Cognitive Modeling and Group Adaptation in Intelligent Multi-Agent Meeting Scheduling *§

Leonardo Garrido and Ramón Brena
Centro de Inteligencia Artificial
Tecnologógico de Monterrey
Monterrey, N.L. 64849 México
{lgarrido, rbrena}@campus.mty.itesm.mx

Katia Sycara
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213 USA
katia@cs.cmu.edu

Abstract

In the framework of meeting scheduling problems, we present an approach where for every agent the behavior of other agents is explained in terms of a common cognitive structure. This structure accounts for the combination of emotional and intellectual factors which produce a particular behavior when confronted to a particular situation generated by an ecology of truly decentralized agents, interacting in a concurrent way. The cognitive structure is translated into a computational architecture intended for empirical experimentation. We present a research perspective aimed to investigate group adaptation and evolution as a consequence of the refinement of every agent cognitive model that each agent maintain.

*This research has been sponsored in part by CONACYT and ITESM Campus Monterrey.
§This paper appears in the proceedings of the First Iberoamerican Workshop on Distributed Artificial Intelligence and Multiagent Systems, LANIA & MIA-UV, 1996.

1 Introduction

In Psychology, two broad streams for studying human mind are usually distinguished: on one hand, the *conductist* approach tries to relate a stimulus with a response, without any assumption concerning the internal structure of the subject; on the other hand, the *cognitive* point of view supposes the presence of information structures and processes explaining social behavior [1]. This distinction has its counterpart in Distributed Artificial Intelligence (DAI) and Multi-Agent Systems (MAS) [2]. In Multi-Agent Systems the conductist view can be better identified with the *reactive agents*, whereas the cognitive point of view can be better identified with the so-called *cognitive agents*, which has an explicit information processing structure [3].

The cognitive point of view has much to offer in the field of MAS. Perhaps the most important aspect is that it offers the possibility of *explaining* the behavior of agents in terms of a basic common cognitive structure. For example, if an agent is refusing your proposal in the context of a negotiation, you may want to know the *reason* of the refusal, so you can be more effective in future proposals.

In cognitive modeling of agents there is an implicit assumption of the existence of a common set of components, in terms of which the explanation of behaviors is formulated. Examples of those components are memory, reasoning process, and so forth.

While there are many cognitive architectures for agents, we want to focus on a particular family of problems, namely the *meeting scheduling problem*, and to devise a cognitive architecture suited to this problem. In this way, we could test and measure the accurateness of our agent models.

Meeting scheduling is an everyday task which is a time-consuming, iterative, and somehow tedious. It can take place between two persons or among several persons. Sometimes, these people only try to schedule one meeting. However, most of the time people need to schedule many meetings at the same time taking into account several constraints. In our daily life, meeting scheduling is a naturally distributed task which, many times, is performed by secretaries via telephone or, many other times, it is performed by ourselves via electronic mail.

Each potential attendee needs to take into account his/her own meeting preferences and calendar availability. Most of the time, each attendee has some uncertain and incomplete knowledge about the preferences and calendar of the other attendees; in fact, agents usually try to keep high degrees of calendar and preference privacy. During the process of scheduling the meeting, all attendees should consider the main group goal (i.e. to schedule a meeting) but they also take into account individual goals (i.e. to satisfy their individual preferences).

We believe that in order to consider realistic scenarios in meeting scheduling (that is,

---

[1] The conductist approach can be better identified by the Watson [61] and Skinner [55] work; and the cognitive approach has been rapidly emerging since the 60's, see [27, 17].

[2] We do not introduce the the DAI/MAS field in this paper, directing the reader to the following references: the classic DAI reference is [1]; another more recent DAI survey paper is [9]; an interesting comparison between DAI and MAS can be found in [11]; the proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95) is [35]; an interesting book chapter describing innovative industrial applications of DAI is [39]; and a couple of very recently published books are [46, 38].

[3] In [65], it can be found an interesting discussion about cognitive and reactive agent architectures.

situations close enough to those found when a group of human beings is engaged in a meeting scheduling process), we need to consider three very important aspects: information privacy, cognitive modeling, and group adaptation [4]. We explain these in more detail in the remaining of the paper.

To begin with, we introduce related work in section 2; next, in section 3, we present a brief description of our work concerning information privacy in meeting scheduling; later, in section 4, we describe our current work in cognitive modeling; next, in section 5, we depict our agent-based architecture; then, in section 7, we go into future work concerning group adaptation; finally, in section 8, we present the conclusions.

2 Related Work

Regarding the meeting scheduling, we can find several commercial products but they are just computational calendars with some special features (e.g. availability checkers, meeting reminders); in [59], a review of several of these products can be found. None of these products is a truly autonomous agent capable of communicating and negotiating with other agents in order to schedule meetings in a distributed way taking into account the user's preferences and calendar availability.

However, there has been much research work in meeting scheduling. Most of the earliest work reached interesting but limited success, see [30, 23, 31] for details.

We can found some interesting approaches in Artificial Intelligence such as [37] and [36]; this work focus in learning user preferences but they do not take to much attention to the social and distributed implications of the distributed meeting scheduling process. Eventually, our agents could collaborate with these systems/agents in order to get the user meeting preferences.

We can find other research work in Distributed Artificial Intelligence. Sen & Durfee's work [47, 48, 49, 50, 51] has been focused on solving the meeting scheduling problem using a centralized host agent capable of communicating with all other agents in order to schedule meetings using a negotiation protocol based on contracts [56]; the main purpose of the host agent is to coordinate the search for a feasible schedule taking into consideration attendees' calendars. However, user preferences are not taken into account during the meeting scheduling process. They have focused their research on several search biases to get different density profiles in agents' calendars.

Other recent work in distributed meeting scheduling is due to Ephrati, Zlotkin and Rosenschein [12]. They present an alternative approach which is economic in flavor. They analyze tradeoffs between the mechanism complexity and information preferences and they introduce the Clarke Tax Mechanism as a method for removing manipulability from them.

In [58], Sycara and Liu present another approach based on modeling and communication of constraints and preferences among the agents. Here the agents are capable of negotiating and relaxing their constraints in order to find and reach agreement on schedules with high joint utility. Using this model, agents also can react and revise the schedule in response to

[4] However this does not mean that the inherent scheduling problem (which is known as an NP-complete problem [19]) does not have enough importance to be considered.

dynamic changes.

Regarding agent cognitive modeling we can find a lot of work in Distributed Artificial Intelligence and Multi-Agent Systems fields. As we will see in next paragraphs, most of that work has been done on modeling knowledge, beliefs, intentions, plans, and goals.

Some general introductions to modeling formalisms can be found in [24, 33, 63]. As introduction references to the DAI work on modeling knowledge and beliefs, the reader can see [13] and [25]. Another interesting reference with complexity results for modal logics of knowledge and beliefs is [26].

There are also other more utilitarian and game-theoretic approaches to agent cognitive modeling; as good examples and for further references see Durfee, Gmytrasiewicz & Vidal's work [10, 22, 60] and Werner's work [62].

Cohen and Levesque [5] proposed agent models based on beliefs and goals. However, perhaps the most common approach in cognitive modeling is to use the beliefs-desires-intentions paradigm, see [45]; however, there have been other approaches, such as that taken by Shoham [53], Singh [54], and Krauss [34].

[64, 66] are very good collections of recent papers in Intelligent Agents whose approach is strongly oriented to agent modeling theories, architectures and languages.

Finally, in regard to group behavior research we can also find much work. In fact, many of this research overlap the cognitive modeling field discussed earlier.

Traditionally, group creation, behavior and adaptation have been usually study by Social Psychology and Sociology. However, recently we can find a lot of work related to group behavior in Distributed Artificial Intelligence and, specially, in Multi-Agent Systems.

Some of this research deals with issues such as commitments, intentions and coordination algorithms (e.g. [5, 8]). There is also much work in group negotiation (e.g. [57]). Furthermore, there is too much effort towards standard communication languages and knowledge sharing mechanisms (e.g. [14, 15]).

Although Fox [18] pointed out the great importance of organizational adaptation, there has been little research on group adaptation, learning, and auto-organization in DAI field. Some of the earliest work in this area are [7, 2, 6, 21].

However, interest in these kind of topics has rapidly been reaching the interest of MAS community. For instance, there has been at least one workshop on these topics this year —Adaptation, Co-evolution and Learning in Multiagent Systems in the AAAI Spring Symposium Series; and two more workshops will be held this year —Learning Complex Behaviors in Adaptive Intelligent Systems in the AAAI Fall Simposium Series and the AAAI-96 International Workshop on Intelligent Adaptive Agents.

3 Meeting Scheduling

In [20], we presented some of our preliminary experimental results in Multi-Agent Meeting Scheduling. In our work, we view meeting scheduling as a distributed task where each agent knows its user preferences and calendar availability in order to act on behalf of its user. Our multi-agent system is based on the communication protocol presented earlier in [58]. Using Allegro CLOS, we implemented our system which consists of truly autonomous software agents running, as independent processes, on different computers.

We have run several experiments in order to explore the tradeoffs between meeting quality and efficiency when varying some experimental parameters. Our experiment results showed how the calendar and preference privacy affect the scheduling efficiency and the meeting joint quality under different experimental scenarios. We believe that these variables play a key role in the distributed meeting scheduling task, specially if we are interested in building distributed systems with truly autonomous and independent agents where there is not a specific centralized host agent.

These experiments have shown that when agents schedule meetings trying to keep their privacy, our multi-agent system behaves in a more stable way with better average fastness reaching the best possible meeting joint qualities. This is of particular interest to us because people always try to keep both kinds of calendar and preference privacies most of the time.

So our multi-agent system shows how we can provide automated support for the meeting scheduling task taking into account user preferences and keeping information privacy. The results obtained through experimentation have confirmed some interesting points because they show some aspects of realistic scenarios where agents try to satisfy their own preferences keeping their information privacy.

However, we have been working under some key assumptions: Our agents accept in advance the protocol and coordination policies. Our agents agree in meeting locations, since we have not model this parameter in our system yet. Our agents are honest; that is, they do not try to take advantage of the exchanged information in order to manipulate the outcome of the negotiation. They always try to reach the optimal joint meeting quality and, as human beings, they first relax their preferences before failing to reach an agreement, as it has been presented in [44].

As we will see in next sections, we intend to let agents learn and infer other agents' mental attitutes and behaviors in order to model more complex and realistic scenarios where agents need to adapt to the others in totally distributed (descentralized) environments.

4 Cognitive modeling

When we try to make cognitive models about other people, the first problem is to identify the factors that are involved in the group problem-solving process. For instance, we frequently take into consideration factors such as the particular event factors, our believes about other agents' attitudes, feelings, and capabilities in addition to our own personality [5].

Figure 1 shows how we can model an agent when it is engaged in a group problem-solving process. As we can see, attitudes, moods [6], event factors and capabilities affect our current dispositions and feelings. Furthermore, our desires and preferences can be affected, on one hand, by the event factors and our capabilities, and, on the other one, by our current dispositions and feelings. This model can be used for defining the operation of an agent, and it also can be used by an agent in order to have a cognitive model about other agents.

[5] We are using here the term *personality* as the set of distinctive and basic human features, such as: attitudes, capabilities, and moods.

[6] We are using here the term *moods* with a very similar connotation to that used for *humors* in the middle age; that is, it is the set of distinctive emotions and feelings of a person or agent.

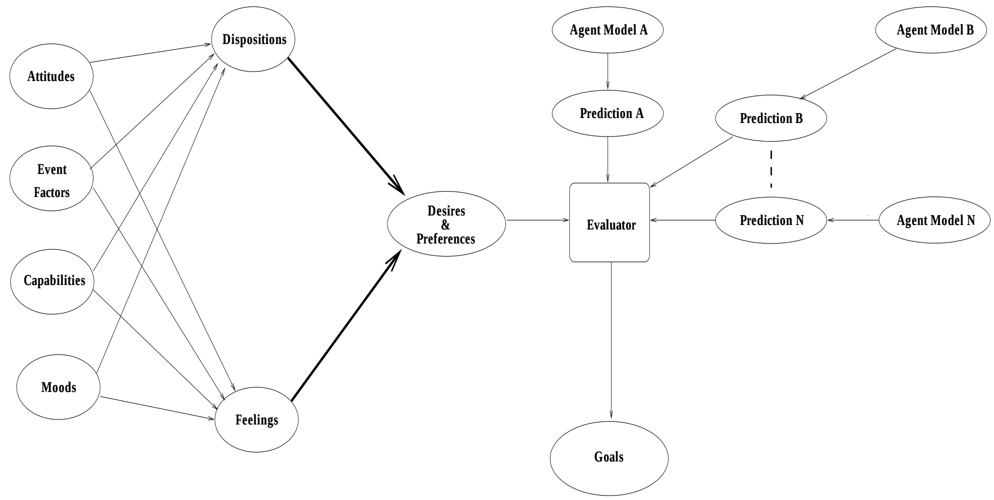

Figure 1: Actors and roles in agent modeling

Knowledge about other agents' personalities (i.e. knowledge about their moods, attitudes, and capabilities) can be very useful for predicting their desires, preferences, intentions, and goals under some particular event. In addition to these predictions, our own desires and preferences can be used as inputs to an evaluation process which gets a set of goals which can be used by other processes such as planning, design, scheduling, group coordination, negotiation, adaptation, and so forth.

Obviously, the quality and effectivity of agent's actions depend on the correctness of the agent's predictions which are based on the agent cognitive models maintained by the agent.

In the following subsections we discuss some particular cognitive factors that we have identified in the meeting scheduling domain. First, we introduce these factors in a completely intuitive way; after this, we will discuss how to deal with their precise meaning.

4.1 Attitudes

By *attitudes* we mean a fixed (that is, relatively invariant with time) bias toward certain kinds of actions. For example, an *individualist* agent tends to give priority to its own interest, in contrast with a *collaborator* agent, which favor the group's interest.

We can detect the following distinctive agent attitudes in a meeting scheduling scenario:

Social attitudes. We categorized here those attitudes related to social or group activities. Some attitudes that we can find in this category are: collaborator, individualist, altruist, selfish, benevolent, malevolent, etc.

Calendar attitudes. In the meeting scheduling domain, we can observe, on one hand, people biased to form calendars with meeting clusters and, on the other one, people biased to form calendars with uniform distributions of meetings. Some attitudes that we can identified under this category are: being biased to meet in the morning, afternoon or nigth; being biased to meet as soon as possible; being biased to meet always in a specific day of the week; being biased to form calendar with uniform meeting distributions; and so forth.

Communication attitudes. We have also observed that there are some attitudes related to the communication process. Some attitudes that can be categorized here are: talk active, passive agents, conformist, nonconformist, etc.

4.2 Event factors

In order to try to differentiate each event that can arise in a meeting scheduling scenario, we can take into account the following factors:

Meeting features. We have identified some relevant meeting characteristics, such as: the meeting goal, compulsory meeting degree (compulsory/optional), meeting attendees, attendees hierarchy, meeting deadline.

Agents' actions/proposals. In the general case, it is necessary to take into account other agents' actions; however, for the time being, we have decided only to consider the agents' proposals and counter-proposals as their observable actions in our meeting scheduling domain.

Temporal Factors. We have identified that there are also other time-related factors, such as: being close to the meeting deadline and being far from the meeting deadline.

4.3 Capabilities

For the time being, we have decided to take into account only the agent's calendar availability as the agent's capability to meet. However, we could consider much more complex capability factors in this and other domains, such as software/hardware failures.

4.4 Moods

As we previously mentioned, we use the term *moods* with a very similar connotation to the medieval *humors*; this is in the sense of being a set of distinctive emotions and feelings of a person or agent. So that here we could take into consideration only those features that traditionally have been known as people *humors*. However, we consider here a lot of feelings and emotions that can be used to distinguish a person or agent, such as:

Affects. These are some special kind of social emotions such as: friendlessness, love, hate, gratitude, admiration, resentment, etc.

Emotions. We distinguish here other more non-social or impersonal feelings such as: sadness, happiness, anger, sorrow, shame, distress, etc.

4.5 Dispositions and feelings

In our model, we use the term *dispositions* as a kind of current temporal attitudes that emerge in each particular situation and are determined by our distinctive attitudes, moods, capabilities, and the particular event factors.

Likewise, we use the term *feelings* as a kind of current temporal moods that also arise in each particular event. These feelings are determined by our distinctive moods, attitudes, capabilities, and the particular factors of the situation.

4.6 Semantics

The precise meaning of the cognitive components is currently being defined following a kind of "operational semantics", that is, a definition of their meaning in terms of the behavior of a given interpreter. The interpreter is to be the algorithm for combining the cognitive components so that a specific preference –and ultimately a behavior– is generated. Of course, this definition does not pretend to account for a general definition of terms like "happiness" but just to be a working definition for the specific context of meeting-scheduling problems.

4.7 The modeling process

To model agents' attitudes, capabilities, dispositions, etc. means to infer them from their observed agents' behavior. This behavior includes mainly the messages sent (or the absence of such messages) by the agent being modelled to other agents. Clearly the inference of other's attitudes from their behavior is not a deterministic activity since a single action could be a consequence of many different cognitive profiles. In this way, modeling can be viewed as an iterative and gradual process, where every new piece of information about a particular agent –that is, its recent behavior– is analyzed in such a way that the model of that agent is further refined. The more the modeling process is refined, the more accurate the models could be.

In Psychology, *attributions* is the term for describing the process of giving an explanation to an observed action. Attributions could be computationally calculated by a backward-chaining process relying on a base of known traits and observed behaviors. This base would account for the cognitive model represented in figure 1.

This is a form of hypothetical reasoning, sometimes referred to as *abduction* [3, 28]. In [41, 42], it is proposed a computational system capable of obtaining the *abducentes* or explanations for a given observation. Other possibility is the use of Bayesian networks [40, 43] to calculate the probability of a certain cognitive profile to explain an observed behavior. Currently we are considering these methods for calculating the cognitive models explaining the agents' behavior.

5 Agent architecture

We have also been working in the design of an agent architecture for supporting further development on our multi-agent meeting scheduling system.

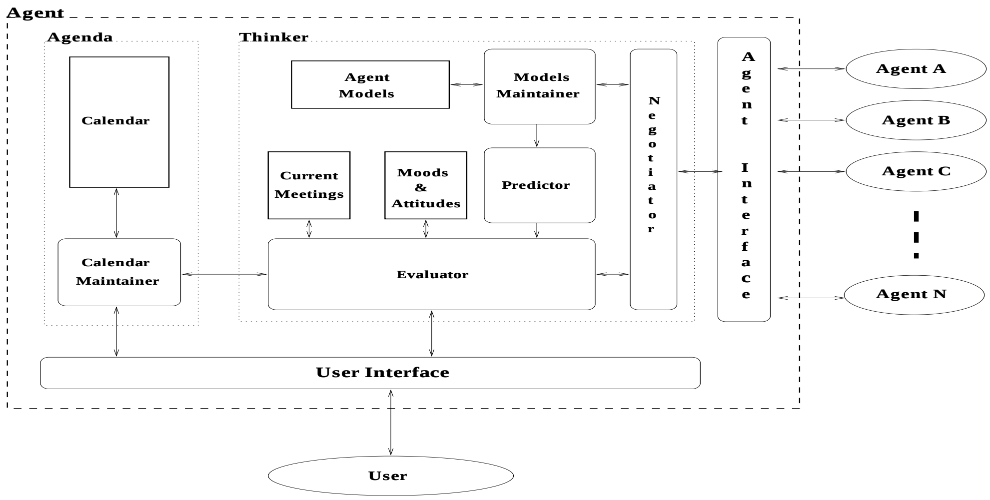

Figure 2: Agent-based architecture

Our agent architecture is an agent-based architecture; that is, we implement each module as an independent process capable of sending and receiving messages in an asynchronic way. Each process/agent is responsible for performing some particular activities. The group activity gives life to the entire meeting scheduling agent. The architecture is shown in figure 2. This figure shows the module/agents (displayed as rounded squares), databases (displayed as squares), and the relationships between them.

A brief description of each component of the agent-based architecture is presented as follows:

User interface. This module is responsible of the interaction with the user. It sends and receives user messages and forwards them to the appropriate module in the adequate format. This module would communicate with the user using any appropriate language such as a graphical and/or natural language.

Agent interface. This module is responsible of the interaction with other agents. It sends and receives other agent messages and forwards them to the appropriate module using the adequate format. In order to ensure agent generality this module can use a standard communication protocol such as KQML [16].

Agenda. This module is composed by two submodules:

Calendar. This is the database where all the agent meeting information is stored.

Calendar manipulator. This module is responsible of managing the calendar database. This module can receive updating and/or retrieval messages from the other modules and it can send the respective answer messages.

Thinker This module is responsible of doing all those cognitive process that gives intelligence to the agent. This module is more complex than the previous one and it is constituted by having several submodules, such as:

Negotiator. Once that this module receive a particular goal from the Evaluator module, the Negotiator module is responsible for performing the appropriate strategies of negotiation, argumentation and persuasion.

Model Maintainer. This module is responsible of inferring, learning, creating and adapting agent models using the agent cognitive model discussed in a previous section.

Predictor. This module is responsible of making predictions about the other agents' desires, preferences and goals based on the models created by the Model Maintainer module.

Evaluator. This module is responsible of evaluating the predictions about other agents, the own moods and attitudes, and the particular meeting features in order to create current goals that are sent to the Negotiator module which takes a specific action (proposal/counter-proposal) or strategy.

Current Meetings. This is a database with all the meeting features information is stored. Remember that each meeting scheduling agent is capable of processing several meeting processes at the same time.

Moods & Attitudes. This is another database where the basic cognitive personality of the meeting scheduling agent is stored.

Agent Models. This database stores all the cognitive agent models that are managed by the Models Maintainer module.

This is an initial architecture which can be easily updated. In fact, this is why we have implemented under an agent-based fashion. We believe that this characteristic gives us modularity and encapsulation enough to be able to develop our architecture towards any appropriate future direction.

6 The Schedule Game

We felt the need of an interesting environment in which to carry out some experiments. Further, we needed a framework rich enough to:

- Allow for cooperative as well as individualistic actions;
- Show or hide agent's private information and goals;
- Define different action styles or "roles" played by the participating agents;
- Make it useful and interesting to discover and model the other agents' roles and characteristics from their behavior.

In order to fulfill the requirements outlined above, we defined a competitive game which can be played by persons as well as by agents.

The goal of the players is to accumulate points, winning the player who accumulates more points in a fixed number of "rounds".

In each round the players "bid" for each one for a specific hour, winning the group of players which totalizes the greatest "count". A count is obtained by simply adding the values that each hour has for each player. Here comes the interesting point: the different hours could have different values for the different players, according to a "role" randomly assigned to each player at the beginning of the game. There are five "roles", with the following points assigned to the hours:

The early-rising , who prefers to schedule meetings as early as possible; the points for each of the hours (from 9 a.m. to 4 p.m.) are 8, 7, 6, 5, 4, 3, 2, 1.

The night owl , who prefers to schedule meetings as late as possible; the points assigned to the hours are 1, 2, 3, 4, 5, 6, 7, 8.

The extreme , who either tries to schedule meetings early in the morning or as late as possible; the values of the hours are 8, 6, 4, 2, 1, 3, 5, 7.

The medium , who prefers to schedule meetings around the middle of the working day; the values of the hours are 1, 3, 5, 7, 8, 6, 4, 2.

The flat , who does not care about the hour of the meeting. The values of the hours are all 4.

The counting procedure is performed by an external "referee", which announces the winner group, but keeps for him the exact amount of points earned by each team and player. Only after, say, 10 rounds, when the game is over, the exact count is announced to every player.

The mechanism outlined above mimics the situation where a group of working people tries to reach an agreement scheduling a meeting, of course, with some major simplifications.

The analogy of the schedule game with the "real" schedule is stressed by the following considerations:

- The players are forced to reach an agreement with other players, because otherwise they could not win.
- For the sake of their own interest, the players should try to maximize his/her group's count.
- The individual utility is also considered, as the players in the winning group receive his/her own contribution, which should then be as high as possible.

Thus, for example, the early-rising tries to schedule the meeting as early as possible, in order both to contribute to the group's count and to receive his/her individual points.

One of the design goals for the schedule game was that it should be useful for the players to guess the others' roles. In the schedule game, each player should try to present a bid which is *interesting for other players*. This is most efficiently done if the player knows –or guesses correctly– the other players' roles. For example if I am a "medium" and I guess that some of the players are early-rising, then I could bid for an hour in the first part of the working day.

The strategies in the schedule game could be fairly complex. Each player bids for the hour s/he believes will get the highest sum, adding all players' bids. But you will not help to win a player you think has already a clear advantage over you, so probably you will prefer to arrive to an agreement with other players.

7 Towards group adaptation

One of the most intriguing aspects of multi-agents systems is how a group behavior emerges from the interaction of individual agents. There are many interesting group behaviors, such as: different mechanisms of group creations, disintegration and modification; groups interacting with other groups; people joining to and leaving groups in a concurrent way; people becoming involved in several groups; formation of coalition and blocks among agents for achieving a common goal, and so forth. And all this is possible with people (agents) usually having individual motivations, beliefs, and goals that can be similar or different along several groups; inside a group, people regularly have different attitudes (e.g. altruist, malevolent, cooperative); sometimes it is possible to see groups with some global goals but with different members' individual goals.

It could seem that if there are many individual and group factors, all of them interacting in a concurrent and *decentralized* way, it would be impossible to see group adaptation and evolution showing better global performance. However, it is not the case. In general, groups evolve and adapt in spite of several apparent negative factors as it can be observed in several human and animal group scenarios.

We can face much of these situations in the decentralized meeting scheduling domain. When we have a truly decentralized multi-agent meeting scheduling scenario (i.e. a situation where there is not a specific host agent nor a fixed centralized protocol), we can face problems such as group adaptation and learning; that is, the group must be able to learn and adapt to itself as a whole in order to develop better performance.

We think (and hope) that our work and results in this domain can be migrated to other domains that also face truly decentralized group problem-solving processes (e.g. distributed decision-making)

As we have already said, we have realized that if we want to tackle more realistic scenarios, agent cognitive models need to take into account other agents' feelings, attitudes, capabilities, traits and moods. In particular, we think that this kind of modeling is specially important in group adaptation.

Although there could be other more reactive or conductist approaches, our hypothesis is that agent cognitive modeling is specially important in order to achieve a successful group adaptation; furthermore, we maintain that besides intentions and goals cognitive models need to take into consideration such things as personality and moods.

Using the cognitive models and agent architecture discussed earlier, we intend to develop several experimental group scenarios in order to test our hypotheses. Our basic research questions are:

- What cognitive factors are critical in group behavior such as group creation, adaptation and dissolution?
- How does agent cognitive modeling help to improve group performance?
- How do personality, moods, attitudes and feelings affect group adaptation?

In the long term, we have some special interest in a particular social phenomenon: the apparent non-explainable global group behavior arising from concurrent and decentralized individual behaviors. This kind of systemic and holistic approach has been treated by others in different domains and contexts, for instance: mind conceptualization in Philosophy [29], the Gestalt approach in Psychology [32], enterprise organization in Business Management [52, 4].

We also think that agent cognitive modeling is specially important in explaining this global and systemic group behavior. Our basic research questions here are:

- How global behavior emerges from individual (cognitive) factors?
- How can we develop social cognitive models from individual agent cognitive ones?

Our final intention is not only to explain group behavior in terms of agent cognitive models but also to develop suitable mechanisms for creating societies of artificial and intelligent agents capable of behave, learn, and adapt in a concurrent, decentralized, and independent way.

8 Conclusions

Modeling is an ubiquitous activity in real-life when you are trying to arrive to an agreement. You have to suppose counterparts' intentions and goals based on what you think their personality is; and the more accurate your predictions are, the more adequate your behavior is.

In this paper we presented a cognitive style modeling aimed to characterize other agents in a meeting scheduling activity. As we have done in our previous work, we will substantiate the adequateness of our cognitive models by doing simulation-based empirical experimentation.

One aspect we will delve into is group adaptation and evolution in meeting scheduling domains as theirs individuals "learn" –that is, model– other agents' cognitive traits.

References

[1] A. H. Bond and Gasser L. *Readings in Distributed Artificial Intelligence.* Morgan Kaufmann, 1988.

[2] S. Cammarata, D. McArthur, and R. Steeb. Strategies of cooperation in distributed problem solving. In *The Eighth International Joint Conference on Artificial Intelligence*, 1983.

[3] E. Charniak and D. McDermott. *Introduction to Artificial Intelligence.* Addison-Wesley, 1985.

[4] Chawla and Renesch, editors. *Learning Organizations: Developing Cultures for Tomorrow's Workplace.* Productivity Press, 1995, 1995.

[5] P. R. snf Levesque H. J. Cohen. Intention is choice with commitment. *Artificial Intelligence*, 42(3):213–261, 1990.

[6] D. D. Corkill and V. R. Lesser. The use of meta-level control for coordination in a distributed problem solving network. In *The Eighth International Joint Conference on Artificial Intelligence*, 1983.

[7] R. Davis and R. G. Smith. Negotiation as a metahphor for distributed problem solving. *Artificial Intelligence*, pages 63–109, 1983.

[8] K. S. Decker and V. R. Lesser. Designing a family of coordination algorithms. In *The First International Conference on Multi-Agents Systems.* AAAI Press, 1995.

[9] E. H. Durfee. The distributed artificial intelligence melting pot. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(6):1301–1306, November/December 1991.

[10] E. H. Durfee and P. J. Gmytrasiewicz. Overeager reciprocal rationality and mixed strategy equilibria. In *The Eleventh Natinal Conference on Artificial Intelligence*, 1993.

[11] E. H. Durfee and J. S. Rosenschein. Distributed problem solving and multi-agent systems: Comparisons and examples. In *The 13th International Workshop on Distributed Artificial Intelligence*, 1994.

[12] E. Ephrati, G. Zlotkin, and J. S. Rosenschein. A non-manipulable meeting scheduling system. In *The 13th International Workshop on Distributed Artificial Intelligence*, 1994.

[13] R. Fagin, J. Y. Halpern, and M. Y. Vardi. What can machines know? on the properties of knowledge in distributes systems. *Journal of the ACM*, 39(2):328–376, 1992.

[14] R. Fikes, M. Cutkosky, T. Gruber, and J. Van Baalen. Knowledge sharing technology project overview. Technical Report KSL 91-71, Stanford University, Knowledge Systems Laboratory, 1991.

[15] T. Finin and R. Fritzon. KQML as an agent communication language. In *The Third International Conference on Information and Knowledge Management (CIKM'94).* ACM Press, 1994.

[16] T. Finin, R. Fritzon, and D. McKay. A knowledge query and manipulation language for intelligent agent interoperability. In *The Fourth National Symposium on Concurrent Engineering, CE & CALS Conference*, 1992.

[17] S. T. Fiske and S. E. Taylor. *Social Cognition.* McGraw Hill, 1991.

[18] M. S. Fox. An organizational view of distributed systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 11:70–70, January 1981.

[19] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* Freeman and Co., 1979.

[20] L. Garrido and K. Sycara. Multi-agent meeting scheduling: Preliminary experimental results. In *The Second International Conference on Multi-Agent Systems.* AAAI Press, 1996. To appear.

[21] L. Gasser, N. Rouquette, R. W. Hill, and J. Lieb. Representing and using organizational knowledge in dai systems. In L. Gasser and Huhns M. N., editors, *Distributed Artificial Intelligence*, volume 2 of *Research Notes in Artificial Intelligence*, pages 55–78. Pitman, 1989.

[22] P. Gmytrasiewicz and E. H. Durfee. Elements of utilitarian theory of knowledge and action. In *The Thirteenth International Joint Conference on Artificial Intelligence*, 1993.

[23] I. Greif. PCAL: A personal calendar. Technical Report TR-213, MIT Laboratory for Computer Science, Cambridge, Mass, 1982.

[24] J. Y. Halpern. Reasoning about knowledge. In *The 1986 Conference on Theoretical Aspects of Reasoning About Knowledge*, 1986.

[25] J. Y. Halpern. Using reasoning about knowledge to analyze distributed systems. *Annual Review of Computer Science*, 2, 1987.

[26] J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.

[27] K. J. Holyoak and P. C. Gordon. Information processing and social cognition. In R. S. Wyer and T. K. Srull, editors, *Handbook of Social Cognition.* Erlbaum, Hillsdale, NJ, 1984.

[28] J. R. Josephson and Josphson S. G. *Abductive Inference: Computation, Philosofy, Technology.* Cambridge University Press, 1994.

[29] I. Kant. *Critiques of Pure Reason.* St. Martin's Press, 1969. Originally published, 1781.

[30] C. Kelley, J. F. and A. Chapanis. How professional persons keep their calendars: Implications for computarization. *Journal of Occupanional Psychology*, 55:141–156, January 1982.

[31] C. Kincaid, P. Dupont, and A. Kaye. Electronic calendars in the office: An assessment of user needs and current technology. *ACM Transactions on Office Information Systems*, 3(1), January 1985.

[32] K. Koffka. *Principles of Gestalt Psychology.* Harcourt, Brace & World, 1935.

[33] K. Konolige. *A Deduction Model of Belief.* Pitman/Morgan Kaufmann, 1986.

[34] S. Krauss and D. Lehmann. Knowledge, belief and time. *Theoretical Computer Science*, 58:155–174, 1988.

[35] V. Lesser and Gasser L., editors. *The First International Conference on Multiagent Systems.* MIT Press, 1995.

[36] P. Maes. Agents that reduce work and information overload. *Communications of the ACM*, 37(7):30–40, July 1994.

[37] T. Mitchell, R. Caruana, D. Freitag, J. McDermott, and D. Zabowski. Experience with a learning personal assistant. *Communications of the ACM*, 37(7):80–91, July 1994.

[38] G. O'Hare and N. Jennings, editors. *Foundations of Distributed Artificial Intelligence.* John Wiley and Sons, 1996.

[39] H. V. D. Parunak. Applications of distributed artificial intelligence in industry. In G. O'Hare and N. Jennings, editors, *Foundations of Distributed Artificial Intelligence*, volume 2. John Wiley and Sons, 1996.

[40] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, CA, 1988.

[41] D. Poole. Explanation and prediction: an architecture for default and abductive reasoning. *Computational Intelligence*, 5(2):97–110, 1989.

[42] D. Poole. A methodology for using a default and abductive reasoning system. *International Journal of Intelligent Systems*, 5(5):521–548, December 1990.

[43] D. Poole. Probabilistic Horn abduction and Bayesian networks. *Artificial Intelligence*, 64(1):81–129, 1993.

[44] H. Raiffa. *The Art and Science of Negotiation.* Harvard University Press, Cambridge, Mass., 1982.

[45] A. S. Rao and M. P. Georgeff. Modeling rational agents within a bdi-architecture. In *The Twelfth International Joint COnference on Artificial Intelligence (IJCAI-91)*, 1991.

[46] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter.* Artificial Intelligence Series. MIT Press, 1994.

[47] S. Sen and E. H. Durfee. A formal study of distributed meeting scheduling: Preliminary results. In *The ACM Conference on Organizational Computing Systems*, 1991.

[48] S. Sen and E. H. Durfee. A formal analysis of communication and commitment. In *The 11th International Workshop on Distributed Artificial Intelligence*, 1992.

[49] S. Sen and E. H. Durfee. The effects of search bias on flexibility in distributed scheduling. In *The 12th International Workshop on Distributed Artificial Intelligence*, 1993.

[50] S. Sen and E. H. Durfee. Adaptive surrogate agents. In *The 13th International Workshop on Distributed Artificial Intelligence*, 1994.

[51] S. Sen and E. H. Durfee. Unsupervised surrogate agents and search bias change in flexible distributed scheduling. In *The First International Conference on Multi-Agents Systems*, 1995.

[52] P. Senge. *The Fifth Discipline: The Art and Practice of the Learning Organization.* Doubleday Currency, 1990.

[53] Y. Shoham. *Reasoning About Change.* MIT Press, 1988.

[54] M. P. Singh. A critical examination of the cohen-levesque theory of intention. In *The Tenth European Conference on Artificial Intelligence (ECAI-92).* John Wiley and Sons, 1992.

[55] B. F. Skinner. Operant behavior. *American Psychologist*, 1963.

[56] R. Smith. The contract net protocol: High-level communication and control in distributed problem solver. *IEEE Transactions on Computers*, 29(12):1104–1113, December 1980.

[57] K. Sycara. Multi-agent compromise via negotiation. In L. Gasser and Huhns M. N., editors, *Distributed Artificial Intelligence*, volume 2, pages 55–78. Pittman, 1989.

[58] K. Sycara and J. Liu. Distributed meeting scheduling. In *The Sixteenth Annual Conference of the Cognitive Society*, 1994.

[59] E. Taub. Sharing schedules. *Mac User*, pages 155–162, July 1993.

[60] J. M. Vidal and E. H. Durfee. Agent modeling methods using limited rationality. In *The Twelfth Natinal Conference on Artificial Intelligence*, 1994.

[61] J. Watson. *Behaviorism.* Norton, New York, 1930.

[62] E. Werner. A unified view of information, intention and ability. In Y. Demazeau and J. P. Müller, editors, *Descentralized AI 2 - Proceedings of the Second European Workshop on Modeling Autonomous Agents and Multi-Agent Worlds (MAAMAW-90).* Elsevier/North Holland, 1991.

[63] M. Wooldridge. The logical modelling of computational multi-agent systems. Technical Report NMU-DOC-94-01, Manchester Metropolitan University, Department of Computing, 1994.

[64] M. J. Wooldridge and N. R. Jennings, editors. *Intelligent Agents.* Lecture Notes in Artificial Intelligence. Springer-Verlag, 1995.

[65] M. J. Wooldridge and N. R. Jennings. Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10(2), 1995. To appear.

[66] M. J. Wooldridge, J. Müller, and M. Tambe, editors. *Intelligent Agents II - Agent Theories, Architectures and Languages*. Lecture Notes in Artificial Intelligence. Springer-Verlag, 1996.