

Generation of Multiple Qualitatively Different Plan Options

Austin Tate, Jeff Dalton and John Levine

Artificial Intelligence Applications Institute
The University of Edinburgh
80 South Bridge, Edinburgh EH1 1HN, UK
{a.tate, j.dalton, j.levine}@ed.ac.uk
<http://www.aiai.ed.ac.uk/~oplan>

Abstract

Work is described which supports multi-agent mixed initiative interaction between a “task assignment” or “command” agent and a planning agent. Each agent maintains an agenda of outstanding tasks it is engaged in and uses a common representation of tasks, plans, processes and activities based on the notion that these are all “constraints on behaviour”. Interaction between the agents uses an exchange of explicit task and option management information. This framework forms the basis for a Web-based demonstration of mixed initiative user/system agents working together to mutually constrain task descriptions and alternative plans and to coordinate the task-oriented generation, refinement and enactment of those plans.

Introduction

Under the University of Edinburgh O-Plan Project (Currie and Tate, 1991; Tate, Drabble and Kirby, 1994), which is part of the DARPA/Rome Laboratory Planning Initiative – ARPI (Fowler *et al.*, 1996; Tate, 1996a), we are exploring mixed initiative planning methods and their application to realistic problems in logistics, air campaign planning and crisis action response (Tate, Drabble and Dalton, 1996). In preparatory work, O-Plan has been demonstrated operating in a range of mixed initiative modes on a Non-Combatant Evacuation Operation (NEO) problem (Tate, 1994; Drabble, Tate and Dalton, 1995). A number of “user roles” were identified to help clarify some of the types of interaction involved and to assist in the provision of suitable support to the various roles (Tate, 1994) – see Figure 1.

The overall concept for our demonstrations of O-Plan acting in a mixed initiative multi-agent environment is to have humans and systems working together in given roles to notionally populate a Course of Action (COA) / Elements of Evaluation comparison matrix. The columns of this matrix are alternative options being explored as a potential solution to a (possibly underspecified) problem and the rows are evaluations of

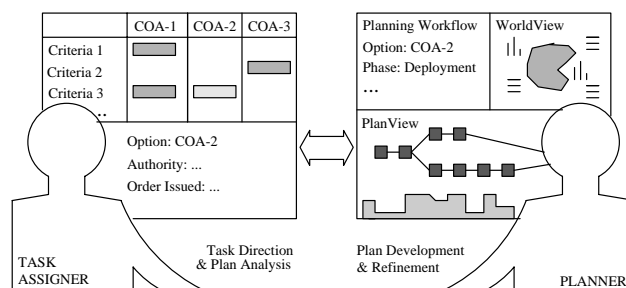


Figure 1: Communication between the Task Assigner and the Planner

the solution being considered. The idea is that the requirements, assumptions and constraints are all refined concurrently using the elements of evaluation (EEs).

We are exploring the links between key user roles in the planning process and automated planning support aids (Tate, 1997). Research is exploring a planning workflow control model using:

- a shared model of mixed initiative planning as “mutually constraining the space of behaviour”;
- the <I-N-OVA> constraint model of activity as the basis for plan communication;
- explicit management between agents of the tasks and options being considered;
- agent agendas and agenda issue handlers;
- explicit provision of authority for an agent to perform its functions.

Agents maintain their own perspective of their part in the task to hand, while cooperating with other agents who may perform parts of the task.

As shown in Figure 1, we envisage two human agents, called the Task Assigner and the Planner, working together to explore possible solutions to a

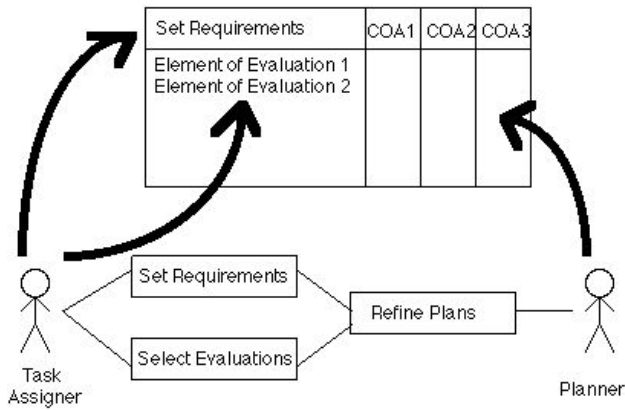


Figure 2: Roles of the Task Assigner and the Planner

problem and making use of automated planning aids to do this. Figure 2 shows how the two human agents work together to populate the COA comparison matrix. The Task Assigner sets the requirements for a particular Course of Action (i.e. what top level tasks must be performed) and selects appropriate evaluation criteria (elements of evaluation) for the resulting plans. The Planner agent acts to refine the resulting plans by adding further constraints and splitting plans to explore two or more possible options for the same COA requirements.

In this paper, we describe our current Web-based demonstration of these ideas, together with the background to this demonstration. We start with the generic systems architecture being used and the architecture of the O-Plan system being used as a *plan server*. We then describe our idea of mixed initiative planning as multiple agents mutually constraining the space of behaviour. The current Web-based demonstration of our ideas is then presented, followed by a summary and future directions.

Generic Systems Integration Architecture

The O-Plan agent architecture to be described in the next section is a specific variant of a generalised systems integration architecture shown in Figure 3. This general structure has been adopted on a number of AIAI projects (Fraser and Tate, 1995). The architecture is an example of a *Model/Viewer/Controller* arrangement. The components are as follows:

Viewers: user interface, visualisation and presentation viewers for the model.

Task and Option Management: the capability to support user tasks via appropriate use of the pro-




Figure 3: Generic Systems Integration Architecture

cessing and information assets and to assist the user in managing options being used within the model. This is sometimes referred to as the *Controller*.

Model Management: coordination of the capabilities/assets to represent, store, retrieve, merge, translate, compare, correct, analyse, synthesise and modify models.

Mediators: Intermediaries or converters between the features of the model and the interfaces of active components of the architecture.

Processing Assets: functional components (model analysis, synthesis or modification).

Constraint Managers: components which assist in the maintenance of the consistency of the model.

Information Assets: information storage and retrieval components.

O-Plan – the Open Planning Architecture

This section describes the O-Plan architecture and the structure of individual O-Plan agents. The components of a single O-Plan agent are shown in Figure 4.




Figure 4: O-Plan Agent Architecture

Task and Option Management

Task and option management facilities are provided by the *Controller* in O-Plan. The O-Plan Controller takes its tasks from an agenda which indicates the outstanding processing required and handles these with its *Knowledge Sources*.

O-Plan has explicit facilities for managing a number of different options which it is considering. O-Plan has an agent level agenda, and agendas which relate to each option it is considering (in fact these are part of the plan representation for these options – the *issues* or 1 part of <I-N-OVA>). Many of these options are internal to the planning agent, and are generated during the search for a solution. Others are important for the interaction between the planner and a user acting as a task assigner.

Abstract Model of Planning Workflow – Plan Modification Operators

A general approach to designing AI-based planning and scheduling systems based on partial plan or partial schedule representations is to have an architecture in which a plan or schedule is critiqued to produce a list of issues or agenda entries which is then used to drive a workflow-style processing cycle of choosing a “plan modification operator” (PMO) to handle one or more agenda issues and then executing the PMO to modify the plan state. Figure 5 shows this graphically.

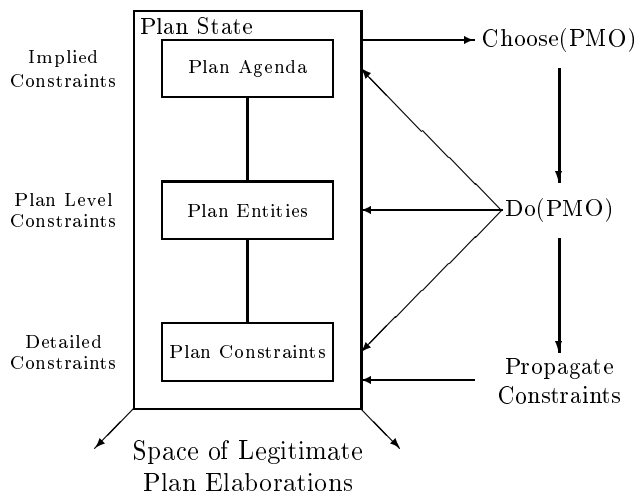


Figure 5: Planning Workflow - Using PMOs to Handle Agenda Issues

This approach is taken in O-Plan. The approach fits well with the concept of treating plans as a set of constraints which can be refined as planning progresses. Some such systems can act in a non-monotonic fashion by relaxing constraints in certain ways. Having the implied constraints or “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 itself.

Representing Plans as a Set of Constraints on Behaviour

The <I-N-OVA> (*Issues – Nodes – Orderings / Variables / Auxiliary*) Model is a means to represent and manipulate plans as a set of constraints.

In Tate (1996b), the <I-N-OVA> model is used to characterise the plan representation used within O-Plan and is related to the plan refinement planning method used in O-Plan. A plan is represented as a set of constraints which together limit the behaviour that is desired when the plan is executed. The set of constraints are of three principal types with a number of sub-types reflecting practical experience in a number of planning systems.

Plan Constraints

- I - Issues (Implied Constraints)
- N - Node Constraints (on Activities)
- OVA - Detailed Constraints
 - O - Ordering Constraints
 - V - Variable Constraints
 - A - Auxiliary Constraints
 - Authority Constraints
 - Condition Constraints
 - Resource Constraints
 - Spatial Constraints
 - Miscellaneous Constraints

Figure 6: <I-N-OVA> Constraint Model of Activity

The node constraints (these are often of the form “include activity”) in the <I-N-OVA> model create the space within which a plan may be further constrained. The I (issues) and OVA constraints restrict the plans within that space to those which are valid. Ordering (temporal) and variable constraints are distinguished from all other auxiliary constraints since these act as *cross-constraints*, usually being involved in describing the others – such as in a resource constraint which will often refer to plan objects/variables and to time points or ranges.

The <I-N-OVA> constraint model of activity allows planning process state as well as the current state of the plan generated to be communicated between agents involved in the planning process. This is done via the Issues part of <I-N-OVA> – which can be used to amend the task and option specific agenda which a planning agent is using for its problem solving.

Authority to Plan

As described in Tate (1993) it is intended that O-Plan will support authority management in a comprehensive and principled way. *Changes* of authority are pos-

sible via Task Assignment agent communication to the Planner agent. This may be in the context of a current plan option and task provided previously or it is possible to give defaults which apply to all future processing by the planner agent. The authorities may use domain related names that are meaningful to the user and may refer to the options, sub-options, phases and levels of tasks and plans known to O-Plan.

Mutually Constraining Plans for Mixed Initiative Planning and Control

Our approach to Mixed Initiative Planning in O-Plan assists in the coordination of planning with user interaction by employing a shared model of the plan as a set of constraints at various levels that can be jointly and explicitly discussed between and manipulated by any user or system component in a cooperative fashion.

The model of Mixed Initiative Planning that can be supported by the approach is *the mutual constraining of behaviour* by refining a set of alternative partial plans. Users and systems can work in harmony though employing a common view of their roles as being to constrain the space of admitted behaviour. Further detail is given in Tate (1994).

Workflow ordering and priorities can be applied to impose specific styles of authority to plan within the system. One extreme of user driven plan expansion followed by system “filling-in” of details, or the opposite extreme of fully automatic system driven planning (with perhaps occasional appeals to an user to take predefined decisions) are possible. In more practical use, we envisage a mixed initiative form of interaction in which users and system components proceed by mutually constraining the plan using their own areas of strength.

Coordination of problem solving must take place between users and the automated components of a planning system. In joint research with the University of Rochester (whose work is described in Allen, Ferguson and Schubert, 1996) we are exploring ways in which the O-Plan controller can be given specific limitations on what plan modifications it can perform, and the specific plan options or sub-options it is working on can be coordinated with those being explored by a user supported by a suitable interface.

A Web-based Demonstration

This section describes our current implementation of these ideas. We have constructed a Web-based demonstration of a task assignment agent working with the O-Plan planning agent to populate and explore different options within a course-of-action matrix. We are using a general-purpose logistics and crisis operations

domain which is an extension of our earlier logistics-related domains (Reece *et al.*, 1993).

This demonstration is a significant milestone on the path towards our stated vision, since it contains many of the elements which have been planned for over the last 3 to 4 years of work and which have been incorporated into O-Plan Version 3.1 since its release in January 1997. These include:

- Multiple option management: exploration of separate options and sub-options.
- Multiple initial conditions: exploration of different initial assumptions about the domain.
- Incremental tasking: adding further requirement constraints to a plan after an initial phase of planning.
- Authority to plan: authorities can be set for any COA investigated allowing for incremental plan refinement alongside user directed addition of planning constraints.
- Plan analysis: facilities for plan analysis/evaluation can be installed which have both brief and longer analysis results to present to the user.
- Evaluation selection: the evaluations presented can be selected to show the ones which are critical.
- Issue maintenance: planning or plan analysis can leave outstanding issues to be addressed, which are summarised and collected to help with planning and coordination workflow.
- Status indication: coloured “traffic lights” are used, as in other ARPI plan visualisation work (Stillman and Bonissone, 1996) to indicate that a chosen plan for a COA is complete (green), has warnings or notes to read (orange) or have issues that need attention (red).

The Web demonstration, Version 3.1 of the O-Plan code and the papers referenced here are available by following links from the O-Plan home page.

The COA Comparison Matrix

The user is initially given a blank COA comparison matrix which is populated by the user and O-Plan during the course of a session (Figure 7). The user acts in the role of the Task Assigner agent, setting the initial assumptions and tasking level requirements for a Course of Action (Figure 8) and selecting elements of evaluation to include in the matrix. The task assigner can split any COA into two or more sub-options and explore further within each. Additional constraints (in

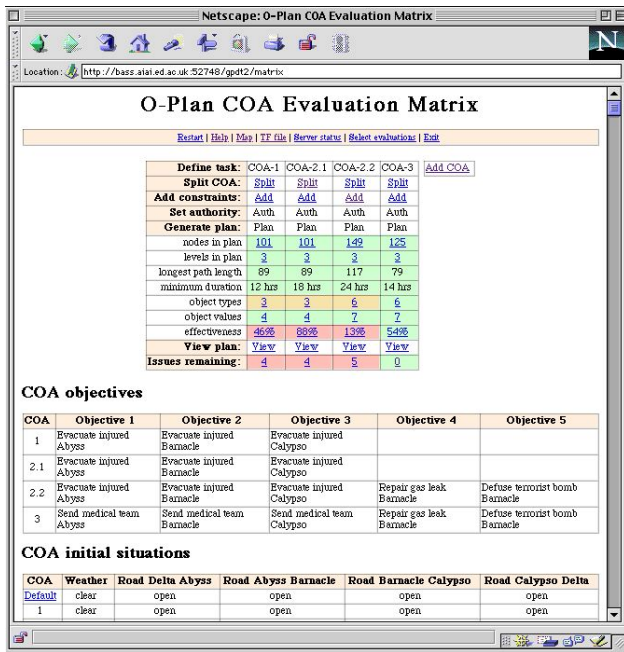


Figure 7: The Course of Action Evaluation Matrix

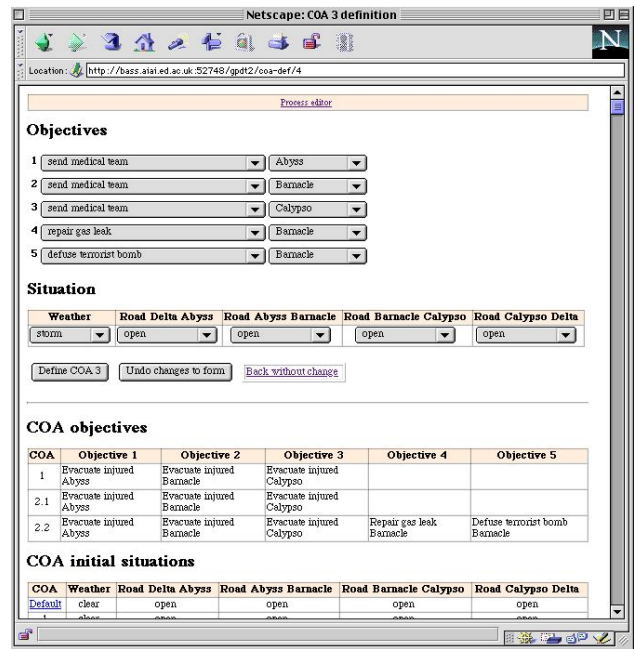


Figure 8: Defining the Requirements for a Course of Action

the form of task level requirements) can be added to any COA. The task assigner can also authorise O-Plan only to plan to a nominated level of detail. Together, these facilities allow for incremental development, exploration and evaluation of multiple qualitatively different plan options.

The COA matrix is an abstract underlying notion and may not appear in a user interface for a completed system. However, it is useful in this demonstration to show our ideas about what is being created and refined as mixed initiative problem solving takes place. In a dialogue system, such as TRAINS (Ferguson, Allen and Miller, 1996), the COA matrix would be the underlying model of the problem solving and the dialogue model would then implicitly refer to this artefact.

“Go Places and Do Things” – The Crisis Operations Domain

We have used a crisis operations domain based on the Pacifica scenarios (Reece *et al.*, 1993; Tate, Drabble and Dalton, 1996) that we call “Go Places and Do Things” (GPDT). This is a three level domain model which closely follows what we observe in large real domain models. The top level is mostly about setting objectives (i.e. COA requirements). The second level is the real planning level and where technological interactions, such as allocating limited resources, need to be resolved. The third level is needed to add detail to the skeleton plans that have been selected.

This domain is a natural extension of our earlier work in the Pacifica Non-combative Evacuation Operations (NEO) domain. In the earlier work, people are evacuated (following some crisis) from a small island using trucks and helicopters. In the new domain, the main goal is to avert a developing crisis in one of the cities on the island, using various vehicles, pieces of equipment and specialist teams. In the crisis domain, unlike previous Pacifica scenarios, the tasks to be performed are complex and may involve plans consisting of hundreds of individual actions.

This domain has been chosen for our current work to demonstrate that O-Plan is sufficiently powerful to be able to cope with these complicated logistical problems and also to provide the O-Plan team with a problem domain which is general enough to allow expansion and experimentation as our ideas and technology develop.

The Scenario The action takes place somewhere in a network of cities, currently on the island of Pacifica (see Figure 9). A number of crisis situations can arise in the cities and on the roads joining them, such as power stations becoming inoperative or people needing medical treatment. The goal of the commander (i.e. the Task Assigner agent) is to respond effectively to the situation so that the immediate crisis situation is dealt with and appropriate repairs are made to restore the status quo.




Figure 9: The Island of Pacifica

World Description The following types of objects exist in this domain:

Cities: these can contain other objects, such as teams, people and equipment.

Roads: these connect some of the cities. They may become blocked to certain classes of vehicle due to weather conditions or landslides. Some may be permanently blocked to certain classes of vehicle (e.g. mud tracks).

Vehicles: these are used to carry equipment, teams and people between cities. There are various types of vehicle which have very different capabilities, such as fast air vehicles of low carrying capacity and slow ground transports capable of carrying large pieces of equipment.

Equipment: there are various pieces of specialist equipment located in the network of cities. These are needed to perform certain tasks, such as repairs at a power station or emergency medical treatment.

Teams: there are also various specialist teams of people located in the cities. These teams perform specialist tasks, such as fast evacuation or building emergency housing.

People: people are located at cities and may need medical treatment or evacuation. As a simplification, we treat people as a single entity to be treated or moved around, rather than counting a specific number.

Weather: the weather may restrict the options available to the planner, such as not allowing helicopters to fly in thunderstorms.

The world state can be described by giving the locations and contents of the vehicles, the locations of the people, teams and pieces of equipment, and the status of the roads, people and weather.

Actions and Plans In this domain, the teams, equipment and people can be moved around using a TRANSPORT action at modelling Level 2:

```
TRANSPORT cargo ITEM using VEHICLE from
CITY to CITY
where ITEM is an object of type team, vehicle,
equipment or people.
```

The result of the action is that the cargo moves from the source to the destination.

Other actions in the domain are dependent on the specific example chosen, but will typically contain around 5 actions at a lower level of detail. Typical examples are:

- Repair a turbine at a crucial power station.
- Give emergency medical treatment to people exposed to toxic fumes.
- Repair a bridge which has been broken in a storm.
- Build emergency housing for refugees.
- Perform emergency operations to make the area safe for a repair team.
- Evacuate the population of one of the cities.

An entire plan will consist of a number of TRANSPORT operations to bring the necessary teams and equipment together, followed by the main tasks. The TRANSPORT operations and main tasks may overlap, as in our demonstration example which follows.

Implementation Status The current O-Plan Task Formalism (TF) file for this domain implements the crisis operations domain for the island of Pacifica, using 12 top level tasks and four cities (Abyss, Barnacle, Calypso and Delta). A Course of Action consisting of 5 tasks at the top level expands to give approximately 30 actions at the second level and 150 tasks at the third level. The exact numbers will depend on the particular Level 1 tasks selected for the Course of Action.

The Demonstration Scenario

The following scenario illustrates how we envisage the system being used and can be used in actual demonstrations of this work.

The task assigner (TA) is told that there are injured people in Abyss, Barnacle and Calypso and that these people need to be treated within the next 18 hours in order to avoid fatalities. The latest weather forecast shows a 50% chance of a storm over Pacifica during the next 24 hours.

The TA decides to try evacuating the injured from all three cities as the first possible plan, using the assumption that the weather is clear. The evaluation criteria are fine and the plan executes within the required deadline. This illustrates how the TA sets up tasks and assumptions within COAs and how the interface displays the elements of evaluation in the matrix.

The TA wants to check that the plan is still OK if the predicted storm occurs. A further COA is added with the tasks being set up as before. This time, the TA sets the weather to “storm”. O-Plan is asked to generate a plan for this new set of COA requirements and finds that the time taken to execute is 18 hours – just on the deadline. This illustrates the basic use of COA columns to compare different courses of action based on different initial assumptions.

However, the TA is now interrupted by a call from the Barnacle field station. Reports are coming in of an explosion at the main Barnacle power station, causing a gas leak. It is thought that this may have been caused by a terrorist bomb. It seems wise to fix the gas leak and send a bomb squad to deal with any other bombs that may have been planted. Meanwhile, the latest weather bulletin indicates that a storm is brewing to the north-east and has a 95% chance of hitting the island within the next 5 hours.

To deal with these turns of events, the TA now splits COA-2 (the realistic weather assumption) into two sub-options and adds two new tasks to one of them – COA-2.2. The new tasks are to repair the gas leak at Barnacle and to defuse other (potential) terrorist bombs at Barnacle. This illustrates the use of plan splitting and addition of new tasks. Unfortunately for the TA, the new plan takes 24 hours, which is 6 hours over the deadline.

The TA now needs to think. The stormy weather prediction has become more definite, so the TA sets the default weather assumption to be “storm”. Then a further COA column is added (COA-3). Since the original task was to simply to treat the injured people at the three cities, evacuation is perhaps an unnecessary luxury. The TA therefore sets up COA-3 to send medical teams to the three cities, repair the gas leak

and defuse the terrorist bomb at Barnacle. Since the default for the weather is “storm”, the TA does not need to note this explicitly. The resulting plan completes within 14 hours, so this new plan seems like the best one so far. The “traffic light” indicators in the matrix show various warnings, mostly concerned with using all available resources of a certain type within the plan. The TA marks all of these as being acceptable and the traffic lights in the column for COA-3 turn green, indicating that the plan is ready to execute.

As a final optimisation, the TA adds another column (COA-4) and sets this up as for COA-3, but with the injured being evacuated from Barnacle rather than a medical team sent (because of the additional danger in Barnacle due to the gas leak and/or terrorist bombs). This plan executes within 17 hours, which is 1 hour less than the deadline.

Future Work

The current demonstration still has some limitations, and we plan to address these in our final project demonstration (due in June 1998). The most important item to be addressed is to add the human planner agent into the demonstration, with the task assigner, planner and O-Plan agent all acting together to explore the plan space in a true mixed initiative interaction. This will require that new facilities be added to support the human planner agent and that communication between agents be provided via Web interaction and teleconferencing. We envisage that the planner agent and the task assigner will have different interface views onto the COA matrix, as illustrated in Figure 1. We also intend to improve the treatment of the crisis operations domain and allow plans to be specified, visualised and refined via a graphical Java-based process editor and plan viewer.

Summary

Five concepts are being used as the basis for exploring multi-agent and mixed-initiative planning involving users and systems: Together these provide for a *shared* model of what each agent can and is authorised to do and what those agents can act upon.

1. *Shared Plan Model* – a rich plan representation using a common constraint model of activity (<I-N-OVA>).
2. *Shared Task Model* – Mixed initiative model of “mutually constraining the space of behaviour”.
3. *Shared Space of Options* – explicit option management.

4. *Shared Model of Agent Processing* – handlers for issues, functional capabilities and constraint managers.
5. *Shared Understanding of Authority* – management of the authority to plan (to handle issues) and which may take into account options, phases and levels.

Using these shared views of the roles and function of various users and systems involved in a command, planning and control environment, we have demonstrated a planning agent being used to support mixed initiative task specification and plan refinement over the world wide web. It has been applied to the generation of multiple qualitatively different courses of action based on emerging requirements and assumptions. The demonstration takes place in a realistic crisis management domain.

Acknowledgements

The O-Plan project is sponsored by the Defense Advanced Research Projects Agency (DARPA) and Rome Laboratory, Air Force Materiel Command, USAF, under grant number F30602-95-1-0022. The O-Plan project is monitored by Dr. Northrup Fowler III at the US Air Force Research Laboratory, Rome. The US Government is authorised to reproduce and distribute reprints for Governmental 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 express or implied, of DARPA, AFRL or the US Government.

References

Allen, J.F., Ferguson, G.M. and Schubert, L.K. (1996). Planning in Complex Worlds via Mixed-Initiative Interaction. In *Advanced Planning Technology*, 53–60, (Tate, A., ed.), AAAI Press.

Currie, K.W. and Tate, A. (1991). O-Plan: the Open Planning Architecture. *Artificial Intelligence*, 51(1), Autumn 1991, North-Holland.

Drabble, B., Tate, A. and Dalton, J. (1995). Applying O-Plan to the NEO Scenarios, in *An Engineer's Approach to the Application of Knowledge-based Planning and Scheduling Techniques to Logistics*. Appendix O, USAF Rome Laboratory Technical Report RL-TR-95-235, December 1995.

Ferguson, G.M., Allen, J.F. and Miller, B.W. (1996). TRAINS-95: Towards a Mixed-Initiative Planning Assistant. *Proceedings of the Third International Conference on AI Planning Systems (AIPS-96)*, 70–77, (Drabble, B., ed.), AAAI Press.

Fowler, N., Garvey, T.D., Cross, S.E., and Hoffman, M. (1996). Overview: ARPA-Rome Laboratory Knowledge-Based Planning and Scheduling Initiative (ARPI). In *Advanced Planning Technology*, 3–9, (Tate, A., ed.), AAAI Press.

Fraser, J. and Tate, A. (1995). The Enterprise Tool Set – An Open Enterprise Architecture. *Proceedings of the Workshop on Intelligent Manufacturing Systems, International Joint Conference on Artificial Intelligence (IJCAI-95)*, Montreal, Canada, August 1995.

Reece, G.A., Tate, A., Brown, D. and Hoffman, M. (1993). The PRECIS Environment. Paper presented at the ARPA-RL Planning Initiative Workshop at AAAI-93, Washington D.C., July 1993.

Stillman J. and Bonissone, P. (1996). Technology Development in the ARPA/RL Planning Initiative. In *Advanced Planning Technology*, 10-23, (Tate, A., ed.), AAAI Press.

Tate, A. (1993). Authority Management – Coordination between Planning, Scheduling and Control. *Workshop on Knowledge-based Production Planning, Scheduling and Control at the International Joint Conference on Artificial Intelligence (IJCAI-93)*, Chambery, France, 1993.

Tate, A. (1994). Mixed Initiative Planning in O-Plan2. *Proceedings of the ARPA/Rome Laboratory Planning Initiative Workshop*, 512–516, (Burststein, M., ed.), Tucson, Arizona, USA, Morgan Kaufmann.

Tate, A. (1996a) (ed.). *Advanced Planning Technology*. 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)*, 221–228, (Drabble, B., ed.) Edinburgh, Scotland, AAAI Press.

Tate, A. (1997). Mixed Initiative Interaction in O-Plan. *Proceedings of AAAI Spring 1997 Symposium on Computational Models for Mixed Initiative Interaction*, Stanford University, March 1997.

Tate, A., Drabble, B. and Kirby, R. (1994). O-Plan2: an Open Architecture for Command, Planning and Control. In *Intelligent Scheduling*, (eds, M.Zweben and M.S.Fox), Morgan Kaufmann.

Tate, A., Drabble, B. and Dalton, J. (1996). O-Plan: a Knowledge-Based Planner and its Application to Logistics. In *Advanced Planning Technology*, 259–266, (Tate, A., ed.), AAAI Press.