Chapter 2: Agent Communications for Chance Discovery

Peter McBurney[1] and Simon Parsons[1 2]

[1] Department of Computer Science
University of Liverpool
Chadwick Building, Peach Street
Liverpool, L69 7ZF, UK
`p.j.mcburney@csc.liv.ac.uk`
[2] Department of Computer and Information Science
Brooklyn College, City University of New York
2900 Bedford Avenue,
Brooklyn, NY 11210, USA
`parsons@sci.brooklyn.cuny.edu`

Abstract. This chapter considers chance discovery and management in a community of intelligent, autonomous, software agents, where agents may have differing beliefs and intentions. For such a community of agents, we derive a set of five requirements for the design of languages and protocols for communications between the agents when discussing chance discovery and management. We then use these requirements to assess two proposals in the multi-agent systems community for agent communications: generic languages, such as the FIPA ACL, and dialogue game protocols. The latter are found to have greater potential capability to support dialogues over chance discovery and management between autonomous agents.

1 Introduction

Chance discovery and management in many real-world situations involves more than one entity. Identifying the causes of failure in a telecommunications network, for example, may involve co-ordination between a number of companies and organizations, each owning physical facilities or providing services over part of the network concerned. Each company in the network may have only a partial view of the whole network, with no one having an entire view. Moreover, each may have information about its portion of the network which is commercially sensitive, and so therefore cannot be shared with the other participants. In such circumstances, relevant knowledge may be distributed and decision-making may require collaboration between the various entities involved. To enable effective chance discovery and management here the entities involved need to be able to communicate with one another. If processes of chance discovery and management are to be automated, then communications languages and protocols are necessary for the computational entities involved. In this chapter, we present the current state of research on communications languages in the multi-agent systems community, and discuss its relevance for chance discovery and management.

The chapter begins in Section 2 with a discussion of the requirements which chance discovery and management (CDM) place on an languages and protocols for agent communications. We then consider, in Section 3, one of the main proposals for a generic Agent Communications Language, FIPA ACL, against these requirements. This is followed, in Section 4, with a similar discussion of dialogue game protocols, which researchers in AI have adopted recently from philosophy as the basis for agent communications. Both sections 3 and 4 also serve as introductions to the specific proposals, for those unfamiliar with these areas of AI research. Section 5 concludes the chapter with a discussion of future work.

2 CDM Requirements

We begin by assuming we have a community of autonomous computational agents, each with a partial view of some domain. Our notion of agent is that presented in [47]: an agent is situated in a specific environment and there exhibits four characteristics: autonomy of decision-making; social awareness; reactivity to events in its environment; and proactivity, as it seeks to actively achieve some goals within its environment. A community of such agents is a multi-agent system where two or more autonomous computational agents interact together. An example could be a network of water data monitors, each responsible for water and flood control in part of some geographic region. In this example, each monitor may be able to collect local information, send and receive information to and from other monitors, and take local action, such as opening a dam sluice to release excess water. Actions may be taken independently or in collaboration with other monitors.

Given such an agent community, we will assume throughout that the relevant domain of knowledge is partitioned, with each agent in the community having only knowledge of, or power in, some proper subset of the domain. For example, the partition of the domain could be based on geography, as in the case of distributed water control monitors, or based on time, as when earth-based radio transceiver stations are only able to observe orbiting satellites for several hours each day, or based on a conceptual hierarchy, as when telecommunications service providers do not have access to lower or higher layers in the physical network on which they operate. In addition, the participating agents may have different goals and intentions, which may arise as a result of different perceptions of their economic interests. These goals may conflict. Even if all the agents in the community are rational, in the economic-theoretic sense that each seeks to maximize its expected utility [6], their goals may be still be in conflict. A provider of value-added telecommunications services such as voicemail, for example, may only be able to increase its profit margin at the expense of the provider of the physical infrastructure underlying the value-added service.

Imagine such an agent community is engaged in the discovery and management of chance events, such as risks or opportunities. What requirements does this mission place on the communications languages and protocols between the participating agents? Firstly, the language would need to be able to transmit domain-specific information in an appropriate form between participants. But simply being able to transmit information is not sufficient. Because the members of the community each have only partial views of the domain, what appears to one agent as a chance event, a risk of flood say, may not appear that way to other agents in the community. Therefore agents also need the ability to question each other about the information transmitted, to give reasons and for their respective beliefs and intentions, and to challenge and contest these. In other words, the participants need the ability to argue with each other concerning the messages they transmit, and this is the second requirement of an agent communications language for CDM.

Thirdly, the ability to present arguments and counter-arguments means that any communications protocol used by the agents must encode some theory of debate or argument, what is called a logic of argumentation [13]. Under what circumstances, for example, must an agent who asserts that some belief is true be then forced to defend this belief against questions or attacks by other agents? How many such questions may be asked, if the discussion is not to run forever? When must a debate end, and a resolution be sought? It is issues such as these which a theory of argument addresses. As the philosopher Jürgen Habermas wrote [19, p. 22]: *"The logic of argumentation does not refer to deductive connections between the semantic units (sentences) as does formal logic, but to nondeductive relations between the pragmatic units (speech acts) of which arguments are composed."* In a later chapter in this volume, we discuss logics of argumentation in more detail (Chapter 18 by Parsons and McBurney).

Being able to persuade another agent to adopt a new intention is of little value if the communications language does not permit that agent to express any changes to its beliefs or intentions. Thus, a fourth requirement is that the communications language should enable the participating agents to articulate relevant changes in their internal states, a process we have called the expression of *self-tranformation* [33], following [18]. This might seem a trivial requirement, but, as will be seen in the next Section, not all agent communications languages facilitate this.

Awareness of the possibility of rare events should lead rational participants to efforts to prevent or encourage their occurrence, and/or to mitigate or enhance their consequences [36]. Thus any discussion between agents about chance discovery leads naturally to discussions about chance management. But chance management may require action by more than one agent in a community, as when several flood control devices need to co-ordinate their actions to prevent a flood. The assumption we have made of of agent autonomy means that, in general, agents cannot be ordered to adopt a specific intention, but must be persuaded.

In earlier work [30], we considered the nature of dialogues about chance discovery and management relative to a standard typology of dialogues in the philosophy of argumentation. This typology, due to Walton and Krabbe [44] identifies six primary types of dialogue, based upon the information the participants have at the commencement of a dialogue (of relevance to the topic of discussion), their individual goals for the dialogue, and the goals they share. **Information-Seeking Dialogues** are those where one participant seeks the answer to some question(s) from another participant, who is believed by the first to know the answer(s). In **Inquiry Dialogues** the participants collaborate to answer some question or questions whose answers are not known to any one participant. **Persuasion Dialogues** involve one participant seeking to persuade another to accept a proposition he or she does not currently endorse. In **Negotiation Dialogues**, the participants bargain over the division of some scarce resource. Here, the goal of the dialogue — a division of the resource acceptable to all — may be in conflict with the individual goals of the participants. Participants of **Deliberation Dialogues** collaborate to decide what action or course of action should be adopted in some situation. Here, participants share a responsibility to decide the course of action, or, at least, they share a willingness to discuss whether they have such a shared responsibility. Note that the best course of action for a group may conflict with the preferences or intentions of each individual member of the group; moreover, no one participant may have all the information required to decide what is best for the group. In **Eristic Dialogues**, participants quarrel verbally as a substitute for physical fighting, aiming to vent perceived grievances.

Most actual dialogue occurrences — both human and agent — involve mixtures of these dialogue types. In our earlier work [30], we observed that dialogues about chance events — chance discoveries — are forms of Inquiries. However, unlike the Inquiries in the Walton and Krabbe classification, they are not disinterested searches for truth; instead, they involve a search overlaid with a value system. Thus, a discussion about possible risks of some system involves a search not for all possible outcomes, but only for those with negative consequences. Thus, chance discovery dialogues differ from Inquiries. Chance management dialogues, in contrast, are dialogue about what to do to prevent, encourage, mitigate or enhance a chance event, and so may be viewed as Deliberation dialogues. However, as with most Deliberations, they may have Information-seeking, Inquiry, Persuasion or Negotiation sub-dialogues embedded within them.

Despite a full discussion and exchange of arguments and counter-arguments between participants, agents may still disagree about what to believe and about what actions to take, if any. If a decision must be made by the community about action, then the community needs some means to resolve differences between its members. Such a mechanism could be as simple as a voting process, as proposed in [23], or an argument-classification system, as in [31].[1] Alternatively, a resolution mechanism may rely on a more complex process of collaborative deliberation, as in [21]. The resolution process adopted will embed or reflect a particular structure of social and political relationships between the participants, such as democracy [10] or an organizational hierarchy [38]. Whatever political structure and resolution mechanism adopted by the agent community, the communications language needs to support this; if, for example, a voting mechanism is used, then agents need to be able to express their votes subject to whatever other conditions, such as confidentiality and verifiability, are desired. This is the fifth requirement of an agent communications language for CDM: that it enable an appropriate mechanism for resolution of differences of beliefs or intentions of the participants.

Finally, as with the design of any artificial language, there may be requirements regarding computational simplicity, non-redundancy, transparency, etc; we are not concerned with these issues in this chapter. It is also worth noting that in classical mathematical communications theory the semantics of any information transmitted is ignored in the design and analysis of the communications mechanism. For example, Shannon states [43, p. 31]: *"Frequently the messages have* meaning; *that is they refer to or are*

---

[1] The argumentation-classification approach is explained in Chapter 18 by Parsons and McBurney.

correlated according to some system with certain physical or conceptual entities. These semantic aspects of communication are irrelevant to the engineering problem." By contrast, our discussion above reveals this domain as one where the design of the communication mechanism is necessarily influenced by the meaning and purposes — the semantics — of the messages to be transmitted.

3 Generic Agent Communications Languages

In this section we consider a generic Agent Communications Languages from the perspective of the requirements listed above. Over the last decade two major proposals have been advanced for agent communications languages, the Knowledge Query and Manipulation Language (KQML) [28], which arose from work sponsored by the U.S. Government's Defense Advanced Research Projects Agency (DARPA), and the Foundation for Intelligent Physical Agents' Agent Communications Language [15], the FIPA ACL, which arose from attempts to develop an industry standard for agent communications. Initially, KQML existed without a defined semantics, criticism of which led to the FIPA ACL being defined from 1995. Since the two languages are similar [24], we focus attention here only on the FIPA ACL. We first describe the language, and then assess it against the CDM requirements of the previous section.[2]

3.1 FIPA ACL

The FIPA ACL has a two-level hierarchy for communications messages. At the lower, or "inner" level, the content of messages may be expressed in any logical language agreeable to the participants. Such content is wrapped in one of 22 possible locutions, the FIPA Communicative Act Library (CAL), which together comprise an "outer" level. For instance, the *inform* locution has the following syntax [15, p. 11]:

```
(inform
    :sender (agent-identifier :name j)
    :receiver (set (agent-identifier :name i))
    :content
        "weather (today, raining)"
    :language Prolog)
```

In this example the content of the message is *"weather (today, raining)"*, while the locution in which this content is wrapped is *inform*. The syntax of the locution indicates that the language in which the message content is expressed is Prolog. The syntax also indicates that the sender of the message is identified by agent-identifier symbol j, while the intended recipient is identified by symbol i.

We now list the names of the 22 locutions of the FIPA CAL, together with the summary descriptions of their intended purposes taken from the current FIPA specification [15]. Where the syntax of the locution differs from its name, the syntax is included in parentheses after the name.

---

[2] More details regarding FIPA ACL and KQML may be found in [24, 28, 46].

Accept Proposal (accept-proposal) The action of accepting a previously submitted proposal to perform an action.

Agree The action of agreeing to perform some action, possibly in the future.

Cancel The action of one agent informing another agent that the first agent no longer has the intention that the second agent perform some action.

Call for Proposal (cfp) The action of calling for proposals to perform a given action.[3]

Confirm The sender informs the receiver that a given proposition is true, where the receiver is known to be uncertain about the proposition.

Disconfirm The sender informs the receiver that a given proposition is false, where the receiver is known to believe, or believe it is likely that, the proposition is true.

Failure The action of telling another agent that an action was attempted but the attempt failed.

Inform The sender informs the receiver that a given proposition is true.

Inform If (inform-if) A macro action for the agent of the action to inform the recipient whether or not a proposition is true.

Inform Ref (inform-ref) A macro action for send to inform the receiver the object which corresponds to a descriptor, for example, a name.

Not Understood (not-understood) The sender of the act (for example, i) informs the receiver (for example, j) that it perceived that j performed some action, but that i did not understand what j just did. A particular common case is that i tells j that i did not understand the message that j has just sent to i.

Propagate The sender intends that the receiver treat the embedded message as sent directly to the receiver, and wants the receiver to identify the agents denoted by the given descriptor and send the received *propagate* message to them.

Propose The action of submitting a proposal to perform a certain action, given certain preconditions.

Proxy The sender wants the receiver to select target agents denoted by a given description and to send an embedded message to them.

Query If (query-if) The action of asking another agent whether or not a given proposition is true.

Query Ref (query-ref) The action of asking another agent for the object referred to by an [*sic*] referential expression.

Refuse The action of refusing to perform a given action, and explaining the reasons for the refusal.

Reject Proposal (reject-proposal) The action of rejecting a proposal to perform some action during a negotiation.

Request The sender requests the receiver to perform some action. One important class of uses of the request act is to request the receiver to perform another communicative act.

---

[3] Note that, as throughout this list, the wording here is that of the FIPA CAL specifications [15]. The two instances of the word *"action"* in this statement have two different meanings, the first referring to communicative actions in the FIPA Communicative Act Library executed in the course of dialogues by agents adhering to the FIPA ACL, while the second refers to actions undertaken by the agents outside a dialogue. Strangely, these different meanings are not made explicit in the specifications.

Request When (request-when) The sender wants the receiver to perform some action when some given proposition becomes true.

Request Whenever (request-whenever) The sender wants the receiver to perform some action as soon as some proposition becomes true and thereafter each time the proposition becomes true again.

Subscribe The act of requesting a persistent intention to notify the sender of the value of a reference, and to notify again whenever the object identified by the reference changes.

The use in these summaries of the word *"act"* to describe locutions is not accidental. The FIPA ACL has been given an axiomatic semantics, called *SL* (for *Semantic Language*), based on speech act theory [15, Informative Annex A].[4] This theory, due originally to a philosopher of language, John Austin [7], and extended by his student, John Searle [42], considers spoken human utterances as actions, in so far as they may change the state of the world. Some utterances do this demonstrably as when countries issue *declarations of war* on other countries. Other utterances change the world unobservably, as when one person informs another of something he or she does not already know, and thereby causes a revision to the listener's beliefs. Austin and Searle classified spoken utterances by their intended and actual effects on the world (including the internal mental states of those hearing the utterances), and developed pre-conditions for those effects to be realized. Speech act theory was then used to provide a semantics for agent utterances, particularly in work by Philip Cohen, Hector Levesque and Raymond Perrault [11, 12]. Drawing on this work, the FIPA ACL Semantic Language *SL* is a modal logic formalism which defines pre- and post-conditions for the 22 locutions of the FIPA Communicative Act Library in terms of the beliefs, uncertain beliefs and desires of the agents participating in the dialogue. For example, the *inform* locution, in which one agent tells another some proposition, may only be uttered if the first agent believes the proposition to be true. This is a semantic condition for the locution.[5]

3.2 Assessment of FIPA ACL

What is one to make of the FIPA ACL as a protocol for chance discovery and management? We answer this question by considering FIPA ACL against each of the five CDM requirements presented in Section 2.

The first requirement is that the language support the transmission of domain-specific information. This FIPA ACL does, and it does so with no limitations on the content of messages. However, it is worth noting that the 22 locutions of the FIPA Communicative Act Library are not situated in any conversational context. It is as if the agents using them have always been connected to each another, are so now, and always

---

[4] For a description of different types of program semantics, see, e.g., [34].

[5] Following Searle [42], this is termed a *sincerity* condition. Because beliefs and desires are internal states of the agents concerned, expressions by an agent of its beliefs, desires or intentions cannot be verified by other agents without access to the code of the speaker. This is not possible in most open agent systems, and so the issue of verifiability of agent communications languages — determining that participants actually conform to the semantic requirements of the language — is problematic [45].

will be, and every so often one agent sends a message to one or more of the others, an act which may initiate an exchange of further messages. The image that comes to my mind is that of a group of prisoners chained together for a lengthy period in a remote location, with only each other for company; consequently, they know each other's topics of conversation in depth.[6] In such a context there is no need for locutions which express a desire to enter or leave a dialogue, or for introductions, or for statements of what the conversation will be about. Only one locution, *cfp*, appears to initiate a conversation on a specific topic, in this case a call for proposals to undertake some action. This absence of entry and withdrawal locutions may or may not be appropriate for specific CDM domains.

The second CDM requirement is that the language support the questioning and defense of claims, and the giving and receiving of reasons for beliefs and intentions. Against this requirement, FIPA ACL is inadequate. Both this language and KQML grew initially from efforts by DARPA to develop technologies for *knowledge sharing* [24]. Such a conceptual paradigm explains why fully eleven of the 22 locutions seek to request or send information (*cfp, confirm, disconfirm, inform, inform-if, inform-ref, propagate, proxy, query-if, query-ref, subscribe*). Despite this, the language has not been designed with the possibility that such information may be rationally questioned or challenged. An agent receiving an *inform*(ϕ) message who is unsure about the truth of its contents ϕ, or who does not hold the belief that ϕ is true, has few options to express these views. That agent may utter *disconfirm*(ϕ), which indicates disagreement, or perhaps *not-understood*(*inform*(ϕ)), which indicates that the second agent did not understand the *inform*(ϕ) message. For the second agent to question the first agent, the *disconfirm*(ϕ) response may be too strong, while the *not-understood*(ϕ) response may be too weak; how may the second agent simply ask the reasons for believing ϕ to be true? And, how may the first agent respond with an argument for ϕ? In any case, there is no obligation on the first agent to explain or justify its endorsement of or belief in ϕ upon such requests. The design of the language appears to have been undertaken without thought that an agent may seek to persuade another to change its beliefs.

Thus, although FIPA ACL permits domain-specific information to be transmitted, thus satisfying the first CDM requirement, the language does not facilitate the questioning and contestation of such information, or the presentation of arguments for or against specific statements. Thus, the second CDM requirement is not satisfied. Likewise, the absence of rules regarding connecting locutions, such as the obligations of agents who assert some statement to defend it, mean that the third CDM requirement is also not satisfied, the requirement for a theory of argument. Users of FIPA ACL may, of course, agree themselves such a theory and use this over and above the language, in the same way that FIPA protocols for auctions have been defined, for example [16, 17], but the language has not been designed with such higher-level protocols in mind. The absence of a logic of argumentation for the FIPA ACL locutions means that any locution may follow any other. In addition to allowing belligerent or malicious participants to disrupt conversations, this also complicates the task of analysis of utterances and conversations. After all, any given sequence of utterances in the FIPA ACL may be followed by any one of the 22 locutions, and each of them in turn followed by one of 22 locutions, and

---

[6] Jean-Paul Sartre's play *"No Exit"* describes just such a situation [41].

so on. This creates a state-explosion problem for participants analysing a sequence of utterances to decide what locution to utter next, or seeking to infer the future utterances of other participants.

The fourth CDM requirement is that the language permits agents to express *self-transformation*. The sincerity condition on the *inform* locution in FIPA ACL means that agents may only utter propositions they believe to be true. Changes of belief in a statement ϕ could only be expressed by successive utterances, as in:

inform(ϕ)
inform($\neg\phi$)

But there is no requirement or even possibility for an explicit retraction of the first statement before utterance of the second, since there are no explicit retraction locutions. If other agents receive these two *inform* statements, are they to conclude the speaker has changed its belief about ϕ, or that the speaker is merely expressing inconsistent beliefs? Indeed, hearing such successive utterances may be indicative of other states in the speaker's mind: madness, or malicious behaviour, or simply buggy code. The absence of any argumentation-theoretic rules in FIPA ACL mean that all of these conclusions are possible, and so further conversational overhead is required for listeners to determine that the speaker has indeed undergone a self-transformation. Moreover, FIPA ACL assumes that beliefs are two-valued (true or false), and so agents cannot express degrees of belief. In addition, the making of tentative suggestions or uncertain proposals does not appear possible. In short, FIPA ACL does not readily provide the means by which participants may express any self-transformation they may undergo as a result of discussion with other agents.

In Section 2, we cited Habermas's characterization of a logic of argumentation as a set of relations between speech acts. Looking at the FIPA ACL locution definitions, we see (as mentioned above) that eleven of the 22 locutions seek to request or transmit information; and that four involve negotiation (*accept-proposal, cfp, propose, reject-proposal*); six involve the performance of actions (*agree, cancel, refuse, request, request-when, request-whenever*) and two error-handling (*failure, not-understood*).[7] Our conclusion, looking at these locutions, is that this language has been designed primarily to enable information transfers and purchase transactions, not the rational discussions about beliefs and intentions required for dialogues over chance discovery and management. Purchase transactions may be initiated by a *cfp(.)* locution, which may elicit a *propose(.)* response and then a series of counter-proposals, and/or requests for information, etc. Following successful conclusion of these exchanges, the participating agents may report on the action they agreed to in the purchase transaction by means of the various action locutions. However, the absence of a theory of argument means that rational discussion about these transactions is not possible.

The fifth CDM requirement is that the language support appropriate resolution mechanisms between agents. FIPA ACL does not include any such mechanisms in its definition, but these could readily be defined to overlay the language, in the same way that certain auction protocols have been defined, e.g., [16, 17]. There would appear to

---

[7] We count *cfp* in both the information and the negotiation categories. This categorization of locutions is based on that of [46].

be no obstacles, other than the weaknesses identified above, which would preclude such overlaid protocols.

In conclusion, we believe the FIPA ACL as it currently stands is unsuitable as a protocol for rational discussion between autonomous agents for chance discovery and management. It provides limited support for the rational questioning of others' beliefs or intentions, or the means to persuade them to change their beliefs and intentions; it provides limited means for agents whose beliefs change to express those changes; it provides no means to express tentative suggestions, group intentions or arguments for positions. In other words, the protocol does not readily support Inquiry, Persuasion or Deliberation dialogues or Negotiation dialogues other than simple purchase transactions. Since, as we noted in the previous section, Chance Discovery and Chance Management dialogues are types of Inquiry and Deliberation dialogues respectively, FIPA ACL, as it stands, does not support CDM.

4 Dialogue Game Protocols

In this section, we consider another approach to the design of agent communications languages which has recently received attention from researchers in AI. This involves the use of formal dialogue games, which are interactions between two or more players where each player "moves" by making utterances, according to a defined set of rules. Although their study dates at least from Aristotle [5], they have found recent application in philosophy, computational linguistics and Artificial Intelligence (AI). In philosophy, dialogue games have been used to study fallacious reasoning [20, 27] and to develop a game-theoretic semantics for intuitionistic and classical logic [26] and quantum logic [35]. In linguistics, they have been used to explain sequences of human utterances [25], with subsequent application to machine-based natural language processing and generation [22], and to human-computer interaction [8]. Within computer science and AI, they have been applied to modeling complex human reasoning, for example in legal domains [9], and to requirements specification for complex software projects [14]. Dialogue games differ from the games of economic game theory [37] in that payoffs for winning or losing a game are not considered, and because there is no use of uncertainty measures, such as probabilities, to model the possible moves of opponents. They also differ from the abstract games recently used as a semantics for interactive computation [1], since these latter games do not share the rich rule structure of dialogue games, nor are they intended to have themselves a semantic interpretation involving the sharing of beliefs and intentions among a group of agents.

4.1 Formal Dialogue Games

We now present a model of a generic formal dialogue game in terms of the components of its specification, taken from [32]. We first assume that the topics of discussion between the agents are represented in some logical language, whose well-formed formulae are denoted by the lower-case Roman letters, p, q, r, etc. A dialogue game specification then consists of the following elements:

Commencement Rules: Rules which define the circumstances under which the dialogue commences.

Locutions: Rules which indicate what utterances are permitted. Typically, legal locutions permit participants to assert propositions, permit others to question or contest prior assertions, and permit those asserting propositions which are subsequently questioned or contested to justify their assertions. Justifications may involve the presentation of a proof of the proposition or an argument for it. The dialogue game rules may also permit participants to utter propositions to which they assign differing degrees of commitment, for example: one may merely *propose* a proposition, a speech act which entails less commitment than would an *assertion* of the same proposition.

Combination Rules: Rules which define the dialogical contexts under which particular locutions are permitted or not, or obligatory or not. For instance, it may not be permitted for a participant to assert a proposition p and subsequently the proposition $\neg p$ in the same dialogue, without in the interim having retracted the former assertion.

Commitments: Rules which define the circumstances under which participants express commitment to a proposition. Typically, the assertion of a claim p in the debate is defined as indicating to the other participants some level of commitment to, or support for, the claim. Since the work of philosopher Charles Hamblin [20], formal dialogue systems typically establish and maintain public sets of commitments, called *commitment stores*, for each participant; these stores are usually non-monotonic, in the sense that participants can also retract committed claims, although possibly only under defined circumstances.[8]

Termination Rules: Rules that define the circumstances under which the dialogue ends.

Dialogue game protocols have been articulated for each of the rule-governed primary types of dialogues in the typology of Walton and Krabbe: e.g., information-seeking dialogues [22, 44]; inquiries [31]; persuasion dialogues [2]; negotiation dialogues [22, 29]; and deliberations [21]. However, as mentioned earlier, most real-world dialogues (whether human or agent) involve aspects of more than one of these primary types. Two formalisms have been suggested for computational representation of combinations of dialogue: the *Dialogue Frames* of Chris Reed [39], which enable iterated, sequential and embedded dialogues to be represented; and our own *Agent Dialogue Frameworks* [32], which permit iterated, sequential, parallel and embedded dialogues to be represented. Both these formalisms are neutral with regard to the modeling of the primary dialogue types themselves, allowing the primary types to be represented in any convenient form, and allowing for types other than the six of the Walton and Krabbe typology to be included.

As an example of a dialogue game protocol, we present in outline a protocol developed by one of us with Leila Amgoud and Nicolas Maudet [2]. This protocol is based on James MacKenzie's philosophical dialogue game *DC* [27], a game for two players,

---

[8] It is worth noting here that more than one notion of *commitment* is present in the literature on dialogue games.

both subject to the same rules. *DC* enables the participants to argue about the truth of a proposition and was designed to study the fallacy of begging the question (*petitio principii*, or circular reasoning). The agent interaction protocol of [2] based on *DC* has four distinct locutions: *assert, accept, question* and *challenge*; these can be instantiated with a single proposition, and also, for the locutions *assert* and *accept*, with a set of propositions which together constitute an argument for a proposition. Thus the participants may communicate both propositional statements and arguments about these statements, where arguments may be considered as tentative proofs (i.e., logical inferences from assumptions which may not all be confirmed). The locutions of this protocol are similar to those of *DC* except that they do not include a locution for retraction of assertions (called *withdrawal* in *DC*). As with MacKenzie's game, the protocol of [2] establishes commitment stores which record, in full public view, the statements each participant has asserted. The syntax for this protocol has only been provided for dialogues between two participants, but could be readily extended to more agents, as the same authors did subsequently in [3].

Amgoud *et al.* demonstrate that their system enables persuasion dialogues, inquiry dialogues and information-seeking dialogues [2]. However, as the authors note, to permit negotiation dialogues, the protocol requires additional locutions.[9] These are presented in a subsequent paper by the same authors [4], in which three additional locutions are proposed, *request, promise* and *refuse*, making seven in all. In addition to instantiation with propositions and with arguments for propositions, several of these locution can also be instantiated with a two-valued function expressing a relationship between two resources. For example, the locution *promise*$(p \Rightarrow q)$ indicates a promise by the speaker to provide resource q in return for receiving resource p.

4.2 Assessment of Dialogue Games

Given the number and variety of dialogue game protocols which have been proposed, it is not possible to assess each of them against the CDM requirements for agent communications languages of Section 2. However, we may attempt an assessment of generic game protocols against each of the five requirements.

The first CDM requirement is that the language support the transmission of domain-specific information. Because all formal dialogue games proposed so far in AI have adopted the same two-layer structure of *message contents* wrapped inside *locutions*, then this requirement is generally satisfied. As with FIPA ACL and KQML, the message contents may be expressed in any formalism agreed between the participants, and so may support the transmission of domain-specific information. The second CDM requirement is that the language support the questioning and defense of claims, and the giving and receiving of reasons for beliefs and intentions. Dialogue game protocols have their origin in the philosophy of argumentation, and so this requirement has typically been satisfied.[10] The same conclusion applies to the third CDM requirement, which is that the agent communications language encode some theory of argument.

---

[9] It may also require additional locutions for deliberation dialogues, although the authors suggest otherwise.

[10] It is interesting to observe that the main application of formal dialogue games before modern times was their use for oral disputation between philosophers [20].

The fourth requirement is that the language permits agents to express self-transformation. Here, the assessment depends very much on the locutions permitted in the game, and their syntactic, semantic and pragmatic rules of use. For example, in [33], we showed that the extent to which certain dialogue game protocols in AI permit self-transformation differs from one protocol to another. The fifth CDM requirement is that the language support appropriate resolution mechanisms between agents. In general, as with FIPA ACL, dialogue game protocols are sufficiently flexible to permit any agreed resolution method to be overlaid.

We conclude from this brief analysis that well-designed dialogue game protocols are better able to support agent communications for chance discovery and management dialogues than are generic agent communications languages such as FIPA ACL.

5 Conclusions

In this chapter, our focus has been on the design of languages and protocols for communications between autonomous software agents engaged in joint chance discovery and management over some domain. We have assumed that the agents may have only partial views of the domain and may have differing interests and objectives. Thus, their beliefs and intentions may differ, and may even be in conflict. Consequently, any discussion between such a community of agents concerning possible chance events and any chance management actions may require considerable debate and argument. Arising from this conclusion, we proposed five requirements that any agent communications language would need to satisfy in order to fully enable agent dialogues about chance discovery and chance management.

We then considered two broad categories of proposals for agent communications languages from the multi-agent system community with respect to these five requirements: generic agent communications languages and formal dialogue game protocols. For the former category, we took the Agent Communications Language of the Foundation for Intelligent Physical Agents, FIPA ACL, as a typical example. Although it met two of the requirements, we concluded that FIPA ACL was inadequate to the task of supporting dialogues over chance discovery and management. In contrast, dialogue game protocols have much greater potential to support CDM discussions, although this will depend on the specific locutions and rules of the dialogue game in any particular case.

Chance discovery and management in a multi-agent context leads to additional issues and questions. For example, multi-agent chance discovery may be computationally more complex than chance discovery by a single, ideal, agent having full knowledge of the domain. Similar issues have recently been considered in the case of multi-agent diagnosis [40], where each agent had only a partial view of the domain. However, this work did not consider that agents may have possibly-differing objectives, which is likely in any real-world application of multi-agent chance discovery. We hope to consider such issues in future work.

References

1. S. Abramsky. Semantics of interaction: an introduction to game semantics. In A. M. Pitts and P. Dybjer, editors, *Semantics and Logics of Computation*, pages 1–31. Cambridge University Press, Cambridge, UK, 1997.
2. L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In E. Durfee, editor, *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS 2000)*, pages 31–38, Boston, MA, USA, 2000. IEEE Press.
3. L. Amgoud and S. Parsons. Agent dialogues with conflicting preferences. In J-J. Meyer and M. Tambe, editors, *Pre-Proceedings of the Eighth International Workshop on Agent Theories, Architectures, and Languages (ATAL 2001)*, pages 1–14, Seattle, WA, USA, 2001.
4. L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In W. Horn, editor, *Proceedings of the Fourteenth European Conference on Artificial Intelligence (ECAI 2000)*, pages 338–342, Berlin, Germany, 2000. IOS Press.
5. Aristotle. *Topics*. Clarendon Press, Oxford, UK, 1928. (W. D. Ross, Editor).
6. K. J. Arrow. *Social Choice and Individual Values*. Wiley, New York, 1951.
7. J. L. Austin. *How To Do Things with Words*. Oxford University Press, Oxford, UK, 1962. Originally delivered as the William James Lectures at Harvard University in 1955.
8. T. J. M. Bench-Capon, P. E. Dunne, and P. H. Leng. Interacting with knowledge-based systems through dialogue games. In *Proceedings of the Eleventh International Conference on Expert Systems and Applications*, pages 123–140, Avignon, France, 1991.
9. T. J. M. Bench-Capon, T. Geldard, and P. H. Leng. A method for the computational modelling of dialectical argument with dialogue games. *Artificial Intelligence and Law*, 8:233–254, 2000.
10. J. Bohman and W. Rehg, editors. *Deliberative Democracy: Essays on Reason and Politics*. MIT Press, Cambridge, MA, USA, 1997.
11. P. R. Cohen and H. J. Levesque. Rational interaction as the basis for communication. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*, pages 221–255. MIT Press, Cambridge, MA, USA, 1990.
12. P. R. Cohen and C. R. Perrault. Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:177–212, 1979.
13. F. H. van Eemeren, R. Grootendorst, F. S. Henkemans, J. A. Blair, R. H. Johnson, E. C. W. Krabbe, C. Plantin, D. N. Walton, C. A. Willard, J. Woods, and D. Zarefsky. *Fundamentals of Argumentation Theory*. Lawrence Erlbaum Associates, Mahwah, NJ, USA, 1996.
14. A. Finkelstein and H. Fuks. Multi-party specification. In *Proceedings of the Fifth International Workshop on Software Specification and Design*, Pittsburgh, PA, USA, 1989. ACM Sigsoft Engineering Notes.
15. FIPA. Communicative Act Library Specification. Technical Report XC00037H, Foundation for Intelligent Physical Agents, 10 August 2001. Available from: `www.fipa.org`.
16. FIPA. Dutch Auction Interaction Protocol Specification. Technical Report XC00032F, Foundation for Intelligent Physical Agents, 10 August 2001.
17. FIPA. English Auction Interaction Protocol Specification. Technical Report XC00031F, Foundation for Intelligent Physical Agents, 10 August 2001.
18. J. Forester. *The Deliberative Practitioner: Encouraging Participatory Planning Processes*. MIT Press, Cambridge, MA, USA, 1999.
19. J. Habermas. *The Theory of Communicative Action: Volume 1: Reason and the Rationalization of Society*. Heinemann, London, UK, 1984. Translation by T. McCarthy of: *Theorie des Kommunikativen Handelns, Band I, Handlungsrationalitat und gesellschaftliche Rationalisierung.* Suhrkamp, Frankfurt, Germany. 1981.
20. C. L. Hamblin. *Fallacies*. Methuen, London, UK, 1970.

21. D. Hitchcock, P. McBurney, and S. Parsons. A framework for deliberation dialogues. In H. V. Hansen, C. W. Tindale, J. A. Blair, and R. H. Johnson, editors, *Proceedings of the Fourth Biennial Conference of the Ontario Society for the Study of Argumentation (OSSA 2001)*, Windsor, Ontario, Canada, 2001.
22. J. Hulstijn. *Dialogue Models for Inquiry and Transaction*. PhD thesis, Universiteit Twente, Enschede, The Netherlands, 2000.
23. L. Hunsberger and M. Zancanaro. A mechanism for group decision making in collaborative activity. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI-2000)*, pages 30–35, Menlo Park, CA, USA, 2000. AAAI Press.
24. Y. Labrou, T. Finin, and Y. Peng. Agent communication languages: The current landscape. *IEEE Intelligent Systems*, 14(2):45–52, March/April 1999.
25. J. A. Levin and J. A. Moore. Dialogue-games: metacommunications structures for natural language interaction. *Cognitive Science*, 1(4):395–420, 1978.
26. P. Lorenzen and K. Lorenz. *Dialogische Logik*. Wissenschaftliche Buchgesellschaft, Darmstadt, Germany, 1978.
27. J. D. MacKenzie. Question-begging in non-cumulative systems. *Journal of Philosophical Logic*, 8:117–133, 1979.
28. J. Mayfield, Y. Labrou, and T. Finin. Evaluating KQML as an agent communication language. In M. J. Wooldridge, J. P. Müller, and M. Tambe, editors, *Intelligent Agents II*, Lecture Notes in Artificial Intelligence 1039, pages 347–360, Berlin, Germany, 1996. Springer.
29. P. McBurney, R. M. van Eijk, S. Parsons, and L. Amgoud. A dialogue-game protocol for agent purchase negotiations. *Journal of Autonomous Agents and Multi-Agent Systems*, 2002. *In press*.
30. P. McBurney and S. Parsons. Chance discovery using dialectical argumentation. In T. Terano, T. Nishida, A. Namatame, S. Tsumoto, Y. Ohsawa, and T. Washio, editors, *New Frontiers in Artificial Intelligence: Joint JSAI 2001 Workshop Post Proceedings*, Lecture Notes in Artificial Intelligence 2253, pages 414–424. Springer, Berlin, Germany, 2001.
31. P. McBurney and S. Parsons. Representing epistemic uncertainty by means of dialectical argumentation. *Annals of Mathematics and Artificial Intelligence*, 32(1–4):125–169, 2001.
32. P. McBurney and S. Parsons. Games that agents play: A formal framework for dialogues between autonomous agents. *Journal of Logic, Language and Information*, 11(3):315–334, 2002.
33. P. McBurney, S. Parsons, and M. Wooldridge. Desiderata for agent argumentation protocols. In C. Castelfranchi and W. L. Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002), Bologna, Italy*, pages 402–409, New York City, NY, USA, 2002. ACM Press.
34. B. Meyer. *Introduction to the Theory of Programming Languages*. Prentice Hall, New York City, NY, USA, 1990.
35. P. Mittelstaedt. *Quantum Logic*. D. Reidel, Dordrecht, The Netherlands, 1979.
36. Y. Ohsawa. Chance discoveries for making decisions in complex real world. *New Generation Computing*, 20(2), 2002.
37. M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, MA, USA, 1994.
38. P. Panzarasa and N. R. Jennings. The organisation of sociality: a manifesto for a new science of multi-agent systems. In *Modelling Autonomous Agents in a Multi-Agent World: Proceedings of the Tenth European Workshop on Multi-Agent Systems (MAAMAW-01)*. Annecy, France, 2001.
39. C. Reed. Dialogue frames in agent communications. In Y. Demazeau, editor, *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS-98)*, pages 246–253. IEEE Press, 1998.

40. N. Roos, A. ten Teije, A. Bos, and C. Witteveen. An analysis of multi-agent diagnosis. In C. Castelfranchi and W. L. Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002), Bologna, Italy*, pages 986–987, New York City, NY, USA, 2002. ACM Press.
41. J.-P. Sartre. *No Exit and Three Other Plays*. Vintage International, USA, 1989.
42. J. Searle. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, UK, 1969.
43. C. E. Shannon. The mathematical theory of communication. In C. E. Shannon and W. Weaver, editors, *The Mathematical Theory of Communication*, pages 29–125. University of Illinois Press, Chicago, IL, USA, 1963. Originally published in the *Bell System Technical Journal*, October and November 1948.
44. D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. SUNY Series in Logic and Language. State University of New York Press, Albany, NY, USA, 1995.
45. M. J. Wooldridge. Semantic issues in the verification of agent communication languages. *Journal of Autonomous Agents and Multi-Agent Systems*, 3(1):9–31, 2000.
46. M. J. Wooldridge. *Introduction to Multiagent Systems*. John Wiley and Sons, New York, NY, USA, 2002.
47. M. J. Wooldridge and N. R. Jennings. Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10(2):115–152, 1995.