

A Survey of Research in Distributed, Continual Planning

Marie E. desJardins,* Edmund H. Durfee,[†] Charles L. Ortiz, Jr.,* and Michael J. Wolverton*

Abstract

Complex, real-world domains require a rethinking of traditional approaches to AI planning. Planning and executing the resulting plans in a dynamic environment requires a continual approach in which planning and execution are interleaved, there may be uncertainty in the current and projected world state, and replanning may be required when the situation changes or planned actions fail. Furthermore, complex planning and execution problems may require multiple computational agents and human planners to collaborate on a solution.

In this article, we describe a new paradigm for planning in complex, dynamic environments, which we term *distributed, continual planning* (DCP). We argue that developing DCP systems will be necessary in order for planning applications to be successful in these environments. We give a historical overview of research leading up to the current state of the art in DCP, and describe research in distributed and continual planning.

The increasing emphasis on real-world applications has led AI planning researchers to develop algorithms and systems that more closely match realistic planning environments, in which planning activity is often distributed, and plan generation may happen concurrently with plan execution. These new research directions have increased the need for distributed systems of cooperating agents for continuous planning and decision support. Multiagent planning architectures, distributed planning, mixed-initiative planning, distributed scheduling, and workflow management methods are currently active areas of research.

We argue that a new paradigm for planning is needed in order to support this style of planning, and that such a paradigm is beginning to emerge within the planning research community. We characterize this new ap-

proach to planning as *distributed, continual planning* (DCP). “Distributed planning” refers to an environment in which planning activity is distributed across multiple agents, processes, or sites. “Continual planning” refers to an ongoing, dynamic process in which planning and execution are interleaved.¹

An agent should engage in distributed planning when planning knowledge or responsibility is distributed among agents, or when the execution capabilities that must be employed to successfully achieve objectives are inherently distributed. An agent should plan continually when aspects of the world can change dynamically beyond the control of the agent, when aspects of the world are revealed incrementally, when time pressures require execution to begin before a complete plan can be generated, or when objectives can evolve over time. Thus, distributed continual planning is important for an agent whose success and efficiency depends on how its current actions affect its future choices, but which must operate in a complex, dynamic, and multi-agent environment.

We first present a historical overview of planning research that emphasizes how the introduction of more realism into the problems faced by planning systems has led to current approaches that fit the DCP paradigm more closely. We then survey related research in distributed and continual planning. We conclude by outlining some important future directions that are needed for true DCP to become a reality.

DCP: A Historical Perspective

In this section, we present a framework for categorizing planning-and-execution systems that centers around the ways in which each system handles the context in which its planning and execution are carried out. We place a range of systems within this framework, and discuss where continual planning, distributed planning, and distributed, continual planning arise in this framework.

¹The term “continuous planning” is also used to describe this style of planning, but since the latter term is also sometimes used to refer to planning for continuous problem spaces (e.g., using continuous models of time and space), we use the less ambiguous term “continual planning.”

*SRI International, 333 Ravenswood Ave., Menlo Park CA 94025; marie@erg.sri.com, ortiz@ai.sri.com, mjlw@erg.sri.com

[†]EECS Department, University of Michigan, Ann Arbor MI 48109; durfee@umich.edu

The context-centered framework gives us a structure that we use to provide a historical perspective on research in this area, which we divide into overlapping stages that are characterized by increasingly sophisticated treatment of planning and execution context. The earliest stage consists of systems that *ignore* context in order to simplify the planning problem. In the next stage, we begin to see systems that *tolerate* context by responding to or anticipating events that occur outside of the planning-and-execution system. The third stage includes systems that *exploit* context, where a system takes advantage of its ability to control the world and interact with other agents. Finally, the most sophisticated systems take the initiative to *establish* context by actively interacting with and configuring the environment.

The characterizations we provide are a simplification of the issues and accomplishments being addressed in each of the stages. In this section, we refer to survey treatments and illustrative examples of work, rather than providing a comprehensive bibliography of the rich planning literature. Our purpose is to illustrate the changing perspectives in the general planning community that have led up to the current interest in DCP.

Ignore Planning and Execution Context

Planning is difficult. Even the simplest planning problem, that of determining how an agent can move from the current world state to a world state that satisfies its preferences, is intractable in the general case. For that reason, *traditional AI planning* research (Hendler *et al.* 1990, Russell and Norvig 1995) has introduced assumptions and simplifications in order to make planning feasible. An agent is typically assumed to know everything that is relevant to the planning problem, and to know exactly how its available actions can change the world from one state into another. The planning agent is assumed to be in control of the world, so that the only changes to the state are due to the agent's deliberate actions. The agent's preferred states of the world are also constant throughout a planning episode (planning problem or set of objectives)—it will not “change its mind” about what goals to achieve in the midst of planning or while executing the plan.

These simplifications allow the planning problem to be serialized: a planning agent first formulates a plan and then executes it. It is assumed that the planning and execution for one episode have no bearing on the planning and execution done for previous or future episodes. The context in which planning and execution are done is effectively ignored during the planning and execution processes.

Tolerate Planning and Execution Context

In the real world, plans do not always proceed as expected. In dynamic, uncertain environments, an agent cannot in general make accurate predictions about the outcomes of its actions. One approach to handling uncertainty is to enumerate the possible states (contin-

gencies) that might arise at execution time and plan for each of them, constructing a (possibly large) *conditional plan* that provides alternative courses of action for each contingency (Schoppers 1987, Boutilier *et al.* 1999).

If the knowledge available to the agent is insufficient or suggests an intractably large set of contingencies, a better approach is *plan monitoring and repair* (Ambros-Ingerson and Steel 1988, Doyle *et al.* 1986). In this approach, the agent formulates a nominal plan that looks best given what (little) the agent knows. As the plan is executed, progress is monitored to ensure that the predicted or assumed conditions actually hold. When deviations are detected, the agent halts execution and revisits its planning decisions, creating a revised plan. In effect, the episodic nature of “plan-then-execute” is maintained, but episodes can be terminated prematurely (when a deviation is detected) and results from one episode can influence another (such as when the previous plan is repaired and reused). If no deviations occur, this approach is equivalent to traditional planning.

In complex environments, however, the planning and execution context might change in ways that suggest a change in future plans without necessarily violating a currently executing plan. Unexpected changes in the world might provide opportunities to accomplish goals more efficiently or effectively. Alternatively, an agent's goals might themselves change, so that although current plan could still be carried out, the motivation for doing so is lost (Levesque *et al.* 1990). When goals and aspects of the world can evolve continuously rather than being fixed throughout a planning episode, an agent should continually evaluate and revise its plans.

Continual planning recognizes that plan revision should be an ongoing process, rather than being triggered only by failure of current plans. It also adopts the perspective of not planning too much detail too far into the future, since evolving circumstances can render such details obsolete. Continual planning therefore tolerates the planning and execution context by maintaining flexibility and opportunism.

Exploit Planning and Execution Context

The approach to continual planning that we have outlined suggests an environment in which an agent is working towards accomplishing its objectives, despite the world dynamically changing in unpredictable ways that interfere with its plans. Such a pessimistic mindset need not be adopted, however, if the agent has knowledge, and the ability to use it, about its relationship to the other agents in the world. In this case, an agent should not simply tolerate the presence of other agents, but should seek to exploit their presence by cooperating with them.

One particularly useful source of interaction is human participants in the planning process. Humans often have expertise and capabilities that computational agents lack. Rather than formulating weak plans

based on its limited expertise and knowledge, therefore, an agent should whenever possible work with people to formulate more efficient and robust plans. A planning process that is distributed between computational and human agents in this way is referred to as *mixed-initiative planning* (Myers 1996, Veloso 1996, Ferguson *et al.* 1996).

A planning agent could also work with other computational agents. By bringing to bear complementary capabilities, a set of agents can collectively formulate a concerted plan of action for one or more of them to carry out (Wilkins and Myers 1998, Kambhampati *et al.* 1993). Similarly, by considering the capabilities of other agents, agents can accomplish their objectives by exploiting each other's strengths through cooperative execution.

Even if the agents do not explicitly cooperate by sharing expertise or capabilities, by exploiting the fact that changes to the world are often a result of other agents' planned actions, an agent can reduce its uncertainty, and increase the quality of its local plan, by influencing the plans that others adopt.

Establish Planning and Execution Context

The approaches to cooperative and negotiated distributed planning that were alluded to in the previous section are essentially distributed versions of traditional "plan-then-act" approaches. For a system of agents to engage in *distributed, continual planning*, however, each must be a continual planning agent, and more. As each agent elaborates its ongoing abstract plans into actions, the elaboration decisions should be compatible, and preferably mutually supporting. The agents' distributed planning mechanisms should continually evaluate the relationships between refinement decisions and redirect decisions in cooperative (or at least compatible) ways.

The idea of allowing agents to exploit the larger multi-agent context for planning (working together to plan) and for execution (taking actions towards common goals) opens the door to purposely establishing such a context in order to improve what agents can accomplish (Tambe 1997, Shoham and Tennenholtz 1992, Stone and Veloso 1999). Agents that have formulated abstract plans can analyze potential relationships between their possible plans, commit to particular constraints on how they will realize these plans, and then incorporate these influences in their elaboration decisions in a decentralized way (Clement and Durfee 1999).

Distributed Planning

The problem of constructing plans in a distributed environment has been approached from two different directions. One approach has begun with a focus on "planning," and how it can be extended into a distributed environment, where the process of formulating or executing a plan could involve actions and interactions of a number of participants. The other approach has begun with an emphasis on the problem of controlling and

coordinating the actions of multiple agents in a shared environment, and has adopted planning representations and algorithms as a means to that end. While these two approaches have led to common ground in distributed planning, and indeed many researchers embrace both approaches (making unequivocal partitioning into one camp or the other impossible), we believe that the distinctions still help in understanding the state of the field.

We refer to the first approach as "Cooperative Distributed Planning" (CDP). Because it places the problem of forming a competent (sometimes even optimal) plan as the ultimate objective, CDP is typically carried out by agents that have been endowed with shared objectives and representations, typically by a single designer or team of designers. While in some cases the purpose of the agents is to form a central plan, more generally the purpose is that the distributed parts of the developing plan will jointly execute in a coherent and effective manner. Thus, in CDP, agents typically exchange information about their plans, which they iteratively refine and revise until they fit together well.

The second approach, which we refer to as "Negotiated Distributed Planning" (NDP) has a different focus. From the perspective of an individual agent engaged in NDP, the purpose of negotiating over planned activities is not to form good collective plans, but rather to ensure that the agent's local objectives will be met by its plan, when viewed in a global context. The emphasis is therefore not on cooperatively defining and searching the space of joint plans to find the "best" group plan as in CDP, but rather on an agent trying to provide enough information to others to convince them to accommodate its preferences.

When individual preferences are aligned in ways where cooperation is in the self-interests of all concerned, CDP and NDP come together; thus, a hard-and-fast partitioning between the two approaches is impossible, and a number of efforts embrace both perspectives. For example, STEAM, which we discuss in the NDP section, draws heavily on planning research; on the other hand, both DIPART and Partial Global Planning, which are described in CDP, have a strong "agent" flavor. Our purpose in making the loose division between CDP and NDP is not to pigeonhole any individual piece of research, but instead to provide some structure to the research landscape, ranging from extreme CDP (where the purpose of distribution is simply to allow parallel computation of plans) to extreme NDP (where the purpose of planning is simply to find resource conflicts) with a rich space of systems in the middle, involving both sophisticated joint planning and the pursuit of self-interest on the parts of agents.

Cooperative Distributed Planning

In this section, we summarize past and current work in cooperative distributed planning (as defined previously). Most of the research we describe here involves two or more computer processes ("agents") cooperat-

ing to build either a single plan or multiple interacting plans. However, we will also mention other related research—for example, a single planner that has special capabilities for multi-agent execution.

In building a distributed planning system, some of the key questions to address include:

- How are plans or partial plans represented?
- What is the basic plan generation method?
- How are tasks allocated to agents?
- How do agents communicate with one another during planning?
- How are the agents' actions coordinated?

We discuss alternative approaches to these problems below.

Plan Representation and Generation Much of the research in cooperative distributed planning is built around notions of abstract plan decomposition, such as hierarchical task network (HTN) planning (Erol *et al.* 1994). An abstraction-based plan representation allows a distributed planning agent to successively refine its planning decisions as it learns more about other agents' plans.

Corkill's distributed version of Sacerdoti's NOAH planner (Sacerdoti 1977) was one of the earliest efforts in distributing an HTN planning algorithm (Corkill 1979). The basic planning procedure and plan representation used by Corkill is the same as that of NOAH: planning proceeds through a hierarchy of plan levels, where at any plan level a partial plan is a partial order of goals and primitive actions. Each distributed agent solves its goals in the same way as NOAH, at each level expanding each unplanned goal by finding an applicable operator (called a *SOUP procedure*) that solves it. NOAH's plan representation is extended to include a special node type representing a placeholder for another agent's plan, as well as special primitive actions for synchronizing multiagent execution.

Just as the SIPE-2 planner (Wilkins 1988) is conceptually descended from NOAH, DSIPE (Distributed SIPE-2) (desJardins and Wolverton, this volume) is conceptually descended from Corkill's distributed NOAH. Both are distributed versions of HTN planners, and the two systems use similar approaches to some of the key representational problems to distributed planning, especially in how the multiple agents maintain a consistent (albeit incomplete) picture of each other's plans. In DSIPE, the local view of the other agents' subplans is called a *skeletal plan*.

Similarly, in Durfee and Lesser's Partial Global Planning (PGP) (Durfee and Lesser 1991), each agent maintains a *partial global plan*, which stores its own partial picture of the plans of all the members of the group. PGP focuses on distributed execution and run-time planning and uses a specialized plan representation, where a single agent's plan includes a set of objectives, a long-term strategy (ordered lists of planned-actions)

for achieving those objectives, a set of short-term details (primitive problem-solving operations), and a number of predictions and ratings. Planning—both the expansion of objectives into planned-actions and the ordering of those actions—is based on a set of heuristics. PGP and its descendants, including Generic Partial Global Planning (Decker and Lesser 1992), have been applied to problems that include distributed acoustic monitoring (Durfee and Lesser 1991), cooperative information gathering (Decker *et al.* 1995), and cooperative robotics (Durfee and Montgomery 1991, Clement and Durfee 1999).

Task Allocation In cooperative planning systems, there is typically an explicit—often centralized—mechanism for assigning subgoals to the agents that will be responsible for planning them.

Distributed NOAH and DSIPE have no automated methods for allocating tasks among the planning agents; the user specifies which agent solves which top-level goals. PGP allows agents to send potential group plans to each other, essentially allowing agents to propose (and counter-propose) an assignment of activities (tasks) to each other (Durfee and Lesser 1989). Since agents can choose to participate or not in such plans, such assignments involve mutual selection between agents (Smith and Davis 1983).

DIPART (Pollack 1996) is an experimental platform for analyzing a wide range of planning and execution mechanisms, including approaches to distributed planning. One of the key areas studied in DIPART is load balancing (i.e., task allocation) among multiple planning agents. The load balancing methods developed are high-level approaches that are independent of any particular planning representation or algorithm. Initial allocation of tasks to agents is specified by the user, but tasks can be automatically reallocated at runtime according to load balancing considerations.

Lansky's COLLAGE (Lansky 1994, Lansky and Getoor 1995) is not a multi-agent planning system *per se*, but it represents one promising way of approaching the task assignment problem in distributed planning. COLLAGE uses a technique called *localization* to decompose a planning problem into subproblems called *regions*. Localizations in COLLAGE can be automatically generated based on abstraction levels or scope. A set of heuristics are used to find regions that minimize the number of interactions between regions, thus permitting COLLAGE to solve the subproblems in each region (relatively) independently. Partitioning subgoals among distributed planning agents according to localization would likewise improve the overall efficiency of the multi-agent planning system, by minimizing conflicts and interactions between agents (and reducing the need to check for them).

Communication Cooperative distributed planners use a variety of techniques for agents to share information.

At the conclusion of each planning level, Distributed

NOAH applies an extended version of NOAH's planning critics to coordinate the planning agents. Each agent sends the other agents a record of the important effects of the newly-expanded portions of its plan (though Corkill does not make it clear which effects are considered to be important). Each agent inserts a record of all new effects it receives into its local plan.

DSIPE extends Distributed NOAH's approach in a number of areas, most notably by expanding the types of planning information that are shared among the planning agents, and by developing irrelevance-based filtering methods for reducing communication requirements (Wolverton and desJardins 1998).

In PGP, communication goes through the channels of a *meta-level organization*, which specifies the roles and responsibilities of planning agents in maintaining plan coordination (e.g., which agents have authority to reconcile conflicts among agents' plans), much like a regular organization specifies how the agents should be organized to solve a domain-level problem. The work of Jennings (Jennings 1995) codifies such responsibilities into communication *conventions* that all agents follow to ensure sufficient awareness about the pursuit of collective plans.

Communication among planning agents to support load balancing in DIPART is generally done by multicasting to all members of the agent's (user-assigned) subgroup. However, in heavy load situations, agents may trade off complete knowledge of system load in exchange for preservation of bandwidth by selective unicasting (called "focused addressing" in (Smith and Davis 1983)). In this case, they must rely on estimates of other nodes' loads based on their histories.

Coordination For a group of cooperative distributed planners to reach a state in which the individual subplans jointly achieve the common objectives, the agents must coordinate their decision making.

In Distributed NOAH, when an agent's critics detect a conflict between one of its own actions and those of another agent, instead of adding an ordering constraint as it would if were not distributed, it inserts a special execution-time coordination action: either "Signal," which will inform another agent of the completion of an action, or "Wait," which waits for a signal from another agent.

PGP achieves coordination and task assignment by a process of negotiation. When agent *A* informs agent *B* about some portion of its local plan, *B* merges the new information about *A* into its partial global plan. It then searches for ways to improve the global plan, e.g., by eliminating redundancy or by better utilizing the group's computational resources. It proposes these improvements to some subset of the other agents. Each other agent may accept the new global plan and incorporate the changes made to its own local plan, it may reject the new global plan, or it may respond with a counterproposal. Throughout the process, each agent will continue to act on its current best assessment of

the plan it should pursue, rather than waiting to converge on a plan (which, due to dynamics, might be quite ephemeral anyway). Thus, this approach is designed for domains where temporary incoherence among agents is tolerated as information about unanticipated changes to plans is propagated. It is not designed for domains with irreversible actions, where conflicts among agents' plans should be resolved before any step is taken that could cause other agents' plans (and thus the system as a whole) to fail.

Coordination between planning agents in DIPART is accomplished in two ways: (1) Incremental merging of individual agent subplans (Ephrati *et al.* 1995a), and (2) *multi-agent filtering* (Ephrati *et al.* 1995b), in which each agent commits to the goals it has already adopted, but bypasses new options that are incompatible with any other agent's known or presumed goals.

Negotiated Distributed Planning

Much of the research into the design of autonomous agents has had as its goal a formal account of mental state that would serve to explain the manner in which the attitudes of belief, desire, and intention (BDI) engender behavior, both at the individual agent level and at the group level. A BDI theory could be viewed as explanatory to the extent that one felt comfortable in ascribing some sense of "rationality" to the behaviors predicted by the theory that were causally connected to individual mental states; in this respect, the term *rational balance* as applied to an agent's mental state and as coined by Nils Nilsson, is an apt one. For example, an intention can be viewed as reflecting an agent's commitment to perform an action; such a commitment is assumed to remain in force until satisfied or until the agent no longer believes the action is possible (Bratman 1987, Cohen and Levesque 1990, Shoham 1990, Wooldridge and Jennings). By virtue of its persistence, an intention will thereby constrain the options an agent will be disposed to consider; a rationally balanced agent would then be one that was not continuously reevaluating each possible option in light of new intentions or goals. Within such models, plans are derivative structures built up from elements of the agent's mental state (Pollack 1990).

The BDI decomposition strategy can also serve as an implementation-neutral way to architect an agent and to also analyze both individual and collective behavior. In contrast to the above research, which has largely been formal in nature, the design of rational-agent architectures (Haddadi and Sundermeyer 1996, Bratman *et al.* 1988) has been guided primarily by computational concerns centered on an agent's inherent resource limitations (Simon 1983, Russell and Subramanian 1995). Agent architectures augment an agent's mental world with a set of processes, each associated with a particular architectural module; processes can either supply inputs for other modules or modify the agent's repository of beliefs, desires, and intentions. As an example, an architecture might include a delibera-

tion module whose role is to weigh alternatives in light of existing commitments (intentions) and new observations (beliefs). A connection between formal theories of mental state and agent architectures (or implementations) can be drawn by reifying mental actions, each representing the behavior of an architectural module (Ortiz 1999).

Collaboration Theories of mental state for individual agents must be extended in important ways in order to model collaborations between agents. The problem is that a group's plan to collaborate on some task is not simply the sum of the agents' individual plans (Grosz and Kraus 1996, Bratman 1992, Searle 1990). In particular, some notion of an individual's commitment to (parts of) the group activity must be captured. One way of accomplishing this is by introducing the idea of a joint intention (Levesque *et al.* 1990, Sonenberg *et al.* 1992), which is meant to represent the group's commitment towards some group goal. One can also augment an agent architecture to include processes that capture behavioral conventions for modifying commitments and for supporting coordinated behavior (Jennings 1995). One problem with the joint intentions model, however, is that communication points must be built in and cannot be derived automatically.

The theory of SharedPlans is an alternative approach that builds group plans out of ordinary beliefs and intentions, together with a new attitude called an *intention-that* (Grosz and Kraus 1996, Grosz *et al.*, this volume). Intentions-that model commitments to states of the world instead of actions. Such intentions can represent, among other things, an individual agent's commitments to aspects of the world that will ensure the success of the group activity; various stages of partiality of a shared plan can thereby be modeled (corresponding to, for example, partial beliefs about how to perform an action, the identity of appropriate agents for a group task, and the knowledge preconditions for an action). Theories based on this idea grew out of work in natural language understanding (Grosz and Sidner 1986), and have been applied to discourse understanding (Lochbaum 1998) and the development of collaborative user interfaces (Ortiz *et al.* 1999). Unlike the irreducible joint-intentions model, communications are dynamically triggered as a consequence of failed intentions-that.

Some systems have been developed that combine both approaches; for example, while the STEAM system (Tambe 1997) is based on the joint intentions model, it also borrows some ideas from SharedPlans for representing partial plans, and integrates ideas about teamwork capabilities into the design of individual agent architectures.

Decision theory provides a computational framework for making decisions about coordination and communication in a distributed planning context. Both STEAM and Gmytrasiewicz *et al.*'s Recursive Modeling Method (RMM) approach (Gmytrasiewicz *et al.* 1991) have

used models of the costs and benefits of communication and applied decision-theoretic methods to determine what and when to communicate. Boutilier (this volume) describes approaches for integrating decision-theoretic planning and multiagent systems.

Negotiation Given a representation for group plans, there remains the question of how agents should negotiate to reach an agreement over distribution of tasks and resources. One way to ensure coordination of agent activities is to embed them in environments in which certain social laws must be adhered to (such as the traffic laws in (Shoham and Tennenholtz 1992)). This builds coordination into the environment by focusing on *public behavior* rather than on private preferences. Other approaches make use of ideas from game theory, such as notions of solution stability and equilibrium, and assume that agents are developed independently and are self-interested: that is, they maximize individual utility. Since agents are not necessarily assumed to be cooperative, much of the effort is directed to developing protocols that will prevent lying. One attractive strategy is to characterize different domains and identify appropriate protocols for each type of domain (Rosenschein and Zlotkin 1994).

Other ideas that have been adapted from the field of economics include negotiation strategies that take time into account, explicitly handling the resource-constrained nature of agents (Osborne and Rubinstein 1990, Kraus *et al.* 1992). Agents are assumed to have a common meta-goal; to minimize the amount of work that each agent must do. In the case of agents who do not have a common, fixed goal during negotiation, an alternative is for agents to construct arguments that they can use to influence the decision making of other agents (Sycara 1990).

Voting schemes, also borrowed from the field of economics, have been applied to achieving consensus. For example, in one multi-agent scheduling system, agents first express their individual preferences, and then a voting mechanism is used to determine group choice (Sen *et al.*). When the assumption of cooperation is dropped, systems can become susceptible to lying. Many approaches around this problem have been proposed: one, for example, requires that agents pay a tax (such as the Clarke tax (Ephrati and Rosenschein 1991)) that, in effect, transfers utility outside the system. Agents are still assumed to be self-motivated and maximize their own utility, but it can be shown that there is only one dominant strategy: telling the truth.

Contract nets (Smith and Davis 1983) represent a market-based approach to negotiation. Contract nets provide a high-level communication protocol in which tasks are distributed among nodes in a system. Nodes are classified as either contractor or manager nodes, and contracts to perform tasks are established through a bidding process. A similar approach is the "consumers and producers approach" (Wellman 1992), in which agents are classified as either consumers or producers,

and each type of agent tries to maximize its own utility. Goods have auctions associated with them, and agents can acquire goods by submitting bids in the auction for that commodity. A more dialogue-oriented approach is found in the COllaborative Negotiation System based on Argumentation (CONSA) (Tambe and Jung, this volume), which uses negotiation methods based on argumentation structures (Toulmin 1958) to resolve conflicts.

A major challenge that remains is the development of models that encompass aspects of the partiality of both mental state and the planning process itself. Methods must be developed for adapting the various approaches in a way that is consistent with the resource-constrained nature of planning agents: planning should be a continuous, incremental process at both the individual and group level.

Continual Planning

An agent working in a world where unexpected hindrances or opportunities could arise should continually be on the lookout for changes that could render its planned activities obsolete. Because its plans can undergo continual evaluation and revision, such an agent will be continually planning, interleaving planning with execution.

Reactive planning systems (Agre and Chapman 1987) can be viewed as a special case of continual planning in which “planning” looks ahead only to the next action. Because of their near-term focus, such systems generally do not handle problems that require complex orderings of actions to accomplish tasks. The assumption in these systems is that such orderings should not be pursued anyway because the later actions are inappropriate by the time they are reached.

Flexible plan execution systems such as PRS (Georgeff and Lansky 1986) and RAPS (Firby 1987) strike a compromise between looking ahead to sequences of actions and avoiding commitments to specific future actions by exploiting hierarchical plan spaces. Rather than refine abstract plan operators into a detailed end-to-end plan, these systems interleave refinement with execution. Plan refinement is delayed as long as possible, so that detailed decisions are made with as much information as possible. However, this can mean that refinement decisions at abstract levels are made and acted upon before all of the detailed refinements have been made. If these abstract refinements introduce unresolvable conflicts at lower levels, the plan may fail in the middle of execution. It is therefore critical that the specifications of abstract plan operators be rich enough to summarize all (or at least most) of the relevant refinements in order to anticipate and avoid such conflicts (Clement and Durfee 1999).

The Continuous Planning and Execution Framework (CPEF—Myers, this volume) continually constructs and revises plans that evolve in response to a dynamically changing environment. CPEF integrates HTN planning techniques, plan execution and monitoring,

and dynamic plan repair methods. Users can interact closely with the system through the Advisable Planner (AP). The execution component can execute plans at multiple levels of abstraction, permitting an open-ended planning process in which the level of detail of plans can vary depending on the timing and nature of the particular planning task being considered.

The costs of execution monitoring and opportunistic replanning must be considered by continual planning systems. In many respects, execution monitoring is simpler: as an agent steps through its planned actions, it periodically checks to make sure that the state of the world that it has reached is consistent with the state that was predicted by the plan. The trick is to keep the costs of monitoring low enough so that they do not consume too much time and effort on the part of the agent, either by devising efficient perceptual strategies (Doyle *et al.* 1986, Musliner *et al.* 1995) or by performing the checks less frequently. Looking for new opportunities is more problematic: unlike simple monitoring, where the agent knows what to look for (which conditions are necessary for the rest of the plan to succeed), finding opportunities is generally a more open-ended problem. Researchers have explored tradeoffs between having *bold* agents, who seldom reconsider their plans, and *cautious* agents, who frequently reevaluate their plans (Kinny and Georgeff 1991).

Different continual planning approaches implement ongoing monitoring and replanning tasks differently. For example, UM-PRS (Lee *et al.* 1994) maintains the hierarchy of plans currently in progress, and before each action traverses this hierarchy from the top down to ensure that the motivations for the next action are still in force. If during this process it encounters a plan abstraction that is now deemed inappropriate, it will consider alternative ways of achieving the goal it had been working on, and will look for opportunities for directing its attention elsewhere. In contrast, the Soar architecture (Laird *et al.* 1987) can also perform continual planning, but does not explicitly reason over a hierarchy of intentions. Instead, it has been extended to generate rules that explicitly monitor the current plan’s context (Wray and Laird 1998). Because these rules are woven into Soar’s rule network, they do not require explicit periodic polling, but instead trigger interrupts to Soar when appropriate.

Future Research Directions

While distributed planning has been explored by a few researchers for many years, the recognition of its importance in current and future applications involving networked agents has triggered a more broad-based, concerted investigation into this still-maturing field. Reasoning and negotiation techniques are needed that permit distributed planning agents to understand the aspects of the distributed plan that are relevant for their decisions, identify and pursue opportunities for coordination, asynchronously refine their plans as other

agents' plans are evolving, and modify their plans in response to information that arrives from other agents.

In the area of continual planning, open-ended planning, in which plans can be continually refined to varying levels of abstraction as the planning horizon is extended, will be a key capability. Plan repair methods in current systems are typically either at the reactive level or at the generative level; methods for smoothly integrating plan repair techniques at multiple levels of abstraction and varying time scales are needed.

Several papers in this volume describe ongoing research efforts that are focused on true DCP—that is, methods for distributed *and* continual planning in dynamic, uncertain domains. Pollack and Harty (this volume) describe a plan management system that goes beyond simply interleaving decisions about planning and execution by incorporating techniques for monitoring the environment, assessing alternative plans, iteratively elaborating abstract partial plans, controlling the meta-level planning process, and coordinating distributed planning agents. Durfee (this volume) describes methods for continual, cooperative elaboration and revision of distributed partial plans in a realistic application domain of semi-autonomous ground vehicles. Myers (this volume) presents a framework for continuous planning and execution that is partially distributed (that is, the planning and execution capabilities are distributed, though the planning process itself is not).

However, most of this research is still in its early stages, and many research challenges remain. Both distributed planning and continual planning need to be better understood for DCP systems to be built; yet simply combining distributed and continual planning methods that have been independently developed may not be sufficient. The synchronization problem in distributed planning becomes even greater when the distributed agents' plans are being executed concurrently, and better models of the overall planning and execution process are needed. Similarly, execution of plans is complicated by the presence of other agents, requiring not only methods for distributed plan execution but distributed plan repair.

References

- P. E. Agre and D. Chapman. Pengi: An implementation of a theory of activity. In *IJCAI-87*, pages 268–272, 1987.
- J. A. Ambros-Ingerson and S. Steel. Integrating planning, execution, and monitoring. In *AAAI-88*, pages 81–88, 1988.
- Craig Boutilier, Tom Dean, and Steve Hanks. Decision theoretic planning: Structural assumptions and computational leverage. *Journal of AI Research*, 1999. To appear.
- Michael E. Bratman, David J. Israel, and Martha E. Pollack. Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4:349–355, 1988.
- Michael Bratman. *Intentions, Plans, and Practical Reason*. Harvard University Press, 1987.
- Michael E. Bratman. Shared cooperative activity. *The Philosophical Review*, 101(2):327–341, 1992.
- Bradley J. Clement and Edmund H. Durfee. Top-down search for coordinating the hierarchical plans of multiple agents. In *Third Conference on Autonomous Agents*, 1999.
- Philip Cohen and Hector Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.
- Daniel D. Corkill. Hierarchical planning in a distributed environment. In *IJCAI-79*, 1979.
- Keith S. Decker and Victor R. Lesser. Generalizing the partial global planning algorithm. *International Journal of Intelligent and Cooperative Information Systems*, 1:319–346, 1992.
- Keith Decker, Victor Lesser, M.V. Nagendra Prasad, and Thomas Wagner. MACRON: An architecture for multi-agent cooperative information gathering. In *Proceedings of the CIKM Workshop on Intelligent Information Agents*, December 1995.
- R. J. Doyle, D. J. Atkinson, and R. S. Doshi. Generating perception requests and expectations to verify the execution of plans. In *AAAI-86*, pages 81–88, 1986.
- Edmund H. Durfee and Victor R. Lesser. Negotiating task decomposition and allocation using partial global planning. In M. Huhns and L. Gasser, editors, *Distributed Artificial Intelligence*, volume 2. Morgan Kaufmann, 1989.
- Edmund H. Durfee and Victor R. Lesser. Partial global planning: A coordination framework for distributed hypothesis formation. *IEEE Transactions on Systems, Man, and Cybernetics*, KDE-1:63–83, 1991.
- Edmund H. Durfee and Thomas A. Montgomery. Coordination as distributed search in a hierarchical behavior space. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-21(6):1363–1378, November 1991. Special Issue on Distributed Artificial Intelligence.
- E. Ephrati and J. Rosenschein. The Clarke Tax as a consensus mechanism among automated agents. In *Proceedings of AAAI-91*, pages 173–178, 1991.
- Eithan Ephrati, Martha E. Pollack, and Jeffrey S. Rosenschein. A tractable heuristic that maximizes global utility through local plan combination. In *Proceedings of the 1st International Conference on Multi-Agent Systems (ICMAS-95)*, 1995.
- Eithan Ephrati, Martha E. Pollack, and Sigalit Ur. Deriving multi-agent coordination through filtering strategies. In *Proceedings of the 14th Joint International Conference on Artificial Intelligence (IJCAI-95)*, 1995.
- K. Erol, D. Nau, and J. Hendler. HTN planning: Complexity and expressivity. In *AAAI-94*, 1994.

- G. Ferguson, J. Allen, and B. Miller. TRAINS-95: Towards a mixed initiative planning assistant. In *Third International Conference on Artificial Intelligence Planning Systems*, 1996.
- R. J. Firby. An investigation into reactive planning in complex domains. In *AAAI-87*, pages 202–206, 1987.
- Michael P. Georgeff and Amy L. Lansky. Procedural knowledge. *Proceedings of the IEEE*, 74:1383–1398, 1986.
- Piotr J. Gmytrasiewicz, Edmund H. Durfee, and David K. Wehe. The utility of communication in coordinating intelligent agents. In *AAAI-91*, pages 166–172, 1991.
- Barbara J. Grosz and Sarit Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(1):269–357, 1996.
- Barbara J. Grosz and Candace Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3):175–204, 1986.
- Afsaneh Haddadi and Kurt Sundermeyer. *Belief-Desire-Intention Agent Architectures*, pages 169–186. John Wiley and Sons, Inc, 1996.
- James Hendler, Austin Tate, and Mark Drummond. AI planning: Systems and techniques. *AI Magazine*, 11:61–77, 1990.
- N. Jennings. Controlling cooperative problem in industrial multi agent systems using joint intentions. *Artificial Intelligence*, 75(2):195–240, 1995.
- S. Kambhampati, M.R. Cutkosky, J.M Tenenbaum, and S.H. Lee. Integrating general purpose planners and specialized reasoners: Case study of a hybrid planning architecture. *IEEE Transactions on Systems, Man, and Cybernetics*, 23:1503–1518, 1993.
- David N. Kinny and Michael P. Georgeff. Commitment and effectiveness of situated agents. In *IJCAI-91*, pages 82–88, 1991.
- Sarit Kraus, Jonathan Wilkenfeld, and Gilad Zlotkin. Multiagent negotiation under time constraints. Technical Report CS-TR-2975, University of Maryland, College Park, 1992.
- J. E. Laird, A. Newell, and P. S. Rosenbloom. SOAR: An architecture for general intelligence. *Artificial Intelligence*, 33:1–64, 1987.
- Amy L. Lansky and Lise C. Getoor. Scope and abstraction: Two criteria for localized planning. In *IJCAI-95*, 1995.
- Amy L. Lansky. Located planning with diverse plan construction methods. Technical Report FIA-TR-9405, NASA Ames Research Center, 1994.
- Jaeho Lee, Marcus J. Huber, Edmund H. Durfee, and Patrick G. Kenny. UM-PRS: An implementation of the Procedural Reasoning System for multirobot applications. In *AIAA/NASA Conference on Intelligent Robotics in Field, Factory, and Space*, pages 842–849, 1994.
- H. Levesque, P. Cohen, and J. Nunes. On acting together. In *Proceedings of AAAI-90*, pages 94–99, 1990.
- Karen E. Lochbaum. A collaborative planning model of intentional structure. *Computational Linguistics*, 34(4):525–572, 1998.
- David J. Musliner, Edmund H. Durfee, and Kang G. Shin. World modeling for the dynamic construction of real-time control plans. *Artificial Intelligence*, 74:83–127, 1995.
- Karen L. Myers. Advisable planning systems. In Austin Tate, editor, *Advanced Planning Technology: Technological Achievements of the ARPA/Rome Laboratory Planning Initiative*. AAAI Press, 1996.
- Charles Ortiz, Barbara Grosz, and Nathan Scales. Interpreting information requests in context: a collaborative web interface for distance learning. *Autonomous Agents and Multi-agent Systems*, 1999. Submitted.
- Charles L. Ortiz. Introspective and elaborative processes in rational agents. *Annals of Artificial Intelligence and Mathematics*, page To appear, 1999.
- M. J. Osborne and A. Rubinstein. *Bargaining and Markets*. Academic Press, California, 1990.
- Martha Pollack. *Plans as Complex Mental Attitudes*, pages 77–103. MIT Press, 1990.
- Martha E. Pollack. Planning in dynamic environments: The “DIPART” system. In Austin Tate, editor, *Advanced Planning Technology: Technology Achievements of the ARPA/Rome Laboratory Planning Initiative*. Morgan Kaufmann, 1996.
- Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter*. MIT Press, 1994.
- Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
- S. J. Russell and D. Subramanian. Provably bounded-optimal agents. *Journal of Artificial Intelligence Research*, 2:575–609, 1995.
- Earl D. Sacerdoti. *A Structure for Plans and Behavior*. American Elsevier, 1977.
- Marcel J. Schoppers. Universal plans for reactive robots in unpredictable environments. In *IJCAI-87*, pages 1039–1046, 1987.
- J. Searle. *Collective Intentions and Actions*, pages 401–415. MIT Press, 1990.
- Sandip Sen, Thomas Haynes, and Neeraj Arora. Satisfying user preferences while negotiating meetings. *International Journal of Human-Computer Systems*, page ???, ???
- Y. Shoham and M. Tennenholtz. On the synthesis of useful social laws for artificial agent societies. In *Proceedings of AAAI-92*, pages 276–281, 1992.
- Yoav Shoham. Agent-oriented programming. Technical Report STAN-CS-90-1335, Stanford University, 1990. Computer Science Department.
- H. Simon. *Alternative Visions of Rationality*. Stanford University Press, 1983.

R.G. Smith and R. Davis. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence*, pages 63–109, 1983.

E. Sonenberg, G. Tidhar, E. Werner, D. Kinny, M. Ljungberg, and A. Rao. Planned team activity. Technical Report 26, Australian Artificial Intelligence Institute, 1992. Revised JJanuary 1994.

Peter Stone and Manuela Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence*, 1999. To appear.

Katya Sycara. Persuasive argumentation in negotiation. *Theory Decision*, 28:203–242, 1990.

Milind Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.

S. Toulmin. *The uses of argument*. Cambridge University Press, 1958.

Manuela M. Veloso. Towards mixed-initiative rationale-supported planning. In Austin Tate, editor, *Advanced Planning Technology: Technological Achievements of the ARPA/Rome Laboratory Planning Initiative*, pages 277–282. AAAI Press, 1996.

Michael Wellman. A general-equilibrium approach to distributed transportation planning. In *Proceedings of AAAI-92*, pages 282–289, 1992.

David E. Wilkins and Karen L. Myers. A multiagent planning architecture. In *AIPS-98*, pages 154–162, 1998.

David E. Wilkins. *Practical Planning: Extending the Classical AI Planning Paradigm*. Morgan Kaufmann, 1988.

Michael J. Wolverton and Marie desJardins. Controlling communication in distributed planning using irrelevance reasoning. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*. AAAI Press, 1998.

M. J. Wooldridge and N. R. Jennings. Agent theories, architectures, and languages: A survey. ???, page ???, ???

Robert Wray and John Laird. Maintaining consistency in hierarchical reasoning. In *AAAI-98*, 1998.