intention to do X"; $C_{Th}(\phi, \psi)$ is interpreted as: "If ϕ becomes true, adopt an intention that ψ hold."[2]

We assume that if a group GR mutually believe both that (1) an agent G intends—in the event of ϕ—to adopt some other intention, and (2) that G believes ϕ, then GR mutually believe that G actually adopts the new intention, as represented by the following axiom schemata:

$$(R_1) \quad MB(GR, Int.To(G, C_{To}(\phi, X)) \wedge Bel(G, \phi)) \Rightarrow MB(GR, Int.To(G, X))$$

$$(R_2) \quad MB(GR, Int.To(G, C_{Th}(\phi, \psi)) \wedge Bel(G, \phi)) \Rightarrow MB(GR, Int.Th(G, \psi))$$

Performatives

The decision-making mechanism uses two performatives: $Declare$ and $Assert$.[3] We model performatives as primitive actions of the form $Perf(G, \phi, GR^-)$, where $Perf$ is either $Declare$ or $Assert$, G is an agent, GR^- is a group of agents that does not include G, and ϕ is a proposition. The following axiom schema represents our simplifying assumption that performatives have been executed (i.e., "done") if and only if the group $GR = GR^- \cup \{G\}$ mutually believe they have been executed:

$$(P) \quad Done(G, Perf(G, \phi, GR^-)) \Leftrightarrow MB(GR, Done(G, Perf(G, \phi, GR^-)))$$

Thus, for example, G makes a declaration ϕ to GR^- if and only if GR mutually believe G has made such a declaration. For convenience, we use the following abbreviations:

$$Declared(G, \phi, GR^-) = Done(G, Declare(G, \phi, GR^-))$$
$$Asserted(G, \phi, GR^-) = Done(G, Assert(G, \phi, GR^-))$$

Declarations. In an arbitrary context, someone's declaring "I hereby dub thee *Sir Lancelot*" may have no implications; however, if Queen Elizabeth II declares it in the context of a royal ceremony, it will result in someone's being knighted. Agents use the $Declare$ performative to establish certain propositions as mutually-believed institutional facts. Following Searle (1998), we use constitutive rules of the form, X counts as Y in the context C, to define the force of declarative speech acts. We represent these rules as axioms of the form, $X \wedge C \Rightarrow Y$. For example, we might define the force of the dubbing declaration as follows:

$$Declared(G_1, Knighted(G_2, \texttt{SirLancelot}), \{G_2\}) \wedge Royalty(G_1) \Rightarrow MB(\{G_1, G_2\}, Knighted(G_2, \texttt{SirLancelot}))$$

Assertions. Agents use $Assert$ to report their beliefs and intentions to other agents. We assume that agents are truthful in their assertions.[4] We also assume that if an agent G asserts to some other agents GR^- that it holds a particular belief or intention, then the result of that assertion is that they all mutually believe G holds that belief or intention, as represented by the following:

$$(S_{Bel}) \quad Asserted(G, Bel(G, \psi), GR^-) \Rightarrow MB(GR, Bel(G, \psi))$$

$$(S_{To}) \quad Asserted(G, Int.To(G, X), GR^-) \Rightarrow MB(GR, Int.To(G, X))$$

$$(S_{Th}) \quad Asserted(G, Int.Th(G, \psi), GR^-) \Rightarrow MB(GR, Int.Th(G, \psi))$$

It follows from (S_{Bel}) and the definition of mutual belief that if *each* member of GR asserts to the other agents in GR that it believes some ψ, then the entire group mutually believe ψ:

Lemma 1. Given axiom schemata (S_{Bel}), the following is valid for all GR and ψ:

$$[(\forall G \in GR) Asserted(G, Bel(G, \psi), GR^-)] \Rightarrow MB(GR, \psi)$$

where GR^- (i.e., the "rest of the group") implicitly depends on G.

Throughout the rest of this paper, we assume a theory that includes all instances of all axiom schemata introduced to this point.

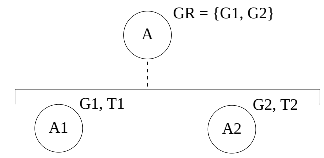

Figure 1: A method $\mathcal{M}_A$ for GR doing A

Methods and Full SharedPlans

We define a *method* $\mathcal{M}$ for a group of agents GR doing a multi-agent action A as a triple: $\langle R, \mathcal{S}, \mathcal{B} \rangle$, where R is a fixed recipe for doing A, $\mathcal{S}$ is a complete set of agent assignments for the subacts in R, and $\mathcal{B}$ is a complete set of bindings for all the parameters of A and R. For example, Figure 1 gives a schematic representation of a method $\mathcal{M}_A$ that specifies that the two-agent group, $GR = \{G_1, G_2\}$, can do the multi-agent action A using the recipe $\{A_1, A_2\}$, where G_1 does the primitive subact A_1 at time T_1, and G_2 does A_2 at T_2.

To *adopt* (or commit to) a fully-specified SharedPlan (FSP) means to *establish* the individual intentions and mutual beliefs specified in the definition of the FSP meta-predicate (Grosz and Kraus 1999; Hunsberger 1999). The FSP requirements depend on the method being used. We write $\mathcal{F}_{\mathcal{M}}$ for the set of FSP requirements corresponding to the method $\mathcal{M}$.

[2] The $Adopt.Int.To$ and $Adopt.Int.Th$ actions correspond to a subset of the functions of Ortiz' $Update$ action (Ortiz 1999b).

[3] Our specification of speech acts borrows from Cohen and Levesque (1990). Cohen and Levesque (1997) derive the semantics for request and commissive speech acts from the definition of an "attempt." We take a more direct approach. Ortiz (1999b) gives an alternative definition of an attempt.

[4] Perrault(1990) addresses the problem of agents making assertions they do not believe.

For the method $\mathcal{M}_A$ depicted in Figure 1, the FSP requirements (i.e., $\mathcal{F}_{\mathcal{M}_A}$) reduce to the following:[5]

F_1: $MB(GR, Is.Recipe(\{A_1, A_2\}, A))$
F_2: $MB(GR, CBA(G_1, A_1))$
F_3: $MB(GR, CBA(G_2, A_2))$
F_4: $MB(GR, Int.Th(G_1, Done(GR, A)))$
F_5: $MB(GR, Int.Th(G_2, Done(GR, A)))$
F_6: $Int.To(G_1, A_1) \wedge MB(GR, Int.To(G_1, A_1))$
F_7: $MB(GR, Int.Th(G_2, CBA(G_1, A_1)))$
F_8: $Int.To(G_2, A_2) \wedge MB(GR, Int.To(G_2, A_2))$
F_9: $MB(GR, Int.Th(G_1, CBA(G_2, A_2)))$

In general, the FSP requirements corresponding to some method $\mathcal{M}$ may be partitioned into two subsets which we denote $\mathcal{F}^0_{\mathcal{M}}$ and $\mathcal{F}^I_{\mathcal{M}}$. The requirements in $\mathcal{F}^0_{\mathcal{M}}$ involve mutual beliefs pertaining only to the validity of the method (i.e., that the recipe is valid and that the agents are able to do their assigned subacts). Agents holding such mutual beliefs need not be committed to doing anything. For the example above, $\mathcal{F}^0_{\mathcal{M}_A} = \{F_1, F_2, F_3\}$. In contrast, the requirements in $\mathcal{F}^I_{\mathcal{M}}$ involve agent commitments and mutual beliefs about those commitments. For the example above, $\mathcal{F}^I_{\mathcal{M}_A} = \{F_4, F_5, F_6, F_7, F_8, F_9\}$.

Invoking a Method

When collaborating on some group activity, agents typically begin with a partial plan that, over time, they elaborate into a complete plan. Grosz and Kraus (1996, 1999) argue that certain planning (i.e., group decision making) processes should be modeled as fixed, fully-specified SharedPlans (FSPs). In a subsequent section, we present a generic group-decision-making method. In this section, we describe how a group of agents can adopt an FSP corresponding to such a method. In particular, we describe how—given an arbitrary proposition ϕ and a decision-making method $\mathcal{M}$—a group's adoption of an FSP corresponding to the method $\mathcal{M}$ may be triggered by the group's mutual belief that ϕ holds. We call ϕ the triggering condition and we say that the group's mutual belief of ϕ *invokes* the method (or invokes the FSP corresponding to the method). Sufficient requirements are:

(1) the group GR mutually believe that the method is valid (i.e., the mutual beliefs in $\mathcal{F}^0_{\mathcal{M}}$ hold); and
(2) the group GR mutually believe that they hold a certain set of background commitments (dependent on both $\mathcal{M}$ and ϕ, as described below) that may be interpreted as their willingness to adopt the FSP when triggered.

We first describe the triggered invocation of an FSP corresponding to the method $\mathcal{M}_A$ depicted in Figure 1. Let ϕ be an arbitrary triggering condition. Let $\mathcal{C}_{\mathcal{M}_A}(\phi)$ be a set containing the following conditional commitments (dependent on both $\mathcal{M}_A$ and ϕ):

C_1: $Int.To(G_1, C_{Th}(\phi, Done(GR, A)))$
C_2: $Int.To(G_1, C_{To}(\phi, A_1))$
C_3: $Int.To(G_1, C_{Th}(\phi, CBA(G_2, A_2)))$
C_4: $Int.To(G_2, C_{Th}(\phi, Done(GR, A)))$
C_5: $Int.To(G_2, C_{To}(\phi, A_2))$
C_6: $Int.To(G_2, C_{Th}(\phi, CBA(G_1, A_1)))$

For example, C_3 represents G_1's commitment (conditioned on ϕ) to adopt an intention that G_2 be able to do A_2—which is precisely what must be mutually believed to satisfy the FSP requirement F_9.

The following theorem says that if the group GR mutually believe both (1) that they hold the conditional commitments in $\mathcal{C}_{\mathcal{M}_A}(\phi)$ and (2) that the triggering condition ϕ holds, then all of the FSP requirements in $\mathcal{F}^I_{\mathcal{M}_A}$ necessarily hold; hence, if they also mutually believe that the method is valid (i.e., the FSP requirements in $\mathcal{F}^0_{\mathcal{M}_A}$ hold), then the FSP requirements corresponding to the method $\mathcal{M}_A$ necessarily hold (i.e., the group has adopted the FSP corresponding to the method $\mathcal{M}_A$). In the theorem, the brackets around $\mathcal{C}_{\mathcal{M}_A}(\phi)$, $\mathcal{F}_{\mathcal{M}_A}$, $\mathcal{F}^I_{\mathcal{M}_A}$ and $\mathcal{F}^0_{\mathcal{M}_A}$ are used to represent the conjunction of all the clauses in the bracketed set.

Theorem 1 (Special Case: $\mathcal{M} = \mathcal{M}_A$). The following is valid for all GR and ϕ:

$$MB(GR, \phi \wedge [\mathcal{C}_{\mathcal{M}_A}(\phi)]) \Rightarrow [\mathcal{F}^I]; \text{ and hence:}$$
$$MB(GR, \phi \wedge [\mathcal{C}_{\mathcal{M}_A}(\phi)]) \wedge [\mathcal{F}^0] \Rightarrow [\mathcal{F}_{\mathcal{M}_A}].$$

[All proofs are omitted due to space limitations.]

Theorem 1 (General Version). Theorem 1 is easily generalized to cover the triggered invocation of an FSP based on an arbitrary method $\mathcal{M}$ whose recipe includes any number of primitive subacts and which involves a group of arbitrarily many agents. The voting-based group-decision-making mechanism described in the next section is based on such a method. The more subacts and agents involved in the method, the more numerous are the requirements in $\mathcal{F}^0_{\mathcal{M}}$ and $\mathcal{F}^I_{\mathcal{M}}$, and the more numerous are the conditional commitments in $\mathcal{C}_{\mathcal{M}}(\phi)$; the basic idea, however, is the same.

Invoking an FSP to Make a Group Decision

In this section, we show how an agent can invoke a decision-making FSP that a group of agents may use to make a group decision. We describe a single voting-based decision-making method $\mathcal{M}_V$; however any decision-making method representable by an FSP may be similarly treated. Thus, the presentation in this section has wide applicability to group decision making in multi-agent systems. Without loss of generality, we describe a scenario in which a particular agent (G_1) invokes the method.

A schematic for the voting method $\mathcal{M}_V$ is shown in Figure 2. The method involves a group GR of n agents: $G_1, \ldots, G_n$. The recipe has $n + 1$ subacts, each of which is a declarative speech act. D_1 is the declaration used by G_1 to invoke the method; $V_2, \ldots, V_n$ are declarations by the rest of the agents either to accept or reject the proposal contained in the invoking declaration; D_a is G_1's declaration

[5] $Is.Recipe$ models that the given set of subacts constitutes a recipe for the given action; CBA ("can bring about") models an agent's ability to do an action (Grosz and Kraus 1999).

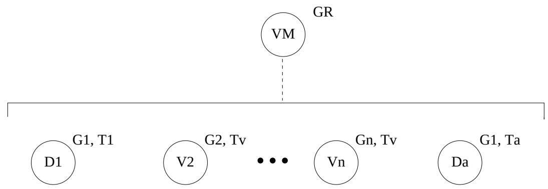

Figure 2: A schematic for the voting method $\mathcal{M}_V$

announcing the result of the voting. The axioms associated with the declarations in $\mathcal{M}_V$ (given below) ensure that if G_1 announces that the group has accepted the proposal, then the group's decision shall necessarily be established as a mutually believed fact.

Let $\mathcal{F}^0_{\mathcal{M}_V}$ and $\mathcal{F}^I_{\mathcal{M}_V}$ be the FSP requirements corresponding to the method $\mathcal{M}_V$. We assume that the group mutually believe that the voting method is valid (i.e., that the mutual beliefs in $\mathcal{F}^0_{\mathcal{M}_V}$ hold). Let $\mathcal{C}_{\mathcal{M}_V}(\phi_V)$ be the set of conditional commitments corresponding to the method $\mathcal{M}_V$ and the triggering condition ϕ_V, as described in the previous section. (The triggering condition ϕ_V is defined below.) We assume that the agents mutually believe that they hold the conditional commitments in $\mathcal{C}_{\mathcal{M}_V}(\phi_V)$. Thus, by Theorem 1 (general version), mutual belief in ϕ_V is sufficient to trigger the group's adoption of the FSP corresponding to $\mathcal{M}_V$.

Without loss of generality, we consider a "simple selection" decision problem in which the group must select a single item σ from a fixed set Σ.

Invoking the Voting Method. G_1's initial declaration is: $Declare(G_1, Invoked\_VM(G_1, \mathcal{M}_V, \pi), GR^-)$, where $\mathcal{M}_V$ includes the following parameters:

group: $GR = \{G_1, \ldots, G_n\}$	invocation time: T_1
voting time interval: T_v	announcement time: T_a

and π includes the following additional parameters:

decision problem: $(\texttt{SelectItem}, \Sigma)$

proposal: $\delta = (\texttt{Select}, \sigma_3)$

The force of G_1's declaration is defined by the following axiom schema:

$$(S_1) \quad Declared(G_1, Invoked\_VM(G_1, \mathcal{M}_V, \pi), GR^-) \Rightarrow MB(GR, Invoked\_VM(G_1, \mathcal{M}_V, \pi))$$

Let ϕ_V be the following triggering condition:

$$Invoked\_VM(G_1, \mathcal{M}_V, \pi)$$

By schema (S_1), G_1's initial declaration is sufficient to ensure that the group mutually believe ϕ_V. Thus, by Theorem 1 (general version) the group necessarily adopts the voting-mechanism FSP in response to G_1's invocation.

The Voting Phase. During the time interval T_v (a parameter in $\mathcal{M}_V$), the agents $G_2, \ldots, G_n$ must do the voting actions $V_2, \ldots, V_n$, respectively. Each vote is a declaration of accepting or rejecting the proposal δ (a parameter in π). Axiom schema (S_2) defines the force of a vote to accept δ. The force of a vote to reject δ can be defined analogously.

$$\begin{aligned}(S_2) \quad & Declared(G, Accepted(G,\delta), \{G_1\}) \\ & \wedge MB(\{G, G_1\}, Invoked\_VM(G_1, \mathcal{M}_V, \pi)) \\ & \wedge G \in GR \\ & \wedge G \neq G_1 \\ & \Rightarrow MB(\{G, G_1\}, Accepted(G, \delta))\end{aligned}$$

The Announcement. After the voting interval, G_1 announces the results of the voting to the rest of the group by making the declaration: $Declare(G_1, \chi, GR^-)$, where χ is either $GroupAccepted(GR, \delta)$ or $GroupRejected(GR, \delta)$. Due to space limitations, we only describe the announcement to accept.

We provide a set of three constitutive rules—(S_3), (S_4) and (S_5) below—sufficient to ensure that if G_1 announces the group decision to accept δ, then that announcement establishes the group's decision as a mutually-believed fact (formalized in Theorem 2 below). We use three constitutive rules rather than a single rule to explicitly show the assumptions sufficient to generate the desired result.

The rules have the following form:

(S_3) $\quad X \wedge C_1 \Rightarrow Y_1$;

(S_4) $\quad X \wedge (C_1 \wedge C_2) \Rightarrow Y_2$, where $Y_1 \Rightarrow C_2$; and

(S_5) $\quad X \wedge (C_1 \wedge Y_2) \Rightarrow Y_3$.

In each rule, X stands for G_1's announcement. The result (Y_1) of the first rule enriches the context for the second rule; the result (Y_2) of the second rule enriches the context for the third rule.

In the first rule (S_3), G_1's announcement, in the context of the group mutually believing that G_1 invoked the voting mechanism, counts both as an assertion that each member agent voted to accept the proposal and as establishing the group's mutual belief that G_1 accepted the proposal.

$$\begin{aligned}(S_3) \quad & Declared(G_1, GroupAccepted(GR, \delta), GR^-) \\ & \wedge MB(GR, Invoked\_VM(G_1, \mathcal{M}_V, \pi)) \\ & \Rightarrow Asserted(G_1, Bel(G_1, Accepted(G_2, \delta)), GR^-) \\ & \vdots \\ & \wedge Asserted(G_1, Bel(G_1, Accepted(G_n, \delta)), GR^-) \\ & \wedge MB(GR, Accepted(G_1, \delta))\end{aligned}$$

From Lemma 1, these assertions entail that the group GR mutually believe that G_1 believes that everyone *else* voted to accept:

$$(\forall G \in GR^-) MB(GR, Bel(G_1, Accepted(G, \delta))).$$

In the second rule (S_4), the context is enriched by the mutual beliefs resulting from the first rule. In this context, G_1's announcement counts as the group mutually believing that everyone voted to accept δ:

$$\begin{aligned}(S_4) \quad & Declared(G_1, GroupAccepted(GR, \delta), GR^-) \\ & \wedge MB(GR, Invoked\_VM(G_1, \mathcal{M}_V, \pi)) \\ & \wedge (\forall G \in GR^-) MB(GR, Bel(G_1, Accepted(G, \delta))) \\ & \wedge MB(GR, Accepted(G_1, \delta)) \\ & \Rightarrow (\forall G \in GR) MB(GR, Accepted(G, \delta))\end{aligned}$$

In the third rule (S_5), the context includes the result of the second rule. In this context, G_1's announcement counts as establishing the group's decision as a mutually believed fact.

$$(S_5) \quad \begin{aligned} &Declared(G_1, GroupAccepted(GR, \delta), GR^-) \\ &\wedge MB(GR, Invoked\_VM(G_1, \mathcal{M}_V, \pi)) \\ &\wedge (\forall G \in GR) MB(GR, Accepted(G, \delta)) \\ &\Rightarrow \ MB(GR, GroupAccepted(GR, \delta)) \end{aligned}$$

Theorem 2. Given axiom schemata, $S_2, \ldots, S_5$, the following is valid:

$$\begin{aligned} &Declared(G_1, GroupAccepted(GR, \delta), GR^-) \\ &\wedge MB(GR, Invoked\_VM(G_1, \mathcal{M}_V, \pi)) \\ &\Rightarrow \ MB(GR, GroupAccepted(GR, \delta)) \end{aligned}$$

Applying the Mechanism to Collaborative Planning

Agents collaborating on some group activity often cannot adopt a fully specified plan for that activity in a single stroke. Instead, they must make numerous decisions (e.g., which recipe to use, which parameter values to use, or which agents or subgroups to assign to the various subtasks) as they elaborate their possibly-hierarchical partial plan into a more complete plan. In this section, we describe how the voting mechanism from the previous section may be applied to the decisions encountered by a group of agents collaborating on some group activity. The key idea is that the definition of the force of a vote to accept a proposal (given in schema (S_2) above) must be augmented according to the context determined by the type of decision being voted on. Thus, for example, a vote to use a particular recipe represents one set of conditional commitments, while a vote to assign a fellow agent to a particular subtask represents another set of commitments.

We illustrate the use of the voting mechanism as a general-purpose decision-making tool by showing how it may be used by a group both to adopt a minimal SharedPlan (i.e., a plan for which no recipe has been selected and no parameters have been bound) and to select a recipe for that plan (which introduces new decision problems).

To simplify the presentation, we assume a weaker model of the group's commitment to the plan-elaboration process than that used by Grosz and Kraus.[6] In particular, for each decision problem facing the group, we require only that the group mutually believe that each agent intends that the group find a way of resolving that problem.

Deciding to Adopt a Minimal SharedPlan (MSP). Suppose G_1 invokes the voting mechanism as described previously, but with $GR = \{G_1, G_2\}$ and π containing the following different parameters:

decision problem: $\texttt{Adopt\_MSP}(A)$
proposal: $\delta_1 = \texttt{YES}$ (i.e., adopt the MSP)

Thus, G_1 proposes that the group adopt a minimal SharedPlan (MSP) to do the multi-agent action A.

To adopt such a plan means to establish a subset of the intentions and mutual beliefs in the definition of the PSP (Partial SharedPlan) meta-predicate (Grosz and Kraus, 1999; Hunsberger, 1999). In this case, these requirements reduce to the following (for each G_i in GR):

$$(M_1) \quad MB(GR, Int.Th(G_i, Done(GR, A)))$$

$$(M_2) \quad MB(GR, Int.Th(G_i, Recipe.Selected(GR, A)))$$

Let $\xi_1 = GroupAccepted(GR, \delta_1)$. An agent voting to adopt an MSP counts as that agent asserting that it intends, in the event of ξ_1 (i.e., in the event that the group decides to accept the proposal), to adopt the intentions required of it in (M_1) and (M_2), as follows.

$$(S_6) \quad \begin{aligned} &MB(GR, Accepted(G, \delta_1)) \\ &\wedge DecisionProblemType(\pi, \texttt{Adopt\_MSP}(A)) \\ &\Rightarrow Asserted(G, I_1 \wedge I_2, GR^-) \end{aligned}$$

where I_1 and I_2 are given by:

$$I_1 = Int.To(G, C_{Th}(\xi_1, Done(GR, A)))$$

$$I_2 = Int.To(G, C_{Th}(\xi_1, Recipe.Selected(GR, A)))$$

Theorem 3 states that if a group of agents using this decision-making mechanism decide to adopt a minimal SharedPlan, then they will in fact adopt it.

Theorem 3. Given schemata, $S_2, \ldots, S_6$, the following is valid:

$$\begin{aligned} &Declared(G_1, GroupAccepted(GR, \delta_1), GR^-) \\ &\wedge MB(GR, Invoked\_VM(G_1, \mathcal{M}_V, \pi)) \\ &\wedge DecisionProblemType(\pi, \texttt{Adopt\_MSP}(A)) \\ &\Rightarrow MSP(GR, A) \end{aligned}$$

Deciding to Select a Recipe. Suppose G_1's intention that the group select a recipe leads[7] G_1 to invoke the voting mechanism with a proposal that the group use the recipe $\{A_1, A_2\}$ (which introduces the unbound parameters T_1 and T_2). To incorporate this decision into their existing plan requires that they establish the following additional mutual beliefs (for each G_i in GR).

$$(M_3) \quad MB(GR, Int.Th(G_i, Agent.Assigned.To(A_1)))$$

$$(M_4) \quad MB(GR, Int.Th(G_i, Agent.Assigned.To(A_2)))$$

$$(M_5) \quad MB(GR, Int.Th(G_i, Params.Bound(\{T_1, T_2\})))$$

These mutual beliefs may be established by a schema nearly identical to S_6 above.

Other Types of Decisions. Making other types of decisions (e.g., to select an agent to do a subact or to bind a parameter) have slightly different requirements, but the general

[6]In addition, we conjecture that this weaker model, which Grosz and Kraus (1999) use to model the group's commitment to finding values for the parameters of the group action, but not their commitment to selecting a recipe or assigning agents to subtasks, may be sufficient in many scenarios.

[7]We use "leads" in the same sense as Grosz and Kraus (1999) in their axioms of intention-that.

procedure remains the same. For example, if the group decides that G_1 should do A_1, then they must establish clauses F_2, F_6 and F_7 (discussed previously), which may be ensured by augmenting the definition of the force of an accept vote to make it count as an assertion of belief in G_1's ability to do A_1 and an assertion of conditional commitment corresponding to C_2 (for G_1's vote) or C_6 (for G_2's vote).

Related Work

Many researchers are actively investigating frameworks for reasoning about collective activity in multi-agent systems. The role of communication is often recognized as crucial, but there are few formal studies of mechanisms for group decision making.

Werner (1990) distinguishes directive and informative speech acts. A "pragmatic interpretation" of high-level messages is used to transform the information and intentional states of agents. Werner also discusses the "institutional effects" of certain "representative declarative" speech acts, and gives examples of how directive and informative speech acts may be used in "social cooperation."

Cohen and Levesque (1997) derive the semantics for request and commissive speech acts from the definition of an "attempt" such that certain requests followed by certain commissive actions result in the formation of joint commitments. In contrast, we propose a general mechanism for group decision making and derive the adoption of a SharedPlan as a special case. Cohen and Levesque (1997) do not discuss declarative speech acts.

The set of background conditions that enable certain speech acts to directly invoke the voting mechanism in our work is comparable to the locker-room agreements presented by Stone and Veloso (1999). Both approaches can be viewed as mutually-believed commitment by the group members to "do the right thing at the right time." The main difference is that the locker-room agreement is used to execute a plan in the absence of reliable communication, whereas the voting mechanism in our work is used to dynamically establish decisions needed in elaborating a partially-specified plan.

Conclusions

The aim of the paper is to more fully specify the dynamic evolution of a partial SharedPlan to a more complete plan. The SharedPlans formalization of collaboration (Grosz and Kraus 1999) stipulates that collaborating agents must commit to certain decision-making processes (namely, adopting the initial commitment, selecting a recipe, assigning agents to subtasks, and identifying action parameters) but does not specify those processes. We have presented a mechanism for group decision making that may be applied to all of these types of decision.

The decision-making mechanism is modeled as a fixed, fully-specified SharedPlan (FSP) that can be directly invoked. We have provided conditions sufficient to ensure that the group will automatically adopt such an FSP as soon as they come to mutually believe a given triggering condition holds. We provided the $Declare$ speech act to enable agents to establish mutual belief of a triggering condition. The definition of the force of the $Declare$ speech act was specified using Searle's constitutive rule (i.e., "the performance of X in the context C counts as Y"). The $Assert$ speech act was used to define the consequences of certain declarations.

We illustrated how the decision-making mechanism can be used to adopt a SharedPlan as well as to make the various decisions needed to elaborate a partially-specified plan, provided that the context C in the constitutive rules is properly managed. Furthermore, the techniques presented in this paper provide a solid foundation for formalizing the process of group decision making based on any mechanism representable by an FSP.

Acknowledgments

The work presented in this paper was supported by NSF grants IIS-9978343, IRI-9618848, and CDA-94-01024. The authors thank Barbara J. Grosz for her invaluable guidance in the preparation of this paper.

References

Cohen, P., and Levesque, H. 1990. Rational interaction as the basis for communication. In Cohen et al. (1990). chapter 12.

Cohen, P. R.; Levesque, H. J.; and Smith, I. 1997. On team formation. In Holmstrom-Hintikka, G., and Tuomela, R., eds., *Contemporary Action Theory*. Kluwer.

Cohen, P.; Morgan, J.; and Pollack, M., eds. 1990. *Intentions in Communication*. MIT Press.

Grosz, B. J., and Kraus, S. 1999. The evolution of SharedPlans. In Wooldridge, M., and Rao, A., eds., *Foundations of Rational Agency*, number 14 in Applied Logic Series. Kluwer.

Hunsberger, L. 1999. Making SharedPlans more concise and easier to reason about. In Muller et al. (1999), 81–98.

Muller, J.; Singh, M.; and Rao, A., eds. 1999. *Intelligent Agents V*, volume 1555 of *Lecture Notes in Artificial Intelligence*. Heidelberg: Springer.

Ortiz, C. L. 1999a. A commonsense language for reasoning about causation and rational action. *Artificial Intelligence* 111(1–2).

Ortiz, C. L. 1999b. Introspective and elaborative processes in rational agents. *Annals of Mathematics and Artificial Intelligence* 25(1–2).

Perrault, C. R. 1990. An application of default logic to speech act theory. In Cohen et al. (1990). chapter 9.

Searle, J. R. 1998. *Mind, Language and Society*. New York: Basic Books.

Stone, P., and Veloso, M. 1999. Task decomposition and dynamic role assignment for real-time strategic teamwork. In Muller et al. (1999).

Werner, E. 1990. Cooperating agents: A unified theory of communication and social structure. In *Distributed Artificial Intelligence: Volume II*. Morgan Kaufmann. 3–36.