

# Competing Agents in Agent-Mediated Institutions

Enric Plaza, Josep Lluís Arcos, Pablo Noriega, Carles Sierra  
IIIA, Artificial Intelligence Research Institute  
CSIC, Spanish Council for Scientific Research  
Campus UAB, 08193 Bellaterra, Catalonia (Spain).  
{enric | arcos | pablo | sierra}@iiia.csic.es  
<http://www.iiia.csic.es>

## Abstract

Social processes and agent interaction always take place in a specific context. A school of thought in social studies analyzes them in the framework of institutions [23]. We present in this paper the notion of *agent-mediated institutions* and show how it is relevant for multiagent systems (MAS) in general and, more specifically, for MAS that include human agents and software agents involved in socio/economic interactions. We will show how the social interactions of human and software agents taking place in the COMRIS<sup>1</sup> Project can be described as such an institution, the Conference Centre institution.

## 1 Introduction

Dialogues are pervasive and, to a large degree, unavoidable. Legal arguments, political debate, domestic disputes, didactic explanations, interviewing, psychotherapy, coordination of actions, negotiation, all tend to involve some form of dialogical interaction. But dialogues are also unavoidable since what is accomplished through them cannot be accomplished in a strictly “monological” setting, because some fundamental ontological, rhetorical or epistemic features would be lost. Thus, dialogues may be worth studying, although they are not simple entities. Certainly not from a formal perspective:

On one hand, dialogues involve multiple participants, who exchange illocutions in rich and complex languages. Thus, classical —*i.e.* monological, truth-semantic, non-dynamic— formal devices, are inadequate to deal formally with these complexities. But in addition, dialogues are typically *situated* (or “opaque” or “unstructured”), in the sense that participants need to react to the illocutions, depending on the conditions or elements present in a given “context” or “situation”. In many dialogical situations meaning is not necessarily established in an objective, *a priori*, form; nor are interventions subject to an objective, *a priori*, clearly expressible protocol. In typical dialogues, participants confirm, adjust, refine or establish their own meanings, intentions, beliefs and

---

<sup>1</sup> COMRIS stands for Co-Habited Mixed-Reality Information Spaces. More information is available at URL <http://arti.vub.ac.be/~comris/>.

actions according to their individual interpretation of what is happening and what the other participants are saying. The first kind of complexity has been addressed through ad-hoc “dialogical” structures (cf. Hamblin [14], Rescher [26], or Hintikka [15] for different approaches), the second one has been the object of increasing attention, mostly from the idea of “a situation” (like Barwise and Perry’s [4]), but also from the notions of “*speech acts*” and “*conversation*” (c.f.[30, 8])

Certainly there are dialogical situations which are irremissibly opaque. That is the case, for instance, of psychotherapy, domestic disputes or everyday conversation, where meaning and commitments are mostly established through highly unstructured dialogical interactions. But then there are other contexts—for example, auctions and other forms of mediated trading and structured negotiation, medical diagnosis or some legal procedures—where *a priori* univocal shared interpretations (“transparency”) and more or less strict interaction protocols are not only desired but enforced.

From a formal perspective, such highly structured contexts still contain intensional, structural and functional elements that are characteristic of complex dialogical contexts, but lend themselves to a much more straight-forward treatment because of their explicitness. For that reason we take them as the starting point for a systematic study of dialogical systems.

This focusing in highly structured dialogical settings is not only convenient from a formal standpoint, but it is advisable as well from a strictly practical point of view, since many real-world situations that have empirical interest can be construed as structured dialogical interactions. That is, in particular, the case of many of the agent-mediated applications that have been studied in recent years such as agent-mediated trading [42, 38, 29, 35, 9, 7, 27], structured negotiation [16, 17, 45, 44, 24] and cooperative problem solving in general [28, 36]. All of these are reifiable as structured *dialogical interactions* for whose description, modeling, study or utilization adequate formal, methodological and technical devices are worth having at hand.

In this paper we will discuss a formal construct that we have found useful for addressing some of the complexities inherent in those highly structured dialogical settings that involve autonomous agents: the notion of *dialogical institution*. In this occasion we will use a specific example—the COMRIS “scheduling scenes”—to illustrate how social conventions can be made explicit, and to motivate the underlying intuitions concerning the social aspects of an institution. The focus on deontological social aspects will force us to deal with the ontological aspects of the institution and those requirements that the institution imposes on individual participants in a rather sketchy fashion here, however, a more thorough treatment of dialogical institutions and other dialogical constructs can be found in our previous publications: [20, 32, 33, 22].

The notion of institution we develop here is based on North’s economic institutions [23], however we adopt a **dialogical stance** by which the institution itself regulates essentially the illocutory exchange of participating agents. This regulation is achievable because of the following assumptions: participants are always agents (human or software) whose only interactions are illocutions, and all whose *public commitments* are linked to an illocution. Intuitively then, for us, an institution is an entity that makes explicit—and enforces—some restrictions on the illocutions some agents may exchange. These restrictions apply to the language participants use, its intended interpretation and the conditions for

utterance and reception. Dialogical restrictions are dependent on the *role* those agents play and are linked to the specific locations where those illocutions are being exchanged. The institution then can be thought of as a pragmatic device which allows the performance of a dialogue among agents who ought to behave according to that institution’s rules.

This dialogical stance is based on the more classical dialectic traditions (e.g., [1, 18, 14]) and has some affinity with the concerns of the “computational dialectics” proposals (cf. Van Eemeren et al.[41] and Praaken ([25]), although none of the typical “refutational” aspects of dialectics bear upon the notion of institution.

The notion of illocution, as we use it in this paper, inherits the pragmatic concerns of Austin and Searle although we do not need to commit to any specific illocutionary pragmatics. Hence, proposals such as Vandervecken’s [30], Singh’s [34] or Belnap’s [5] could be used to formalize the illocutory aspects we require. And the resulting notion of institution could then adopt a pragmatic interpretation in line with Dignum and Van Linder’s (as advanced in, say, [10]).

We should also acknowledge that, in retrospect, we have become aware of the similitude our proposal for “performative structure” has with Flores, Winograd, Medina-Mora and Ludlow’s notion of conversation and conversation flow ([12, 43, 19]), and in general with some of the issues addressed by the “conversationalist” views of coordination and the so-called “Language-Action-Perspective” as seen in for example Dignum [11], Haddadi [13], Barbuceanu and Fox [2, 3] and, to some degree, with those stated by Traum [37] and Sadek [6].

## 1.1 COMRIS Conference Centre

A conference takes place in a physical setting, the conference centre, where different activities take place in different locations by people that adopt different roles (speaker, session chair, attendant, organization staffer, etc.). During the conference, people pursue their interests moving around the physical locations, and engaging in different activities. In a moment in time people are physically distributed along the conference, possibly interacting with other people: this *physical space* is shown on the right of Figure 1. We can easily think about the spatial proximity relations that exist among people in this space. However, if we think about an *informational space* where the past background and current interests of the conference attendants are represented, we could think of a new kind of *proximity* relation that is a function of the similarity among people’s interests and backgrounds. This *informational* or *virtual space* is shown to the left of Figure 1.

The COMRIS project [39] is about studying and exploiting the synergy of these two spaces, and their relationship, with the objective of supporting and improving the achievement of an attendant’s interests while attending a conference. The approach of COMRIS is to develop software agents inhabiting the virtual space that take up some specific activities on behalf of some interest of an attendant in the conference. Specifically, a Personal Representative Agent (PRA) is an agent inhabiting the virtual space that is in charge of advancing some particular interest of a conference attendant by searching for information and talking to other software agents. An example of a PRA for setting up appointments about a particular topic of interest is later described in §3.

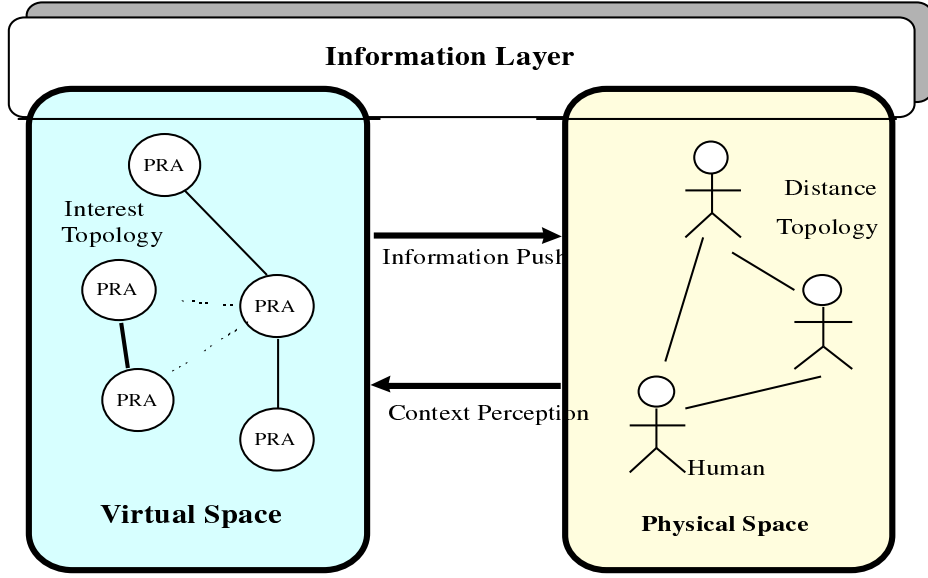


Figure 1: A depiction of the virtual interest-based space and the physical proximity-based space of COMRIS.

An attendant may have several PRAs pursuing her different interests. Attendants have to instruct their PRAs specifying a presentation (e.g. a topic and possibly a collection of related subtopics), an appearance (a collection of features describing the view an agent wants to offer to the other agents) for interacting with other PRAs, and a collection of actions (e.g. meeting people, attending to talks, making appointments, etc) in which the PRA can participate for achieving the attendant’s interests. The collection of actions is provided by the Conference Centre definition as a set of scenes and roles in which a PRA can participate.

When a PRA has gathered information that may be of interest to an attendant it will try to *push* that information to that attendant<sup>2</sup>. The attendants wear a computer device, nicknamed *the parrot*, that runs a speech-generation program and this is able to unobtrusively convey (*push*) that information to him or her using spoken language. It is clear, however, that not any information is relevant for the attendant at any point in time. Rather, the physical situation of the attendant and the concrete activity in which the attendant is involved will determine—or at least bias—whether a particular information is relevant for him at a particular moment in time<sup>3</sup>.

Moreover, several PRAs may compete for having access to a person’s attention in their “information push” activity. COMRIS develops a *competition for attention* mechanism [40] that deals with both sides of the equation: several information reports that compete to be pushed to the attendant and the physical situation of the attendant that determines which of these information reports is more relevant for him or her at that moment in time. We can picture this

<sup>2</sup>As shown in the rightward arrow of Figure 1.

<sup>3</sup>This relevance is shown as the leftward arrow of Figure 1.

competition for attention mechanism as a process that, biased by the *context information*<sup>4</sup> of the physical space selects a “winning” PRA and pushes the information report to his/her *parrot* where it will be orally rendered. For each conference attendant there is a Personal Assistant (PA) agent in the virtual space that embodies this process and “owns” the channel communicating with the attendant’s parrot.

## 1.2 Interactions, Rules, and Institutions

Determining the relevance is only a part of the issue: not only people are involved in located activities in the physical space—attending the talks of a thematic session, meeting a colleague to discuss about a topic, etc—the software agents are also meeting among them in pursue of some person’s interest. In order to analyze this complex network of agents interacting we will presently introduce the notion of an agent-mediated institution [22].

However, it is first appropriate to consider the kinds of activities that take place into a regular Conference Centre: a typical attendant pays the registration, registers at the information desk, assists to thematic sessions, makes appointments and later attends or skip them, etc. Although we can imagine whatever activity taking place in a Conference Centre it is true nonetheless that some of them occur regularly, and the people involved know some rules of behavior for these activities—even if most of this rules are tacit. In fact, a conference may be considered as an *institution* both because these recurring activities occur in it and because conferences are indeed set up for supporting and favoring most of these activities—while other activities are discouraged or plainly forbidden.

These activities the Conference Centre is willing to favor and uphold are the ones COMRIS is interested in analyzing (with agent-mediated institutions) and in providing a computational framework that supports and facilitates them (with the Personal Representative Agents inhabiting the virtual space).

## 2 Agent-Mediated Institutions

Intuitively, an agent-mediated institution is the computational realization of a set of explicit enforceable restrictions imposed on a collection of dialogical agent types that concur in space and time to perform a finite repertoire of satisfiable actions.

The advantages of using this notion of agent-mediated institution are the following:

- it allows to describe in a comprehensive framework the roles and interactions of both human and software agents in a specific setting (the institution);
- it makes explicit the relationship between the computational framework developed by COMRIS and the existing organization of a conference (the institution we will call the Conference Centre);
- it clarifies the difference between:

---

<sup>4</sup>Currently, context information is elaborated by the parrot that determines which other parrots (and hence, attendants) are in its immediate physical surroundings.

1. the “rules of the game” enforced by the Conference Centre upon human and software agents behavior, and
2. the particular strategies the human and software agents may recur to in pursuing individual goals.

The notion of agent-mediated institution is founded on three pillars which contain the elements already mentioned:

1. A *dialogic framework*. Some aspects of an institution are fixed, constituting the context or framework of interaction amongst agents. In a dialogic institution, agents interact through illocutions. Institutions define which are the acceptable illocutions, which is the ontology—including roles, place and time—the common language to refer to the “world”, a common language for communication, and a common metalanguage for, let’s say, establishing differences. An institution can often consider social relationships between participating agents as relevant—in terms of authority, for instance—or some “personality” feature which can affect the interactions. All of these contextual features are what we call dialogic framework.
2. A *performative structure*. Interactions between agents are articulated through agent group meetings, which we call scenes, with a well-defined communication protocol. The Protocol of a scene defines the possible dialogues agents may have.
3. *Behavior rules*. Agent actions in the context of an institution—as far as we are concerned, dialogic actions—have consequences, usually in the shape of compromises which impose restrictions on dialogic actions of agents in the scenes they will act in the future. behavior rules affect the performative structure, putting limits to it.

The intuition behind the notion of institution is that it is a very useful mechanism in specifying autonomous interactive agents. In the following, we will introduce some notation and definitions concerning agent-mediated institutions.

## 2.1 Dialogical Framework

Adapting the definitions of [21, 31] we conclude that the basic notion for defining the ontologic elements is a dialogical framework:

**Definition 1** A *Dialogical Framework*,  $\mathcal{DF}$ , is a tuple

$$\mathcal{DF} = \langle Roles, L, ML, CL, T \rangle,$$

where,

1. Roles is a set of agent types,
2. L is an object language (an ontology),
3. ML is a metalanguage,
4. CL is a communication language, and

5.  $T$  is a model of time.

In other words, concerning the Conference Centre institution, we can summarily exemplify of a dialogical framework  $\mathcal{DF}$  as follows:

1.  $Roles$  is the set of PA's and PRA's types, as well as the types of registered persons (presenters, listeners, invited speakers, etc) and the types of conference staff members (PC member, session chair, information desk staffers, etc)
2.  $L$  is an ontology containing the elements of which the agents talk about, e.g. agendas, time-slots, projects, consortiums, topics of interests, thematic sessions, etc.
3.  $ML$  is the (meta) language used to describe the restrictions the institution imposes on the roles of attendants.
4.  $CL$  is a communication language, that is to say the set of illocutions (performatives) the agents can exchange.
5.  $T$  is the kind of model of time used for the formalization.

## 2.2 Performative Structure

A *performative structure*,  $\mathcal{PS}$ , is the one in charge of modeling the dialogue's dynamics. The definition is based on the notion of scene we are presenting next. The whole scene has an initial state in which illocution exchange between participants in the scene starts. At any moment, the scene is in such state from which it can go out through the transitions that lead to new states. A transition happens if the illocution which labels it is produced. A scene is considered to be finished when it is in one of the possible terminal states. The set of transitions and states determines which is the scene's protocol, i.e. the set of possible dialogic interactions.

**Definition 2** *Given a dialogic framework  $\mathcal{DF} = \langle Roles, L, ML, CL, T \rangle$ , a scene is a tuple  $s = \langle A, R, W, w_0, W_f, \rho \rangle$ , where,*

1.  $A$  is a finite agent variables set.
2.  $R : A \rightarrow Roles$  assigns a role to each agent variable in the scene.
3.  $W$  is a state set.  $w_0 \in W$  is the scene's initial state and  $W_f \subseteq W$  the final states of the scene.
4.  $\rho : W \times CL \rightarrow W$  is a transition function between states, where transitions are labeled with an illocution.

A scene can be “carried out” when we have assigned concrete agents (i.e. programs or humans) for each and every variable of  $A$ . When the assignation has been completed, the scene can start its execution, beginning with the initial state until ending in one of the final states. When a scene is finished, the agents taking part in it will move, as we will see in the definition of performative structure. Some may remain in the scene, though, in order to re-execute it with new agents.

A complex society structure, as is the case of an institution, is formed by multiple interrelated scenes which are executed concurrently or sequentially according to the case. The main point to take into account so as to establish their interrelation is determining which agents leave which scene, in which state they leave it, and which scene they join. Restrictions imposed to this agent flow will greatly determine the flexibility of the formalism. There is a necessary restriction from our point of view, though, and it refers to the integrity/soundness of the agents taking part of a scene. That is, for a transition to be possible between states of a given scene, all the agents of the scene must be present. In other words, if an agent leaves the scene this will remain frozen—no illocutions are broadcasted—until the agent returns. Bearing these restrictions in mind, we can modelize the notion of transition between scenes in a way similar to transitions modeled in a Petri net: with a prior incidence and a post incidence functions. The formalism allows a specific agent to act in more than one scene at a time. This provides the necessary flexibility to modelize the fact that an agent can participate in more than one dialogue at a time.

It is worth mentioning that the notion of scene represents the concept of specifications of a set of agents which follow a given dialogue protocol. Nevertheless, at a given moment an institution can allow for the simultaneous executing of more than one particularization with different actors.

We will note down with sub-index  $S$  the joining for all scenes of the  $S$  set with different components, for instance,  $W_S = \{w | \langle A, R, W, w_0, W_f, \rho \rangle \in S \text{ and } w \in W\}$ .

**Definition 3** *Given a dialogic framework  $\mathcal{DF} = \langle Roles, L, ML, CL, T \rangle$  a performative structure is a tuple  $\mathcal{PS} = \langle S, T, \alpha, \beta \rangle$ , where,*

1.  $S$  is a finite yet not empty set of scenes;
2.  $T$  is a finite yet not empty set of scenes between transitions;
3.  $\alpha : W_S \times T \rightarrow 2^{A^s}$  is the prior incidence function;
4.  $\beta : T \times W_S \rightarrow 2^{A^s}$  is the post incidence function.

verifying that

1. there is an integrity of agents: for every  $t \in T$ ,  $a \in \alpha(w, t)$  for some/any  $w \in W_S$  if and only if  $a \in \beta(t, w')$  for some/any  $w' \in W_S$ .
2. there is an only active state in a scene: for every  $t \in T$ ,  $w, w' \in W_{s'}$  such that  $w \neq w'$ , if  $\beta(t, w) \neq \emptyset$  and  $\beta(t, w') \neq \emptyset$  then  $s \neq s'$ .

## 2.3 Rules of Behavior

Even though scene protocols are necessary to describe agent interactions, they may not be sufficient to make fully explicit the “rules of the game” which all participating agents are supposed to follow in a given institution. These rules will be defined as the *individual rules of behavior* of each agent role. Such rules are, ideally, part of the internal model of each participating agent, but may be required and/or enforced by the institution.






Figure 2: A depiction of part of the performative structure involved in making appointments. Thin boxes are synchronizing points; dot circles are final states.

## 2.4 Institution

Finally, we can group the components of an agent-mediated institution (AMI):

**Definition 4** *An agent-mediated institution,  $I$ , is a 3-tuple  $AMI = \langle \mathcal{DF}, \mathcal{PS}, \mathcal{BR} \rangle$ , where,*

1.  $\mathcal{DF}$  is a dialogical framework
2.  $\mathcal{PS}$  is a performative structure,
3.  $\mathcal{BR}$  are the rules of behavior to which participating agents are subject to.

## 3 Appointment Scenes: An example of agent interaction protocol

We will describe in this section the set of scenes needed for arranging appointments among two conference attendants. The actors involved in these scenes are two conference attendants (noted  $Att_1, Att_2$ ), their personal agents (noted  $PA_1, PA_2$ ), and one personal representative agent for each attendant (noted  $PRA_1, PRA_2$ ).

Appointment scenes require that other scenes are played before (like the Participant Registration Scene) that are not described here. These previous scenes have modeled the state of the different agents:

- each attendant has provided a profile of personal interests that is accessible by their PRAs.
- $PRA_1$  and  $PRA_2$  have established a preliminary contact in the interest-based navigation (IBN) scene. IBN is the scene by which agents identify other agents with similar interests in virtual space.
- each attendant has an individual agenda with a finite number of time slots that she has cleared to be used by her PRAs for making appointments (at most one per time-slot)

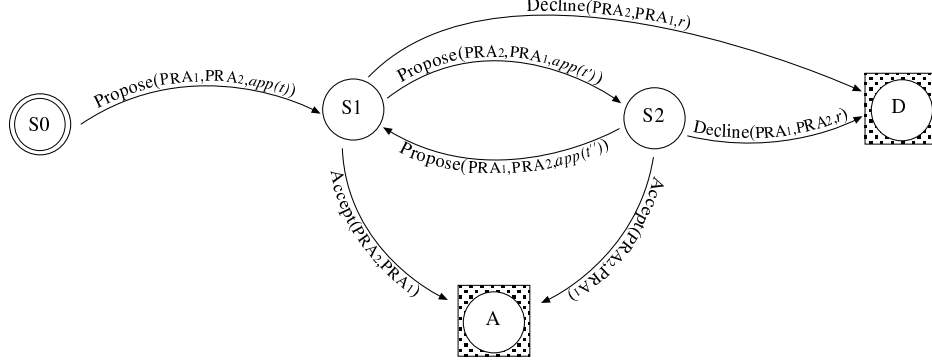


Figure 3: A depiction of a simple interaction protocol for the Appointment Proposal Scene.

Arranging appointments is composed then of two main scenes: the Appointment Proposal Scene (APS) in which PRAs negotiate a set of common topics for discussion (the meeting content); and the Appointment Coordination Scene (ACS) in which PRAs negotiate about the appropriate meeting schedule. When PRAs reach an agreement in APS, they move, together with PAs and attendants into the competition for attention scene. There, PRAs compete for the attendant’s attention informing their PAs about the agreement reached in APS. If both attendants accept the proposal of meeting content, the ACS scene starts for negotiating time-slots. Otherwise ACS is triggered to notify a decline.

### 3.1 Appointment Proposal Scene

The participants of this scene are the personal representative agents  $PRA_1$  and  $PRA_2$ . The goal of the scene is to agree upon a set of topics for discussing in the appointment—represented in Figure 3 as  $app(t)$ ,  $app(t')$ ,  $\dots$  where the argument of  $app$  is the set of topics. Following the interaction protocol shown in Figure 3, the scene is played as follows:

1. one of the PRAs takes the initiative and sends an appointment proposal to the other PRA with a set of initial interest topics. We will refer to the initiating agent  $PRA_1$  and to the other  $PRA_2$ . This interest topics set is intended to be a subset of the attendant’s posted profile of interests.
2.  $PRA_2$  evaluates the proposal and can either (i) accept, (ii) decline, or (iii) send a counter proposal to  $PRA_1$  with a (partially) different set of topics.
3. in turn, when  $PRA_1$  receives the counter proposal of  $PRA_2$  evaluates it and can also either accept, decline, or send a counter proposal to  $PRA_2$ . This negotiation phase finish when an agreement on topics is reached or one of them decides to withdraw it.

When PRAs reach the acceptance final state, the scene finishes and each PRA move to a competition for attention scene with their corresponding PAs and attendants. Remark that, up to this point, no time commitment was

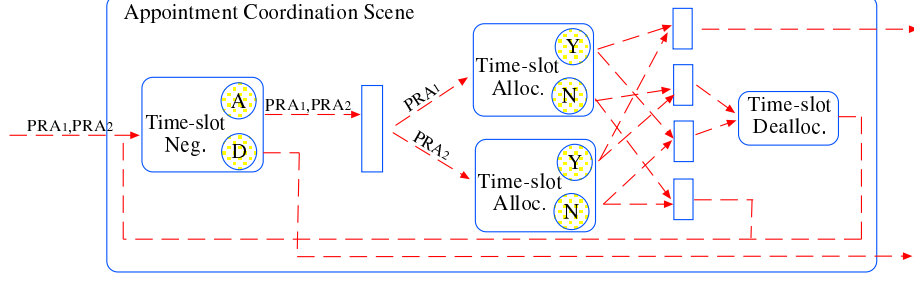


Figure 4: A depiction of the interaction protocol of the Appointment Coordination (sub)scenes.

made—thus no change in the agenda of attendants have yet taken place. Nevertheless, the acceptance of appointment contents in CFA scene involves the commitment of accepting the appointment if PRAs reach a time-slot agreement in ACS scene. We adopt this lazy commitment strategy in order to facilitate the management of the agenda.

### 3.2 Appointment Coordination Scene

The precondition for this scene is that attendants have accepted the appointment content. ACS is decomposed in two (sub)scenes: time-slot negotiation and time-slot allocation. During the first subsceine PRAs negotiate about time-slots without blocking them. When a given time-slot is accepted by both, each of them tries to block the time-slot in its agenda. When both PRAs succeeded in blocking the time-slot the appointment is fixed. Otherwise, blocked time-slots are unblocked and the time-slot negotiation scene starts again. Figure 4 shows the scene’s decomposition.

The scene is composed of the following phases:

1.  $PRA_1$  and  $PRA_2$  negotiate a time-slot following the same negotiation protocol used in negotiating the appointment topics (see Figure 3).
2. When PRAs reach the decline final state, the scene finishes and the negotiation ends.
3. When PRAs reach the acceptance final state, the scene finishes and each PRA move to a time-slot allocation scene. In this scene whether the time-slot is unblocked in its agenda, it is blocked; otherwise allocation fails.
4. if both PRAs have achieved in blocking the time-slot the appointment is fixed and the notification will be sent to attendants.
5. if only one PRA fails in allocating the time-slot, the other PRA unblock the time-slot in its agenda and both PRAs move again into the time-slot negotiation scene.
6. when neither  $PRA_1$  and  $PRA_2$  achieve to block the time-slot, they move directly to the time-slot negotiation scene.

## 4 Other Scenes

We are currently specifying other scenes and implementing the PRAs that will carry out the conference attendants' interests in those scenes.

1. *Proximity Alert*: in this scene an attendant is informed that she is physically near to another person with similar interests. Variants of this include proximity to exhibitor booths or thematic sessions with topics similar to the user's.
2. *Commitment Reminder*: this is a type of alert that reminds an attendant of the proximity of a deadline she is committed to (e.g. an appointment, but also commitments with the Conference organization, like chairing a session that is about to start, or boarding a bus that is about to leave for a tour the user has paid for).
3. *Consortium and Project Formation*: one of the more important tasks in conferences like the ESPRIT Week. Similar to appointment scenes but involves several people with different interests that are relevant to separate aspects of a project.
4. *Propagandist*: in this scene a PRA from someone like a thematic session chairperson or an exhibition booth responsible tries to find and attract persons that might be interested in the topics of the session or of the exhibition. Similar to the Appointment scene but there is no need to negotiate a time-slot (session is fixed and exhibitor would prefer to remind the user when he is near and not in a particular time-slot—also similar to Proximity Alert).

## 5 Discussion

We have described a couple of scenes for making appointments that involved Personal Representative Agents inhabiting the virtual space. These scenes also involved the people inhabiting the physical space through the mechanism of competition for attention. It is worth noticing that the two protocols of interaction described determine basically the illocutions that the agents use (e.g. **propose** **accept**, **decline**) and the meaningful relations among them by means of a finite state machine (e.g. a **decline** makes sense only after a **propose** has been issued). The *propositional content* of each illocution is left open, and even the representation in which this content is conveyed is left open. Indeed a particular implementation of the agent-mediated Conference Centre institution would have to decide on the ontology and the representation (XML, KIF, Horn clauses, etc) to be used; however the point is that the current level of abstraction is the adequate to analyze the situations in which we want to analyze a social scenario where human and software agents interact in a meaningful way.

The notion of agent-mediated institution also leaves open the internals of specific agents: the preferences and decision procedure by which a software agent makes a particular action (e.g. a particular **propose**). However only in a particular framework of interaction shared by the agents involved these decision processes can take place, and this is precisely what agent-mediated institutions has to offer.

## Acknowledgements

The research reported in this paper is supported by the ESPRIT LTR 25500-COMRIS *Co-Habited Mixed-Reality Information Spaces* project.

## References

- [1] Aristotle. *Prior Analytics*. Number 391 in Loeb Classical Library. Harvard University Press, 1960.
- [2] Mihai Barbuceanu and Mark S. Fox. Cool: A language for describing coordination in multi agent systems. In *ICMAS-95*, 1995.
- [3] Mihai Barbuceanu and Mark S. Fox. The architecture of an agent building shell. In *ATAL-95*, pages 235–250, 1996.
- [4] Jon Barwise and John Perry. *Situations and Attitudes*. MIT Press, 1983.
- [5] N. Belnap and M. Perloff. In the realm of agents. *Annals of Mathematics and Artificial Intelligence*, 9:25–48, 1993.
- [6] P. Bretier and M. D. Sadek. A rational agent as the kernel of a cooperative spoken dialogue system: Implementing a logical theory of interaction. In J. Müller, M. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III (LNAI Volume 1193)*, LNCS, pages 189–203. Springer Verlag, 1996.
- [7] A. Chavez, D. Dreilinger, R. Guttman, and Pattie Maes. A real life experiment in creating an agent marketplace. In *Proc. of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'97)*, 1997.
- [8] William Clancey, Stephen Smoliar, and Mark Stefik. *Contemplating Minds. A forum for Artificial Intelligence*. MIT Press, 1994.
- [9] Clearwater. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific, 1995.
- [10] F. Dignum and B. van Linder. Modeling social agents: Communication as action. In J. Mueller, M. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III (LNAI Volume 1193)*, LNCS, pages 205–218. Springer Verlag, 1996.
- [11] Frank Dignum, Jan Deitz, Egon Verharen, and Hans Weigand. Communication modelling, The Language Action Perspective. In *Proceedings of the First International Workshop on Communication Modelling, Tilburg, The Netherlands 1-2 July 1996*, eWIC. Springer, 1996.
- [12] Carlos F. Flores. *Inventando la Empresa del Siglo XXI*. Hachette, 1989. (A revised translation of "Management and Communication in the Office of the Future". Unpublished Ph. d. dissertation, Univ. of California at Berkley. 1981).

- [13] A. Haddadi. A hybrid architecture for multi-agent systems. In S. M. Deen, editor, *Proceedings of the 1993 Workshop on Cooperating Knowledge Based Systems (CKBS-93)*, pages 13–26. DAKE Centre, University of Keele, UK, 1993.
- [14] C. L. Hamblin. *Fallacies*. Methuen, 1970.
- [15] J. Hintikka. *Logic, Language Games and Information*. Clarendon Press, 1973.
- [16] N. R. Jennings, P. Faratin, M. J. Johnson, T. J. Norman, P. O’Brien, and M. E. Wiegand. Agent-based business process management. *International Journal of Cooperative Information Systems*, 5(2&3):105–130, 1996.
- [17] S. Kraus, J. Wilkenfeld, and G. Zlotkin. Multiagent negotiation under time constraints. *Artificial Intelligence*, 75:297–345, 1995.
- [18] Paul Lorenzen. *Normative Logic and Ethics*. Bibliographisches Institut, Zurich, 1967.
- [19] R. Medina Mora, Terry Winograd, Rodrigo Flores, and Fernando Flores. The action workflow approach to workflow management technology. In *CSCW 92*, 1992.
- [20] P. Noriega and C. Sierra. Towards layered dialogical agents. In J. Müller, M. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III (LNAI Volume 1193)*, LNCS, pages 157–171. Springer Verlag, 1996.
- [21] P. Noriega and C. Sierra. Towards layered dialogical agents. In *Proceedings of the ECAI’96 Workshop ATAL’96*, number 1193 in Lecture Notes in Artificial Intelligence, pages 157–171. Springer-Verlag, 1996.
- [22] Pablo Noriega. *Agent-Mediated Auctions: The Fishmarket Metaphor*. PhD thesis, Universitat Autònoma de Barcelona, 1997.
- [23] D. C. North. *Institutions, Institutional Change and Economic Performance*. Cambridge Univ. Press, Cambridge, U.K., 1990.
- [24] S. Parsons and N. R. Jennings. Negotiation through argumentation—a preliminary report. In *Proc. Second Int. Conf. on Multi-Agent Systems, ICMAS’96*, pages 267–274, Kyoto, Japan, 1996.
- [25] H. Praaken and G. Sartor. A dialectical model of assessing conflicting arguments in legal reasoning. *Artificial Intelligence and Law Journal*, Vol 3(3-4):pp 331–368, 1996.
- [26] N. Rescher. *Dialectics: A controversy-oriented approach to the theory of knowledge*. SUNY, 1977.
- [27] Juan A. Rodríguez-Aguilar, Francisco J. Martín, Pablo Noriega, Pere Garcia, and Carles Sierra. Towards a test-bed for trading agents in electronic auction markets. *AI Communications*, 11(1):5–19, 1998.
- [28] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter*. The MIT Press, Cambridge, USA, 1994.

- [29] Rina Schwartz and Sarit Kraus. Bidding mechanisms for data allocation in multi-agent environments. In *ATAL97*, 1997.
- [30] John R. Searle and Daniel Vanderveken. *Foundations of illocutionary logic*. Cambridge University Press, 1985.
- [31] C. Sierra, N. R. Jennings, P. Noriega, and S. Parsons. A framework for argumentation-based negotiation. In M. P. Singh, A. Rao, and M. S. Wooldridge, editors, *Intelligent Agents IV*, number 1365 in Lecture Notes in Artificial Intelligence, pages 177–192. Springer-Verlag, 1997.
- [32] Carles Sierra, Nick R. Jennings, Pablo Noriega, and Simon Parsons. A framework for argumentation-based negotiation. In A. Rao M. P. Singh and M. J. Wooldridge, editors, *Intelligent Agents IV (LNAI Volume 1365)*, LNCS, pages 177–192. Springer Verlag, 1997.
- [33] Carles Sierra and Pablo Noriega. A formal framework for accountable agent interactions. In *BISFAI '97*, pages 23–24, 1997.
- [34] M. P. Singh. A semantics for speech acts. *Annals of Mathematics and Artificial Intelligence*, 8:47–71, 1993.
- [35] R. G. Smith and R. Davis. Frameworks for cooperation in distributed problem solving. *IEEE Trans on Systems, Man and Cybernetics*, 11(1):61–70, 1981.
- [36] P. Stone and M. Veloso. Multiagent systems: A survey from a machine learning perspective. In *Submitted to IEEE Transactions on Data Engineering(TKDE)*, 1996.
- [37] D. R. Traum. A reactive-deliberative model of dialogue agency. In J. Müller, M. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III (LNAI Volume 1193)*, LNCS, pages 157–171. Springer Verlag, 1996.
- [38] M. Tsvetovatyy and M. Gini. Towards a virtual marketplace: Architecture and strategies. In *Proceedings of PAAM96*, 1996.
- [39] Walter Van de Velde. Co-habited mixed reality. In *Proceedings of IJCAI'97 workshop on Social Interaction and Communityware*, 1997.
- [40] Walter Van de Velde, Sabine Geldof, and Ronald Schrooten. Competition for attention. In *Proceedings ATAL-97*, 1997.
- [41] H. van Eemeren, R. Grootendorst, and F. Henkemanns. *Fundamentals of Argumentation Theory, A Handbook of Historical Backgrounds and Contemporary Developments*. Lawrence Erlbaum Associates, 1996.
- [42] M. P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–22, 1993.
- [43] T. Winograd and F. Flores. *Understanding Computers & Cognition*. Ablex, 1986.

- [44] Peter R. Wurman, , Michael P. Wellman, and William E. Walsh. The michigan internet auctionbot: A configurable auction server for human and software agents. In *Second International Conference on Autonomous Agents (AGENTS'98)*, 1998.
- [45] G. Zlotkin and J.S.Rosenchein. Mechanisms for automated negotiation in state oriented domains. *Journal of Artificial Intelligence Research*, 5:163–238, 1996.