Mixed Initiative in Interactions between Software Agents

Daniel Rousseau

Knowledge Systems Laboratory
Gates Computer Science Bldg 2A
Stanford, CA 94305-9020, USA
Email: rousseau@hpp.stanford.edu
Phone: (415) 723-0948

Bernard Moulin

Computer Science Department
Laval University
Sainte-Foy (Quebec), G1K 7P4, Canada
Email: moulin@ift.ulaval.ca
Phone: (418) 656-5580

Abstract

We have been working during the past several years on techniques for modeling the way that software agents can take and release the initiative while interacting together. We are interested in building multiagent systems composed of software agents that can interact with human users in sophisticated ways which are analogous to human conversations. In this paper, we describe two projects we have worked on: a multiagent approach for simulating conversations between software agents, and the Virtual Theater.

1. Introduction

The need for software agents that assist users in achieving various tasks, collaborate with them, entertain them, or even act on their behalf is getting greater. Software agents are computer systems that exploit their own knowledge bases, have their own goals and their own capabilities, perform actions, and interact with other agents as well as with people. Autonomy is an essential characteristic of such agents, which they express when they take the initiative.

Agents take the initiative when they decide to act on their own because another agent asked them to do something or to provide some information, or because of a personal goal to be satisfied. In fact, software agents are not very different from human beings at this level, because they have their own individuality, and they can also share goals and form teams in order to satisfy common goals.

We have been working during the past several years on techniques for modeling the way that software agents can take and release the initiative while interacting together. We are interested in building multiagent systems composed of software agents that can interact with human users in sophisticated ways which are analogous to human conversations. In this paper, we will describe how we deal with initiative in two projects we have worked on:

- an approach based on human conversations for modeling interactions between software agents;
- the interaction between synthetic actors able to receive directions and improvise their behavior in multimedia environments.

2. Modeling conversations in a multiagent universe

In order to explore new interaction modes for software agents which need to be more sophisticated than simple exchanges of messages, we have analyzed human conversations and elaborated an interaction approach for software agents based on a conversation model. Using this approach, we have developed a multiagent system that simulates conversations involving software agents.

Initiative is an important aspect of human conversations. Interlocutors take the initiative alternatively to inform, ask questions, order, commit, express their feelings, etc. Sometimes, a speaker keeps the initiative most of the time, as when a researcher gives a formal talk, or a pastor preaches about the gospel. In other circumstances, any of the interlocutors can take the initiative when he or she feels like talking, as when friends are involved in an informal discussion.

In interactions between synthetic agents, or between synthetic agents and persons, it is the same. Interactions are richer when the agents can all take the initiative when they feel like doing it, rather than being limited to reacting to what another agent does.

Human conversations are usually characterized by mixed initiative, either if they are dialogues or group conversations. It is not always the same person who speaks, but usually people do not speak at the same time. They know when it is time to speak and when it is better to listen to another person. We have applied the same mechanisms to interactions between synthetic and human agents. Those agents have a common goal, to have a satisfying conversation, but they also have their own goals that they would like to satisfy through the conversation. Taking the initiative allows them to satisfy such goals.

We developed a conceptual framework for modeling and simulating conversations which integrates several techniques found in conversation analysis, speech act theory, computational linguistics, artificial intelligence, multiagent systems, planning, and cognitive psychology. This framework is plausible according to socio-linguistic analyses of conversations, and is adequate for simulating

sophisticated interactions which take place between software agents. In such a framework, we consider that all software agents are autonomous and able to reason on mental states such as goals, beliefs, and capabilities, plan their behavior and perform actions in order to influence other agents' mental models or to modify the environment. Software agents build a model of each conversation they are involved in. Such a model contains the knowledge structures describing different types of information exchanged by the interlocutors, that we call *locutor-agents* (Moulin & Rousseau 1996).

We consider that a conversation unfolds with respect to different levels of communication. Locutor-agents perform communicative acts in order to take the initiative, to transfer concepts, to manage the subjects discussed in a conversation as well as the quality of the information, and to maintain the communication. Communicative acts are composed of speech acts or gestures that influence the agents' mental models. We assume that software agents can interpret and plan communicative acts using *life cycles* for each of the different levels of communication. A life cycle for a given level of communication shows the states and the transitions between states that are allowed during a conversation for such a level. In the rest of this section, we describe the *initiative* life cycle. The negotiation life cycle, regarding the concepts that are transferred, and the conversation life cycle, regarding the global states of a conversation, are described in (Moulin & Rousseau 1996).

Managing the initiative was studied by conversation analysts, such as Sacks et al. (1978) and Francis and Hunston (1992). They noticed several ways for a locutor to take or release the initiative during a conversation. Researchers in artificial intelligence have also considered such an aspect of a conversation, including Cawsey (1992), and Traum and Hinkelman (1992). We used all those works to create a life cycle that would model the conversation protocol regarding the initiative during a conversation between software agents. Transitions in this life cycle are realized by the performance of certain kinds of actions. For instance, agents can take the initiative if they speak. They can release it if they stop speaking. They can give the initiative to a specific agent if they ask that agent a question. Such an approach can be applied to dialogues as well as to group conversations. It reflects the fact that locutor-agents negotiate and compete in a conversation.

In our approach, each locutor agent keeps track of the state of the initiative for each participant of the conversation in a special knowledge structure called *initiative agenda*. Such a structure indicates for each agent the current state of the initiative, according to the life cycle shown in Figure 1. We describe each transition and identify it by a number set between parentheses.

The locutor-agent that starts the conversation takes the initiative by performing a communicative act (2) while the other agents are waiting for the initiative (1). Waiting for the initiative implies listening to the speaker and waiting for it to finish speaking before saying something. An agent can perform several communicative acts while it has the initiative (8). It can release the initiative by offering it to a specific agent through a question or an order (12), or by leaving it to any agent that wants to speak (23). This last case occurs when a locutor-agent stops speaking or lowers its intonation, giving the opportunity to other agents to take the initiative. When another agent takes the initiative, the agent that has just released the initiative usually waits for it again (13, 19). It can also choose to leave the conversation without listening to what the speaker has to say (14, 21). If the initiative is released and no one else takes it, the agent that released the initiative may take it again (12, 24).

A locutor agent can wait for the initiative while one or more agents perform one or more communicative acts (3). The initiative is usually taken when an agent who is waiting for the initiative considers that it has the right to speak (5) because the former speaker has just released the initiative (12, 23). An agent who is waiting for the initiative can also ask the speaker for it by, for instance, raising its hands (4). Afterward, the initiative may be granted (6) or denied to the agent that requested it. The initiative can also be interrupted if an agent speaks before the speaker releases it (16). The interrupted speaker is allowed to ask for the initiative (15), take it back by speaking (10), wait for it (17) or just leave the conversation (22).

A prototype, PSICO, has been implemented to test our approach. PSICO is composed of several software agents able to plan, perform, and interpret communicative acts with respect to the different life cycles.

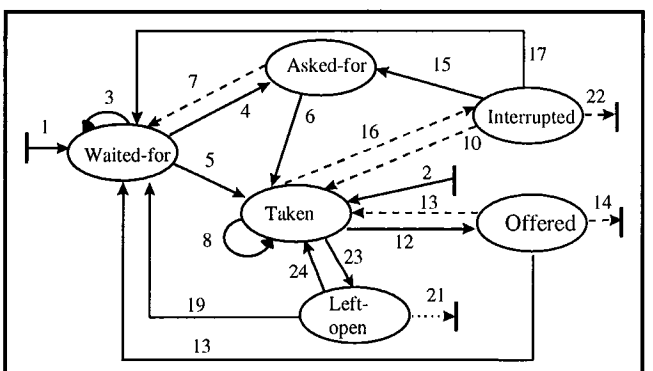

Figure 1: Initiative Life Cycle

3. Interaction in the Virtual Theater

In collaboration with Barbara Hayes-Roth at Stanford University, we are working on an interactive system called the Virtual Theater (Hayes-Roth & van Gent 1996). The Virtual Theater is a system providing autonomous, synthetic actors and avatars portraying characters in a textual or graphical environment. Users play a given role in the context of a scenario, and interact with the other actors through their avatars in order to build stories interactively. Synthetic actors are able to improvise their behaviors with respect to the directives they receive from a predefined script or from users. An example of application is the Cybercafé (Rousseau & Hayes-Roth 1996), in which an autonomous agent plays the role of a waiter, and a user portrays a customer through his or her avatar. The number of autonomous agents and avatars varies depending on the application. Each user has his or her own screen containing a menu that presents the actions that can be selected for his or her avatar in the current context, and either a textual description or an animation of the story.

Life cycles are used to define the possible states and transitions of a character at an abstract level. For instance, a customer of the Cybercafé has a life cycle specifying the possible states and transitions concerning the order and serving of food, and another life cycle to know when to walk, stand, or sit. A life cycle specifies how each transition can be realized, who should realize it, and which actions are not allowed when the context is in a given state. In the life cycle specifying whether a customer is standing or sitting, the customer who is sitting cannot walk or sit, but will be standing if he or she stands up.

In the context of improvisation, users and autonomous actors must decide when to lead and when to let the others take the initiative. Most of the initiative management is performed through life cycles in the Virtual Theater. We identify three types of situations regarding the initiative:

- *A user who wants to communicate with another user through his or her avatar.*
  He or she can do it at any time by typing what he or she wants to communicate in a special window. The message is transmitted as it is to the addressee.
- *A user who wants to interact with an autonomous actor.*
  He or she can select options corresponding to actions in a menu. Usually, it is done while autonomous actors are waiting for an avatar's action. The actions that are available in the menu shown to a user depend on the current context according to the avatar's life cycles.
- *An autonomous actor who wants to communicate with another actor.*
  Each actor has life cycles that helps itself to decide if it should perform an action or wait. Those life cycles take into account turn taking by considering when it is time to act and when it is time to wait for another agent's reaction. For instance, a waiter who asks a customer if he or she is ready to order knows that he should wait for the customer's action before acting again.

Allowing those three types of situations encourages the users and the autonomous actors to improvise.

4. Conclusion

We have described two computer applications in which we consider initiative in interactions between software agents using life cycles. The first application, based on human conversations, allows the simulation of conversations between two or more agents. It includes a specific life cycle to deal with turn taking. The second application, the Virtual Theater, integrates initiative management in different life cycles. Those life cycles are used by actors able to improvise their behavior.

Both approaches allow any agent to take the initiative and prescribe when it is acceptable to do so. The first approach is more flexible, because it identifies clearly the different situations regarding the initiative at an abstract level. The second approach is not as general, but requires identifying the transitions at a less abstract level. Deciding which approach is best depends on the desired strategy.

References

Cawsey, A. 1992. *Explanation and Interaction: The Computer Generation of Explanatory Dialogues.* The MIT Press, Cambridge.

Francis, G.; Hunston, S. 1992. Analysing Everyday Conversation. *Advances in Spoken Discourse Analysis*, M. Coulthard ed., Routledge, 123-161.

Hayes-Roth, B.; van Gent, R. 1996. *Story-Making with Improvisational Puppets and Actors.* Technical Report KSL-96-05, Knowledge Systems Laboratory, Stanford.

Moulin, B.; Rousseau, D. 1996. An Approach for Modelling and Simulating Conversations. *Essays in Speech Act Theory*, D. Vanderveken and S. Kubo eds., John Benjamins Publishing Company. Also Technical Report DIUL-RR-9402, Laval University, May 1994.

Rousseau, D.; Hayes-Roth, B. 1996. *Personality in Synthetic Agents.* Technical Report KSL-96-21, Knowledge Systems Laboratory, Stanford.

Sacks, H.; Schegloff, E. A.; Jefferson, G. 1978. A Simplest Systematics for the Organization of Turn Taking for Conversation. *Studies in the Organization of Conversational Interaction*, J. Schenkein ed., Academic Press, New York, 7-55.

Traum, D. R.; Hinkelman, E. A. 1992. Conversation Acts in Task-Oriented Spoken Language, *Computational Intelligence* 8: 3.