A Formal Semantics for Proxy Communicative Acts

Marcus J. Huber, Sanjeev Kumar, Philip R. Cohen

Computer Science and Engineering
Oregon Graduate Institute
20000 NW Walker Road
Beaverton, OR 97006, USA
{marcush,skumar,pcohen}
@cse.ogi.edu

David R. McGee

Battelle, Pacific Northwest
National Laboratory
Richland, WA 99352
david.mcgee@pnl.gov

Abstract

Mediation services are becoming increasingly important in multiagent systems. An agent that can act on behalf of another agent is one important example of mediation functionality commonly required. Within this paper, we define and analyze PROXY and PROXY-WEAK communicative acts that formally specify semantics for interacting with middle agents that provide proxy services. These two communicative acts are shown to have a distinctly different impact upon the mental state of the agents involved and impose significantly different levels of commitment upon the middle agents.

1 Background

A common design attribute in many general-purpose multiagent software architectures and distributed computing environments is agents or processes whose sole purpose is to help locate other agents to find and communicate with them. These mediators or middle agents include the "facilitator" agents in the FIPA [5] and OAA [8] architectures, the "proxy" agents of the DARPA CoABS Project's Grid, and the proxy web servers on a network. Within KQML [4], it is common for an agent to use agents to "recruit" other agents that can provide it services.

These software proxies are an increasingly important aspect of distributed computational systems and are being introduced in a wide range of domains. In many cases, the notion entails a computational process that acts on behalf of, and typically assumes full responsibility for, the activities of another computational process (which ostensibly has some limitation that requires the use of the proxy). In other cases, the proxying entity should have no real responsibility. Unlike KQML's forward [4], proxying does not always entail simply passing on messages between entities that might not otherwise have contact with each other.

As an example of the difference in commitment levels that we are trying to support, recall the Watergate affair. President Nixon wanted the special prosecutor Archibald Cox to be fired, and asked Elliott Richardson (the Attorney General) to do so. Nixon wanted Richardson to take responsibility for the firing rather than just

have Richardson tell Cox that Nixon is firing Cox.[1]

Whereas we spend the beginning of the paper on details of the notation that we employ and the definition of key concepts, the focus of the paper is devoted to the definition and analysis of two new communicative acts, PROXY and PROXY-WEAK. Their employment should facilitate the deployment of mediation agents that provide proxying services. Our analysis below will show that the two acts result in the middle agents having significantly different levels of commitments relative to the final agents, where PROXY imposes significant and PROXY-WEAK imposes relatively little responsibility upon the middle agents. Our definition of PROXY-WEAK facilitates *third-party performative* semantics; we show that satisfaction and successful discharge of the PROXY-WEAK speech act is semantically equivalent to the sending agent performing a speech act directly on the final target agent even while going through an intermediary. As part of our analysis, we show that FIPA's two speech acts of the same names suffer from a number of deficiencies, including misrepresentation of the sender's true intentions and failure to provide third-party performative semantic support.

2 Background Concepts

We adopt an attempt-based semantics [2][3][12] for communication performatives in the following definitions. We refer to this prior work for many of our background definitions. Although in [7] it was shown that the attempt-based semantics of this earlier work can be extended to groups of agents, we will use the symbols α, β, γ, δ, and τ to represent single agents in the following definitions.

We use a modal language with the usual connectives of a first order language with equality, as well as operators for propositional attitudes and event sequences. Full details of this modal language and semantics can be found in [2][3]. In summary, BEL has weak S5 semantics and GOAL has system K semantics, with the possible-world model definitions of [2]. KNOW is defined as true knowledge, i.e., (p $\wedge$ (BEL τ p)). (HAPPENS *a*) and (DONE *a*) say that a sequence of actions described by the action expression *a* will happen next or has just happened, respectively. (HAPPENS τ *a*) and (DONE τ *a*) also specify the agent for the action sequence that is going to happen or has just happened. BEFORE and AFTER are defined in terms of HAPPENS. (UNTIL q p) says p will remain true at least until q is true. (PRIOR p q) says that proposition p will become true no later than proposition q. An action expression is built from variables ranging over sequences of events using constructs of dynamic logic: a;b is action composition and *p*? is a test action. We treat BMB between two agents as a semantic primitive in this paper, as in [7]. Mutual belief between two agents α and β is defined in terms of unilateral mutual belief as (BMB α β p) $\wedge$ (BMB β α p) [3]. In our model, BMB can be established by default [1][6].

Definition 1. *PGOAL (Persistent Goal)*

$$(\text{PGOAL } \tau\ p\ q) \equiv (\text{BEL } \tau \neg p) \wedge (\text{GOAL } \tau \Diamond p) \wedge (\text{KNOW } \tau\ [\text{UNTIL } [(\text{BEL } \tau\ p) \vee (\text{BEL } \tau \Box\neg p) \vee (\text{BEL } \tau \neg q)]\ (\text{GOAL } \tau \Diamond p)]\ ).$$

[1] As a point of interest, Richardson resigned to avoid being Nixon's proxy.

Persistent goal formalizes the notion of commitment. An entity τ having a persistent goal p is committed to that goal. The entity τ cannot give up the goal that p is true in the future, at least until it believes that one of the following is true: p is accomplished, or is impossible, or the relativizing condition q is untrue. In this paper, we often leave the q term unspecified when it has the value of the constant *true*.

Definition 2. *INTEND (Intention)*
(INTEND τ *a q*) $\equiv$ (PGOAL τ [HAPPENS τ (BEL τ (HAPPENS *a*))?;*a*] *q*)

Intention to do an action *a* is a commitment to do the action knowingly. The entity τ is committed to being in a mental state in which it has done the action *a* and, just prior to which, it believed that it was about to do the intended action next.

Definition 3. *ATTEMPT*
(ATTEMPT τ *e* ϕ ψ *t*) $\equiv$
t?;[(BEL τ $\neg\phi$) $\wedge$
(GOAL τ (HAPPENS *e*;$\Diamond\phi$?)) $\wedge$
(INTEND τ *t*?;*e*;ψ? (GOAL τ (HAPPENS *e*;$\Diamond\phi$?)))]?;*e*

An attempt to achieve ϕ via ψ is a complex action expression in which the entity τ is the actor of event *e* at time *t* and, just prior to *e*, the actor chooses that ϕ should eventually become true and intends that *e* should produce ψ relative to that choice. So, ϕ represents some ultimate goal that may or may not be achieved by the attempt, while ψ represents what it takes to make an honest effort.

Definition 4. *PWAG (Persistent Weak Achievement Goal)*
(PWAG τ_1 τ_2 *p q*) $\equiv$
[$\neg$(BEL τ_1 *p*) $\wedge$ (PGOAL τ_1 *p*)] $\vee$
[(BEL τ_1 *p*) $\wedge$ (PGOAL τ_1 (MB τ_1 τ_2 *p*))] $\vee$
[(BEL τ_1 $\Box\neg p$) $\wedge$ (PGOAL τ_1 (MB τ_1 τ_2 $\Box\neg p$))] $\vee$
[(BEL τ_1 $\neg q$) $\wedge$ (PGOAL τ_1 (MB τ_1 τ_2 $\neg q$))]

This definition, adapted from [11], states that an entity τ_1 has a PWAG with respect to another entity τ_2 when the following holds: (1) if entity τ_1 does not believe that *p* is currently true, it will have a persistent goal to achieve *p*, (2) if it believes *p* to be either true, or to be impossible, or if it believes the relativizing condition *q* to be false, then it will adopt a persistent goal to bring about the corresponding mutual belief with entity τ_2.

Definition 5. *SINCERE*
Entity α is sincere with respect to entity β and proposition *p* if, whenever α wants β to come to believe *p*, it wants β to come to know *p*. Agents may not in fact be sincere but, as with all the mental states discussed here, they are responsible for being sincere.

(SINCERE α β *p*) $\equiv$ $\forall$e (GOAL α (HAPPENS *e*;(BEL β *p*)?)) $\supset$
(GOAL α (HAPPENS *e*;(KNOW β *p*)?))

Definition 6. *TRUST*
Entity α trusts entity β for proposition *p* if whenever α believes that β believes *p*, α also believes *p*.

(TRUSTS α β *p*) $\equiv$ (BEL α (BEL β *p*)) $\supset$ (BEL α *p*)

Definition 7. *HAPPENING*
An action expression *a* is happening if one of the following is true (1) *a* has just been done, or (2) *a* is going to happen next (i.e. *a* is just starting), or (3) there exists some initial subsequence of *a* (represented by *e*) that has just been done but *a* is not yet done.

(HAPPENING *a*) ≡(DONE *a*) ∨
(HAPPENS *a*) ∨
[∃ *e* (*e* ≤ *a*) ∧ (DONE *e*) ∧ ¬(DONE *a*)]

2.1 REQUEST

We use the single-agent version of the definition of the REQUEST performative that is defined in [7]. Here, α is the entity performing the REQUEST, γ is the intended recipient (the intended actor), *e* is the event of performing the REQUEST, *a* is the action to be done, *q* is a relativizing condition, and *t* is the time point of the utterance.

Definition 8. *REQUEST*

(REQUEST α γ *e a q t*) ≡ (ATTEMPT α *e* ϕ ψ *t*)
where ϕ = (DONE γ *a*) ∧
[PWAG γ α (DONE γ *a*) (PWAG α γ (DONE γ *a*) *q*)]
and ψ = [BMB γ α (BEFORE *e*
[GOAL α
(AFTER *e*
[PWAG α γ ϕ *q*])])]

Intuitively, this definition says that in making a request of addressee γ, the requestor α is trying to get γ to do the action *a*, and to form the commitment to do *a* relative to the requester's commitment that it do it. More formally, by substituting for ϕ and ψ in the definition of ATTEMPT (Definition 3), we obtain the goal and the intention of the REQUEST respectively. The goal of REQUEST is that the intended actor γ eventually do the action *a* and also have a PWAG with respect to the requester α to do *a*. The intended actor's PWAG is with respect to the requester's PWAG (towards γ) that γ does the action *a*. The requester's PWAG is itself relative to some higher-level goal *q*. The intention of REQUEST is that the recipient γ believes there is a mutual belief between the recipient and the requester that before performing the REQUEST the requester α had a goal that, after performing the REQUEST, α will have a PWAG with respect to the intended actor γ about the goal ϕ of the request.

2.2 INFORM

The definition of the INFORM speech act that we use within this paper is shown below (derived from [12]). Here, α is the entity performing the INFORM, γ is the intended recipient, *e* is the event of performing the INFORM, *p* is the proposition being informed, and *t* is the time point of the utterance.

Definition 9. *INFORM*

(INFORM α γ *e p t*) ≡ (ATTEMPT α *e* ϕ ψ *t*)
where ϕ = [BMB γ α *p*]

and ψ = [BMB γ α
(BEFORE *e*
[GOAL α
(AFTER *e*
[BEL γ
(BEFORE *e*
[BEL α *p*])])])]

In the definition of INFORM, the sender α has the goal that the intended recipient γ come to believe that there is mutual belief about *p*. The intention of INFORM is that the recipient γ believes there is a mutual belief between the recipient and the informer that before performing the INFORM, the informer α had a goal that after performing the INFORM, the intended recipient γ would believe that, before performing the INFORM, α believed proposition *p*.

3 PROXYING

The notion of proxying involves one entity's asking another entity to do something on its behalf and typically taking responsibility for the action it was asked to do. For this paper, we define two speech acts, PROXY and PROXY-WEAK, that facilitate agents asking other agents to perform speech acts on their behalf. As we will show, the PROXY speech act imposes significant commitments upon the intermediate agent, while the PROXY-WEAK speech act greatly reduces the burden placed upon this proxying agent. Both PROXY and PROXY-WEAK are speech acts based upon REQUEST that take a speech act as an argument – an example of composability of speech acts.

3.1 NOTATION

In the definitions to follow, we need to specify how rewriting occurs for embedded speech acts so we first introduce some notation. In our discussions, we use the following schematic variables: *sact* ranges over speech act types, α is the performer of the speech act, γ is the intended recipient of the speech act, δ is the intended recipient of the speech act performed by the middle agent, *e* is an event type, *a* is an action, *q* is a relativizing condition, and *t* refers to a time point.

We use a parameter substitution function that, when applied to a speech act, replaces all occurrences of the schematic variable representing the specified speech act parameter by the given value. For the speech acts defined within this paper, we use the following abbreviations for speech act parameters: sender (s), intended-recipient (i), distribution (final) recipient (d), event (e), action (a), proposition (p), constraint condition (c), relativizing condition (q), and time (t).

For example, if

sact = (INFORM α γ *e on-vacation*(α) *t*)

we can specify a new speech act *sact*' using our substitution function

sact' = (*sact* s/γ i/δ e/*e*' t/*t*')

which represents the original INFORM speech act with all occurrences of the sender parameter replaced by γ, all occurrences of the intended recipient parameter replaced by δ, etc. The parameter substitution function is position independent since

the same parameter may occur in different positions in the various speech acts' parameter lists. In such an expression, all unreferenced speech act parameters are left unchanged. In the definitions below in which we use this function, substitution is performed automatically and parameters specified within the embedded speech act do not need to be specified by the uttering agent.

3.2 PROXY

When performing a PROXY speech act, the sender α wants the intended recipient γ to perform the embedded speech act to δ. This descriptive definition is very similar to that for the FIPA PROXY communicative act [5], which states, "The sender wants the receiver to select target agents denoted by a given description and to perform an embedded message to them."[2] The PROXY speech act lends itself to the deployment of middle agents within multiagent domains that can be a fully responsible proxy for other agents. As we will show, we have defined PROXY in such a manner that the individual proxying the embedded communicative act must conform to its logical preconditions.

In the definition of PROXY below, *sact* is an embedded speech act that α wants γ to perform to δ and *c* is a constraint condition for distributing the embedded communicative act (*e.g.* a time deadline).

Definition 10. *PROXY*

(PROXY α γ *e* δ *c* (*sact* s/γ i/δ) *q t*) ≡
(REQUEST α γ *e* (*c*?; (*sact* s/γ i/δ)) *q t*)

PROXY is defined as a request by the sender α for an intermediary entity γ to perform a specified speech act to a final target entity δ if the condition c is met. *sact* is any speech act but the performer of *sact* will be γ and the final recipient will be δ (these substitutions are shown explicitly above but do not need to be performed by the original utterer, α). In the definition above, *e* is the event of the PROXY action at time *t* and *q* is the relativizing condition of the embedded REQUEST.

An example of PROXY in use is shown below (assume a FIRE performative with intuitively defined semantics – by its utterance, the receiver is fired), where Nixon tells Richardson to fire Cox. If Richardson honors Nixon's PROXY, Cox will be fired and Richardson will be the party responsible for Cox's firing.

(PROXY
s/Nixon
i/Richardson
e/*e*
d/Cox
(FIRE

---

[2] We reproduce FIPA's definition here for convenience but we do not have space in which to explain all of the syntax:

<*i*, proxy(*j*, *Ref x* δ(*x*), <*j*, *cact*>, ϕ)> ≡
<*i*, inform(*j*, $I_i((\exists y)(B_j$ (*Ref x* $\delta(x) = y) \wedge$ Done(<*j*, *cact*(y)>, $B_j \phi$))))>
FP : $B_i \alpha \wedge \neg B_i$ ($Bif_j \alpha \vee Uif_j \alpha$)
RE : $B_j \alpha$

where $\alpha = I_i((\exists y)$ (B_j (*Ref x* $\delta(x) = y) \wedge$ Done(<*j*, *cact*(y)>, $B_j \phi$)))

s/Richardson
i/Cox
e/*e*'
t/t'
)
q/true
t/t
)

We can now establish the following results about the mental states of the middle agent.

Theorem 1a: After the middle agent γ honors a PROXY of a REQUEST to do action *a*, it becomes committed to the final recipient doing action *a*. Formally,

|= (DONE (PROXY α γ *e* δ *c* SACT *q t*);*c*?;SACT)
∧ (SINCERE γ δ [PWAG γ δ ϕ *q*])
⊃ (PGOAL γ (DONE δ *a*) *Q*)

where,

SACT = (REQUEST γ δ *e*' *a q*' *t*'),
Q is the relativizing condition defined below,
and ϕ is the goal of the PROXY (definitions 10, 8).

Proof sketch: The middle agent γ has just honored the proxy. Therefore, from the antecedent, (DONE (REQUEST γ δ *e*' *a q*' *t*')) is true. From the definition of REQUEST as an ATTEMPT (Definition 8), the intention of this REQUEST is ψ, where

ψ = [BMB δ γ
(BEFORE *e*'
[GOAL γ
(AFTER *e*'
[PWAG γ δ ϕ' *q*])])]

ϕ' = (DONE δ *a*) ∧ *P*,

and P represents the PWAG conjunct in ϕ above of PROXY (Definition 8, 10).

From the definition of ATTEMPT (Definition 3), we see that (INTEND γ t?;*e*';ψ? …) must have been true just before γ did the REQUEST action. In other words, γ must have had an intention to bring about a BMB between the recipient and itself that before γ made the request γ had the goal that after the REQUEST is done, it will have a PWAG with the final recipient δ about ϕ'. Assuming that agents are sincere in their communication, γ must have the PWAG with final recipient δ about ϕ' after it does the REQUEST action because sincere agents cannot intend to bring about a BMB about a proposition they believed to be false. Therefore, (PWAG γ δ ϕ' *q*) is true after the REQUEST *e*' is done. Since γ has just done the REQUEST action, it does not yet believe that the final recipient δ has done the action *a*. That is, ¬(BEL γ (DONE δ *a*) ∧ *P*) is true. Therefore, from the definition of PWAG (Definition 4), we see that the first disjunct

[¬(BEL γ *p*) ∧ (PGOAL γ *p*)]

is true, where

p = (DONE δ *a*) ∧ *P*

Substituting for p in the PGOAL conjunct above, we get

$(\text{PGOAL } \gamma\,(\text{DONE } \delta\ a) \wedge P)$

By definition, if an agent is committed to the conjunction $p1 \wedge p2$, it must be committed to each of $p1$ and $p2$ relativized to the original commitment. Therefore,

$(\text{PGOAL } \gamma\,(\text{DONE } \delta\ a) \wedge P) \supset (\text{PGOAL } \gamma\,(\text{DONE } \delta\ a)\ Q)$
where, $Q = (\text{PGOAL } \gamma\,(\text{DONE } \delta\ a) \wedge P)$

This proves the desired result. □

The ramification of Theorem 1a is that the middle agent γ of a PROXY personally acquires not only a commitment towards α to perform the embedded REQUEST, but a commitment towards the final agent δ as well. This imposes a significant responsibility upon the middle agent and is not something that all middle agents will wish to accept. Later in this paper we will introduce the PROXY-WEAK speech act that greatly reduces the responsibilities of the middle agent in the case of embedded REQUEST actions.

Theorem 1b: Just before middle agents honor a PROXY of an INFORM for some proposition p, they are required to believe p. Formally,

$|= (\text{DONE } (\text{PROXY } \alpha\ \gamma\ e\ \delta\ c\ \text{SACT}\ q\ t);c?;\text{SACT})$
$\wedge\ (\text{SINCERE } \gamma\ \delta\ p)$
$\supset\ (\text{BEFORE } e'\ [\text{BEL } \gamma\ p])$

where,
$\text{SACT} = (\text{INFORM } \gamma\ \delta\ e'\ p\ t')$

Proof sketch: We use similar arguments as in the proof of theorem 1a. γ has just honored the PROXY by performing the embedded INFORM. From the definition of INFORM as an ATTEMPT (Definition 9), the intention part of INFORM is

$\psi =$ [BMB δ γ
(BEFORE e'
[GOAL γ
(AFTER e'
[BEL δ
(BEFORE e'
[BEL γ p])])])]

Since the INFORM has just been done, the middle agent γ must have had the intention to bring about BMB that γ believed p before performing the INFORM. Therefore, by the sincerity assumption, γ must have believed p (i.e., (BEFORE e' [BEL γ p]) is true). This proves the desired result. □

Honoring a PROXY of an embedded INFORM imposes a significant responsibility upon the middle agent such that if the middle agent cannot verify the proposition's truth value it simply cannot honor the PROXY from α. While there are many situations where the middle agents can satisfy such a strong requirement there are also many situations where the middle agent should not be forced and cannot be expected to believe the embedded proposition. We believe that we can accommodate both situations with our semantics. The PROXY-WEAK speech act that we introduce below provides more options for the middle agent as it removes this strong responsibility to believe the proposition.

3.3 PROXY-WEAK

Next, we define a form of proxy that we call PROXY-WEAK that removes the "strong" requirement of precondition conformance upon the intermediate agent γ and, upon satisfaction and successful performance, provides third-party speech act semantics. Unlike PROXY of a REQUEST to do an action *a*, the PROXY-WEAK of a REQUEST should not commit the middle agent to the final recipient doing action *a*. And, unlike PROXY of an INFORM for proposition *p*, the PROXY-WEAK of an INFORM should not require the middle agent to believe *p*.

Perhaps most importantly, PROXY-WEAK should support the requirements of a third-party performative [3] – the successful execution of the PROXY-WEAK and subsequent embedded speech act should be equivalent to the sender's performing a speech act directly to the final agent, *even when going through the proxy*. We start with a definition corresponding essentially to FIPA's definition of PROXY-WEAK [5][3] and, after finding that it has significant limitations, define a version that we believe captures the key missing aspects.

Definition 11a. *PROXY-WEAK* (incorrect)

(PROXY-WEAK α γ *e* δ *c* sact *q t*) ≡
  (REQUEST α γ *e*
    [*c*?;(INFORM γ δ *e*' (GOAL α ◇(DONE γ sact)) *t*')] *q t*)

for sender α, intended (proxying) recipient γ, event *e*, final target δ, condition *c*, relativizing condition *q*, time *t*. Furthermore, sact may be any speech act, but the sender will be γ and the final recipient will be δ, i.e. sact = (*sact* s/γ i/δ).

Essentially, this definition of PROXY-WEAK has the originating agent α saying to the intermediate agent γ, "When *c* is true, perform the INFORM to δ regarding my wanting you to perform the indicated speech act". The middle agent γ then is supposed to say to δ, "α wants me to do *sact* to you." Since the middle agent γ always perform an INFORM in honoring PROXY-WEAK, from Theorem 1b the following is true:

(BEFORE *e*'
  [BEL γ
    (GOAL α
      ◇[DONE γ (*sact* s/γ r/β i/δ)])])

That is, before performing the INFORM to δ, γ must believe that α wanted it to perform *sact*. However, by performing a PROXY-WEAK, the goal of α was *not* that γ does *sact*, but rather that γ perform an INFORM regarding the *sact*. This misrepresents α's goals to δ and is therefore incorrect.

---

[3] We reproduce FIPA's definition here for convenience but we do not have space in which to explain all of the syntax:

<*i*, proxy(*j*, *Ref x* δ(*x*), <j inform<y, I_i Done(<*i*, cact(y)>)>, ϕ)> ≡

<*i*, inform(*j*, I_i((∃y)(B_j (*Ref x* δ(*x*) = y) ∧ Done(<j inform<y, I_i Done(<*i*, cact(y)>)>, B_j ϕ))))>

FP : B_i α ∧ ¬B_i (Bif_j α ∨ Uif_j α)

RE : B_j α

where α= I_i((∃y) (B_j (*Ref x* δ(*x*) = y) ∧ Done(<j inform<y, I_i Done(<*i*, cact(y)>)>, B_j ϕ)))

The above definition also does not result in performance of a third-party performative by the proxying agent. To illustrate this point, consider the Nixon example given earlier. Suppose Nixon (α) performs the equivalent of a PROXY-WEAK to Richardson (γ) with *sact* being the performative for 'fire' and Cox being the target agent (δ). According to the above definition of PROXY-WEAK, Richardson can satisfy Nixon's PROXY-WEAK by performing an INFORM to Cox corresponding in natural language to Richardson's saying to Cox, "Nixon wants me to fire you". However, this INFORM does not result in Cox's getting fired by Nixon. The key here is that performatives are accomplished in virtue of their being uttered and here Richardson's utterance does not result in 'fire' being performed by Nixon. The next definition addresses this limitation.

Definition 11b. *PROXY-WEAK*

(PROXY-WEAK $\alpha\ \gamma\ e\ \delta\ c$ *sact* $q\ t$)

$\equiv$ (REQUEST $\alpha\ \gamma\ e$ [c?;(INFORM $\gamma\ \delta\ e'\ \theta\ t'$)] $q\ t$)

where,

θ = (HAPPENING sact) and

sact = (*sact* s/α i/δ e/$e;e'$ t/t')

In other words, PROXY-WEAK of a speech act *sact* is a REQUEST to INFORM that the speech act *sact* is happening using the two acts – the sender's and the intermediary's. The two actions are $e;e'$, where e' is the very act of informing this fact – hence γ's act of performing the INFORM also completes α's speech act to δ.

Using this definition of PROXY-WEAK, Richardson *will* satisfy Nixon's PROXY-WEAK by saying, in natural language, "Nixon hereby fires you". Here 'fires' is used as a third party performative – it is a performative because saying so in the right situation makes it so. We note that by the definition of PROXY-WEAK as a REQUEST, when the middle agent γ accepts the REQUEST, it has a PWAG with sender α about performing the INFORM act with respect to the sender's PWAG that γ does the INFORM. From the definition of PWAG (Definition 4), γ will establish a mutual belief to that effect after performing the requested INFORM. The PROXY-WEAK is discharged successfully when this mutual belief is established. This is evident from the next two theorems.

Theorem 2a: When the middle agent successfully discharges a PROXY-WEAK performed to it, the original sender believes that it has performed the embedded speech act to the target even though it may not have observed the middle agent's act directly and only knows that it was done. Formally,

$\models$ [DONE ([PROXY-WEAK $\alpha\ \gamma\ e\ \delta\ c$ *SACT* $q\ t$];
[MB $\alpha\ \gamma$ (DONE [INFORM $\gamma\ \delta\ e'$ (HAPPENING *SACT*) t'])]?
)]

$\wedge$ (SINCERE $\alpha\ \gamma$ [PWAG $\alpha\ \gamma\ \phi\ q$])

$\supset$ (BEL α (DONE α *SACT*))

where,

SACT = (*sact* s/α i/δ e/$e;e'$ t/t') and

ϕ is the goal of the PROXY-WEAK (Definitions 11b, 8).

Proof sketch: By performing a PROXY-WEAK, the sender α requested the middle agent γ to inform the final recipient δ that (HAPPENING *SACT*). (1) From the usual assumption of sincerity, α cannot make that REQUEST unless α believes that p. (2)

When γ establishes the mutual belief that the INFORM has been done, α believes that the event *e*' (i.e. the INFORM event) has been done. From (1) and (2), α believes that the event *sequence e;e*' is happening and also believes that the event *e*' has just been done. So α believes that the event sequence *e;e*' has just been done and hence believes that the action *SACT* represented by the event sequence *e;e*' has just been done. This establishes the desired result. □

Theorem 2b: When the middle agent satisfies a PROXY-WEAK performed to it, the final recipient will come to believe that the original sender has performed the embedded speech act to it provided that it trusts the middle agent. Formally,

|= [DONE (PROXY-WEAK α γ *e* δ *c SACT q t*);*c*?;(INFORM γ δ *e*' *p t*')]
  ∧ (TRUSTS δ γ *p*)
  ∧ (SINCERE γ δ *p*)
  ⊃ (BEL δ (DONE *SACT*))
where
  SACT = (*sact* s/α i/δ e/*e;e*' t/*t*') and
  p = (HAPPENING *SACT*)

Proof sketch: By assumption of sincerity, the middle agent γ believes the proposition being informed. The final recipient δ trusts the middle agent γ. (1) Therefore, δ also believes the proposition being informed i.e. δ believes that *p*. (2) The final recipient δ has just received the INFORM from the middle agent. Therefore, δ believes that the event *e*' (i.e. the INFORM event) has been done. From (1) and (2), δ believes that the event sequence *e;e*' is happening and also believes that the event *e*' has just been done. So it believes that the event sequence *e;e*' has just been done and hence believes the action *SACT* represented by the event sequence *e;e*' has just been done. This establishes the desired result. □

Theorem 3a: After a middle agent honors a PROXY-WEAK of a REQUEST to do action *a*, it does *not* become committed to the final recipient doing action *a*. Formally,

|≠ [DONE (PROXY-WEAK α γ *e* δ *c* SACT *q t*);*c*?;(INFORM γ δ *e*' θ *t*'))
  ⊃ (PGOAL γ (DONE δ *a*) *Q*)
where,
  SACT = (REQUEST α δ *e*' *a q*' *t*'),
  θ = (HAPPENING SACT), and
  Q is a relativizing condition.

Proof sketch: From definition 11b, note that the middle agent γ performs an INFORM • γ never performs the embedded speech act no matter what it is. Therefore, when *sact* is a REQUEST, γ does not have the goal and intentions of (REQUEST α δ *e*' *a q*' *t*'). In particular, from the definition of INFORM (Definition 9), γ does not have a PWAG with δ for doing *a* and hence is not committed to δ doing *a*. The relativizing condition *Q* does not come into play and, without loss of generality, may have any value other than false. □

Theorem 3b: After the middle agent honors a PROXY-WEAK of an INFORM for some proposition *p*, it is *not* required to have believed *p*. Formally,

|≠ [DONE (PROXY-WEAK α γ *e* δ *c* SACT *q t*);*c*?;(INFORM γ δ *e*' θ *t*')]
  ⊃ (BEFORE *e*' [BEL γ *p*])
where,

SACT = (INFORM α δ *e'* *p* *t'*) and
θ = (HAPPENING SACT)

Proof sketch: Similar to the proof of Theorem 3a where *sact* is an INFORM. In this case, the middle agent γ performs an INFORM with the propositional content of (HAPPENING *sact*) rather than an INFORM with the propositional content *p* directly. The middle agent γ therefore must believe (HAPPENING *sact*) just prior to honoring the PROXY-WEAK, but not necessarily *p*. □

4. Discussion

Our analysis of the semantics of PROXY and PROXY-WEAK above shows that middle agents can be deployed in multiagent systems and flexibly accommodate a wide range of domains and scenarios, some in which the middle agents must take full responsibility for their actions and some in which the middle agents act more as if they were simple couriers. Because of the strong semantic definitions involved, an agent faced with a decision to perform a PROXY or a PROXY-WEAK speech act can reason about the burden placed upon the middle agent and choose between them knowingly. Similarly, a middle agent receiving one of these speech acts can reason about the level of responsibility expected of it and make a knowledgeable decision about whether to honor the PROXY or PROXY-WEAK. We can demonstrate through a set of proofs similar to that used in [7] that these semantics will also hold for groups of agents as targets of PROXY, PROXY-WEAK, and the embedded speech acts.

Prior work on agent communication languages (e.g., [4][5][9][10]) either lack support for middle agent speech acts or lack the strength and depth of semantics as we have introduced above. Furthermore, there has been no work on agent communication languages that has successfully defined a speech act that supports third-party semantics until now. The prior work most similar to that presented within this paper has been performed by FIPA [5]. The FIPA standards body has defined proxy communicative acts with an intent similar to ours [5]. There are several significant differences between our approach and that of FIPA, however. First, FIPA's PROXY is defined as an INFORM between the originating agent and the middle agent, while ours is defined using a REQUEST to the middle agents. This is significant in that the middle agent within FIPA need not be expected to do anything, while in our definition the middle agent is expected to perform a subsequent speech act if it agrees to honor the REQUEST. Second, FIPA defines their equivalent of PROXY-WEAK in terms of an INFORM of an intent by the original sender to have the middle agent do the embedded communicative act. Because the middle agent will only ever perform an INFORM and never the embedded speech act directly, the FIPA definition therefore misrepresents the sending agent's intentions to the target agent (see the discussion of Definition 11a).

In summary, the PROXY and PROXY-WEAK communicative acts defined in this paper provide speech acts that support agents interacting with middle agents that can act on their behalf. Our analysis has shown that the two acts result in the middle agents having significantly different levels of commitments relative to the final group, where PROXY imposes significant and PROXY-WEAK imposes very little responsibility upon the middle agents. We also have shown that the PROXY-WEAK speech act results in the correct embodiment of third party performative

semantics, where we obtain the equivalence of the sending agents performing a speech act directly on the final target agents even while going through proxies.

Acknowledgement

We gratefully acknowledge the DARPA CoABS Program (contract F30602-98-2-0098, A0 G352) for supporting the research presented in this paper. Pacific Northwest National Laboratory is operated by Battelle for the U.S. Department of Energy under contract DE-AC06-76RLO 1830. We would also like to thank the anonymous reviewers for their excellent comments.

References

[1] Cohen, P.R. On Knowing What to Say: Planning Speech Acts. *Ph.D. Thesis*, Department of Computer Science, University of Toronto, 1978.

[2] Cohen, P.R., and Levesque, H.J. Intention is Choice with Commitment, *Artificial Intelligence*, 42(3), 1990.

[3] Cohen, P.R., and Levesque, H.J. Performatives in a Rationally Based Speech Act Theory. *Proc. of the 28th Annual Meeting of the Association for Computation Linguistics*, 1990.

[4] Finin, T., and Labrou, Y. A proposal for a new KQML Specification, TR CS-97-03, CSEE Dept., UMBC, Baltimore, MD, 1997.

[5] FIPA ACL Specification. http://www.fipa.org/, 2000.

[6] Katagiri Y. Belief Coordination by Default. *Proceedings of the 2nd International Conference on Multi-Agent Systems*, pages 142-149, 1996.

[7] Kumar, S., Huber, M.J., McGee D.R., Cohen, P.R., and Levesque, H.J. Semantics of Agent Communication Languages for Group Interaction. *Proceedings of the 16th National Conference on Artificial Intelligence (AAAI)*, pages 42-47, 2000.

[8] Martin, D.L., Cheyer, A.J., and Moran, D.B., The Open Agent Architecture: A Framework for Building Distributed Software Systems, *Applied Artificial Intelligence*, vol. 13, pages 91-128, 1999.

[9] Pitt J. and Mamdani. A.. Designing Agent Communication Languages for Multi-Agent Systems. In F. Garijo and M. Boman (eds.): *Multi-Agent System Engineering: Proceedings MAAMAW'99*, LNAI1647, Springer-Verlag, pp102–114, 1999.

[10] Singh, M.P. A Social Semantics for Agent Communication Languages. *Proceedings of the IJCAI 2000 Workshop on Agent Communication Languages*, 2000.

[11] Smith, I.A., and Cohen, P.R. Toward a Semantics for an Agent Communications Language Based on Speech-Acts. *Proceedings of the 13th National Conference on Artificial Intelligence*, 1996.

[12] Smith, I.A., Cohen, P.R., Bradshaw, J. M., Greaves, M., and Holmback, H. Designing Conversation Policies Using Joint Intention Theory. *Proceedings of the 3rd International Conference on Multi-Agent Systems*, 1998.