Toward A Suite of Performatives Based upon Joint Intention Theory

Marcus J. Huber[1], Sanjeev Kumar[2], David McGee[3]

[1]Intelligent Reasoning Systems, San Diego, CA
[2]Oregon Health and Science University, Beaverton, OR
[3]Natural Interaction Systems, Portland, OR
marcush@marcush.net, skumar@cse.ogi.edu,
dmcgee@naturalinteraction.com

Abstract. Agent communication languages defined using joint intention theory have enjoyed a long research history. A number of performatives have been defined and refined in this literature with particular emphasis on the basic performatives of REQUEST and INFORM, which subsequently have many subtle versions. Even these less common performatives have been extended and refined multiple times. In many cases the underlying definitions upon which the various performatives are based have been modified as well. While working toward implementing a multi-agent system with communications based upon joint intention semantics, it quickly became apparent that it was going to be difficult to identify a single set of performatives with correct and compatible definitions. We also realized that a set of performatives with enough breadth to cover the needs of real fielded multi-agent systems has not yet been defined. We intend this paper to provide in a single place a broadly applicable set of compatible performatives defined using joint intention semantics. Many of the performatives previously defined in the literature have been brought to the same semantic basis while we have also defined a number of new performatives to increase the breadth of performatives available to agent developers.

1   Introduction

This paper defines a broad range of Agent Communication Language (ACL) performatives based upon the semantics of joint intention (JI) theory [2,4,5,6,7,10,11,13,14,15,20,25]. A great deal of prior research on specifying "speech acts" [27] for agent communication languages based on joint intention theory has already been performed [2,4,5,6,7,12,16,17,20,28,29].

Researchers fielding multi-agent systems using joint intention theory and performatives based upon it [13,14,15,24,31,33] run into several issues however, due primarily because joint intention theory has such a rich research history:

- Performative definitions are spread throughout the literature, with no single paper to refer to.
- Base-level semantic definitions related to joint intentions have changed slowly over time in the research, both in its semantics and in its notation (as limitations are eliminated or extensions made).
- Performative definitions have changed to match changes in these basic definitions, but not all performatives previously defined are updated with each underlying definition change.
- Performative definitions have been modified over time even when the underlying semantic definitions have remained constant (again, ostensibly to remove limitations, make extensions, etc.)

- The complexity of the definitions may sometimes cause confusion when selecting the correct performative in multi-message exchanges.
- Not all of the performatives that might be considered necessary for fielding a multi-agent system have been defined, notably "utility" performatives, both those implicitly required by joint intention theory and those not so required but found to be useful when fielding systems based on ACLs with other semantics [8,19,34].

To address these issues, we extensively reviewed the semantic definitions of performatives based on Joint Intention theory, particularly those based on the work of Cohen and Levesque, and defined a broad range of performatives with a single, consistent, unified semantic basis with more explicit historical context. This has the benefit of 1) a single semantic basis for all of the performatives, 2) a single place to refer to performative definitions, 3) a much broader selection of performatives than has been specified to date, and 4) more direct and clearer applicability to a much broader set of singe and multi-message interactions than the performatives defined to date.

The performative defined in the paper fall into the following general categories:
- Core: INFORM, REQUEST, SHOUT
- Specialized Core: SUBSCRIBE, QUERY, PROXY, PROXY-WEAK, STANDINGOFFER
- Team-oriented: AGREE, REFUSE, CANCEL, FAILURE, ACCEPT, REJECT, WITHDRAW, ORDER
- Utility: IMPOSSIBLE, RELFAIL, SUCCESS, ACKNOWLEDGE, NOTUNDERSTOOD

Core performatives are the most basic performatives, defined directly using the basic definition of ATTEMPT. The Core performatives are those that all of the other performatives are based, either using specialization or composition. The Specialized Core performatives are Core performatives that have particular message content expressions. Team-oriented performatives exploit the semantics of Joint Intention theory to form and dissolve teams of agents under various circumstances. Utility performatives are also specializations of Core performatives, but are in general used in service of the semantic ramifications of the other performatives. This categorization of the performatives is arbitrary to some extent and we could have categorized them differently. For example, the performatives in the Utility category could be considered Specialized or Team-oriented. Nevertheless, we have found it useful to group them as shown above.

2 Agent Communication Language Components

2.1 Basic Semantic Notation and Definitions

A summary of the basic notation and definitions follows. Full details of this modal language, including formal models, are beyond the scope of this paper, but can be found in [3] and [4].

e, e′, etc. are events

a, a′, etc. are actions (complex event expressions)

p, q, etc. are propositions (where q, q′, etc. are used as relativizing conditions)

t, t′, t_1, t_2, etc. are time points

x, y, z are agents

• = always

• = eventually

e < e′ says that e occurs before e′
(BEL x p) say that p follows from x's beliefs.
(GOAL x p) say that p follows from x's choices.
(HAPPENS a), (HAPPENED a), and (DONE a) say that a sequence of actions described by the action expression a will happen next, happened sometime in the past, or has just happened, respectively. (HAPPENS x a) and (DONE x a) also specify the agent for the action sequence that is going to happen or has just happened.
(UNTIL p q) says that q holds up to the time that p becomes true.
BEFORE and AFTER are defined using HAPPENS and DONE.
(EARLIER p) ≡ ¬p ∧ ∃e (HAPPENED p?;e)

An action expression is built from variables ranging over sequences of events using constructs of dynamic logic: a;b is action composition, p? is a test action, | specifies non-deterministic choice, and ‖ indicates concurrent actions. Mutual belief is defined in terms of unilateral mutual belief (BMB) [16].

In some of the definitions to follow, we need to specify how rewriting occurs for embedded speech. We use a parameter substitution function $subst_{perf}$ such that $subst_{perf}$(param/val) replaces all occurrences of the schematic variable param representing a specified parameter of performative *perf* by the given value val. For the speech acts defined within this paper, we use the following abbreviations for speech act parameters: speech act (*sa*), sender (*s*), intended-recipient (*i*), distribution (final) recipient (*d*), event (*e*), action (*a*), proposition (*p*), constraint condition (*c*), relativizing condition (*q*), and time (*t*). All unreferenced speech act parameters are left unchanged.

For example, if sact = (INFORM x y e on-vacation(x) t), a fully specified speech act, we can specify a new speech act sact' using our substitution function

sact' = $subst_{sact}$(s/y i/z e/e' t/t')

which represents an INFORM speech act with all occurrences of the sender parameter replaced by y, all occurrences of the intended recipient parameter replaced by z, etc.

Definition: HAPPENING

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

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 [13].

Definition: PGOAL (Persistent Goal)

(PGOAL x p q) ≡ (BEL x ¬p) ∧ (GOAL x •p) ∧
                (UNTIL [(BEL x p) ∨
                    (BEL x •¬p) ∨
                    (BEL x ¬q)]
                GOAL x •p)

A persistent relativized goal formalizes the notion of commitment. An agent x having a persistent goal p is committed to that goal. The agent x 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, is impossible, or the relativizing condition q is untrue [3]. Note that we assume that agents are competent with respect to their commitments [18].

Definition: INTEND (an action)

(INTEND x a q) ≡ (PGOAL x
         [HAPPENS x
            (BEL x (HAPPENS a))?;a] q)

Intention to do an action a is a commitment to do the action knowingly. The agent x 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 [3].

Definition: ATTEMPT

(ATTEMPT x e   t) ≡
    t?;[(BEL x ¬) ∧
      (GOAL x (HAPPENS e;• ?)) ∧
      (INTEND x t?;e; ?(GOAL x (HAPPENS e;• ?)))]?;e

An attempt to achieve via is a complex action expression in which the agent x is the actor of event e at time t and, just prior to e, the actor chooses that should eventually become true and intends that e should produce relative to that choice. So represents some ultimate goal that may or may not be achieved by the attempt, while represents what it takes to make an honest effort [4].

Definition: PWAG (Persistent Weak Achievement Goal)

(PWAG x y p q) ≡
  [¬(BEL x p) ∧ (PGOAL x p q)] ∨
  [(BEL x p) ∧ (PGOAL x (MB x y p) q)] ∨
  [(BEL x •¬p) ∧ (PGOAL x (MB x y •¬p) q)] ∨
  [(BEL x ¬q) ∧ (PGOAL x (MB x y ¬q))]

This definition, recently revised in [17], states that an agent x has a PWAG with respect to another agent y when the following holds: (1) if agent x believes that p is not 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 agent y. A PWAG expression is often used in the performatives below that are intended to create a joint, or social, commitment.

2.2 ACL Performative Semantics and Descriptions

Definition: Inform

(INFORM x y e p t) ≡ (ATTEMPT x e   t)
 where
    = (BMB y x p)
 and
    = (BMB y x
         (BEFORE e
           [GOAL x
             (AFTER e
               (BEL y
                 [BEFORE e
                   (BEL x p)
                 ]))
            ]))

In this performative, the sender x has the goal that the intended recipient y come to believe that there is mutual belief that y believes p. The intention of INFORM is that

the y comes to believe there is a mutual belief between the y and x that before sending the INFORM, x had a goal that after sending the INFORM the intended recipient y would believe that, before sending the INFORM, x believed proposition p (most recently from [17]).

The INFORM performative is a general-purpose speech act suitable for any communication related to propositional belief. Many of the following speech acts specialize this speech act or us it in more complex action expressions.

Definition: Request

(REQUEST x y e a q t) ≡ (ATTEMPT x e   t)
  where
    = (DONE y a) ∧ [PWAG y x (DONE y a)
                              (PWAG x y (DONE y a) q)∧q]
  and
    = (BMB y x (BEFORE e
             [GOAL x
              (AFTER e
                (BEL y [PWAG x y   q]))]))

We use the single-agent version of the definition of the REQUEST performative that is defined in [16] and later refined in [17]. Intuitively, this definition says that in making a request of addressee y, the requestor x is trying to get y to do the action a, and to form the commitment to do a relative to the context q and the requester's commitment that it do it.

The goal of the requester is that the intended actor y eventually do the action a and have a PWAG with respect to the requester x to do a. The intended actor's PWAG is with respect to (i.e., relativized to) the requester's PWAG (towards y) that y does the action a and also with respect to q. The requester's PWAG is itself relative to some higher-level goal q.

The intention of REQUEST is that the recipient y believe there is a mutual belief between the recipient and the requester that before performing the REQUEST, the requester x had a goal that after performing it, the requester x will have a PWAG with respect to the intended actor y about the goal   of the request.

The REQUEST performative is a general-purpose speech act useful whenever one agent wants another agent to do something. Many of the following speech acts specialize this speech act or us it in more complex action expressions.

Definition: Subscribe

(SUBSCRIBE x y e $\prod$ q t) ≡ (REQUEST x y e α q t)
where,
  α = •(∀t_1,t_2,t_3.(t < t_1< t_2 < t_3).∀e′,∀τ.
      (DONE (t_1?;¬$\prod$(τ)?;t_2?;$\prod$(τ)?;(INFORM y x e′ $\prod$(τ) t_3) |
            t_1?; $\prod$(τ)?;t_2?;¬$\prod$(τ)?;(INFORM y x e′ ¬$\prod$(τ) t_3))))?
  where τ represents a grounded object and
  $\prod$ is a unary predicate, indicating the information of interest to the sender, that evaluates to true or false for objects (τ).

A new JI performative, a SUBSCRIBE is a REQUEST for the sender to INFORM the recipient whenever the indicated predicate changes state for any objects known to y that satisfy $\prod$. Unless initially REFUSEd or subsequently CANCELed (see below), the recipient y should INFORM $\prod$(τ) or ¬$\prod$(τ) whenever the truth value of the predicate changes from false to true or true to false, respectively.

The SUBSCRIBE performative is useful in situations where an agent needs to be kept up to date on another agent's beliefs over an extended period of time. For

example, $\prod$=OnTable, y might come to inform x that OnTable(blockA), another message to inform x that ¬OnTable(blockB), etc. In the CIANC project, in which autonomous vehicles might be used for surveillance [33], $\prod$ can be the predicate HostileWithinRange(), such that an agent subscribing to a surveillance agent using that predicate will be kept advised on hostile units that come into and leave weapons range.

Definition: Query
(QUERY x y e p q t) ≡ (REQUEST x y e α q t)
  where
   α = •(∃e',t', t'>t, e'>e.(DONE (INFORM y x e′ p t') |
                              (INFORM y x e′ ¬p t')))?

In this performative, the sender asks the recipient about their belief in the truth value of a particular proposition (similar to the "yes-no question" performative in [7]). In other words, y is requested to INFORM x, at some future time with a distinct event, of its belief in the value of p. The QUERY performative is useful in situations when only the current belief value of a single proposition is required by the querying agent.

Definition: Proxy
(PROXY x y z e c sact q t) ≡
  (REQUEST x y e
            (c?;•[∃e',t'.(DONE $subst_{sact}$(s/y i/z e/e' t/t'))]?) q t)

PROXY is defined as a request by the sender x for an intermediary agent y to perform a specified speech act, sact, to a final target agent z if the condition c is met (modified from the [12], which gives details of this speech acts' 3rd party semantics). sact can be any fully-specified speech act, but the sender of sact (performed as a distinct event at a future time) will be y and the final recipient will be z. With the PROXY performative, the "middle agent" y assumes all responsibility associated with uttering the proxied performative sact to z. The PROXY performative provides agents the important capability of using intermediary agents [8,9,19,21,30,34] to do speech acts on its behalf.

Definition: Proxy-Weak
(PROXY-WEAK x y z e c sact q t) ≡
  (REQUEST x y e
            (c?;•[∃e',t'.(DONE (INFORM y z e' θ t'))]?) q t)
  where
  θ = (HAPPENING $subt_{sact}$(s/y i/z e/e;e' t/t'))

This performative is a weaker version than PROXY, where the middle agent y sends an INFORM message about sact rather than perform sact itself. This protects the middle agent from bearing the obligations associated with performing sact directly and at the same time results in a state *as if* x had performed sact directly to z (modified from [10], which gives details of this speech acts' 3rd party semantics). The PROXY-WEAK performative provides agents the important capability of using intermediary agents [8,9,19,21,30,34] to pass on requests or information to third agents.

Definition: Shout

Given that ∏ is a unary group membership predicate indicating who the intended recipients are,

(SHOUT x ∏ e p t) ≡ (ATTEMPT x e φ ψ t)
  where
    φ = ∀y. ∏(y) [BMB y x p]
  and
    ψ = ∀y. ∏(y). [BMB y x
                (BEFORE e
                  [GOAL x
                    (AFTER e
                      [BEL y
                        (BEFORE e
                          [BEL x p]
)])])]

The sending agent x holds that some proposition is true and intends that the receiving agents, specified as those satisfying the predicate ∏, also come to believe that the proposition is true. This is an extended form of INFORM to support making an utterance to multiple, incompletely-specified recipients, a capability not supported by the ACLs of KQML, Singh, FIPA, and others [8,19,27,34]. In this SHOUT performative, the sending agent is essentially performing an INFORM to each of the agents satisfying ∏. Examples of ∏ in real applications might be of the form contractors() (to represent all agents the sending agent has contracts with), inferiors() (to represent all agents that the agent considers as inferiors), platoon-members() (to represent all agents that are associated with supporting a platoon), etc.

   A group-theoretic version of this performative based on [16] would provide more flexible semantics (in that the speaker would not have to know anything about the recipients beliefs, unlike in this definition in which it eventually does) but creating a complete suite of performatives with the group-theoretic work remains for future work.

Definition: StandingOffer
(STANDINGOFFER x y e a q t) ≡ (INFORM x y e θ t)
where
  θ = ∀e′,t′.(DONE
        (INFORM y x e′ (PWAG y x (DONE x a) q)) t′) ⊃
      (DONE e′;(PWAG x y (DONE x a)
                        (PWAG y x (DONE x a) q)∧q)?)

In this performative (most recently from [17]), the sender x is making a standing offer to the recipient y to form a commitment (PWAG) to do something for the recipient in the future. This obligation on x's part will only arise, however, when y sends it an INFORM message saying that it has a commitment and wants x to adopt the commitment relative to y's. As shown in [18], the STANDINGOFFER followed by the definition's INFORM results in a Joint Intention team just as if y had sent a REQUEST and x had replied with an AGREE.

Definition: Order
Within the confines of this paper, we define the following relationship definitions for use in the ORDER performative:
OBEDIENT-TO
(OBEDIENT-TO x y a q) ≡

$\forall$e,t.[(DONE y (REQUEST y x e a q t)) $\supset$
    (PWAG y x (DONE x a) q) $\wedge$
    (PWAG x y (DONE x a)
                [PWAG y x (DONE x a) q]$\wedge$q)]

Agent x is obedient to y with respect to action a and relativizing condition q when for every REQUEST by y, x will always adopt a PWAG to do a with respect to y's PWAG that x do a. A synonym to OBEDIENT-TO is INFERIOR (i.e., if y is OBEDIENT-TO x, then y is INFERIOR to x).

AUTHORITY-OVER
(AUTHORITY-OVER x y a q) $\equiv$ (OBEDIENT-TO y x a q)
Agent x has authority over agent y with respect to a and q. A synonym to AUTHORITY-OVER is SUPERIOR (i.e., if x has AUTHORITY-OVER y, then x is SUPERIOR to y).

Then, the definition of ORDER is:

(ORDER x y e a q t) $\equiv$ (SUPERIOR x y a q)?;(REQUEST x y e a q t)

In this new performative, we extended the single-agent version of the definition of the REQUEST performative to support authority relationships, wherein the sender has some recognized authority to unilaterally task the recipient. Here, x is the agent performing the ORDER, y is the intended recipient (the intended actor), e is the event of performing the ORDER, a is the action to be done, q is a relativizing condition, and t is the time point of the utterance. An ORDER with the appropriate SUPERIOR relationship automatically results in nested persistent goals (PWAGs) as if an explicit AGREE (see below) had been performed by the recipient.

Definition: Agree
(AGREE x y e a q t) $\equiv$
   ($\exists$e$'$, t$'$.(EARLIER (DONE (REQUEST y x e$'$ a q t$'$)))?;
   (INFORM x y e (PWAG x y (DONE x a)
                 (PWAG y x (DONE x a) q)$\wedge$q) t)

This performative is an agreement to perform an action requested by another agent. This is similar to the CONFIRM performative of [29] and [17], but revised extensively to require the historical context (i.e., an earlier REQUEST). As shown in [17], this performative is sufficient to create a team with interlocking PWAGs when in response to a prior REQUEST.

Definition: Refuse
(REFUSE x y e a q t) $\equiv$
   ($\exists$e$'$, t$'$.[EARLIER
        (DONE (REQUEST y x e$'$ a q t$'$))])?;
   (INFORM x y e $\neg$[PWAG x y (DONE x a)
                (PWAG y x (DONE x a) q)$\wedge$q] t)

The sender will NOT adopt the recipient's goal to do an action from a prior request. This is similar in some respect to that of the REFUSE performative of [29] and [17]. Our formulation is more specific in the context term (a REQUEST compared to a more general PWAG expression). In addition, their formulation uses •$\neg$[PWAG…], indicating that the refusing agent would never perform the action for the other agent. However, we feel that this is not desirable as the agent that is currently refusing might accept to do the action at a later time. The definition of [28] has the same

¬[PWAG …] semantics but does not include the historical context (i.e., the EARLIER term).

Definition: Cancel
(CANCEL x y e a q t) ≡
  (∃e′,t′.[EARLIER
               (DONE (REQUEST x y e′ a q t′))])?;
  (INFORM x y e ¬(PWAG x y (DONE y a) q) t)

In the CANCEL performative, the sender no longer has the goal that the recipient perform an action (accepted in response to a prior REQUEST by the sender). This has been revised from [17] to specify REQUEST in the context rather than a PWAG expression. Not that this performative can also be used to dissolve a team in certain circumstances, such as when a REQUEST/AGREE performative sequence was performed between x and y.

Definition: Failure
(FAILURE x y e a q t) ≡
  (∃e′,t′.(EARLIER [DONE (REQUEST y x e′ a q t′)]))?;
  (INFORM x y e
         (DONE x (¬p ∧ (INTEND x ¬p?;a;p?);a;¬p?)?) t)

In this new performative, the sender tells the recipient that the action a was attempted but was not successfully completed. The DONE expression indicates that the sender intended that the action would have certain expected results but that after performing the action the results were not what was expected. This is not sufficient to terminate a team of agents; if the sender is part of a team with respect to performing action a, it will persist in pursuing the action until the team's PWAG is eventually satisfied. Even without the power to dissolve a team, this message is often useful as a status message between agents in fielded multi-agent systems.

Definition: Accept
(ACCEPT x y e a q t) ≡
  (∃e′,t′.(EARLIER [DONE (STANDINGOFFER y x e′ a q t′)]))?;
  (INFORM x y e (PWAG x y (DONE y a) q) t)

This performative is an acceptance of the other agent standing offer to perform action a (significantly modified from [17] by adding the explicit historical context term). Note that an additional difference between ACCEPT and AGREE is that the ACCEPT's PWAG is only relative to q while the AGREE's PWAG is relative to the other agent's PWAG and q.

Definition: Reject
(REJECT x y e a q t) ≡
  (∃e′,t′.(EARLIER
            [DONE (STANDINGOFFER y x e′ a q t′)]))?;
  (INFORM x y e ¬(PWAG x y (DONE y a) q) t)

This performative is the opposite of an ACCEPT, in that the sending agent tells the recipient that it will not be taking up the standing offer for y to perform action a. This has been modified significantly from [17] by adding the explicit historical context term and using ¬(PWAG…) rather than •¬(PWAG …). Note that the difference between REFUSE and REJECT is that the REJECT's PWAG is only

relative to q while the REFUSE's PWAG is relative to the other agent's PWAG *and* q.

Definition: Withdraw
(WITHDRAW x y e a q t) ≡
   (∃e′,t′.(EARLIER [DONE (STANDINGOFFER x y e′ a q t′)]))?;
   (INFORM x y e ¬θ t)
where

$$\theta = \forall e''.t''.(\text{DONE (INFORM } y \; x \; e'' \; (\text{PWAG } y \; x \; (\text{DONE } x \; a) \; q)) \; t'') \supset$$
$$(\text{DONE } e'';(\text{PWAG } x \; y \; (\text{DONE } x \; a)$$
$$(\text{PWAG } y \; x \; (\text{DONE } x \; a) \; q) \wedge q)?)$$

This performative provides the means by which an agent can remove its STANDINGOFFER (the semantics of which is represented by θ) to do an action for another agent (revised from [17]). Note that the recipient of a STANDINGOFFER need do nothing to "remove" itself from a STANDINGOFFER relationship, as it has no outstanding commitment to the original sender.

Definition: Impossible
(IMPOSSIBLE x y e a t) ≡
   (∃e′,t′.(EARLIER
           (DONE [(REQUEST x y e′ a q t′) |
                (REQUEST y x e′ a q t′)])))?;
   (INFORM x y e •¬(DONE a) t)

In this new performative, the sender tells the recipient that it is no longer possible to perform action a. Per the definition of PWAG, this is sufficient to dissolve a "team" of agents with interlocking PWAGs and will typically be used only (but often) for this purpose. Note that either agent involved in the team may send this message to satisfy the PWAG and thereby dissolve the team.

Definition: RelFail
(RELFAIL x y e q t) ≡
   (∃e′,t′.(EARLIER
           (DONE [(REQUEST x y e′ a q t′) |
                (REQUEST y x e′ a q t′)])))?;
   (INFORM x y e ¬q t)

In this new performative, the sender tells the recipient that the relativizing condition of an earlier REQUEST is no longer valid. Per the definition of PWAG, this is sufficient to dissolve a "team" of agents with interlocking PWAGs and will typically be used only (but often) for this purpose. Note that either agent involved in the team may send this message to satisfy the PWAG and thereby dissolve the team.

Definition: Success
(SUCCESS x y e a t) ≡ (INFORM x y e (DONE x a) t)

In this new performative, the sender tells the recipient that action a has been performed successfully. Per the definition of PWAG, this is sufficient to dissolve a "team" of agents with interlocking PWAGs (assuming such communication achieves a state of mutual belief). This is a very simple performative specialization, but since

it is a very common message in fielded multi-agent systems, particularly those based upon joint intention theory, the resulting simplified parsing and interpretation will result in significant time and computational resources savings over time.

Definition: Acknowledge

(ACKNOWLEDGE x y e a q t) $\equiv$
   (INFORM x y e
        ($\exists e',t'$.[EARLIER
              (DONE y
                  (REQUEST y x e' a q t'))]) t)

An agent will use this performative to acknowledge that a prior REQUEST has been made. Among other uses, this can act as a courtesy message between agents. Note that this is a new performative definition and is completely different than the ACKNOWLEDGE performative in [29]. This utility performative fills a vital role in fielded agent systems.

Definition: NotUnderstood

(NOTUNDERSTOOD x y e a p t) $\equiv$ (INFORM x y e θ t)
 where
   $\theta = \exists e',t'$.(EARLIER (DONE (sact y x e' t')) $\wedge$
             $\neg intentof$(sact y x e' t') $\wedge$
             $\neg goalof$(sact y x e' t')
   sact is any fully specified speech act,
   intentof is an operator that extracts the goal portion of sact, and
   goalof is an operator that extracts the intention portion of sact

In this new performative, the sender x informs y, the sender of an earlier message (sact), that it does not understand something about the message by expressing that it did receive the message but that neither the goal nor intention portion of the speaker's speech act was successful. Note that this performative is not yet semantically correct because, strictly speaking, the intention portion of our speech acts cannot be untrue (by definition of speech acts, they are automatically true by their utterance) but might be completed using 2^{nd}-order logic. Even though the semantics of this performative remains an area of future work, we found it useful as agent implementers to include this performative in fielded multiagent applications.

3 Discussion

In this paper, we have defined a broad number of performatives based upon a single, coherent semantics of Joint Intention theory definitions. While we have based many of the performatives upon a variant found in prior literature, most of those have received an update to the latest semantics of this paper. Many of the other performatives in this paper are completely novel.

    We recognize that agents will often interact in some form of typical or common pattern of performative exchange (so-called conversation policies or interaction protocols [1,8,17,29]) and many of the performatives have been designed with this in mind. However, a significant amount of communications will be performed outside of the auspices of any *explicit* interaction pattern. These "standalone" performatives, INFORM, REQUEST, SHOUT, and ORDER, will certainly be used within conversation policies, but are more likely than the other performatives here to be used alone.

    Oftentimes an *implicit* protocol will be employed in practical, fielded, multiagent systems [15,31,33]. For example, the performatives NOTUNDERSTOOD might be returned after any utterance, and ACKNOWLEDGE might be returned after any

REQUEST, irrespective of any explicit conversation policy. The REQUEST performative itself holds a special place in Joint Intention theory in that its utterance implies certain subsequent behavior (the establishment of mutual belief, in particular) on the part of the requesting and requested agent, much of which can be satisfied by messaging. We created the SUCCESS, RELFAIL, and IMPOSSIBLE performatives to fill the messaging needs of the agents involved in a REQUEST utterance, as each of these performatives can be used to establish mutual belief regarding one aspect of the PWAGs in the definition of REQUEST. So, while an explicit conversation policy may be in place upon the utterance of a REQUEST, the implicit messaging requirements can be satisfied by our suite of performatives.

The SUBSCRIBE and QUERY performatives also have a simple implicit protocol associated with them. By definition, utterance of a SUBSCRIBE will be followed by one or more INFORM messages should the recipient AGREE to the SUBSCRIBE request. To terminate a SUBSCRIBE performative, a CANCEL performative must be sent by the agent that performed the SUBSCRIBE. Similarly, the QUERY performative will be followed by one INFORM should the recipient of the QUERY decide to honor the request. Interaction protocols that make the messages subsequent to SUBSCRIBE and QUERY are certainly not prohibited.

The remaining performatives were defined with the expectation of using them primarily, if not solely, within interaction protocols. For example, we designed a number of other performatives to complement the REQUEST performative within an interaction protocol. These specifically are AGREE, REFUSE, CANCEL, SUCCESS, FAILURE, IMPOSSIBLE, and RELFAIL. AGREE and REFUSE can immediately follow a REQUEST in order to establish or prevent establishment of a Joint Intention team of agents, respectively. Following an AGREE message, the requesting agent can utter a CANCEL to dissolve the team. A FAILURE message, not strictly required by Joint Intention theory, can be used by the requested agent to indicate a particular state of progress toward performing the requested task; the FAILURE message indicates that the tasked agent is actively trying to do the task and is having some setbacks (but that the task might yet be performed, as otherwise it would have sent an IMPOSSIBLE performative). A SUCCESS performative can be used to indicate that the action was completed successfully and will dissolve a JI team. The IMPOSSIBLE performative indicates from either requesting or requested agents that the task requested can never be accomplished and will result in dissolution of a JI team. RELFAIL, similarly, indicates that the context under which a REQUESTed task is relevant has become false and use of this performative will also result in the dissolution of a JI team.

With regards to the PROXY and PROXY-WEAK performatives, their utterance will be followed, by definition, by another speech act, assuming that the middle agent decided to honor the sender's request. Whether or not the PROXY and PROXY-WEAK performatives are included within explicit interaction protocols is a domain specific decision, but they will always engender the possibility of future speech acts due to their utterance.

The ACCEPT, REJECT, and WITHDRAW performatives are all designed to be complementary performatives to the STANDINGOFFER performative. Once an agent performs a STANDINGOFFER speech act, the recipient can perform and ACCEPT or REJECT performative to establish or prevent establishment of a JI team, respectively. Subsequent to an ACCEPT, the agent performing the STANDINGOFFER can utter a WITHDRAW to terminate the terms of the offer.

The ORDER performative stands out from the others defined in this paper due to its use of agent-to-agent relationships. Inclusion of this performative within the paper was motivated by the CIANC project [33], in which agents related to human operators interact with each other while controlling and interacting with agent-based autonomous robotic vehicles within a military environment. We have found the need to make explicit the authority relationships between agents as well as the "rules of engagement" (authorized to do a, prohibited from doing a, obliged to x to do a, etc. much like deontic logic relationships [32]) that the agents are constrained by. As

such, we have subsequently found it useful to develop performatives specialized to these authority and deontic relationships. The ORDER primitive is just one example of a set of performatives that will be developed along these lines.

This paper presents significant progress towards our goals of a single semantic basis for all of the performatives, a single place to refer to performative definitions, a much broader selection of performatives than has been specified to date, and clearer applicability to a much broader set of single and multi-message interactions than has been previously defined.

Much remains to be done in the field of developing an ACL based on Joint Intention Theory, however. For example, the semantics of the NOTUNDERSTOOD performative, which is important in fielded systems, has yet to be correctly defined in JI semantics. Of even broader impact, we need to apply the developments of group-theoretic semantics of [16] to all of the performatives to provide the flexibility of groups of agents as senders and/or recipients. There are a large number of applications where this capability is useful in multiagent systems, when agents acting as representatives of institutions and organizations (e.g. sales and purchasing agents) interact with each other.

4   Acknowledgements

References

1. Bauer B.; Muller, J.P.; Odell, J. An Extension of UML by Protocols for Multiagent Interaction. In Proceedings of the Fourth International Conference on MultiAgent Systems (ICMAS'00), pages 207-214, Boston, Massachussetts, 2000.
2. Cohen, P.R.; Levesque, H.J. Rational Interaction as the Basis for Communication. In P.R. Cohen, J. Morgan, and M.E. Pollack, Eds. Intentions in Communication, System Development Foundation Benchmark Series, MIT Press, Cambridge, MA, 221-256, 1990.
3. Cohen, P.R.; Levesque, H.J. Intention is Choice with Commitment, Artificial Intelligence, 42(3):213-261, 1990.
4. Cohen, P.R.; Levesque, H.J. Performatives in a Rationally Based Speech Act Theory, In Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics, Pittsburgh, Pennsylvania, 1990.
5. Cohen, P.R.; Levesque, H.J. Confirmations and Joint Action, In Proceedings of the 12th International Joint Conference on Artificial Intelligence, San Mateo, California, 1991.
6. Cohen, P.R.; Levesque, H.J. Teamwork. Nous, 25(4): 487-512. 1991.
7. Cohen, Philip; Levesque, Hector. Communicative Actions for Artificial Agents, in Proceedings of the First International Conference on Multi-Agent Systems,   San Francisco, CA, 65-72, 1995.
8. Finin, T.; Labrou, Y.; Mayfield, J. KQML as an Agent Communication Language. Software Agents, Bradshaw ed., MIT Press, 1995.
9. Giampapa, J.; Paolucci, M.; Sycara, K. Agent Interoperation Across Multi-Agent System Boundaries, Proceedings of the Fourth International Conference on Autonomous Agents (Agents 2000), pgs 179-186, 2000.
10. Grosz, Barbara J.; Kraus, Sarit. Collaborative Plans for Complex Group Action. Artificial Intelligence 86(2):269-357. 1996.
11. Grosz, Barbara; Hunsberger, Luke; Kraus, Sarit. Planning and Acting Together, AI Magazine 20(4):23-34, Winter 1999.
12. Huber, Marcus; Kumar, Sanjeev; Cohen, Philip; McGee, David. A Formal Semantics for Proxy Communicative Acts, In Proceedings of Agent Theories, Architectures, and Languages (ATAL), 221-234. Seattle, WA, August, 2001.

13. Jennings, N. R. Commitments and Conventions: The Foundation of Coordination in Multi-Agent Systems, The Knowledge Engineering Review, 8(3):223-250. 1993.

14. Jennings, N.R. Specification and Implementation of a Belief-Desire-Joint-Intention Architecture for Collaborative Problem Solving, Intelligent and Cooperative Information Systems, 2(3):289-318. 1993.

15. Jenning, N.R. Controlling Cooperative Problem Solving in Industrial Multi-Agent Systems using Joint Intentions, Artificial Intelligence, 75(2):195-240. 1995.

16. Kumar, S.; Huber, M.J.; McGee D.R.; Cohen, P.R.; Levesque, H.J. Semantics of Agent Communication Languages for Group Interaction. In Proceedings of the 16th National Conference on Artificial Intelligence, 42-47, 2000.

17. Kumar, Sanjeev; Huber, Marcus; Cohen, Philip R. Representing and Executing Protocols as Joint Actions, In Proceedings of the First International Conference on Autonomous Agent and MultiAgent Systems, 543-550, Bologna, Italy, June, 2002.

18. Kumar, Sanjeev; Huber, Marcus; Cohen, Philip; McGee, David. Toward a Formalism for Conversational Protocols Using Joint Intention Theory, Journal of Computational Intelligence (Special Issue on Agent Communication Language), Brahim Chaib-draa and Frank Dignum (Guest Editors), 18(2):174-228, 2002.

19. Labrou, Y.; Finin, T. A Semantic Approach for KQML - A General Purpose Communication Language for Software Agents. In Proceedings of the Third International Conference on Information and Knowledge Management, 447-455. 1994.

20. Levesque, H.J.; Cohen, P.R.; Nunes, J.H.T. On Acting Together. In Proceedings of AAAI-90, Boston, 1990.

21. 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.

22. Pitt, J.; Kamara, L.; Artikis, A. Interaction Patterns and Observable Commitments in a Multi-Agent Trading Scenario. In Proceedings of Agents 2001, Montreal, 2001.

23. Pitt, J.; Mamdani, A. A Protocol-Based Semantics for an Agent Communication Language. In Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, Stockholm, Sweden, 486-491, 1999.

24. Pynadath, D.V.; Tambe, M.; Chauvat, N.; Cavedon, L. Toward Team-Oriented Programming. In Intelligent Agents VI: Agent Theories, Architectures and Languages, 233-247. N.R. Jennings and Y. Lespérance, eds., Springer-Verlag, 1999.

25. Rao, A.S.; Georgeff, M.P. Social plans: Preliminary report. In E. Werner and Y. Demazeau, editors, Decentralized AI 3. 57-76. Elsevier Science Publishers B.V. Amsterdam, Netherlands, 1992.

26. Searle J.R. Speech Acts. Cambridge University Press, 1969.

27. Singh, M.P. Agent Communication Languages: Rethinking the Principles, Computer 31(12):40-47, IEEE Computer Society, December 1998.

28. Smith, I.A., Cohen, P.R. Toward a Semantics for an Agent Communications Language based on Speech Acts. In Proceedings of AAAI-96, 24-31, Portland, Oregon, 1996.

29. Smith, I.A.; Cohen, P.R.; Bradshaw, J.M.; Greaves, M.; Holmback, H. Designing Conversation Policies using Joint Intention Theory. In Proceedings of the Third International Conference on MultiAgent Systems, 269-276, Paris, France, IEEE Press, 1998.

30. Sycara, K., Decker, K., and Williamson, M., Middle-Agents for the Internet, In Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence, pgs 578-583, January, 1997.

31. Tambe, M., Towards Flexible Teamwork. Journal of Artificial Intelligence Research, 7:83-124. 1997.

32. von Wright, G.H. Deontic Logic, Mind. 60:1-15, 1951.

33. Wood S.; Zaientz J.; Beard J.; Frederiksen R.; Huber M. CIANC3: An Agent-Based Intelligent Interface for Future Combat Systems Command and Control, In Proceedings of the 2003 Conference on Behavior Representation in Modeling and Simulation, 2003.

34. www.fipa.org/specifications.