Modeling Discourse Acts in Computer-Assisted Collaborative Decision Making

Afzal Ballim and Nikos Karacapilidis
MEDIA Research Group, LITH-DI, Swiss Federal Institute of Technology
IN Ecublens, 1015 Lausanne, Switzerland
{Afzal.Ballim,Nikos.Karacapilidis}@epfl.ch

Abstract

This paper presents a model of discourse acts that agents use to communicate their attitudes to each other, or affect the attitudes of others, in a multi-agent decision making process. Such processes usually raise a lot of intricate debates and negotiations among participants, where conflicts of interest are inevitable and support for achieving consensus and compromise is required. Furthermore, they are often performed in the presence of ill-structured information, brought up by parties with different backgrounds and interests. Focusing on the argumentation process itself, and in the context of an already implemented system, the paper also comments on the associated machinery needed in a computer-assisted decision making environment.

1 Introduction

Traditional decision making techniques, coming from areas such as mathematical economics, operations research, game theory and statistics, build on a probabilistic view of uncertainty, where possible actions are evaluated through their expected utility. The use of such crisp values has been extensively criticized; the specification of the complete sets of probabilities and utilities required renders such approaches impractical for the majority of decision making tasks that involve common sense knowledge and reasoning [Tan and Pearl, 1994]. On the other hand, Artificial Intelligence (AI) approaches basically attempt to reduce the burden of numerical information required, while pay much attention to the automation of the process.

Furthermore, traditional approaches are built on the assumption of a predefined set of alternatives and criteria, and provide methods to quantify and aggregate subjective opinions (consider, for instance, the Analytic Hierarchy Process [Saaty, 1980]). Everyday practices, however, make obvious that there is a lot of room for debate here. We view multi-agent decision making as a collaborative process, where agents have to follow a series of *communicative actions* in order to establish a common belief on the dimensions of the problem. Such dimensions may concern the choice criteria, the existing or desired alternatives, or the objective function, to mention some. Issues of knowledge elicitation and representation are inherent in these environments and an appropriate machinery is needed.

More specifically, approaches to collaborative decision making (CDM) may be divided into two large classes. In the first one, a set of alternatives is determined *a priori* and the task is one of deciding between them. In the second class, an *ideal case* is decided upon first, and a subsequent task is to find a real case that best approximates the ideal. In both approaches, however, there is a number of common elements:

- an overall task goal is specified;
- a set of alternatives is selected;
- a collection of choice criteria must be settled upon by the participants;
- a decision function must be composed which combines criteria to decide between alternatives.

The overall task goal is generally not a subject of debate, although it may be ill-defined and the decision process may involve sub-processes to clarify the goal. These sub-processes may be also considered as CDM processes themselves. In fact, each element of the decision process may itself be the subject of a sub-decision process. Decision making can, therefore, be recursive.

Proc. of the 2nd Int. Conf. on Practical Aspects of Knowledge Management (PAKM98)
Basel, Switzerland, 29-30 Oct. 1998, (U. Reimer ed.)

http://sunsite.informatik.rwth-aachen.de/Publications/CEUR_WS/

The set of alternatives can be a predetermined, closed set (no further alternatives can later be considered), a predetermined open set (leeway is given to allow integration of new alternatives), or a postdetermined set (in the case of finding a match to an ideal case). Interesting conclusions concerning the implicit goals, *a priori* positions, and biases of the participants in the process may often be inferred from the manner in which they present alternatives. It is often the case that participants have applied unspecified choice criteria before proposing alternatives (eliminating what they consider to be useless alternatives). In cases where the participants are of unequal stature in the CDM process (as, for instance, when a mixture of middle and upper management are involved), this can have a profound implicit effect on the collaborative aspects of the decision process; modeling of the hierarchical relations between participants may be necessary to understand what may appear to be illogical or contradictory decisions.

The choice criteria are the basis of any decision process. They provide the metrics upon which alternatives are compared, and accepted or rejected. As the foundations of the CDM process, they may be the subject of much debate. The inclusion of particular criteria may cause one to consider alternatives that would otherwise not figure in the process, while the exclusion of certain criteria may automatically eliminate certain alternatives that would *prima facie* be included. They can, therefore, be a preliminary battleground for power struggles between factions involved in the process.

The decision function, however, is where most argumentation is centred, since it is here that the relative value of choice criteria is established and applied to select between the alternatives. The argumentation used is often authoritative or based on voting.

2 Background and Motivation

Research on various aspects of CDM has been receiving growing interest in the AI and CSCW community in recent years (see, for instance, [De Michelis and Grasso, 1994], [Di Eugenio et al., 1997]). Computer tools to aid the CDM process vary from simple classical tools (e-mail mailing lists, dedicated newsgroups and news servers) and web-based discussion forums, to more dedicated systems that address the needs of a user to interpret and reason about knowledge during a discourse. For instance, *QuestMap* [Conklin, 1996] (based on the *gIBIS* hypertext groupware tool [Conklin and Begeman, 1988]) captures the key issues and ideas during meetings and creates shared understanding in a knowledge team. All the messages, documents, and reference material for a project can be placed on the system's ''whiteboard'', and the relationships between them can be graphically displayed. Users end up with a ''map'' that shows the history of an on-line conversation that led to key decisions and plans. Among other systems that address interesting knowledge management and representation issues, we only mention here *Euclid* [Smolensky et al., 1987], *Janus* [Fischer et al., 1989] and *Belvedere* [Suthers et al., 1995] (for a comparative analysis, see [Karacapilidis and Papadias, 1998a]).

Although this category of systems provides a cognitive argumentation environment that stimulates discussion among participants, it lacks decision making capabilities. On the contrary, the HERMES system [Karacapilidis and Papadias, 1998a, 1998b] focuses on aiding decision makers reach a decision, not only by efficiently structuring the discussion, but also by providing reasoning mechanisms for it. Following Conklin's terminology [Conklin, 1992], HERMES not only ''captures the informal organizational memory'' embodied in such environments, but also helps the users during the decision making process itself by integrating features based on concepts from well-established areas such as Knowledge Representation, Decision Theory, Non-Monotonic Reasoning, Constraint Satisfaction, and Truth Maintenance. Note that HERMES is intended to act as an *assistant* and *advisor*, by facilitating communication and recommending solutions, but leaving the final enforcement of decisions and actions to the agents.

The argumentation framework of HERMES is a variant of the informal IBIS model of argumentation [Rittel and Webber, 1973] (as is the case for *QuestMap* and *gIBIS*). HERMES supports as argumentation elements *issues, alternatives, positions,* and *preferences*. Figure 1 illustrates an instance of the system's *Discussion Forum* that concerns a real discussion about the planning of cyclepaths in the city of Bonn. The agents involved in the discussion, namely the representatives from the Cyclists Union and the related City Hall departments, bring up the necessary argumentation in order to express their interests and perspectives. As shown, HERMES maps a multi-agent decision making process to a discussion graph with a hierarchical structure.

While this latter system provides significant automatisation of the decision making process, it is still relatively low-level as far as the *discourse acts* that can be represented are concerned; further refinement of those already integrated is needed. Moreover, there are at least two areas in which further automatisation would be desirable:

- an *argument assistant* that can follow and advise on the *details* of an argument, and not just its form;
- an *argument support tool* that can peruse a document collection. The goal of such a tool would be to find excerpts from the collection that can be used as referal material by the agent to support or refute a given argument.

A prerequisite for such tools is the ability for the computer to understand (at least partially) the dialogue in a decision-related argument between people, and the discourse structure used in presenting supportive material in a document. This requires a computational model of the discourse acts which are used in these cases. Although there has been work in AI on dialogue and discourse in collaboration and negotiation, that work is not sufficient for modeling dialogues in the collaborative decision making process.

For instance, Sidner [1994] presents a model of collaborative negotiation based on the idea of establishing *mutual beliefs*, that is, things that we hold in common. This model

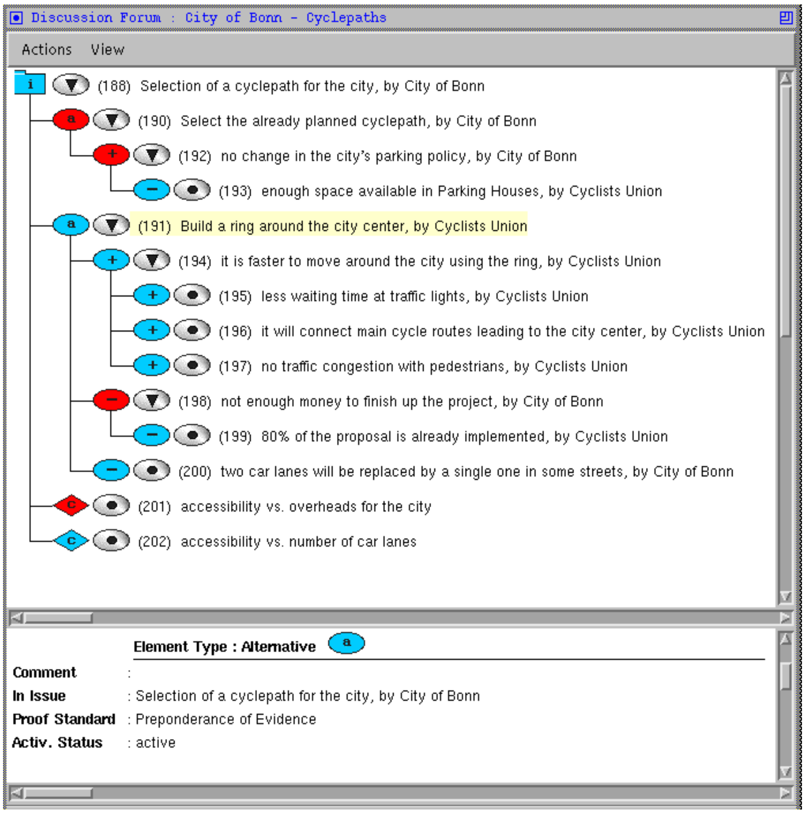

- i (188) Selection of a cyclepath for the city, by City of Bonn
  - a (190) Select the already planned cyclepath, by City of Bonn
    - \+ (192) no change in the city's parking policy, by City of Bonn
      - − (193) enough space available in Parking Houses, by Cyclists Union
  - a (191) Build a ring around the city center, by Cyclists Union
    - \+ (194) it is faster to move around the city using the ring, by Cyclists Union
      - \+ (195) less waiting time at traffic lights, by Cyclists Union
      - \+ (196) it will connect main cycle routes leading to the city center, by Cyclists Union
      - \+ (197) no traffic congestion with pedestrians, by Cyclists Union
    - − (198) not enough money to finish up the project, by City of Bonn
      - − (199) 80% of the proposal is already implemented, by Cyclists Union
    - − (200) two car lanes will be replaced by a single one in some streets, by City of Bonn
  - c (201) accessibility vs. overheads for the city
  - c (202) accessibility vs. number of car lanes

Element Type : Alternative a

Comment :
In Issue : Selection of a cyclepath for the city, by City of Bonn
Proof Standard : Preponderance of Evidence
Activ. Status : active

Figure 1: *An instance of the Hermes Discussion Forum*

rests upon the absence of deception, and appears fragile in the presence of mutual misunderstanding. The work of Cohen and Levesque [1990] and of Smith and Cohen [1996] is very similar to Sidner's work, but relies in addition on the primitive notion of *joint goals*. Based on Searle's idea [Searle, 1969] that *requesting* something means that one is attempting to get an agent to perform an action, they define all *illocutionary acts* in terms of agent's mental states (illocutionary is an act performed as the result of a speaker making an utterance; the effect is called a *perlocutionary act*).

Core and Allen [1997] have introduced a scheme for annotating communication acts in dialogue; this scheme ignores the formation of opinions by hearers about speakers, and gives a single coding for each utterance. It is generally the case, however, that utterances can and should be considered to perform multiple functions.

We argue that an understanding of the implications in the CDM process requires a model of the *mental attitudes* of the agents involved (their beliefs, desires, intentions, goals, etc.) as they pertain to the task at hand (see also [Jennings, 1993], [Grosz and Sidner, 1990]). Further, it requires a model of the particular form of *discourse acts* that agents use to communicate their attitudes to each other, and affect the attitudes of others. In addition, it requires a model of the actions that relate to the argument process itself (what we shall refer to as the *internal acts*). The goal of this paper is to propose an acceptable model for these cases. The various approaches described above all take this *rational agent* approach, as do we. The model presented here, however, is oriented explicitly towards CDM.

3 A Model of Discourse Acts in Collaborative Decision Making

Collaborative decision making generally involves much debate between the participants, with negotiations and the formation of power blocks and voting cells (subgroups of individuals who act together with a common strategy or vote as an individual). Conflicts of opinion and interest are inevitable, and people use differing strategies to resolve these conflicts (see [Merin, 1997]). Generally, it is presumed that all participants are being cooperative and honest. Agents act optimally with respect to the information at hand, and do not consider how others might interpret their actions (they do not work to "please the boss").

3.1 Objects and Relations, States and Actions

The first question to consider is the primitive components of the model. This includes objects of discourse such as the alternatives among which we must choose, the criteria for evaluating these alternatives, and the method for applying those criteria (decision function). The participants in the decision process are also "objects" since we must model the interaction between them. We must also allow comparisons between criteria and alternatives along numerical scales, so it must be possible to represent notions such as *"cost is more important than speed"* or *"the total cost must be less than 5,000"* in a car purchase discourse. The basic language that we propose to use is a sorted first-order predicate calculus with extensions for numerical comparison. We will use terms such "*agt1*" or "*agtN*" to denote individual participants, and "*group*" to denote them all.

To understand the dynamics and evolution of a decision, we must have notions of time and states, and of actions that change states. *States* represent a coherent situation. They may be already past or current. In order to consider alternatives to the existing state of affairs, we introduce the notion of a hypothesis, which is an artificial state of the world that may not exist, and may not be possible. It provides a framework for hypothetical reasoning by an agent. If agent *agt1* hypothesizes the state `S1`, then we will write this as:

(HYP agt1 `S1`*)*

Actual states, future states, hypothetical states and temporal intervals are reified objects in our proposal, and all propositions must have them; however, for simplicity, they are only included in the examples here when pertinent.

Actions are defined in terms of the changes that they evoke in states. The simplest acts in our system are the actions of telling something to an agent, and its complement of hearing what an agent is saying:

(TELL agt1 agt2 `P`*), (HEAR agt2 agt1* `P`*)*

In order to describe an action, we will talk about the *effects* it will have. These effects are the difference between a start state and the new state that the action produces. We represent that an act `A` in state `S0` results in state `S` as:

(EFFECT `A S0 S`*)*

States are typified by a non-exhaustive set of assertions, so an action such as turning a switch might be defined as:

(EFFECT Turn(switch) [is(switch,off)] [is(switch,on)]) ∨
(EFFECT Turn(switch) [is(switch,on)] [is(switch,off)])

3.2 Primitive Mental Attitudes

When agents interact with each other in a dialogue, there is a number of mental attitudes that dictate the form of the interaction as well as the long-term behaviour of the agents throughout the dialogue. The attitudes can concern states of the world or actions that effect the world.

3.2.1 Beliefs

The most basic of these is the attitude of belief. In the context of AI and in this article, belief is used as a global term that covers the notion of knowledge as well as the notion of information in which we do not have complete confidence. Presuming that this information can be expressed as a proposition in, say, a first-order predicate calculus, then belief can be considered to be a modal operator relating an agent to a belief (see [Ballim, 1992] and [Ballim and Wilks, 1991] for a more detailed description). So, for example, the belief held by the agent *agt1* that ISDN lines are fast might be expressed as:

$$(BEL\ agt1\ \forall x[\mathtt{ISDN}(x) \rightarrow \mathtt{fast}(x)])$$

3.2.2 Mutual Belief

It is often convenient to refer to some belief as being commonly held between a group of agents. Where `P` is such a belief, this will be expressed as:

(MB agt1 ... agtN `P`*)*

3.2.3 Desires

Agents act in a world that is not always the same as the world in which they would like to be. Desires are used to express the wishes of the agent about the state of the world. They may also be expressed as modal operators over propositions. Thus, the desire that MBONE software be used (for some purpose), might be expressed in the following way:

$$(DES\ agt1\ \exists x[\mathtt{use}(x) \wedge \mathtt{MBONE}(x)])$$

3.2.4 Goals

Although agents might desire a particular state of affairs to exist, they are not obliged to act upon those desires. For example, we might all desire a care-free life, but our responsibilities would stop us acting upon this desire by (for example) dropping out of society. State of affairs that we wish to exist, and, further, towards which we actively aspire are referred to as *goals.* If an agent had as a goal that ISDN lines be used in the context of a decision process, then this might be written as:

$$(GOAL\ agt1\ \exists x[\mathtt{use}(x) \wedge \mathtt{ISDN}(x)])$$

3.2.5 Intentions

The essential difference between desires and goals is that there is the notion of having an intention to act to achieve the a goal, while that is not necessarily the case for desires. While it might be useful to thus define *intention* with respect to states of world ("I will work to achieve a certain state of the world") and then define a goal as desiring a state and intending that state, in our framework it is more reasonable that intention is related to actions, so that we intend to perform an action. So, if an agent intends to call another agent, we might represent this as:

(INT agt1 `call(agt1,agt2)`*)*

The predicate `call` above should be interpreted as an action. The relation between our goals and the methods of achieving these goals is a *plan.* A simple plan is a sequence of actions, but more complex plans can be defined which depend on contingencies and options.

3.2.6 Support

For our domain, it is necessary to consider the relationship that agents believe hold between propositions, so we introduce the notion of *support.* Agent *agt1* believes that `P1` supports `P2` will be written as:

$$(\mathit{BEL}\ \mathit{agt1}\ (\mathit{SUPP}\ \mathtt{P1}\ \mathtt{P2}))$$

3.2.7 Refute

The opposition of support is that a belief refutes another. Agent *agt1* believes that `P1` disproves `P2` will be written as:

$$(\mathit{BEL}\ \mathit{agt1}\ (\mathit{REF}\ \mathtt{P1}\ \mathtt{P2}))$$

Such definitions are common in AI. While incomplete, they are sufficient to allow us now to describe the model of dialogue and internal acts in collaborative decision making.

3.3 Dialogue Acts

The interactions between agents in a collaborative decision making process are codified using *dialogue acts*. These can be used to interpret the discourse occurring between agents during the decision process, as well as to make inferences about their attitudes and to predict their likely future behaviour. The model that we present here is not a complete model of human discourse, but an interesting subset which allows an analysis of the group decision process. In the following, we define necessary conditions for a dialogue act to have occurred. Multiple dialogue acts can be associated with a single utterance.

3.3.1 Consider

The start of any group decision process involves proposing a topic for consideration. To simplify matters, we will presume that the topic is always an action to be performed (e.g., buying a car, choosing between different video conferencing systems).

$$(\mathit{CON}\ \mathit{agt1}\ \mathit{group}\ \mathtt{A}) \models (\mathit{BEL}\ \mathit{agt1}\ (\mathit{EFFECT}\ \mathtt{A}\ \mathtt{S0}\ \mathtt{S})) \wedge (\mathit{HYP}\ \mathit{agt1}\ \mathtt{S}) \wedge (\mathit{TELL}\ \mathit{agt1}\ \mathit{group}\ \mathtt{A}) \wedge (\mathit{GOAL}\ \mathit{agt1}\ (\mathit{HYP}\ \mathit{group}\ \mathtt{S}))$$

3.3.2 Inform

Agent *agt1* says something to somebody that he believes to be true.

$$(\mathit{INFORM}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{P}) \models (\mathit{BEL}\ \mathit{agt1}\ \mathtt{P}) \wedge (\mathit{TELL}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{P})$$

3.3.3 Request

A general dialogue act is that of requesting something. We split up requesting into three sub-cases of agent *agt1* making a demand, and later we will explain more complex requesting in terms of these.

- additional information is requested concerning a belief `P`:

$$(\mathit{ASKINFO}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{P})$$

- an opinion is demanded with respect to a belief `P`:

$$(\mathit{ASKOP}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{P})$$

- an act `A` is requested concerning a belief `P`; for instance, *agt1* wants *agt2* to agree or disagree with a belief `P`:

$$(\mathit{ASKACT}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{A}\ \mathtt{P})$$

In the last case, when *agt1* does not specify the act he wants *agt2* to perform (that is, requests all possible acts), we write: $(\mathit{ASKACT}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{P})$

3.3.4 Compare

It is often necessary to compare beliefs or to compare beliefs with values. This is necessary, for instance, for data consistency purposes. In the dialogue, an agent may declare that they believe one criteria to be more important than another, or that some criteria must have values in a particular range. As such cases are very frequent, we will introduce a *compare* dialogue act to treat them.

$$(\mathit{COMP}\ \mathit{agt1}\ \mathit{agt2}\ \Re\ \mathtt{B1}\ \mathtt{B2}) \models (\mathit{BEL}\ \mathit{agt1}\ (\Re\ \mathtt{B1}\ \mathtt{B2})) \wedge (\mathit{TELL}\ \mathit{agt1}\ \mathit{agt2}\ (\Re\ \mathtt{B1}\ \mathtt{B2}))$$

where $\Re$ is a comparison relation between beliefs (e.g., “is more important than”). For values, we have

$$(\mathit{COMP}\ \mathit{agt1}\ \mathit{agt2}\ \Re\ \mathtt{C}\ \mathtt{V}) \models (\mathit{BEL}\ \mathit{agt1}\ (\Re\ \mathtt{C}\ \mathtt{V})) \wedge (\mathit{TELL}\ \mathit{agt1}\ \mathit{agt2}\ (\Re\ \mathtt{C}\ \mathtt{V}))$$

where $\Re$ is a comparison of criteria (`C`) with value (`V`).

3.3.5 Propose

Agent *agt1* proposes that belief `P` be accepted by another agent.

$$(\mathit{PROP}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{P}) \models (\mathit{DES}\ \mathit{agt1}\ (\mathit{BEL}\ \mathit{agt2}\ \mathtt{P})) \wedge (\mathit{INFORM}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{P})$$

3.3.6 Agree

Agent *agt1* express his agreement with the belief `P` that was proposed by *agt2*.

$$(\mathit{AGR}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{P}) \models (\mathit{PROP}\ \mathit{agt2}\ \mathit{agt1}\ \mathtt{P}) \wedge (\mathit{INFORM}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{P})$$

3.3.7 Acknowledge

Agent *agt1* accepts that belief `P` may be valid, but that he requires further proof before agreeing to accept it.

$$(\mathit{ACK}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{P}) \models (\mathit{PROP}\ \mathit{agt2}\ \mathit{agt1}\ \mathtt{P}) \wedge (\mathit{ASKINFO}\ \mathit{agt1}\ \mathit{agt2}\ P)$$

3.3.8 Disagree

Agent *agt1* expresses that he disagrees with a belief `P`.

$$(\mathit{DISAGR}\ \mathit{agt1}\ \mathit{agt2}\ \mathtt{P}) \models (\mathit{PROP}\ \mathit{agt2}\ \mathit{agt1}\ \mathtt{P}) \wedge (\mathit{INFORM}\ \mathit{agt1}\ \mathit{agt2}\ \neg\mathtt{P})$$

3.3.9 Corroborate

Agent *agt1* agrees with a belief `P1` and wishes to give it further credence by proposing `P2`.

$$(\mathit{CORR}\ agt1\ group\ \mathtt{P1}\ \ \mathtt{P2}) \models (\mathit{AGR}\ agt1\ group\ \mathtt{P1}) \wedge (\mathit{INFORM}\ agt1\ group\ (\mathit{BEL}\ agt1\ (\mathit{SUPP}\ \mathtt{P2}\ \ \mathtt{P1})))$$

3.3.10 Challenge

Agent *agt1* disagrees with belief `P1`, and wishes to further undermine it by proposing `P2`.

$$(\mathit{CHALL}\ agt1\ group\ \mathtt{P1}\ \ \mathtt{P2}) \models (\mathit{DISAGR}\ agt1\ group\ \mathtt{P1}) \wedge (\mathit{INFORM}\ agt1\ group\ (\mathit{BEL}\ agt1\ (\mathit{REF}\ \mathtt{P2}\ \ \mathtt{P1})))$$

3.3.11 Discard

Proposal by agent *agt1* to another agent *agt2* to stop considering a belief `P` that has been proposed.

$$(\mathit{DISCARD}\ agt1\ agt2\ \mathtt{P}) \models \neg(\mathit{BEL}\ agt1\ \mathtt{P}) \wedge (\mathit{INFORM}\ agt1\ agt2\ (\mathit{GOAL}\ agt1\ \neg(\mathit{HYP}\ agt2\ [\mathtt{P}])))$$

3.3.12 Clarify

Agent *agt1* proposes (to *agt2*) a replacement belief (`P1`) for consideration as a clarification and replacement for belief `P2` which was previously proposed (`P2` may have been proposed by any agent).

$$(\mathit{CLA}\ agt1\ agt2\ \mathtt{P1}\ \ \mathtt{P2}) \models (\mathit{PROP}\ agt1\ agt2\ \mathtt{P1}) \wedge (\mathit{INFORM}\ agt1\ agt2\ (\mathit{GOAL}\ agt1\ \neg(\mathit{HYP}\ agt2\ [\mathtt{P1}]))) \wedge (\mathit{PROP}\ agt1\ agt2\ \mathtt{P2})$$

3.3.13 Counter-offer

Agent *agt1* is not ready to accept the belief `P1` proposed by *agt2* and, without explicitly disagreeing to `P1`, he proposes a belief `P2`.

$$(\mathit{COFF}\ agt1\ agt2\ \mathtt{P1}\ \ \mathtt{P2}) \models (\mathit{PROP}\ agt2\ agt1\ \mathtt{P1}) \wedge (\mathit{INFORM}\ agt1\ agt2\ (\mathit{GOAL}\ agt1\ \neg(\mathit{HYP}\ agt2\ [\mathtt{P1}]))) \wedge (\mathit{PROP}\ agt1\ agt2\ \mathtt{P2})$$

4 An Example Dialogue

Consider the discourse transcription illustrated in Figure 2, which took place among executives of a european research consortium and concerns the purchase of an appropriate videoconferencing tool for their needs. Communication was performed through electronic mail; each individual's message was delivered to all executives through an electronic mailing list. The example comprises portions of the overall discourse and serves the illustration of the computational model suggested in the previous section.

The segmentation of the e-mail messages exchanged, as appears in Figure 2, was made in a way to clearly present the different types of acts involved in such discourses. Messages have been very slightly edited to improve their comprehensibility, that is we almost quote here their most representative parts.

The links appearing in the figure simply denote a *follow up* relation and help the reader to comprehend the evolution of the discourse. In other words, the father node of such a link precedes (has been asserted before) a child node.

For instance, the A16 discourse portion follows up what has been declared in A15, which in turn follows up E4 and so on. Each discourse portion is accompanied by a header that denotes the agent who brought it up and the time order of the portion concerning this very agent (e.g., C3 should be interpreted as the third such portion expressed by agent C). Actually, according to the communication practice followed by the agents, sets of these portions happened to be included in the same e-mail message (e.g., B1, B2 and B3 composed a reply message of agent B to agent A).

Agent A initiates the discourse

A1: "I am writing to you about the choice of a video conferencing tool to be purchased for the purposes of our consortium",

remindseveryone of some conditions holding

A2: "I remind you that the budget is around 30 KECUs and that the choice has to be made by February 15, 1998),

provides a list of questions

A3: "Here is a list of questions I have in order to identify the needs in a more accurate way"

and potential products

A4: "Here is a list of potential products I have already identified", respectively).

and requests some acts from the other agents by a certain date

A4: "Please send your answers and comments before January 25, 1998".

Agents B, C, D and E then discuss the various choice criteria and alternatives initially suggested. Exhorted by the request of agent A for additional products

A9: "Could you please send me additional suggestions of products that either you have heard about or you have experimented yourself."

E brings up in the sequel of the discourse a new alternative

E4: "We have also some recent experience with videoconferencing tools from the MERCI suite",

as agent A also does later

A16: "We could also consider the use of a SMART Board".

In the rest of the section, we express these discourse portions according to the suggested model. When necessary, the expressions are further explained. Note that, whenever

Figure 2: A discourse example.

an act is addressed to the whole group, we omit the addressee parameter (for instance, *(CON agt1* `A`*)* is the shortened writing for *(CON agt1 group* `A`*)*). Also for brevity reasons, an overlined state will denote an *achieved* one (depending on the context of the discourse, it may reads better if we use instead the terms *solved*, *decided*, *done* and so on). Finally, an underlined belief will denote a *selected* one.

A1: *(CON agtA* `purchase`*)*

A2: *(COMP agtA* equal `budget 30Kecu`*)*

(INFORM agtA [(GOAL agtA $\overline{\texttt{purchase}}$*)* before `15.02.98`*])*

Having initiated the discourse, agent A asserts a constraint concerning the budget available and the time period for the solution of the problem (when temporal intervals are pertinent, we mention them explicitly like in the last expression above).

A3:(PROP agtA `question_list`*)* $\wedge$
(BEL agtA (EFFECT (ASKACT agtA `question_list`*)* $\neg\overline{\texttt{accurate}}$ $\overline{\texttt{accurate}}$*))*

Agent A proposes a list of questions related to the problem (concerning choice criteria) with the (implicit) belief that acting upon them will render the problem accurately defined.

A4:[(ASKINFO agtA `question_list`*)* before `25.01.98`*]*

[(ASKACT agtA `question_list`*)* before `25.01.98`*]*

A5:(ASKOP agtA `technology`*)*

(BEL agtA (EFFECT $\underline{\texttt{ISDN}}$ $\neg\overline{\texttt{technology}}$ $\overline{\texttt{technology}}$*))*

(BEL agtA (EFFECT $\underline{\texttt{MBONE}}\ \neg\overline{\texttt{technology}}\ \overline{\texttt{technology}}$*)) (ASKINFO agtA* technology*)*

Agent A introduces the *"technology"* issue asking the opinion of the group. At the same time, A lists two alternative selections while leaves room for suggestions of additional ones (the last act above expresses the *"..."* part of the portion).

A6: (ASKOP agtA connections*)*

A7: (ASKOP agtA platform*)*

(BEL agtA (EFFECT $\underline{\texttt{PC}}\ \neg\overline{\texttt{platform}}\ \overline{\texttt{platform}}$*))*

B1: ((AGR agtB ISDN*)* ∧ *(HYP agtB* $\overline{\overline{\texttt{ISDN}}\supset\overline{\texttt{quality}}}$*)* ∧

(HYP agtB $\underline{\texttt{quality}}$*)* ∧ *(COMP agtB* better ISDN any_technology*))*

((CORR agtB MBONE nice_experiences*)* ∧ *(HYP agtB* $\overline{\overline{\texttt{bandwidth}}\supset\overline{\texttt{MBONE}}}$*)* ∧

(HYP agtB $\overline{\texttt{bandwidth}}$*))*

Agent B agrees with the *"ISDN"* alternative and compares it to the other ones in case that *"quality"* is an issue. B then also agrees with the *"MBONE"* alternative and gives it further credence by the *"nice experiences"* B has got (note that the third component of a hypothesis should always be a state).

C1: (AGR agtC MBONE*)*

D1: (CORR agtD ISDN safe*)* ∧ *(CORR agtD* ISDN multipoint_conferences*)*

B2: (PROP agtB 2*)* ∧ *(BEL agtB (EFFECT* $\underline{2}\ \neg\overline{\texttt{connections}}\ \overline{\texttt{connections}}$*))*

(PROP agtB 3*)* ∧ *(BEL agtB (EFFECT* $\underline{3}\ \neg\overline{\texttt{connections}}\ \overline{\texttt{connections}}$*))*

(COMP agtB equal 2 3*)*

B proposes that 2 or 3 people will simultaneously be connected (both alternatives are equally weighed and believed as a solution of the issue of *"connections"*).

C2: (HYP agtC $\underline{\texttt{MBONE}}$*)* ∧ *(DISCARD agtC* connections*)*

B3: (PROP agtB Solaris*)* ∧ *(BEL agtB (EFFECT* $\underline{\texttt{Solaris}}\ \neg\overline{\texttt{platform}}\ \overline{\texttt{platform}}$*))* ∧

(COMP agtB better Solaris any_platform*)*

(PROP agtB IRIX*)* ∧ *(BEL agtB (EFFECT* $\underline{\texttt{IRIX}}\ \neg\overline{\texttt{platform}}\ \overline{\texttt{platform}}$*))* ∧

(COMP agtB better IRIX any_platform*)*

(COMP agtB equal Solaris IRIX*)*

Similarly to the B2 discourse portion, agent B proposes *"Solaris"* or *"IRIX"* to the *"platform"* issue. The two proposals are equally weighed.

C3: (CLA agtC op_sys platform*)*

(PROP agtC Unix/Linux*)* ∧ *(BEL agtC (EFFECT* $\underline{\texttt{Unix/Linux}}\ \neg\overline{\texttt{op\_system}}\ \overline{\texttt{op\_system}}$*))*

(PROP agtC Win/MacOS*)* ∧ *(BEL agtC (EFFECT* $\underline{\texttt{Win/MacOS}}\ \neg\overline{\texttt{op\_system}}\ \overline{\texttt{op\_system}}$*))*

(COMP agtC better Unix/Linux Win/MacOS*)*

(COMP agtC equal Unix Linux*)*

(COMP agtC equal Win MacOS*)*

Agent C proposes to replace the *"platform"* issue with the *"operating system"* one and gives alternative solutions and comparisons among them.

A8: (PROP agtA product_list*)* ∧

(BEL agtA ($\exists$ product_list(x) *(EFFECT* $\underline{\texttt{product\_list(x)}}\ \neg\overline{\texttt{product}}\ \overline{\texttt{product}}$*)))*

Agent A proposes a list of potential products with the belief that (the selection of) one of those will solve the *"product"* issue. Obviously, in the above expression, we can use directly $\overline{\texttt{purchase}}$ instead of $\overline{\texttt{product}}$.

A9: (ASKACT agtA (PROP group addit_products*))*

(ASKOP agtA product_list*)*

A10: (PROP agtA DialView*)* ∧*(BEL agtA (EFFECT* $\underline{\texttt{DialView}}\ \neg\overline{\texttt{product}}\ \overline{\texttt{product}}$*))*

Portions A11-A14 are similar to A10.

E1: (CORR agtE LiveLAN good_reviews*)*

D2: (AGR agtD LiveLAN*)* ∧ *(COMP agtD* better LiveLAN any_product*)*

(AGR agtD TriniCom*)* ∧ *(COMP agtD* better TriniCom any_product*)*

(COMP agtD equal LiveLAN TriniCom*)*

E2: (CORR agtE RSVP updated_version*)*

(ASKOP agtE RSVP*)*

Agent E corroborates the *"RSVP"* product and (implicitly) asks for the group's opinion.

E3: (CHALL agtE NTE tried*)*

E4: (PROP agtE MERCI_tools*)* ∧ *(BEL agtE (EFFECT* $\underline{\texttt{MERCI\_tools}}\ \neg\overline{\texttt{product}}\ \overline{\texttt{product}}$*))*

(CORR agtE MERCI_tools experience1*)*

(HYP agtE $\overline{\texttt{more\_details}}$*)*

Having proposed the *"MERCI tools"* and further supported them with his experience, agent E makes the hypothesis that *"more details"* can be given (by him).

A15: (ASKACT agtA agtE (INFORM agtE agtA `more_details`*))*

The above act of agent A is addressed only to agent E (all other acts in the example are addressed to the whole group).

A16: (PRO agtA SMART*)* ∧ *(BEL agtA (EFFECT* $\underline{\texttt{SMART}}$ $\neg\overline{\texttt{product}}$ $\overline{\texttt{product}}$*))*

(CORR agtA SMART `experience2`*)*

(CORR agtA SMART `price(software)`*)*

(CHALL agtA SMART `price(projector)`*)*

(HYP agtA `exists(projector)`*)* ∧ *(DISCARD agtA* `price(projector)`*)*

Agent A proposes the *"SMART tool"*, speaks in favour (*"price of software"*) and against (*"price of projector"*) it, while finally proposes to discard the *"price of projector"* issue under the hypothesis that *"the projector is already available"*.

5 Discussion

Work on CDM should exploit recent advances in AI and information technology, in order to:

- remove the communication impediments among spatially dispersed agents;
- efficiently elicitate and represent the domain knowledge;
- develop efficient mechanisms to structure and consistently maintain the decision analysis, and
- automate the decision making process itself.

An analysis such as that above is one step in this process. It further requires an interpretation of the above in terms of the evaluation of the decision function. This interpretation introduces the notion of *internal acts* which reconcile the dialogue acts with the argumentation process in an automatic fashion. Our ultimate goal is the development of an advanced framework for such a system. Various internal acts have been already integrated in HERMES, thus enhancing collaborative decision making by supporting argumentative discourse among decision makers [Karacapilidis and Papadias, 1998a]. The system is implemented in Java applets and runs on the Web, thus provides relatively inexpensive access to a broad public. It can be used for distributed, asynchronous or synchronous collaboration, allowing decision makers to surpass the requirements of being in the same place and working at the same time. As mentioned in Section 2, HERMES allows participants to structure their arguments, and then it checks the form of the argument in applying a decision function. We discuss below the major subtasks involved in this process.

5.1 Consistency Checking

This is related to the path consistency problem between the ordering of the asserted beliefs. Each belief has a *consistency label,* which can be *consistent* or *inconsistent*. Every time a *comparison* act is asserted (similar to the *COMP* in the proposed model), HERMES checks if both beliefs of this newcoming act exist in another, previously inserted, one. If yes, the new act is considered either *redundant*, if it also has the same preference relation, or *conflicting*, otherwise. A redundant act is ignored, while a conflicting one is grouped together with the previously inserted (now conflicting, after the assertion of the new act) ones in an - automatically created by the system - issue for consideration, the rationale being to gather together conflicting beliefs and stimulate further argumentation on them until only one becomes active. If both beliefs of such a new act do not exist in a previously inserted constraint, its consistency is checked against previous active and consistent acts of the same type that refer to the same element. Such processes involve *path consistency* algorithms [Mackworth and Freuder, 1985]. In HERMES, the issue is handled by a polynomial algorithm ($O(N^3)$ where N is the number of the associated beliefs).

5.2 Contradictions

Consider the simple case where an agent believes `P` while another `¬P`. A CDM system should be capable to detect such cases from the dialogue acts suggested above. The notion of hypothesis is crucial here for reasoning with such contradictions.

5.3 Truth Maintenance

HERMES performs a labelling function in order to characterize beliefs as *active* or *inactive*. Active beliefs may have been supported or corroborated by an agent while inactive ones would have been refuted. Such truth maintenance forms the basis for coherency checking.

5.4 Coherence Checking

The dialogue acts are presented as individual objects that are related to utterances. Over the course of a dialogue for decision making, it is necessary to establish that the coherency of the dialogue is being respected. This involves checking that, for instance, if an agent proposes a belief `P`, and believes that `Q` refutes `P`, then the agent cannot reasonably propose `Q` without first discarding `P`.

5.5 User Queries

The system should efficiently respond to user queries of the type: *"give me all beliefs* `P` *which have been proposed by the agtX supporting a belief* `Y`*"*. Because the system cannot be expected to make complete judgement decisions itself, the aim of the system must be to evaluate the arguments being presented by participants, to extract as much information concerning their beliefs and argumentation as possible, and to present such information to the user upon demand (cf. [Labrou and Finin, 1994], [Lochbaum et al., 1990]).

5.6 Conclusion

It is our belief that the integration of the dialogue acts presented in this paper with systems such as HERMES is the most reasonable way to cross the divide between, on the one hand, the simple systems for remote CDM that provide no more than a communication channel and archiving facility between participants (e.g., newsgroups and webforums) and, on the other hand, fully-blown reasoning systems that attempt to automatically solve the decision making problem. This latter approach appears to us to be infeasible at the moment.

However, the mix of human and machine reasoning that forms the basis of our proposition we believe to be feasible. In this method, the dialogue acts presented in this paper are derived from the dialogue. Human intervention would be necessary to determine certain specifics of the propositions being asserted, and relations holding between them. Accepting that the process of interpreting dialogue acts from utterances is fallible and non-exclusive, the objective of the system is to establish the most likely and coherent set of dialogue acts associated with each intervention by a participant, and to maintain the set of hypotheses being supported or considered by participants (see also [Buckingham Shum, 1997]). The participants can use this to better understand the arguments of their fellow decision makers, and advance more pointed argumentation of their own.

5.6.1 References

Ballim, A. (1992). ViewFinder: A Framework for Representing, Ascribing and Maintaining nested Beliefs of Interacting Agents. *Ph.D. Thesis No. 2560*, Université de Genève, Switzerland.

Ballim, A. and Wilks (1991). *Artificial Believers.* Lawrence Erlbaum Associates, Hillsdale, NJ.

Buckingham Shum, S. (1997). Balancing Formality with Informality: User-Centred Requirements for Knowledge Management Technologies. In *Proceedings of the AAAI Spring 1997 Symposium on Artificial Intelligence in Knowledge Management (AIKM'97)*, AAAI Press.

Cohen, P.R., and Levesque, H.J. (1990). Intention is choice with commitment. *Artificial Intelligence,* 42(3), 213-261.

Conklin, J. (1996). Designing Organizational Memory: Preserving Intellectual Assets in a Knowledge Economy. GDSS Working Paper. Available at: `http://www.gdss.com/DOM.htm`

Conklin, J. (1992). Capturing Organizational Memory. In D. Coleman (ed.) *Proceedings of GroupWare'92*, Morgan Kaufmann, San Mateo, CA, 133-137.

Conklin, J. and Begeman, M.L. (1988). gIBIS: A Hypertext Tool for Exploratory Policy Discussion. *ACM Transactions on Office Information Systems*, 6(4), 303-331.

Core, M.G., and Allen, J.F. (1997). Coding Dialogs with the DAMSL Annotation Scheme. In *Proceedings of AAAI-97 Fall Symposium on Communicative Action in Humans and Machines,* AAAI Press.

De Michelis, G., and Grasso, M. A. (1994). Situating conversations within the language/action perspective: the Milan Conversation Model. In *Proceedings of CSCW-94,* ACM Press, 89-100.

Di Eugenio, B., Jordan, P.W., Thomason, R.H., and Moore, J.D. (1997). Reconstructed Intentions in Collaborative Problem Solving Dialogues. In *Proceedings of AAAI-97 Fall Symposium on Communicative Action in Humans and Machines,* AAAI Press.

Fischer, G., McCall, R., and Morch, A. (1989). JANUS: Integrating Hypertext with a Knowledge-based Design Environment. In Proceedings of the *Hypertext'89 Conference*, New York: ACM, 105-117.

Grosz, B.J., and Sidner, C. L. (1990). Shared plans in discourse. In Cohen, P., Morgan, J., and Pollack, M. (eds.) *Intentions in Communication,* MIT Press.

Jennings, N. (1993). Specification and Implementation of a Belief-Desire-Joint-Intention Architecture for Collaborati-ve Problem Solving. *Int. Journal of Intelligent and Cooperative Information Systems*, 2(3), 289-318.

Karacapilidis, N.I., and Papadias, D. (1998a) HERMES: Supporting Argumentative Discourse in Multi-Agent Decision Making. In *Proceedings of AAAI-98,* AAAI/MIT Press, 827-832.

Karacapilidis, N.I., and Papadias, D. (1998b). A Computational Approach for Argumentative Discourse in Multi-Agent Decision Making Environments. *AI Communications Journal* 11(1), 1998.

Labrou, Y., and Finin, T. (1994). A semantics approach for KQML: a general purpose communication language for software agents. In *Proceedings of CIKM-94,* ACM Press.

Lochbaum, K., Grosz, B. J., and Sidner, C. L. (1990). Models of plans to support communication: An initial report. In *Proceedings of AAAI-90,* AAAI Press, 485-490.

Mackworth, A., and Freuder, E. (1985). The Complexity of some Polynomial Network Consistency Algorithms for Constraint Satisfaction Problems. *Artificial Intelligence* 25, 65-74, 1985.

Merin, A. (1997). Communicative Action as Bargaining: Utility, Relevance, Elementary Social Acts. In *Proceedings of AAAI-97 Fall Symposium on Communicative Action in Humans and Machines,* AAAI Press.

Rittel, H., and Webber, M. (1973). Dilemmas in a General

Theory of Planning. *Policy Sciences*, 155-169.

Saaty, T. (1980). *The analytic hierarchy process.* Mc-Graw Hill, New York.

Searle, J.R. (1969). *Speech acts: An essay in the philosophy of language.* Cambridge University Press.

Sidner, C.L. (1994). An Artificial Discourse Language for Collaborative Negotiation. In *Proceedings of AAAI-94,* AAAI/MIT Press, 814-819.

Smith, I.A., and Cohen, P.R. (1996). Toward a semantics for an agent communications language based on speech-acts. In *Proceedings of AAAI-96,* AAAI/MIT Press, 24-31.

Smolensky, P., Fox, B., King, R., and Lewis, C. (1987). Computer-aided Reasoned Discourse, or How to Argue with a Computer. In Guindon, R. (ed.) *Cognitive Science and its Applications for Human-Computer Interaction*, Hillsdale, NJ: Erlbaum, 109-162.

Suthers, D., Weiner, A., Connelly, J., and Paolucci, M. (1995). Belvedere: Engaging Students in Critical Discussion of Science and Public Policy Issues.In *Proceedings. of the 7th World Conference on AI in Education*, 266-273.

Tan, S.W, and Pearl, J. (1994). Qualitative Decision Theory. In *Proc..of AAAI-94*, AAAI/MIT Press, 928–933.