

The dynamics of intention in collaborative activity

Action editors: Luca Tummolini and Cristiano Castelfranchi

Barbara J. Grosz^a, Luke Hunsberger^{b,*}

^a DEAS, Harvard University, Maxwell-Dworkin 249, 33 Oxford Street, Cambridge, MA 02138, USA

^b Computer Science Department, Vassar College, Box 444, 124 Raymond Avenue, Poughkeepsie, NY 12604, USA

Received 31 March 2005; accepted 7 November 2005

Available online 3 March 2006

Abstract

An adequate formulation of collective intentionality is crucial for understanding group activity and for modeling the mental state of participants in such activities. Although work on collective intentionality in philosophy, artificial intelligence, and cognitive science has many points of agreement, several key issues remain under debate. This paper argues that the dynamics of intention – in particular, the inter-related processes of plan-related group decision making and intention updating – play crucial roles in an explanation of collective intentionality. Furthermore, it is in these dynamic aspects that coordinated group activity differs most from individual activity. The paper specifies a model of the dynamics of agent intentions in the context of collaborative activity. Its integrated treatment of group decision making and coordinated updating of group-related intentions fills an important gap in prior accounts of collective intentionality, thus helping to resolve a long-standing debate about the nature of intentions in group activity. The paper also defines an architecture for collaboration-capable computer agents that satisfies the constraints of the model and is a natural extension of the standard architecture for resource-bounded agents operating as individuals. The new architecture is both more principled and more complete than prior architectures for collaborative multi-agent systems.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Collective intentionality; Multi-agent systems; Agent architectures; Collaborative planning

1. Introduction

There is broad agreement in philosophy, artificial intelligence and cognitive science (Bratman, 1999; Gilbert, 2000; Grosz & Kraus, 1996; Grosz & Kraus, 1999; Grosz & Sidner, 1990; Kinny et al., 1994; Levesque, Cohen, & Nunes, 1990; Searle, 1990; Tuomela, 1995; Tuomela, 2004b) that the collective, joint activity of a group is more than the simple sum of the domain-oriented actions performed by the individuals in that group: coordinating activities, typically including some communicative actions, are required. There is also cross-field agreement that the plans that underlie the collective activity of a group involve more than the simple sum of the plans of the individual

group members and that the simple sum of the intentions of the individuals toward their own actions does not fully capture collective intentionality. However, there is a marked divergence in claims about the constructs needed to fill the gap between the whole and any simple collection of single-agent actions, intentions and plans.

Several philosophers and cognitive scientists have argued that collective intentionality requires a different kind of mental-state construct, namely an intentional attitude that although individually held is different from and not reducible to an “ordinary” intention. For instance, Searle (1995) claims that, “In addition to singular intentionality there is also collective intentionality,” where “the intentionality that exists in each individual head has the form ‘we intend’.” He argues that collective intentionality “is a biologically primitive phenomenon that cannot be reduced to or eliminated in favor of something else.” Similarly, Tuomela (1995) argues for individually held

* Corresponding author.

E-mail addresses: grosz@eecs.harvard.edu (B.J. Grosz), hunsberg@cs.vassar.edu (L. Hunsberger).

“we-intentions” which he describes as “agreement-based social intentions that agents have in situations of joint action” with content that “can be taken to be something like ‘to do X jointly’ or ‘we to do X jointly’.” He also distinguishes intending in the *we-mode* from intending in the *I-mode*, and argues that “the we-mode is not reducible to the I-mode and these modes may be in conflict” (Tuomela, 2004a). Gilbert (2000) presents a “plural subject” account of collective intentionality in which agents form a joint commitment to “intend as a body”. She argues that the individual “commitment shares” are not, and do not require the existence of, personal commitments. In the remaining discussion, we will use the term “plural stance” to distinguish approaches like these, which require a separate, new category of intentions that have plural subjects.

Several of the accounts taking a plural stance argue further that agreements and the obligations they entail play a central role in distinguishing the mental state required for collective intentionality from that required for individual activity. In particular, they claim that obligations serve to bind the participants of a group activity together in coordinating and pursuing their collaborative endeavor. For instance, in Tuomela’s account of “full-blown, agreement-based joint intention” (Tuomela, 1995), agreements entail obligations on the participants to adopt certain specified relevant intentions. Similarly, in Gilbert’s account (Gilbert, 2000), joint commitments have inherent obligations on agents to provide appropriate action. Castelfranchi (1995) argues for “social commitments”, which involve similar obligations, rights and entitlements.

Some computational approaches to formalizing collaboration or designing computational agents capable of collaboration (Jennings, 1995; Levesque et al., 1990; Tambe, 1997) also take the plural stance. They employ notions of joint intention that although reducible to sets of individually held attitudes, reduce to something other than “ordinary” intentions.

In contrast to accounts taking the plural stance, Bratman (1999) claims that no new kind of intention is required for characterizing collective action and intentionality and that collaboration can occur in the absence of any initiating agreement or mutual obligation. He argues that an interlocking web of beliefs, mutual beliefs and ordinary intentions is sufficient. To achieve the coordination and commitment needed for collective, cooperative activities, his account requires that the content of each participant’s group-related intention “that we [perform the joint action] J” include the modifier “in accordance with and because of [our intentions] and meshing subplans of [our intentions].” Some computational approaches to formalizing collaboration and constructing collaborative multi-agent systems also deploy solely ordinary intentional attitudes (Grosz & Kraus, 1996; Grosz & Kraus, 1999; Grosz & Sidner, 1990; Kinny et al., 1994). In the remaining discussion, we will use the term “singular stance” to refer to approaches like these that use only intentions having a singular subject. In singular-stance approaches, the object of an

intention may be a group action; that is, a singular subject intends, but the content of the intention – i.e. what is intended – may have a plural subject (e.g. that we do something).

This paper argues that the singular and plural stances can be reconciled, but only through careful consideration of the dynamics of intentions in collaborative activities. An adequate account of these dynamics requires not only the modification of theories of collaboration but also the redesign of architectures for collaboration-capable, resource-bounded computer agents. To address these needs, the paper extends the SharedPlans formalization of collaborative activity (Grosz & Kraus, 1996; Grosz & Kraus, 1999) to accommodate more completely the dynamics of intention. In particular, it defines a model, the “coordinated cultivation of SharedPlans” (CCSP) model, which provides a more comprehensive treatment than prior work of the interactions of decision making and intention updating in group activity. It then articulates an architecture for collaboration-capable computer agents based on the CCSP model. The major theoretical contribution of the CCSP model is to enable an explanation of collective intentionality that rests on purely singular-stance constructs while capturing the essential properties argued for in accounts taking the plural stance. The major contribution from a computational perspective is to provide a more general architecture for collaboration-capable agent design.

The next section of the paper briefly describes the SharedPlans formalization of collaborative activity. Section 3 presents the CCSP model and the CCSP-based architecture for collaboration-capable computer agents. Of special note, Section 3.2 highlights important differences between the dynamics of intention in purely single-agent activity and in collective group activity. Section 4 shows that the CCSP model meets the requirements for collective intentionality identified in prior singular- and plural-stance approaches and compares the CCSP-based architecture to existing architectures for collaborative, multi-agent systems.

2. Background

2.1. SharedPlans

Agents, whether people or computer systems, are resource bounded and the worlds in which they operate are dynamic (Bratman, Israel, & Pollack, 1988; Grosz & Kraus, 1996). Thus, for both individual and multi-agent planning, it must be possible to form initial, incomplete plans which may be expanded and changed over time. Partial plans typically include intentions that are under-specified. For example, if a cook has not yet decided whether to get tomatoes from the grocery store or the farmer’s market, his plan to make Caprese salad will be incomplete and his intention “to make a Caprese salad” under-specified. Likewise, if two friends have agreed to go to the movies, but not yet decided which movie to see, their plan to go to the movies will be incomplete, and

- Agents have a SharedPlan to do α if and only if they hold the following beliefs and intentions:
- (1) individual intentions that the group perform α ;
 - (2) mutual belief of a (partial) recipe for α ;
 - (3) individual or group plans for the sub-acts in the (partial) recipe;
 - (4) intentions that the selected agents or subgroups succeed in doing their sub-acts (for all sub-acts that have been assigned to some agent or group); and
 - (5) [in the case of a partial SharedPlan] subsidiary commitments to group decision-making processes aimed at completing the group's partial plan.

Fig. 1. The beliefs and intentions required for a SharedPlan.

whatever intentions represent their commitment “to go to the movies” will be under-specified. As a result of such incompleteness, plans and intentions related to actions, whether of individuals or groups, require that means-ends reasoning, or planning,¹ be done and that decisions be made about various aspects of the doing of those actions.

SharedPlans is a formal specification of the mental-state requirements of participants in a collaborative, group activity that addresses the need for partiality in plans and the requisite commitment of agents to group decision-making processes.² The specification is given in terms of the intentions, beliefs and mutual beliefs of the participants. It deploys two intentional attitudes, *intending to* (do an action) and *intending that* (a proposition hold). Intentions-to are used to represent an agent's commitments to its own actions, whether for individual ends or as part of a group activity. Such commitments include commitments toward reasoning about ways of accomplishing its intended actions—that is, to planning. Intentions-that are used to represent an agent's commitments toward a group (or joint) activity and the actions of its fellow participants in service of that activity. Such commitments lead an agent to engage in a particular kind of plan-based reasoning, called *intention cultivation* (Grosz & Kraus, 1999), that includes reasoning about other participants' actions and intentions and ways the agent could contribute to their success in the context of the group activity.

The main constituents of the SharedPlans specification are given informally in Fig. 1. The term “recipe”, used in the figure, refers to a way of doing an action (Pollack,

1990). More specifically, a recipe for an action α is a set of actions (or “sub-acts”) and constraints such that doing those actions under those constraints constitutes doing α .

Each participant's intention “that the group do α ” (stated in Clause 1 of the figure) is of special importance to collective action. In the remainder of this paper, we refer to such intentions as *group-activity-related intentions* (or, *GAR intentions*). The intention-cultivation process associated with an agent's GAR intention generates a set of decisions that the group must make if they are to complete their plan for doing α . These “decision problems” include choosing a recipe for the group activity, determining the agents or subgroups that will be responsible for performing the sub-acts in the recipe, and selecting appropriate values for various other action parameters (e.g. the time at which the action is to be done or the resources to be used in doing it). The SharedPlans specification explicitly requires collaborating agents to be committed to solving these decision problems together; that is, they must be committed to reaching agreement (Clause 5). The SharedPlans formalization also includes a set of axioms that stipulate various properties of intentions and constraints among them (Grosz & Kraus, 1996; Grosz & Kraus, 1999). These axioms, together with intention-cultivation processes deriving from agents' GAR intentions, engender the cooperation and coordination required for collective group activity.

2.2. The reasoning-choosing-updating cycle

Agents with a partial plan engage a centrally important three-phase cycle of *reasoning*, *choosing* and *updating*. The SharedPlans formalization focuses primarily on specifying the commitments of collaborating agents toward the reasoning and choosing phases of this cycle, but does not explicitly address the operational dynamics of making choices and subsequently updating intentions. In the course of implementing collaboration-capable computer agents according to the SharedPlans specification (Grosz, Hunsberger, & Kraus, 1999), the second author recognized the need to more directly address these operational dynamics (Hunsberger, 2002b). First, there must be a formal

¹ The terms “means-ends reasoning” and “planning” are used by different fields and authors to refer to essentially the same process, namely the process of determining how to achieve a goal or objective (“the ends”) or do an action. This process is sometimes characterized as determining the “sub-acts” that must be done (“the means”) to perform a high-level action (“the ends”). In this paper we use the terms interchangeably, typically choosing the one that best matches the most related research.

² SharedPlans was initially presented at a 1986 workshop as reported in a paper (Grosz & Sidner, 1990) in the volume resulting from that workshop (Cohen, Morgan, & Pollack, 1990). The formalization was significantly generalized and revised in a subsequent set of papers (Grosz & Kraus, 1993; Grosz & Kraus, 1996; Grosz & Kraus, 1999; Hunsberger, 1999; Lochbaum, 1994). We use “SharedPlans” in this paper to refer to the formalization as refined through this progression.

specification of group decision making that makes explicit what it means to reach agreement. Second, the participants must be committed not only to reaching agreement on decisions, but also to updating their intentions to be in accordance with the agreed upon decisions; that is, the results of group decisions made in expanding a partial plan to a more complete one must eventually be reflected in the individual intentions of each agent.

Although a significant body of research has addressed the dynamics of intentions in individual activity (Lansky & Georgeff, 1987; Meyer, van der Hoek, & van Linder, 1999; Ortiz, 1999; Rao & Georgeff, 1991, 1992; Rao & Georgeff, 1995; Rao & Georgeff, 1995), the inter-related dynamics of decision making and intention updating in collaborative, multi-agent activity has received scant attention. Computational approaches to the problem of how participants initially form a collaborative team (Cohen, Levesque, & Smith, 1997; Hunsberger & Grosz, 2000; Kinny et al., 1994; Panzarasa, Jennings, & Norman, 2001, inter alia) address the decision-making and intention-update question in that context. Various implementations of collaborative multi-agent systems (Giampapa & Sycara, 2002; Jennings, 1995; Pynadath, Tambe, Chauvat, & Cavedon, 2000; Tambe, 1997, inter alia) have recognized the gap in coverage of dynamics in theories of collaboration, with different systems highlighting different problems. However, the mechanisms used by these implementations to compensate for the lack of theory have been limited in scope.

3. The coordinated cultivation of SharedPlans

The coordinated cultivation of SharedPlans (CCSP) model of the dynamics of intention in collaborative, group activity provides a uniform treatment of group decisions and extends the SharedPlans formalization to accommodate the inter-related and dynamic processes of group decision making and intention updating. In so doing, the CCSP model specifies the way in which agents must coordinate the cultivation of GAR intentions in group activity, and it provides the basis for resolving the conflicts between the singular- and plural-stance accounts of collective intentionality.

Fig. 2 sketches the CCSP model's treatment of group decision making and intention updating in collective activity, enabling an overview as well as providing a framework for the more detailed discussion that follows. As this sketch shows, a GAR intention, together with constraints introduced by the collective-action context in which it occurs, lead an agent to commit to participating in group decision-making processes. The group decision-making processes themselves require not only that agents be committed to participating, but also that they have certain reasoning and communicative capabilities and that they have well defined mechanisms (or rules) for establishing group decisions. The italicized font in the "capabilities" and "group decision-making mechanisms" boxes is used to distinguish the separate situation for people, who come readily equipped with such capabilities and are able to

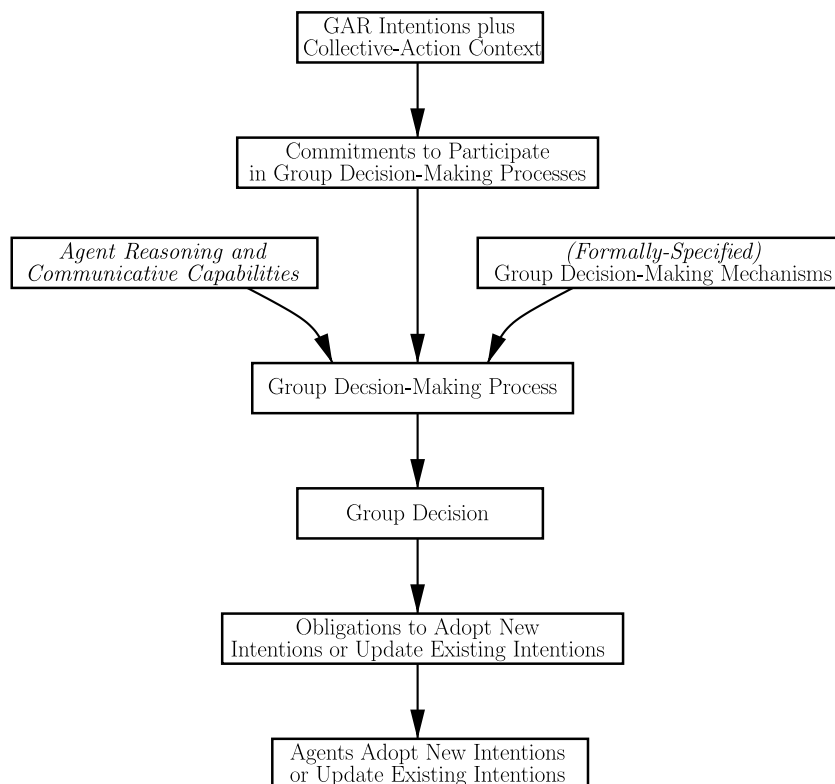


Fig. 2. An overview of group decision making in the CCSP model.

function with informal decision processes, from the situation for computer agents, which require computational subsystems for reasoning and communication as well as formally specified decision-making mechanisms that are amenable to automated reasoning algorithms. The result of a group decision-making process is the establishment of a group decision. In the context of collective activity, a group decision is an agreement that involves the making of a choice (e.g. choosing a recipe or allocating a task to an agent or subgroup). As such, a group decision entails obligations on the participants. In particular, each agent is obliged to adopt new intentions or modify existing intentions to be in accordance with the group decision. Typically, agents respond to such obligations by adopting new intentions or modifying their existing intentions.

Regardless of the particular issue at stake, all group decisions have certain characteristics in common. The CCSP model makes the commonalities clear and explicit by treating all group decisions uniformly. The advantages of this approach include the following: (1) the source of an agent's commitment to engage in group decision-making mechanisms (for any type of decision problem) is derived from general constraints on collective activity which, in turn, leads to a simpler specification of such commitments; (2) the differences among group decisions are characterized in terms of the different sets of intention updates normally associated with each kind of group decision; and (3) the formal specification of mechanisms for generating group decisions may be abstracted away from the semantics of group decisions. These important features of the CCSP model are described further below.

The remainder of this section presents the CCSP model and the architecture based on it. Section 3.1 explains the source of agents' commitments to group decision-making processes. Section 3.2 describes the essential differences in the reasoning–choosing–updating cycle as it occurs in collective activity as compared with single-agent activity. Section 3.3 specifies the obligations entailed by group decisions. Section 3.4 revisits the decision-making flow sketched in Fig. 2 to provide a more complete picture of the dynamics of intentions in collective activity using the concepts developed in Sections 3.1–3.3. Section 3.5 then discusses the formal specification of group decision-making mechanisms in the context of the CCSP model. Section 3.6 presents an architecture for a collaboration-capable agent based on the CCSP model.

3.1. The commitment to participating in group decision-making processes

In purely single-agent activity, an agent's intention to do some action α_i entails a commitment to engage in certain planning activity.³ This commitment derives directly from

the axiomatic requirement that an agent's plan be means-end coherent (Bratman, 1997). The agent must have or be committed to determining a recipe for α_i . Although the agent may involve others in the reasoning phase of its decision-making process – for instance if they can help it construct a recipe – the decisions (i.e. the choices in the choosing phase) are the agent's own.

In collective activity, each agent holds a GAR intention “that we do α_g ” (Clause 1 of Fig. 1), and, in this situation too, the requirement that an agent's intentions be means-end coherent entails a commitment toward decision making. However, the collective-activity context imposes two additional constraints: the common-content constraint (CCC) and the coordinated-cultivation requirement (CCR).

- Common-content constraint (CCC): The GAR intentions held by the participants in a collective, group activity that are related to that activity must all have the same content.
- Coordinated-cultivation requirement (CCR): The participants in a collective, group activity must update their GAR intentions related to that activity only in accordance with decisions of the group.

It is important to recognize the essentially different nature of these constraints. The CCC constrains the content of GAR intentions, requiring uniformity in their content. The CCR constrains the way in which changes may be made to GAR intentions. Thus, the CCC is a static constraint whereas the CCR is a dynamic one.

Violating the CCC makes a group's plan incoherent. If one agent intends “that we cook dinner today at seven o'clock” while another intends “that we cook dinner tomorrow at 9 o'clock”, then (at least on the basis of these intentions) these two agents will not act collectively; they will not cook dinner together. However, the CCC by itself is not sufficient, because the content of GAR intentions will not remain uniform if updating is not done “in sync”. It is this synchronizing of updating that the CCR ensures.

Thus, the CCR (operating on the dynamic aspect of intentions) ensures the persistence of the common-content property as a group's plan evolves over time. It explicitly rules out an agent's unilaterally updating its GAR intentions, with the term “unilateral” applying not only to individual participants, but also to proper subgroups. According to the CCSP model, agents who do not adhere to both constraints do not meet the conditions of collaborative activity. The two constraints working in concert are required for agents to coordinate the cultivation of their GAR intentions.

Means-end coherence requires an agent to eventually update its GAR intentions, but the CCR only allows updates in accordance with decisions of the group. Thus, an agent's commitment to decision making – which is inherent in any action-related intention – is transformed by the collective-action context into a commitment to participate in group decision-making processes.

³ When it is important to distinguish single-agent (or individual) actions from multi-agent (or group) actions, we use the subscripts i for individual, g for group.

3.2. Dynamics of single-agent and collective activity: Essential differences

The CCC and CCR highlight two essential differences between single-agent and group activity with respect to the reasoning–choosing–updating cycle. First, the two kinds of activity differ in terms of the locus of control in the choosing and updating phases. Second, they differ with respect to whether the entire cycle is primarily an internal, mental activity or has essential social components. We consider each of these differences in turn, for each looking first at individual activity and then at the situation for participants in group activities.

In the case of single-agent activity, the locus of control for all means-ends reasoning, decision making and intention updating is the (single) agent performing the action. For instance, in the salad-making example, the cook decides where to get his tomatoes and then updates his plan to include obtaining tomatoes from that source. Although various considerations of other people may play a role in his reasoning, the decision is his. Once he has made a choice, he simply updates his intention in accordance with his decision. Were he not to do so, he would be considered to be irrational or otherwise to lack typical reasoning capabilities.

In group activity, the locus of control is the group. Although responsibility for decision making is typically distributed, intention updating requires coordination on the part of all participants. Agents may reason separately, and in a distributed fashion, but they must eventually choose collectively. Furthermore, as described in Section 3.1, the actual change in their intentions is constrained by the group context. In effect, the CCC and CCR enable the group to exert influence over each participant's internal intention-update process, rather than its being solely under the agent's own control. Even in organizations with rigid command structures, where the reasoning and choosing is left to a team leader, the rest of the agents must nonetheless abide by that leader's choices when updating their intentions. More generally, and crucially, collaborating agents must coordinate the updating of their GAR intentions. Were a participant in such an activity not to adhere to this stricture, that participant would be considered uncooperative or in some other way to be violating social norms.⁴

In short, in the case of individual activity, a single agent typically controls all phases of the reasoning–choosing–updating cycle, and intention updating follows rules of individual utility and rationality. In contrast, in the case of group activity, intention updating is constrained by various social norms, expectations, and constraints.

In the single-agent case, typically the entire reasoning–choosing–updating cycle is internal to the agent. In particular, choosing and updating are both internal, mental

actions that are very tightly coupled; the choosing directly effects the intention updating.⁵ An individual's decision is a mental action whereby an agent (internally) commits itself to some state of affairs, and individual decisions typically involve the immediate adoption of new intentions or modification of existing intentions.

In contrast, in the group case, even though the end result of a group decision also typically involves the adoption of new intentions or the modification of existing intentions by the participants, there is a much clearer dividing line between the choosing and updating phases. The choosing phase is an external, or social, act, and only the updating is internal, or mental.⁶ The CCSP model requires that a group decision-making mechanism (as described in Section 3.5) be specified for the social, choosing phase of the cycle, to reflect this key characteristic. For people, the decision-making mechanism might be informal; for instance, two friends might establish their decision to drive to NY together by making certain declarations and shaking hands (to “seal the deal”). Once made, a group decision establishes a new fact, external to the members of the group (i.e. not simply elements of their mental states), namely the proposition that the group has decided thus-and-so. For instance, the proposition that “the group has decided to drive to NY together” becomes true if a group makes that decision. Of course, in addition to the external fact, each of the participants in the group has a set of beliefs related to the decision, including ones about the decision itself.

In the CCSP model, the updating phase is coordinated by obligations. A group decision is an agreement which, like any agreement, entails certain obligations. These obligations typically lead the participants to adopt new intentions or to update their existing GAR intentions to be in accordance with that decision. The actual updating is an internal, mental activity of each individual. However, as a result of all of the participants updating their GAR intentions “in sync”, the distributed set of such updated GAR intentions may be seen to be the group correlate to the internal sense of updating in the individual case.

3.3. Decisions and obligations

Different kinds of group decisions engender different types of obligations. For example, when a group *GR* makes a group decision to form a new collaborative team for doing some activity α , each member of *GR* becomes obliged to adopt a new GAR intention (“that *GR* does α ”) and, furthermore, to constrain the updating of that GAR intention as stipulated by the CCR. However, such other kinds of group decisions as selecting a recipe and deciding which

⁵ These mental actions may be supported by external, physical actions. For instance, the agent might flip a coin or ask someone else to tell him what to do.

⁶ Singh (1998) has highlighted the importance of distinguishing the social and mental aspects of commitments.

⁴ As much work in economics argues, what is individually rational may, in fact, not be best for the group, and social norms are sometimes introduced to affect the balance between individual and group good.

agent will do a particular sub-act occur in the context of existing GAR intentions and yield revisions to them.

Decisions about which agent will do a sub-act (or, task), which are typically called “task-allocation decisions”, are a special case in that they require not only the updating of GAR intentions, but also the adoption of new, subsidiary intentions. For example, when a task β_i is assigned to an agent G , that agent must adopt a subsidiary intention to do β_i . Crucially, this new intention is not subject to the CCR, because the action is an action which agent G alone is doing, and, as a result, G is the only relevant decision maker. However, the other agents in the group must adopt subsidiary intentions whose content may be glossed as “that G is able to do β_i ”.⁷ Such intentions ensure that the other agents do not hamper G ’s efforts to do β_i and may lead those agents to offer collaborative support (Grosz & Kraus, 1999). Similarly, when allocating a task β_g to a subgroup SG , the members of SG must form a collaborative team whose goal is that they do β_g . In particular, each member of SG must adopt a subsidiary GAR intention (“that we [i.e. SG] do β_g ”). This new GAR intention is, of course, also subject to the CCR: it may be updated only in accordance with decisions made by the subgroup SG . Thus, the allocation of tasks to agents or subgroups includes the delegation of decision-making authority.

Fig. 3 lists the obligations entailed by each of a core class of group decisions that arise in collaborative group activity, including decisions to establish a new collaborative group or to coordinate the updating of related GAR intentions in an existing collaboration. The parameter-binding, recipe-selection and task-allocation decisions oblige the participants to update their relevant GAR intentions.⁸ In the figure, the particular updates are specified by giving the relevant portion of the content of the GAR intentions both before and after the group decision. For task-allocation decisions, the subsidiary intentions that the participants are obliged to adopt are derived directly from the SharedPlans specifications. In line with the discussion above and the original SharedPlans formalization, the intentions adopted by the agent(s) to whom a task is allocated and those adopted by the rest of the agents in the group are different.

It is important to note that the CCSP model is not limited to the types of group decisions shown in Fig. 3, but rather it provides for the uniform treatment of a wide variety of group decisions. For example, agents seeking to coordinate their activities in the presence of temporal constraints may need to add new temporal constraints to their

- | |
|---|
| <ul style="list-style-type: none"> • Group Decision to Form a Team GR to do α Collaboratively:
Obligations: Each agent is obliged to adopt a new GAR intention (constrained by the CCR) “that GR does α.” • Group Decision to Bind a Parameter p to the value v:
Precondition: Prior GAR intention contains an unbound parameter p.
Obligations: Each agent is obliged to update its GAR intention by replacing all occurrences of p by v. • Group Decision to Select a Recipe R_α with sub-acts β_1, \dots, β_n:
Precondition: Prior GAR intention includes “that GR does α.”
Obligations: Each agent is obliged to update its GAR intention to include “that some agent or subgroup does β_1, \dots, and that some agent or subgroup does β_n.” • Group Decision to Allocate a Sub-Act β_i to an Agent G:
Precondition: Prior GAR intention includes “that some agent or subgroup does β_i.”
Obligations: Each agent is obliged to update its GAR intention to include “that G does β_i.”
G is obliged to adopt a new, subsidiary intention to do β_i.
Each agent other than G is obliged to adopt a new intention that G be able to do β_i. • Group Decision to Allocate a Sub-Act β_g to a Subgroup SG:
Precondition: Prior GAR intention includes “that some agent or subgroup does β_g.”
Obligations: Each agent is obliged to update its GAR intention to include “that SG does β_g.”
Each agent in SG is obliged to adopt a new, subsidiary GAR intention (constrained by the CCR with respect to the subgroup SG) “that SG does β_g.”
Each agent not in SG is obliged to adopt a new intention “that SG be able to do β_g.” |
|---|

Fig. 3. The obligations entailed by group decisions.

GAR intentions (Hunsberger, 2002a). Decisions about timing would oblige the participants to update the content of their GAR intentions in accordance with these decisions.

Finally, the CCSP model also accommodates another important aspect of group decision making, namely, that decisions about group activities are frequently interdependent and, thus, agents may need to combine multiple, related decisions. For example, the result of a lengthy multi-agent negotiation might be a bundle of interdependent task-allocation and parameter-binding decisions. In the CCSP model, the set of obligations entailed by a bundle of decisions is simply the union of the sets of obligations individually entailed by the decisions in that bundle.

3.4. Dynamics revisited

The CCSP approach to group decision making in collective activity, as previously illustrated in Fig. 2, explains the way in which agents holding GAR intentions constrained by the CCR become committed to participating in group decision-making mechanisms. Subsequent participation in such mechanisms typically leads to the establishment of group decisions which, being agreements, entail certain obligations on the participants. Group decisions to select

⁷ See the original SharedPlans formalization (Grosz & Kraus, 1996) for the rationale for the “able to” formulation.

⁸ Pasula (1996) highlighted the need for participants in a group activity to coordinate their decisions about action parameters, especially those that might be used across sub-acts being done by different agents; Hadad (2002) addressed the need for coordination with respect to temporal constraints. The treatment here is not only more general, allowing both kinds of decisions to be handled uniformly, but also explicitly represents the social obligation aspect of intention-update in collective activity.

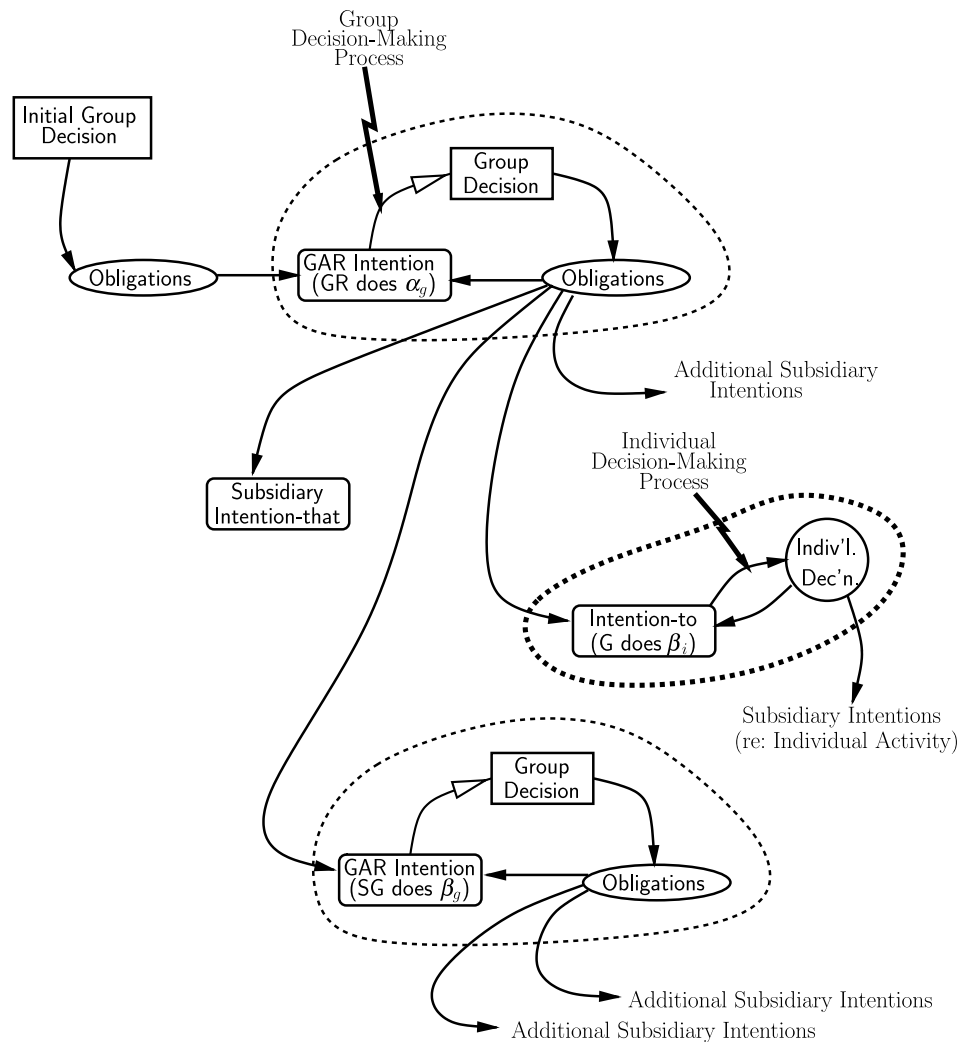


Fig. 4. A participant's view of the dynamics of a hierarchical SharedPlan.

a recipe, bind a parameter or allocate a task entail obligations on participants to update their existing GAR intentions. Task-allocation decisions additionally entail obligations on participants to adopt new, subsidiary intentions. Since these subsidiary intentions may involve further commitments to decision-making processes and intention updating, the typical result is a complex hierarchy of intentions which, together with related mutual beliefs, constitute the group's SharedPlan. Fig. 4 illustrates a participant's view of this hierarchy.

As described in Section 3.3 the subsidiary intentions come in three varieties. First, when a task β_g is allocated to a subgroup SG , each agent in that subgroup is obliged to adopt a GAR intention "that SG does β_g ", which is subject to the CCR with respect to that subgroup. Just as the agents in GR coordinate the cultivation of their GAR intentions concerning α_g , the agents in SG coordinate the cultivation of their GAR intentions concerning β_g . Second, when a task β_i is allocated to an individual G , that agent typically adopts a subsidiary intention to do β_i , in which case G is individually responsible for completing its sub-

plan for β_i .⁹ Finally, whether a task is allocated to an individual or a subgroup, the rest of the agents in the group are obliged to adopt subsidiary intentions that the selected individual or subgroup be able to do the assigned task.

To summarize, the process of incrementally completing a plan by moving from a set of GAR intentions constrained by the CCR, to a new group decision, to updated GAR intentions and newly adopted subsidiary intentions is recursive. The recursion stops when all intentions in the hierarchy have been fully specified and the single-agent actions associated with intentions at the bottom of the hierarchy have all been successfully executed (i.e. when the SharedPlan is complete).

The entire process is typically initiated by a group decision to form a collaborative group to do some α_g . However, as Bratman (1999) has argued, collaborative activity may also arise spontaneously without any initiating agreement or mutual obligations. The CCSP model's characterization

⁹ The incremental refinement of single-agent subplans has been thoroughly addressed in the literature and will not be discussed further here.

of collaborative activity in terms of GAR intentions and the coordinated cultivation requirement is consistent with Bratman's view. However, this paper focuses on collective activity initiated by an explicit group decision.

3.5. Group decision-making mechanisms

As observed earlier, the CCSP model's uniform treatment of group decisions allows the semantics of group decisions to be distinguished from the mechanisms used to generate group decisions. This section addresses the formal specification of such mechanisms.

People are quite adept at using informal mechanisms for establishing group decisions. For example, Bill and Ted might come to an agreement by nodding their heads and winking meaningfully. However, computer agents, which use automated reasoning systems, require formally specified mechanisms for establishing group decisions.¹⁰ Thus, the CCSP model requires that group decision-making mechanisms (GDMMs) be formally specified to enable agents to reliably establish group decisions for the purpose of coordinating the updating of their GAR intentions. However, to increase agent flexibility, the model leaves open the particular kinds of GDMMs agents might use to generate group decisions in any given circumstance.

Different GDMMs may operate in very different ways. For example, a *Unanimous Consent* mechanism might generate a group decision only if every participant explicitly accepts a given proposal. Alternatively, a mechanism for allocating tasks might take the form of an auction. With some GDMMs, the parties to a group decision may always learn of the generation of a new decision simultaneously; with others, they may learn of a new decision at different times. However, all of the agents must eventually be able to verify that a group decision has in fact been established.

In previous work (Hunsberger, 2002b), we argued that the specification of a GDMM must include (1) the possible inputs an agent can make into the mechanism; (2) the conditions under which agents may make those inputs; (3) rules for determining which combinations of agent inputs establish group decisions; and (4) a method for making the new decision known to all the members of the group. We presented a general framework for formally specifying GDMMs using Dynamic Deontic Linear Time Temporal Logic (DDLTLB) (Dignum & Weigand, 1995a, 1995b). In that work, agent inputs into a mechanism are in the form of declarative speech-acts (Austin, 1962; Searle, 1995). The conditions under which certain inputs are allowed, the rules for how certain combinations of inputs establish group decisions, and the resulting obligations are all easily expressed in DDLTLB. We illustrated the

use of that framework by specifying a sample *Unanimous Consent* mechanism and proved theorems characterizing the robustness and reliability of that mechanism.

In a related paper (Hunsberger & Grosz, 2000), we presented an auction-based mechanism that agents can use when faced with some proposed group activity to decide whether to initiate a collaborative effort. That mechanism, which can also be used to generate bundles of task-allocation decisions for a pre-existing collaborative effort, allows the participants to protect any private, pre-existing commitments they might have by including temporal constraints in their bids. Elsewhere (Hunsberger, 2002b), we have provided algorithms that agents can use to facilitate the bid-generation process needed to participate in such mechanisms.

3.6. The CCSP agent architecture

The CCSP architecture for a collaboration-capable agent, depicted in Fig. 5, extends the IRMA architecture for resource-bounded agents developed by Bratman et al. (1988). In the figure, ovals represent collections of information (e.g. beliefs, desires, intentions, decisions or obligations) loosely referred to as “databases”, whereas boxes represent computational modules that process the information in those databases. We use the term “component” to refer to either databases or computational modules. The solid arrows represent the flow of control from a computational module to a database; the dashed arrows represent the flow of information from a database to a computational module. The components of the IRMA architecture appear on the left side of the figure; additions for collective activity appear on the right.

The IRMA architecture addresses the ways in which means-ends reasoning and the weighing of alternative courses of action interact when agents are resource bounded. It deals with the dynamics of intentions for single-agent activities. In particular, its design embodies the constraining roles of plans and intentions. New opportunities, whether internally or externally generated, and options generated by means-ends reasoning processes are filtered for compatibility with existing plans and intentions. This filtering process eliminates options for future activity that would conflict with existing intentions, reflecting the focusing effect of plans (which is necessary because agents are resource bounded), but it allows an agent to change its mind should a new, high-priority option arrive. As a result, an agent's partial plans are refined by means-ends reasoning, but in ways that are compatible with the full set of active intentions and plans. An agent is committed to what “it is doing” and the “characteristic roles” of those commitments are to drive means-ends reasoning and to constrain the set of options that the agent considers (Bratman, 1997).

The IRMA architecture reflects the tight link between the choosing and intention-updating phases in single-agent activity that was described in Section 3.2. In particular,

¹⁰ Of course, even people employ formal mechanisms for generating group decisions when the stakes are high (e.g. as when a buyer, a seller, their attorneys, and a bank together agree to the conditions for the sale of a house).

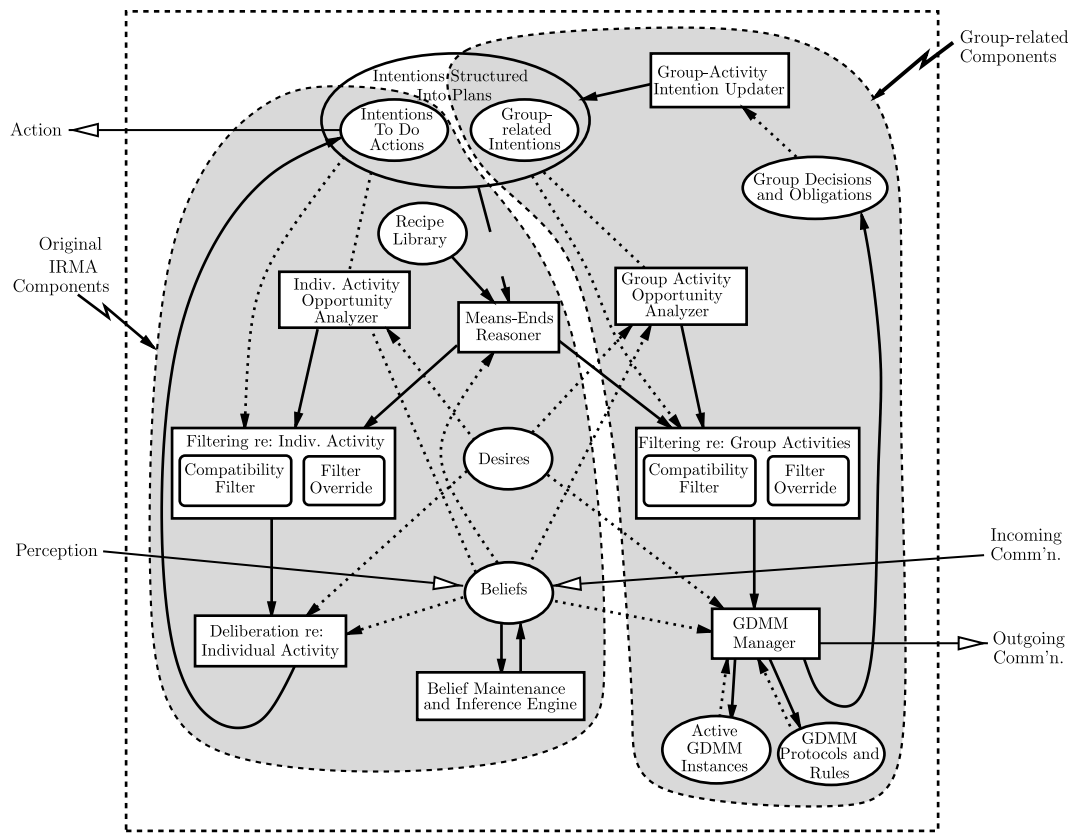


Fig. 5. The CCSP architecture for a collaboration-capable agent.

single-agent decisions are not explicitly represented in the IRMA architecture.

To extend this architecture to treat collaborative, group activities requires adding components to handle an agent's participation in complex group decision-making processes. These new components reflect that agents not only are resource bounded, but also, when acting collaboratively, are constrained by the CCR from making certain decisions unilaterally and, in addition, that group decisions oblige agents to adopt new intentions or update existing intentions. Group decisions and their corresponding obligations are explicitly represented on the "group activity" side of the architecture making explicit the important distinction between the social act of choosing and the internal, mental act of updating one's intentions. Furthermore, whereas the "individual activity" side of the architecture contains a single "Deliberation" module, the "group activity" side contains both a "GDMM Manager" and a "Group-Activity Intention Updater". The "GDMM Manager" handles the agent's interactions with other agents as they deliberate together to refine their partial plans and further specify their GAR intentions. The "Group-Activity Intention Updater" meets obligations by adopting new intentions or updating existing intentions.

To further illustrate the roles of the new components and their operation during the evolution of a typical collaborative plan, we will use the example of a group of musicians collaborating to perform a concert (or "gig"). The

"Group-Activity Opportunity Analyzer" is responsible for generating new candidates for collaborative group activity by monitoring the database of beliefs (which includes beliefs arising from incoming communication or other perception of the external environment). For example, an agent might learn of a new opportunity for playing a gig from an incoming email message. When such an opportunity is generated, it is sent to the "Filtering re: Group Activities" module which determines whether that option would be compatible with the other activities that the agent is already engaged in, which are represented in the database labeled "Intentions Structured into Plans". For instance, if the agent is already scheduled to play a gig on Friday night, an opportunity to play another gig that same night on the other side of the country would probably be discarded. However, the "Filter Override" module can allow certain high-priority opportunities to survive the filtering process even if they might conflict with an agent's pre-existing commitments.

The group-related compatibility and override filters differ from those concerning purely single-agent activities by possibly taking into account the preferences and obligations not only of the agent, but also of other participants or the group as a whole. The process of reconciling potential conflicts among intentions is more complex in the context of group activities because it requires the weighing of trade-offs between individual and group good (Grosz, Kraus, Sullivan, & Das, 2002; Grosz, Kraus, Talman,

Stossel, & Havlin, 2004). For instance, an agent might determine that the needs of the group were sufficiently great that it should reconsider its personal intention that led to a compatibility conflict.

If the new opportunity survives the filtering process, the “GDMM Manager” uses the “GDMM Protocols and Rules” database, which includes the specifications of all of the mechanisms known to the agent, to select a particular GDMM and initiate a group decision-making process using that GDMM. For example, the agent might decide to use a *Unanimous Consent* mechanism and send an email message to its fellow musicians proposing that they agree to play this new gig on Friday night. The newly initiated mechanism is referred to as an “instance” of a GDMM and is recorded in the “Active GDMM Instances” database.

Replies from other agents (“incoming communication”) enter the “Beliefs” database and are processed by the “GDMM Manager” which updates the status of the GDMM instance in the “Active GDMM Instances” database. If all the other agents agree, then the agent who originated the proposal announces the group decision (“outgoing communication”) and records it in its own “Group Decisions and Obligations” database. Participation in mechanisms initiated by other agents is handled similarly, except that the “GDMM Manager” might learn of the group’s decision from an announcement made by another agent.

Group decisions to engage in new group activities typically lead the “Group-Activity Intention Updater” to create a new GAR intention and enter it into the database of “Intentions Structured Into Plans”. When all of the agents in the group have adopted corresponding GAR intentions, each constrained by the CCR, the group is, for the first time, in the state of having a SharedPlan.

The new GAR intention requires that certain decisions be made (e.g. how to get to the gig and what equipment to bring). As in the single-agent case, the “Means-Ends Reasoner” generates potential solutions to such decision problems. However, in the multi-agent case, these potential solutions may involve actions to be performed with or by other agents (e.g. Bill and I borrow a car or Charlie brings the high-powered equipment). Each potential solution must be filtered to ensure that any action to be performed by this agent, whether individually or with others, is compatible with its existing intentions and plans. Each potential solution that survives this filter is sent to the “GDMM Manager”, which determines possible communications the agent might make to the rest of the group (in the context of a GDMM) to get them to consider that potential solution.

When the agent learns from other agents about new group decisions, the agent’s “GDMM Manager” records the new decision (e.g. to rent a van) and associated obligations into the “Group Decisions and Obligations” database. Such obligations typically lead the “Group Activity Intention Updater” to enter new intentions into the database of “Intentions Structured into Plans” or to

update intentions already in that database. In the case of a task-allocation decision, say, “Bob to drive van”, the updater enters new, subsidiary intentions concerning the driving of the van. If “Bob” refers to this agent, then the intention is an intention to drive the van, which may lead to subsidiary single-agent planning activity (on the left side of the diagram). If “Bob” refers to a different agent, then the new intention is an intention that Bob be able to drive the van.¹¹

4. Discussion and related work

The CCSP model draws on and contributes to two different, but complementary areas, one related to the features of the model itself and the second to the architecture based on it. In this section, we first show that the CCSP model satisfies the requirements for an account of collective intentionality argued for by others, including both plural- and singular-stance approaches. In particular, it meets the plural-stance approaches’ criteria through the use of constraints on intention updating (represented in the CCC and CCR), rather than through the introduction of a new kind of plural-subject intentional attitude. We then compare the CCSP-based architecture to alternative collaborative multi-agent architectures.

Tuomela (2004a) argues that collaborating agents employ a we-mode of intending that is different from and not reducible to the ordinary I-mode of intending. For example, he claims that in the group case, “the goal state or event comes about due to the collective effort by, or at least under the collective guidance of, the group members.” Elsewhere (Tuomela, 1995) he stipulates that we-intenders must be disposed to engage in certain forms of practical reasoning, for example, to derive their individual contributions to the collaborative activity; and that collaborating agents require a “socially-existing authority system” (or group-will formation system) and must be motivated to use it. The CCSP model, though its use of GAR intentions constrained by the CCC and CCR, meets all of these requirements. Furthermore, it highlights the CCR as the source of an agent’s motivation to participate in group decision-making mechanisms, provides criteria for formally specifying such mechanisms, and specifies the particular obligations that are entailed by several common types of group decisions.

Searle (1990) argues that “the notion of a we-intention, of collective intentionality, implies the notion of *cooperation* [italics his].” The following features of the CCSP model handle crucial elements of Searle’s account: (1) an agent’s motivation to participate in group decision-making processes; (2) the group’s coordinated updating of corresponding GAR intentions; (3) the persistence of the common-content property and, thus, the coherence of a

¹¹ Because the cultivation of this latter form of intention is not the focus of this paper, such intentions are not shown in the database of “Intentions Structured into Plans”.

group's evolving SharedPlan; and (4) the commitment of agents to avoid interfering with the efforts of their fellow participants. In addition, hierarchical SharedPlans represent the relationships between GAR intentions and subsidiary plans, including in settings in which the higher-level intention is cooperative while the lower-level intentions are competitive (as happens, for example, in organized sports).

Searle also argues that it is important to be able to distinguish a set of agents acting as a collective from their acting as individuals who just happen to accomplish the same goal. He presents as an example the case of a group of agents collectively intending to run to a common point from the case of a set of agents each independently intending to run to that same point. The CCSP model handles this distinction through its requirement in the collective case that each agent have a GAR intention (“that we run to the common point”) and a derivative intention (“to run to the common point”) that is correctly linked to this GAR intention. In addition, each agent's GAR intention is subject to the CCR, which ensures the coherence of the collective intention. Searle also argues that it is important to be able to distinguish cases where the content of the collective intentionality (“we are running a pass play”) is different from the content of derivative intentions (“I am blocking the defensive end”). The CCSP model accommodates this distinction with the same mechanism: I intend “that we run a pass play” and I have a subsidiary intention “to block the defensive end”.

Finally, Searle presents two versions of a Business School scenario that he claims form a counter-example to any account of collective intentionality based solely on ordinary, singular-subject intentions (which he calls “I-intentions”). In the scenario, graduates of the Business School have learned that they may help humanity by pursuing their own selfish interests. In the first version, each graduate “intends to pursue his own selfish interests without reference to anybody else”, whereas in the second version, the graduates “form a pact to the effect that they will all go out together and help humanity by way of each pursuing his own selfish interests.” The CCSP model, like the SharedPlans formalization on which it is based, distinguishes these two versions by whether or not the graduates all have GAR intentions “that we help humanity by each pursuing our own selfish interests”. As described in an earlier paper (Grosz, 1999), the SharedPlans formalization has mechanisms that handle additional features of these scenarios argued for by Searle.

Gilbert (2000) argues that any model of collective intentionality (or “shared intention”) must account for the following features: (1) that shared intentions entail certain obligations (e.g. “not to act contrary to the shared intention” or “to promote the fulfillment of the shared intention”), as well as corresponding rights and entitlements; (2) that participants in a shared intention “are not in a position to remove its constraints unilaterally”; and (3) “that there could be a shared intention to do such-and-such

though none of the participants personally intend to conform their behavior to the shared intention”. The CCSP model handles the first feature through its requirement that group decisions entail certain obligations, in particular, obligations to update existing GAR intentions or to adopt related subsidiary intentions. Through the CCR constraint against the unilateral updating of GAR intentions, the CCSP model has the second feature. We have not attempted to capture the third feature, because it does not accord with our conception of collective intentionality.

Bratman (1999) presents a model of Shared Cooperative Activity that employs ordinary, singular-subject intentions, but with the highly specialized content: “that we J in accordance with and because of [our intentions] and meshing subplans of [our intentions].” In contrast, GAR intentions in the CCSP model have the much simpler content, “that we do α ”, but their updating is constrained by the CCR. This separation of the content of an intention from constraints on how that content can be updated enables the CCSP agent architecture to employ essentially the same means-ends reasoner used in the single-agent case, while adding a new module – the “GDMM Manager” – to handle the deliberation and communication processing required to manage the agent's participation in group decision-making mechanisms.

A wide-range of multi-agent systems architectures and coordination mechanisms have been developed to enable the construction of computer systems that function like members of a team carrying out a joint activity (Giampapa & Sycara, 2002; Decker & Lesser, 1992; Jennings, 1995; Tambe, 1997, *inter alia*). In the context of this paper, which has focused on the central role of the dynamics of intention in an explanation of collective intentionality, the most relevant architectures are those based on previous models and formalizations of collaboration and teamwork, GRATE* (Jennings, 1995) and STEAM (Tambe, 1997).

GRATE*'s architecture is based on the “Joint Responsibility” model of collaborative activity, an extension of Joint Intentions (Levesque et al., 1990) that includes a requirement that collaborating agents be committed to an agreed-upon “solution” to their joint activity—where a solution includes a recipe (i.e. a set of tasks) plus a complete set of task assignments. As a result, the GRATE* agent architecture incorporates protocols for recipe selection and task allocation. However, like the Joint Intentions formalization on which it builds, GRATE* does not address a group's incremental refinement of their partial plan. STEAM, which draws on both the Joint Intentions and SharedPlans formalizations, incorporates similar constituents through its use of the SharedPlans-specification requirements that agents be committed to choosing recipes and assigning agents to sub-acts in those recipes. Each of these systems also addresses various problems that arise in going from the formalizations to implementations. For instance, STEAM incorporates decision-theoretic mechanisms for deciding when it is worthwhile for an agent to communicate (Tambe, 1997).

These systems treat only particular types of group decisions, and they treat each type separately. The lack of a general treatment of group decision making is not surprising, because the theories on which these systems were based also did not explicitly represent the complex relationships between group decisions, their entailed obligations and the intentions that participants modify or adopt in response to them. In describing the development of STEAM, Tambe (1997, p. 90) remarks on the importance of the formalizations of collaboration saying, “In the absence of both the explicit representation of team goals and plans, as well as commitments and responsibilities they engender, agents are often forced to rely on the problematic domain-specific coordination plans, leading to ... teamwork failures”. The architecture presented in this paper is the first to identify group decision making as a general, pervasively needed capability; to represent explicitly the obligations group decisions entail; and, by embodying the CCC and CCR constraints, to link decision making and intention updating explicitly. It is based on the CCSP model, which is the kind of formalization to which Tambe refers, and it thus provides the basis for more general and robust collaboration-capable agents similar to the way in which the IRMA architecture provides a basis for resource-bounded single-agent systems.

5. Conclusions

This paper defines the coordinated cultivation of SharedPlans (CCSP) model of the dynamics of intention in collaborative activity. The CCSP model makes explicit the connection between group decisions and intention update in collective, group activity through its augmentation of the SharedPlans formalization with the common-content constraint and coordinated cultivation requirement, specification of an explicit role for obligations in mediating between group decisions and intention update, and the requirement that group decision-making mechanisms be formally specified. This integrated treatment of group decision making and intention updating provides the basis for defining a more general architecture for collaboration-capable agents. The paper defines such an architecture, extending the standard IRMA architecture for resource-bounded agents to be an architecture for resource-bounded agents capable of acting collaboratively.

Prior characterizations of collective intentionality from both the singular- and plural-stances have considered only static properties of intentions, essentially asking what the content of these intentions should be. Their concerns have been essentially equivalent to only one of the two constraints in the CCSP model, the common-content constraint. This paper has argued that the dynamics of intentions play a central role in collective intentionality. The second CCSP constraint, the coordinated cultivation requirement, is essential for there to be maintenance of the common-content property. By having both constraints, the CCSP enables an explanation of collective intentionality

that rests on purely singular-stance constructs while capturing the essential properties argued for in accounts taking the plural stance. It thus resolves a long-standing debate about the nature of intentions in collective group activity.

Acknowledgment

The research reported in this paper was supported in part by National Science Foundation Grants IIS-0222892 and IIS-9978343 to Harvard University.

References

- Austin, J. L. (1962). *How to do things with words*. Cambridge, MA: Harvard University Press.
- Bratman, M. E. (1997). *Intentions, plans, and practical reason*. Cambridge, MA: Harvard University Press.
- Bratman, M. E. (1999). *Faces of intention: selected essays on intention and agency*. Cambridge University Press.
- Bratman, M. E., Israel, D. J., & Pollack, M. E. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4, 349–355.
- Castelfranchi, C. (1995). Commitments: from individual intentions to groups and organizations. In V. Lesser (Ed.), *First international conference on multi-agent systems (ICMAS-95)* (pp. 41–48). The MIT Press.
- Cohen, P. R., Levesque, H. J., & Smith, I. (1997). On team formation. In G. Holmstrom-Hintikka & R. Tuomela (Eds.), *Contemporary action theory* (Vol. II, pp. 87–114). Netherlands: Kluwer Academic Publishers.
- Cohen, P. R., Morgan, J., & Pollack, M. E. (1990). In *Intentions in communication*. Cambridge, MA: The MIT Press.
- Decker, K. S., & Lesser, V. R. (1992). Generalizing the partial global planning algorithm. *International Journal of Intelligent Cooperative Information Systems*, 1(2), 319–346.
- Dignum, F., & Weigand, H. (1995a). Communication and deontic logic. In R. Wieringa & R. Feenstra (Eds.), *Information systems, correctness and reusability* (pp. 242–260). World Scientific.
- Dignum, F., & Weigand, H. (1995b). Modelling communication between cooperative systems. In *Conference on advanced information systems engineering* (pp. 140–153).
- Giampapa, J. A., & Sycara K. (2002). Team-oriented agent coordination in the retina multi-agent system. Technical Report CMU-RI-TR-02-34, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, December 2002. Presented at AAMAS 2002 Workshop on Teamwork and Coalition Formation.
- Gilbert, M. (2000). *Sociality and responsibility*. New York: Rowman & Littlefield Publishers Inc.
- Grosz, B. (1999). The contexts of collaboration. In K. Korta, E. Sosa, & X. Arrazola (Eds.), *Cognition, Agency and Rationality* (pp. 175–187). Dordrecht: Kluwer.
- Grosz, B. J., Hunsberger, L., & Kraus, S. (1999). Planning and acting together. *AI Magazine: Special Issue on Distributed, Continual Planning*, 23–34.
- Grosz, B. J., & Kraus, S. (1993). Collaborative plans for group activities. In *Proceedings of the 13th international joint conference on artificial intelligence (IJCAI-93)*. San Mateo, CA, USA: Morgan Kaufmann.
- Grosz, B. J., & Kraus, S. (1996). Collaborative plans for complex group action. *Artificial Intelligence*, 86, 269–357.
- Grosz, B. J., & Kraus, S. (1999). The evolution of SharedPlans. In M. Wooldridge & A. Rao (Eds.), *Foundations of rational agency, number 14 in applied logic series* (pp. 227–262). The Netherlands: Kluwer Academic Publishers.
- Grosz, B. J., Kraus, S., Sullivan, D. G., & Das, S. (2002). The influence of social norms and social consciousness on intention reconciliation. *Artificial Intelligence*, 142, 147–177.

- Grosz, B. J., Kraus, S., Talman, S., Stossel, B., & Havlin, M. (2004). The influence of social dependencies on decision-making: initial investigations with a new game. In *Proceedings of the third international joint conference on autonomous agents and multi-agent systems (AAMAS-2004)* (pp. 782–789). ACM Press.
- Grosz, B. J., & Sidner, C. L. (1990). Plans for discourse. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in communication (Chapter 20)*. Cambridge, MA: The MIT Press.
- Hadad, M. (2002). *Combining Cooperative Planning and Temporal Reasoning in Dynamic Multi-Agent Systems*. Ph.D. Thesis, Bar Ilan University, Ramat Gan, Israel.
- Hunsberger, L. (1999). Making SharedPlans more concise and easier to reason about. In *Agent architectures, theories and languages V. Lecture notes in artificial intelligence* (Vol. 1555, pp. 81–98). Springer-Verlag.
- Hunsberger, L. (2002a). Algorithms for a temporal decoupling problem in multi-agent planning. In *Proceedings of the eighteenth national conference on artificial intelligence (AAAI-2002)*.
- Hunsberger, L. (2002b). *Group Decision Making and Temporal Reasoning*. Ph.D. Thesis, Harvard University. Available as Harvard Technical Report TR-05-02.
- Hunsberger, L., & Grosz, B. J. (2000). A combinatorial auction for collaborative planning. In *Fourth international conference on multi-agent systems (ICMAS-2000)* (pp. 151–158). IEEE Computer Society.
- Jennings, N. R. (1995). Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artificial Intelligence*, 75, 195–240.
- Kinny, D., Ljungberg, M., Rao, A. S., Sonenberg, E., Tidhar, G., & Werner, E. (1994). Planned team activity. In C. Castelfranchi & E. Werner (Eds.), *Artificial social systems. Lecture notes in artificial intelligence*. Amsterdam: Springer-Verlag.
- Lansky, A. L., & Georgeff, M. (1987). Procedural knowledge. Technical Report 411, SRI International, January 1987.
- Levesque, H. J., Cohen, P. R., & Nunes, J. H. T. (1990). On acting together. *Seventh national conference on artificial intelligence* (Vol. 1, pp. 94–99). AAAI Press/MIT Press.
- Lochbaum, K. E. (1994). *Using Collaborative Plans to Model the Intentional Structure of Discourse*. Ph.D. Thesis, Harvard University, October 1994.
- Meyer, J.-J. Ch., van der Hoek, W., & van Linder, B. (1999). A logical approach to the dynamics of commitments. *Artificial Intelligence*, 113.
- Ortiz, C. L. Jr., (1999). Introspective and elaborative processes in rational agents. *Journal of the Annals of Mathematics and Artificial Intelligence*.
- Panzarasa, P., Jennings, N., & Norman, T. J. (2001). Formalising collaborative decision-making and practical reasoning in multi-agent systems. *Journal of Logic and Computation*, 11(6), 1–63.
- Pasula, H. (1996). Design of a collaborative planning system. Seniors Honors Thesis, Harvard University.
- Pollack, M. E. (1990). Plans as complex mental attitudes. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in communication (Chapter 5)*. Cambridge, MA: The MIT Press.
- Pynadath, D. V., Tambe, M., Chauvat, N., & Cavedon, L. (2000). Toward team-oriented programming. In N. R. Jennings & Y. Lesperance (Eds.), *Intelligent agents VI, agent theories, architectures, and languages. Lecture notes in artificial intelligence* (Vol. 1757). Berlin Heidelberg: Springer-Verlag.
- Rao, A. S., & Georgeff, M. P. (1991). Modeling rational agents within a BDI-architecture. In *Proceedings of the second conference on principles of knowledge representation and reasoning (KR'91)* (pp. 473–484). San Mateo, CA: Kaufman.
- Rao, A. S., & Georgeff, M. P. (1992). An abstract architecture for rational agents. In *Proceedings of the third conference on principles of knowledge representation and reasoning (KR'92)* (pp. 439–449). Morgan Kaufman Publishers Inc.
- Rao, A. S., & Georgeff, M. P. (1995). BDI agents: from theory to practice. In V. Lesser (Ed.), *First international conference on multi-agent systems (ICMAS-95)* (pp. 312–319). The MIT Press.
- Rao, A. S., & Georgeff, M. P. (1995). The semantics of intention maintenance for rational agents. In *Proceedings of the 14th international joint conference on artificial intelligence (IJCAI-95)* (pp. 704–710).
- Searle, J. R. (1990). Collective intentions and action. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in communication*. Cambridge, MA: The MIT Press.
- Searle, J. R. (1995). *The Construction of Social Reality (Chapter 19)*. Allen Lane, London.
- Singh, M. P. (1998). Agent communication languages: rethinking the principles. *IEEE Computer*, 31(12), 40–47.
- Tambe, M. (1997). Agent architectures for flexible, practical teamwork. In *Fourteenth national conference on artificial intelligence* (pp. 22–28).
- Tambe, M. (1997). Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7, 83–124.
- Tuomela, R. (2004a). Cooperation and the we-perspective. Paper for the conference *Rationality and Commitment*, St. Gallen, May 2004.
- Tuomela, R. (2004b). Joint action. Invited paper for the Workshop on Holistic Epistemology and Theory of Action, University of Leipzig, June 24–27.
- Tuomela, R. (1995). *The importance of us: A philosophical study of basic social notions*. Stanford, CA: Stanford University Press.