O-Plan: A Situated Planning Agent [1]

Brian Drabble & Austin Tate

AI Applications Institute
University of Edinburgh
80 South Bridge
Edinburgh EH1 1HN, UK
FAX: 44 131 650 6513

B.Drabble@ed.ac.uk & *A.Tate@ed.ac.uk*

Abstract

This paper describes the need to view a planner as situated in an environment dealing with the whole "planning" problem. While a planning agent deals with plan generation aspects, other agents are concerned with aspects such as task elicitation, plan analysis, reactive execution, plan repair, etc. Each of these systems has its own perspective on the planning problem and each of the systems must be capable of communicating in a way which allows other systems to assimilate new information into their perspective of the problem. Within such a collection of agents a situated planner takes task assignments from a superior agent and creates a plan or further elaborates it before passing it to the execution support agents for further processing or enactment. The paper describes the O-Plan system which has been developed as an architecture within which situated agents, such as planning agents, can be created. The paper provides a summary of work to date on the planning and execution agents. The paper then goes on to describe current research involving O-Plan which aims to address the communication between a task assignment agent and a planning agent. It concentrates on three key issues in this area: communication of plans, assessment of the quality of plans and the role of authority in the planning process.

Keywords: Planning, Task Assignmentt, Plan Analysis, Plan Communication

[1] O-Plan work is supported by the US Advanced Research Projects Agency (ARPA) and the US Air Force Rome Laboratory acting through the Air Force Office of Scientific Research (AFSC) under contract F49620-92-C-0042. The view and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of ARPA or the U.S. Government.

1 Introduction

The aim of this paper is to describe the need for AI planners to be viewed as situated agents, where planning is one of a number of tasks involved in dealing with the whole problem of task assignment, planning, execution and control. While the planner deals with the plan generation aspect of the problem other agents deal with problems such as task elicitation, plan analysis, reactive execution, plan repair, etc. Each of these systems has its own perspective on the planning problem and each of the systems is capable of communicating in a way which allows other systems to assimilate new information into their perspective of the problem. Within such a collection of agents a situated planner takes task assignments from a superior agent and creates a plan or further elaborates it before passing it to the execution support agents for further processing or enactment.

The reason for taking this view is that planners cannot be considered as functioning in isolation. In addition to being able to communicate about the overall task being performed, the planner must be able to interact closely with the environment in which it is placed.

In many domains such as manufacturing, construction assembly, logistics, spacecraft control, etc. the planner needs to deal with a changes occurring in two very different ways:

1. **Change of Task and Requirements**:
   The task set to the planner may change as the plan is being generated requiring the planner to:
   - alter its focus, '`plan to move the 82nd airborne now rather than later`',
   - choose alternative methods, e.g. '`move the 82nd by sea rather than air`',
   - abandon the task altogether, '`abandon the deployment task and return the 82nd airborne to their home base`'.
2. **Change in the Environment**:
   Events may occur in the domain which require the plan to be repaired by the insertion of new constraints or activities. In some cases the failure may be so severe that the entire plan needs to be abandoned and an alternative chosen.

Generating plans and reacting to changes in the environment have been the primary focus to date for the O-Plan project [4, 16] as well as a many other researchers. However, the problem of dealing with task assignment and its link to the generative planner has been neglected by planning researchers. In many domains the problems of command, task setting, planning, plan analysis and enactment have been compartmentalised leading to many systems having an inability to assimilate new information into existing plan options. Current research on the O-Plan project aims to address this area and in particular the problem of allowing different situated agents to maintain their own perspective on the planning problem while at the same time allowing plans to be communicated between them. This will make it possible to communicate and use commands, plans, and tasks with improved precision, timeliness and level of detail between a number of situated agents. Research has already begun on addressing three key issue of the task assignment problem:

- communication of plans between agents,
- assessment of the quality of plans being generated by an agent,
- the role of authority in the task assignment process.

Each of the these key issues is dealt with in a separate section.

The structure of the paper is as follows. Section 2 provides an overview of the O-Plan system and the areas of AI planning research which it has covered to date. Section 3 describes the new focus of the O-Plan project and provides details of the research already undertaken in the area of Task Assignment. Section 4 provides a summary of the paper and outlines future work in the area of Task Assignment.

2 O-Plan - the Open Planning Architecture

O-Plan is a command, planning and control architecture which has an open modular structure intended to allow experimentation on or replacement of various components. The research is seeking to isolate functionality that may be generally required in a number of applications and across a number of different command, planning, scheduling and control systems. O-Plan has been applied to logistics tasks in which flexibility of response to changing situations is required.

O-Plan is intended to be a domain-independent, general planning and control framework with the ability to make use of detailed knowledge of the domain. See [2] for background reading on AI planning systems. See [4] for details of the first version of the O-Plan planner which introduced an agenda-based architecture and the main system components. That paper also includes a chart showing how O-Plan relates to other planning systems. The second version of the O-Plan system adopted a multi-agent approach and situated the planner in a task requirement and plan execution setting. The multi-agent approach taken is described in greater detail in [16].

Figure 1 shows the communications between the 3 agents in the O-Plan architecture. A user specifies a task that is to be performed through some suitable interface. We call this process *task assignment*. A *planner* plans to perform the task specified. The *execution system* seeks to carry out the detailed actions specified by the planner while working with a more detailed model of the execution environment.

The O-Plan approach to command, planning, scheduling and control can be characterised as follows:

- successive refinement/repair of a complete but flawed plan or schedule
- least commitment approach
- using opportunistic selection of the focus of attention on each problem solving cycle
- building information incrementally in "constraint managers", e.g.,

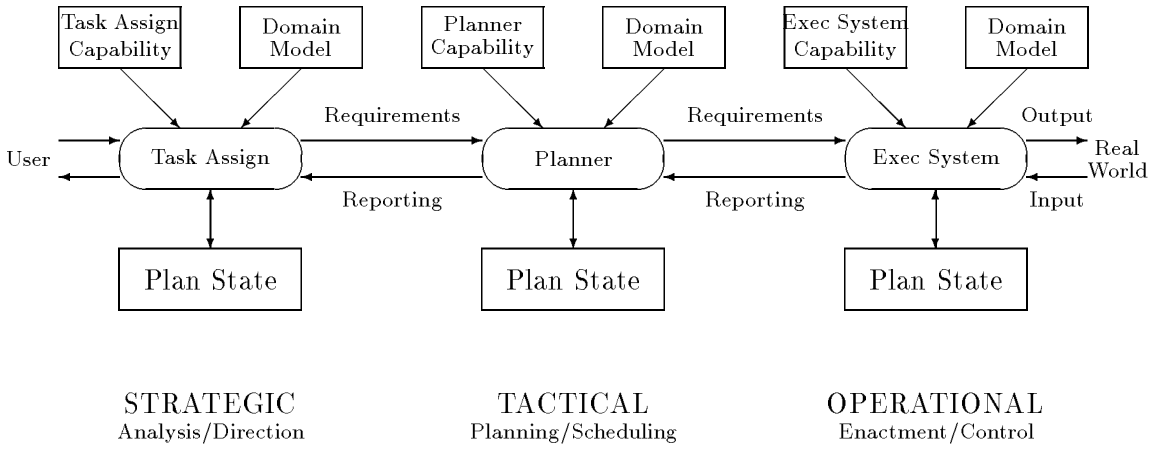

Figure 1: Communication between Strategic, Tactical and Operational Agents

  - object/variable manager
  - time point network manager
  - effect/condition manager
  - resource utilisation manager

- using localised search to explore alternatives where advisable
- with global alternative re-orientation where necessary.

The O-Plan project has sought to identify modular components within an AI command, planning and control system and to provide clearly defined interfaces to these components and modules. The background to this work is provided in [13]. The various components plug into "sockets" within the architectural framework. The sockets are specialised to ease the integration of particular types of component. See Figure 2.

The various components of the architecture for each O-Plan agent are:

PlanWorld Viewers – User interface, visualisation and presentation viewers for the plan – usually differentiated into technical *plan* views (charts, structure diagrams, etc.) and *world* views (simulations, animations, etc.).

Knowledge Sources – Functional components which can analyse, synthesise or modify plans.

Domain Library – A description of the domain and a library of possible actions.

Constraint Managers – Support modules which manage detailed constraints within a plan and seek to maintain as accurate a picture as possible of the feasibility of the current plan state with respect to the domain.

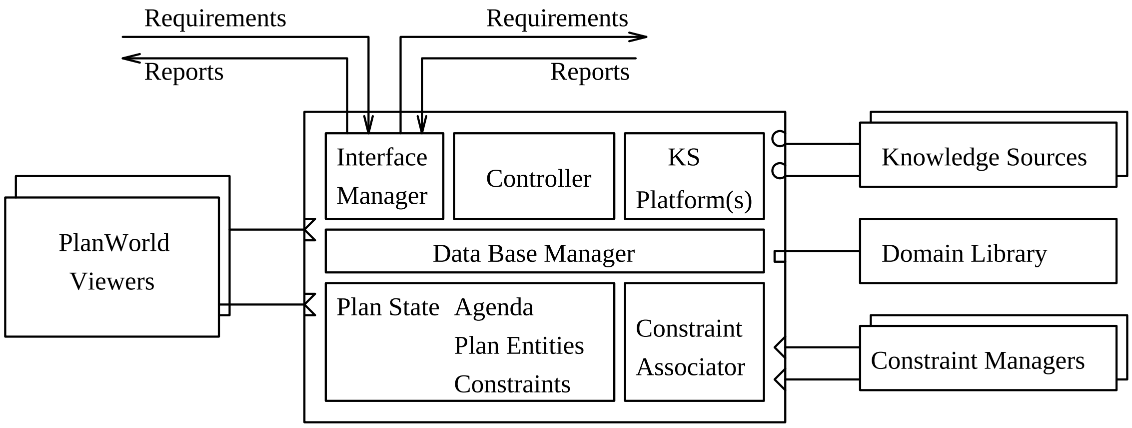

Figure 2: O-Plan Agent Architecture

These plug-in components are orchestrated by an O-Plan agent kernel which carries out the tasks assigned to it via appropriate use of the Knowledge Sources and manages options being maintained within the agent's *Plan State*. The central control flow is as follows:

Interface Manager – Handles external events (requirements or reports) and, if they can be processed by the agent, posts them on the agent *Agenda*.

Controller – Chooses Agenda entries for processing by suitable Knowledge Sources

Knowledge Source Platform(s) – Chosen Knowledge Sources are run on an available and suitable Knowledge Source Platform.

Data Base Manager – Maintains the Plan State being manipulated by the agent and provides services to the Interface Manager, Controller and Knowledge Sources running on KS Platforms to allow this.

Constraint Associator Acts as a mediator between the Plan State maintained by the data base manager and the various Constraint Managers that are installed in the agent. It eases the management of interrelationships between entities and detailed constraints.

3 Task Assignment to a Situated Planning Agent

The aim of this section is to provide an overview of the new focus of the O-Plan project. The main emphasis is on the development of a task assignment agent and the link between it and the planning agent. In developing the designs for a task assignment agent the project has addressed three key issues:

- **Communication**:
  Within a number of situated agents each with their own perspective of the planning

problem it is essential that they can communicate plans, tasks, etc, in a form which can be easily generated and integrated into different perspectives. The <I-N-OVA> [2] constraint model of plans [14] has been developed to address this and a number of other issues related to the communication of plans.

- **Plan Quality**:
  The task assigner needs to analyse the quality of the plans being generated and to provide feedback and direction concerning the options and plans which should be explored further. Joint work with USC/ISI to link O-Plan to their EXPECT system [7] has shown that plans can be generated and analysed to provide valuable feedback to human planners.

- **Role of Authority**:
  The activities of the various situated agents need to be coordinated and authority management is viewed as one of the main constraints to be considered. For example, in generative planning authority is important in considering which options are available and what level of planning detail should be considered. In enactment it is important to identify (and possibly name) which phases of the plans can be executed and which parts should be held back.

Each of these items will be dealt with in a separate subsection.

3.1 Communication of Plans between Situated Agents

The <I-N-OVA> (*Issues – Nodes – Orderings/Variables/Auxiliary*) Model is a means to represent plans 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 to the environments in which they are generated.

As shown in Figure 3 the <I-N-OVA> constraint model underlying plans is intended to support a number of different uses of plan representations:

- automatic manipulation of plans and to act as an ontology to underpin such use.
- human communication about plans.
- principled and reliable acquisition of plan information.
- formal reasoning about plans.

These cover both formal and practical requirements and encompasses the needs of both human and computer-based planning systems.

Our aim is to characterise the plan representation used within O-Plan [4],[16] and to more closely relate this work 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 production planning systems.

[2] <I-N-OVA> is pronounced as in "Innovate".

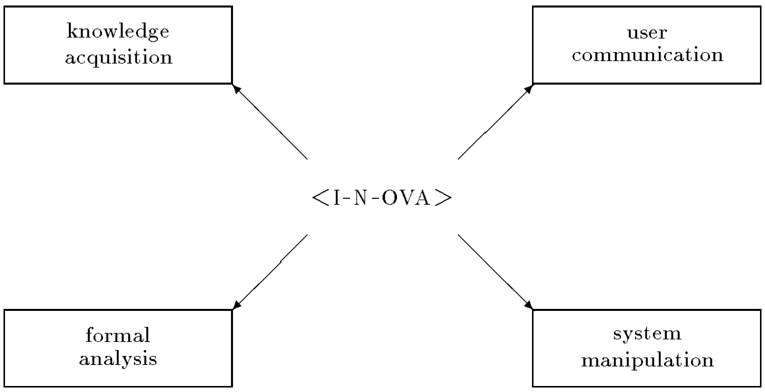

Figure 3: <I-N-OVA> Supports a Number of Requirements

A plan is represented as a set of constraints which together limit the behaviour that is desired when the plan is executed. Work on O-Plan and other practical planners has identified different entities in the plan which are conveniently grouped into three types of constraint. The set of constraints describes the possible plan elaborations that can be reached or generated as shown in Figure 4.

The three types of constraint in a plan are:

1. Implied Constraints or "Issues" – the pending or future constraints that will be added to the plan as a result of handling unsatisfied requirements, dealing with aspects of plan analysis and critiquing, etc. The implied constraints are the issues to be addressed, i.e., the "to-do" list or agenda which can be used to decide what plan modifications should be made to a plan by a planner (user or system).

2. Plan Entities or Plan Node Constraints – the main plan entities related to external communication of a plan. They describe a set of external names associated to time points. In an activity planner, the nodes are usually the actions in the plan associated with their begin and end time points. In a resource-centred scheduler, nodes may be the resource reservations made against the available resources with a begin and end time point for the reservation period.

3. Detailed Constraints – specialised constraints on the plan associated with plan entities. Empirical work on the O-Plan planner has identified the desirability of distinguishing two special types of detailed constraint:

   - Ordering or Temporal Constraints (such as temporal relationships between the nodes or metric time properties).
   - Variable Constraints (co-designation and non-co-designation constraints on plan objects in particular).

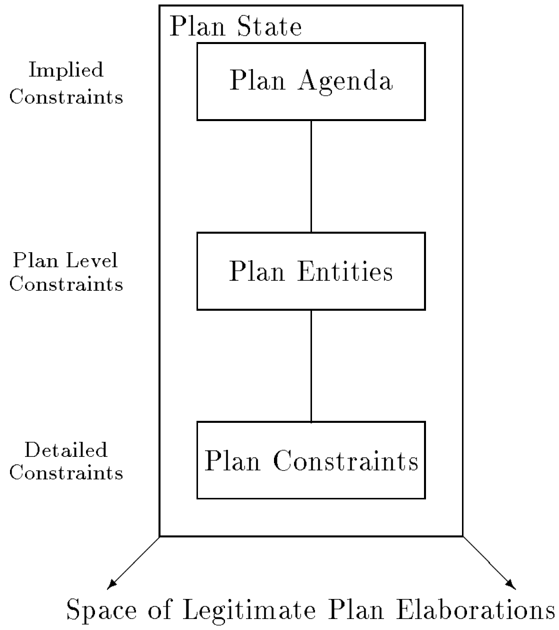

Figure 4: Plan Constraints Define Space of Plan Elaborations

These two constraints are highlighted since they often form part of other constraints within a temporal reasoning domain such as occurs in planning and scheduling problems. Knowing that these constraints have such cross "associations" has been found to simplify planner system design of constraint handling mechanisms and ease implementation issues [13],[15].

Other Detailed Constraints relate to input (pre-) and output (post-) and protection conditions, resources, authority requirements, spatial constraints, etc. These are referred to as:

- Auxiliary Constraints.

Auxiliary Constraints may be expressed as occurring at a time point (referred to as "point constraints") or across a range of the plan (referred to as "range constraints"). Point constraints can be used to express input and output constraints on nodes or for other constraints which can be expressed at a single time point. Range constraints relate to two or more time points and can be used to express protection intervals, etc.

There is a deliberate and direct mapping of the model of plans and activity used within O-Plan and the <I-N-OVA> Constraint Model of Plans to existing structured analysis and diagraming methods such as IDEF, R-Charts, etc. Other researchers have recognised the value of merging AI representation concepts with structured analysis and diagramming techniques for systems requirements modelling (e.g., [3],[11]).

3.2 Integrating Plan Quality Considerations into Planning

In producing plans, human planners take into account a variety of criteria that guide their decisions. Besides constraints imposed by the domain itself, these criteria often express preferences among alternative plans that meet the given requirements. Human planners can use these criteria for two important purposes:

- when asked to generate one plan, human planners are able to discern between an ordinary solution and a better quality one and propose the latter.
- when asked to generate several alternative plans, human planners are able to discern between similar alternative solutions and qualitatively different ones. They can relax different criteria to explore tradeoffs.

Current AI planners are good at generating a solution that satisfies the requirements that they are given. Some planners provide facilities to control the quality of the solution to be returned, by using evaluation functions or search control rules. However, they do not usually integrate plan quality considerations across several plans. In addition, there is not enough data on the adequacy of these representations to reflect the plan quality criteria that are necessary in practice. Often, the quality criteria that human expert planners consider:

- are highly dependent on the situation and the scenario at hand. Some criteria may be more important if there is a certain deadline, or new criteria may need to be considered if new considerations come up.
- include complex factors and tradeoffs that are often not represented by an automatic planner

Research in the area of plan analysis has concentrated on addressing two keys issues:

- to provide a tool – EXPECT – [7] which allows human planners to define criteria for plan quality and for preferences among alternative plans and options.
- to operationalise these criteria to guide a generative planner – O-Plan – in proposing better quality plans [8, 5, 6].

The approach taken has been to combine the O-Plan planner with a knowledge-based system that reasons about plan evaluation EXPECT Figure 5 describes the architecture of the O-Plan and EXPECT systems and the way in which plans and analysis information flows.

Using these two systems it has been possible to build an interface between the planner and the user that provides the following functionality:

- support to the user in defining criteria for evaluating plan quality through a knowledge acquisition tool,
- evaluation of the quality of plans proposed by the planner,

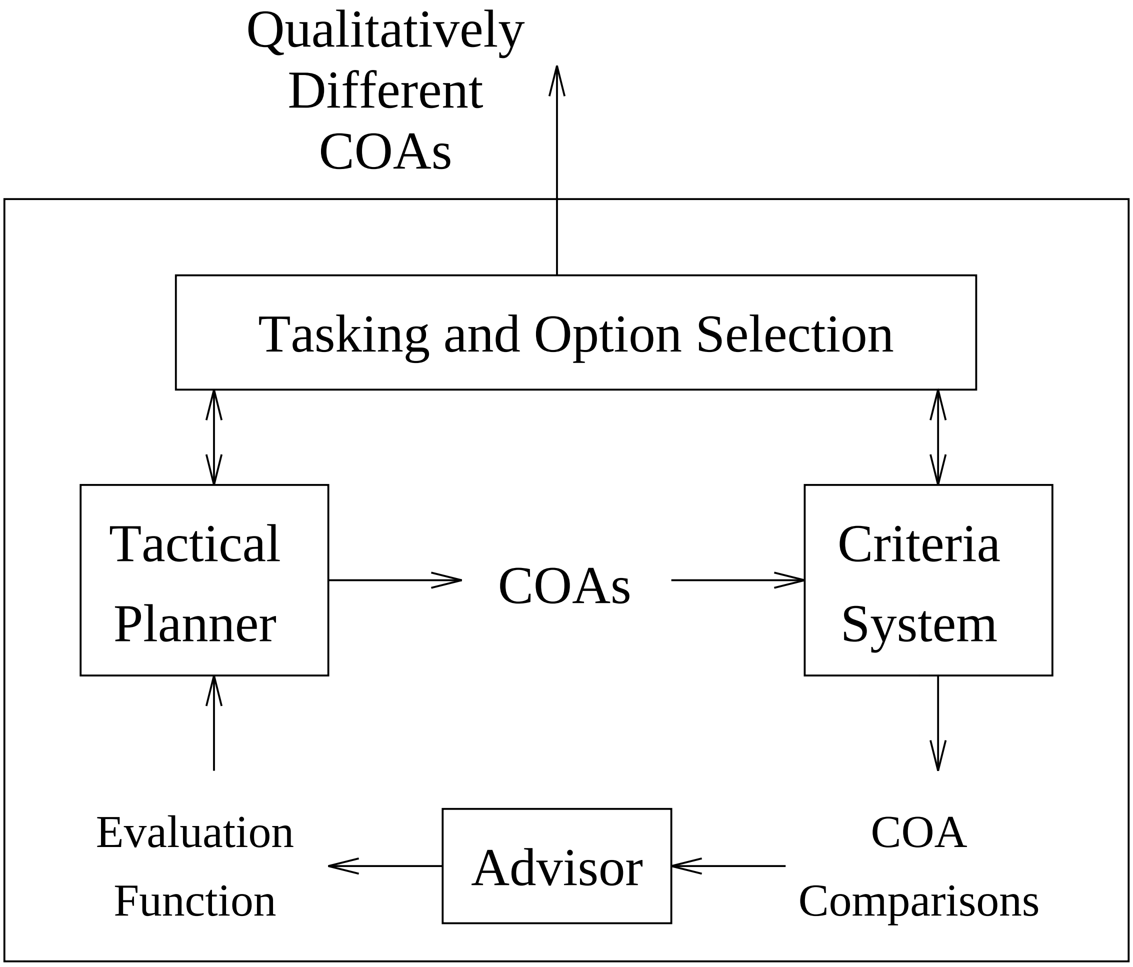

Figure 5: Situated Planner and Plan Analysis Tool

- provision of justifications for good and bad plan quality.

The work on plan analysis is motivated by the transportation planning domain that is the focus of the ARPA/Rome Laboratory Planning Initiative [10]. This domain involves the movement of materials and forces with a mixture of aircraft, ships, trucks and trains. The task being investigated is to generate multiple Courses of Action (COAs) and an evaluation of the tradeoffs among them using the relevant plan quality evaluation factors from a logistics perspective. This allows the human planners to identify those options which are critical to a plan's success and those parts of the plan which need further exploration and refinement. An evaluation matrix from a series of different COAs is described in Figure 6.

To date the O-Plan system is able to generate plans which can be evaluated by the EXPECT system. Work is continuing to extend EXPECT and O-Plan to strengthen the ability to support a user in specifying, comparing and refining the constraints on qualitatively different plans at the task assignment level of a planning support environment

	COA 1	COA 2	COA 3	COA4
A-PORTS:				
- airports	1	1	1	2
- sorties/hr	315	315	315	480
- sq ft ac parking	2M	2M	2M	3M
S-PORTS:				
- seaports	1	1	1	2
- piers	6	6	6	15
- berths	6	6	6	16
- max vessel size in ft	600	600	600	765
- oil facilities	1	1	1	3
CLOSURE DATE	C + 29	C + 22	C + 23	C + 23
LOG PERS	1154	5360	5396	7362
LOCs:				
- number locations	1	5	7	6
- miles max distance	20	99	140	120
- air and sea?	yes	yes	yes	yes

Figure 6: EXPECT's evaluation of several alternative plans generated by O-Plan.

and to allow this information to be used directly by O-Plan in guiding it in its search for the best solution.

3.3 The Role of Authority for a Situated Planning Agent

At the moment, the Task Assignment agent in O-Plan informs the planner and execution agents when they can create a plan for a nominated task and when a plan can be executed. This is done through a simple menu interface. It is intended that O-Plan will support authority management in a more comprehensive and principled way in future [12]. O-Plan will support:

- the notion of separate *plan options* which are individually specified task requirements, plan environments and plan elaborations. The Task Assignment agent can create as many as required. The plan options may contain the same task[3] with different search options or may contain a different task and environmental assumptions. It is possible to have only one plan option as the minimum[4].
- the notion of plan *phases*. These are individually provided actions or events stated explicitly in the top level task description given by the Task Assignment agent. Greater precision of authority management is possible by specifying more explicit phases at the task level. It is possible to have only one "phase" in the task as the minimum[5].

[3] Mutiple conjunctive tasks in one scenario is also possible.
[4] Plan options may be established and explicitly switched between by the Task Assignment agent.
[5] In fact any sub-component of any task schema or other schema included by task expansion in a plan can be referred to as a "phase" within the O-Plan planning agent.

- the notion of plan *levels*. Greater precision of authority management is possible by specifying more explicit levels in the O-Plan domain description language Task Formalism (TF). It is possible to have only one "level" in the domain as the minimum.
- for each "phase", planning will only be done down to an authorised "level" at which point planning will suspend leaving appropriate agenda entries until deeper planning authorisation is given.
- execution will be separately authorised for each "phase".

The Task Assignment agent of O-Plan will support authority management in a task setting framework. To establish an appropriate basis for future developments and allow for some initial internal support for authority management to be incorporated, the current release version of O-Plan has a simple authority scheme and reports this in the Task Assignment menu shown here.

```
Domain: pacifica
Status: plan option 1 - planning ...
Task: Operation_Blue_Lagoon
Authority: plan(all=inf), execute(all=no)
```

This reports that the system is planning for task `Operation_Blue_Lagoon` in the domain `pacifica` and that it is planning within `plan option 1` currently. It is authorised to plan to any level of detail for all phases `(plan all=inf)` but is not yet authorised to execute any actions (`(execute all=no)`.

A prototype HARDY-based[6] user interface for the Task Assignment agent has been created.

4 Summary

There is a need to view AI planning systems as being situated within an environment where there are a number of other agents and systems which deal with the whole planning problem. While the planner is responsible for the plan generation aspects of the problem, other agents should be responsible for dealing with other aspects of the whole "planning" problem, e.g. task elicitation, plan analysis, reactive execution monitoring, etc. This view is motivated by the obvious realisation that planning systems cannot operate in isolation and for a task to be solved successfully its needs to be communicated and reasoned with between a number of systems.

The O-Plan architecture has been designed to support the creation of situated agents and work to date has concentrated on building generative planning agents and an execution agent with links between these two agents. The outcome of this research has been used in a number of systems. For example, the reactive execution agent work of Reece [9] and the Optimum-AIV [1] system developed for Assembly, Integration and Verification of spacecraft at the European Space Agency.

[6] HARDY is a C++ based diagramming aid and hypermedia tool from AIAI.

Future research on the O-Plan architecture will concentrate on its ability to support a task assignment agent and the link between it and the planning agent. This is an area of planning research which has been neglected by planning researchers. However, it is an important aspect of the planning problem as a planner needs to fully understand the requirements set by the task assigner and needs the guidance which this can provide in identifying an appropriate solution. The planner also needs to provide feedback on plan feasibility to the Task Assigner. As part of this research three key issues have already been investigated. The key issues are the representation and communication of plans as sets of constraints, the use of quality criteria to analyse and direct the generative planner and the role of authority in coordinating the activities of a number of situated agents.

References

[1] Aarup, M., Arentoft, M.M., Parrod, Y., Stokes. I., Vadon, H. and Stader, J. *Optimum-AIV: A Knowledge-Based Planning and Scheduling System for Spacecraft AIV*, in Intelligent Scheduling (eds. Zweben, M. and Fox, M.S.), Published by Morgan Kaufmann Inc, San Francisco, USA, 1994.

[2] Allen, J., Hendler, J. and Tate, A., *Readings in Planning*. Published by Morgan-Kaufmann Inc, San Francisco, USA, 1990.

[3] Borgida, A., Greenspan, S. and Mylopoulos, J., Knowledge Representation as the Basis for Requirements Specifications, IEEE Computer Magazine, Special Issue on Requirements Engineering Environments, April 1985.

[4] Currie, K.W. and Tate, A., O-Plan: the Open Planning Architecture, *Artificial Intelligence* 52(1), pp. 49-86, North-Holland, 1991.

[5] Drabble, B., Gil, Y. and Tate, A. *Acquiring Criteria for Plan Quality Control*, in the proceedings of the AAAI Spring Symposium workshop on Integrated Planning Applications, June 1995. Stanford Univerisity, CA, USA. Published by the American Association for Artificial Intelligence, Menlo Park, California.

[6] Drabble, B., Gil, Y. and Tate, A. *Yes, but why is that plan better?*, in the proceedings of the International Conference on Artificial Intelligence in the Petroleum Industry, 13th-15th September 1995, Lillehammer Hotel, Lillehammer, Norway.

[7] Gil, Y. *Knowledge Refinement in a Reflective Architecture*, in the proceedings of the Twelfth National Conference on Artificial Intelligence, Seattle, WA, USA. August 1994. Published by AAAI Press/ The MIT Press Menlo Park, CA, USA.

[8] Gil, Y., Hoffman, M. and Tate, A. *Domain Specific Criteria to Direct and Evaluate Planning Systems*, in proceedings of the ARPA/Rome Laboratory Planning Initiative Workshop, Tucson, February 1994. Published by Morgan Kaufmann, Published Inc, San Francisco, USA, 1994.

[9] Reece, G. A. and Tate A. *Synthesising Protection Monitors from Protection Intervals*, in the proceedings of the Second International Conference on Artificial Intelligence Planning Systems, (eds. K. Hammond), University of Chicago, June 13-15 1994, pp 146-152. AAAI Press, Menlo Park, California.

[10] Reece, G.A., Tate A., Brown, D.I., Hoffman, M. and Bernard, R.E. *The PRECiS Environment*, in the proceedings of the ARPA/Rome Laboratory Planning Initiative Workshop at the National Conference on Artificial Intelligence (AAAI-93), Washington, DC, 1993.

[11] Ramesh, B. and Dhar, V., Representing and Maintaining Process Knowledge for Large-Scale Systems Development, IEEE Expert, pp. 54-59, April 1994.

[12] Tate, A, *Authority Management - Coordination between Task Assignment, Planning and Execution*, in the working papers of the IJCAI-93 workshop on Knowledge-based Production Planning, Scheduling and Control, August 1993.

[13] Tate, A., The Emergence of "Standard" Planning and Scheduling System Components, in *Current Trends in AI Planning*, (eds. Backström, C. & Sandewall, E.), IOS Press, 1993.

[14] Tate, A. Characterising Plans as a Set of Constraints - the <I-N-OVA> Model – a Framework for Comparative Analysis, in the Special Issue on "Evaluation of Plans, Planners, and Planning Agents", ACM SIGART Bulletin Vol. 6 No. 1, January 1995.

[15] Tate, A., Drabble, B. and Dalton, J. Reasoning with Constraints in O-Plan2, Proceedings of the ARPA/Rome Laboratory Planning Initiative Workshop, (ed. Burstein. M.), Tucson, Arizona, USA, Morgan Kaufmann, 1994.

[16] Tate, A., Drabble, B. and Kirby, R.B., O-Plan2: an Open Architecture for Command, Planning and Control, in *Intelligent Scheduling* (eds. Zweben. M. and Fox. M.S.). Published by Morgan Kaufmann Inc, San Francisco, USA, 1994.