

## Computer-Mediated Collaborative Decision Making: Theoretical and Implementation Issues

Nikos Karacapilidis  
*Computer Science Department*  
*Swiss Federal Institute of*  
*Technology (EPFL)*  
*Lausanne, Switzerland*  
*karacapi@di.epfl.ch*

Dimitris Papadias  
*Computer Science Department*  
*Hong Kong University of Science*  
*and Technology*  
*Clearwater Bay, Hong Kong*  
*dimitris@cs.ust.hk*

Costas Pappis  
*Dept. of Industrial Management*  
*University of Piraeus*  
*80 Karaoli and Dimitriou Str.*  
*Piraeus 18534, Greece*  
*pappis@unipi.gr*

### Abstract

*Development of systems for computer-mediated Collaborative Decision Making (CDM) attracts increasing interest from various research areas. Although some approaches provide a cognitive argumentation environment and methods to structure the related discussions, they lack consensus seeking and decision making capabilities. In addition, the majority of them is not based on a well defined set of users' communicative actions. Reporting on HERMES, a fully implemented web-based system that enhances group decision making by providing an argumentation framework to the agents involved, this paper discusses theoretical and implementation aspects of an advanced group decision support system. Argumentation in our framework is performed through a set of discourse acts, especially defined for the CDM context following an artificial intelligence perspective. The proposed system provides the appropriate machinery for automating processes such as discussion structure, consistency checking and reasoning for decision making. Moreover, it includes further assistance modules with information retrieval, natural language processing and argument building features.*

analysis and systematically directing the pattern, timing, or content of the discussion. Major issues arising during the development of such a system include effective work organization in order to improve coordination, and use of communication technology to make decision making more efficient.

CDM usually raises a lot of intricate debates and negotiations among participants. Conflicts of interest are inevitable and support for achieving consensus and compromise is required. Each decision maker may adopt and, consequently, suggest his/her own strategy that fulfills some goals at a certain level. Opinions may differ about the relevance or value of a proposition when deciding an issue. Decision makers may have arguments supporting or against alternative solutions. In addition, they have to confront the existence of insufficient and too much information simultaneously. In other words, for some parts of the problem, relevant information which would be useful for making a decision is missing, whereas for others, the time needed for the retrieval and comprehension of the existing volume of information is prohibitive. Furthermore, participants need appropriate means to assert their preferences, which often are expressed in qualitative terms. Due to the above, proper definition of all acts decision makers perform during such processes and provision of procedures for automation of data processing, especially in data intensive situations, are of high importance.

Decision makers are not necessarily proficient in computer science and information technology; they need appropriate tools in order to easily follow the processes involved. Such tools should stimulate their participation giving them an active role. This parallels the vision of the DSS community pioneers, that is, by supporting and not replacing human judgement, the system comes in second and the users first. CDM admittedly falls in the

### 1. Introduction

Collaborative Decision Making Systems can be defined as interactive computer-based systems which facilitate the solution of ill-structured problems by a set of decision makers, working together as a team [19]. Their main objective is to augment the effectiveness of decision groups through the interactive sharing of information between group members and the computer. This can be achieved by (i) removing communication impediments, and (ii) providing techniques for structuring the decision

category of “wicked” problems [23], a class of problems that can be addressed through discussion and collaboration among the agents involved. Consensus emerges through the process of collaboratively considering alternative understandings of the problem, competing interests, priorities and constraints. The application of more formal modeling and analysis tools is impossible before the problem can be articulated in a concise and agreed upon manner.

Computer-mediated CDM has been receiving growing interest in the last few years. Proliferation of Internet technologies incites development of such systems on the World-Wide Web, mainly due to its platform-independent communication framework and associated facilities for data representation, transmission and access. In particular, attention focuses on the implementation of argumentation support systems for different types of groups and application areas. Such systems address the needs of a user to interpret and reason about knowledge during a discourse. For instance, *QuestMap* [6] 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 *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. *QuestMap* was based on the *gIBIS* hypertext groupware tool [5] which aimed at capturing the rationale of a design process.

*Euclid* [28] is another system that provides a graphical representation language for generic argumentation. *JANUS* [11] is based on acts of critiquing existing knowledge in order to foster the understanding of a design process. *SEPIA* [29] is a knowledge-based authoring and idea processing tool for creating and revising hyperdocuments that views authoring as a design process. Finally, *Belvedere* [30] uses a rich graphical language to represent different logical and rhetorical relations within a debate, originally designed to support students engaged in critical discussion of science issues. Although this category of systems provides a cognitive argumentation environment that stimulates discussion among participants, it lacks decision making capabilities.

Numerous web-based conferencing systems have also been deployed, such as *AltaVista Forum Center*, *Open Meeting*, *NetForum* and *UK Web's Focus*, to mention some. They usually provide means for discussion structuring and user administration tools, while the more sophisticated ones allow for sharing of documents, on-line calendars, embedded e-mail and chat tools, etc. Discussion is structured via a variety of links, such as

simple *responses* or different *comment types* (e.g., *qualify*, *agree*, *example* in *Open Meeting*) to a previous message. However, the above systems merely provide threaded discussion forums, where messages are linked “passively”, which usually leads to an unsorted collection of vaguely associated comments. As pointed out by the developers of *Open Meeting*, there is a lack of consensus seeking abilities and decision making methods [14]. Furthermore, this category of systems is not based on a well defined set of users’ communicative actions.

A prerequisite for computer-mediated CDM 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 Artificial Intelligence (AI) on dialogue and discourse in collaboration and negotiation (see, for instance [8,9,13,22], that work is not sufficient for modeling dialogues in the CDM context. More specifically, it is rather general and not explicitly oriented towards real-life CDM environments.

Section 2 illustrates the model we suggest. In the sequel, Section 3 presents the modules that constitute our framework for computer-mediated CDM. The argumentation-based group decision support module and its associated machinery for aiding decision makers reach a decision are described in Section 4. Automation in our system includes mechanisms that not only structure the related discussion, but also provide reasoning and consistency mechanisms. Section 5 reports on further assistance tools with information retrieval, natural language processing and argument building features. Finally, Section 6 comments on interesting issues in CDM and concludes the paper.

## 2. Modeling the Discourse

Attempts to model the dialogue process, mainly coming from the AI discipline, generally presume 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. For instance, Sidner [26] presents a model of collaborative negotiation based on the idea of establishing *mutual beliefs*, that is, beliefs that agents hold in common. This model rests upon the non-existence of deception, and appears fragile in the presence of mutual misunderstanding. The work of Cohen and Levesque [4] and of Smith and Cohen [27] is very similar to Sidner’s work, but relies in addition on the primitive notion of *joint goals*. Based on Searle’s idea [25] that *requesting* something means that one is attempting to get an agent

to perform an action, they define a set of *illocutionary acts* (the act performed as the result of a speaker making an utterance; its effect is a *perlocutionary act*) in terms of agent's mental states. Core and Allen [7] introduce a scheme for annotating communication acts in dialogue, which 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.

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. Further, it requires a model for the particular form of *discourse acts* that agents use to communicate their knowledge and intentions, and affect the attitudes of others. In addition, it requires a model of the actions that relate to the argument process itself.

### 2.1. 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 agents must choose, the criteria for evaluating these alternatives, and methods for applying these criteria (i.e., the *decision* or *objective function*). Participants in a decision process are also "objects" since their interactions must be modelled. A complete model should also allow comparisons between criteria and alternatives along numerical scales, so it can represent notions such as "*cost is more important than speed*" or "*the total cost must be less than 5,000*", in the context of a car purchase discussion. The basic language that we propose is a sorted first-order predicate calculus with extensions for numerical comparison. We will use terms such as "*agt<sub>1</sub>*" or "*agt<sub>n</sub>*" to denote individual participants, and "*group*" to denote the set of all agents.

To capture the dynamics and evolution of a decision, the model must include 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 *agt<sub>1</sub>* hypothesizes state *SI*, then we write: (*HYP agt<sub>1</sub> SI*).

Actual, future and 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 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 agt<sub>1</sub> agt<sub>2</sub> P*) and (*HEAR agt<sub>2</sub> agt<sub>1</sub> P*). Finally, in order to describe an action, we consider the *effects* it will have. These effects are the difference between a start state and a new state produced by the action. We represent that an act *A* in state *SO* results in state *S* as: (*EFFEFFECT A SO S*).

### 2.2. Primitive Mental Attitudes

When agents interact with each other in a dialogue, there is a number of mental attitudes that dictate the form of their interaction and their long-term behavior throughout the dialogue. The attitudes can concern states of the world or actions that effect the world. 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 [2,3] for a more detailed description). So, for example, the belief held by an agent *agt<sub>1</sub>* that *ISDN* lines are fast might be expressed as: (*BEL agt<sub>1</sub>  $\forall x$  [ISDN(x)  $\rightarrow$  fast(x)]*). It is often convenient to refer to some belief as being commonly held between a group of agents, that is we talk about a *mutual belief*. If *P* is such a belief, this will be expressed as: (*MB agt<sub>1</sub> ... agt<sub>n</sub> P*)

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 agt<sub>1</sub>  $\exists x$  [use(x)  $\wedge$  MBONE(x)]*). Although agents might desire a particular state of affairs to exist, they are not obliged to act upon those desires. 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 is written as: (*GOAL agt<sub>1</sub>  $\exists x$  [use(x)  $\wedge$  ISDN(x)]*).

The essential difference between desires and goals is that there is the notion of having an intention to act to achieve 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 (the predicate *call* should be interpreted as an action), we might represent

this as:  $(INT\ agt1\ call(agt1,agt2))$ . The relation between 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.

For our domain, it is necessary to consider the relationships agents believe that hold between propositions; we introduce the notions of *support* and *refute*. If agent *agt1* believes that *P1* supports *P2*, we write  $(BEL\ agt1\ (SUPP\ P1\ P2))$ , while if *agt1* believes that *P1* disproves *P2*, we write  $(BEL\ agt1\ (REF\ P1\ P2))$ . Such definitions are common in AI (see, for instance, [4,7,15,26]). While incomplete, they are sufficient to allow us to describe the model of dialogue acts in collaborative decision making.

### 2.3. Dialogue Acts

The interactions between agents in a collaborative decision making process are codified using dialogue acts (note that multiple dialogue acts can be associated with a single utterance in a discussion). 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 behavior. The model of discourse acts we propose (including the

necessary conditions for them to occur) is illustrated in Table 1. It should be made clear that this is not a complete model of human discourse, but an interesting subset which allows an analysis of the collaborative decision making process. While a detailed description of the model does not fall in the scope of this paper, in the rest of this section we outline most of the corresponding acts.

The start of any group decision process involves *opening of an issue* or, according to the model's terminology, *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, deciding between different medical treatments for a patient). According to the corresponding act (namely, *consider* - see Table 1), *agt1* proposes (to all decision makers involved) an act *A* which produces a desired state *S*; *agt1* hypothesizes *S* and has as goal that the same will hold for the *group*.

Another basic dialogue act is that of *requesting* something. We split up requesting into the following three sub-cases of agent *agt1* making a demand: additional information is requested concerning a belief *P*, an opinion is demanded with respect to a belief *P*, and an act *A* is requested concerning a belief *P* (e.g. *agt1* wants *agt2* to agree or disagree with *P*).

Description	Dialogue Act	Condition
consider	$(CON\ agt1\ group\ A)$	$(BEL\ agt1\ (EFFECT\ A\ S0\ S)) \wedge (HYP\ agt1\ S) \wedge (TELL\ agt1\ group\ A) \wedge (GOAL\ agt1\ (HYP\ group\ S))$
inform	$(INFORM\ agt1\ agt2\ P)$	$(BEL\ agt1\ P) \wedge (TELL\ agt1\ agt2\ P)$
request information	$(ASKINFO\ agt1\ agt2\ P)$	-
request an opinion	$(ASKOP\ agt1\ agt2\ P)$	-
request an act	$(ASKACT\ agt1\ agt2\ A\ P)$	-
compare two beliefs	$(COMP\ agt1\ agt2\ \mathfrak{R}\ P1\ P2)$	$(BEL\ agt1\ (\mathfrak{R}\ P1\ P2)) \wedge (TELL\ agt1\ agt2\ (\mathfrak{R}\ P1\ P2))$
compare criteria and values	$(COMP\ agt1\ agt2\ \mathfrak{R}\ C\ V)$	$(BEL\ agt1\ (\mathfrak{R}\ C\ V)) \wedge (TELL\ agt1\ agt2\ (\mathfrak{R}\ C\ V))$
propose	$(PROP\ agt1\ agt2\ P)$	$(DES\ agt1\ (BEL\ agt2\ P)) \wedge (INFORM\ agt1\ agt2\ P)$
agree	$(AGR\ agt1\ agt2\ P)$	$(PROP\ agt2\ agt1\ P) \wedge (INFORM\ agt1\ agt2\ P)$
acknowledge	$(ACK\ agt1\ agt2\ P)$	$(PROP\ agt2\ agt1\ P) \wedge (ASKINFO\ agt1\ agt2\ P)$
disagree	$(DISAGR\ agt1\ agt2\ P)$	$(PROP\ agt2\ agt1\ P) \wedge (INFORM\ agt1\ agt2\ \neg P)$
corroborate	$(CORR\ agt1\ group\ P1\ P2)$	$(AGR\ agt1\ group\ P1) \wedge (INFORM\ agt1\ group\ (BEL\ agt1\ (SUPP\ P2\ P1)))$
challenge	$(CHALL\ agt1\ group\ P1\ P2)$	$(DISAGR\ agt1\ group\ P1) \wedge (INFORM\ agt1\ group\ (BEL\ agt1\ (REF\ P2\ P1)))$
discard	$(DISCARD\ agt1\ agt2\ P)$	$\neg (BEL\ agt1\ P) \wedge (INFORM\ agt1\ agt2\ (GOAL\ agt1\ \neg (HYP\ agt2\ S(P))))$
clarify	$(CLA\ agt1\ agt2\ P1\ P2)$	$(PROP\ agt1\ agt2\ P1) \wedge (INFORM\ agt1\ agt2\ (GOAL\ agt1\ \neg (HYP\ agt2\ S(P1)))) \wedge (PROP\ agt1\ agt2\ P2)$
counter-offer	$(COFF\ agt1\ agt2\ P1\ P2)$	$(PROP\ agt2\ agt1\ P1) \wedge (INFORM\ agt1\ agt2\ (GOAL\ agt1\ \neg (HYP\ agt2\ S(P1)))) \wedge (PROP\ agt1\ agt2\ P2)$

Table 1 Discourse acts in collaborative decision making

It is often necessary (e.g., for inconsistency detection) to weigh alternative beliefs or to compare beliefs with values. In the dialogue, an agent may declare that he/she believes one criterion to be more important than another, or that some criteria must have values in a particular range. As such cases are very frequent, we introduce a *compare* dialogue act to treat them. In such cases,  $\mathfrak{R}$  is a comparison relation between beliefs (e.g., of the type “*P1 is more important than P2*”) or between a criterion *C* and a value *V* (e.g., of the type “*C should be equal to V*”).

Using other acts, an agent *agt1* may: *propose* a belief to be accepted by another agent; express his/her *agreement* with a belief that was proposed by another agent; accept that a belief may be valid but he/she requires further proof before agreeing to accept it (*acknowledge* act); express that he/she *disagrees* with a belief; agree with a belief *P1* and wish to give it further credence by proposing *P2* (*corroborate* act); disagree with a belief *P1* and wishes to further undermine it by proposing *P2* (*challenge* act); propose to another agent *agt2* to stop considering a belief that has been proposed (*discard* act); propose (to *agt2*) a replacement belief *P1* for consideration as a clarification and replacement for belief *P2* which was previously proposed (*clarify* act); finally, *agt1* may declare that he/she is not ready to accept the belief *P1* proposed by *agt2* and, without explicitly disagreeing to *P1*, he/she proposes a belief *P2* (*counter-offer* act).

### 3. Building a CDM System

Traditional decision making techniques, coming from areas such as mathematical economics, operations research, game theory and statistics, 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 [24]). 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 appropriate machinery is needed. Furthermore, traditional approaches 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 [31]. On the other hand, AI approaches basically attempt to reduce the burden of numerical information required, while pay much attention to the automation of the process.

Using the model of discourse acts described in the previous section, our ultimate goal is the development of a complete system for supporting CDM in real-life applications. Information technology is essential for achieving proper communication and enabling information access to everybody. The system described in this paper is based on the World-Wide Web platform, and meets three major practical requirements: (i) it provides relatively inexpensive access; (ii) it has intuitive interfaces in order to be easily usable by inexperienced users, and (iii) it can be available on all prominent operating systems and hardware platforms. In order to use the system, one only needs a Web browser and Internet access. We focus on *distributed, asynchronous* collaboration, allowing agents to surpass the requirements of being in the same place and working at the same time. The system 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 system consists of four components (Figure 1). The *Argumentation-based Group Decision Support Module* enhances decision making by supporting argumentative discourse among decision makers (note that, depending on the type of the CDM process, the overall system may be “administrated” by a *discussion moderator* which can intervene when is needed). The module enables decision makers to propose and discuss alternative solutions to a problem by electronically sending various discourse acts to the system’s server (for instance, they can use the *consider* act to open an issue, the *corroborate* act to further support an alternative, etc.). Users are then able to access the structured protocol of a discussion, which is based on a formal argumentation theory. The component is able to elicitate and represent the domain knowledge; moreover, it efficiently structures and consistently maintains the decision analysis, and provides appropriate mechanisms for automating the CDM process itself. All discourse acts are recorded in the system’s relational database (mSQL has been used). Interactions with the database are performed through *Java Database Connectivity* (JDBC) drivers. A first version of this component, namely *HERMES*, which supports a subset of the discourse acts described in the previous section, is fully implemented in Java [16] and has already received positive evaluation feedback.

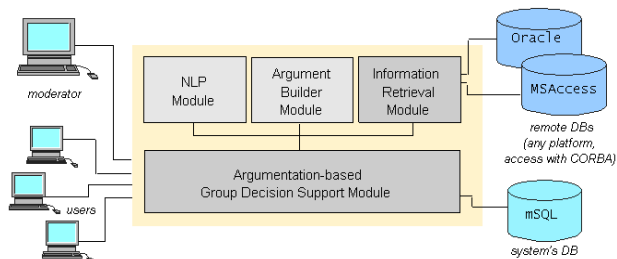


Figure 1 The System

The *Information Retrieval Module* aims at assisting decision makers retrieve and reuse information stored in remote databases. It is able to access databases on any platform (e.g., MSAccess on PCs, Oracle on Unix, etc.) through the CORBA communication protocol. The module is also implemented in Java. The *Natural Language Processing (NLP) Module* provides techniques for the analysis of the messages users assert during a discussion. The idea is to identify the intended discourse acts (according to the model discussed in Section 2) and, in the sequel, automatically extract and link the corresponding argumentation items. Finally, the *Argument Builder Module* concerns perusal of a discussion's patterns, the aim being to advice agents perform discourse acts that best interpret their interests and intentions (for instance, to support or refute a certain argument). Apparently, this module is highly associated with the Information Retrieval one. Having retrieved the required information, decision makers may construct more robust arguments to be then submitted in the corresponding CDM discussion. The NLP and Argument Builder modules are currently under implementation.

#### 4. Argumentative Discourse in CDM

This section discusses concepts and techniques involved in the Argumentation-based Group Decision Support module, in the context of the HERMES system. As said above, HERMES supports a subset of the discourse acts defined in Section 2. However, integration of all such acts is rather straightforward.

The argumentation framework of HERMES is an extension of the ZENO [12], which in turn has its roots to the informal IBIS model of argumentation [23]. HERMES supports as argumentation elements *issues*, *alternatives*, *positions*, and *constraints* representing *preference relations*. We consider here a real medical decision making example concerning the appropriate treatment to be followed for a patient case. Three medical doctors participate in the discussion, namely Dr. Brown, Dr. Clark and Dr. Wadder. Figure 2 illustrates two instances of the corresponding HERMES *Discussion Forum*. As shown, our approach maps a multi-agent decision making process to a discussion graph with a hierarchical structure.

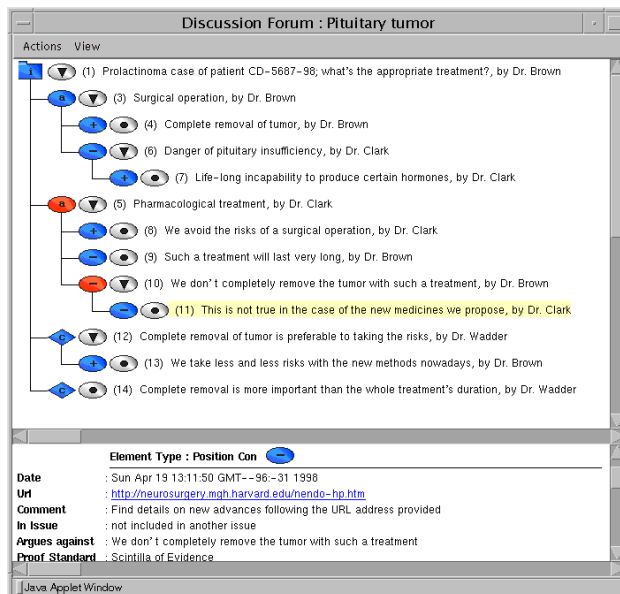


Figure 2 A Discussion Instance

Issues correspond to decisions to be made, or goals to be achieved (e.g., issue-1: "Prolactinoma case of patient CD-5687-98; what's the appropriate treatment?"). They are brought up by agents and are open to dispute. Issues consist of a set of alternatives that correspond to potential choices (e.g., alternative-3: "Surgical operation" and alternative-5: "Pharmacological treatment" both belong to issue-1, and have been proposed by Dr. Brown and Dr. Clark, respectively). Issues can be inside other issues in cases where some alternatives need to be grouped together. For instance, if two alternative pharmacological treatments were proposed, they could be grouped in an internal issue, say the *subissue* "which is the most appropriate pharmacological treatment?". According to the related argumentation, the best of these alternatives will then be compared to alternative-3.

Positions are asserted in order to support the selection of a specific course of action (alternative), or avert the agents' interest from it by expressing some objection. For instance, position-4: "Complete removal of tumor" has been asserted to support alternative-3, while position-6: "Danger of pituitary insufficiency" to express Dr. Clark's objection to it. Positions may also refer to some other position in order to provide additional information about it, e.g., position-11: "This is not true in the case of the new medicines we propose" (arguing against position-10), and position-7: "Life-long incapability to produce certain hormones" (arguing in favor of position-6). A position always refers to a single other position or alternative; supporting positions are shown with "+" in Figure 2, while counter-arguments with "-".

In decision making environments, one has usually to define priorities among actions and weigh different

criteria. Unfortunately, well defined utility and probability functions regarding properties or attributes of alternatives (used in traditional approaches), as well as complete ordering of these properties are usually absent. In HERMES, constraints provide a qualitative way to weigh reasons for and against the selection of a certain course of action. A constraint is a tuple of the form  $[position, preference\ relation, position]$ , where the preference relation can be *more (less) important than* or *of equal importance to*. Constraints may give various levels of importance to alternatives. Like the other argumentation elements, they are subject to discussion; therefore, they may be "linked" with positions supporting or challenging them (see, for instance, position-13). In Figure 2, constraint-12: "*Complete removal of tumor is preferable to taking the risks*" expresses the preference relation "position-4 is more important than position-8", while constraint-14: "*Complete removal is more important than the whole treatment's duration*" represents the relation "position-4 is more important than position-9".

Alternatives, positions and constraints have an *activation label* indicating their current status (it can be *active* or *inactive*). This label is calculated according to the argumentation underneath and the *type of evidence* specified for them. Activation in our system is a recursive procedure; a change of the activation label of an element (alternative, position or constraint) is propagated upwards. When an alternative is affected during the discussion, the issue it belongs to should be updated since a new choice may be made. In general, different elements of the argumentation, even in the same debate, do not necessarily need the same type of evidence. Quoting the well-used legal domain example, the arguments required to indict someone need not be as convincing as those needed to convict them [10]. Therefore, a generic argumentation system requires different proof standards (work on AI and Law uses the term "burdens of proof").

HERMES uses three different proof standards (we do not claim that the list is exhaustive; other standards, that match specific application needs, can be easily incorporated to the system): (i) *Scintilla of Evidence (SoE)*: according to this proof standard, a position is active, if at least one active position argues in favor of it; (ii) *Beyond Reasonable Doubt (BRD)*: according to BRD, a position is active if there are not any active positions that speak against it; (iii) *Preponderance of Evidence (PoE)*: according to this standard, a position is active when the active positions that support it outweigh those that speak against it; in the case of alternatives, PoE produces positive activation label when there are no alternatives with larger score in the same issue. It uses constraints to determine the activation level by

comparing the relative importance of supporting and counter arguments (for details on the associated *scoring mechanisms*, see [16]).

In the discussion instance of Figure 2, the proof standard is SoE for all positions and PoE for alternatives. Position-10 is inactive (the accompanying icons of inactive items are shown in red color) since position-11 is active and speaks against it. On the contrary, position-6 is active (the accompanying icons of active items appear in blue color) since there is at least one active position (position-7) that speaks in favor of it. Active positions are considered "accepted" due to discussion underneath (e.g., strong supporting arguments, no counter-arguments), while inactive positions are (temporarily) "rejected". Similarly, active alternatives correspond to "recommended" choices, i.e., choices that are the strongest among the alternatives in their issue. Note that both alternatives in Figure 2 are actually linked with a position in favor and a position against them (concerning alternative-5, the second position against it, namely position-10, is inactive). Constraint-12 renders a higher score to alternative-3, hence this is the solution indicated (at this state) by the system.

The activation label of constraints is decided by two factors: the discussion underneath (similarly to what happens with positions) and the activation label of their constituent positions (i.e., the positions compared through the constraint). In Figure 2, both constraints are active. If during the evolution of the discussion, a new position inactivates position-4, this will result in the inactivation of both constraints, since position-4 is one of their constituent positions.

Apart from an activation label, each constraint has a *consistency label* which can be *consistent* or *inconsistent*. Every time a constraint is inserted in the discussion graph, the system checks if both positions of the new constraint exist in another, previously inserted, constraint. If yes, the new constraint is considered either *redundant*, if it also has the same preference relation, or *conflicting*, otherwise. A redundant constraint is ignored, while a conflicting one is grouped together with the previously inserted constraint in an issue automatically created by the system, the rationale being to gather together conflicting constraints and stimulate further argumentation on them until only one becomes active. If both positions of the new constraint do not exist in a previously inserted constraint, its consistency is checked against previous active and consistent constraints belonging to the same issue using a path consistency algorithm [17].

The Web platform facilitates access to the current knowledge by making available all relevant data and documents. When highlighting an argumentation item in

the upper pane of the Discussion Forum window, all related information is given in the lower pane. By clicking on the Url entry, the associated HTML file appears in a new Web browser's window. In the example shown in Figure 2, the Url corresponds to a detailed presentation of the effects of the new medicine proposed (position-11 is highlighted). In such a way, participants may "attach" to their discourse elements useful multimedia information (text, images, sound, etc.).

### 5. Providing Further Assistance

This section reports on issues concerning the rest three modules of the CDM system presented in Section 3. While not of minor importance, we view their role as enhancing the Argumentation-based module; they augment its effectiveness providing decision makers with advanced capabilities. To our knowledge, there is no other CDM system that integrates all these features.

#### 5.1. The Natural Language Processing Module

Using only the Argumentation-based module, the user has to choose the location in the graph where he/she wants to add an argumentation element and the type of the discourse act. That is, the module takes into account only the *form* of an argument and not its *content* (this advocates the supervising of such acts by the discussion moderator). The goal of the NLP module is to automate the process of message insertion in the discussion graph using appropriate techniques. For example, in the discussion of Figure 2, assume that a user wants to assert an argument of the type: "I don't think we should perform a surgical operation for the patient since he is over-aged and this admittedly involves serious risks; I can give you extra evidence for that which I have collected from our lab's records". Such a message cannot be attached to any particular element; so it must be analysed to extract the element(s) that it refers to and the type(s) of the discourse act the user has in mind.

Unrestricted natural language understanding is certainly a long way from being solved; however, information extraction methods can work because they depend on strong *a priori* restrictions on the kinds of patterns they need to search for. In our case, the constraining assumption is that the user's sentence carries *themes*, a set of *argumentation items*, and a set of *attitudes*. The strength of this restriction makes the task feasible; concerning users of a CDM system, the restriction is not a burdensome one. In fact, the NLP module could relieve them of much of the tedious burden of manually linking each relation; moreover, it could eliminate misinterpretations and mistakes during the assertion of messages.

Ongoing implementation of the module involves integration of off-the-shelf *segmenters*, *taggers*, *probabilistic partial parsers*, and *broad-coverage grammars*, and development of additional core NLP technologies in the context of CDM (among others, *element extractor*, *phrasal similarity evaluator*, and *thesaural phrasal similarity evaluator*).

#### 5.2. The Information Retrieval Module

In order to reach understanding, negotiate and resolve conflicts through a CDM system, agents may want to consult various types of information to warrant their arguments towards the selection or rejection of a statement or action (cf. [20,21]). Such information may be stored either in external, remote databases with legacy data or in the system's proprietary one. In the CDM domain, case-based reasoning and learning techniques have been particularly useful due to their resemblance to the way people evaluate a potential future action by using past experience, the scarcity (or even absence) of explicitly stated rules, and the ill-structured definitions of the associated problems [1,18].

The Information Retrieval module allows agents to ask various types of queries regarding the discussion and related material. Each time agents want to add a position in the discussion graph, they are able to retrieve auxiliary information from a pre-specified set of databases. Figure 3 shows such a query to an external medical database (this act may correspond to Dr. Brown's wish to further support position-4 by bringing up information concerning previous cases). Note that this applet window is dynamically created, depending on the user's selection of database and table (this is achieved by exploiting *metadata* about the structure of the selected database).

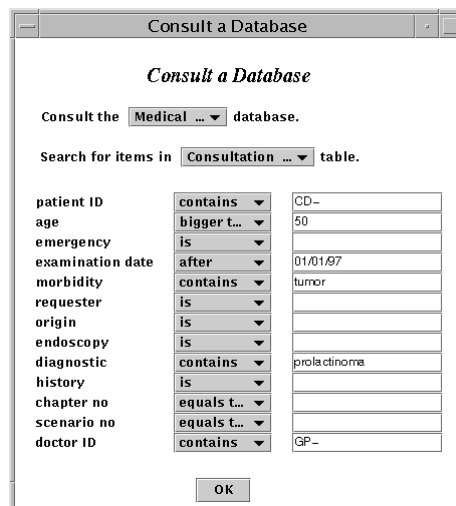


Figure 3 Consulting the Database



### 5.3. The Argument Builder Module

This module aims at assisting agents constructing robust arguments. Since the overall system cannot be expected to make complete judgement decisions itself, the module must be able to evaluate the arguments being presented by participants, to extract as much information concerning their beliefs and argumentation as possible (taking also into account the whole discussion graph), to provoke the appropriate queries to the related databases, and finally, to point out discourse acts that successfully reflect their interests and intentions. The module interoperates with both the NLP and the Information Retrieval modules and should be viewed as a users' "advisor".

The module follows a set of rules, concerning the *grammar* and *syntax* of a CDM-related discourse, in order to match patterns of an on-going discussion with pre-specified ones. For instance, if one's belief is supported by only one piece of evidence, it will suggest him/her to take the appropriate actions to further support it; in cases of conflicts between agents, the module will advise them on actions for conflict resolution (such actions try to achieve a deeper understanding of the issue to be solved).

## 6. Discussion

Approaches to decision making may be divided into two large classes. In the first one, a set of alternatives is determined *a priori* and the task is to select one of 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: (i) an overall *task goal* is specified; (ii) a set of *alternatives* is selected (this set may not be exhaustive); (iii) a collection of *choice criteria* must be determined by the participants; (iv) 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. The sub-processes themselves may be considered to be CDM processes. In fact, each element of the decision process may itself be the subject of a sub-decision process. Decision making can, therefore, be recursive.

The set of alternatives may 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, and 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.

It is our belief that the integration of the modules discussed above with a properly specified model of discourse acts 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 web forums) 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 appears to be feasible. We view human intervention as necessary to determine certain instances 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. The participants can use this to better understand the arguments of others, and advance more pointed argumentation of their own.

## Acknowledgments

The authors thank Joerg Gsponer and Nikos Stergiopoulos for providing the real example used in this paper, and Dekai Wu for his helpful comments on the NLP module. Nikos Karacapilidis is supported by the Commission of the European Communities through the Ercim Fellowship Programme, and Dimitris Papadias by DAG96/97.EG36 and DAG97/98. EG02.

## References

- [1] Aamodt, A., and Plaza, E. (1996). Case-Based Reasoning: Foundational Issues, Methodological Variations and System Approaches. *AI Communications* 7(1).
- [2] 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.
- [3] Ballim, A. and Wilks (1991). Artificial Believers. Lawrence Erlbaum Associates, Hillsdale, NJ.
- [4] Cohen, P.R., and Levesque, H.J. (1990). Intention is choice with commitment. *Artificial Intelligence*, 42(3), 213-261.
- [5] Conklin, J. (1992). Capturing Organizational Memory. In Proceedings of *GroupWare'92*, D. Coleman (ed.), Morgan Kaufmann, San Mateo, CA, 133-137.
- [6] Conklin, J. (1996). Designing Organizational Memory: Preserving Intellectual Assets in a Knowledge Economy. GDSS Working Paper. Available at: <http://www.gdss.com/DOM.htm>
- [7] 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.
- [8] 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.
- [9] 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.
- [10] Farley, A.M., and Freeman, K. (1995). Burden of Proof in Legal Argumentation. In *Proceedings of the ICAIL'95 Conference*, ACM Press, New York, 156-164.
- [11] Fischer, G., McCall, R., and Morch, A. (1989). JANUS: Integrating Hypertext with a Knowledge-based Design Environment. In *Proceedings of Hypertext'89*, ACM Press, New York, 105-117.
- [12] Gordon, T., and Karacapilidis, N. (1997). The Zeno Argumentation Framework. In *Proceedings of the ICAIL'97 Conference*, ACM Press, New York, 10-18.
- [13] 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.
- [14] Hurwitz, R., and Mallery, J.C. (1995). The Open Meeting: A Web-Based System for Conferencing and Collaboration. In *Proceedings of 4th International WWW Conference*.
- [15] Jennings, N. (1993). Specification and Implementation of a Belief-Desire-Joint-Intention Architecture for Collaborative Problem Solving. *Int. Journal of Intelligent and Cooperative Information Systems*, 2(3), 289-318.
- [16] Karacapilidis, N.I., and Papadias, D. (1998a). HERMES: Supporting Argumentative Discourse in Multi-Agent Decision Making. In *Proceedings of AAAI*, AAAI/MIT Press.
- [17] Karacapilidis, N.I., and Papadias, D. (1998b). A Computational Approach for Argumentative Discourse in Multi-Agent Decision Making Environments, *AI Communications* 11(1).
- [18] Karacapilidis, N.I., Trousse, B., and Papadias, D. (1997). Using Case-Based Reasoning for Argumentation with Multiple Viewpoints. In D. Leake & E. Plaza (eds.) *Case-Based Reasoning: Research and Development*, Lecture Notes in AI, Vol. 1266, Springer-Verlag, Berlin, 541-552.
- [19] Kreamer, K.L., and King, J.L. (1988). Computer-based systems for cooperative work and group decision making. *ACM Computing surveys* 20(2), 115-146.
- [20] 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.
- [21] 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.
- [22] Merin, A. (1997). Communicative Action as Bargaining: Utility, Relevance, Elementary Social Acts. In *Proceedings of AAAI Fall Symposium on Communicative Action in Humans and Machines*, AAAI Press.
- [23] Rittel, H., and Webber, M. (1973). Dilemmas in a General Theory of Planning. *Policy Sciences*, 155-169.
- [24] Saaty, T. (1980). *The analytic hierarchy process*. McGraw Hill, New York, 1980.
- [25] Searle, J.R. (1969). *Speech acts: An essay in the philosophy of language*. Cambridge University Press.
- [26] Sidner, C.L. (1994). An Artificial Discourse Language for Collaborative Negotiation. In *Proceedings of AAAI-94*, AAAI/MIT Press, 814-819.
- [27] 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.
- [28] 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*, Erlbaum, Hillsdale, NJ, 109-162.
- [29] Streitz, N., Hannemann, J., and Thuring, M. (1989). From ideas and arguments to hyperdocuments: Travelling through activity spaces. In *Proceedings of Hypertext'89*, ACM Press, New York, 343-364.
- [30] 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.
- [31] Tan, S.W., and Pearl, J. (1994). Qualitative Decision Theory. In *Proc. of AAAI Conf.*, AAAI Press, 928-933.