O-P^3 – Open Planning Process Panels

Austin Tate, John Levine, Jeff Dalton and Stuart Aitken
Artificial Intelligence Applications Institute
The University of Edinburgh
80 South Bridge, Edinburgh EH1 1HN, UK
{a.tate, j.levine, j.dalton, s.aitken}@ed.ac.uk
`http://www.aiai.ed.ac.uk/~oplan`

Abstract

This paper introduces Open Planning Process Panels (O-P^3). These panels are based on explicit models of the planning process and are used to coordinate the development and evaluation of multiple courses of action. We describe the generic ideas behind O-P^3 technology, a general methodology for building O-P^3 interfaces and two applications based on O-P^3 technology – the Air Campaign Planning Process Panel (ACP^3) and the O-Plan two-user mixed-initiative planning Web demonstration. This work has an impact on a number of important research areas outside planning, including Computer Supported Cooperative Work (CSCW) and workflow support.

Introduction

Real world planning is a complicated business. Courses of action to meet a given situation are constructed collaboratively between teams of people using many different pieces of software. The people in the teams will have different roles, and the software will be used for different purposes, such as planning, scheduling, plan evaluation, and simulation. Alternative plans will be developed, compared and evaluated, and more than one may be chosen for briefing. In general, planning is an example of a multi-user, multi-agent collaboration in which different options for the synthesis of a solution to given requirements will be explored.

The process of planning is itself the execution of a plan, with agents acting in parallel, sharing resources, communicating results and so on. This planning process can be made explicit and used as a central device for workflow coordination and visualisation.

We have used this idea to create Open Planning Process Panels (O-P^3). These panels are used to coordinate the workflow between multiple agents and visualise the development and evaluation of multiple courses of action (COAs). The generic notion of O-P^3 has been used to implement two real applications – the Air Campaign Planning Process Panel (ACP^3) and the O-Plan two-user mixed-initiative planning Web demonstration. In the former, O-P^3 is used to build a visualisation panel for a complex multi-agent planning and evaluation demonstration (TIE 97-1) which uses 11 different software components and involves several users. In the latter, O-P^3 technology is used to enable the development and evaluation of multiple COAs by a commander, a planning staff member and the O-Plan automated planning agent.

O-P^3 technology could have an impact on several important research areas:

- Automated planning: O-P^3 shows how automated planning aids such as AI planners can be used within the context of a wider workflow involving other system agents and human users.
- Computer-supported cooperative work (CSCW): O-P^3 uses explicit models of the collaborative planning workflow to coordinate the overall effort of constructing and evaluating different courses of action. This is generalisable to other team-based synthesis tasks using activity models of the task in question (e.g. design or configuration).
- Multi-agent mixed-initiative planning: O-P^3 facilitates the sharing of the actions in the planning process between different human and system agents and allows for agents to take the initiative within the roles that they play and the authority that they have (Tate, 1993).
- Workflow support: O-P^3 provides support for the workflow of human and system agents working together to create courses of action. The workflow and the developing artefact (i.e. the course of action) can be visualised and guided using O-P^3 technology.

The kind of planning system that we envisage O-P^3 being used for is one in which the planning is performed by a team of people and a collection of computer-based planning agents, who act together to solve a hard, real world planning problem. Both the human and the system agents will act in given roles and will be constrained by what they are authorised to do, but they will also have the ability to work under their own initiative and volunteer results when this is appropriate. When the planning process is underway, the agents will typically be working on distinct parts of the plan synthesis in parallel. The agents will also be working in parallel to explore different possible courses of action; for example, while one COA is being evaluated, another two may be in the process of being synthesised.

This paper introduces O-P^3 technology. It begins with a description of the generic O-P^3 ideas, based on the central notion of an explicit shared model of the activities involved in creating a plan – the planning process. We then describe the two applications which have been based on O-P^3 – ACP^3 and the O-Plan Web demonstration. We conclude with a summary and future directions for O-P^3.

Generic O-P^3 Technology

The generic O-P^3 is based on an explicit model of the planning process, which would be encoded using an activity modelling language such as IDEF3. This represents the planning process as a partially-ordered network of actions, with some actions having expansions down to a finer level of detail (i.e. to another partially-ordered network).

The purpose of O-P^3 is to display the status of the nodes in the planning process to the users, to allow the users to compare the products of the planning process (i.e. the courses of action) and to allow the users to control the next steps on the "workflow fringe" (i.e. what actions are possible next given the current status of the planning process). In the context of creating plans, O-P^3 is designed to allow the development of multiple courses of action and the evaluation of those courses of action using various plan evaluations.

A generic O-P^3 panel would have any of a number of "sub-panels", which can be tailored to support specific users or user roles. These include:

- A course of action comparison matrix showing:
  - COAs vs elements of evaluation, with the plan evaluations being provided by plug-in plan evaluators or plan evaluation agents;
  - the steps in the planning process (from the explicit process model), the current status of those steps (the *state model*), and control for the human agent of what action to execute next;
  - the *issues* outstanding for a COA that is being synthesised and which must be addressed before the COA is ready to execute;
- a graphical display showing the status of the planning process as a PERT chart, which is a useful alternative view of the planning process to that given by the tabular matrix display;
- other visualisations, such as bar charts, intermediate process product descriptions, and textual description of plans.

The generic O-P^3 methodology for building Open Planning Process Panels consists of the following steps:

- Consider the agents (human and system) who are involved in the overall process of planning. Assign roles and authorities to these agents.
- Construct an activity model of the planning process, showing the partial ordering and decomposition of the actions and which agents can carry out which actions. This activity model could be represented using an activity modelling language such as IDEF3.
- Build a model of the current state of the planning process and an activity monitor which will update this state model as actions in the planning process take place.
- Construct appropriate O-P^3 interfaces for each of the human agents in the planning process, taking into account the role which they play in the interaction. This means that each different user role will have a O-P^3 interface which is tailored to the overall nature of their task.

Generic O-P^3 design rules are used to inform the construction of the O-P^3 interfaces:

- Each user role in the planning process is provided with a panel which is tailored to activities and needs of that role.
- Each user role is assigned a colour to distinguish between the roles. This is used, for example, as a background colour for the header of the panel. Since a given user may act in more than one distinct user role, this acts as a useful visual cue as to which user role is being enacted at any one time.
- The generic O-P^3 panel consists of three parts: a graph sub-panel (PERT chart), a matrix sub-panel (COA comparison matrix) and other sub-panels (e.g. information on assumed environmental conditions). The graph sub-panel and the other sub-panels are optional items (depending on how useful they are for a given application).

- The graph sub-panel contains a partially-ordered graph showing the activity model of the planning planning process. Since the activity model may be large and may apply for each COA being developed, it may not be possible to show the whole network, so some sort of navigation based on decompositions and switching between COAs may be needed.

- The actions shown in the graph sub-panel are annotated with colours to show their current status in the *state model* (see above). The colours used are adapted from other ARPI plan visualisation work (Stillman and Bonissone, 1996).

- The matrix sub-panel is a table which contains two types of rows and and two types of columns. The rows are process steps (verb phrases) and COA descriptors (noun phrases). The process steps labels are coloured with the user role background colour and the COA descriptors are white. The columns are the individual COAs being developed (labelled COA-N) and a column reflecting the overall workflow (labelled "Overall").

- The process steps in the matrix sub-panel are an appropriately flattened form of the activity model of the planning process. The status of the actions can be shown using the same colours as are used in the graph sub-panel. The currently active workflow fringe (i.e. what can be done next) is shown using active hyperlinks – clicking on a hyperlink initiates the action.

- The rows are arranged in three parts, running from top to bottom. The first section is concerned with process steps prior to plan synthesis, such as setting the COA requirements. The middle section consists of the COA descriptors and is filled out when a COA has been synthesised. The final section consists of process steps which come after plan synthesis, such as addressing any outstanding issues and viewing the resulting COA in various ways.

- The COA descriptors relate to the COA products produced by the steps of the planning process, such as the minimum duration of the plan and the effectiveness. These can be provided by separate plan evaluators, simulators, etc. The COA descriptors can be selected by the users to show only the critical elements of evaluation. Colours are used to show whether the result is acceptable and raises no issues (green), is possibly acceptable but has some issues to note (orange) or is not acceptable unless the user is prepared to relax the initial requirements (red).

- The other sub-panels can contain other useful information such as tables showing the COA objectives and assumed environmental conditions for each COA.

The O-P^3 agent interfaces then allow the human agents to play their part in the overall planning process, alongside the system agents, which will be AI planners, schedulers, plan evaluators and so on. This is illustrated in Figure 1.

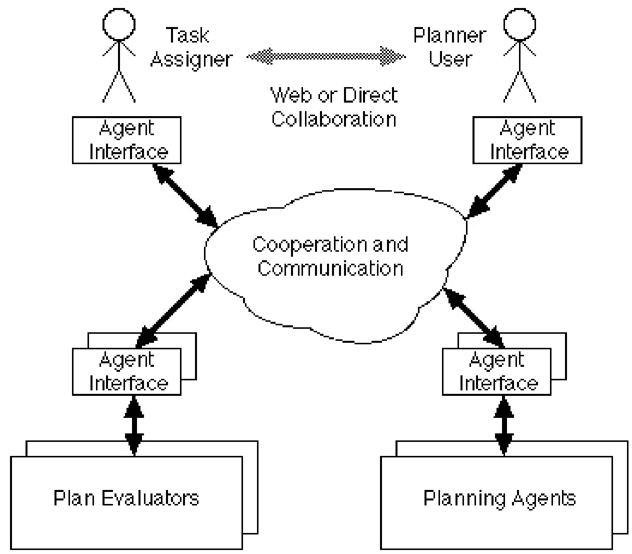

Figure 1: Using O-P^3 Interfaces

Application 1 – ACP^3

The ARPI TIE 97-1 demonstration brings together eleven, separately developed, software systems for planning and plan evaluation. When the demonstration is run, these systems work together to create and evaluate multiple courses of action in the domain of Air Campaign Planning. The systems communicate with each other by exchanging KQML messages. Finding out what is happening at any given time could (in theory) be done by watching these KQML messages, but this was obviously less than ideal as these messages use technological terms which are far removed from the terminology used by the user community.

Our aim was to use O-P^3 technology to build a visualisation component for this demonstration which would allow the target end users to view the current state of the planning process in process terms they are familiar with. This has resulted in ACP^3 – the Air Campaign Planning Process Panel.

Modelling the Planning Process

The software components of TIE 97-1 can be described as performing activities such as planning, scheduling, simulation and plan evaluation. Going into more detail, we can talk about hierarchical task network planning and Monte Carlo simulation methods. However, end users are more likely to conceive of the processes of Air Campaign Planning in more general, domain-related terms, such as “develop JFACC guidance” and “create support plan”. The gaps in terminology and in levels of description can be bridged by building models of the planning process which are rooted in established ACP terminology. We have therefore made use of the previously elicited and verified ACP process models of Drabble, Lydiard and Tate (1997) as our source of terminology and as the basis of our IDEF3 models of the planning process for TIE 97-1. The full models used for building ACP^3 are described in Aitken and Tate (1997).

Building ACP^3

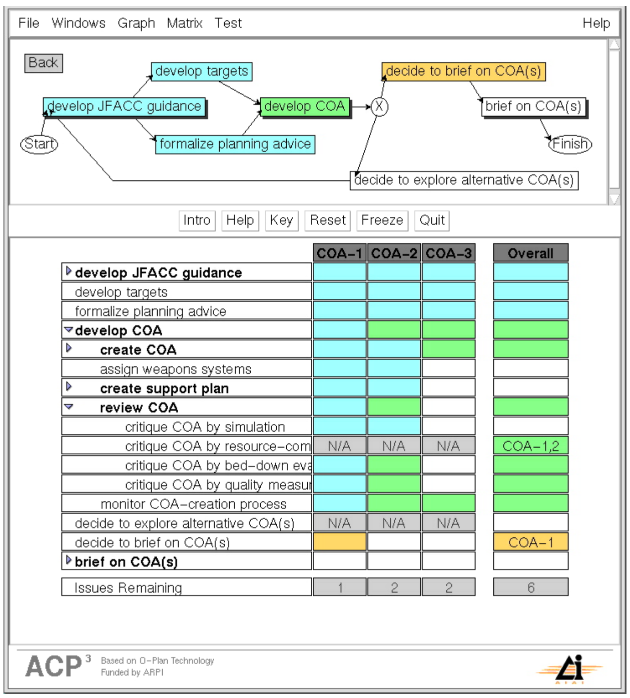

Figure 2: The ACP^3 Viewer

The ACP^3 viewer is shown in Figure 2. The purpose of ACP^3 is to track the overall planning process and display this to the viewers of the ARPI TIE 97-1 demonstration in a meaningful way using appropriate military process terminology. The planning process is shown in two separate sub-panels. The tabular COA comparison matrix shows COAs being developed (columns) against a tree-based view of the planning process. The graph viewer sub-panel shows the planning process as a PERT network. Since the planning process consists of many nodes with expansions, the graph viewer can only display one individual graph from the planning process for one COA. Other graphs may be reached by clicking on nodes with expansions, and the end user can choose which COA to view.

The two views are required because the planning process in TIE 97-1 is a complex artefact. It is possible to see the whole process for every COA in the COA matrix, but information about the partial ordering of the actions in a graph is lost when the graph is converted to a tree structure. The graph viewer shows the full partial ordering but space considerations mean that only a single graph for a single COA can be shown at one time.

The ACP^3 process monitor works by watching for certain KQML messages which it can relate to the status of certain nodes in the ACP process models. As the demonstration proceeds, the status of actions in the model progress from white (not yet ready to execute), to orange (ready to execute), then to green (executing) and finally blue (complete). The final column in the COA matrix is labelled “overall” and summarises the overall status of the COA creation and evaluation process.

The panel is written entirely in Java to form the basis for future Web-based process editors and control panels.

Application 2 – O-Plan

The current O-Plan project (Tate, Drabble and Dalton, 1996; Tate, Dalton and Levine, 1998) is concerned with providing support for mixed-initiative planning. The current demonstration shows interaction between two human agents and one software planning agent (the O-Plan plan server). 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 to populate the COA matrix component of the O-P^3 interface.

As shown in Figure 3, we envisage two human agents acting in the user roles of Task Assigner and Planner User, working together to explore possible solutions to a problem and making use of automated planning aids to do this. Figure 4 shows how the two human agents work together to populate the matrix. The Task Assigner sets the requirements for a particular course of action (i.e. what top level tasks must be performed), selects appropriate evaluation criteria for the resulting plans and decides which courses of action to prepare

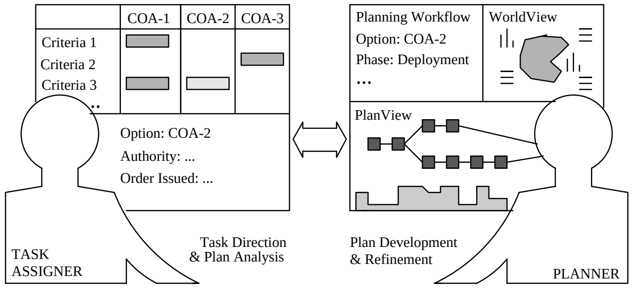

Figure 3: Communication between TA and Planner

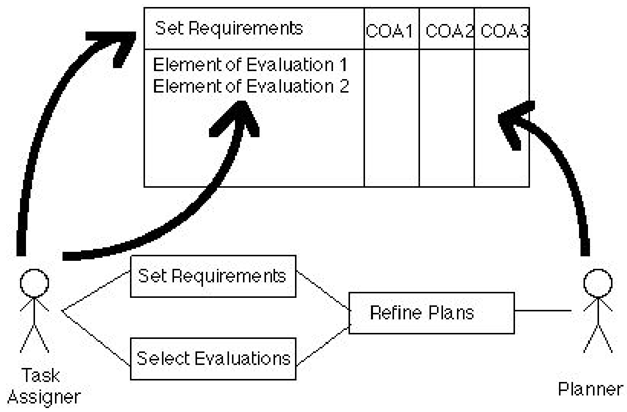

Figure 4: Roles of the Task Assigner and the Planner

for briefing. The Planner User works with O-Plan to explore and refine the different possible course of action for a given set of top level requirements. The two users can work in parallel, as will be demonstrated in the example scenario.

The overall planning task is thus shared between three agents who act in distinct user and system roles. The Task Assigner (TA) is a commander who is given a crisis to deal with and who needs to explore some options. This person will be given field reports on the developing crisis and environmental conditions. The Planner User is a member of staff whose role is to provide the TA with plans which meet the specified criteria. In doing this, the Planner User will make use of the O-Plan automated planning agent, whose role is to generate plans for the Planner User to see. The Planner User will typically generate a number of possible course of action using O-Plan and only return the best ones to the TA.

For our current demonstration, we are using a general purpose logistics and crisis operations domain which is an extension of our earlier Non-Combative Evacuation Operations (NEO) and logistics-related domains (Reece *et al.*, 1993). This domain, together with the O-Plan Task Formalism (TF) implementation, is described in detail by Tate, Dalton and Levine (1998).

The two human users are provided with individual O-P^3 panels which are implemented using a CGI-initiated HTTP server in Common Lisp and which therefore run in any World Wide Web browser – the Common Lisp process returns standard HTML pages. This way of working has many advantages:

- the two users can be using different types of machine (Unix, PC, Mac) and running different types of Web browser (Netscape, Internet Explorer, Hotjava, etc.);
- the only requirement for running O-Plan is a World Wide Web connection and a Web browser (i.e. no additional software installation is needed);
- the two users can be geographically separate – in this case, voice communication via the telephone or teleconferencing is all that is required in addition to the linked O-P^3 interfaces.

The planning process for the TA and the Planner User is made explicit through the hypertext options displayed in the process parts of the O-P^3 panels. These are either not present (not ready to run yet), active (on the workflow fringe) or inactive (completed). Further parts of the planning process are driven by *issues* which O-Plan or the plan evaluation agents can raise about a plan under construction and which can be handled by either or both of the human agents. Because the planning process is made explicit to the two users through these two mechanisms, other visualisations of the planning process itself are not required. However, the products of the planning process (the courses of action) are complex artefacts for which multiple views are needed. In the current version, the courses of action can be viewed as a PERT network, as a textual narrative, or as a plan level expansion tree (all at various levels of detail).

The user roles are arranged such that the TA has authority over the Planner User who in turn has authority over O-Plan. This means that the TA defines the limits of the Planner User's activity (e.g. only plan to level 2) and the Planner User then acts within those bounds to define what O-Plan can do (e.g. only plan to level 2 and allow user choice of schemas). Other aspects of what the two users are authorised to do are made explicit by the facilities included in their respective panels.

The COA Comparison Matrix

The two panels for the Task Assigner and Planner User are shown in Figures 5 and 6. Each user has control

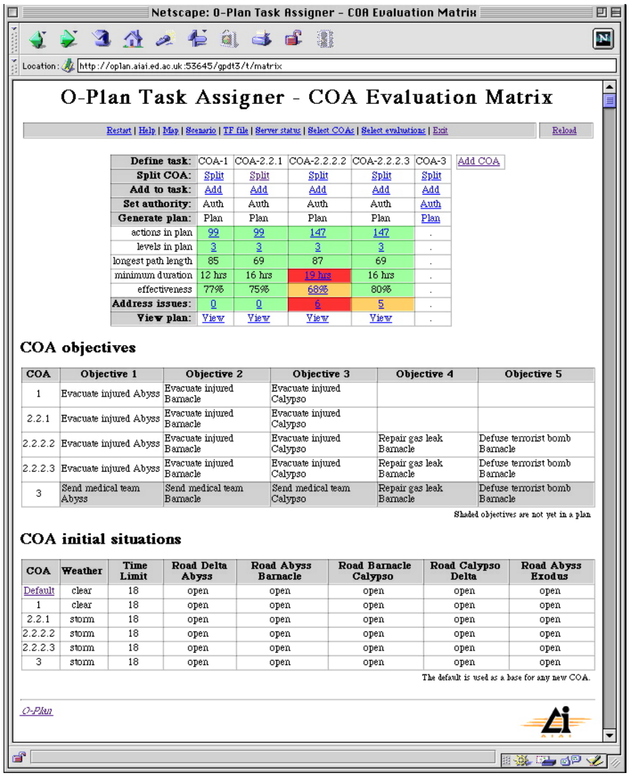

Figure 5: The Task Assigner's Panel

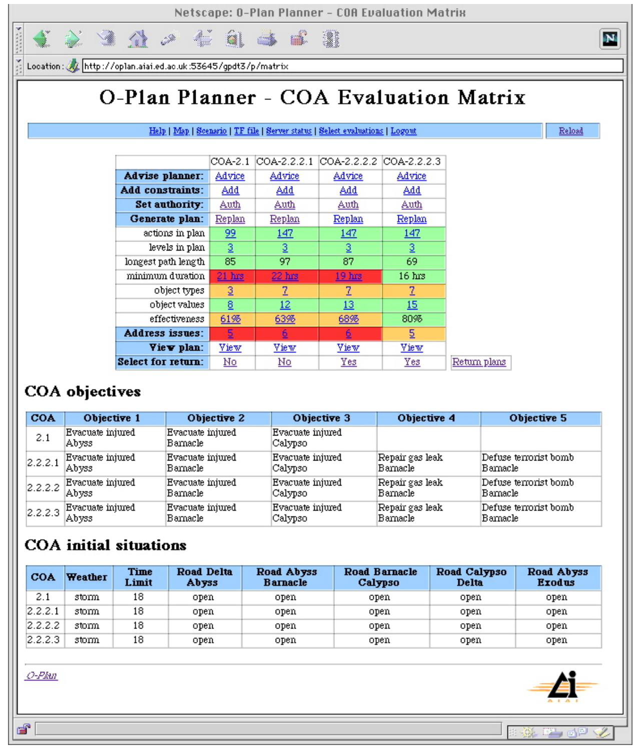

Figure 6: The Planner User's Panel

over the plan evaluation elements which are shown, to enable the critical elements of evaluation to be chosen. In the example scenario given later, the TA is only interested in the minimum duration and the effectiveness, so only these are selected. On the other hand, the Planner User wants a variety of data to pick the best COA, so all evaluations are shown.

The role of the TA is to set up the top level requirements for a course of action. Once this is done, the COA is passed across to the Planner User, whose matrix is initially blank. The Planner User then explores a range of possible COAs for the specified requirements and returns the best ones to the TA. When the Planner User returns a COA to the Task Assigner, the column for that COA appears in the Task Assigner's matrix. The Planner User and the Task Assigner can be working in parallel, as demonstrated in the scenario.

The Demonstration Scenario

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

Initial situation: the action takes place on the island of Pacifica, with emergencies being planned for at the cities of Abyss, Barnacle and Calypso. The TA is told to deal with injured civilians at Abyss, Barnacle and Calypso within the next 18 hours. Plans are only acceptable if their effectiveness is 75% or greater. The weather forecast gives a 50% chance of a storm within the next 24 hours (Figure 7).

Initial preparations: The TA sets up the default situation, setting the time limit to 18 hrs. The weather and road situations are left with their default values pending more accurate reports.

COA-1: The TA first explores the option of evacuating the injured from all three cities in clear weather. The COA requirements are passed directly to the planner user. A plan is generated which executes in 12 hrs and has an effectiveness of 77%, which is acceptable. The plan has 3 issues outstanding. The planner user addresses these and returns the plan to the TA.

COA-2: The TA then sets up a second COA with the same evacuation tasks but this time assuming stormy weather, to check for all eventualities. This new set of COA requirements is passed to the planner user. The first plan generated takes 21hrs and has an effectiveness of 61%, both of which are unacceptable. The planner asks the O-Plan planner for an alternative plan. The new plan (COA-2.2) executes in 16 hrs and has an effectiveness of 75%, both of which are acceptable. The planner user returns COA-2.2 to the TA

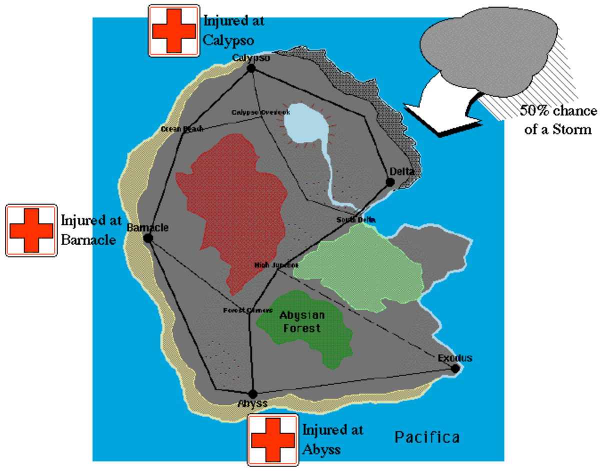

Figure 7: The Initial Situation

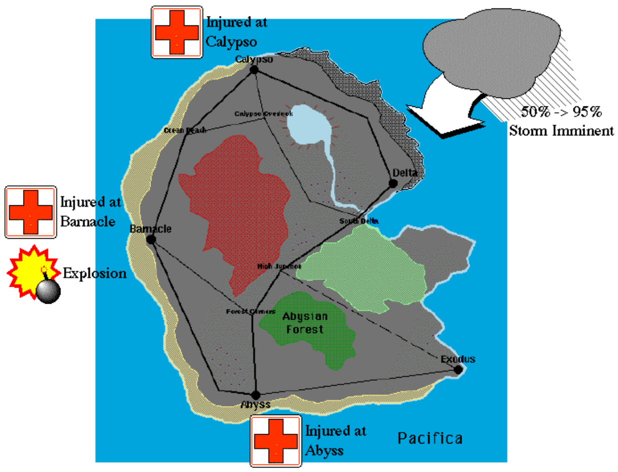

Figure 8: The Developing Situation

and deletes COA-2.1. At this point, the TA has an acceptable plan for both clear and stormy conditions.

Developing situation: the TA is now contacted by the Barnacle field station. Reports are coming in of an explosion at the power station, causing a gas leak. It is thought that this is due to a terrorist bomb, so it seems wise to fix the gas leak and send a bomb squad to defuse any remaining bombs. Meanwhile, the latest weather report indicates that a storm is brewing and has a 95% chance of hitting the island (Figure 8).

COA-2.2.2: to deal with this turn of events, the TA splits COA-2.2 (the realistic weather assumption) into two sub-options and adds two new tasks to COA-2.2.2, to repair the gas leak at Barnacle and send a bomb squad to Barnacle. COA-2.2.2 is now passed to the planner user. Since the original COA-2.2 took 16 hrs, the planner user switches schema choice on, to have fine control of the addition of the two new tasks to the existing plan. The planner user is given the option of using fast or slow vehicles for the two tasks and chooses fast vehicles. However, this plan takes 22 hrs and has an effectiveness of 63%. The planner user replans and chooses a mixture of fast and slow vehicles for the "repair gas leak" task and a fast vehicle for the "defuse terrorist bomb" task. While better, the new plan takes 19 hrs and has an effectiveness of only 68%. The TA is getting impatient and tells the planner user "this is taking too long. Just give me the best one so far." The planner user returns COA-2.2.2.2, keeping COA-2.2.2.1 for further back office work.

COA-3: The TA decides to try sending medical teams to the three cities to deal with the injured civilians rather than evacuating them. After updating the default situation to reflect the weather report, the TA starts to set up COA-3 with these tasks, and so begins to define the requirements on the screen.

COA-2.2.2.3: Meanwhile, the planner user has continued to explore the possibilities for COA-2.2.2. The plan was improved when the planner user used some slow vehicles in the plan, so it seems likely that this is because the limited number of fast vehicles are being used repeatedly, resulting in a longer (i.e. more linear) plan. The planner user presses "replan" and chooses to use a slow vehicle in the "defuse terrorist bomb" task – since sending the bomb squad is only a precaution, using the limited number of fast vehicles for evacuating the injured and fixing the known gas leak seems like a good idea. The planner user was right – the resulting plan executes in 16 hrs and has an effectiveness of 80%. Viewing the plan at level 2 displays that this plan has good parallelism. The planner user now addresses the issues raised by COA-2.2.2.3 and returns this plan to the TA, saying "I think I've fixed the problem with COA-2.2.2".

Back to COA-3: The TA sees the new plan. "That looks good, now see what you can do with COA-3 as an alternative". The planner user (still in "ask user" schema selection mode) selects the fast vehicle option for 4 of the tasks, but selects a slow vehicle for the "defuse terrorist bomb" task. The resulting plan executes in 12 hrs and has an effectiveness of 79%.

Choice of COA: The TA now has a choice between COA-2.2.2.3 and COA-3. While COA-3 takes 4 hrs less, it is slightly less effective, and more importantly, it only sends medical teams to the three cities rather than evacuating the injured people. The TA could now examine other details of the two plans, using the plan views and the other elements of evaluation, in order to make an informed choice between the two or plan further.

O-Plan – Summary

The O-Plan Web demonstration illustrates mixed-initiative interaction between two human agents and one system planning agent engaged in the process of developing multiple qualitatively different courses of action. O-P^3 interfaces are provided for the two human users which are tailored to their individual user roles.

Summary of O-P^3 Technology and Future Applications

In this paper, we have introduced the generic notion of Open Planning Process Panels (O-P^3). These panels are used to coordinate the workflow between multiple agents and visualise the development and evaluation of multiple courses of action (COAs). We have described how O-P^3 technology has been used to implement two real applications – the Air Campaign Planning Process Panel (ACP^3) and the O-Plan two-user mixed-initiative Web demonstration of crisis response planning.

Both of these systems have an explicit notion of the planning process, which is a multi-agent interaction. The agents in both systems are assigned with roles which relate to the actions the users can carry out in the planning process. Both systems use the notion of a COA matrix which shows possible steps in the planning process for each course of action being developed. In ACP^3, this is used as a visualisation device. In the O-Plan demonstration, the population of this matrix is central to the mixed-initiative interaction between the Task Assigner, Planner User and O-Plan.

A number of other applications of O-P^3 technology are envisaged. An O-P^3 panel for the US DARPA Genoa program's intelligence gathering process is under investigation. This panel, termed G-P^3, would include the matrix sub-panel and the graph sub-panel from O-P^3. However it is thought that G-P^3 would also include new sub-panels to provide a "process product" perspective (showing the status of various information products under development) and new panels intended to give more role specific workflow status for a number of types of user. The main innovation in G-P^3 would be hooks to allow intelligent planning technology (e.g. provided by O-Plan) to be used to dynamically generate and adapt workflows and the planning process to accommodate changing requirements and situations. Such an "Intelligent Workflow Planning Aid" using O-Plan has already been demonstrated for Air Campaign Planning process (Drabble, Tate and Dalton, 1996).

Acknowledgements

The O-Plan project is sponsored by the Defense Advanced Research Projects Agency (DARPA) and the US Air Force Research Laboratory at Rome (AFRL), Air Force Materiel Command, USAF, under grant number F30602-95-1-0022. The O-Plan project is monitored by Dr. Northrup Fowler III at AFRL. 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

Aitken, S. and Tate, A. (1997). Process Modelling of the TIE 97-1 Demonstration: Modelling Complex Techniques Using ACP Terminology. ISAT Technical Report ISAT-AIAI/TR/6, Version 1, December 1997.

Drabble, B., Tate, A. and Dalton, J. (1996). ACP Process Management: O-Plan IFD-5 Qualifier. O-Plan Technical Report ARPA-RL/O-Plan/TR/30, Version 1, November 1996.

Drabble, B., Lydiard, T. and Tate, A. (1997). Process Steps, Process Product and System Capabilities. ISAT Technical Report ISAT-AIAI/TR/4, Version 2, April 1997.

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., 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.

Tate, A., Dalton, J. and Levine, J. (1998). Generation of Multiple Qualitatively Different Plan Options. *Proceedings of AIPS-98*, Pittsburgh, USA.