the situation in the medical intern market, where medical graduates are generally homogeneous in terms of their skills and in terms of the entry level position to which they are applying. Sailors, on the other hand, look for appropriate billets based on their rate and rating, which refer to the pay grade and occupation specialty, respectively. The use of eligibility criteria in the detailing process decreases the potential for conflict between preferences and helps the LP method create more satisfying and even stable matchings.

3 Decentralized Matching

Solving the decentralized general assignment problem has been investigated by [1, 8, 2, 18], where often there exists a unilateral preference of the agents on one side over those on the other. In this paper, we focus on the design of decentralized matching schemes for the detailing process. The market includes a set of Sailors and a set of billets with a bilateral preference over each other. In the centralized deferred acceptance (DA) or linear programming (LP) algorithm, Sailors and Commands (for their billets) must decide what preference lists to submit. After the submission of preferences, offers, acceptances, and rejections are carried out automatically in a central clearinghouse. A decentralized matching scheme does not require participants to submit their rank order lists (ROLs) before the matching; instead they can decide what to do at each stage, based on current and past matching information. These random updates resolve disagreements and conflicts over the course of time and ultimately are resolved in satisfying matchings.

We use the following additional notations:

$p_{ij} =$ the ask price of Sailor i applies for the billet j,

$\bar{p}_{ij} =$ the upper limit of willing-to-pay of billet j for Sailor i,

$r_{ij} =$ the reservation value of Sailor i over billet j.

$\epsilon =$ pre-defined decrement in ask price.

Before their application, Sailors evaluate a variety of aspects related to each billet, such as location, duty type (sea or shore), and potential promotions. The reservation value r_{ij} is the minimum acceptable payment for Sailor i to take billet j. It reflects the Sailor's preference over the billet. The more disagreeable the billet, the higher the reservation value. The ask price p_{ij} is the payment Sailor i expects to get by taking billet j. The maximal willing-to-pay $\bar{p}_{ij}$ represents the Command's evaluation over the eligibility of Sailor i based on his rate and rating. The more qualified and suitable a Sailor is, the higher payment he expects to get.

We present two different decentralized matching schemes. The first one is basically a decentralized version of the LP method (we call it *decentralized linear programming* (DLP)) where Sailors apply billets through a detailer who is aware of Commands' preferences. The detailer plays a role as a monitor or coordinator of the matching process carried out in multiple stages. In each stage the detailer collects applications and determines a provisional match. The match ends when no Sailor submits an application. In the second matching scheme, there is no detailer who acts as the middle man between Sailors and Commands. A Sailor contacts a Command directly to negotiate over possible matchings. The Sailor and the Command are tentatively committed to each other if the Sailor's application is accepted at the current stage. The Sailor might be rejected later when the Command chooses a more qualified candidate. He then applies for an alternative billet. The match ends when there is no Sailor left who has not already been rejected by all of his achievable Commands. This scheme does not require deterministic preference lists from the agents. We call it *decentralized random matching* (DRM).

3.1 Decentralized Linear Programming Matching

There are certain matching rules to obey in the decentralized linear programming (DLP) method, which are discussed as follows.

3.1.1 Applying Rules

Sailors participate in the matching process by logging onto the network through a computer so that they can be identified by the matching system. There is a detailer who monitors and coordinates the matching process. Commands submit their $\bar{p}_{ij}$ to the detailer before the assignment begins. Sailors' reservation values over billets r_{ij} are private information and are not revealed to the detailer and Commands. Sailors communicate with the detailer but not with each other. All the prices and valuations are assumed to be integers. The matching is conducted progressively within a fixed time window in multiple stages.

The detailer announces the initial ask price p_{ij} for all the Sailors in the matching system, which are set much larger than r_{ij} based on conjectured reservation values from previous experience. Each Sailor then decides which billets he wants to bid at the initial ask price. We assume that Sailors follow the *myopic best response* (MBR) policy. That is, each Sailor bids the m billets whose value to him exceeds the corresponding reservation value by the most amount. No manipulative behavior is considered in this study where a Sailor may choose a billet that is not within his list of the m most preferred billets. As discussed below, the detailer collects all the applications and determines the best provisional matching, ensuring that each billet is assigned to at most one Sailor and each Sailor is matched with at most one billet. If it is possible to assign each billet to a Sailor who demands it, the matching already has the desired equilibrium. If not, there must exist a set of billets such that the number of Sailors demanding billets only in this set is greater than the number of billets in the set. Such a set is called an *over-demanded* set. At this point any unmatched Sailor i decreases his ask prices p_{ij} for all the m billets applied in a previous stage by a constant ϵ. Sailor i then reevaluates the billets based on the updated p_{ij}, resubmits the application (including the updated p_{ij} to the detailer), and expects to be matched in the next stage. If Sailor i is matched with a billet j in the next stage, he is allowed to keep the same ask prices for the billets to which he applied. Informed by the detailer, Sailor i', who previously matched with billet j, becomes unassigned. Following the same rules, Sailor i' decreases his ask price by ϵ and resubmits his application, and so on. It is clear that the matching process cannot go on indefinitely because once the ask price p_{ij} becomes lower than the reservation value r_{ij}, billet j will not be preferable to Sailor i any more, and eventually it cannot belong to any over-demanded set. The matching is terminated when no Sailor submits an application and the provisional matching remains the same as in the previous stage. It follows that for some set of ask prices there will be an equilibrium matching.

3.1.2 Winner Determination

After announcing the initial ask price and gathering the applications from all the Sailors, the detailer needs to solve the winner determination problem to determine the matching between

Sailors and billets. The goal is to provide a feasible provisional matching, maximized in total payoff for both Sailors and Commands. An exact solution can be obtained by solving the following mathematical formulation.

$$
\begin{aligned}
Max \quad & \sum_i \sum_j (\bar{p}_{ij} - p_{ij}) x_{ij} & \\
s.t. \quad & \sum_i x_{ij} \leq 1 \quad \forall\, j & (4) \\
& \sum_j x_{ij} \leq 1 \quad \forall\, i & (5) \\
& x_{ij} \geq 0 \quad \forall\, i, j. & (6)
\end{aligned}
$$

The objective coefficients $v_{ij} = \bar{p}_{ij} - p_{ij}$ are functionally equivalent to the weights a_{ij} in the centralized LP algorithm, where v_{ij} can be considered as the payoff Sailor i and billet j can achieve by matching each other. However, as distinguished from the uniform assigned weights a_{ij} that are only based on the rank order, the payoff v_{ij} is "discriminatory" over different billets and Sailors. Commands specify their preferences over Sailors by paying them different salaries. Sailors rank billets by showing different minimal acceptable payments. Billets with the same rank order but various requirements in moving cost, skills and training level and etc., can be differentiated by $\bar{p}_{ij}$ as well. In the detailing process this would be the expected situation.

The detailer needs to solve the winner determination problem in each round. To save this demanding computational work, a heuristic is developed to approximate the optimal winner determination solution. It is basically a greedy heuristic for solving the set covering problem, which is described as follows.

Step 0: The detailer sorts all Sailors on v_{ij} in an ascending order for each billet.

Let $j = 1$.

Step 1: For Sailors i and k, and for billet j,

If $\mu(i) = 0$ and $v_{ij} \geq v_{kj}, k \neq i$,

$\mu(i) = j$.

Step 2: If $\mu(i) \neq 0$ for all i or $j = M$,

stop and the matching is completed;

else,

$j = j + 1$ and go to Step 1.

This decentralized linear programming (DLP) mechanism is appealing to the Sailors because it does not require them to decide in advance exactly what their bidding behaviors and preference lists would be. Instead, at each round a Sailor can make use of present and past stages of the matching process to decide his next bid. The final ask price for a billet will differ from the minimum equilibrium price by at most $k\epsilon$, where k is the minimum of the number of Sailors and billets [2]. Since the maximum willing-to-pay $\bar{p}_{ij}$ is fixed during the auction, the final payoff of each Sailor-billet match will differ from the minimum equilibrium payoff by at most $k\epsilon$. It is worth noting that the outcome approaches the optimal solution for small values of ϵ.

Proposition 2 *For the simple Navy detailing process, the final payoff of any Sailor-billet match determined by the decentralized linear programming matching will differ from the unique minimum equilibrium payoff by at most $k\epsilon$.*

<u>Proof</u> This result is shown in [2] for the descending auction protocol in an exchange economy, where buyers want no more than one item over a set of available goods. Such is the case for the simple one-to-one Navy detailing process. □

3.1.3 Addressing the Navy Detailing Requirement

The DLP is in essence the decentralized version of the linear programming method that achieved field success in Britain. Therefore, DLP still retains the advantage of reaching high fill rates for the Sailors and Commands. It is clear that the higher v_{ij} is, the more likely $x_{ij} = 1$. Therefore, Commands can ensure that priority billets are matched with qualified Sailors by setting the $\bar{p}_{ij}$ of priority billet j relatively higher than those of the other regular billets. A similar approach can also be applied to undesirable billet j so that incentives in the monetary form can be added to $\bar{p}_{ij}$ to attract more applications.

The problem of married couple looking for co-locations is addressed in [19] where couples, who have preference over pairs of billets, are treated differently from the single sailors. Similarly, we make modifications on the definitions of prices, valuations and application rules. Sailors are divided into singles and married couples. The matching process for singles remains the same as described above. Each married couple has reservation values over pairs of billets considered acceptable to the couple. We call such pairs bundled billets. On the other hand, Commands have their willing-to-pay for the bundled billets of different married couples. One member of each married couple, representing the couple, participates in the matching process by submitting ask prices to the detailer, who then determines the provisional matching by allocating single billets and bundled billets to the applicants. The only complexity lies in the winner determination by the detailer, which can be solved in exactly the same way as in combinatorial auctions [18, 8].

3.2 Decentralized Random Matching

The DLP method satisfies most requirements for the Navy detailing process. However, there exists a central controller involved in the matching process preventing the direct negotiation between the Sailor and the Command. The main goal of the DoN's proposal to reform the Navy's enlisted distribution system is to reduce human intervention and subjectivity. Therefore, it is necessary to circumvent any central controller and allow bilateral negotiation between the agents. We believe that the *decentralized random matching* (DRM) scheme incorporating *preference randomness* and bilateral negotiation improves the current detailing practice in the right direction.

3.2.1 DRM Algorithm

To fully understand the DRM algorithm, we introduce some terminology. For a given matching μ, a Sailor s and a billet n is a *blocking pair* if they are not matched to each other, but they each prefer one another to their current match in μ. A matching is a *stable matching* if no blocking pair exists. If (s', n') is a blocking pair for a matching μ, we say that a new matching μ' is obtained from μ by *satisfying* the blocking pair if s' and n' are matched to each other at μ', their matches at μ (if any) are unmatched at μ', and all other agents are matched to the same agents at μ' as they were at μ.

It is shown that the centralized DA algorithm creates nonempty stable matchings. Blocking pairs are satisfied one by one until a stable matching is reached. A natural question is whether there exists at least one such path leading to a stable matching from any initial matching and any preference lists of the agents. The question is addressed in [15] and [16] with a theorem presented as follows.

Proposition 3 *(Roth and Vande Vate) Let μ be an arbitrary matching for (S, N), then there exists a finite sequence of matchings $\mu_1, ..., \mu_k$, where μ_1 is the initial matching and μ_k is stable, for each $i = 1, ..., k-1$, there is a blocking pair (s_i, n_i) for μ_i such that μ_{i+1} is obtained from μ_i by satisfying blocking pair (s_i, n_i).*

The implication of this theorem for the Navy detailing process is that starting from any arbitrary matching, a new matching can be achieved by randomly satisfying a blocking pair through the bilateral negotiation between the Sailor and the Command, eventually converging to a stable matching. Table 1 presents a decentralized random matching algorithm for the detailing process. Any available Sailor who is unmatched and has not been rejected by all the Commands he prefers, attempts to apply to a Command for his preferred billet. If the billet that the Sailor applies for is unfilled and he is qualified for the position, then the billet is tentatively assigned to the Sailor. If the billet is matched with some other Sailor who is more qualified, then the Command will reject the Sailor's application. Otherwise, the Command will accept the application and the previously matched Sailor becomes unassigned. The process continues until there is no unmatched Sailor who has not been rejected by all his preferred Commands or the matching deadline is exceeded. The algorithm in Table 1 is a Sailor proposing version. A Command proposing matching algorithm can be derived accordingly.

Proposition 4 *The decentralized random matching procedure without matching time limitation creates the same stable matching as the outcome of the centralized deferred acceptance procedure. In particular, it creates the Sailor-optimal stable matching with the revealed preference.*

Proof The argument is similar to the standard proof from the centralized DA algorithm. Regardless of the random elements involved in Step 1 in Table 1, the matching created by DRM is stable with respect to the revealed preference because there can be no blocking pairs for the final matching. If a Sailor prefers a billet to his current position, he must have already applied to the billet and been rejected.

We show the matching is Sailor-optimal by proving that no Sailor is ever rejected by a billet to which he could be matched at some stable matching. It can be shown by induction. Assume that for a given step in the procedure no Sailor has yet been rejected by a billet that is achievable for him. At this step, assume that billet n rejects Sailor s. If n rejects s as unacceptable, then n is not achievable for s, and we are done. Suppose billet n rejects Sailor s in favor of s' who is assigned the billet, then n prefers s' to s. We must show that n is not achievable for s. We know s' prefers n to any billet except for those who have rejected him, and which are therefore, by inductive assumption, not achievable for him. Consider a hypothetical matching μ that matches s with n and everyone else to an achievable agent. Then s' prefers

Inputs:	Randomly selected ask price p_{ij} and willing-to-pay $\bar{p}_{ij}$. Detailing process deadline $\bar{t}$.
Output:	A stable matching μ between Sailors and billets.
Preconditions:	All billets unfilled, all Sailors are unmatched.
Postcondition:	A matching filling priority billets, undesirable billets, possible high fill rate and high satisfaction for both Sailors and Commands.
Algorithm:	
Step 0:	Initialization: Random ask price p_{ij} for each Sailor i over billet j; Random willing-to-pay $\bar{p}_{ij}$ for each billet j over Sailor i. Initial matching: $\mu = \emptyset$. $t = 0$.
Step 1:	Any Sailor i with $\mu(i) = 0$ and there is at least one billet that has not rejected him? If no, go to Step 2; else, Set $t = t + 1$. If $t \geq \bar{t}$, go to Step 3; else, Sailor i attempts to apply to the billet j with highest $\bar{p}_{ij} - r_{ij}$ who has not rejected him. If $v_{ij} \leq 0$, where $v_{ij} = \bar{p}_{ij} - p_{ij}$, Go to Step 1a. If $\mu(j) = 0$, $\mu(i) = j$; else if $\mu(j) = k$ for some other Sailor k, and $v_{ij} > v_{kj}$, $\mu(i) = j$; else $\mu(i) = 0$.
Step 1a:	Consider the next billet that has not rejected i. Go to Step 1.
Step 2:	STOP. In this case the final outcome is the matching μ that matches each Sailor to the billet (if any) he prefers.
Step 3:	After-market negotiation when deadline $\bar{t}$ is exceeded.

Table 1: A decentralized random matching algorithm for the Navy detailing process

n to his matched billet in μ. So the matching μ is unstable and blocked by (s', n). Therefore there is no stable matching that matches s and n, and they are not achievable for each other. □

3.2.2 Addressing the Navy Detailing Requirement

As shown above, the DRM algorithm can achieve stable matchings by bilateral negotiations between a Sailor and Command with random preference lists. Since market participants make assignment decisions based on the matching payoff v_{ij}, similar approaches used in decentralized linear programming (DLP) to satisfy priority billets and force Sailors to undesirable ones can be applied to decentralized random matching (DRM) as well, by adjusting a Command's willing-to-pay $\bar{p}_{ij}$. In addition, direct communication between a Sailor and a Command enable possible leeway in eligibility criteria and compensation rules. A high fill rate is more likely to be achieved. A centralized enhanced deferred acceptance (DA) algorithm is provided in [19] for the Navy detailing process with both singles and couples. The approaches for handling the couples in that algorithm can be used in the DRM. However, further investigation is required to resolve stability issues in the resultant matchings.

4 Conclusion

We investigated the design of decentralized matching algorithms. As distinguished from the centralized matching algorithm in [3], [14] and [19], where agents must decide what preference lists to submit, a decentralized matching scheme does not require participants to determine their applying behavior and submit their preferences before the matching; instead, they can decide what to do at each stage based on the current and past matching information. In the centralized matching procedure a clearinghouse carries out offers, acceptances and rejections automatically. On the other hand, using bilateral negotiations, decentralized matching resolves disagreements and conflicts between agents from opposite market sides over the course of time.

	Centralized Matching		Decentralized Matching	
	DA	LP	DLP	DRM
Stable Matching	Yes	No	No	Yes
Random Preference	No	No	Yes	Yes
Bilateral Negotiation	No	No	No	Yes
Matching Couples	Yes	No	Yes	Yes
Matching Priority Billets	No	No	Yes	Yes
High Fill Rate	No	Yes	Yes	Yes
Matching Undesirable Billets	No	No	Yes	Yes

Table 2: Feature summary of different matching algorithms for the Navy detailing process

Agents finally converge in satisfying matches.

We present two decentralized matching schemes for the Navy detailing process in which Sailors apply for available billets. In the first scheme, DLP, Sailors submit applications to a detailer who determines the tentative matching in each stage. An exact winner determination model and a heuristic are provided. A near-optimal matching is shown to be achievable by DLP. The second scheme, DRM, allows agents to process random preferences and to make decisions based on the current market information. No central controller exists in the market in DRM; instead, Sailors and Commands negotiate directly with each other. We show that it produces the same stable matching as that produced by the centralized deferred acceptance algorithm.

Table 2 presents the performance comparison of different matching schemes for the Navy detailing process. Among the centralized matching algorithms, LP can achieve high fill rate for Sailors. DA can handle the matching market with both singles and couples. However, they cannot address other special requirements in the detailing process effectively. DLP can deal with most market complications. However, it does not allow bilateral negotiations and no stable matching is ensured. To create a stable matching with random preference through direct negotiation between market participants and accommodate the operational requirements, we believe that DRM is the best alternative for the Navy detailing process.

Acknowledgements

The authors would like to thank Michael D. Rectenwald for his editorial assistance and Cuihong Li for her helpful suggestions. This work was supported in part by the Navy Personnel Research, Studies and Technology (NPRST) under Navy contract number N6610-98-D-9501 and by a subcontract with Resource Consultants, Inc (RCI), subcontract number 03-177-7048/086.

References

[1] D.P. Bertsekas, D.A. Castanon, and H. Tsaknakis. Reverse auction and the solution of inequality constrained assignment problems. *SIAM Journal on Optimization*, 3:268–299, 1993.

[2] G. Demange, D. Gale, and M. Sotomayor. Multi-item auctions. *The Journal of Political Economy*, 94(4):863–872, August 1986.

[3] D. Gale and L. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, January 1962.

[4] B. Gates and M. Nissen. Two-sided matching agents for electronic employment market design: Social welfare implications. Technical report, Naval Postgraduate School, Monterey, California, 2003.

[5] J.H. Ho and E.H. Low. Two-sided matching for the us navy's enlisted detailing process: A comparison of deferred acceptance and linear programming via simulation. Master's thesis, Naval Postgraduate School, Monterey, California, 2002.

[6] G.K.W. Koh. A redesign of the navy enlisted personnel distribution process. Master's thesis, Naval Postgraduate School, Monterey, California, 2002.

[7] Department of the Navy. Fiscal year (fy) 2001 budget estimates. Technical report.

[8] D.C. Parkes. *Iterative Combinatorial Auctions: Achieving Economics and Computational Efficiency*. PhD thesis, University of Pennsylvania, 2001.

[9] P. Robards. Applying two-sided matching processes to the united states navy enlisted assignment process. Master's thesis, Naval Postgraduate School, Monterey, California, 2001.

[10] A. Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, 92(6):991–1016, 1984.

[11] A. Roth. A natural experiment in the organization of entry-level labor market: Regional markets for new physicians and surgeons in the united kingdom. *The American Economic Review*, 81(3):415–440, June 1991.

[12] A. Roth. The economist as engineer: Game theory, experimentations, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, July 2002.

[13] A. Roth and E. Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *The American Economic Review*, 89(4):748–780, September 1999.

[14] A. Roth and M. Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modelling and Analysis.* Econometric Society Monograph Series. Cambridge University Press, 1990.

[15] A. Roth and J. Vande Vate. Random paths to stability in two-sided matching. *Econometrica*, 58(6):1475–80, 1990.

[16] A. Roth and X. Xing. Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists. *The Journal of Political Economy*, 105(2):284–329, April 1997.

[17] M.U. Unver. On the survival of some unstable two-sided matching mechanisms in entry-level medical labor markets. Department of Economics, Koc University, November 2003.

[18] M.P. Wellman, W.E. Walsh, P.R. Wurman, and J.K. Mackie-Mason. Auction protocols for decentralized scheduling. *Games and Economic Behavior*, 35:271–303, 2001.

[19] W. Yang, J. Giampapa, and K. Sycara. Two-sided matching for the navy detailing process with market complication. CMU-RI-TR-03-49, November 2003.