Concern-based Segmentation for Agents Societies to Facilitate Negotiation Over Service Exchange

Sameh Abdel-Naby, Paolo Giorgini, and Bruno Beaufils

Laboratoire d'Informatique Fondamentale de Lille
Université Lille 1
{sameh.abdel-naby,bruno.beaufils}@lifl.fr

Dipartimento di Ingegneria e Scienza dell'Informazione
Universitá degli Studi di Trento
paolo.giorgini@disi.unitn.it

Abstract. Different communities of software agents may address different predefined concerns, some of these concerns may also be the result of combining two or more sub-concerns. In our research, we look at a concern as the abstract concept a community of agents is continually supporting as long as certain objectives are being accomplished. Besides, a sub-concern is a generically narrower concept of a community's abstract concern, which is satisfied whenever a specific set of objectives is being achieved.
In this paper, we introduce a formal description of the setting wherein we expect software agents to negotiate over service exchange, (i.e., acquisition and provision). We also formalize the negotiation issue that any two interacting agents in the context of our research are expected to resolve while characterizing the negotiation parties - *Agents*. The purpose of the formalization we introduce is to facilitate the achievement of goals that service-driven software agents are interacting to reach.

1 Introduction

In Distributed Artificial Intelligence (DAI), several of the research conducted in Multi-Agent Systems (MASs) addresses the mutually beneficial agreements that a group of interacting autonomous agents are expected to reach. In our research, we look at agents as the transportable software packets that each represents a set of needs a user demands from a service provision entity.

However, when a set of software agents attempt to reach an agreement, a certain level of cooperation must be reached first. In order for this cooperation to be reached a level of organization must be put into consideration. This organization can be made according to needs, or characteristics, or even concerns.

In multi-agent systems, several research efforts are addressing the negotiation of agents in different contexts and, for different purposes. Literature contributions, such as those of [15], [13], and [7], are presenting negotiation models that address specific situations wherein computing machines or robots interacting with each other to achieve certain predefined goals.

Agents' deployments in service provisioning applications are continually growing and, related research are relatively expanding, for an overview: [10], [9], and [6]. Therefore we are motivated to contribute in enhancing the methods used to employ these service-driven software agents. In particular, the literature of Multi-Agent Systems (MAS) is also witnessing the success of delivering services to users of computing devices, such as those presented in Kore [1] and mySAM [2]. These applications apply several of DAI's approaches and take advantage of agent-oriented software engineering methodologies to build goal and service-driven architectures that assist users on-the-go.

This paper is organized as follows: Section 2 introduces the abstract concern-driven setting. Section 3 introduces our notion of the negotiation issue that we expect to negotiating agents to interact upon. Section 4 introduces our approach to formalizing the negotiating agents' requests and tradeoffs. Section 5 discusses some of the related work, and section 6 concludes our paper.

2 The Negotiation Setting

In an automated warehouse, when a group of robots are concerned with placing all of the received objects in dedicated spaces then this is a *concern* for these robots. However, a sub-concern emerges when a subgroup of these robots is concerned with organizing - *only* - the north part of this warehouse. A possible sub-subconcern occurs while two robots of the north-part subgroup are concerned with organizing the red objects only. However, robots operating in a warehouse together with the robots operating in a nearby automobile manufacturer are forming a *society* of robots.

Definition 1 ***A Agents Society****: is a set of agents located in a space wherein different interests' agents are encountering.*

As definition 1 outlines, and figure 1 depicts, when a group of agents come into a common space and, within this group; a number of agents are assigned to completing different abstract concerns, together they form what we call an agents' society. To better elaborate on this, we should think of an agents' society the same as we think of all robots in factories of a specific industrial zone. For example, the industrial zone in Milan has different factories that each has a number of operating robot agents, therefore, all agents in all factories of Milan are forming the Milan's society of industrial robot agents, even though each of these robots is having tasks with different natures to achieve.

We breakdown a society into sub-societies in definition 2. So, within all factories of Milan, robots involved in car manufacturing, and those of washing machines production, together they are forming two different communities of robot agents, but yet they both belong to the Milan's society of industrial agents. However, the classification of societies and their communities are affected by the perspective a problem is tackled from. For instance, from a different perspective, industrial agents of north Milan can also be considered as a society by itself, and every set of similar robots can form a society's possible community. This

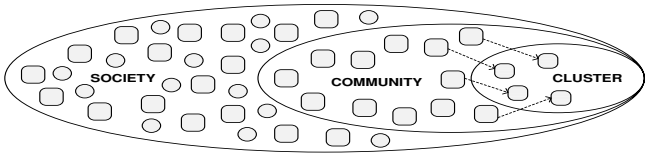

Fig. 1. Agents' Society, Community, and Cluster.

classification can also be made according to robots colors, types, or names, and so on.

Definition 2 ***A Agents Community****: is a subset of an agents' society where a common interest is shared among all of its participants.*

In a community of agents, if a group of agents come into agreement about completing a sub-concern of their community's abstract concern, then we call this group an *agents' cluster*, (Definition 3). For an agents' cluster to be formed, the achievement of a common task must be shared among this cluster's parties. Meaning, this cluster's parties are uniting to achieve a task, (e.g., two agents in a warehouse: organizing boxes). However, for a union to occur, prospective agents must first agree on forming this union, and for agents to agree they must negotiate.

Definition 3 ***An Agents Cluster****: is a subset of an agents' community wherein all parties have came into a mutually beneficial agreement that satisfies their predefined needs.*

The total number of concerns a specific society addresses can be descendingly placed on a pyramid of concerns. In figure 2, this concerns' pyramid has a society's very abstract concern on top of it, then this concern's sub-concerns in less-abstract levels. However, these sub-concerns can also be considered as abstract concerns for sub-communities that play different roles inside the larger community. Depending on the size and number of responsibilities a community has, breaking down the concerns into sub-concerns can be carried on within N levels of descending abstraction, until the least community of a society is defined.

Since the least sub-community in a chain of a larger communities is the one that has an abstract concern and a set of indivisible sub-concerns. Then - *at the bottom of our pyramid* - any set of sub-concerns that is linked to an earlier level concern are, together, reflecting a specific community where agents' clusters may exist. Depending on the number of involved agents, different sizes of clusters may exist within a community. However, in the setting we are concerned about we consider a specific type of clustering in which only two agents are involved, which we call it a *Union*.

The reason we limit the size of a *cluster* to two negotiating parties only is due to the fact that this model is a step on addressing bilateral negotiation of two software agents representing users of pocket computing devices in acquiring

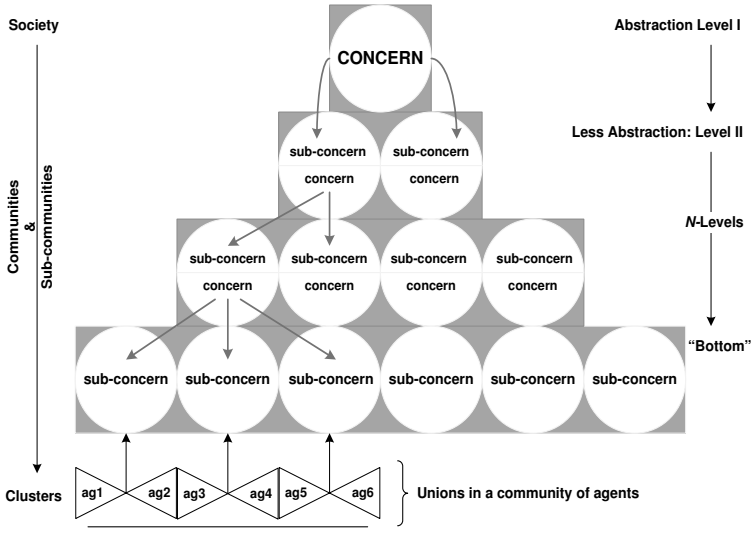

Fig. 2. Community's concerns, sub-concerns and agents' unions

a service. However, this consideration varies according to a *concern*'s needs and *society* structure.

Example 1 (A University, A Nearby Factory, and The Province.).

People working for the same university may each have a different role, but together they form one community that satisfies the university's main concern, which is providing education. On a higher level of abstraction, the university's community and the community of a nearby factory, together, they form a new community of workers, which satisfies a different concern, such as developing the province. However, a professor or more of the same university may have their own community-related concerns, (e.g., become a dean, get funds). Therefore, in a community, each played role can be associated with new concerns that are different from the community's concern, but yet are subs of it, *sub-concerns*.

In order for this professor to satisfy his personal concern - *community's sub-concern* - a set of actions needs to be taken, (e.g., write proposals, dine with key contacts). However, each of these actions would cost its taker something in return. For instance, a professor would tradeoff some of his time with the action of writing a project proposal. Since different actions are likely required to satisfy a single concern, then the value of the tradeoff associated with this concern is expected to increase, (e.g., lots of time), or a number of different tradeoffs would emerge, (e.g., time and money). Therefore, a community's sub-concern is associated with one or more action(s), and a cost that is equal to one or more tradeoff(s). Thus far, we could possibly say that the set of sub-concerns a community searches to satisfy corresponds to a set of tradeoffs members of this community are ready to do.

At a certain time, a professor's insistence to submit a project proposal, and his lack of tradeoffs (e.g., no enough time), may push him to negotiate with his colleagues the idea of establishing a union. Then, the professor might have a chance to tradeoff something that he does not lack at this particular moment, (e.g., fund sharing). In this situation, the negotiation process a professor would carry out is affected by a set of subsequent requests that he previously prepared to eventually ensure the satisfaction of his personal concern.

From this example, we can conclude that For every community of people there is a main concern. For every member of this community there is a number of personal concerns, which are also sub-concerns of a community's main concern. There is a cost for every action taken by a member to satisfy a community's sub-concern. This cost is a set of tradeoffs a member will do in order to perform this concern's related actions. Consequently, there are two possibilities for a community member to satisfy his concern, either 1) all tradeoffs are available and given by this member unaccompanied, or 2) a mutually beneficial union is established between a group of community members in order to make all concern's tradeoffs available.

Making the later possibility occur will require the existence of a negotiation process between a member and a potential union partner. A member's decision whether to *accept* a union or no will depend on the responses he receives to the requests he asks to the same union's potential partners. Therefore, a set of requests corresponds to the tradeoffs a member is willing to do while taking place in a union.

As figure 3 depicts, every community's concern can be described by means of a number of of smaller sub-concerns wherein each of these sub-concerns can be either directly assigned to an entity that is committed to satisfying it or, each sub-concern is in turn divided into a less-abstract set of sub-concerns. Therefore, assuming that there are N concerns, $\mathbf{Concerns} = \{C_1, ..., C_N\}$, that are distributed among an M levels of abstraction, $\mathbf{Levels} = \{L_1, ..., L_M\}$ in which every $l \in Levels$ represents a class of concerns that is more abstract than its subsequent one, then, elements of the **Concerns** are distributed among all **Levels**. Therefore, the number of concerns at L_1 is less than, or equal to, the number of concerns at L_2, and the number of concerns at L_2 is less than, or equal to, the number of concerns at L_M.

Consequently, the way our model is seen is at the top level of abstraction, $L_1 \in$ **Levels**, contains one main concern. The bottom level of abstraction, $L_b \in$ **Levels**, contains the set of concerns, (i.e., sub-concerns), that cannot be broken-down. As a result, as shown in figure 3, every subset of concerns are associated with one concern of its earlier level of concerns, which is one possible bottom edge of a concerns' hierarchy. Hence, this bottom edge represents also a society's specific community, its main concern, and its members' sub-concerns.

Definition 4 ***A Union****: is the agreement of two agents of the same community to fulfill part or all of each others requests.*

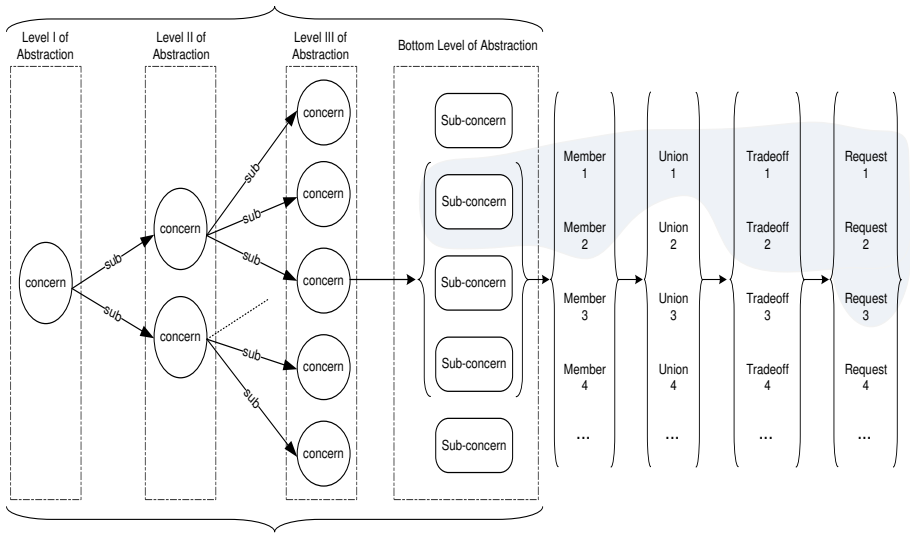

Fig. 3. Community's concern, subconcerns, members, unions, requests & tradeoffs

In every community there is set of members, **Members** $= \{M_1, ..., M_s\}$. Every subset of **Members** represents a community's possible *union* - Definition 4. Since in this article we assume that a community's sub-concern can be achieved by means of members' unioning, then, all of the community's sub-concerns are associated with all possible unions that can are likely to be established in a community, **Unions**.

Therefore, a $u \in$ **Unions** contains a subset of two **Members**. Members of every possible union will be trading off something that they have in order to eventually satisfy the sub-concern they are assigned to. Consequently, a set of tradeoffs is associated with all possible unions within a community.

Definition 5 ***A Tradeoff****: is a part, or the whole, of what an agent is ready to give in exchange of fulfilling one or more of service requirements.*

In order for a member to agree about trading off one or more of his belongings, and then establish a union with other members, he needs to guarantee his share of the union benefits. Therefore, a potential union partner will attempt to maximize his union benefits by negotiating the fulfillment of a set of requests with the other potential partner of the same union. The requests of this member, the requests of other potential union partners, and all requests of a community's possible unions' partner, together, they will form a community's set of requests, **Requests**.

Definition 6 ***A Request****: is one of the characteristics describing the general demand an agent foresees fulfilled in a prospective union.*

We conclude the depictions of figure 3 by locating a possible relationship between a community's sub-concern, members of this community, the union

these members may form, the tradeoffs they will make, and the requests each member will impose in order to ensure that a union's benefits is worth giving a specific tradeoff. This possible set of relationships are the highlighted area in figure 3.

Following the example we mentioned earlier in this chapter, (i.e., A University, A Nearby Factory, and The Province), we could possibly assume here that "*providing education*" is a main concern for a University. Therefore, a set of sub-concerns can also be taking into account, (e.g., hiring enough lecturers). The University, its main concern, and its sub-concerns, are all of a *community* that is part of a bigger picture, which is a *society*.

Formally, this can possibly look like that: *providing education* can be positioned at $C_3 \in$ **Concerns** referring to the fact that this concern is at the third level of abstraction within a *society*, meaning at $(L_3) \in$ **Levels** of abstraction. Therefore, the same university's sub-concerns, (e.g., increase research funds, increase classrooms), are located on the subsequent level of abstraction, $L_3 + 1$.

As described earlier, in each community there is a set of members who share one common concern and, also unite to resolve this concern's sub-concerns. Then, within the community of a university, $M_1 \in$ **Members** can be a Professor that was assigned to satisfying the sub-concern "Increase research funds".

However, to better elaborate on this interconnected relations between Members, Concerns, Requests and Tradeoffs, we will assume that M_1 are lacking the time to achieve the delegated task. Consequently, M_1 attempts to negotiate the idea of establishing a union, $U_1 \in$ **Unions**, with M_2, so that the time they both have will be enough to write a project proposal and increase research funds. In this situation, time is tradeoff, $F_1, F_2 \in$ **Tradeoffs**, both members will have to exchange with the fact of writing a project proposal.

On the other hand, M_1 and M_2 cannot tradeoff their time unless union benefits are guaranteed for each, (e.g., fund commission, promotion, PhD students). Therefore, a community's sub-concern is completed if $M_1, M_2 \in$ **Members** fulfill their requests, $\{R_1, R_2, R_3\} \in$ **Requests**, by means of joining possibly $U_1 \in$ **Unions** while $F_1, F_2 \in$ **Tradeoffs** are the union conditions.

Definition 7 ***An Instance:*** *is the specific concern an agent attempts to complete by means of negotiating the establishment of a union with one of the same community's agents.*

We conclude this section by highlighting the distinction between any of the community's sub-concerns and the sub-concern that a specific agent addresses. We do that by using the word "*instance*" - *Definition 7* - to refer to the single concern a specific agent attempts to complete while playing a certain role in a specific community, and "*instances*" to describe all the same agent's concerns. Consequently, an "*instance*" is a concern for an agent but a sub-concern for its community.

3 The Negotiation Issue

In this section, we define the issue which two agents of this community are going address in their negotiation, and probably agree on its realization. Depending on the sets of requests each negotiating agent is seeking to fulfill, a deal between two agents can be reached under different conditions. Therefore, we then link between an agent's possible situations of acceptance with its requests. We conclude this section by putting together the sets of requests of an agent and the tradeoffs it is ready to offer and, linking them with the different cases wherein a successful negotiation may occur.

Definition 8 ***A Service-Centric Community****: is the set of agents interacting with the intention to fulfill the abstract objective of acquiring a predefine service.*

Here, we consider the negotiation between two agents that are members of the same service-centric community, which we outline in Definition 8. Therefore, all agents are aware of the community's abstract concern / provided service, (e.g., dating, or ridesharing, or bartering). Depending on the kind of service a community is concerned with, a *union* between two agents reflects the completeness of a unique community's sub-concern, (e.g., date(john,sara)). Therefore, a community's sub-concern is created once an agent is searching for a union partner so that together they fulfill each others' requests.

In a service-centric community, we assume the existence of a central agent that we call it a *head-agent*, Definition 9. This *head-agent* is the managing authority of a community, (e.g., a multi-agent platform).

Definition 9 ***A Head-Agent****: is the central and managing member of a community, which is responsible of applying a community's common regulations.*

For example, within a community of dating service, although all agents are seeking to get a date and somehow pay for it, yet; the *head-agent* will be responsible of putting all male agents in one category and doing the same for all female agents. Besides, the *head-agent* will ensure that any male agent that searches for a union partner is actually looking into the category of female agents, and the vice versa. The same applies for a ridesharing community. All ride-giver agents will be separated from ride-seeker agents, even though both categories contain agents of the same type. By "*same type*" we mean; all agents are searching to acquire a service and give something in return.

In a service-centric community, two negotiating agents will discuss the formation of a mutually beneficial *union*. Meaning, a successful negotiation should lead agentX to employ a number of its capabilities (*Tradeoffs*) in order to satisfy a set of needs (*Requests*) agentY has, while agentY is doing the same for agentX.

As figure 4 depicts, for two agents to complete their *instances*, they must unite. For a *union* to occur, agents must tradeoff something that they are capable of providing with the completeness of these *instances*. In order for two agents to agree to tradeoff something they have, they must first be persuaded with the

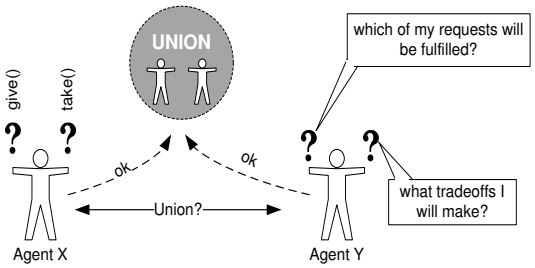

Fig. 4. The Negotiation Issue: A *Union*

benefits of this prospective union while being involved in a negotiation session - *Definition 10*.

Definition 10 ***A Negotiation Session*** *is the time space in which two agents are negotiating the formation of a union.*

An agent gets into a *negotiation session* following the *head-agent*'s task of applying the community's common rules. However, all agents of a specific community are having their own description of the service they are searching to acquire and what they are willing to give in return, (e.g., if I get a blond or curly female from 20 to 25 I would give either a dinner or flower).

For every agent, this general description is broken down into a set of requests and tradeoffs (Definition 6 & Definition 5), (e.g., hair = blond, hair.alternative = curly, age 20, age.alternative = 25, tradeoff1=dinner, tradeoff2=flower). Eventually, an agent negotiate with its potential union partner the possibility to satisfy a set of requests with respect to the associated tradeoffs.

Depending on the nature of requests each agent in a service-centric community is searching to fulfill, the service a community makes available to its members may have different forms. From agents' perspective, the different forms a service takes correspond to different levels of satisfaction an agent may attain while acquiring a service. The highest level of satisfaction an agent could possibly attain is associated with the fulfillment of a specific set of requests. The lowest level of satisfaction an agent may reach to - *not satisfied* - corresponds to the situation where none of the requests' subsets can possibility be fulfilled.

Example 2 (Levels of satisfaction in a ridesharing service).

In a community wherein a ridesharing service is made available throughout its members' interactions, we assume that there is a ride seeker agent that is called **AgentS** and, the highest level of satisfaction **AgentS** may attain is when a negotiation with **AgentG.1** - *that is a ride giver agent* - leads to forming a union in which **AgentS** will be: 1) picked from home, 2) at 14:30, 3) dropped by the post office, 4) with no stops in-between and, 5) the cost is $5.

The lowest level of satisfaction **AgentS** may attain is when none of the negotiations performed within the available time has led to any union formation

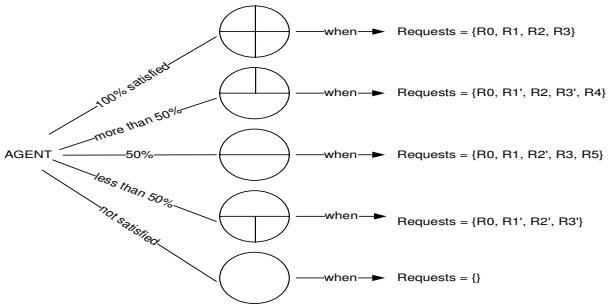

Fig. 5. Levels of satisfaction Vs. Sets of Requests

and, consequently, **AgentS**'s *instance* was not completed. **AgentS** may also be 50% satisfied if a negotiation with **AgentG.2** has led to union in which a car ride with couple of stops are made in-between the departure and the arrival points.

An agent's *instance* can take different forms wherein each reflects a different level of satisfaction an agent may attain. Reaching each of these forms is associated with the fulfillment of a different subset of agent's requests. Therefore, the total number of forms and agent *instance* may take correspond to a large set of different requests. The optimal form of an agent's instance, and its highest level of satisfaction, is obtained when a specific set of requests is fulfilled; a *key-set*.

Definition 11 ***A Key-Set****: is a specific subset of requests that an agent attempt to fulfill in order to obtain the optimal form of its instance.*

In figure 5, we use different types of circles to simplify our notion of an *instance*'s different forms. The highest level of satisfaction an agent may attain is the optimal case of an agent's *instance*, which is represented by means of a crossed-circle. For an agent's *instance* to become a crossed-circle, a set of specific requests must be fulfilled, **Requests** = {R0, R1, R2, R3}. Therefore, the set of requests that leads to this particular shape is the *key-set*.

Following the depictions of the same figure, a number of less optimal forms of the same agent's *instance* can be obtained when different sets of requests are relatively fulfilled. These emerging sets of requests may contain an agent's new types of requests or variants of the *key-set* requests. We use the blank circle to symbolize the case when none of the requests are fulfilled.

However, we conclude the depictions of figure 5 by highlighting the fact that any agent in a service-centric community can be either completely satisfied, not satisfied at all, or having a level of satisfaction that is neither optimal nor insufficient.

One or more of the items available in an agent's list of tradeoffs is associated with the realization of its *key-set*, which is in return associated with an agent's optimal *instance*'s shape. However, since an agent's *instance* may take different

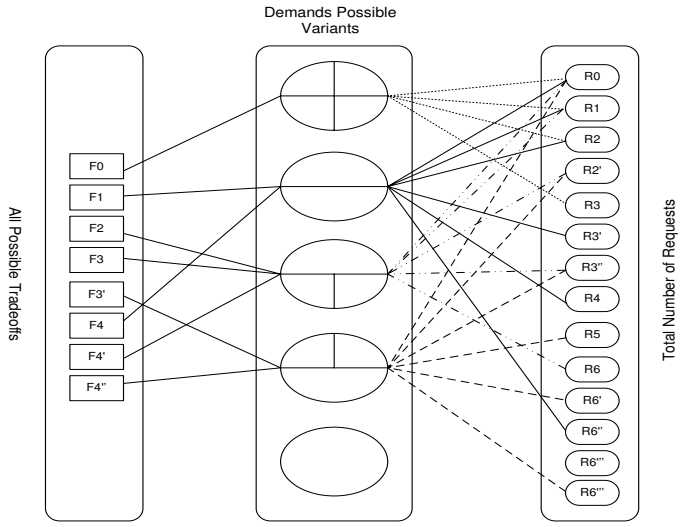

Fig. 6. Satisfaction Levels, Sets of Requests, and Sets of Tradeoffs

forms in which one is optimal, then, different sets of requests may emerge in order to define the other instance's forms. Plus, each of the emerging sets of requests become linked to different tradeoffs or, variations of the *key-set*'s item of an agent's list of tradeoffs or, combinations between all possible tradeoffs.

Figure 6 depicts the relationship between each possible instance form (agent's levels of satisfaction), **Requests** subsets, and **Tradeoffs**. We use the crossed-circle, again, to symbolize the optimal form of an agent's instance, the centrally divided circle refer to 50% fulfillment of the same agent's instance, the circle with its bottom side vertically divided refer to a level of fulfillment that is greater than 50% but less than 100%, and the circle with its upper part vertically divided refer to a level of fulfillment that is greater than 0 but less than 50%.

As figure 6 shows, the *key-set* = {R0, R1, R2, R3} is the one associated with the crossed-circle (optimal instance shape), which is in turn linked to the uppermost item of the tradeoffs - $F0 \in$ **Tradeoffs**.

Following the depictions of figure 6, we then show the case in which a less optimal condition of the same agent's instance is linked to a subset of requests that is not the *key-set*, but yet it includes some of its elements. In addition, we show the case in which a possible instance's form is associated with more than one tradeoff, and the resulted set of tradeoffs includes a variant of an already included tradeoff, (e.g., a book's soft or hard copy).

Among other several scenarios that figure 6 may show, we would like to highlight the existence of an empty circle that we involve to symbolize the case of total incompletion of an agent's *instance*, and consequently it is not connected to any requests' subset or associated with any tradeoffs, but yet it is likely occurring.

Definition 12 *A possible agent's* ***View*** *is the combination of an agent's possible subset of requests plus its associated tradeoffs.*

Since in our research every agent may have a number of requests and a number of tradeoffs and, since all available requests are typically associated with one or more tradeoff. Then, an agent may have an unspecific number of *instances* in which all of them are located in-between the optimal and the minimal forms of satisfaction.

Putting together any satisfaction instance, its specific set of requests, plus their associated set of tradeoffs, we reach a particular combination that we call it an agent's *View*. According to definition 12, an agent may have more than a single *View*, and these *Views* vary according to agent's interests. Consequently, an agent may succeed to find another agent that is capable of fulfilling one of its Views.

A possible *View* of an agent in a community where rideshare service is applicable could be the combination of 1) a set of requests, **Requests** = {start.trento, end.povo, route.nostops, time.1530} and, 2) a set of tradeoffs, **Tradeoffs** = {euro.10} and, 3) an instance of this agents satisfaction, **Satisfaction** = 100%.

4 The Negotiating Agents

Assuming that there are N autonomous agents, **Agents** = $\{A_1, ..., A_N\}$. These agents are bilaterally negotiating to resolve the issue of establishing a mutually beneficial union in order to fulfill each other's service requests, form a union, which we mentioned in section 3.

We reflect an agent's requests, its *instance*'s forms (levels of satisfaction), and list of tradeoffs using a matrix. This matrix's horizontal edge represents an agent's requests, which are in turn describe the characteristics of the service sought, **Requests** = $\{R_1, ..., R_M\}$. The same matrix's vertical edge represents an agent's set of tradeoffs, which is in turn describe the characteristics of the payment, **Tradeoffs** = $\{F_1, ..., F_K\}$. The relations between the elements of these two sets are identified through the existence of an intersection between one another.

Definition 13 *A* ***Set-Cell*** *of an agent's matrix is the intersection between its* $r_0 \in$ ***Requests*** *and its* $f_0 \in$ ***Tradeoffs****.*

In this matrix, a cell wherein the initial point of its vertical column and the initial point of its horizontal row meet exists, and we refer to it as the *set-cell*, definition 13. Since $r_0 \in$ **Requests** contains the uppermost element of an agent's set of requests, and $f_0 \in$ **Tradeoffs** contains the uppermost element an agent's list of tradeoffs, then, a matrix's *set-cell* is also the starting point of the *key-set* a specific matrix represents.

Example 3 (Bob & Alice).

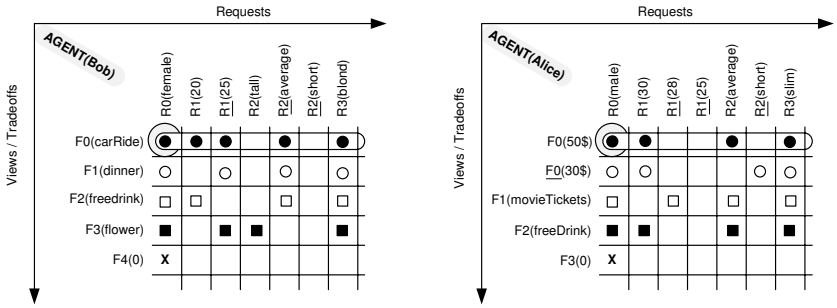

Fig. 7. An example of two agents' service matrixes

In a dating community of agents, agent Bob represents a male that is searching for a female and, agent Alice represents a female that is searching for a male.

In figure 7, agent Bob and agent Alice are both searching for dates. Bob's initial request is indicating the fact that he is searching for a female date, R0(female). Bob would like his date to be either 20 years old - R1(20), or 25 years old - R1(25). Since Bob is an understanding average height person, he prefers his date to be either Tall, Average, or Short - R2(tall), R2(average), or R2(short), but he likes only blonds - R3(blond).

However, Bob is willing to offer different tradeoffs in exchange of the fulfillment of every different combination of his service requests. Meaning, for every satisfied *View* there is a tradeoff item(s) associated with it.

A carRide - `F0(carRide)` - is assigned by Bob to the *key-set*. A dinner - `F1(Dinner)` - is assigned to his date if she is 25 years old, has an average height, and blond. A free drink - `F2(freedrink)` - if she is 20 years old, has an average height, and yet a blond. Eventually, Bob will give nothing - `F4(0)` - if none of his service requests are fulfilled.

On the other hand, similar to the matrix notion of agent Bob; agent Alice has associated the fulfillment of diverse combinations of her service requests - *Views* - with a list of tradeoffs that she is capable of offering. This list contains, 1) different amounts of money, 2) tickets to watch a movie, or 3) a free drink. Therefore, if her prospect date is a male that is between 30 and 25 years old, has either short or average height, and slim, he is most likely to fit.

In this example, if Bob's age, height, and weight are similar to those of Alice's preferences, and Alice's characteristics are matching those Bob is searching to fulfill, then; a negotiation between these two agents may lead to establishing a *Union*. This *Union* is expected to satisfy a possible *View* of the requests and tradeoffs that each of them has previously combined.

5 Related Work

The art of negotiation, as introduced in [11], has been always including attractive research arguments of great scholarly interest. From the literature of

Multi-Agent Systems, many research efforts have been approaching differently the problem of resolving complex situations among interacting agents by means of self-organization as presented in [5], and others by means of argumentation [4], and also by means of cooperation as presented in [3]. However, negotiation, as another alternative for resolving complex situations among agents, is the focus of our research work.

In [14], a model for coalition formation was proposed to enable each agent to select individually its allies. The model supports the formation of any coalition structure, and it does not require any extra communication or central coordination entity. Similarly, the definition of an optimal coalition in [12] is based on Pareto dominance and distance weighting algorithm.

In [8], a model for automated negotiation is proposed for mobile agents to achieve complex tasks in mobile web commerce, which is repeatedly a client to server approach. They introduced a definition of a user and purchase profile that are used mainly in supporting the buying decisions prior to actual product purchase. However, they have assumed that all the heavy computation is performed in the fixed network, by the *CallApplication*, (i.e., one of the assumed linked applications that is responsible of communicating product info with end users), and the agent platform. Similar to our approach, authors of this paper have considered the fact that users are always required to be connected to the managing platform while the negotiation taking place.

6 Conclusions

In this paper, we started by introducing our view of an agents' *Society*, *Community*, and a *Cluster*. Then we introduced our notion of *Concerns* and the fact that a concern can also be perceived as a sub-concern from less abstract level of the society tree, and so on until a sub-concern is linked to a member of a *society* that ends at that level of abstraction: then it is that member's concern and, therefore, his responsibility to complete. Then we gave an example of mapping concerns onto society's levels.

We then introduced our approach of defining a *Request*, a *Tradeoff*, and a *Union* between two agents, which we linked to a *Society*'s concerns and levels of abstraction afterwards. Then, we distinguished between any *Society*'s *Community* and a *Service-Centric Community*. Then, we defined what we call the *Head-Agent*, which we assume to be responsible of putting together members of a service-centric community in order to get their service requests fulfilled.

Within a *service-centric community*, each member has a set of requests to be fulfilled and in return a set of *tradeoffs* must be made. Therefore, we then linked a members' different satisfaction levels, (i.e., *views*), with the types of tradeoffs he/she ready to make and, the different subset of requests that will consequently be fulfilled. A member attempts to fulfill any of his *views* by negotiating the establishment of a *Union* with one of the same community's members.

References

1. Bombara, M., Calì, D., Santoro, C.: KORE: A multi-agent system to assist museum visitors. In: WOA. pp. 175–178. Villasimius, CA, Italy (September 2003)
2. Bucur, O., Boissier, O., Beaune, P.: A context-based architecture for learning how to make contextualized decisions. In: Proceedings of the First International Workshop on Managing Context Information in Mobile and Pervasive Environments. Ayia Napa, Cyprus (May 2005)
3. Doran, J.E., Franklin, S., Jennings, N.R., Norman, T.J.: On cooperation in multi-agent systems. The Knowledge Engineering Review 12(3), 309–314 (1997)
4. Dung, P.M.: On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. Artificial Intelligence. 77(2), 321–357 (1995)
5. Gleizes, M.P., Léger, A., Athanassiou, E., Glize, P.: Abrose: Self-organization and learning in multi-agent based brokerage services. In: IS&N 1999: Intelligence in Services and Networks Paving the Way for an Open Service Market. pp. 41–54. Springer, Barcelona, Spain (April 1999)
6. Jennings, N.R., Crabtree, B.: The practical application of intelligent agents and multi-agent technology. Applied Artificial Intelligence 11(5), 3–4 (1997)
7. Kraus, S.: Strategic negotiation in multiagent environments. MIT Press, MA, USA (September 2001)
8. Matos, F.M., Madeira, E.R.M.: An automated negotiation model for m-commerce using mobile agents. In: ICWE'03: Web Engineering. pp. 313–328. Springer, Heidelberg (January 2003)
9. Mckean, J., Shorter, H., Luck, M., Mcburney, P., Willmott, S.: Technology diffusion: analysing the diffusion of agent technologies. Autonomous Agents and Multi-Agent Systems 17(3), 372–396 (2008)
10. Munroe, S., Miller, T., Belecheanu, R.A., Pěchouček, M., McBurney, P., Luck, M.: Crossing the agent technology chasm: Lessons, experiences and challenges in commercial applications of agents. The Knowledge Engineering Review 21(4), 345–392 (2006)
11. Raiffa, H.: The Art and Science of Negotiation. Belknap Press of Harvard University Press, Cambridge, MA, USA (March 1985)
12. Scully, T., Madden, M.G., Lyons, G.: Coalition calculation in a dynamic agent environment. In: ICML'04: Proceedings of the twenty-first international conference on Machine learning. p. 93. ACM, NY, USA (2004)
13. Smith, R.G.: The contract net protocol: High-level communication and control in a distributed problem solver. IEEE Transactions on Computers C-29(12), 1104–1113 (1981)
14. Wanyama, T., Far, B.H.: Negotiation coalitions in group-choice multi-agent systems. In: AAMAS'06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems. pp. 408–410. ACM, New York, NY, USA (2006)
15. Zlotkin, G., Rosenschein, J.S.: Negotiation and task sharing among autonomous agents in cooperative domains. In: Proceedings of the Eleventh International Joint Conference on Artificial Intelligence. pp. 912–917. ACM, San Mateo, CA (1989)