

Planning and Acting Together

Barbara J. Grosz, Luke Hunsberger, and Sarit Kraus

■ People often act together with a shared purpose; they collaborate. Collaboration enables them to work more efficiently and to complete activities they could not accomplish individually. An increasing number of computer applications also require collaboration among various systems and people. Thus, a major challenge for AI researchers is to determine how to construct computer systems that are able to act effectively as partners in collaborative activity.

Collaborative activity entails participants forming commitments to achieve the goals of the group activity and requires group decision making and group planning procedures. In addition, agents must be committed to supporting the activities of their fellow participants in support of the group activity. Furthermore, when conflicts arise (for example, from resource bounds), participants must weigh their commitments to various group activities against those for individual activities.

This article briefly reviews the major features of one model of collaborative planning called SHAREDPLANS (Grosz and Kraus 1999, 1996). It describes several current efforts to develop collaborative planning agents and systems for human-computer communication based on this model. Finally, it discusses empirical research aimed at determining effective commitment strategies in the SHAREDPLANS context.

Pollack (1998, 1992) has argued that “there’s more to life than making plans.” First, planning agents inhabit worlds that are constantly changing, worlds in which actions have uncertain outcomes; they must be able to adapt to these changes and deal with these uncertainties. Second, planning most often must be done incrementally and must be interleaved with execution. Third, in addition to planning sequences of actions, agents must manage their commitments and monitor their environments effectively, assess alternative plans and possibilities, and coordinate with other agents (Pollack 1998; see the article by Pollack and Horty, also in this issue).

There’s not only more to life than making

plans, there’s more to making plans than making them alone, and there’s more to managing plans than managing them alone. Multiagent planning involves more than the coordination of individual plans. It requires group decision making (for example, to select a high-level decomposition for the group action or to assign agents and subgroups to constituent tasks). Similarly, multiagent plan management requires group processes for plan evaluation, plan backtracking, and commitment monitoring. Although these group processes have individual-plan analogs, their operation is necessarily complicated by the multiagent context. Furthermore, the processes agents use to manage their individual plans must take into account their commitments to any group activities.

This article describes research focused on a particular kind of distributed, continual plan management, namely, plan management in the context of multiagent collaboration. Collaboration is distinguished from other kinds of multiagent interaction (for example, cooperation, coordination, and competition) by the existence of a shared goal and the agents’ commitments to this shared goal. Our work attempts to identify the capabilities needed by agents if they are to plan and act collaboratively. In this work, the groups that collaborate may persist over long periods of time (as do orchestras, sports teams, and systems administration groups), form spontaneously for a single group activity (as when a group forms for a programming project), or come together repeatedly (as do surgical teams and airline crews).

We have developed a formalization of collaboration, called SHAREDPLANS (Grosz and Kraus 1999, 1996), which provides both a specification for the design of collaboration-capable agents and a framework for identifying and investigating fundamental questions about collaboration. A recent paper (Hadad and

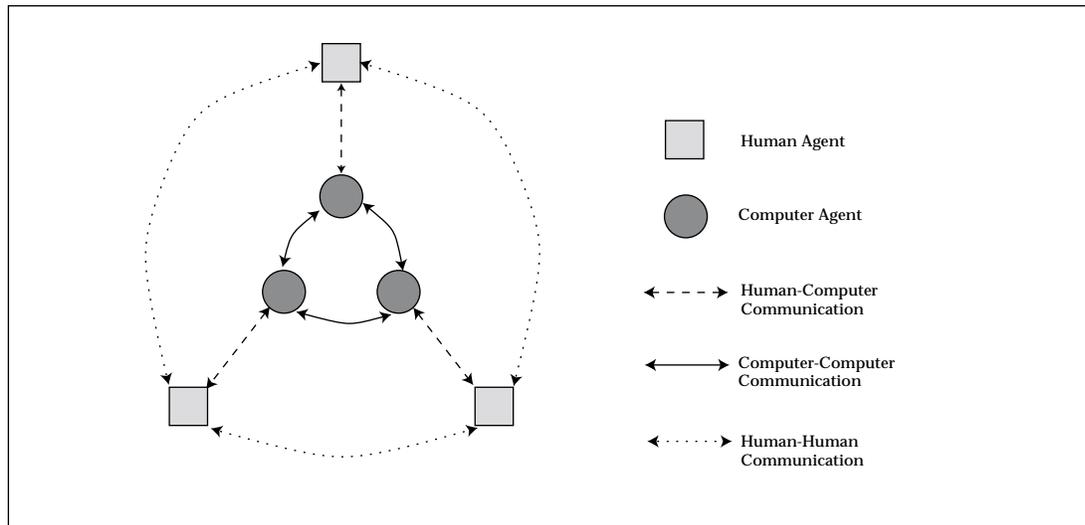


Figure 1. GIGAGENTS for the Musical Gig Domain.

Kraus 1999) presents several examples illustrating that the use of SHAREDPLANS leads to better use of resources, better coordination of tasks, and higher chances of fulfilling agents' goals. The particular properties of SHAREDPLANS that contribute to these improvements include provision for agents to interleave planning and acting, inclusion of commitments that can lead agents to behave helpfully, and constraints that prohibit the adoption of conflicting intentions.

We are currently developing three systems based on the SHAREDPLANS theory: (1) GIGAGENTS, a multiagent system for collaborations of heterogeneous groups of people and computer agents; (2) WEBTRADER, a multiagent system for electronic commerce; and (3) DIAL (which stands for "distributed information access for learning"), a system that provides a collaborative interface for distance learning.

The GIGAGENT system has an architecture that supports participation in several different kinds of collaborative activity involving teams of both people and computer systems. The name GIGAGENT was motivated by a sample application involving a group of human musicians who form SHAREDPLANS to perform musical gigs. Figure 1 provides a schematic of a GIGAGENT system that might support this kind of collaboration, a collaboration in which the main activity (playing the gig) is an action to be done by a group of people, but computer systems are involved in various supporting activities. In such settings, each person involved in the collaboration has a dedicated computer assistant. Should information about a potential gig arise (for example, over the internet or by e-mail), the group of computer

agents could handle much of the work of determining whether the gig might be feasible and, if so, establishing the framework of a SHAREDPLAN for doing the gig. For example, computer GIGAGENTS might handle such tasks as making arrangements for the group to rent a van or sending e-mail confirming the gig contract.

Another type of GIGAGENT system is one in which a combined team of people and computer systems performs the overall collaborative activity. We are currently implementing the GIGAGENT system in a computer systems administration domain that requires this kind of organization (figure 2). In this domain, people perform such actions as installing new software, and systems perform such actions as running backup programs.

The WEBTRADER system (Hadad and Kraus 1999) was designed to operate in the electronic commerce domain. Figure 3 gives a schematic view of a typical WEBTRADER configuration. In this domain, there are several enterprises, each employing several buyer and seller agents. A collaborative transaction occurs when a buyer agent and a seller agent, or a human buyer and a seller agent, form a SHAREDPLAN for the buyer to purchase some item from the seller.

In both the GIGAGENT and WEBTRADER systems, we use SHAREDPLANS not only explicitly to model the collaborations among the various agents but also implicitly in the design of the requisite human-computer interfaces. The goal is to make the interfaces collaborative and, hence, more effective. In this effort, we are drawing on our experience in designing DIAL, a system that provides a collaborative interface for distance learning.

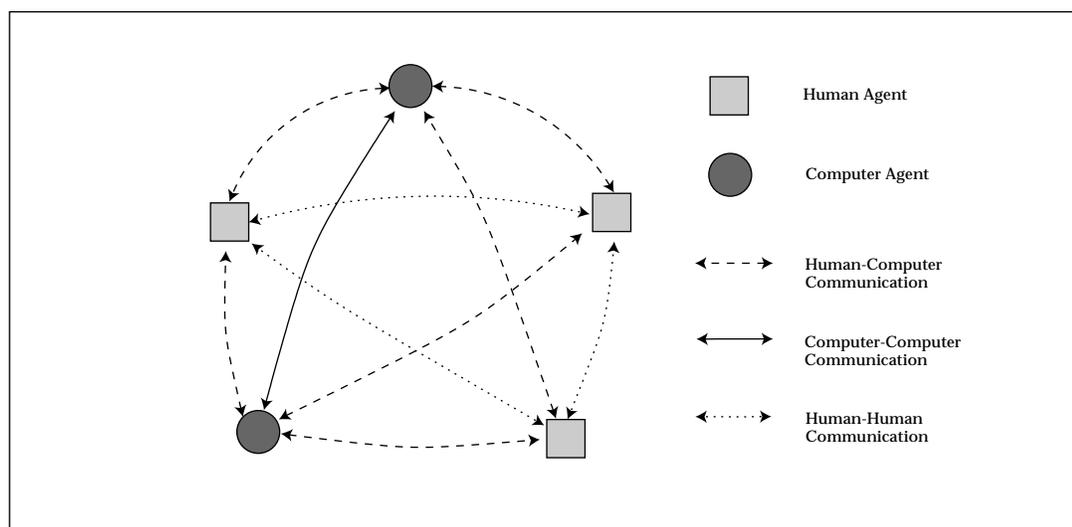


Figure 2. GIGAGENTS for the Systems Administration Domain.

The next sections of this article present brief overviews of the SHAREDPLANS formalization and the basic processes required in collaborative systems based on it. Subsequent sections provide additional details of the GIGAGENT, WEB-TRADER, and DIAL systems. The final section describes empirical research on the problem of intention-reconciliation in a group context; this work was prompted by both the formalization and the implementation efforts.

The SHAREDPLANS Formalization of Collaboration

Pollack (1990) argued that agents typically have different beliefs about the ways to do an action or to achieve some desired state and that, as a result, it was necessary to move away from a view of plans as primarily data-structure-like to a view of plans in terms of mental state. In this view, an agent has a plan to do some action if it has a certain set of beliefs and intentions pertaining to that action. Bratman (1987) also argued for a mental-state view of plans, emphasizing the roles of intentions in focusing means-ends reasoning, in constraining the other intentions an agent might subsequently adopt, and in guiding replanning.

The SHAREDPLANS formalization of collaborative planning is based on a mental-state view of plans. It was originally conceived to provide the basis for modeling the intentional structure of discourse (Grosz and Sidner 1990). We subsequently generalized the formalization to accommodate more than two participating agents and to support the construction of teams of collaboration-capable agents (Grosz

and Kraus 1999, 1996; Hunsberger 1999). Lochbaum (1998, 1995) demonstrated the power of using SHAREDPLANS as the basis for intentional structure in dialogue. Others adapted this use of SHAREDPLANS to the design of graphically based human-computer interfaces (Rich and Sidner 1998). Tambe (1997) incorporated some elements of SHAREDPLANS in STEAM (which stands for “shell for teamwork”), which he used to build systems for RoboCup soccer and military simulations. A comparison of the SHAREDPLANS formalization with alternative computational accounts (Kinny et al. 1994; Cohen and Levesque 1991; Cohen, Levesque, and Nunes 1990) and philosophical theories (Bratman 1992; Searle 1990) can be found in earlier papers (Grosz and Kraus 1999, 1996).

The SHAREDPLANS formalization treats complex actions and partial plans, models the possibility of agents contracting out actions to other agents, provides for the interleaving of planning and execution, and distinguishes the information needed by agents directly involved in the planning of some constituent action from the information the group at large must have about that action. The model presumes general procedures for reconciling conflicts among intentions and requires a variety of group decision-making processes.

The formalization distinguishes two kinds of intention: (1) *intention to* do an action and (2) *intention that* a proposition hold. An agent intending to do some action must not only be committed to doing the action, it must also have appropriate beliefs about its ability to do the action, and it must have some knowledge about how to do it or about how to figure out how to do it. An agent intending that some

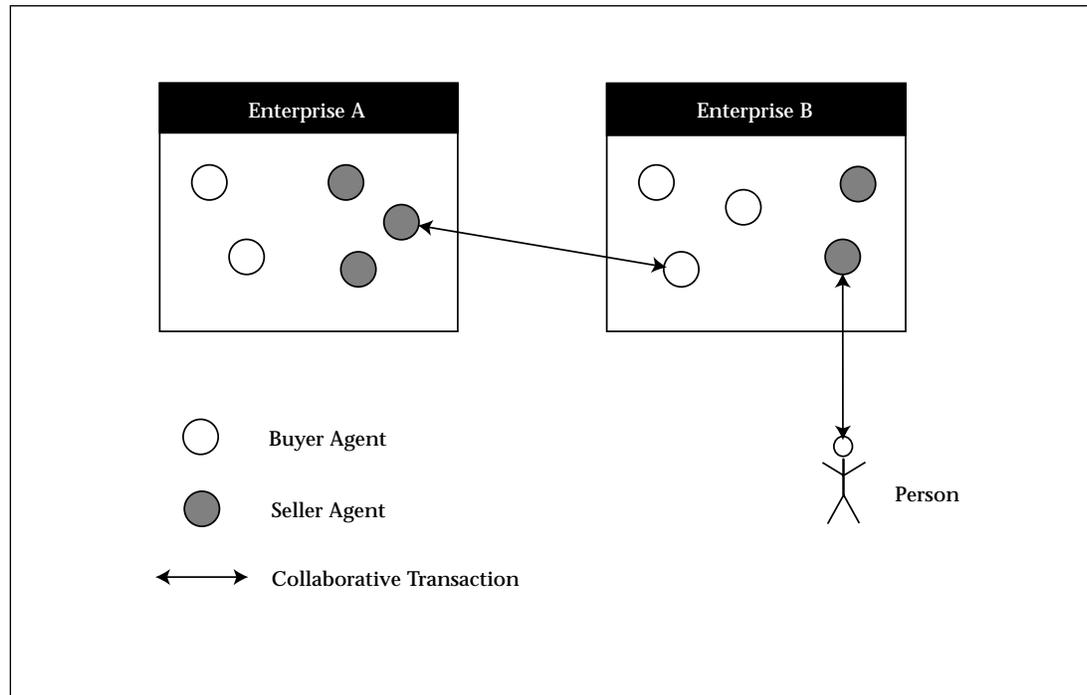


Figure 3. WEBTRADER: SHAREDPLANS in an Electronic Commerce Domain.

proposition hold must be committed to doing what it can, if anything, to bring about the proposition, but it need not necessarily be able to do anything. Both forms of intention preclude agents (knowingly) adopting conflicting intentions. In SHAREDPLANS, intentions—that are used to represent agents' commitments to their group activity, to specify the collaborative support participants offer one another, and to provide a basis for agents to form the mutual beliefs necessary for collaboration (Grosz and Kraus 1999).

The mental-state view of plans is crucial to the SHAREDPLANS formalization. It enables the coordination and integration of SHAREDPLANS for group actions with individual plans for single-agent constituent actions of these plans because the same kinds of belief and intention play a role in each. Furthermore, in a multi-agent context, plans must be ascribable to agents and groups of agents—either theoretically, from an omniscient-observer perspective, or from the perspective of a fellow collaborator. Because ascribing a plan to an agent or group of agents depends on a determination that the relevant beliefs and intentions hold, specifying the conditions under which a plan can be so ascribed requires a mental-state view of plans. In addition, adopting the mental-state view has enabled the proof of several theorems relating agent knowledge and beliefs to the existence of plans (Hunsberger 1999).

Plan Definitions

The SHAREDPLANS formalization distinguishes partial plans and complete plans. A *full SHAREDPLAN* (FSP) is a complete plan in which agents have fully determined how they will perform an action. Because most of the time that agents are involved in a collaborative activity their plans are partial, the FSP definition serves as a target for their planning activity—agents know they must reach the mental states specified by the FSP definition to have a complete plan. The *partial SHAREDPLAN* (PSP) definition provides a specification of the minimal mental-state requirements for collaboration to exist and gives criteria governing the process of completing the plan.

Agents involved in a SHAREDPLAN to do some action are required to have certain individual and mutual beliefs about how the action and its constituent subactions are to be done. These beliefs are specified in terms of recipes. A (*full*) *recipe* for doing some action A_α is a set of actions and constraints such that the doing of these actions under these constraints constitutes the doing of A_α . A *partial recipe* is a set of actions and constraints that can be extended to a full recipe.

Ultimately, actions must be performed by individual agents. Thus, SHAREDPLANS contain individual plans of the participating agents as constituents. Extending the work of Pollack (1990), the SHAREDPLANS formalization specifies

For a group GR to have a Full SHAREDPLAN (FSP) to do a complex action A_α requires that:

- GR mutually believe that each member agent intends that the group do A_α .
- GR mutually believe that they have a full recipe for doing A_α .
- Each action in that recipe be fully resolved.

A basic action A_β is *fully resolved* if:

- Some agent G_β in GR intends to do A_β .
- GR mutually believe that G_β intends to do A_β and is able to do A_β .
- GR mutually believe that each member agent intends that G_β be able to execute A_β .

Similarly, a complex action A_κ is *fully resolved* if:

- Some subgroup GR_κ in GR have a full plan to do A_κ .
- GR mutually believe that GR_κ has a full plan to do A_κ and is able to do A_κ .
- GR mutually believe that each member agent intends that GR_κ be able to execute A_κ .

Figure 4. Informal Definition of a Full SHAREDPLAN (FSP).

that for an agent to have a *full individual plan* to do some (single-agent) action, it must have knowledge of a recipe for the action, believe that it will be able to perform the actions in the recipe, and have intentions to do the constituent actions. Because individual plans, like SHAREDPLANS, are partial much of the time, the formalization also defines a *partial individual plan* in which an agent may have merely a partial recipe for doing the action and may have merely partial plans for doing the actions in the recipe.

For both individual and group plans, *elaborating* a partial plan into a full plan involves choosing a full recipe or extending (perhaps incrementally) a partial recipe. It also involves establishing subordinate plans for each subaction in the recipe. In a full plan, the group action is fully decomposed into *basic* actions (that is, actions that an individual agent can execute at will under appropriate conditions).

The FSP and PSP definitions are presented informally in figures 4 and 5.¹ The numbers in the figures indicate some of the essential elements of a SHAREDPLAN: (1) commitment to the group activity; (2) commitment to the actions of others in service of the group activity; (3) the existence of, and commitment to, a group

process for selecting a recipe for the group activity; and (4) the existence of, and commitment to, a group process for assigning agents or subgroups to constituent actions. As the definitions make evident, planning collaboratively introduces several complexities not present in the single-agent case. First, collaborating agents typically have access to different information and thus hold different sets of beliefs about actions. For example, only the agents directly involved in the planning and execution of some subaction typically need to know the details of how the subaction will be done. Second, plan elaboration in a multiagent context requires the assignment of agents or subgroups to the constituent individual and multiagent actions (respectively) in the recipe. Third, agents must form commitments not only to doing their own actions but also to supporting the actions of others and supporting the group activity as a whole. Finally, the decision making required for plan elaboration is a group process and thus requires negotiation and group decision-making protocols.

Plan-Management Processes

A group of collaborating agents utilize a vari-

For a group GR to have a Partial SHAREDPLAN (PSP) to do a complex action A_α requires that:

- GR mutually believe that each member agent intends that the group do A_α . (1)
- Either GR mutually believe that they have a full recipe for doing A_α , or they mutually believe that they have a partial recipe that can be extended into a full recipe they can use to do A_α and they have a full plan to select such a recipe, and
- Each action in the (possibly partial) recipe be either *at-least-partially-resolved* or *unresolved*.

A basic action A_β is *at least partially resolved* if:

- Some agent G_β in GR intends to do A_β .
- GR mutually believe that G_β intends to do A_β .
- GR mutually believe that each member agent intends that G_β be able to execute A_β , and each agent in GR intends that GR mutually believe that G_β is able to execute A_β . (2)

Similarly, a complex action A_κ is *at least partially resolved* if:

- Some subgroup GR_κ in GR have a plan to do A_κ .
- GR mutually believe that GR_κ have a plan to do A_κ .
- GR mutually believe that each member agent intends that GR_κ be able to execute A_κ and each agent in GR intends that GR mutually believe that GR_κ are able to execute A_κ . (2)

A basic action A_ε is *unresolved* if:

- GR mutually believe that some member of GR could do A_ε .
- GR have a full plan to select such an agent. (3)

Similarly, a complex action A_μ is *unresolved* if:

- GR mutually believe that some subgroup of GR could do A_μ .
- GR have a full plan to select such a subgroup. (4)

Figure 5. Informal Definition of a Partial SHAREDPLAN (PSP).

ety of group processes to manage their multi-agent plan. As a group, they must elaborate their partial plan into a full plan, which in turn requires that they complete the recipe if it is partial, select agents or subgroups to do constituent actions, and establish a variety of individual commitments and mutual beliefs. Although the general structure of these group processes is similar to their individual-plan counterparts, they differ in their fundamental reliance on group decision making. Taken together, these planning processes are the means by which partial plans evolve into full plans; as such, they constitute necessary elements of the plan-management repertoire of any group of collaboration-capable agents.

The SHAREDPLANS definitions constrain the group decision-making processes in many ways. For example, a group decision to extend

a partial recipe is constrained by the requirement that the new subactions in the recipe extension meet, at a minimum, the specifications for unresolved subactions in the PSP definition (figure 5). Similarly, a group decision to assign an agent or subgroup to some constituent action is constrained by the specifications for at-least-partially-resolved subactions in the PSP definition. Although the SHAREDPLANS definitions constrain these group planning processes, they do not fully determine them. We are using the GIGAGENT system to explore candidate algorithms for these group processes, as discussed in the next section.

In addition to group plan-management processes, the SHAREDPLANS formalization requires each agent to have various individual plan-management processes. Some of these individual plan-management processes are

associated with intentions—that arising from group activities. The role of these *intention-cultivation* processes is to determine what is necessary to monitor the intended proposition; determine possible courses of action to make the proposition hold; and decide which courses of action, if any, should eventually be carried out (Grosz and Kraus 1999). Other individual plan-management processes are associated with intention-conflict resolution. We are examining the effects of various intention conflict-resolution strategies in our empirical research, as discussed later in this article. In addition, a major emphasis of our work in constructing collaboration-capable systems is the investigation of candidate algorithms for intention-cultivation processes.

Manifesting Collaboration

The SHAREDPLANS formalization places numerous constraints on the design of a collaborative agent. First, an agent intending to do some action must commit to means-ends reasoning (as represented by a plan to elaborate a partial plan). Second, for a group to have a SHAREDPLAN, they must have agreed about certain decision-making procedures (for example, to select recipes, to assign agents and subgroups to subactions). Third, because a rational agent may not simultaneously hold conflicting intentions, a collaboration-capable agent must have some means of detecting conflicting intentions and reconciling conflicts when they occur. Fourth, as noted previously, an agent participating in collaborative activity must have processes with which to manage its various intentions—that—from intentions that certain mutual beliefs are established to intentions that various subgroups be able to carry out their assigned tasks.

In this section, we describe three systems currently under development: (1) WEBTRADER, a system for electronic commerce (Hadad and Kraus 1999); (2) GIGAGENTS, a collaboration-capable multiagent system; and (3) DIAL, a system for collaborative distance learning. In the first two systems, elements of the SHAREDPLANS specifications are directly reflected in the system architectures. In the third system, the SHAREDPLANS formalization has been used to inform the system design, but the architecture does not directly reflect the SHAREDPLANS specifications.

WEBTRADER: Electronic Commerce Application

The goal of the WEBTRADER application is to incorporate collaborative features into a com-

puterized system for buying and selling items such as books, clothes, and furniture on the internet. In the WEBTRADER environment, several enterprises, each with several kinds of goods to sell to each other or to individual (human) buyers, have intelligent buyer and seller agents. The job of *buyer agents* is to purchase goods that are missing from the stocks of their enterprises. The job of *seller agents* is to sell their enterprises' goods.

In the WEBTRADER environment, agents' benefits are measured in terms of income. SHAREDPLANS can be formed between agents belonging to the same enterprise who want to maximize their enterprise's income. They can also be formed among agents interested in collaboration only as a means of maximizing their individual benefit. For example, suppose a buyer agent wants to purchase an item from a seller agent of another enterprise. Although each agent wants to maximize its own enterprise's benefit, and the agents might have certain conflicting interests, they share the goal of successfully completing the transaction. By forming a SHAREDPLAN, they increase their chances of completing a mutually beneficial transaction. For example, having a SHAREDPLAN might motivate the seller to notify the buyer should delivery problems arise.

The use of SHAREDPLANS in the system design can also prove mutually beneficial when an automated seller interacts with a human buyer. For example, tracking the intentional context of an interaction (as we also do in the DIAL system, described later) can make it easier for the seller to help the buyer identify an item to purchase. Thus, if a buyer has selected an out-of-stock CD and the intentional context is that the buyer wanted this CD because it contained a particular song, then the seller might suggest another CD containing a different version of the same song. In contrast, if the intentional context is that the buyer wanted the CD because it contained songs recorded by a particular singer, then the seller might suggest some other CD recorded by the same singer.

In the electronic commerce domain, agents begin with partial plans and develop them over time into complete plans. For example, at the beginning of an interaction, a buyer might only be able to characterize his needs approximately; later on, with the help of a seller agent—who might, for example, present information about available items and their properties—they might be able to find a specific item to satisfy the buyer's needs. The electronic commerce domain is also dynamic: Things change, possibly in the middle of the planning process, while the plan is still partial. In addi-

The goal of the WEBTRADER application is to incorporate collaborative features into a computerized system for buying and selling items such as books, clothes, and furniture on the internet.

*The
GIGAGENT
system is a
multiagent
system that
supports the
collaboration
of any
number of
people and
computer
agents, each
of whom can
participate
simultane-
ously in
activities
involving
several
different
groups.*

tion, an agent's beliefs about the results of its own actions might be uncertain or faulty, and its knowledge about the world or other agents might be incomplete. For example, a buyer and a seller might have agreed on the date of payment for a particular item, but as the date approaches, the buyer might realize that he is unable to pay for the item as originally planned. Thus, while an agent is planning or acting on the basis of a partial plan, the plan might have to be revised. By providing for both individual and collaborative plans to be partial and specifying the processes by which they can be elaborated, the SHAREDPLANS formalization enables us to develop agents that are able to act in the dynamic and uncertain electronic commerce environment.

In contrast to other agent-based markets or retail outlets on the web (Walsh, Welman, and Wurman 1998; Chavez et al. 1997; Doorenbos, Etzioni, and Weld 1997; Albayrak et al. 1996; Chavez and Maes 1996; Klaus et al. 1996; Schrooten 1996; Takahashi et al. 1996), which focus on either extraction of information or negotiation strategies, our system supports cooperative interactions of buying and selling goods. The current prototype of the WEBTRADER agent is able to construct and carry out simple plans for buying and selling items in a simulated environment (which we implemented) but is not yet able to buy and sell on the web. It is able to cooperate simultaneously with more than one agent. To enable such cooperation, we developed and implemented data structures and control mechanisms that enable the agent to handle several complex actions at a time. For example, a WEBTRADER agent is able to interrupt negotiations with one agent, start new negotiations with another agent, and then interrupt these negotiations to return to the first agent.

GIGAGENTS: Domain-Independent Collaborative Multiagent System

The GIGAGENT system is a multiagent system that supports the collaboration of any number of people and computer agents, each of whom can participate simultaneously in activities involving several different groups. Even within a single group activity, hierarchical task expansion can lead to a single agent being involved simultaneously in several active, subordinate SHAREDPLANS. The participation of people in the GIGAGENT system is facilitated by agent wrapper software that handles most of the routine chores required of collaborating agents (for example, sending and receiving messages and keeping track of commitments, proposals, pending decisions, and their interrelationships).

The purpose of the GIGAGENT system is to provide a platform for studying multiagent collaboration. We are using GIGAGENTS to test and evaluate various strategies and algorithms for team formation, commitment generation, plan elaboration, and conflict resolution (especially conflicts arising from the interaction of commitments to individual and group activities). In the current implementation, we are simulating a computer systems administration domain in which heterogeneous groups of computer and human agents work together on tasks such as upgrading software and hardware, restoring deleted files from backups, keeping web pages up to date, and checking system security.

Group decision making is a significant component of GIGAGENT activity. Agents must agree on how they will select recipes, bind various parameters of their group actions, and assign agents or subgroups to various constituent tasks. Furthermore, this group decision-making activity typically happens incrementally. The GIGAGENT architecture supports the incremental execution of planning actions (and domain actions, if desired), thereby enabling each agent to participate in several ongoing collaborative negotiations pertaining to different decisions required by various group activities. The system neither imposes nor presumes any particular organizational or management structure among collaborating agents. A range of possibilities—from completely democratic protocols (one agent, one vote) to autocratic ones—can be incorporated and tested. An initial voting-based protocol has been implemented to enable us to test the overall system architecture and infrastructure.

We plan to use the GIGAGENT system to test and evaluate strategies for conflict resolution coming out of our empirically based SPIRE research project (discussed later). We are also beginning to investigate algorithms for the single-agent cultivation processes associated with intentions-that held by an agent. These cultivation processes, implemented as incrementally executable actions, are responsible for monitoring the status of their associated intentions-that, considering ways of bringing about the intended propositions, and determining any potential conflicts. Although we expect the cultivation processes to rely in part on domain-dependent reasoning, our current research aims to uncover the domain-independent portions of such processes.

Finally, in addition to the explicit collaboration among agents in the GIGAGENT system, there is implicitly modeled collaboration in the implementation of the agent wrapper

software that facilitates the participation of human agents in the system. To make this interface more effective, we are investigating how to increase its collaborative nature by basing its operation on general principles gleaned from our work on the SHAREDPLANS formalization and the human-computer interface systems we have built based on SHAREDPLANS. The next section discusses this use of SHAREDPLANS in more detail.

SHAREDPLANS for Interface Design

We have used SHAREDPLANS in the design of a collaborative human-computer interface system. DIAL provides for multimedia interactions with a complex information system. Our goal in building DIAL was to demonstrate the efficacy of the SHAREDPLANS model, and collaboration more generally, in providing people with natural, flexible, and effective means of communicating with computer systems. A central aim of the work has been to develop a system that enables users to obtain the information they need without having to specify the way in which the system should find it. DIAL is active in working with users to identify information relevant to their needs or tasks rather than providing the narrow input-output window of current-generation interfaces.

The current DIAL implementation operates in the domain of information support for distance learning by students in an introductory programming class. It allows students to access course notes, videos, and assignments, as well as reference materials and teaching fellows, when undertaking such tasks as studying for exams, reviewing lecture materials, and working on assignments. In this domain, the collaborative plans formed by DIAL and its users are *information-locating plans*, that is, plans to locate information relevant to some task of the user (for example, studying for an exam, preparing a lecture). These collaborative plans are in service of the user's individual intentions and plan to carry out the task. Typically, the initial SHAREDPLAN will spawn many subsidiary SHAREDPLANS (Lochbaum 1998, 1995), which will also be information-locating plans. For example, a user's individual intention to study for a midterm exam might lead her to engage DIAL to form a SHAREDPLAN with the system to find information relevant to the exam. In such a case, DIAL would present a menu indicating the types of information available (for example, assignments and lecture notes). At this point, were the user to select "review an assignment," she and DIAL would then form a subsidiary SHAREDPLAN to locate assignments and assignment information relevant to the exam.

DIAL builds a representation of context based on the subsidiary relationships between the various SHAREDPLANS it forms with the user. It utilizes this contextual information to reduce the user's communication burden. In particular, information requests are interpreted relative to the graphically displayed (and, hence, mutually believed) prevailing intentional context, which may be manipulated by the user as well as the system. As a result, requests may be expressed more succinctly. For example, if a user asks to see videos in the context of reviewing some topic (say, sorting algorithms), DIAL would interpret the request as a request to see videos on that topic. DIAL has been designed to use information about the intentional context and also to respond and act collaboratively rather than in the master-slave style typical of most current human-computer interfaces. For example, it presents choices that the user might be unaware of as requests are pursued. The system has also been designed to use intentional context to plan alternative courses of action, but such plan-revision capabilities are not yet integrated into the implementation.

DIAL does not explicitly reason about beliefs and intentions but does incorporate the specifications of SHAREDPLANS listed in figure 6. To simplify the construction and use of DIAL, we stipulated which actions the user and the system could perform and fixed the subtask assignments rather than leaving them to a dynamic decision-making process.

Empirical Studies of Decision Making in a Collaborative Context

In addition to providing a foundation for constructing collaborative systems, the SHAREDPLANS formalization establishes a framework for identifying and investigating fundamental questions about collaboration. In this section, we describe an empirical study of one such question.

Agents involved in a collaborative activity may need to reconcile their intentions to do group-related actions with other, conflicting intentions. For example, an agent committed to perform a certain task as part of a group activity might be given the opportunity to do another, unrelated activity that occurs at the same time as the group-related task. Because agents are presumed to act in an individually rational manner, seeking to maximize their (or their owner's) utility rather than the utility of the group, they need to be able to weigh the costs and benefits of the two options in the context

The SHARED-PLANS formalization establishes a framework for identifying and investigating fundamental questions about collaboration.

- *Intention-that the group perform the group action:* DIAL's commitment to the collaborative information-locating activities it undertakes with a user are manifest in its tracking of the intentional context and its use of the context to guide the interpretation of user queries and selections and to constrain its replies.
- *Constructing and extending (partial) recipes:* The user and DIAL incrementally compose recipes for their group activities throughout any given session.
- *Mutual belief of subtask assignments and action capabilities:* The visual display of intentional structure provides common ground for establishing these mutual beliefs.
- *Intentions-that other participants be able to do subactions in the recipe to which they are committed:* DIAL constrains user choices to those actions that can be performed in the current context.

Figure 6. Specifications of SHAREDPLANS Incorporated into the DIAL System.

of their commitment to the group. Intention reconciliation in such contexts is complex because agents need to consider how other group members will react if they fail to honor their commitments. Because agents might want to collaborate in the future with other members of the group, there can be significant costs to defaulting on group-related tasks.

To study the problem of intention reconciliation in the context of collaboration, we have constructed the SPIRE (SHAREDPLANS intention-reconciliation experiments) simulation system (Sullivan et al. 1999). SPIRE is general enough to allow us to model agents from a large set of problem domains; it enables us to consider the impact of group norms and environmental factors on agents faced with conflicting intentions as well as the effectiveness of different intention-reconciliation strategies that agents might adopt. The group norms include penalties imposed when agents default on group-related tasks. The environmental factors include the size of the group and the number of tasks in each time period. We aim to provide insight into the types of factors that affect individual and group behavior and outcomes and thus derive principles for designers of collaboration-capable agents.

Initial pilot studies in the SPIRE framework (Sullivan et al. 1999) have established that the number of times agents default on group tasks drops off as agents give more weight to group factors in their decision making. The number

of defaults decreases as the penalty for defaulting increases, and individual income and group income increase as the number of defaults decreases. Although these results might be unsurprising, they indicate that the group factors and norms we have designed are reasonable and provide a concrete demonstration of how group-related factors can affect the decision making of self-interested agents. A study that examined various aspects of *social consciousness* (agents being willing to sacrifice short-term personal gain for longer-term group good) has shown that as agents become more socially aware, defaults decrease and more group-related tasks are completed, but a moderate degree of social consciousness functions better than an extreme degree when the combination of group-task and individual-task income is considered (Glass and Grosz 1999).

Conclusions

The SHAREDPLANS formalization, which we overviewed in this article, specifies the minimal mental-state requirements for a group of agents to have a plan for collaborative activity, constraints on agents' beliefs and intentions as they initiate and expand partial plans, and stopping conditions for planning processes. It provides a means of representing the commitments of participants in collaborations to their group activities and treats the partiality that arises naturally in most planning situations.

We have designed an architecture for collaborative agents based on this formalization and built two systems based on this architecture. One, called GIGAGENTS, models and supports explicit collaboration in planning and acting among both human and computer agents. The second, a system for electronic commerce, called WEBTRADER, supports the cooperative processes of buying and selling goods on the web. The formalization has also proved useful in the development of a collaborative interface to a system for distance learning, called DIAL.

The development of the SHAREDPLANS formalization and systems based on it have uncovered several fundamental questions in multiagent collaborative planning. In this article, we briefly described one such problem area: the problem of reconciling intentions in the context of group activities. We developed SPIRE, an empirical framework that enables different decision-making processes to be simulated and studied. The article presented the results of early studies using this framework.

Our current efforts focus on continued empirical studies of intention-reconciliation, incorporation of the results of these studies into the collaboration-capable systems we are developing, and development of more sophisticated group decision-making algorithms for use in these systems.

Acknowledgments

This research has been supported in part by National Science Foundation grants IRI 95-25915, IRI-9618848, and CDA-94-01024. The authors thank David Sullivan for his helpful comments on earlier drafts and for his assistance with the section on empirical studies.

Note

1. To simplify the presentation, single-agent groups are permitted. Fomal definitions can be found elsewhere (Hunsberger 1999).

References

- Albayrak, S.; Meyer, U.; Bamberg, B.; Fricke, S.; and Tobben, H. 1996. Intelligent Agents for the Realization of Electronic Market Services. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multiagent Technology*, 11–23. London: Practical Application.
- Bratman, M. E. 1992. Shared Cooperative Activity. *The Philosophical Review* 101(2): 327–341.
- Bratman, M. E. 1987. *Intention, Plans, and Practical Reason*. Cambridge, Mass.: Harvard University Press.
- Chavez, A., and Maes, P. 1996. KASBAH: An Agent Marketplace for Buying and Selling Goods. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multiagent Technology*, 75–90. London: Practical Application.
- Chavez, A.; Dreilinger, D.; Guttman, R.; and Maes, P. 1997. A Real-Life Experiment in Creating an Agent Marketplace. In *Software Agents and Soft Computing: Toward Enhancing Machine Intelligence: Concepts and Applications*, eds. H. S. Nwana and N. Azarmi, 160–178. Lecture Notes in Computer Science, Volume 1198. New York: Springer.
- Cohen, P. R., and Levesque, H. J. 1991. Teamwork. *Noûs* 25(4): 487–512.
- Cohen, P. R.; Levesque, H. J.; and Nunes, J. H. T. 1990. On Acting Together. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, 94–99. Menlo Park, Calif.: American Association for Artificial Intelligence.
- Doorenbos, R. B.; Etzioni, O.; and Weld, D. S. 1997. A Scalable Comparison-Shopping Agent for the World-Wide Web. In *Proceedings of the First International Conference on Autonomous Agents*, 39–48. New York: Association of Computing Machinery.
- Glass, A., and Grosz, B. J. 1999. Socially Conscious Decision Making. Paper presented at the Sixth Annual Bar Ilan Symposium on the Foundations of Artificial Intelligence (BISFAI-99), 23–25 June, Ramat Gan, Israel.
- Grosz, B. J., and Kraus, S. 1999. The Evolution of SHAREDPLANS. In *Foundations of Rational Agency*, eds. A. Rao and M. Wooldridge, 227–262. Dordrecht, The Netherlands: Kluwer Academic.
- Grosz, B. J., and Kraus, S. 1996. Collaborative Plans for Complex Group Action. *Artificial Intelligence* 86(2): 269–357.
- Grosz, B. J., and Sidner, C. L. 1990. Plans for Discourse. In *Intentions in Communication*, eds. P. Cohen, J. Morgan, and M. Pollack, 417–444. Cambridge, Mass.: Bradford Books.
- Hadad, M., and Kraus, S. 1999. SHAREDPLANS in Electronic Commerce. In *Intelligent Information Agents*, ed. M. Klusch, 204–231. Heidelberg, Germany: Springer-Verlag.
- Hunsberger, L. 1999. Making SHAREDPLANS More Concise and Easier to Reason About. In *Intelligent Agents V—Agent Theories, Architectures, and Languages (ATAL-98)*, eds. J. P. Muller, M. P. Singh, and A. S. Rao, 81–98. Lecture Notes in Artificial Intelligence 1555. Heidelberg, Germany: Springer-Verlag.
- Kinny, D.; Ljungberg, M.; Rao, A. S.; Sonenberg, E.; Tidhar, G.; and Werner, E. 1994. Planned Team Activity. In *Artificial Social Systems*, eds. C. Castelfranchi and E. Werner, 227–256. Lecture Notes in Artificial Intelligence 830. Amsterdam, The Netherlands: Springer-Verlag.
- Klaus, F.; Muller, J.; Heimig, I.; and Scheer, A.-W. 1996. Intelligent Agents in Virtual Enterprises. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multiagent Technology*, 250–223. London: Practical Application.
- Lochbaum, K. E. 1998. A Collaborative Planning Model of Intentional Structure. *Computational Linguistics* 24(4): 525–572.
- Lochbaum, K. E. 1995. The Use of Knowledge Pre-

conditions in Language Processing. In Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95), 1260-1266. Menlo Park, Calif.: International Joint Conferences on Artificial Intelligence.

Pollack, M. E. 1990. Plans as Complex Mental Attitudes. In *Intentions in Communication*, eds. P. Cohen, J. Morgan, and M. Pollack, 77-103. Cambridge, Mass.: Bradford Books.

Pollack, M. E. 1998. Plan Generation, Plan Management, and the Design of Computational Agents (Abstract). In Proceedings of the Third International Conference on Multiagent Systems (ICMAS-98), 8-9. Washington, D.C.: IEEE Computer Society.

Pollack, M. E. 1992. The Uses of Plans. *Artificial Intelligence* 57(1): 43-68.

Rich, C., and Sidner, C. L. 1998. COLLAGEN: A Collaboration Manager for Software Interface Agents. Technical Report 97-21a, Mitsubishi Electric Research Laboratories, Cambridge, Massachusetts.

Schrooten, R. 1996. Agent-Based Electronic Consumer Catalog. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multiagent Technology*, 543-556. London: Practical Application.

Searle, J. R. 1990. Collective Intentions and Actions. In *Intentions in Communication*, eds. P. Cohen, J. Morgan, and M. Pollack, 401-415. Cambridge, Mass.: Bradford Books.

Sullivan, D.; Glass, A.; Grosz, B. J.; and Kraus, S. 1999. Intention Reconciliation in the Context of Teamwork: An Initial Empirical Investigation. In *Cooperative Information Agents III*, eds. M. Klusch, O. Shehory, and G. Weiss, 149-162. Lecture Notes in Artificial Intelligence 1652. Berlin: Springer-Verlag.

Takahashi, K.; Nishibe, Y.; Morihara, I.; and Hattori, F. 1996. Collecting Shop and Service Information with Software Agents. In *Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multiagent Technology*, 587-595. London: Practical Application.

Tambe, M. 1997. Agent Architectures for Flexible, Practical Teamwork. In Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-97), 22-28. Menlo Park, Calif.: American Association for Artificial Intelligence.

Walsh, W. E.; Wellman, M. P.; and Wurman, P. R. 1998. The Michigan Internet AUCTIONBOT: A Configurable Auction Server for Human and Software Agents. In Proceedings of the Second International Conference on Autonomous Agents, 301-308. New York: Association for Computing Machinery.



Barbara J. Grosz is Gordon McKay Professor of Computer Science in the Division of Engineering and Applied Sciences at Harvard University. Grosz has written several seminal papers in discourse processing and

collaborative planning, developed the discourse component of a number of natural language processing systems, and is widely regarded as having established the research field of computational modeling of discourse. She has also published papers on natural language interfaces to databases. A fellow of the American Association for the Advancement of Science and the American Association for Artificial Intelligence (AAAI), she received the University of California at Berkeley Computer Science and Engineering Distinguished Alumna Award in 1997 and the AAAI Distinguished Service Award in 1999. She was president of AAAI; a member and former chair of the Board of Trustees of the International Joint Conferences on Artificial Intelligence (IJCAI), Inc.; chair of IJCAI-91; and a member of the Computer Science and Telecommunications Board of the National Research Council. Grosz's research interests include the development of models of collaborative planning for joint human-computer problem solving, design of collaborative human-computer interfaces, and determining the effects of discourse structure on intonation. Before joining the faculty at Harvard, she was director of the Natural Language Program at SRI International and cofounder of the Center for the Study of Language and Information. Grosz received an A.B. in mathematics from Cornell University and a Ph.D. in computer science from UC Berkeley.



Luke Hunsberger is a Ph.D. candidate in computer science at Harvard University. He received a B.A. and an M.A. in mathematics from the University of Oregon and subsequently taught in the Mathematics

Department at Harvard. His current research interest is multiagent collaborative planning. His e-mail address is luke@eecs.harvard.edu.



Sarit Kraus is an associate professor of mathematics and computer science at Bar-Ilan University and adjunct associate professor at the Institute for Advanced Computer Studies at the University of Maryland

at College Park. She graduated from the Hebrew University of Jerusalem with a Ph.D. in computer science in 1989 and subsequently spent two years at the University of Maryland at College Park as a postdoctoral fellow before joining the faculty at Bar-Ilan University.

Kraus has worked extensively in the following areas: automated negotiation and cooperation among automated agents, development of intelligent agents, electronic commerce, information agents, distributed AI, information broadcasting, and security of databases and agents. She has published over 95 papers in leading journals and major conferences and is coauthor of a forthcoming book on heterogeneous agent systems. She has received substantial grants from a variety of agencies and has served on numerous program committees for major conferences and workshops. In addition, she is on the editorial board of the *Annals of Mathematics and Artificial Intelligence*, the *Journal of Artificial Intelligence Research*, the *Journal of Autonomous Agents and Multiagent Systems*, and *Artificial Intelligence Communications*.

In recognition of her research accomplishments on intelligent agents, Kraus was awarded the 1995 International Joint Conferences on Artificial Intelligence Computers and Thought Award. In 1998, she won the ATAL Best Paper Award.