Delegate: Analyze, Deputize, and Supervise *

Eva Cogan and Simon Parsons

Department of Computer and Information Science,
Brooklyn College, City University of New York,
2900 Bedford Avenue,
Brooklyn, NY 11210, USA.
cogan@sci.brooklyn.cuny.edu, parsons@sci.brooklyn.cuny.edu

Abstract. In any organization (such as a multi-agent system), it is natural for parties to delegate some of their obligations to others. Dialogues achieving delegation are thus fundamental to a multi-agent system. Here we extend the range of argumentation-based dialogues to include those involving delegation of obligations. To the formal framework previously investigated by the authors, we now add *obligations* and two new locutions. The *declare* locution, which makes a proposition true, is used to declare a new obligation. The *subscribe* locution is used to keep the delegator updated with progress. This is necessary since the delegator retains responsibility even though he does not fulfill the obligation directly. A new, minimal protocol, *Delegate*, is proposed to handle the various components of delegation. The correctness of its prerequisites is shown to be both necessary and sufficient for its success. Several examples illustrate the kinds of dialogues that can transpire using the Delegate protocol.

1 Introduction

Agents in multi-agent systems must be provided with the ability to engage in dialogues. Much of the recent work on dialogue (Dignum *et al.* [8], Parsons and Jennings [17], Reed [22], Schroeder *et al.* [26], and Sycara [28]) has considered communication that is *argumentation-based*; that is, it is based on the exchange of reasons for and against particular positions. The advantage of the argumentation-based approach is that it provides a way of making communication *rational* — agents do not have to accept statements unless they are supported by a convincing argument.

Using a formal model of argumentation [3] , we've investigated capturing different types of argumentation-based dialogue — information seeking, persuasion, and inquiry in the terminology of Walton and Krabbe [29]. We examined their basic properties and complexity [19] , their prerequisites [7] and their outcomes [20]. Here we extend this investigation, looking at a new kind of dialogue — *delegation*. To implement delegation, we add obligations to our framework and introduce two locutions (*declare* and *subscribe*).

Delegation dialogues are, as we argue below, particularly important for multi-agent systems. In any organisation or team, it is natural for parties to delegate some of their obligations to others. Once the delegator determines what and to whom he wants to

* Dale Carnegie. *How to Stop Worrying and Start Living*. Simon and Schuster, 2004.

© *A. Gelbukh, C. Yáñez Márquez, O. Camacho Nieto (Eds.)*
Advances in Artificial Intelligence and Computer Science
Research on Computing Science 14, 2005, pp. 47-60

delegate, he must declare that the second party has the obligation, and subscribe to be informed of the progress. He must analyze, deputize, and supervise.

The paper is structured as follows: In Sect. 2 we provide the background to our work by discussing organisations and obligations. Section 3 analyzes delegation. In Sect. 4 we present our dialogue system and describe an approach to formal argumentation on which the dialogue system is built. Section 5 introduces the *declare* and *subscribe* locutions that we need. In Sect. 6 we present the Delegate protocol, discuss what it accomplishes, and give examples of its use. Section 7 concludes.

2 Obligations and Organisations

Any multi-agent system, however *ad hoc*, is an *organisation* in the sense of Santos [25] —"a society of agents ... whose rule-governed interactions are aimed at some specified ... goals." Such an organisation has a hierarchy of power relationships and a specification of what agents are obliged to do. Here we will take these as given, assuming that they have been established by the *coalition formation* process [24, 27] that brought the agents together. When an agent is made responsible for securing some goal, it might just go ahead and achieve it. However, if the goal is one that the agent cannot or doesn want to achieve on its own, its responsibility is to exercise its power and influence to get others to perform the tasks necessary for reaching that goal. In other words, there are two ways in which an agent can "see to it that p." It may just act —this is *direct agency*, formalized by $E_X p$, "agent X brings it about that p." Alternatively, it may ensure that the results are obtained, possibly via other agents — this is *indirect agency*, $G_X p$, "agent X ensures that p."

Norman and Reed [16] distinguish between bringing about a *state* and bringing about an *action*. For our purposes, it is sufficient to rely on Hamblin's [12] "pseudostate", a state in which an action has been done, and deal only with states. Indeed, we can get away with a conjunction of propositions as state and need no explicit representation. We believe that our work complements that of [16] by providing the nuts and bolts from which delegation dialogues can be constructed. We suggest locutions that can be used and show how these result in obligations being delegated. To us this seems to supply a missing component in Norman and Reed's model which, while providing a detailed description of the components of a model for delegation, stops short of identifying the formal locutions that agents engaged in delegation dialogues might use.

O denotes what an agent is obligated to do. For instance, $OG_X p$ means that "agent X is obligated to ensure that p holds," in other words, "X is responsible for p." $\Diamond E_X$ denotes that "X is *able* to $E_X p$" and $\Diamond G_X p$ denotes that "X is *able* to $G_X p$."

Influence channels represent the effective powers (of one agent over another) that are recognized within the organization. $X >_p Y$ denotes the existence of an influence channel from X to Y with respect to p. An influence channel may be used to exercise influence by the "attribution of responsibilities," represented by $E_X O G_Y p$. In other words, X brings it about that Y is obligated to ensure that p.

Our notation is summarized in Table 1.

Table 1. Notation

$E_X p$	Direct agency
$G_X p$	Indirect agency
$X >_p Y$	There is an influence channel from X to Y with respect to p
$\Diamond G_X p$	X is *able* to $G_X p$
$Bel\ X\ p$	$p \in BEL(X)$
$OG_X p$	$p \in OBL(X)$
$Int\ X\ p$	$p \in INT(X)$
$Des\ X p$	$p \in DES(X)$
$CS(X)$	The Commitment Store of X

Some relevant properties of these modalities, which we use in this paper, are as follows:

$$G_X p \rightarrow p$$
$$E_X p \rightarrow G_X p$$
$$E_X O G_Y p \rightarrow X >_p Y$$
$$O G_X p \rightarrow \Diamond G_X p$$

The first three of these come from [25] who builds on [14]. The last is the characteristic deontic formula — an agent cannot be obligated to do what it is unable to do. See, for example, [13].

We have previously [3, 18] considered the agents to be structured using the belief/desire/intention (BDI) model [30]. To incorporate obligations we follow [5] who add O (for obligation) to BDI to obtain the BOID model. It specifies the precedence among the mental states of an agent — Beliefs (information), Desires (internally motivated potential goals), Obligations (externally motivated potential goals), and Intentions (goals to which the agent is committed as a result of deliberation). The set of intentions is a consistent subset of the desires and the obligations. The behavior of a BOID agent is governed by the way in which it resolves conflicts within and between its various mental states. For example, the precedence could be BIOD or BDOI. We will assume that the agents are equipped with a mechanism for resolving such conflicts.

3 Delegation

Delegation is fundamental to multi-agent systems. Castelfranchi and Falcone [6] observe that

> Although there are many definitions of "agent", some of which [are] in full disagreement with each other, the majority of them are based on the notions of task, ... "on behalf of" or ... delegation. ... *[I]n delegation an agent A needs or likes an action of another agent B and includes it in its own plan.* In other words, *A is trying to achieve some of its goals through B's actions.*

What is delegation? The dictionary definition reads:

delegate

noun – a person authorized or sent to act for others; a deputy; an agent; a representative.

v.t. – to authorize, send, or appoint as a delegate

When I delegate to you, I must give you the responsibility to ensure something for which I am responsible[1], and so I need to choose a person who is able to ensure what want and who will accept the responsibility.[2] I set a clear goal but allow you to discover the best way to achieve that goal. I tell you the relevant information or the sources thereof. You find a plan. This, of course, fits in well with the notion of autonomous agency. [3] However, despite delegating, the ultimate accountability still lies with me, the delegator. That's why I must establish appropriate checkpoints, and this, in turn, is why *subscribe* locutions are needed.

It is rational for me to delegate a responsibility to you when the following conditions hold:

1. It is my responsibility, but one of the following applies:
   - I can't do it myself.
   - I can't do it by myself as well as I can with your help.
   - I don't want to do it myself.
2. I believe you are capable.
3. I believe I have the authority to obligate you.
4. I believe you will fulfill your obligation. An agent is autonomous. It "decides" whether or not to fulfill its obligation according to its BOID type. In the deliberation [4] that precedes the dialogue, A chooses to delegate to B because A determines that B will (or with enough probability will) carry out its obligation once it accepts it; that is, it is what [5] call a *social agent.*

So, in delegation, agent A seeks to obligate agent B to ensure that A's obligation fulfilled. A believes that it has the authority to do so and that B is able to do the job.

Falcone and Castelfranchi[10] classify delegation based on the following:

- the interaction between the agents. For example, *strong delegation* is delegation based on explicit agreement.
- the degree of specification of the task. For example, in *open delegation* the agent delegates just a result (state of the world) and leaves it up to the delegatee to decide on a plan. He doesn't specify the individual actions.
- the degree of control. For example, the agent maintains control by requiring *feedback* on whether the state has been realized.

Our analysis of delegation leads us to believe that the three levels just described are appropriate for autonomous agents in a multi-agent system.

A dialogue protocol to implement delegation will be proposed in Sect. 6. But first, we need some machinery.

---

[1] http://www.see.ed.ac.uk/ gerard/Management/art5.html

[2] As Andrew Carnegie said "The secret of success is not in doing your own work but in recognizing the right man to do it."

[3] Our colleague Ira Rudowsky [23] quips that an intelligent agent must "know how with no how."

4 Argumentation Based Dialogues

4.1 Dialogues

A *dialogue* is a sequence of messages exchanged between two agents which all bear upon the same subject. Each agent X has a private knowledge base, Σ_X, partitioned into a set of beliefs, obligations, intentions and desires denoted by $BEL(X)$, $OBL(X)$, $INT(X)$, and $DES(X)$ respectively. We will use the notation $Bel\ X\ p \equiv p \in BEL(X)$, $OG_X p \equiv p \in OBL(X)$, $Int\ X\ p \equiv p \in INT(X)$, and $Des\ X\ p \equiv p \in DES(X)$.

There is also an organizational knowledge base, ORG, accessible by every agent, which contains, among other facts, the influence channels of the organization, i.e., the hierarchy.

Each agent's commitment store, $CS(X)$, contains the commitments made by the agent during the current dialogue. Following Hamblin [11] we take commitments to be propositions that an agent is prepared to defend. Each agent in a dialogue has access to its own private knowledge base and both commitment stores. The union of the commitment stores can be viewed as the public state of the dialogue at any given time.

4.2 Locutions

The following locutions (moves in the dialogue game) are available to the agents. Some of the moves we use here were first introduced in [19] and modified in [21]. Each locution has a rule describing how to update the commitment store after the move. For all moves, player A addresses the ith move of the dialogue to player B, p is a proposition, and S is a set of propositions. The special character $\mathcal{U}$ may also be asserted. It indicates that A cannot give an answer. As soon as $\mathcal{U}$ is asserted, the dialogue terminates.

The first two moves allow propositions to be asserted. An agent uses these locutions to state propositions that it wishes to place "on the record" in the dialogue. Typically these are ones that it wishes the other agent in the dialogue to accept. The next two moves respond to assertions, taking the propositions that another agent has asserted and moving them into the speaker's commitment store. The *question* locution can be used to ask the other player about the truth of any proposition. Since a question makes no commitment, the CS remains unchanged. Finally, *challenge* is a means of asking the other player to state the support of an argument for a proposition.

$assert(p)$	$CS_i(A) = CS_{i-1}(A) \cup \{p\}$
$assert(S)$	$CS_i(A) = CS_{i-1}(A) \cup S$
$accept(p)$	$CS_i(A) = CS_{i-1}(A) \cup \{p\}$
$accept(S)$	$CS_i(A) = CS_{i-1}(A) \cup S$
$question(p)$	$CS_i(A) = CS_{i-1}(A)$
$challenge(p)$	$CS_i(A) = CS_{i-1}(A)$

4.3 Acceptable Arguments

Multi-agent systems are made up of a number of autonomous agents that are brought together in order to achieve some task. We assume that the agents are individually *rational*, by which we mean [30] that they act in their own best interests. We are interested in managing the interactions between the agents in order to ensure that they too are rational in the sense discussed by [15], a sense in which an agent does not have to accept an assertion made by another agent unless the second agent can provide a convincing reason for its assertion. We use a formal system of argumentation both to give agents mechanism for constructing such reasons and to give them a means to assess whether reason is convincing.

We briefly summarize the formal argumentation system of [19, 20] which we use for this purpose. Their full description also deals with preferences between arguments, which, for simplicity, we ignore here. There are, of course, other ways to define a system of argumentation. This is just one approach, based on [1, 2], which itself is based on [9], and which our experience suggests is an adequate framework for handling agent communication. Table 2 lists some of the symbols that we use.

Table 2. Symbols

$\vdash$	$\equiv$	$\neg$	$\rightarrow$	$\wedge$	$\vee$
inference	equivalence	negation	implication	and	or

Each agent X has a knowledge base, Σ_X, which contains formulas of a propositional language.

An *argument* is a 2-tuple (S, p) where

- p is a proposition
- $S \subseteq \Sigma_X$ satisfying all of the following:
  1. S is consistent
  2. $S \vdash p$
  3. S is minimal, so no proper subset of S can $\vdash p$

S is called the *support* of the argument and p is its *conclusion*.

If agent A is engaged in a dialogue with agent B, then $S \subseteq (\Sigma_A \cup CS(A) \cup CS(B))$.

Two arguments may conflict. More precisely arguments may *undercut* one another. An argument is undercut if and only if there is another argument which has as its conclusion the negation of an element of the support for the first argument.

Now, a set of arguments S *defends* an argument A iff for each argument B that undercuts A, there is an argument in S that undercuts B. An *acceptable* argument, is one that is not undercut, or for which there is an acceptable argument that undercuts each of the arguments that undercut A. An acceptable argument is one which is, some sense, proven since all the arguments which might undermine it are themselves

undermined. However, this status can be revoked following the discovery of a new argument (possibly as the result of the communication of some new information from another agent via a dialogue).

The preconditions for the locutions are determined by the knowledge base and what we have previously called the *attitude* of the agent. Different attitudes are appropriate for the different roles played by agents. For example, in court a lawyer may wish to assert propositions that are supported by arguments that he doesn't find persuasive but which he hopes the jury will, while in a discussion with a friend, the same lawyer may only wish to assert propositions supported by arguments which both he and the friend find acceptable. A range of attitudes for the assertion and acceptance of propositions is explored in [21], and here we assume what is called a thoughtful/skeptical agent, one that is *allowed* to *assert* and *accept* only propositions for which it has an acceptable argument.

5 *Declare* and *Subscribe* Locutions

To handle delegation, we introduce the *declare* and *subscribe* locutions.

5.1 Declare

To *declare* is to make it so by saying so. Examples include "I do hereby declare you married (hired, obliged, appointed, Harry, etc.)" and "This meeting is hereby adjourned" (a declaration often not used early enough in the opinion of the authors). The state is brought about *by* the declaration. One can declare a status in society, such as "hired", but not a physical property, such as "pink". A new mathematical definition is a declaration. (By contrast, a dictionary definition is an assertion of current usage.) We need declaration for delegation, and so introduce a locution to capture it:

declare(p) $\quad CS_i(A) = CS_{i-1}(A) \cup \{p\}$

Please note that if B challenges A's declaration, B is really challenging A's authority to declare p. A's response must establish A's authority, not the truth of p. There may also be conditions that must be true for p. For example, in the United States, one cannot declare two people married to each other if at least one of them is already married to someone else. Furthermore, they need to have a marriage licence. These would be adequate grounds to challenge a declaration of marriage (and indeed some marriage services even have a point at which the officiator asks if any such challenges will be put forward). In particular, recall that one cannot be obligated to bring about a state when one isn't able to do so. So, when p is of the form OG_Bq, A needs to establish $A >_q B$ and $\Diamond G_Bq$.

5.2 Subscribe

We also need *subscribe*, as used in KQML. The reason is that one who delegates an obligation is still responsible for it as well. Thus he needs to be notified of the fulfillment of the obligation or of progress towards that goal. The *subscribe* locutions are the

mechanism by which this can be implemented. We distinguish between two kinds subscription:

subscribeAt($\{t_i\}$) A subscribes to be informed at a specific time or at sequence times. B has to report at those times. This may, for example, have the effect of question that gets repeated at regular intervals.

subscribeIf(p) A subscribes to be informed when some proposition p becomes true.

Like *question*, *subscribe* does not affect the commitment store. $CS_i(A) = CS_{i-1}(A)$. We postpone a more formal representation of *subscribe* until future work.

Subscriptions are common. Obvious examples include: stock quotes on a computer screen, alerts, exploration results, and monitoring a patient in a hospital. Both kinds of *subscribe* might be used. A nurse may be asked to update a doctor with information about a patient's condition at regular intervals as well as when some significant event occurs. An investor might subscribe to receive a financial newsletter every month and alerts when there are major moves or threats in the market.

6 The Delegate Protocol

Delegation is a fundamental dialogue type in a multi-agent system, as discussed Sect. 3 . In a delegation dialogue, A delegates to B its obligation to ensure that p. To propose a protocol for delegation dialogues, we will use the following macro.

Challenge and Defense $CD(X, Y, p)$ is a macro for a common sequence of locutions. Suppose agent X has asserted a proposition p for which agent Y has no acceptable argument. Agent Y then challenges p. Agent X attempts to defend p by providing the support of an argument for p. Y may then (when necessary) challenge each element of the defense. If Y accepts the elements of the defense and they do indeed form the support of an acceptable argument for p, then Y can accept p.

CD(X, Y, p)

1. *Y challenges p*
2. $\begin{cases} X \textit{ asserts } S, \text{ the support of an argument for } p & \text{if allowed by its attitude,} \\ \text{the dialogue terminates} & \text{otherwise.} \end{cases}$
3. for each $s \in S$ $\begin{cases} Y \textit{ accepts } s & \text{if allowed by its attitude,} \\ CD(X, Y, s) & \text{otherwise.} \end{cases}$
4. $\begin{cases} Y \textit{ accepts } p & \text{if allowed by its attitude,} \\ \text{the dialogue terminates} & \text{otherwise.} \end{cases}$

6.1 Delegate

Using the machinery we've developed, we now give what seems to us to be the minimum protocol required to achieve the delegation of an obligation. The preconditions for the Delegate protocol follow directly from the analysis of delegation in Sect. 3. A

is responsible to ensure that p. A intends to make B responsible for ensuring p. A believes it has the authority to do so and that B is able to do the job. During a delegation dialogue, A delegates to B its obligation to ensure that p.

Delegate(A, B, p)

preconditions:

- OG_{Ap}
- $Int\ A\ OG_{Bp}$
- $Bel\ A \Diamond G_{Bp}$
- $Bel\ A\ A >_p B$

1. *A asserts* relevant information or sources
2. $\begin{cases} B\ accepts \text{ relevant information or sources} & \text{if allowed,} \\ CD(A, B, \text{relevant information or sources}) & \text{otherwise.} \end{cases}$
3. $\begin{cases} A\ subscribeAts\{t_i\} \text{ to be informed of progress} & \text{or} \\ A\ subscribeIfs\ p & \text{or both} \end{cases}$
4. *A declares* OG_{Bp}
5. $\begin{cases} B\ accepts\ OG_{Bp} & \text{if allowed,} \\ CD(A, B, OG_{Bp}) & \text{otherwise.} \end{cases}$

There are several things to note about this protocol:

1. B could give A the authority to declare an obligation for it even if ORG does not.
2. The agents might negotiate ([18]) how often the update will occur or how detailed it will be. The delegator will keep closer tabs on an agent it doesn't trust as much or when the stakes are higher.
3. As in our earlier protocols, we assume that an agent accepts whenever it can.
4. B might have some questions.
5. While more complex protocols may be considered (or required) in certain situations (something we will investigate in the future), this is sufficient to permit a basic form of delegation.
6. The dialogue succeeds when B accepts the obligation, thus fulfilling A's intention. This corresponds to level 2 below.

6.2 Successes

We want to see that the Delegate protocol accomplishes what we set out to do. There are three levels of accomplishment:

1. ***Creation of obligation.*** Suppose A intends that B be obligated to ensure p. Assume B is able to see to it that p, either directly or indirectly, and there is an influence channel from A to B with respect to p. A's declaration of B's obligation accomplishes A's intention. $\{\Diamond G_{Bp},\ A >_p B,\ A\ declares\ OG_{Bp}\} \vdash OG_{Bp}$. A has created the obligation for B.

2. ***Delegation of obligation.*** If, in addition, p was initially the obligation of A, and B accepts this obligation, then A has successfully delegated its obligation to B. The correctness of A's beliefs is sufficient for A to achieve A's intention and, thus, the *success of the delegation dialogue.*
3. ***Fulfillment of obligation.*** Moreover, if A is right about B being an agent whose BOID type causes it to fulfill its obligations, it will "see to it that p." $\{OG_{BP}, OG_{BP} \rightarrow (E_{BP} \vee G_{BP}), E_{BP} \rightarrow G_{BP}, G_{BP} \rightarrow p\} \vdash p$. So, G_{AP} but $\neg E_{AP}$ and A found a plan to achieve its goal of having its obligation fulfilled through B. This is, of course, beyond the scope of the delegation dialogue.

In addition to being sufficient for success of the dialogue, the correctness of A's two beliefs is also necessary. As was noted in Sect. 3,

- $E_A OG_{BP} \rightarrow A >_p B$. To create an obligation for B, A must have influence over B.
- $OG_{BP} \rightarrow \Diamond G_{BP}$. Deontic Logic and commonsense don't let you hold someone responsible for what he is unable to do.

What's more, I can't delegate an obligation that isn't mine, so OG_{AP} must be true. So the truth of $OG_{AP} \wedge \Diamond G_{BP} \wedge (A >_p B)$ is necessary for the success of the delegation dialogue.

Thus the correctness of the preconditions of the Delegate protocol is both necessary and sufficient for the Delegate protocol to achieve A's intention; in other words, for the protocol to succeed.

6.3 Examples

The following examples illustrate the use of the protocol.

Example 1. Professor A wants to attend a conference. A asks grad student B to teach A's class. Let t = "The class is taught."

$Delegate(A, B, t)$
preconditions: OG_{At} (Normally it's OE_{At}, but if A doesn't intend to E_{At}, it's OG_{At}.), $Int\ A\ OG_{Bt}$, $Bel\ A \Diamond G_{Bt}$, $Bel\ A\ A >_t B$

A $assert(\{$"We're up to Chapter 8.", " 'www.bc.edu/slides' is the URL."$\})$
B $accept(\{$"We're up to Chapter 8.", " 'www.bc.edu/slides' is the URL."$\})$
A $subscribeAt$("After the lecture, let me know who attended.")[4]
A $declare(OG_{Bt})$
B $accept(OG_{Bt})$

This first example is straightforward. Someone who would normally fulfill a responsibility, but doesn't want to do so this time, delegates it to another person. In the next example, the delegator cannot fulfill the obligation by himself as well as he'd like.

[4] After the lecture, for (i = 1 to NumberOfStudents) *question*(student[i] attended).

Example 2. Our professor doesn't feel as comfortable with a complex subject as he'd like. He asks his colleague, who is an expert in the field, to give a guest lecture. His colleague kindly gives him the authority to delegate. The dialog proceeds the same way.

In the next example the delegator cannot fulfill his obligation without help and therefore has no option but to delegate it.

Example 3. Money manager A determines that his client should own a certain stock. A is obligated to act in the best interest of his client and ensure that the client owns the stock. A commissions broker B to buy the stock. Let b = "A's client owns some stock."

$Delegate(A, B, b)$
preconditions: $OG_A b$, $Int\ A\ OG_B b$, $Bel\ A \Diamond G_B b$, $Bel\ A\ A >_b B$

A $assert(\{$"symbol = 'ABC' ", "quantity = 100", "maximum price = 10"$\})$
B $accept(\{$"symbol = 'ABC' ", "quantity = 100", "maximum price = 10"$\})$
A $subscribeIf(b)$
A $declare(OG_B b)$
B $accept(OG_B b)$

In the next example the delegator has to do a little convincing to get the delegatee to accept.

Example 4. Hospital administrator A is approached by Donor D about surgery. He gives enough money to demand the best care. He might make a conditional promise (or a threat). A tells surgeon B to do the surgery. After some discussion, B agrees. Let c = "D gets best care." Let b = "Surgeon B, on the staff of the hospital, is known to be the best in the field." Let p = "B performs the surgery." Let d = "D makes a large donation this year."

$Delegate(A, B, p)$
preconditions: $OG_A p$, $Int\ A\ OG_B p$, $Bel\ A \Diamond G_B p$, $Bel\ A\ A >_p B$

A $assert($"D's phone number is 1(718)555-1234."$)$
B $accept($"D's phone number is 1(718)555-1234."$)$
A $subscribeIf(p)$
A $declare(OG_B p)$
B $challenge(OG_B p)$
A $assert(\{\Diamond G_B p, A >_p B\})$
B $accept(\Diamond G_B p)$
B $challenge(A >_p B)$
A $assert(\{b,\ b \wedge c \rightarrow p,\ \neg c \rightarrow \neg d\})$
B $accept(b)$ ("unreluctantly" accepts the compliment)
B $accept(b \wedge c \rightarrow p)$
B $accept(\neg c \rightarrow \neg d)$
B $accept(A >_p B)$
B $accept(OG_B p)$

In the next example, the delegator attempts to delegate rather than fulfill his obligation himself. However, the delegatee can't accept that he is able do it and the dialogue fails.

Example 5. The situation starts out like the one in [16] but here the professor isn't successful. Professor A doesn't want to spend time making copies of the handouts for his class. A instructs secretary B to do so on his behalf. He is busy preparing the chairperson's report, which has a higher priority. Let c = "The handouts are copied." and r "The report is written."

$Delegate(A, B, c)$
preconditions: $OG_A c$, $Int\ A\ OG_B c$, $Bel\ A\ \Diamond G_B c$, $BelA\ A\ >_c\ B$

A $assert$("The originals are in A's mailbox.")
B $accept$("The originals are in A's mailbox.")
A $subscribeIf(c)$
A $declare(OG_B c)$
B $challenge(OG_B c)$
A $assert(\{\Diamond G_B c, A\ >_c\ B\})$
B $challenge(\Diamond G_B c)$

A can't produce an argument acceptable to B because
$\{OG_B r,\ OG_B r \rightarrow G_B r,\ G_B r \rightarrow \neg\Diamond G_B c\} \subseteq \Sigma_B$.

In the final example, the delegatee challenges the delegator's right to generate obligation. The delegatee does not grant the delegator the needed authority and dialogue terminates unsuccessfully.

Example 6. Our lazy (or overwhelmed) Professor A doesn't want to spend time preparing an exam. A asks secretary B to do it. The task isn't in B's job description and refuses. Let p = "The exam is prepared."

$Delegate(A, B, p)$
preconditions: $OG_A p$, $Int\ A\ OG_B p$, $Bel\ A\ \Diamond G_B p$, $BelA\ A\ >_p\ B$

A $assert$("We covered chapters 1-5.")
B $accept$("We covered chapters 1-5.")
A $subscribeIf(p)$
A $declare(OG_B p)$
B $challenge(OG_B p)$

7 Conclusions

We have argued that, in general, in a multi-agent system some agents will need to fulfill obligations by getting other agents to complete them. The way to do this in an organisation (of which a multi-agent system is an instance) is delegation, and so there is a need for mechanisms for delegation in multi-agent systems. By providing such a mechanism

this paper extends the formal inter-agent dialogues that were studied in [19, 20, 7] to deal with delegation.

The main contributions of this paper are as follows. First we introduced obligations into the formal framework of [19, 20]. This is necessary since obligations are the tokens that are passed during delegation—I delegate p to you by getting you to accept the obligation to ensure p. Second, we introduced two new locutions, *declare* and *subscribe*, which are necessary parts of a delegation. The *declare* locution makes a proposition true , while *subscribe* sets up a reporting structure that ensures that the delegator is updated with the progress towards the fulfillment of the delegated obligation. The first is used to declare an obligation for the delegatee. The second is necessary since the delegator retains responsibility even though he does not do it directly. Finally, we proposed a new protocol, Delegate, that handles the various components of delegation. The correctness of its prerequisites was shown to be both necessary and sufficient for its success. Some examples illustrated the dialogues that can transpire while following the Delegate protocol.

We plan to further investigate the dialogues that the Delegate protocol can support, further evaluate its properties, and develop more complex versions. We also intend to consider how delegation dialogues may be combined with the other kinds of dialogue to manage agent teams.

Acknowledgements This work was made possible by funding from NSF #REC-02-19347, NSF #IIS 0329037, EU FP6-IST 002307 (ASPIC), PSC-CUNY Award #66171-00 35 and PSC-CUNY Award #67437-00 36. We thank Karen Kletter,Peter McBurney, Rohit Parikh, and the anonymous reviewer for their comments.

References

[1] L. Amgoud. *Contribution a l'integration des préferences dans le raisonnement argumentatif.* PhD thesis, Université Paul Sabatier, Toulouse, July 1999.

[2] L. Amgoud and C. Cayrol. On the acceptability of arguments in preference-based argumentation framework. In *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, pages 1–7, 1998.

[3] 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*, pages 31–38, Boston, MA, 2000. IEEE Press.

[4] M. Bratman. *Intentions, plans, and practical reason.* CSLI Publications, Stanford, California, 1999. Originally published by Harvard University Press, Cambridge, MA, in 1987.

[5] J. Broersen, M. Dastani, J. Hulstijn, Z. Huang, and L. der van Torre. The BOID architecture: conflicts between beliefs, obligations, intentions and desires. In J. P. Müller, E. Andre, S. Sen, and C. Frasson, editors, *Proceedings of the Fifth International Conference on Autonomous Agents*, pages 9–16, Montreal, Canada, 2001. ACM Press.

[6] C. Castelfranchi and R. Falcone. Towards a theory of delegation for agent-based systems. *Robotics and Autonomous Systems*, 24:141–157, 1998.

[7] E. Cogan, S. Parsons, and P. McBurney. What kind of argument are we going to have today? In S. Kraus, M. Singh, S. Koenig, and M. Wooldridge, editors, *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems*. ACM Press, 2005.

[8] F. Dignum, B. Dunin-Kęplicz, and R. Verbrugge. Agent theory for team formation by dialogue. In C. Castelfranchi and Y. Lespérance, editors, *Seventh Workshop on Agent Theories, Architectures, and Languages*, pages 141–156, Boston, USA, 2000.

[9] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.

[10] R. Falcone and C. Castelfranchi. The human in the loop of a delegated agent: the theory adjustable social autonomy. *IEEE Transactions on Systems, Man, and Cybernetics, Part* 31(5):406–418, 2001.

[11] C. L. Hamblin. *Fallacies*. Methuen and Co Ltd, London, UK, 1970.

[12] C. L. Hamblin. *Imperatives*. Basil Blackwell, 1987.

[13] J. F. Horty. *Agency and Deontic Logic*. Oxford University Press, 2001.

[14] A.J.I. Jones and M. Sergot. A formal characterisation of institutionalised power. *Journal of the IGPL*, 4(3):427–443, June 1996.

[15] P. McBurney. *Rational Interaction*. PhD thesis, University of Liverpool, 2002.

[16] T. J. Norman and C. Reed. Delegation and responsibility. In C. Castelfranchi and Y. Lespérance, editors, *Intelligent Agents VII*, pages 136–149. Springer Verlag, 2001.

[17] S. Parsons and N. R. Jennings. Negotiation through argumentation — a preliminary report. In *Proceedings of Second International Conference on Multi-Agent Systems*, pages 267-274, 1996.

[18] S. Parsons, C. Sierra, and N. R. Jennings. Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292, 1998.

[19] S. Parsons, M. Wooldridge, and L. Amgoud. An analysis of formal inter-agent dialogues. In *1st International Conference on Autonomous Agents and Multi-Agent Systems*. ACM Press, 2002.

[20] S. Parsons, M. Wooldridge, and L. Amgoud. On the outcomes of formal inter-agent dialogues. In *2nd International Conference on Autonomous Agents and Multi-Agent Systems*. ACM Press, 2003.

[21] S. Parsons, M. Wooldridge, and L. Amgoud. Properties and complexity of formal inter-agent dialogues. *Journal of Logic and Computation*, 13(3):347–376, 2003.

[22] C. Reed. Dialogue frames in agent communications. In Y. Demazeau, editor, *Proceedings of the Third International Conference on Multi-Agent Systems*, pages 246–253. IEEE Press, 1998.

[23] I. Rudowsky. Intelligent agents. Proceedings of AMCIS2004- Americas Conference Information Systems, 2004.

[24] T. Sandholm and V. R. Lesser. Coalitions among computationally bounded agents. *Artificial Intelligence*, 94(1-2):99–137, 1997.

[25] F. Santos. A modal logic framework for organization analysis and design. In John Horty and Andrew J.I. Jones, editors, *Sixth International Workshop on Deontic Logic in Computer Science (DEON'02)*, 2002.

[26] M. Schroeder, D. A. Plewe, and A. Raab. Ultima Ratio: Should Hamlet kill Claudius?. *Proceedings of the 2nd International Conference on Autonomous Agents*, pages 467–468, 1998.

[27] O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1–2):165–200, 1998.

[28] K. Sycara. Argumentation: Planning other agents' plans. In *Proceedings of the Eleventh Joint Conference on Artificial Intelligence*, pages 517–523, 1989.

[29] D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, Albany, NY, 1995.

[30] M. Wooldridge. *Reasoning about rational agents*. The MIT Press, 2000.