A Generalized Matching Framework : Combining Matchmaking and Coalition Formation

Jeroen M. Valk [a] Duco N. Ferro [ab]

[a] *Almende B.V., Westerstraat 50, 3016 DJ Rotterdam*
[b] *Vrije Universiteit Amsterdam, De Boelelaan 1081a, 1081 HV Amsterdam*

Abstract

This paper presents a matchmaking framework for a new matchmaking component in the ASK communication platform (http://www.ask-cs.nl/). The framework allows a variety of matchmaking techniques to be applied in a single automated matchmaking system. A variety of settings are generalized in a way that combines coalition formation with typical matching settings such as perfect/maximal matching and bi/tri-partite matching. Our framework turns out to be useful to categorize existing work on coalition formation and matching, but at the same time allows us to identify new domains to be studied.

1 Introduction

The problem of matching together individuals in an effective way has been studied in a variety of different settings, e.g., perfect/maximal matching [5, 6], bi/tri-partite matching [14, 9, 10], or coalition formation [8, 1, 7, 11]. Originally, matchmaking is concerned with bringing together individuals into couples [19, 2]. This particular setting is commonly represented using a graph. The nodes of the graph are individuals and two individuals that are connected by an edge could potentially form a match. In particular, a special case where the matching graph is bipartite is studied extensively. This so-called bipartite matching problem has many applications.

Many situations occur in everyday life where one group must be mapped on another, e.g., employees/jobs, patients/doctors, consumers/products. At a first glance, it seems logical to model these matching problems using bipartite graphs. However, even in typical bipartite circumstances it is often useful to take additional dimensions into account. A patient, for example, goes to the doctor for certain reason, i.e., a certain expertise is required. This extra dimension should be taken into account, because the quality of the match depends on it. While the latter type of matching is known as 3-dimensional matching, theory on matching in more than three dimensions can, as far as we know, only be found in the area of coalition formation.

Coalition formation was originally studied in cooperative game theory. However, coalitions are not formed to bring together individuals on a certain occasion, but rather to suggest possible cooperations between players in a non-cooperative game. If two players are in the same coalition this means that they are fully cooperative in order to maximize their collective payoff. That is, the members of a coalition act as a single player in the game. Since players in a non-cooperative game can always improve on their collective payoff by cooperating, it was traditionally assumed that players will always form the largest possible coalition [1]. The main concern of cooperative game theory is therefore how to divide the collective gain over the members of the coalition.

More recently, there is an interesting trend in research on coalition formation to dispute the traditional arguments. For reasons such as, e.g., bounded rationality, incomplete information, or communication overhead in large coalitions, it is argued that it is sometimes undesirable to form large coalitions. As a result, there is an increasing interest in games that are not superadditive. In our opinion, this trend makes that matchmaking and coalition formation become more and more related. Some researchers recognized this relationship between matching and coalition formation.

[1] The precise argument in traditional game theory is based on the observation that all games that are put in coalitional form are superadditive.

In [16], for example, pair partnership matching of elements from an unordered group is treated as a variant of the general coalition formation process. In pair coalition formation an agent is satisfied with a coalition of itself and only one other partner.

A comprehensive framework that combines the many aspects of matchmaking and coalition formation is still missing however. The aim of this paper is to fill in this gap. We present a matching framework that concerns matching structures which contain variable-sized groups of individuals and support different (possibly conflicting) opinions on the effectiveness of groups. The framework (i) generalizes existing matchmaking and coalition formation models, (ii) allows new domains to be studied, and (iii) is used for the implementation of a matchmaking system developed at Almende B.V., a research company situated in Rotterdam. The outline of the paper is as follows. First, Section 2 presents the generalized matchmaking framework. Next, Section 3 shows how the framework is related to existing models from the literature. Then, in section 4 we give an illustration of how the framework can be applied in a healthcare domain and we end in Section 5 with the conclusions.

2 A Generalized Matching Framework

In this section we present a matching framework that generalizes existing approaches to matchmaking and cooperative game theory. First we provide some basic notions in our model, i.e. the matching environment, matching evaluation, the matching structure and control mapping. Then we define a notion of (un)stability. Finally we give some properties each representing different aspects of matching.

2.1 The Basic Framework

In many applications, matchmaking is driven by the fact that individuals (e.g., users or other entities) request to be matched to other individuals or items. In the health care domain, for example, matchmaking concerns users requesting a particular medical treatment, preferably, by the best doctor available.

In its most general form, the matching problem we are interested in consists of finding an effective grouping of elements. We distinguish *unordered* and *ordered* groups. Given a countable set E of elements, an unordered E-group is simply a finite subset of E. This means that the group has no structure at all. An ordered E-group, on the other hand, is a vector $(e_1, \ldots, e_i, \ldots, e_n)$ such that each $e_i \in E$. This means that each element has a certain role; the total order which is associated with an ordered group is used to establish the roles of the elements in the group, e.g. buyer or seller, patient or care provider, etc.

Matching elements are either *passive* elements, e.g. a book or a service, or *active* elements that can be considered parts of which a particular individual (agent) has full *control*. Some individuals may be represented by more than one element. These indviduals are said to be in control of these elements. The idea is that active elements, in contrast to passive elements, take into account the opinion, which their controlling individual has about the groups that are formed in the matching process. The active elements of which an individual has control represent the different roles that the individual possibly has, e.g., employee, patient, owner, seller, father, etc. The quadruple, containing all individuals, active elements, passive elements and the control function for active elements, is called the *matching environment* , or sometimes simply called the *environment*.

Definition 2.1 (Matching Environment) *A (matching) environment* $M = (\mathcal{A}, A, B, U)$ *is a quadruple with:*

(i) $\mathcal{A}$ *the set of agents that exercise control, and*

(ii) A *a countable set of all active elements, and*

(iii) B *a countable set of all passive elements, and*

(iv) U *a control mapping, given by the surjective function* $U : A \rightarrow \mathcal{A}$ *that maps active elements to agents.*

Note, we will write E to refer to the set of all elements, both active and passive. If the control function U is injective as well then it is said to be *localized*. In this situation, we will assume U to be the identity mapping on A, i.e., $\mathcal{A} = A$ and $U(a) = a$ for all $a \in A$. This means that every agent is represented by one active element exactly.

Given a matching environment, matchmaking is concerned with finding a particular configuration of elements into groups. The configuration must be such that the commitments of the elements with respect to the group are clear. For this reason, we adopt an assumption that is common in the literature: i.e., each element can belong to only one group at the same time [17, 13]. In this way, the element can be fully committed to the group to which it belongs. Note that, in our framework, this means that agents can still have control over elements in different groups.

Definition 2.2 (Matching Configuration) *Given a matching environment $M = (\mathcal{A}, A, B, U)$, with $E = A \cup B$, a (matching) configuration in M is a set $\mathcal{C}$ of E-groups such that every $e \in E$ occurs in exactly one member of $\mathcal{C}$.*

In the literature, the concept of a matching configuration appears frequently, although under different names. Cooperative game theory, for example, often uses the term *coalition structure* [20]. In combinatorics the matching configuration is known under the name *matching* or *perfect matching* [6].

Not all matching configurations are equally good. In general, there may be different opinions on the value of a matching configuration. We define an M-opinion as follows, assuming that an opinion on a matching configuration is formed by the opinions on the individual groups in the configuration.

Definition 2.3 (M-opinion) *Let $M = (\mathcal{A}, A, B, U)$ be a matching environment. An M-opinion is a pair $(\mathcal{V}, v)$ where $\mathcal{V} \subseteq \bigcup \Gamma$ and $v : \mathcal{V} \to \mathbb{R}$.*

The part of cooperative game theory that deals with so-called transferable utility(TU) games assumes that there is only one opinion and that this single opinion is complete [12]. The reason being that in a TU game all utilities are assumed to be monetary and therefore any disagreement about the value of a coalition can be eliminated using side-payments, i.e., changing the distribution of the collective gain over the members of a coalition.

Not everything can be compensated by side-payments however. In such settings, different opinions on the values of coalitions can be distinguished. In an environment where different opinions exist, what actually happens depends on how control is distributed over the different opinions. We define an M-evaluation as a family of M-opinions where the index set, represented by the set of agents, of the family is used to identify the different opinions.

Definition 2.4 (M-Evaluation) *Let $M = (\mathcal{A}, A, B, U)$ be a matching environment. An M-evaluation is a finite family $V = \{(\mathcal{V}_a, v_a)\}_{a \in \mathcal{A}}$ of M-opinions.*

It is not always possible that all mtaching configurations may form. There may be all sort of constraints on possible combinations ofcoalitions. The collection of all possible combinations of matching configurations that might be formed is called the *skill set* for the environment.

Definition 2.5 (Matching Skills) *A collection Γ of matching configurations in M is called a skill set for M.*

We can now describe the matching problem in terms of matching environments M, skill sets Γ for these environments and M-evaluations.

Definition 2.6 (Matching Problem) *The matching problem is: given a matching instance $I = (M, \Gamma, V)$ with $M = (\mathcal{A}, A, B, U)$ a matching environment, Γ a skill set, and V a M-evaluation, find a matching configuration $C \subseteq \Gamma$ with some desirable properties.*

In the next section, we will introduce one such property, which we left unspecified in the above definition.

	v_X	v_Y	v_x	v_y		v_X	v_Y	v_x	v_y
$\{X\}$	4	2	0	0	$\{X, y\}$	9	4	4	9
$\{Y\}$	2	4	0	0	$\{Y, y\}$	4	8	3	6
$\{x\}$	0	0	4	2	$\{X, x\}$	6	3	8	4
$\{y\}$	0	0	2	4	$\{Y, x\}$	0	2	2	0

Table 1: an M-evaluation

2.2 Strongly Unstable Matching Configurations

An interesting property of a matching configuration is its (un)stability. Intuitively, a matching configuration is unstable, if there exists a coalition of agents that has enough influence over the agents outside the coalition to deviate from the matching configuration in such a way that each agent in the coalition is individually better off. This means that unstability depends on the possibility of a transition $\mathcal{C} \stackrel{\mathcal{A}_0}{\to} \mathcal{C}'$ influenced by $\mathcal{A}_0 \subseteq \mathcal{A}$ from matching configuration $\mathcal{C}$ to $\mathcal{C}'$. It is difficult to provide an exact characterization of the above transition relation, because its depends on the strategy adopted by the agents outside the coalition whether the transition will succeed or not. Nevertheless, sufficient conditions can be given for which the transition is guaranteed to succeed under the assumption of rational agents.

A sufficient condition on unstability that is often provided in the literature can be formulated as follows. Let $\mathcal{A}_0 \subseteq \mathcal{A}$ be a coalition of agents. We now only consider transitions to matching configurations where only isolated regroupings among the active elements under control of $\mathcal{A}_0$ are allowed. That is, if $\mathcal{C} \stackrel{\mathcal{A}_0}{\to} \mathcal{C}'$ then there exists a subset $\mathcal{C}_0 \subseteq \mathcal{C}'$ such that for every $e \in E$, e occurs in $\mathcal{C}_0$ iff e is an active element under control of $\mathcal{A}_0$, i.e., $U(e) \in \mathcal{A}_0$. Under the above restriction, $\mathcal{A}_0$ can guarantee that agents outside the coalition can only regroup the partial matching configuration $\mathcal{C}' \setminus \mathcal{C}_0$. This means that $\mathcal{A}_0$ can guarantee a new coalition structure that belongs to $\Gamma_0 = \{\mathcal{C} \mid \mathcal{C}_0 \subseteq \mathcal{C}\}$. If in all these possibilities, each member of $\mathcal{A}_0$ is better off then $\mathcal{C}$ must have at least one outgoing transition, and therefore, $\mathcal{C}$ is unstable.

The above considerations can be summarized in the following very strong notion of unstability.

Definition 2.7 (Strong Unstability) *Let $I = (M, \Gamma, V)$ be a matching instance with $M = (\mathcal{A}, A, B, U)$. A matching configuration $\mathcal{C} \in \Gamma$ is called I-unstable if there exists a subset $\mathcal{A}_0 \subseteq \mathcal{A}$ and a partition $\mathcal{C}_0$ of $A_0 = \{a \in A \mid U(a) \in \mathcal{A}_0\}$ such that $\Gamma_0 = \{\mathcal{C}' \in \Gamma \mid \mathcal{C}_0 \subseteq \mathcal{C}\}$ is non-empty and:*

$$\sum_{C \in \mathcal{C} \cap \mathcal{V}_a} v_a(C) < \min_{\mathcal{C}' \in \Gamma_0} \sum_{C \in \mathcal{C}' \cap \mathcal{V}_a} v_a(C),$$

for every $a \in \mathcal{A}_0$.

Let us give a simple example to illustrate that stability definitions in existing literature [4, 15] provide a strong notion of unstability of matching configurations.

Consider a kind of stable marriage problem on a deserted island with only two men X and Y, and two women x and y. The possible couples that can be formed are $\{X, x\}$, $\{X, y\}$, $\{Y, x\}$, and $\{Y, y\}$. Man X, for instance, has different preferences for $\{X, x\}$ and $\{X, y\}$, but considers Y to be his friend as well. This means that, although X prioritizes its own partner selection, he also cares about whether $\{Y, x\}$ or $\{Y, y\}$ will be formed. Now we assume that: (i) the same holds for Y, (ii) the women's situation is symmetric to the men's situation, and (iii) each man/woman considers the situation of his/her friend to be at most half as important as his/her own situation. Under these assumptions, an interesting example of a matching model is as follows.

Example 2.8 (Stable Marriage) *Consider the matching instance $I = (M, \Gamma, V)$ where (i) $M = (\mathcal{A}, A, B, U)$ with $A = B = (\{X, Y, x, y\} and\, \Gamma = \{\{X\}, \{Y\}, \{x\}, \{y\}, \{X, x\}, \{X, y\}, \{Y, x\}, \{Y, y\}\}$, (ii) U is the identity mapping on A, and (iii) V is the family of M-opinions shown in Table 1. Note that for every $C \in \Gamma$, if x_1 and x_2 are friends and x_1 occurs in C then $v_{x_1} \geq 2 \cdot v_{x_2}$.*

Now, consider the matching configuration $\mathcal{C} = \{\{A, a\}, \{B, b\}\}$. Note that, $v_A(\mathcal{C}) = v_A(\{A, a\}) + v_A(\{B, b\}) = 6 + 4 = 10$, and $v_b(\mathcal{C}) = v_b(\{A, a\}) + v_b(\{B, b\}) = 4 + 6 = 10$. To check whether A and b have an incentive to deviate from $\mathcal{C}$, according to Definition 2.7, we need take into account the *minimum* of A's and b's evaluations of $\Gamma_0 = \{\{\{A,b\},\{a,B\}\},\{\{A,b\},\{a\},\{B\}\}\}$. This yields $\min\{v_A(\{A,b\}) + v(\{a,B\}), v(\{A,b\}) + v(\{a\}) + v(\{B\})\} = \min\{9+0, 9+0+2\} = 9$. But $10 \geq 9$, which means that the formation of $A_0 = \{A,b\}$ is not sufficient to show $\mathcal{C}$'s unstability. However, this conclusion is based on the formation of $\{\{A,b\},\{a,B\}\}$, and A could know that B and a would prefer to stay alone then to join together. This, $\{\{A,b\},\{a\},\{B\}\}$ is actually the only matching configuration that is feasible within Γ_0. Based on this observation, it is not difficult to see that $\mathcal{C}$ should be classified as unstable. Therefore, we may conclude that Definition 2.7, in some circumstances, provides a notion of unstability that is too strong.

2.3 Properties of the Framework

Given a triple $I = (M, \Gamma, V)$ with M a matching environment, Γ the skill set, and V an M-evaluation, the matching problem in general searches for a matching configuration with some desirable properties. It depends on the structure of I which are the relevant properties that the matching problem should take into account:

(i) a matching instance I is called *ordered/unordered*, if all the E-groups in the members of Γ are *ordered/unordered*; otherwise I is called *mixed*.

(ii) a matching instance I is called n-partite, if there exists a partition $\{E_i\}_{1\leq i\leq n}$ of E such that every E-group belonging to one of the members of Γ contains exactly one element from every E_i for $1 \leq i \leq n$;

(iii) given subset $N \in \mathbb{N}$, a matching instance I is called is called N-sized, if for every E-group that occurs in Γ has a size that belongs to N;

(iv) a matching instance I is *globalized* if for every $a_1, a_2 \in A$, $U(a_1) = U(a_2)$, and I *localized* if U is injective;

(v) a matching instance I is *bounded*, if for every E-group C in V_a, there exists a $b \in C \cup A$, such that $U(b) = a$; alternatively if I is not bounded, we call I *social*.

As a convention we assume the n-partite matching instance is either ordered or unordered. Note that a strict n-partite matching instance is also singleton N-sized, where $N = \{n\}$. These above properties allow us to categorize existing approaches to matching, but at the same time allows us to identify new domains to be studied.

3 Matching Framework & Related Work

As mentioned earlier, in existing literature, different terms are used to refer to the matching framework we defined. It is interesting to see that the model we present is able to capture essentials of different approaches to matching. In this section, we give a description of existing approaches to matching and coalition forming and show how it relates to our work. First, we present the relation between stable marriage and our framework. Second, we discuss matching as it is often approached in combinatorial mathematics. Third, we illustrate how cooperative game theory commonly approaches the matching problem.

3.1 Matching : Stable Marriage

Stable marriage can be considered as a special case of a matching instance $I = (M, \Gamma, V)$. A stable marriage setup is typically represented by two finite sets A_1 and A_2 of n elements each, and for each member in the one set, a preference relation on the members in the other set. In terms of our framework this boils down to a matching instance I that is bipartite, localized, bounded and where E is finite.

Traditionally stable marriage uses total orders instead of M-opinions to represent the preferences of individuals [3]. This means that a family $\{\preceq_x\}_{x\in A}$ is given, where $y_1 \preceq_x y_2$ iff the $v_x(\{x, y_1\}) \leq v_x(\{x, y_2\})$. Using the total orders, the stability condition is defined as follows.

Definition 3.1 (Unstable Marriage) *Given a matching instance $I = (M, \Gamma, V)$, a matching configuration C in I is called unstable, if there existst two E-groups $\{X, x\}$ and $\{Y, y\}$, such that $x \preceq_X y$ and $Y \preceq_y X$.*

It can be proven that the avove definition is a special of Definition 2.7.

A variant on stable marriage is the, non-bipartite, stable roommate problem. We can represent this variant by a matching instance that is $\{2\}$-sized but not bipartite. For the stable roommate problem there does not always exist a stable matching configuration.

3.2 Matching : Maximum Cardinality

Combinatorial matching algorithms usually deal with matching problems of forming unordered pairs as groups. Garey and Johnson, for example, define the matching problem as follows [6].

Definition 3.2 (Matching Problem(Garey and Johnson)) *Given a graph $G = (V, E)$, find a set $E' \subseteq E$ of edges such that every vertex $v \in V$ is incident with at most one edge in E'; if every vertex is incident with at least on edge in E' as well, then a matching M of E' is called perfect.*

It can be seen that the matching problem can be modelled using an unordered, $\{2\}$-sized matching instance. Since the matching problem searches a so-called perfect matching, where every individual is matched in a pair, there are no M-opinions at all. That is, every element is passive.

An optimization variant of the matching problem is called maximal matching. In the maximal matching, the matching instance is unordered as well, but $\{1, 2\}$-sized: we have both the singleton sets and unordered pairs as E-groups. The maximum matching searches for a maximum cardinality set of matched pairs. This is modelled by a globalized control mapping and a single M-opinion that assigns 1 to every unordered pair and 0 to every singleton set.

Another variant on this matching problem is known as weighted matching [18]. A graph can be extended by labeling each edge E of the graph G with a weight. Here, the weights can express a distribution over the preference of a vertex for its incident edges. The goal is to (globally) optimize the sum of weights. Similarly to the maximum cardinality matching, the weights can easily be incorporated in the M-opinions.

3.3 Matching : Coalition Formation

In game theory, the approach to the matching problem is different from the combinatorial view. In cooperative game theory, for example, games are considered in their so-called *coalitional form.* The relation of matching to coalition formation is already noticed in [16] where pair partnership matching of elements from an *unordered* group is esentially treated as a variant to the general coalition formation process. In pair coalition formation an agent is satisfied with a coalition of itself and only one other partner.

Let $N = \{1, 2, \ldots, n\}$ be the set of players. Any non-empty subset S of N is a called a coalition. A coalition S can obtain a total utility v. This utility can be distributed among the members(players) of S. Collectively the utilities of all coalitions are represented by a value function.

Definition 3.3 (Value Function) *The value function in a n-person game is given by $v : 2^A \to \mathbb{R}$; For each coalition $S \subseteq N$, $v(S)$ representing the amount of utility that the members of S, collectively, can gain from the coalition.*

The idea is to find a partition $\mathcal{C} = \{A_1, \ldots, A_n\}$ of A (i.e., the coalition structure) such that the sum of all corresponding values are maximized, i.e., $v(A_1) + \cdots + v(A_n)$ is maximized.

A game in coalitional form can be represented by a unordered, globalized, bounded matching instance such that E is finite and Γ contains all partitions of 2^A.

4 Application Scenario: Home Health Care

The generalized matching framework presented in this paper has been applied to further develop a matchmaking component, called the Matcher, within the ASK communication platform. The development and implementation of this novel matching component is carried out as a part of the research within Almende. In particular, we focus on a health care scenario that is studied in the context of the "Luister" project. In "Luister", the aim is to realize a personalized and flexible health care service at Humanitas, a national member-based institute for social services and community development. The matchmaking of health care demand and provision is dependent on the availability, skills and preferences of *patients* and *carers*. Patients can articulate their health care demand through a speech interface, i.e., voice recognition on a telephone connection. Consequently the patient uttering a demand is matched to a suitable and available carer. This results in an appointement with the carer, who will be notified of this commitment automatically.

The participants in a match are enabled to express their preferences on the actual experienced health care provision. After the scheduled service is delivered, feedback from both patient and carer is requested. One of the projects research questions is how to use feedback to improve service quality through personalization.

Initially, we use an ordered bipartite matching instance I to model this domain where the partition $\{A_p, A_c\}$ consists of a *finite* set of patients A_p and a *finite* set of carers A_c. We assume that (i) I is localized, (ii) the patients are active elements, and (iii) the carers can be both active and passive. We could further refine the model of the scenario by partitioning the carers into two groups: professionals $A_{c,p}$ and volunteers $A_{c,v}$; the Matcher is designed to make these kind of refinements on-the-fly while the system is running.

Now let us say that professionals are not allowed to have preferences for patient (professional ethics does not allow them to refuse treatment to an individual in need for care), i.e. every element $a_i \in A_{c,p}$ is *passive*, while preferences of volunteers are taken into account in the matching procedure, i.e. every element $a_j \in A_{c,v}$ is *active*, since a pleasant match may motivate them to provide as much care as possible.

Using the flexibility of our framework, the matching component can dynamically add other matching classes or individuals. Additional context that is relevant for the matching can be included whenever this information becomes available. The additional context can be practically everything: e.g., place, time, language, etc. In this way, the system is able to deal dynamically with multi-dimensional matching in variable dimensions.

5 Conclusion

In this paper, we presented a general matching framework covering some representative matching approaches in literature. The main goal of this model is to combine the properties of different matching approaches. The need for such a general model arises in the light of a real-life application scenarios such as presented in section 4.

We have defined a generalized model and criteria for determining the quality of a matching configuration. We related these criteria to the ones used in combinatorial matching algorithms and cooperative game theory. One of the interesting aspects of our model is that our matching configuration can contain groups of variable size. It allows contextual matching, where certain dimensions of the context can be included or discarded depending on the type of match. In addition, our approach carefully takes into account different, possibly conflicting, opinions on the quality of a match, i.e., the effectiveness of a group.

Several aspects of our model are interesting for further research. In the near future, we want to explore new domains with unbounded control mapping. Research will focus on the effect of opinions of an individual about matches in the matching configuration in which it is not involved.

Next to this, we would like to explore techniques for forecasting of missing evaluations. In society people can base their (local) decision making on external information. Friends, colleagues, like-minded people or others they trust can make recommendations or suggestions for good restaurants, interesting books or a pleasant dentist. An individual consults this group of acquaintances for recommendations on possible matches. In particular, we are interested in systems where the forecasting is localized and distributed over a (possibly large) collection of personal agents.

References

[1] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, Second Edition*. MIT Press and McGraw-Hill, 2001.

[2] K. Decker, K. Sycara, and M. Williamson. Middle-agents for the internet. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, Nagoya, Japan, 1997.

[3] R. Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, New York, third edition edition, 1997.

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

[5] Z. Galil. Efficient algorithms for finding maximum matchings in graphs. *ACM Computing Surveys*, 18:23–38, 1986.

[6] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, New York, NY, USA, 1979.

[7] M. Klusch and A. Gerber. Dynamic coalition formation among rational agents. *IEEE Intelligent Systems*, 17(3):42–47, 2002.

[8] M. Klusch and O. Shehory. Coalition formation among rational information agents. In *Seventh European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, Eindhoven, The Netherlands, 1996.

[9] D.E. Knuth. Stable marriage and its relation to other combinatorial problems. In *CRM Proceedings and Lecture Notes*, volume 10. American Mathematical Society, 1997.

[10] C. Ng and D. Hirschberg. Three-dimensional stable matching problems. *SIAM Journal On Discrete Mathematics*, 4:245–252, 1991.

[11] E. Ogston, B. Overeinder, M. van Steen, and F. Brazier. A method for decentralized clustering in large multi-agent systems. In *AAMAS Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 789–796, New York, NY, USA, 2003. ACM Press.

[12] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.

[13] T. Rahwan and N. R. Jennings. Distributing coalitional value calculations among cooperative agents. In *Proceedings of the Third European Workshop on Multi-Agent Systems*, page 499, Brussels, Belgium, 2005. Koninklijke Vlaamse Academie van Belie voor Wetenschappen en Kunsten.

[14] M. Riedel. Online matching for scheduling problems. *Lecture Notes in Computer Science*, 1563:571–580, 1999.

[15] A. Roth and M. Sotomayor. Two sided matching. *Econometric Society Monographs*, 18, 1989.

[16] D. Sarne and S. Kraus. Time-variant distributed agent matching applications. In *AAMAS*, pages 168–175, 2004.

[17] O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2):165–200, May 1998.

[18] J. Silc and B. Robic. Processor allocation based on weighted bipartite matching scheduling. Technical Report CSD-TR-97-1, 9, 1997.

[19] F. Stolzenburg, J. Murray, and K. Sturm. Multi-agent matching algorithms with and without coach. In *Proceedings First German conference on Multiagent System Technologies(MATES)*, pages 192–204. Springer-Verlag, Berlin Heidelberg, Germany, 2003.

[20] F. Tohme and T. Sandholm. Coalition formation processes with belief revision among bounded-rational self-interested agents. *Journal of Logic and Computation*, 9(6):793–815, 1999.