

ARTIMIS: Natural Dialogue Meets Rational Agency

M.D. Sadek, P. Bretier, & F. Panaget

France Telecom — CNET

DIH/RCP

Technopole Aiiiticipa - 2, avenue Pierre Marzin - 22307 Lannion Cedex France

Emails: {sadek, bretier, panaget}@lannion.cnet.fr

Abstract

We present an effective generic communicating rational agent, ARTIMIS, and its application to cooperative spoken dialogue. ARTIMIS' kernel is the implementation of a formal theory of interaction. This theory involves a set of generic axioms which models, in a homogeneous logical framework, principles of rational behaviour, communication, and cooperation. The theory is interpreted by a specifically designed reasoning engine. When applied to the context of natural dialogue, ARTIMIS includes specialised components for speech and natural language processing.

1 Introduction

Natural human-computer dialogue is obviously a challenging application for intelligent systems: a system designed ultimately to replace a human operator must satisfy stronger requirements than most computer-based tools. For example, cooperative behaviour, such as providing suggestive answers, is an obvious prerequisite for acceptance by humans. This alone favours a generic intelligent-agent-based approach to automated dialogue. Furthermore, designing artificial autonomous agents to be kernels of intelligent systems requires the development of formal theories of reasoning and interaction.

In this paper, we describe the basic components of an effective generic communicating rational agent, ARTIMIS, applied to cooperative spoken dialogue [Sadek *et al.*, 1996]. ARTIMIS' kernel is the implementation of a formal theory of interaction. This theory involves a set of generic axioms that models, in a homogeneous logical framework, principles of rational behaviour, communication, and cooperation. It thus supports the rational unit of an autonomous communicating agent. It is expressed in a first-order (multi)modal logic of mental attitudes (belief, uncertainty, and intention) and actions.

Section 2 gives the context of the application, illustrated by a sample dialogue. Section 3 describes the heart of the system, a *rational unit*, which is a generic homogeneous logical theory interpreted by a specifically designed inference engine. This rational unit embodies

a real demonstration of the applicability of our generic AI approach to dialogue. Section 4 describes specialised components devoted to deal with some of the problems relative to natural language input and output processing. Section 5 presents a reasoning example. Finally, ARTIMIS is compared to related work in the field of intelligent agents.

2 Application and sample dialogue

ARTIMIS is currently being used on the AGS domain [Sadek *et al.*, 1996], i.e., the spoken query, over the telephone, of information voice services directory hosted by France Telecom (Audiotel services). We have chosen to begin with the following restricted domain and semantic representation level: (1) the system only handles two major classes of services: weather forecasts and employment, and (2) the services are described by relevant attributes: theme, geographic area, etc. With this kind of semantic information as a background, the system accepts nearly unconstrained spontaneous language as input, and reacts in a cooperative way (with natural-sounding output). Below is a sample dialogue transcription. S stands for the system, U for the user, R for the output of the speech recogniser, and XXX for "semantic garbage".

- S₁: *Bienvenue sur le serveur A.G.S ARTIMIS. Que puis-je pour vous ?*
Welcome to the A.G.S. ARTIMIS server. How can I help you?
- U₁: *Je recherche un emploi a Calais. Est ce que il y a des serveurs ?*
I'm looking for a job in the Calais area. Are there any servers?
- R₁: *je XXX emploi a Calais est ce que il y a XXX*
Non, il n'y a pas de serveurs d'emploi pour Calais. Par contre, il y a un serveur d'emploi pour le Pas-de-Calais et un serveur d'emploi pour Lille, Lun d'eux vous convient-il ?
No, there aren't any employment servers for Calais. However, there is an employment server for Pas-de-Calais and an employment server for Lille. Are you interested in one of these?
- U₂: *Pour le Nord pas Calais.*
For Nord, not for Calais.
- R₂: *XXX nord pas Calais.*
- S₃: *Le numero du serveur de petites annonces d'emploi pour le Nord est le 36 68 74 44. Desirez-vous un autre renseignement ?*

The phone number of the server of employment advertisements for Nord is 36 68 74 44. Do you want any other information?

- U3: Est-ce qu'il est accessible de Calais ?
Is it accessible from Calais ?
- R3: est ce que XXX accessible de Calais
- S4: Oui, il est accessible de Calais. Desirez-vous un autre renseignement ?
Yes, it is accessible from Calais. Do you want any other information?
- U4: Je veux aussi la meteo.
I also want the weather forecasts.
- R4: XXX XXX meteo
- S5: Etes-vous interesse par la meteo generate ou par la meteo specialisee ?
Are you interested in general weather forecasts or in specialised weather forecasts?
- U5: La meteo specialisee.
Specialised weather forecasts.
- R5: XXX meteo specialisee
- S6: Un serveur de quelle meteo specialisee ?
A server for which specialised weather forecasts?
- U6: Que me proposez vous ?
What can you suggest?
- R6: XXX proposez XXX
- S7: Il y a huit types de meteo specialises. Les deux premieres sont la meteo aerienne et la meteo cotiere. Vous recherchez un serveur de quelle meteo specialisee ?
There are eight kinds of specialised weather forecasts. The two first are atmospheric weather forecasts and coastal weather forecasts. Which specialised weather forecast are you interested in?
- U7: Cotiere pour le Nord.
Coastal, for the Nord.
- R7: cotiere XXX nord.
- S8: Le numero du serveur de meteo cotieres pour le Nord est le 36 68 08 59. Desirez-vous un autre renseignement ?
The phone number of the server of coastal weather forecasts for Nord is 36 68 08 59. Do you want any other information?
- U8: Quel est son fournisseur ?
Who is the provider?
- R8: quel XXX fournisseur
- S9: Le fournisseur du 36 68 08 59 est Meteo-France. Desirez-vous un autre renseignement ?
The provider of 36 68 08 59 is Meteo-France. Do you want any other information?
- U9: ...

3 The Rational Unit

The central theme of our approach is that an intelligent dialogue system has to be an intelligent system first of all. Furthermore, the commonly accepted prototypical intelligent behavior is rational behavior. In a simplified way, to behave rationally is to be permanently driven, at a certain representation level, by principles that optimally select the actions leading to those futures in conformity with a given set of motivations and desires (see, e.g., [Cohen and Levesque, 1990a; 1990b] and [Sadek, 1991a; 1994]). It is at this (hypothesised) Knowledge Level that the concepts of mental attitudes and intentional actions are relevant.

The rational unit is the kernel of an intelligent agent. It gives the system its dialogue abilities, which result

from explicit reasoning processes, modelled by the formal theory of rational interaction. The theory itself will only be outlined here (for more details, see [Sadek, 1991a; 1991b; 1994; 1992; Bretier, 1995]).

In regards to the formal approach for knowledge representation, the logic framework is adequate, for various reasons: homogeneity, genericity (due to its large coverage), ability to intuitively account for mental attitudes (which makes it easy to maintain), and its potential usability, both as a modelling and an implementation tool.

Based on an integrated formal model of mental attitudes and rational action, the theory provides a unified account, expressed in a homogeneous logical framework, for the different constituents and capabilities involved in (cooperative) communication, such as the management of the logical relationships over the system's mental attitudes (beliefs, uncertainties, intentions, and plans). Communicative acts are introduced as regular actions, which can be recognised and planned using general principles of rational behaviour. In this framework, dialogue is viewed as a derivable activity, which relies on more primitive abilities: it dynamically results from explicit reasoning processes motivated by the observation and the planning of communicative acts. Due to the genericity of the principles, this approach achieves the robustness required by an intelligent dialogue system: to soundly react to complex situations, possibly incompletely specified when the system has been designed.

3.1 Logical Framework

The theory is expressed in a logic of attitudes and actions (or events), formalised in a first-order modal language (close in several aspects to that of [Cohen and Levesque, 1990a]). Let us briefly outline the part of formalism used in this paper. The symbols p , q , and r are taken to be (closed) formulae, ϕ and ψ formula schemata, and i and j schematic variables denoting agents. Logical implication will be written using \Rightarrow . We will write $\models \phi$ to mean that ϕ is valid. The constraints imposed on the semantic model will be introduced as **axioms**.

To allow reasoning about action, the universe of discourse involves sequences of events, in addition to individual objects and agents. Events can be combined to form action expressions, such as sequential actions $a_1; a_2$, or nondeterministic choice actions $a_1 | a_2$. where a_1 and a_2 denote primitive sequences of events, or action expressions. The language involves variables e , e_1 , ... ranging over the set of event sequences, variable schemata a , a_1 , ... for action expressions, and the operators *Feasible* and *Done*. The formula *Feasible(a, p)* means that a can take place and if it does p will be true after that action. The formula *Done(a, p)* means that a has just taken place and p was true before that action. The abbreviations *Feasible(a)*, *Done(a)* and *Possible(ϕ)*, will be used for *Feasible(a, true)*, *Done(a, true)*, and $(\exists e) \text{Feasible}(e, \phi)$, respectively.

The mental model of an agent is based on three attitudes: *belief*, *uncertainty*, and *choice* (or, to some extent, *goal*). The attitude of *intention* is defined as a compos-

ite concept defined in terms of belief and choice [Sadek, 1991a; 1992]. In this paper, we only use belief and intention (intention is taken as a macro-attitude). These two attitudes are formalised by the modal operators B and I , respectively. Formulae such as Bp , and Up can be read as "i (implicitly) believes (that) p" and "i intends to bring about p", respectively. The modelled agents are taken to be fully introspective and to have consistent beliefs. Thus, the logical model for the operator B is a *KD45-model*.

3.2 Communicative Acts

We concentrate on communicative acts but the analysis applies to rational (or intentional) actions in general. By communicative act (CA), we mean any action performed by an agent with the intention that it be observed by (at least) one other agent.

As for an ordinary rational action, the components of a CA model that are involved in a planning process must characterise, on the one hand, the reason why this act is "selected," and, on the other hand, the conditions that have to be satisfied in order to plan the act. We call the first component the rational effect (RE) of the act and the second one its feasibility preconditions (FPs). It is worth noting that the RE is the effect intended by the agent of the act; and whether it will really take place or not, only concerns the recipient and has not to (and cannot validly) be captured within the act model.

We introduce here without elaborated discussion simplified, yet operational, versions of the action models that we use later (see [Sadek, 1991a; 1991b] for a detailed analysis). The models in question are those of the act of i Informing j that ϕ , and the (abstract) act of i Informing j of a (referent of the term) δ (or *the* δ , if the operator t for definite description is available):

$$\begin{aligned} < i, Inform(j, \phi) > \\ & FP: B_i \phi \wedge \neg B_j \phi \\ & RE: B_j \phi \\ < i, Informref(j, \delta(x)) > \\ & FP: (\exists x) B_i \delta(x) \wedge \neg B_i (\exists x) B_j \delta(x) \\ & RE: (\exists x) B_j \delta(x) \end{aligned}$$

Rational action models, and particularly CA models, are implemented in this theory using the axioms of logical rationality presented in the next section.

3.3 Principles of Logical Rationality

The first property we introduce follows directly from the definition of intention (see, e.g., [Sadek, 1992]) and expresses the fact that an agent can adopt the intention to bring about a state of affairs only if she believes that this state does not currently hold:

$$\text{Axiom A1: } \models I_i \phi \Rightarrow B_i \neg \phi.$$

Axiom A2 provides an agent with the capability to plan an act, whenever she intends to achieve its RE. The agent's intention to achieve a goal generates the intention to perform one of the acts (1) that have the goal as RE

and (2) the agent has no reason for not carrying it out.¹

Axiom A2: $\models I_i p \Rightarrow I_i Done(a_1 | \dots | a_n)$, where a_k , $k \in \{1 \dots n\}$, are the actions such that: (1) p is the RE of a_k , and (2) $\neg I_i \neg Possible(Done(a_k))$.

Axiom A3 requires an agent to search for the satisfiability of the FPs of an act whenever she selects that act (by axiom A2):

$$\text{Axiom A3: } \models I_i Done(a) \Rightarrow (BiFeasible(a) \vee I_i B_i Feasible(a)).$$

Obviously, an action is believed feasible by an agent when all of its FPs (from the agent's point of view, i.e., what the agent believes to be the action's FPs) are satisfied. This is expressed by the following axiom:

$$\text{Axiom A4: } \models B_i (Feasible(a) \Leftrightarrow p), \text{ e } p \text{ are the FPs of action } a.$$

Whenever an agent considers that a given event has just occurred, she is necessarily committed to believe that its effects and persistent FPs hold (from the agent's viewpoint):

Axiom A5: $\models B_i (Done(a) \Rightarrow p)$, where p is an effect or a persistent FP (such as certain mental preconditions for CAs) of action a .²

The kernel of the process which enables an agent to derive a plan is specified by axioms A2 and A3. Thus, the planning process is merely a regular consequence of the rational-behaviour axioms, hence completely deductive: no extra planning algorithm is needed to implement the CA models.

3.4 Mental-Attitude Transfer & Cooperative Behaviour

In the part of the theory that we have outlined so far, nothing constrains an agent to adopt the beliefs of another agent, or her intentions in order to be cooperative towards her (such as merely answering a question for her). For this to be possible, we need axioms to allow the transfer of mental attitudes.

A form of *belief transfer* appears when an agent i thinks that an agent j wants to communicate to her some mental attitude $\phi(j)$. In this case, and this is the *first preliminary principle for cooperative behaviour*, agent i comes to believe what she thinks that j wants to communicate to her, unless she believes the contrary:

$$\text{Axiom A7: } \models B_i (I_j B_i \phi(j) \wedge \neg B_i \neg \phi(j) \Rightarrow B_i \phi(j)).$$

The *second preliminary principle for cooperative behaviour* states that an agent i will adopt some intention

¹ This property extends and formalises the *principle of rationality* proposed by Newell [1982] p. 102: *if an agent has knowledge that one of its actions will lead to one of its goals, then the agent will select that action.*

² Concerning communicative acts, the effect in question is the *intentional effect*, i.e., the intention an agent attempts to make public in performing the act. For example, the intentional effect of agent i asserting p towards agent j is $I_i B_j | B_j p$.

(or goal) of an agent j , if she does not have the opposite intention:

Axiom A8: $\models B_i(I_j\phi \wedge \neg I_i\neg\phi \Rightarrow I_iB_j\phi)$.

The theory provides several other axioms and derived properties for cooperative behaviour.³ The following axiom, which formalised the harmony with other agents, will be used later:

Axiom A9: $\models \Gamma \Rightarrow \neg I_i\neg B_j\phi$ where T characterises the "right" circumstances, which, in some cases (depending on the instantiation of ϕ), can be merely reduced to *true*.

3.5 Implementation

To implement this kind of theory, two approaches can be adopted: either automate the inference process formalised by the theory, or use the theory as the formal specification of an effective system. We have adopted the first approach since it satisfies, in a more direct way, adequacy, genericity, and maintainability criteria. The inference engine [Bretier, 1995; Bretier and Sadek, 1996], based on a first-order modal logic, namely *quantified-KD45* (with respect to the belief modality), involves other modalities representing uncertainty, intention and action. Inference is automated using an extended resolution method, where formulae are represented in their syntactical form and where the instantiation of axiom schemata uses sub-formulae unification.

4 Specialised Components

When applied to natural dialogue, the ARTIM1S system includes two isolated subsystems to deal with natural language input and output. Currently, algorithmic bottlenecks (especially the size of the search space for the input side) prevent the direct logical specification of these processes in our forward-chaining theorem-proving framework. However, both are meant to be eventually integrated as logical theories in the rational unit.

4.1 Natural Language Input

The natural language interpretation subsystem features both syntactic and semantic robustness using island-driven parsing and semantic completion. Island-driven parsing simply means finding small syntactic structures in the text, with as few long-range dependencies as possible. Example:

Input sentence:

I'd like the weather forecast for the Cotes-d'Armor area

Recognized:

LdJike X weather forecast X X cotes d armor area

Concepts:

I_u weather .forecast cotes_d_armor

The result is a set of *mentioned concepts*, or a list of possible alternatives when overlapping phrases yield nondeterminism. Each of these hypotheses is then fed

³ For example, axiom $\models I_iB_j\phi \Rightarrow B_i\phi \vee I_iB_i\phi$ constrains an agent to behave sincerely; axiom $\models B_i(\phi \wedge B_j\neg\phi) \Rightarrow I_iB_j\phi$ leads an agent to try to correct what she thinks to be erroneous beliefs of another agent (by corrective answers).

into the semantic completion process, which builds a well-formed logical formula out of the mentioned items. The fundamental hypothesis here is that of *semantic connectedness*: the user never mentions two concepts without (at least implicitly) linking them by some relation(s) or other concept(s).

Example: (mentioned) *weather_forecast cotes_d_armor*
 (default) *phones(x)*
 (inferred) *phone(z,x) servers(z)*
weather_forecast(y) theme(z,y)
domain(y,cotes_d_armor)

What we get is a minimal contextually consistent first-order formula containing the mentioned objects. We then obtain a full-fledged communicative act (hopefully) conveying the user's request whose logical form is suited for direct use by the rational unit:

$\langle u, Inform(s, I_u Bref_u(\lambda x phones(x) \wedge$
 $(\exists y) weather_forecast(y) \wedge (\exists z) server(z) \wedge$
 $domain(y, cotes_d_armor) \wedge$
 $phone(z, x) \wedge theme(z, y))) \rangle$

Informally, this formula means that the user (u) informs the system (s) that she wants (I_u) to know ($Bref_u$) the (λx) phone number of a weather forecast server for Cotes d'Armor.

4.2 Natural Language Output

The natural language generation subsystem [Panaget, 1996] contains a *linguistic act planner* and a *linguistic realiser* (see also [Appelt, 1985]). The former is concerned with how the intentions of the system (i.e., communicative acts) are communicated to the user. The latter is concerned with realising the acts specified by the planner as a well-formed utterance.

Linguistic acts serve as an abstract representation of utterances or parts of utterance. Two families are distinguished: *surface speech acts* and *referring acts*. Surface speech acts are used to deal with the fact that: (1) different communicative acts can be realised by the same utterance and (2) a single utterance can verbalise a complex sequence of communicative acts. Linguistic acts are defined in terms of their effect and their preconditions of feasibility. The effect of a surface speech act is the verbalisation of communicative acts whereas the effect of a referring act is the reference to objects of the world. Preconditions are used to express that: (1) the accomplishment of a linguistic act must be relevant to the context in which the utterance will be used and (2) there exists a well-formed part-of-utterance. An interesting point is that the second type of precondition establishes explicitly the communication channel from the planner to the realiser.

The linguistic realiser satisfies the second type of precondition by constructing the most relevant parts-of-utterance and utterances according to the context and to its linguistic knowledge. During the construction, when the realiser requires a noun phrase, a pronoun or a proper name, it asks the planner to select a referring act. This mechanism establishes the second direction of the communication channel.

For example, if the system wants to inform the user that there is a "provider" relation between the Cotes-d'Armor weather forecast server and Meteo-France, the generator can produce, according to the context, either a simple declarative sentence with a proper name and a noun phrase, a positive answer with a pronoun and a proper name or an elliptical sentence:

"Meteo-France is the provider of the Cotes-d'Armor weather forecast server."

"Yes, it is provided by Meteo-France."

"Meteo-France."

5 Reasoning Example

Let us focus now on an example of CA planning, effectively produced by ARTIMIS. Suppose that the system, s , has just observed the CA, $a_1(x)$, built in section 4.1. Let us use the following abbreviations:

$$\begin{aligned}
 p(x) &\equiv (\exists y) \text{weather_forecast}(y) \wedge (\exists z) \text{server}(z) \\
 &\quad \wedge \text{domain}(y, \text{cotes.d.armor}) \wedge \text{phone}(z, x) \\
 &\quad \wedge \text{theme}(z, y) \\
 q(x) &\equiv (\exists x) B_u \text{phones}(x) \wedge B_u p(x) \\
 \neg q(x) &\equiv (\forall x) \neg B_u \text{phones}(x) \vee \neg B_u p(x) \\
 r(x) &\equiv (\exists x) B_s \text{phones}(x) \wedge B_s p(x) \\
 \neg r(x) &\equiv (\forall x) \neg B_s \text{phones}(x) \vee \neg B_s p(x) \\
 a_1(x) &\equiv \langle u, \text{Inform}(s, I_u q(x)) \rangle \\
 a_2(x) &\equiv \langle s, \text{Informref}(u, \text{phones}(x) \wedge p(x)) \rangle \\
 c &\equiv 36\ 68\ 02\ 22
 \end{aligned}$$

The inference process performs the following steps. The first column displays the clause numbers as generated by our program, the second displays the clauses (or resolvents), and the third displays the numbers of the clauses and axioms used to produce the resolvents:

1.	$B_s \text{Done}(a_1(x))$	<i>input act</i>
6.	$B_s I_u q(x)$	1; (A5)
16.	$B_s \neg q(x)$	6; (A1)
19.	$I_s q(x) \vee I_s \neg q(x)$	6; (A8)
24.	$I_s q(x)$	19; (A9)
28.	$I_s \text{Done}(a_2(x))$	24; (A2)
37.	$\neg I_s \text{Done}(a_2(x)) \vee I_s (\neg B_s \phi(x) \wedge r(x)) \vee$ $B_s \text{Feasible}(a_2(x))$	(A3)
50.	$B_s \text{nt}(c)$	<i>KB</i>
51.	$B_s p(c)$	<i>KB</i>
53.	$\neg B_s p(c) \vee B_s \text{Feasible}(a_2(c)) \vee B_s q(c)$	50; (A4)
64.	$\neg B_s p(c) \vee B_s \text{Feasible}(a_2(c))$	16, 53
71.	$B_s \text{Feasible}(a_2(c))$	51, 64

Once the KB is saturated, agent s (i.e., the system) performs those acts that it intends and believes are feasible, namely, in the example, $a_2(c)$:

$$\begin{aligned}
 &\langle s, \text{Informref}(u, \text{uc phones}(c) \wedge \\
 &\quad (\exists y) \text{weather_forecast}(y) \wedge (\exists z) \text{server}(z) \wedge \\
 &\quad \text{domain}(y, \text{cotes d armor}) \wedge \\
 &\quad \text{phone}(z, c) \wedge \text{theme}(z, y)) \rangle
 \end{aligned}$$

Then the NL generation subsystem verbalises this act by the following utterance:

s: The phone number of a weather-forecast server for the Cotes-d'Armor area is 36 68 02 22.

6 Comparison to related work

6.1 Formal semantics frameworks

Several formalisations of intention, based on a possible-worlds semantic model, have been proposed in the literature, but they differ in their interpretation of what

a possible world represents. In [Cohen and Levesque, 1990a], and also in [Konolige and Pollack, 1993], a world is a time-line representing a sequence of events, temporally extended infinitely into the past and the future. In [Rao and Georgeff, 1991], a world is a branching-time structure with a single past and multiple futures. In both, a temporal index is used to identify a particular point in the course of events. In our framework, a possible world is a single point in the course of events, and the possible worlds are related to each other by event-dependent accessibility relations. The resulting structure is a branching-time structure with a single past and multiple futures. Even though these three possible-worlds models seem to be nearly equivalent, they determine the *intuitive interpretation* of the formalisation.

Concerning the definition of mental attitudes, our model is similar to Cohen & Levesque's [1990a]. However, there are some significant differences between the two models, particularly in the definition of intention. For example, in C&L's model, an agent can have an intention to bring about a proposition without necessarily having this proposition as an achievement goal. So, the primary interest of an agent is not the end result of her intention but that she will achieve this result by herself. In our definition of intention, the primary interest of an agent is the end result of her intention (the definition explicitly imposes on the intended property to be a persistent goal). Moreover the achievement of the proposition in question may result from a multi-agent sequence of events, that is, it should be possible that the actions of other agents be a part of the agent's plan.

6.2 Theories of rational agency

Significant developments in rational agency have been elaborated in [Cohen and Levesque, 1990a; Rao and Georgeff, 1991; Konolige and Pollack, 1993; Shoham, 1993; Traum, 1996; Wooldridge, 1996]. The methodological frameworks proposed are of great interest for analysing the concept of intentional action. However, except for [Cohen and Levesque, 1990a], none of them provides an explicit formalisation of action, or of communication and cooperation principles. Moreover, while C&L's account of communicative act models is intended to be part of a theory of rational (inter)action, it lacks needed precision on when and how the act models can be used by an agent when carrying out a planning process.

In our case, we show that rationality principles can and need to be clearly specified within a unified logical framework of a rational action theory, in order for communicative acts to be planned as regular actions. Moreover, we show that (interagent communication) planning can be a fully deductive process, which naturally derives from the mere application of rationality principles to (communicative) act models, without any external devices.

6.3 Implementations

Several agent-oriented tools have been implemented. The AGENTO language [Shoham, 1993] allows the pro-

gramming of agents in terms of commitment rules. In [Lesperance *et al*, 1996], GOLOG, an agent programming language based on situation calculus is introduced. Other researchers directly implement their logical theories with theorem provers, thus giving the implemented systems a precise semantics. Concurrent METATEM [Wooldridge, 1996] is a multi-agent language based on the execution of a temporal belief logic by a tableaux-based decision procedure. Rao & Georgeff [Rao and Georgeff, 1995] also present tableaux-based provers for BDI logics.

However, except for Rao & Georgeff, none of these analyses the interaction of the modal operators. Our system is specifically designed to implement a first-order modal logic which includes axioms formalising the interaction between modal operators for belief, intention, and action. Moreover, our system is the only one which does not need external devices to implement planning or communicative/cooperative behaviour. These features are naturally obtained by formalising mental attitudes, actions, and rationality principles, in the same homogeneous logical framework.

7 Conclusion

The ARTIMIS system runs on a single Sun UltraSPARC 1 workstation. An ISDN interface handles phone calls, getting the standard **8kHz μ -law** speech input. Speech processing is realised by a HMM-based speaker-independent, continuous-speech recogniser, using a bigram language model [Dupont, 1993; Juvet *et al*, 1991] and a PsOLA-based TTS system [Bigorgne *et al*, 1993].

ARTIMIS displays advanced human-computer cooperative spoken dialogue abilities. This dialogue technology, based on the concept of communicating rational agent, is expected to be experimented, in the short or middle term, in context of real services. The results of an evaluation of the first version of the system can be found in [Sadek *et al*, 1996].

References

- [Appelt, 1985] D. E. Appelt. *Planning English Sentences*. Cambridge University Press, 1985.
- [Bigorgne *et al*, 1993] D. Bigorgne, O. Boeffard, B. Cherbonnel, F. Emerard, D. Larreur, J-L. Le Saint-Milon, I. Metayer, C. Sorin, and S. White. PSOLA text-to-speech system. In *ICASSP*, 1993.
- [Bretier and Sadek, 1996] 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 [Miiller *et al*, 1996].
- [Bretier, 1995] P. Bretier. *La communication orale cooperative : contribution d la modelisation logiquet et a la mise en oeuvre d'un agent rationnel dialoguant*. These de Doctorat Informatique, Universite de Paris XIII, France, 1995.
- [Cohen and Levesque, 1990a] P. R. Cohen and H. J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42(2-3):213-262, 1990.
- [Cohen and Levesque, 1990b] P. R. Cohen and H. J. Levesque. Rational interaction as the basis for communication. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in communication*. MIT Press, 1990.
- [Dupont, 1993] P. Dupont. Dynamic use of syntactical knowledge in continuous speech recognition. In *EUROSPEECH*, 1993.
- [Juvet *et al*, 1991] D. Juvet, K. Bartkova, and J. Monne. On the modelisation of allophones in an HMM-based speech recognition system. In *EUROSPEECH'91*, 1991.
- [Konolige and Pollack, 1993] K. Konolige and M. E. Pollack. A representationalist theory of intention. In *IJCAI*, 1993.
- [Lesperance *et al*, 1996] Y. Lesperance, H.J. Levesque, F. Lin, D. Marcu, R. Reiter, and R.B. Scherl. Foundations of a logical approach to agent programming. In *Intelligent Agents II*, LNAI volume 1037, 1996.
- [Miiller *et al*, 1996] J. P. Miiller, M. J. Wooldridge, and N. R. Jennings, editors. *Intelligent Agents III— Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL)*, LNAI, 1996.
- [Newell, 1982] A. Newell. The knowledge level. *Artificial Intelligence*, 18:87-128, 1982.
- [Panaget, 1996] F. Panaget. *D'un systeme generique de generation d'annonces en contexte de dialogue oral a la formalization logique des capacites linguistiques d'un agent rationnel dialoguant*. These de Doctorat Informatique, Universite de Rennes 1, France, Mars 1996.
- [Rao and Georgeff, 1991] A. S. Rao and M. P. Georgeff. Asymmetry thesis and side-effect problems in linear-time and branching-time intention logics. In *IJCAI*, 1991.
- [Rao and Georgeff, 1995] A. S. Rao and M. P. Georgeff. formal models and decision procedures for multi-agent systems. Technical Note 61, AAI1, Sydney, Australia, 1995.
- [Sadek *et al*, 1996] M. D. Sadek, A. Ferrieux, A. Cozannet, P. Bretier, F. Panaget, and J. Simonin. Effective human-computer cooperative spoken dialogue: The ags demonstrator. In *ICSLP*, 1996.
- [Sadek, 1991a] M. D. Sadek. *Attitudes mentales et interaction rationnelle : vers une theorie formelle de la communication*. These de Doctorat Informatique, Universite de Rennes I, France, Juin 1991.
- [Sadek, 1991b] M. D. Sadek. Dialogue acts are rational plans. In *Proceedings of the ESCA/ETRW Workshop on the structure of multimodal dialogue*, pages 1-29, Maratea, Italie, September 1991.
- [Sadek, 1992] M. D. Sadek. A study in the logic of intention. In *Proc. of the Conf. on Principles of Knowledge Representation and Reasoning*, 1992.
- [Sadek, 1994] M. D. Sadek. Communication theory = rationality principles + communicative act models. In *AAAI-94 Workshop on Planning for Interagent Communication, 1994*.
- [Shoham, 1993] Y. Shoham. Agent-oriented programming. *Artificial Intelligence*, 60(1):51-92, 1993.
- [Traum, 1996] D. R. Traum. A reactive-deliberative model of dialogue agency. In [Miiller *et al*, 1996].
- [Wooldridge, 1996] M. J. Wooldridge. A knowledge-theoretic semantics for concurrent metatem. In [Miiller *et al*, 1996].