Chapter 13

Using Planning to Adapt to Dynamic Environments

Austin Tate

Contents

13.1 Introduction

Planning is about much more than solving specifically stated problems whereby some goal state is reached from some initial state as efficiently as possible. The real world is a messy place. The current state of the world may be only partially known or observable. The goal, objective, or mission itself may be imprecisely stated, and the agents available to carry out the activities involved may be only partially specified. The model we have of the state, objectives, and agent capabilities may be imperfect. People and systems often should work in harmony as a team to solve problems and accommodate the roles, capabilities, and preferences of the various agents. The real world is also dynamic and changing—the state of the environment, the objectives, and the agents or their capabilities can be in a dynamic state of flux.

Artificial Intelligence (AI) planning and knowledge-rich plan representation techniques have been developed to generate, refine, and adapt plans in highly dynamic situations to provide resilience. They seek to address some of the real messy problems in the world.

Realistic planning systems must allow users and computer systems to cooperate and work together using a "mixed initiative" style. Black box or fully automated solutions are not acceptable in many situations. Studies of expert human problem solvers in stressful or critical situations (Klein, 1998) show that they share many of the problem-solving methods employed by some of the methods studied in AI planning to address these issues (Tate, 2000b).

There is also a need to model domains in which planning takes place, understanding the roles and capabilities of the various human and system agents involved in the planning process and in the domain in which plans are executed, and allow for communication of information about tasks, plans, intentions, and effects between those agents.

This paper argues that a hierarchical task network (HTN) least commitment planning and plan refinement approach—as used for many years in practical planning systems, such as NOAH (Sacerdoti, 1977), Nonlin (Tate, 1977), SIPE (Wilkins, 1988), O-Plan (Currie and Tate, 1991; Tate et al., 1994b) and SHOP (Nau et al., 2005)—provides an intelligible framework for mixed-initiative multi-agent human/system planning environments. When joined with a strong underlying constraint-based plan representation it can provide a framework in which powerful problem solvers, based on search and constraint reasoning methods, can be employed to work in highly dynamic situations and still retain human intelligibility.

I-Plan and its underlying <I-N-C-A> (Issues – Nodes – Constraints – Annotations) ontology is a planner created in the I-X intelligent systems framework, which follows these principles.

13.2 Development of a Flexible AI Planning Approach

Realistic planning systems must allow users and computer systems to cooperate and work together using a "mixed initiative" style. Black box or fully automated solutions are not acceptable in many situations, in which human responsibility is paramount. Highly dynamic environments demand adaptable solutions. Studies of expert human problem solvers in stressful or critical situations show that they share many of the problem-solving methods employed by hierarchical planning methods studied in AI. But powerful solvers and constraint reasoners can also be of great help in parts of the planning process. A more intelligible approach to using AI planning is needed that can use the best "open" styles of planning based on shared plan representations and HTNs and which still allow the use of powerful constraint representations and solvers.

The field of AI planning—that is, reasoning about the activity necessary to achieve stated goals—has a long and distinguished history (Allen et al., 1990; Tate, 1996a). Notwithstanding its successes, most work is based on simplifications and unrealistic general assumptions, which restrict the application of planning algorithms to specific problems under specific conditions. These unrealistic assumptions can be summarized as follows: (a) the presence of an omniscient agent able to formulate centralized, all-encompassing plans; (b) action schemata that capture the totality of conditions under which they are applicable and of effects they bring about; and (c) an environment that is unaffected by external agency, being changed only by the projected actions in a plan.

While research into specialized algorithms has continued, often leading to notable improvements, a shift of emphasis is needed to support planning in dynamic environments and in cooperation with human planners addressing real tasks. One of the key insights is to recognize the value of AI work in the representation of plans rather than in any particular algorithm and that real planning is as much a social activity as a computational task.

This insight guided the development of the open planning architecture (O-Plan) (Currie and Tate, 1991) and its development into one of the first web-based task-support applications (Tate, 1996b; Tate et al., 2003; Tate and Dalton, 2003) and the gradual distillation and refinement of previous plan representations into the <I-N-C-A> ontology (Tate, 1998; Tate, 2003). This model can be used to describe not only plans, but also the planning process itself and, hence, to communicate aspects of this task, raising it to the level of a collaborative social activity in an approach we term intelligible planning (Tate, 2000b).

To encourage and support this shift from automated reasoning to distributed collaboration, a generic set of software tools and documentation, collectively called the I-X intelligent systems suite, has been developed (Tate et al., 2002). I-Plan is a planning system based on these principles. It is part of the I-X suite of intelligent tools. I-Plan is modular and can be extended via plug-ins of various types. It is intended to be a "lightweight" planning system that can be embedded in other applications. In its simplest form, it can provide a small personal planning aid that can be deployed in portable devices and other user-orientated systems to add planning facilities into them. In its more developed forms, it can have the power of longer-established generative HTN AI planners, such as O-Plan.

13.3 I-X: Intelligent Systems Architecture

The I-X approach has five aspects:

1. Systems integration: A broad vision of an open architecture for the creation of intelligent systems that support the "process" for the synthesis of a result or "product." It is based on a "two-cycle" approach, which uses plug-in components to "handle issues" and to "maintain the domain model."
2. <I-N-C-A> ontology: A core notion of the representation of a process or plan as a set of nodes making up the components of the process or plan model along with constraints on the relationship between those nodes. It includes a set of outstanding issues and can maintain annotations for various purposes, including rationale capture.
3. Reasoning: The provision of plug-in reasoning capabilities in the form of "issue handlers" and "constraint managers."
4. Viewers and user interfaces: To support various roles of users performing activities and to provide modules that present the state of the process they are engaged in and the status of the products they are working with.
5. Applications: Work in various application sectors that seeks to create generic approaches (I-Tools) for the various types of tasks in which users may engage. One important application is I-Plan for planning tasks. See Figure 13.1.

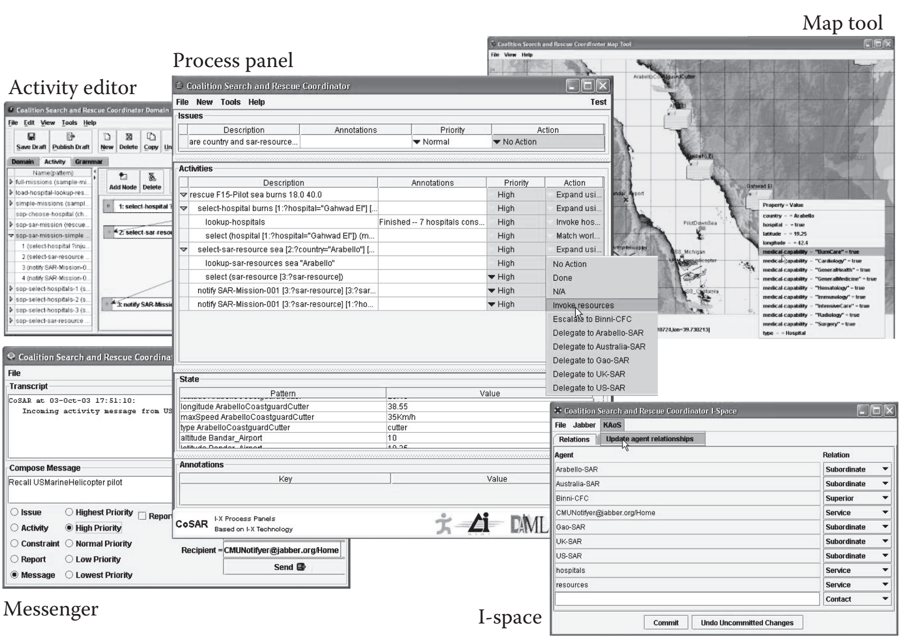

Figure 13.1 I-X task support tools.

13.4 Features of the "Intelligible Planning" Approach

There are a number of features that can encourage an approach to planning that is intelligible to the people responsible for the process and involved in planning and execution:

- Expansion of a high-level abstract plan into greater detail when necessary.
- High-level "chunks" of procedural knowledge (standard operating procedures, best practice processes, tactics, techniques, and procedures, etc.) on a human scale—typically five to eight actions—can be manipulated within the system.
- Ability to establish that a feasible plan exists, perhaps for a range of assumptions about the situation, while retaining a high-level overview.
- Analysis of potential interactions as plans are expanded or developed.
- Identification of problems, flaws, and issues with the plan.
- Deliberative establishment of a space of alternative options perhaps based on different assumptions about the situation involved, of special use ahead of time, in training and rehearsal, and to those unfamiliar with the situation or utilizing novel equipment.
- Monitoring of the execution of events as they are expected to happen within the plan, watching for deviations that indicate a necessity to re-plan (often ahead of this becoming a serious problem).
- Represent the dynamic state of the world at points in the plan and use this for "mental simulation" of the execution of the plan.
- Pruning of choices according to given requirements or constraints.
- Situation-dependent option filtering (sometimes reducing the choices normally open to one "obvious" one).
- Satisficing search to find the first suitable plan that meets the essential criteria.
- Heuristic evaluation and prioritization of multiple possible choices within the constraint search space.
- Uniform use of a common plan representation with embedded rationale to improve plan quality, shared understanding, etc.

The previously described features describe many aspects of problem-solving behavior observed in expert humans working in unusual or crisis situations (Klein, 1998). But they also describe the hierarchical and mixed-initiative approach to planning in AI developed over the last four decades.

13.5 A More Intelligible Framework for AI Planning: The I-X Approach

The I-X approach involves the use of shared models for task-directed communication between human and computer agents who are jointly exploring (via some "process") a range of alternative options for the synthesis of an artifact, such as a design or a plan (termed a "product"). It allows for two levels:

- Outer level: human-relatable plan representations and HTN planning style for outer level.
- Inner level: detailed search, constraint solvers, analyzers, and simulations act in this framework to provide feasibility checks, detailed constraints, and guidance.

It also provides for

- Sharing of issues, processes, and process products between humans and systems described via <I-N-C-A>
- Secure policy-managed communications, reporting, logging
- Context, environment, and agent capability–sensitive option generation
- Links between informal/unstructured outline planning and more structured, detailed planning

The I-X system or agent has two processing cycles (see Figure 13.2):

- Handle issues
- Respect domain constraints

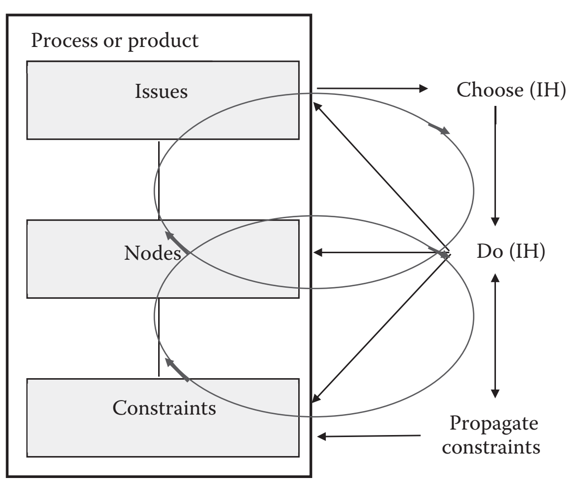

Figure 13.2 I-X approach two cycles: handle issues, propagate constraints.

An I-X system or agent carries out a (perhaps dynamically determined) process, which leads to the production of (one or more alternative options for) a synthesized artifact.

An I-X system or agent views the synthesized artifact as being represented by a set of constraints on the space of all possible artifacts in the domain.

13.6 I-Plan

The I-Plan design provides an extensible framework for adding detailed constraint representations and reasoners into planners. These can be based on powerful automated methods. But this can be done in a context that provides overall human intelligibility.

The I-Plan design is based on two cycles of processing. The first addresses one or more "issues," and the second ensures that constraints in the domain in which processing takes place are checked and respected. So the processing cycles can be characterized as "handle issues, respect constraints."

The emerging partial plan is analyzed to produce a further list of issues and added constraints. A choice of the issues to address is used to drive a workflow-style processing cycle of choosing "issue handlers" and then executing them to modify the emerging plan state. Checks are then made on the sets of constraints available to check their validity, to add further deduced constraints via propagation, and to signal any indicated or potential constraint violations. In some cases, sophisticated constraint managers can give "maybe" answers when constraints are added, giving vital information on possible fixes or alternatives for adding constraints such that the set of constraints can be made consistent again and problem solving can continue (Dalton et al., 1993; Tate et al., 1994a).

This approach is taken in systems like O-Plan, OPIS (Smith, 1994), DIPART (Pollack, 1994), TOSCA (Beck, 1993), etc. The approach fits well with the concept of treating plans as a set of constraints that can be refined as planning progresses.

Some such systems can also act in a nonmonotonic fashion by relaxing constraints in certain ways. Having the implied constraints or an "agenda" as a formal part of the plan provides an ability to separate the plan that is being generated or manipulated from the planning system and process itself, and this is used as a core part of the I-Plan design.

Mixed initiative planning approaches, for example in O-Plan (Tate, 1994), improve the coordination of planning with user interaction by employing a clearer shared model of the plan as a set of constraints at various levels that can be jointly and explicitly discussed between and manipulated by the user or system in a cooperative fashion. I-Plan adopts this approach.

13.7 <I-N-C-A>

The <I-N-C-A> model is a means to represent plans and activity as a set of constraints. By having a clear description of the different components within a plan, the model allows for plans to be manipulated and used separately from the environments in which they are generated. The underlying thesis is that plans can be represented by a set of constraints on the behaviors possible in the domain being modeled and that plan communication can take place through the interchange of such constraint information.

The <I-N-C-A> representation is intended to utilize a synergy of practical and formal approaches that are stretching the formal methods to cover realistic representations as needed for real problem solving and can improve the analysis that is possible for practical planning systems.

The <I-N-C-A> constraint model provides support for a number of different uses:

- For automatic and mixed-initiative generation and manipulation of plans and other synthesized artifacts and to act as an ontology to underpin such use
- As a common basis for human and system communication about plans and other synthesized artifacts
- As a target for principled and reliable acquisition of plans, process models, and process product information
- To support formal reasoning about plans and other synthesized artifacts

These cover both formal and practical requirements and encompass the requirements for use by both human and computer-based planning and design systems.

When first designed (Tate, 1996b), <I-N-C-A> was intended to act as a bridge to improve dialogue between a number of communities working on formal planning theories, practical planning systems, and systems engineering process management methodologies. It was intended to support new work then emerging on automatic manipulation of plans, human communication about plans, principled and reliable acquisition of plan information, and formal reasoning about plans. It has since been utilized as the basis for a number of research efforts, practical applications, and emerging international standards for plan and process representations. For some of the history and relationships between earlier work in AI on plan representations, work from the process and design communities and the standards bodies, and the part that <I-N-C-A> played in this see Tate (1998).

In Tate (2000a), the <I-N-C-A> model is used to characterize the plan representation used within O-Plan and is related to the plan refinement planning method used in O-Plan. The <I-N-C-A> work is related to emerging formal analyses of plans and planning. This synergy of practical and formal approaches can stretch the formal methods to cover realistic plan representations as needed for real problem solving and can improve the analysis that is possible for practical planning systems.

We believe the <I-N-C-A> approach is valid in design and synthesis tasks more generally; we consider planning to be a limited type of design activity.

13.8 I-X Approach and I-X Process Panels

I-X and <I-N-C-A> provide a shared intelligible, easily communicated, and extendible conceptual model for objectives, processes, standard operating procedures and plans:

- I Issues
- N Nodes/activities
- C Constraints
- A Annotations

Intelligent activity planning, execution, monitoring, re-planning, and plan repair is supported via I-Plan and I-P2 (I-X Process Panels); I-P2's aim is a workflow and messaging "catchall" and can take *any* requirement to

- Handle an issue
- Perform an activity
- Respect a constraint
- Note an annotation

I-X deals with these via

- Manual activity
- Internal capabilities
- External capabilities
- Reroute or delegate to other panels or agents
- Plan and execute a composite of these capabilities

It receives reports and interprets them to

- Understand the current status of issues, activities, and constraints
- Understand the current world state, especially the status of process products
- Help the user control the situation

It maintains the current status, models and knowledge.

It copes with partial knowledge of processes and organizations.

It uses representation and reasoning together with the state to seek to present current, context-sensitive, options for action.

It supports a mixed-initiative collaboration model of "mutually constraining things."

13.9 Applications of I-X

- Disaster planning, evacuation operations, military operations in urban terrain, search and rescue
- Rapidly-deployed coalition operations support
- Help desk support
- Computer and systems configuration
- (Multilingual) maintenance procedures aid
- Unusual and emergency procedures assistant

13.10 Comparing the Intelligible Planning Approach to Studies of Expert Human Planners in Dynamic Situations

The features of the "Intelligible Planning" approaches used in O-Plan and I-Plan and described in Section 13.4 are similar to the approaches observed in expert human problem solvers performing in dynamic, stressful, or unusual situations. These observations were made in studies over many years by Klein (1998), and he contrasts these with some automated "black box" AI and algorithmic techniques.

There are different types of planning technology available from the AI community. This is not restricted to a simple kind of search from some known initial state to some final desired state seeking the best solution according to some predefined criteria. Gary Klein's book (1998) on how people make decisions in situations such as military operations, fire fighting, or other life-threatening environments provides a rich set of case studies to show that, in relatively few situations, were deliberative planning techniques in obvious use. People just seemed to be making the "right" choices—or a choice that worked, which was all that was required. They attributed their rapid selection of a suitable course of action to training, experience, or even ESP!

They stated, "I don't plan, I just know what to do." When options were deliberated over and evaluated, the situation for those involved was novel or unusual to their previous experience.

Klein's studies show how people in stressful environments select a course of action and adapt it as circumstances alter. Many of the decisions made by the subjects relate to issues that AI planning researchers are addressing. However, they are far removed from the traditional search style of deliberative plan generation. The AI techniques we describe to address planning in dynamic environments are of much wider utility than simple, fully automated search methods. The planning requirements seen in human problem solving in dynamic situations can be mapped to some of the AI representation and reasoning methods we are bringing to bear on practical planning problems.

- Overall management of the command, planning, and control process steps to improve coordination.
- Expansion of a high-level abstract plan into greater detail when necessary.
- High-level "chunks" of procedural knowledge (standard operating procedures, best practice processes, tactics, techniques, and procedures, etc.) at a human scale—typically five to eight actions—can be manipulated within the system (Klein, p. 52 and p. 58).
- Ability to establish that a feasible plan exists, perhaps for a range of assumptions about the situation, while retaining a high-level overview. (Klein, p. 227, "Include only the detail necessary to establish a plan is possible—do not fall into the trap of choreographing each of their movements").
- Analysis of potential interactions as plans are expanded or developed (Klein, p. 53).
- Identification of problems, flaws, and issues with the plan (Klein, p. 63 and p. 71).
- Deliberative establishment of a space of alternative options perhaps based on different assumptions about the situation involved, of special use ahead of time, in training and rehearsal, and to those unfamiliar with the situation or utilizing novel equipment (Klein, p. 23).
- Monitoring of the execution of events as they are expected to happen within the plan, watching for deviations that indicate a necessity to re-plan (often ahead of this becoming a serious problem) (Klein, pp. 32–33).
- AI planning techniques represent the dynamic state of the world at points in the plan and can be used for "mental simulation" of the execution of the plan (Klein, p. 45).
- Pruning of choices according to given requirements or constraints (Klein, p. 94, "singular strategy").
- Situation-dependent option filtering (sometimes reducing the choices normally open to one "obvious" one (Klein, pp. 17–18).
- Satisficing search to find the first suitable plan that meets the essential criteria (Klein, p. 20).
- Anytime algorithms that seek to improve on the best previous solution if time permits.
- Heuristic evaluation and prioritization of multiple possible choices within the constraint search space (Klein, p. 94).
- Repair of plans while respecting plan structure and intentions.
- Uniform use of a common plan representation with embedded rationale to improve plan quality, shared understanding, etc. (Klein, p. 275, seven types of information in a plan).

Gary Klein was asked to comment upon the "Intelligible Planning" approach in Tate (2000b, Appendix) and described in Section 13.4 as compared to his observations of natural problem solving and decision making in humans operating in stressful situations and dynamic environments. He observed the following in this edited personal communication (from Tate, 2000b, Appendix):

1. I felt a strong kinship with what you are attempting. The effort to use satisficing criteria, the use of anytime algorithms to permit continual improvement, the shift from abstract to detailed plan when necessary, the analysis of interactions in a plan, the identification of flaws in a plan, the monitoring of execution, the use of mental simulation, the representation of a singular strategy, heuristic evaluation, plan repair, and so forth are all consistent with what I think needs to be done.
2. My primary concern is how you are going to do these things … The discipline of AI can provide constraints that will help you understand any of these strategies in richer detail. But those constraints may also prevent you from harnessing these sources of power.
3. Your slogan, "Search and you're dead," seems right. Unconstrained search is a mark of intellectual cowardice. And it is also not a useful strategy.

Personal communication reply from Austin Tate to Gary Klein (from Tate, 2000b, Appendix):

I want to clarify my use of the slogan, "Search and you're dead," over the last 20 years. This is the headline, but I then clarify what I mean as "(Unconstrained) search and you're dead."

I have found this to be a useful slogan to express my general approach, and it makes for good knockabout fun on panels at conferences. The idea should be to richly describe the constraints known using whatever knowledge is available about the problem and then to seek solutions in that constrained space. We seek to use knowledge of the domain to constrain the use of blind search or "black box" automated methods in ways that are intelligent and intelligible (to humans).

In reality, all planning systems we build have sophisticated search and constraint management components, and it is an aim of our research to be able to utilize the best available in an appropriate context. Search can be a useful tactic in situations in which you are under-constrained and stuck. AI has made enormous advances in constraint management using search and other methods—so much so that some of its proponents argue that we do not need to bother with domain expertise or being knowledge-based about many of the problems we are addressing. It's this latter overenthusiasm for one approach that I seek to counter. Even very powerful search can be made more useful if put into a sensible knowledge-based context. This is, of course, more relevant when humans are involved in the decisions as then a more naturalistic style of mutually progressing toward a solution becomes a key to successful use of the technology.

13.11 Summary

HTN planning could be a useful paradigm to allow for agent operations that can dynamically adapt to a specific context and allow the following:

- Composition of workflows from requirements and component/template libraries
- Coverage of simple through to very complex (preplanned) components
- Execution support, reactive repair, recovery, etc.
- Mixed initiative (people and systems) planning and execution
- Provision of a framework within which more detailed specialized solvers, optimizers, and simulators work

I-X technology and its underlying <I-N-C-A> ontology to represent processes and plans that can act as a flexible, extendable, and intelligible framework to deploy such an approach.

Acknowledgments

The O-Plan and I-X projects were sponsored by the Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory Command and Control Directorate under various programs, including the Planning Initiative and Agent-Based Computing Programs, by the UK Defence Evaluation Research Agency (DERA), and by others. The University of Edinburgh and research sponsors are authorized to reproduce and distribute reprints for their purposes notwithstanding any copyright annotation hereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of other parties.

References

Allen, J., Hendler, J. and Tate, A. (1990) *Readings in Planning*, Morgan Kaufmann, San Mateo, CA.

Beck, H. (1993) TOSCA: A Novel Approach to the Management of Job-shop Scheduling Constraints, *Realising CIM's Industrial Potential: Proceedings of the Ninth CIM-Europe Annual Conference*, pp. 138–149, (eds. Kooij, C., MacConaill, P.A., and Bastos, J.).

Currie, K.W. and Tate, A. (1991) O-Plan: The Open Planning Architecture, *Artificial Intelligence* 52(1), Autumn 1991, North-Holland.

Dalton, J., Drabble, B. and Tate, A. (1993) The O-Plan Constraint Associator, *Proceedings of the 13th UK Planning and Scheduling SIG*, Strathclyde University, UK, 14–15 September 1994.

Klein, G. (1998) *Sources of Power—How People Make Decisions*, 1st Edition, MIT Press, 1998.

Nau, D.S., Au, T.-C., Ilghami, O., Kuter, I., Muñoz-Avila, H., Murdock, J.W., Wu, D. and Yaman, F. (2005) Applications of SHOP and SHOP2, *IEEE Intelligent Systems* 20(2) pp. 34–41, Mar.–Apr. 2005.

Pollack, M. (1994) *DIPART Architecture*, Technical Report, Department of Computer Science, University of Pittsburgh, PA 15213, USA.

Sacerdoti, E. (1977) A Structure for Plans and Behaviours. *Artificial Intelligence* series, publ. North Holland.

Smith, S. (1994) OPIS: A Methodology and Architecture for Reactive Scheduling, in *Intelligent Scheduling*, (eds., Zweben, M. and Fox, M.S.), Morgan Kaufmann, Palo Alto, CA., USA.

Tate, A. (1977) Generating Project Networks, *Proceedings of the Fifth International Joint Conference on Artificial Intelligence (IJCAI-77)*, pp. 888-893, Boston, Mass. USA, August 1977. Reprinted in “Readings in Planning” (eds. Allen, J., Hendler, J. and Tate, A.), Morgan-Kaufmann, 1990.

Tate, A. (1994) Mixed Initiative Planning in O-Plan2, *Proceedings of the ARPA/Rome Laboratory Planning Initiative Workshop*, (ed. Burstein, M.), Tucson, Arizona, USA, Morgan Kaufmann, Palo Alto.

Tate, A. (ed.) (1996a) *Advanced Planning Technology - Technological Achievements of the ARPA/Rome Laboratory Planning Initiative (ARPI)*, AAAI Press.

Tate, A. (1996b) Representing Plans as a Set of Constraints—the <I-N-OVA> Model, *Proceedings of the Third International Conference on Artificial Intelligence Planning Systems (AIPS-96)*, pp. 221–228, (Drabble, B., ed.) Edinburgh, Scotland, AAAI Press.

Tate, A. (1998) Roots of SPAR—Shared Planning and Activity Representation, *Knowledge Engineering Review*, Vol. 13, No. 1, March 1998. See also http://www.aiai.ed.ac.uk/project/spar/

Tate, A. (2000a) *<I-N-OVA> and <I-N-CA>—Representing Plans and other Synthesized Artifacts as a Set of Constraints*, AAAI-2000 Workshop on Representational Issues for Real-World Planning Systems, at the National Conference of the American Association of Artificial Intelligence (AAAI-2000), Austin, Texas, USA, August 2000.

Tate, A. (2000b) Intelligible AI Planning, in Research and Development in Intelligent Systems XVII, *Proceedings of ES2000, The Twentieth British Computer Society Special Group on Expert Systems International Conference on Knowledge Based Systems and Applied Artificial Intelligence*, pp. 3–16, Cambridge, UK, December 2000, Springer.

Tate, A. (2003) <I-N-C-A>: A Shared Model for Mixed-initiative Synthesis Tasks, *Proceedings of the Workshop on Mixed-Initiative Intelligent Systems (MIIS) at the International Joint Conference on Artificial Intelligence (IJCAI-03)*, pp. 125–130, Acapulco, Mexico, August 2003.

Tate, A. and Dalton, J. (2003) O-Plan: A Common Lisp Planning Web Service, *Proceedings of the International Lisp Conference 2003*, October 12–25, 2003, New York, NY, USA, October 12–15, 2003.

Tate, A., Dalton, J. and Stader, J. (2002) I P2—Intelligent Process Panels to Support Coalition Operations, *Proceedings of the Second International Conference on Knowledge Systems for Coalition Operations (KSCO-2002)*, pp. 184–190, Toulouse, France, 23–24 April 2002.

Tate, A., Drabble, B. and Dalton, J. (1994a) Reasoning with Constraints within O-Plan2, *Proceedings of the ARPA/Rome Laboratory Planning Initiative Workshop*, (ed. Burstein, M.), Tucson, AZ, USA, Morgan Kaufmann, Palo Alto, CA, USA.
Tate, A., Drabble, B. and Kirby, R. (1994b) O-Plan2: An Open Architecture for Command, Planning and Control, in *Intelligent Scheduling*, (eds., Zweben, M. and Fox, M.S.), Morgan Kaufmann, Palo Alto, CA., USA.
Tate, A., Levine, J., Dalton, J. and Nixon, A. (2003) Task Achieving Agents on the World Wide Web, in *Spinning the Semantic Web*, Fensel, D., Hendler, J., Liebermann, H. and Wahlster, W. (eds.), Chapter 15, pp. 431–458, MIT Press, 2003.
Wilkins, D. (1988) *Practical Planning*, Morgan Kaufmann, Palo Alto, 1988.