Using Shared Models of Activity to Underpin Coalition Planning

Austin Tate, Jeff Dalton, John Levine and Peter Jarvis
Artificial Intelligence Applications Institute
The University of Edinburgh
80 South Bridge, Edinburgh EH1 1HN, UK
{a.tate, j.dalton, j.levine, p.a.jarvis}@ed.ac.uk
`http://www.aiai.ed.ac.uk/project/ix`

Abstract

A shared understanding, based on a shared underlying conceptual model, for objectives, the world state, activities, plans, the planning process, and inter-agent responsibilities and authorities can provide a framework for systems to support Coalition Operations. Such a shared model can have value in supporting separate organisational entities with their own internal or legacy systems, but coming together into "virtual oraganisations" for specific missions and to achieve specific objectives.

This paper describes research which it is proposed is conducted within the US DARPA Control of Agent-Based Systems (CoABS) Program.

Introduction

We intend to address the issue of how to make productive use of human and computer-based agents who are cooperating to carry out one or more tasks where multiple options for a solution are being considered. We will explore mixed initiative interaction between autonomous human and software agents in an environment where the agents may play a variety of roles, have varying capabilities and availability (time and resource limitations), and are constrained by the authorities they have been given. Crucially, many of these aspects will be dynamically determined via agent interactions within the agent organization.

We will contribute towards the development of agent interoperability standards, develop prototype tools, and ensure that the concepts are effectively taken up by providing readily available demonstrations of the concepts in a simple framework. The work plan will seek to ensure that the project team is able to engage with those organizations and individuals capable of transferring the concepts and technology into productive military use.

Our proposal is focussed on Team Coordination. To address the creation of agents that can play roles as team members, we will investigate both the agents that can fit into teams and the organizational structure of the teams themselves. It is possible for such a "team player" agent to act as a controller, maintaining coordination with further computational resources, agents or humans. Team coordination will play a central role in agent-based organizations, including cases that involve specialized human agents. Moreover, team structures will continue to be useful for software agents, just as they have continued to be useful for humans, even when there are large numbers of self-organizing adaptive agents.

Our work on "Team Player Agents" will begin by using a particular team structure which we believe is well-suited to cooperative planning and control and to the envisaged demonstrations of the team coordination approach in a suitable challenge problem domain. However, we believe that general principles will emerge from this structure.

Our approach allows agents to take initiative within the authorities or autonomy they have been allowed. This can be dynamically and precisely defined. We will employ a set of related models shared by agents to underpin purposeful coordination. The models provide:

- **Shared Task Model** – a mixed initiative model of the cooperative planning and control tasks as "mutually constraining the space of behavior".
- **Shared Plan Model** – a rich plan representation using a common constraint model of activity.
- **Shared Space of Options** – explicit inter-agent option management.
- **Shared Model of Agent Processing Capabilities** – as handlers of outstanding issues.
- **Shared Understanding of Authority** – management of the authority to do work (to handle issues).

In particular, the work will explore authority management between agents, and the content of such communication. This central aspect makes use of the other models shared by the agents. A workflow management approach will be the basis for each agent's understanding of its tasking environment. Each agent will have a set of agendas reflecting its links to other agents, and a set of outstanding issues within its problem solving state for the range of options it is working on.

Explicit management of inter-agent process coordination will be used as a basis for the creation of prototype configurable process state visualization and management control tools – which we refer to as "Intelligent Process Panels" – I-P^2. Hierarchical abstraction will be employed to ensure communication at an appropriate level to the agents and humans involved. The panels will be used for experiments with a variety of team coordination aids. Earlier work on a "COA Evaluation Matrix" interface to a web-accessible O-Plan Planner Agent, and on an Air Campaign Planning "Process Panel" (ACP^3) within ARPI's TIE 97-1, will be used as a basis for this work.

The project's deliverables will include agent control concepts, re-usable shared models of agent organizations suited to cooperative planning and control, re-usable prototype tools and packaged web-accessible demonstrations on chosen challenge problems.

The research could demonstrate the following benefits for future military systems involving teams of human and software agents:

- Utility as a unifying principle for inter-agent and tool coordination and for user interface design.
- Management of "to do" lists for software agents and humans.
- Allowance for configuration or re-configuration of process flow to suit user needs and to respond to circumstances.
- Integration basis for future inter-agent mixed initiative planning framework for co-operative problem solving, intelligent workflow, distributed event handling and process management.
- Looking towards future active Intelligent Task and Option Management rather than simple workflow step enactment.

The need for aids to support team coordination becomes more important as the scenario environment gets more ambitious. It becomes vital when multiple options are being concurrently developed by teams in which the members are changing over time (such as in hand overs and shift working).

Our approach is compatible with the "Model–Viewer–Controller" style of systems integration architecture being promoted by the JTF ATD Reference Architecture and its successors leading up to the AITS Reference Architecture. We are promoting the use of explicit management of intelligent authorizable agents within a task-driven environment as a basis for mixed initiative interaction supportable by the controller in the AITS architecture and related systems in use in JFACC and a range of DARPA-sponsored programs. The research proposed will allow us to engage with the communities developing standards for the architecture, for the representations used within it, and for the component technologies that it will integrate.

We intend to develop links with other CoABS program participants to develop a suitable challenge problem or scenario. This would form a concrete example to refine our own concepts and ensure that they are usable in a variety of agent integration architectures.

It has been suggested by the DARPA CoABS Program management that a scenario relating to Coalition Operations could form a suitable basis for our work.

Innovative Claims

We will explore a number of interrelated ways to maintain effective control of agent systems and to enable agents in a team to cooperate purposefully to solve problems requiring multiple skills and needing multiple solution options. We will investigate both the agents that can fit into teams and the organizational structure of the teams themselves, bringing together work in AI Planning with recent advances in process management, workflow and organizational modeling.

We will address a problem typical of that found in many realistic military planning environments. A mixture of human specialists and computational agents will work together to explore multiple plan options to deal with a dynamically evolving situation.

To address this scenario we will create novel ways to represent and use authority among agents, and we will explore the contents of such authority descriptions. The capabilities and roles of the agents involved and their relationships to one another will be modeled. AI planning technology will be used to assist agents in planning their own or a team's cooperative processes, in seeking sub-processes to respond to dynamically assigned tasks, and to adapt to changing circumstances, or to problems that emerge with a previously selected process plan.

Our approach to effective inter-agent cooperation can be outlined as follows:

- Inter-related, shared models provide a sound basis for agent cooperation.
- Authority management among the agents in a team is an effective way to control behavior and prevent

conflicts.

- Clear definition of the roles and capabilities of humans and systems assists in agent communication and task-driven cooperation.
- Team activity monitoring and visualization via configurable "Planning Process Panels" helps human users to control and understand agent system behavior.
- Team activity and individual agent workflow planning and re-planning allow agents and teams to better manage their own activities.
- Agent agendas support inter-agent tasking, keeping track of agent activity, and can support explanation of what's happening.

A possible extension would be to record and make use of process design decision rationale, alongside the complementary dependency and causal rationale which is traditionally found in process and activity models. Communication of this knowledge would be used to support decision-making by promoting a shared understanding of each agent's objectives and the outstanding issues along with the options being considered to address them. It would connect agents to their specific roles and responsibilities in the planning and control processes.

Vision

Agent Organization

Our proposal addresses the coordination of work between members of a team of specialist human and software agents to realize the benefits of an agent approach to defining an overall system which can act purposefully. Our approach sets in place an *organization* of agents much as is found in normal enterprises. Clear tasking, an understanding of authority management, and shared models of how agents perform their functions and interact are the basic building blocks to be explored in the proposed research.

An organization contains a set of agents. These agents may be individual humans or system tools, or they may act as a representative, or a coordinator, of a team of other agents[1]. Each plays one or more roles which in some cases may be dynamically determined. Depending on their role in the organization, some are able to accept tasking from other agents or from outside the organization. Each agent has a set of capabilities which determine the functions it can perform[2] and the quality and responsiveness it can achieve in the performance of those functions. The availability of agents in time and space, and the cost of using those agents, will be variable (for both human and software agents).

Cooperative Planning and Execution Tasks

We will show (i) the selection of suitable processes for performing planning and control tasks that involve dynamic configuration of an agent team, (ii) team cooperation on the task, and (iii) the responsive reorganization of the processes and team as circumstances vary. We will demonstrate that the inter-agent models enable robust and purposeful cooperative exploration of multiple options for behavior in the domain for which the plans are being created.

Each agent will maintain its own unique perspective on the shared tasks being undertaken by the organization. We will create "Planning Process Panels" to act as a human agent's window on the status of the organization or agent team's work. Such panels will also support an agent in calling up sub-processes suited to the particular tasks in hand, and the resources that can be used to perform these tasks at any time – using an intelligent workflow planning aid [*Drabble et. al., 1997a*]. Dynamic reconfigurability to repair processes "on-the-fly" [*Drabble et. al., 1997b; Reece et. al., 1993*] will be incorporated. For agents with suitable roles and authority, this same panel will act as a "Process Control Panel". Each panel will be customizable to reflect the level of abstraction which the agent or user wishes to see, the role they play, the permissions they have to task other agents, the capabilities they have available, and so on.

Inter-Agent Task Driven Cooperation

We have elected to work with just two human agents in order to simplify the research infrastructure and tools needed to demonstrate the concepts to be explored during the research[3]. We will use a scenario of two human agents and a collection of software agents (some of which may have direct user interfaces involving the two "agent" humans or others interacting in this limited

[1] An example of such a coordinator might be a "room representative" in Mitre's Collaborative Virtual Workspace.

[2] These could also be dynamically determined, but in the proposed research they will be considered fixed to make the areas to be explored tractable.

[3] This is a limitation for the demonstrations and evaluation experiments only. The team coordination framework to be explored will allow for multiple human or software agents.

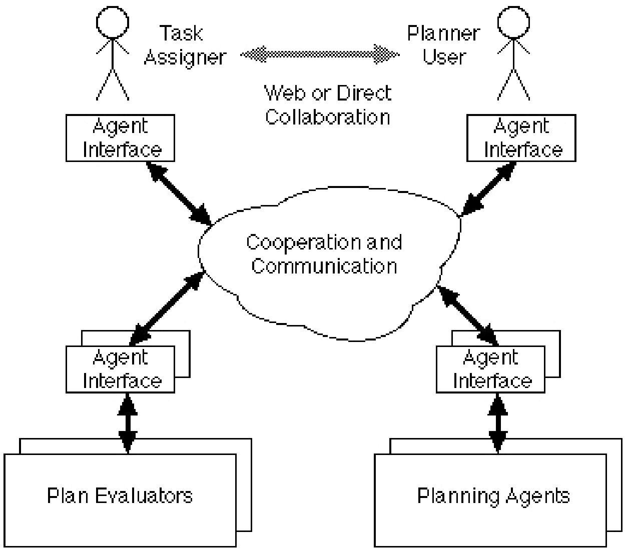

Figure 1: User Collaboration and Agent Cooperation

way). This is shown in Figure 1. The human and system agents will have a range of different roles and have differing responsibility (authority) and technical capabilities to perform their joint tasks. The two humans would play roles equivalent to a commander and a military staff member. The system tools and agents would provide automated or semi-automated planning, plan evaluation and execution support. We envisage that between 5 and 10 such automated agents will be used, dependent on the scenarios chosen. The human and software agents within the environment would interface via a suitable agent "wrapper" through the communications network (local or wide area). However, the project will take as an assumption that direct human-to-human communication is possible via other routes (via collaboration software or directly), and that some tools may be used directly by users outside of the scope of the communicated agent-to-agent interactions. This is a realistic situation in many organizations. We will also assume that some of the software agents "wrap" legacy applications and cannot therefore be assumed to be able to provide processing state feedback unless the agents themselves communicate such information [4].

While only a small number of agents are involved, the rich possible interactions between groups of such agents will be represented in the work. The two humans will communicate using readily available (and low cost) collaboration tools (initially conferencing tools within Netscape 4 or Internet Explorer 4). However, we believe that the concepts we will explore will be relevant to more than two humans, and will be compatible with more sophisticated collaboration tools such as MCC's Collaboration Management Infrastructure or Mitre's Collaborative Virtual Workspace. On TIEs within and beyond the ABS program, we hope to work with these groups to achieve a synergy of approaches to add task-driven cooperation based on shared models into the emerging collaboration tools.

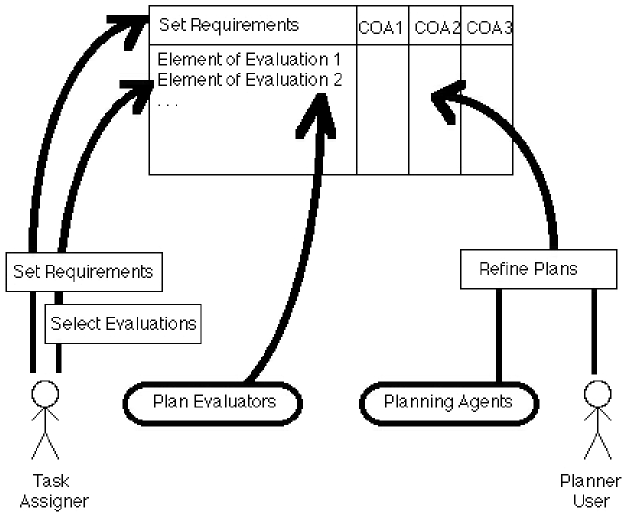

Figure 2: Roles of the Human and Software Agents

The human agents and some of the software agents would be truly mixed initiative, and able to take the initiative within their (dynamically assigned) authorities and capabilities in the collaboration. Figure 2 illustrates the individual roles that could be assigned to the agents for a shared model of the task of providing a set of alternative Courses Of Action (COAs) for some task and providing comparative evaluations of those COAs. In this scenario, there is a "Task Assigning" commander, a human planning staff member, and a set of cooperating computer-based planning and plan evaluation agents. The human and software agents can cooperate by having a shared understanding of their roles and authorities in the development of the "COA Matrix".

Technical Rationale

This proposal addresses the control of agents that can be members of a team and can coordinate their efforts with other human and software agents. A shared view of the cooperation and authority relationships between agents and the roles of those agents is used to guide or

[4] These are assumptions used in an approach the project team has taken before in realistic enterprise process support for a number of practical applications [*Fraser and Tate,1995*]

constrain agent activity and to ensure that it is productively channeled towards task achievement. The approach involves work in a number of related areas which are described in the following sections:

Shared Models and Standards

We will investigate the use of shared models for task-directed communication between human and software agents who are jointly exploring a range of alternative options for activity (COAs). Activity planning is chosen as a demonstration area for the proposed research in order to make maximum use of previous research assets of the project team. However, the approach is intended to apply across a range of other types of co-operative tasks.

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

1. **Shared Task Model** – a mixed initiative model of the cooperative planning task as "mutually constraining the space of behavior" [*Tate, 1994*].
2. **Shared Plan Model** – a rich plan representation using a common constraint model of activity (<I-N-CA>) - based on our earlier work on <I-N-OVA> [*Tate, 1996*].
3. **Shared Space of Options** – explicit option management [*Tate et. al., 1998*].
4. **Shared Model of Agent Processing Capabilities** – handlers for issues, agent functional capabilities and constraint managers [*Tate et. al., 1996*].
5. **Shared Understanding of Authority** – management of the authority to do work (to handle issues) which may take into account options, phases and levels of abstraction of the work processes [*Tate, 1993; Jarvis et. al., 1999*].

The proposed research will develop all aspects of these important and related shared models, but will particularly seek to make significant progress on the authority management aspect. This central aspect in turn involves communication about all the other shared models.

Using these shared views of the roles and function of various users and systems involved in a command, planning and control environment, we will demonstrate a number of planning agents being used to support mixed initiative task specification and plan refinement over the world wide web.

I-X - I-Technology

Previous work on O-Plan (Open Planning Architecture [*Currie and Tate, 1991; Tate et. al., 1994*]) has explored mixed initiative planning methods and their application to realistic problems in logistics, air campaign planning and crisis action response [*Tate et. al., 1996*] and Non-Combatant Evacuation Operations (NEOs) [*Drabble et. al., 1995; Tate, 1994*]. 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*]. A demonstration environment has been created which uses the World Wide Web to allow users access from any web browser to an O-Plan planning agent (see Figure 3 [*Tate et. al., 1998*]).

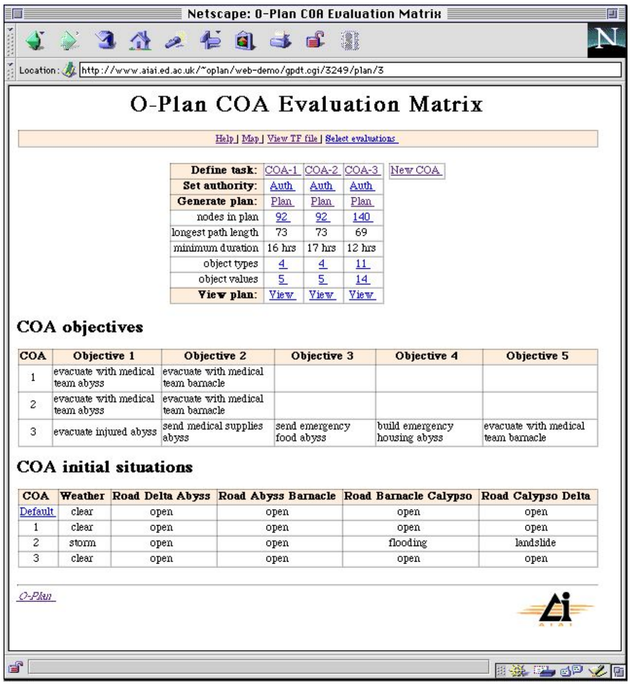

Figure 3: O-Plan Pacifica "COA Matrix" Demonstration

A number of aspects of the O-Plan Technology base have become more important to our work in explicit process management and workflow support for communities of agents. These include a process editor and its library, process and workflow control panels, and user interface technology for displaying the status of cooperation in a mixed initiative environment. We are proposing to drawn on the existing O-Plan technology and systems to further develop some of these aspects, to make the underlying concepts clearer, and to better package them for use by others. We also plan to implement much of the new systems in Java. We call the systems collectively *I-X - I-Technology* - the "I" standing for "Intelligent" amongst other things. See the web page at `http://www.aiai.ed.ac.uk/project/ix` [*Tate, 1999*] for more details.

We propose to conduct the I- Technology work of the project in such a way that the system and its demonstrations will be available via web delivery either running on Edinburgh servers for demonstration and occasional use, or available for freely accessible download for use by others in an embedded or packaged way in their own systems. *Lightweight* Java-based tools are to be created to support the work[5]. Java-based process/task management control panels will be used to act as a focus for the investigation of suitable views onto the shared models and controls that are helpful to the different roles of the users involved.

While the I-Technology package prototyping environment includes modules for each necessary part of the system, it is not proposed that each be developed to the same level of detail. The eventual aim is that the concepts and modules simply act as a framework to show how shared protocols and models can be interchanged between tools that perform the specific roles of each module included. Importantly, it is also not anticipated that the tools to be developed act as an alternative to system integration frameworks (such as the JTF ATD Reference Architecture or its successors, the GCCS LES Framework and the AITS Reference Architecture). Rather, the research proposed and the individual modules will act as guides to next-generation possibilities for parts of these environments. It is anticipated that this will allow convergence with next generation COTS and GOTS tools and protocols – which by then may have been influenced by the proposed research.

There will be 6 parts in the I-Technology Package:

1. I-PE – a process editor that can be used to create and maintain the process models used for inter-agent activity. Import and export to and from other formats will be supported. The eventual aim is for the development of the Process Editor to be taken much further by incorporating research in process rationale capture and explicit design process issue recording in such a way that it can be used in all other aspects of this research.

2. I-PL – an interface to a process librarian for storing and retrieving process models for use elsewhere in the agent environment. It is anticipated that this will be quite simple in its implementation. The research on this aspect will rather concentrate on the interfaces to such a library as it is anticipated that object-oriented standard librarians will be created by others in the DARPA and military program community.

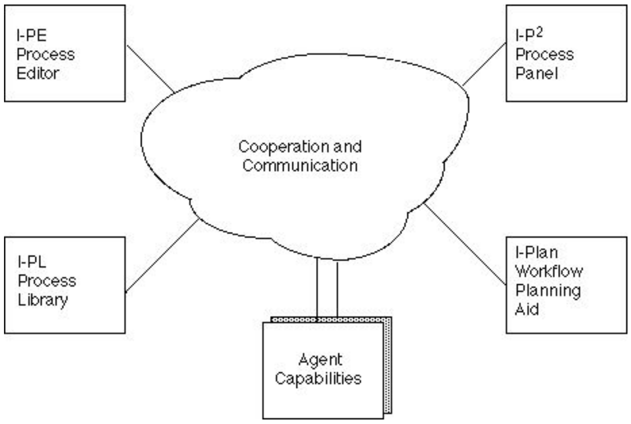

Figure 4: Components of the I-Technology Package

3. I-P^2 – an intelligent workflow controller technology base that will allow process status visualization or process control panels to be created for individual human agents. The technology will support alternative views of the status of the process, the agents involved and the process products that are being created or modified. Our emerging model of inter-agent authority management in a mixed initiative environment will be supported. The technology will be designed to be incorporated alongside COTS or GOTS workflow technologies, or to be an experimental base for novel explorations of new workflow control principles.

4. I-Plan as a workflow planning aid – as a demonstration of a specialist process an agent management aid, an O-Plan agent (using existing O-Plan technology), will be incorporated. This will be based on earlier work on using O-Plan as an ACP Process Workflow Planning Aid, performed in a "qualifier" for an earlier proposed demonstrator (IFD-5) under ARPI [*Drabble et. at., 1997*].

5. I-Plan as a system component or agent capability – as a demonstration of an automated authorizable mixed initiative planning agent within the framework. It is anticipated that cooperation with others under the CoABS program will lead to the characterization of their technologies in such a way that

[5] In a similar way to AIAI's earlier C++ based HARDY Meta-CASE and Model-to-Web publishing tools released freely and in widespread use – e.g. for ACP Process Modeling under ARPI – see `http://arpi.isx.com/community/doma.html`.

these can be defined as such capabilities. Our work to define a Process Panel for TIE 97-1 under ARPI has allowed us to perform some initial work on such a task involving 7 DARPA-sponsored research tools.

6. A set of demonstration and briefing materials to make use of the above Technology Package. This will allow for the incorporation of the domain demonstration materials if they are not considered sensitive.

As mentioned earlier, each part of this framework will not be developed to the same level of detail. Some (such as the Process Librarian) will be simple versions that act as stubs for more comprehensive and standardized services (such as CORBA-compliant object systems) that will perform these roles in future military systems. We will use standard basic protocols at a transport layer that are already widely adopted (KQML is likely, though this would be easily reconfigured).

We will particularly concentrate on the creation of the configurable Intelligent Process Panel Technology (I-P^2), and a Process Editor that can support configurable and adaptable process models which can record process causal and dependency rationale (and with an optional work package – process design rationale [*Polyak and Tate, 1998*]).

I-P^2 – Process Panel User Interface Technology

Process status visualization tools have already been explored by us as part of ARPI's TIE 97-1. This used user interface and process display concepts that have emerged from O-Plan research on mixed initiative multi-agent workflow to support crisis management planning and control [*Tate et. al., 1998*]. An example Air Campaign Planning Process Panel (ACP^3) proposed for monitoring agent interactions (human and 7 different systems communicating both manually and via KQML) is shown in Figure 5. Such tools have been crafted specifically for the demonstration needs to date and have used a mixture of CGI scripting, a Lisp-based HTTP server, and hand crafted Java code. The experience gained will now enable us to package the concepts for more flexible re-use and configurability – using a wider range of process status displays, e.g.,

- process perspective
- process product or results perspective
- task/objective status perspective
- option under exploration perspective
- agent or system tool perspective

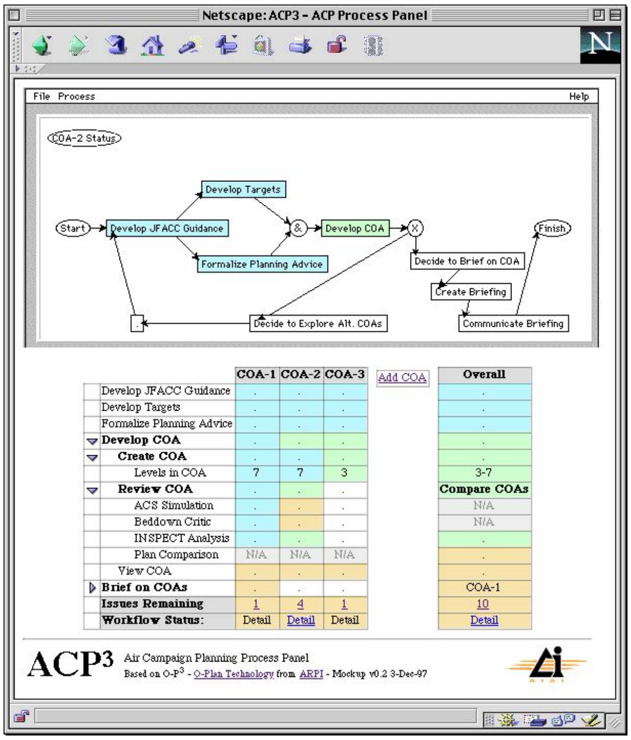

Figure 5: ACP^3 – Air Campaign Planning Process Panel from ARPI TIE 97-1 and EFX-98

The design incorporates facilities to enable such a panel to act as an *Process Control Panel* for an agent to interact in a mixed initiative way with others. The model of Mixed Initiative Interaction for a Planning and Control environment that can be supported by the approach is the *mutual constraining of behavior by refining a set of alternative partial plans*. Humans and software systems can work in harmony through employing a shared view of their roles as being to constrain the space of admitted behavior. Workflow ordering and priorities can be applied to impose specific styles of authority to plan within the system. Two extremes are possible: user driven plan expansion followed by system "filling-in" of details; or fully automatic system-driven planning (with perhaps occasional appeals to an user to take predefined decisions). In more practical use, we envisage a mixed initiative form of interaction in which humans and software systems proceed by mutually constraining the plan using their own areas of strength.

I-PE – Multi-Perspective and Issue-Based Process Editor

A process editor will be created (as a Java-based web accessible facility) to support the initial creation, and subsequent maintenance of, the shared models of the domain which are required.

Issue-based reasoning will be used to capture the state of modeling especially when multiple domain experts and modelers are involved and there are inconsistencies in the information available. We call this approach *Multi-Perspective Modeling*. An initial methodology for planning domain modeling (the *O-Plan TF Method*) will be developed further. We will extend this process/plan domain modeling methodology to allow for the design of structures for teams of agents, involving such things as establishing roles and authorities. Simple tools currently being used to validate a domain model when loaded into a planner (as in the *O-Plan TF Compiler*) will be elaborated to provide greater support during domain construction and maintenance. Import and export facilities from and to emerging standard representations such as IDEF0, PIF, NIST PSL, HPKB's PSM Language and SPAR will be provided in a robust way.

As shown in Figure 6, AIAI has been moving towards an approach to support modeling where we use a range of methods, and create a single shared model based on one or more small generic ontologies or conceptual models. The terminology in the model is anchored in a lexicon, which itself can be developed during modeling. Modeling support is provided by the creation of an *agenda* of outstanding modeling *issues*. These will eventually be handled by seeking *issue handlers* which will select appropriate methods or tools to help the modeler address the outstanding issues.

AIAI has done practical modeling work using this technique (e.g. on the ARPI ISAT Project Air Campaign Planning Process and with the UK Search and Rescue Coordination Center). We have also used Multi-Perspective Approaches to reasoning about plans in the O-Plan work under ARPI.

Typical multiple perspective modeling approaches used by us have created a single ontologically underpinned model using our work on ontologies for processes and for enterprises and utilizing several modeling methods from Europe (such as CommonKADS - itself combining a number of perspectives), the UK (such as Role/Activity Diagrams), and the USA (such as IDEF-3). Multiple-view or multiple perspective modeling is viewed in the software engineering and requirements capture communities as a valuable technique.

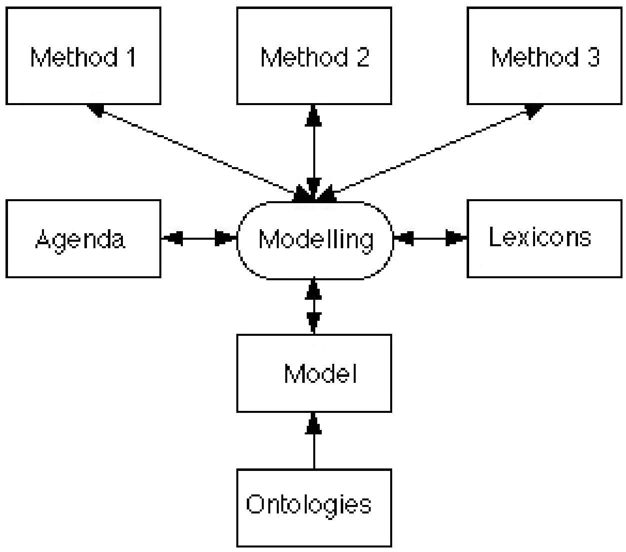

Figure 6: Multi-Perspective Modeling Approach

Inter-Agent Rationale Communication

In the proposed optional work package (WP7), we wish to renew work (which has languished since the mid 1980s) on the capture of design and decision rationale in plans and process descriptions [*Daniel, 1983*]. In this earlier work, decisions were captured during planning via decision graphs. These graphs were used to maintain the dependencies between a planning agent's decisions and the resultant structure of the plan. This facilitated intelligent re-planning since failures in plan execution could be traced back to the original decision. Only plan elements that depended on this decision were removed from the plan. This provided value to automated planning, but applies to human cooperation aspects as well.

This type of knowledge can be used to build a bridge between work on design rationale capture and the use of such information in process management and planning. One planning perspective considers a plan to be a specialized type of design [*Tate, 1996a*]. Given this viewpoint, we can incorporate design rationale research by managing a plan or plan fragment as a designed artifact. A rich corpus of methods, representations, notations, etc. from the design research community [*Moran and Carroll, 1996*] can be drawn upon to effectively and efficiently utilize this knowledge.

This decision rationale is complementary to the dependencies and causal rationale that is maintained in most AI planners today. We recently completed a review of these aspects of planning rationale [*Polyak and Tate, 1998*]. In this review, we point out the fact that decision rationale in AI planning is still in its infancy and we report on some of our initial work in incorporating a design rationale approach. The design rationale will be utilized across the proposal framework, and its utility will be evaluated.

Evaluation Plan

The overall aim of this work is to have human and software agents working together to construct shared plans and to explore a range of alternative plan options. Our proposal is that the coordination of the effort will be facilitated by the twin notions of authority and user roles. This is analogous to the authority models and user roles adopted by human teams to ensure that members work together efficiently while cooperating and negotiating on areas of mutual interest.

We propose four related strands of evaluation work: experimentation with authority models, tackling the ABS program challenge problems that are appropriate to our work, qualitative comparison of each of our demonstration deliverables [D1, D2 and D3] against our proposed vision; and extension of our existing software instrumentation and testing to deal with plans requiring the cooperation of more than one planning agent. Each of these strands is described in the sections which follow.

Authority Models and User Roles

Our team coordination approach allows multiple planning agents, plan evaluation agents and human users to work together to explore a range of courses of action to be used in response to a developing situation.

We will experiment with different authority models to see which features produce the best results for various problems, including challenge problems selected from the set proposed for the ABS program. A hypothesis for this has been provided through task analyses in a range of realistic problems (see [*Tate, 1994*]). We will experiment with different styles of interaction between the computer and human agents, based on the explicit and shared notion of an authority to perform some action (such as the authority to plan) and with the human agents being assigned user roles (such as task assignment agent), with appropriate interfaces being provided to the human agents.

There are a number of trade-offs to be explored in this area. A strict model of authority would result in no conflict between agents, but might also result in a low quality solution (due to not allowing agents to do things that might improve the plan), under-use of agents capabilities, or tedious and trivial user interaction. The opposite model, of allowing all agents to do everything they are capable of with no agent having authority over another, could result in an unnecessary amount of conflict, including deadlocked situations. Any model of authority proposed needs to be tested for the quality of the solution, the time taken to reach that solution, the amount of time spent in unproductive conflict, perceived support to some types of user, and so on.

Coalition Challenge Problems

Well-designed challenge problems are an important part of our evaluation strategy. It is our wish to select a challenge problem that requires support to Coalition military operations in areas such as Operations Other Than War (OOTW). We are already familiar with operations such as Non-combatant Evacuation Operations (NEOs), so these might be one areas to explore with potential collaborators.

In tackling these challenge problems, we hope to demonstrate two important aspects of the proposed work. The first is the basic technical capability of the software planning and plan evaluation agents to work together to construct and criticize a plan in these complex domains. The second is the ability of the system of agents as a whole, including the human agents, to effectively and efficiently explore a space of possible plan options to respond to an emerging situation. This will involve experimentation using simulated user trials at AIAI, with the human agents being appropriately briefed and then given a scenario to respond to. The subjects for these trials would be members of AIAI staff and students from the Division of Infomatics. These trials would not only give a measure of the success in tackling the challenge problems, but would also give valuable feedback for the system developers on (for example) the style of graphical interaction with the system (i.e. the interfaces provided for the human agents to interact with the other agents present).

Comparison with Proposed Vision

Each of our demonstration deliverables represents a step towards our proposed vision. It is therefore essential to compare each one against our final goal to see both what we have achieved so far and which aspects of the vision are easier or harder than we expected. In doing this, we will also try to determine which technological features, such as the use of clear and explicit user roles or the specific types of authority description, are crucial in supporting the various features of the vision, such as agent organization and inter-agent task driven cooperation.

Software Instrumentation and Testing

The current O-Plan system accommodates a number of facilities to allow for development testing, evaluation experimentation, and software instrumentation. An auto-testing capability allows a package of domain descriptions and tasks to be provided to the planner, one or more plans to be generated for each task in the selected domains, and the results to be compared to previously generated plans. A "sanity checker" for plans automates much of the inspection that would otherwise be needed to check that the plans generated are valid against a set of criteria. These facilities allow for repeatable testing of an O-Plan Planning Agent against test suites and provide "hooks" for extending the types of automatic testing that can be done on the plans generated by O-Plan.

For our I-X and I-Plan systems, we propose to use this test suite on our individual planning agents and construct an additional test suite of plans to be constructed by a number of planning agents working together. Since a multi-agent system is, in essence, a complex piece of software, it is even more important to apply these testing techniques to the system of software agents as a whole.

The compiler for the O-Plan Task Formalism (TF) domain description language already provides good levels of diagnostic support to domain writers and will be extended during the proposed project to provide feedback through I-PE (Process Editor) to domain experts who are describing processes and plans. The Task Formalism Compiler can be run in a "checking mode" separately to the O-Plan planner to provide such information for use in the Process Editor.

Acknowledgements

The O-Plan project is sponsored by the Defence Advanced Research Projects Agency (DARPA) and the U.S. Air Force Research Laboratory (AFRL) under grant number F30602-99-1-0024. Mr. Wayne Bosco at AFRL (Rome) monitors the O-Plan project. Dr. Doug Dyer, Dr. Mark McHenry and Dr. Jim Hendler at DARPA, and Mr. Ken Sharpe of SAIC also guide the project. The U.S. Government and the University of Edinburgh are authorized to reproduce and distribute reprints for their purposes notwithstanding any copyright annotation hereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either express or implied, of DARPA, AFRL, the US Government or the University of Edinburgh.

References

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

[*Daniel, 1983*] Daniel, L., Planning and Operations Research, Artificial Intelligence: Tools, Techniques and Applications, Harper and Row, New York, 1983.

[*Drabble et. al., 1997a*] Drabble, B., Tate, A. and Dalton, J., ACP Process Management: O-Plan IFD-5 Qualifier, O-Plan Technical Report, ARPA-RL/O-Plan/TR/3, AIAI, University of Edinburgh, January 1997.

[*Drabble et. al., 1997b*] Drabble, B., Dalton, J. and Tate, A., Repairing Plans On-the-fly, Proceedings of the NASA Workshop on Planning and Scheduling for Space, October 1997, Oxnard, CA, USA. October 1997.

[*Drabble et. al., 1995*] 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.

[*Fraser and Tate, 1995*] 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.

[*Jarvis et. al., 1999*] Jarvis, P., Stader, J. Macintosh, A., Moore. J., and Chung, P (1999) What Right Do You Have to Do That? Infusing Adaptive Workflow Technology with Knowledge about the Organisational and Authority Context of a Task, in Proceedings of the First International Conference on Enterprise Information Systems (ICEIS-99), Setubal, Portugal. `http://www.aiai.ed.ac.uk/project/tbpm`

[*Moran and Carroll, 1996*] Moran, T.P. and J.M. Carroll, J.M. (eds.) Design Rationale: Concepts, Techniques, and Use, Lawrence Erlbaum Associates, Mahwah, New Jersey, 1996.

[*Polyak and Tate, 1998*] Polyak, S. and Tate, A., Rationale in Planning: Causality, Dependencies and Decisions, Knowledge Engineering Review, Cambridge University Press, 1998.

[*Reece et. al., 1993 ??????*] Reece, G., Tate, A., Brown, D. and Hoffman, M., The PRECiS Environment, Paper to the ARPA/Rome Laboratory Planning Initiative Workshop at the National Conference on Artificial Intelligence (AAAI-93), Washington D.C., USA. ARPI Report ARPA-RL/CPE/Version 1, ISX Corporation. Also available as AIAI-TR-140, AIAI, University of Edinburgh, August 1993.

http://www.aiai.ed.ac.uk/~oplan/pacifica

[*Tate, 1993*] Tate, A., 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, 1994*] Tate, A., Mixed Initiative Planning in O-Plan2, Proceedings of the ARPA/Rome Laboratory Planning Initiative Workshop at Tucson, Arizona, USA, (ed. M. Burstein), Morgan-Kaufmann, 1994.

[*Tate, 1996a*] Tate, A., Towards a Plan Ontology, AI*IA Notiziqe (Quarterly Publication of the Associazione Italiana per l'Intelligenza Artificiale), Special Issue on Aspects of Planning Research, 9,, 1, pp.19-26, March 1996.

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

[*Tate, 1999*] Tate, A., I-X: Technology for Intelligent Agents and Tools, Web Page, March 1999. http://www.aiai.ed.ac.uk/project/ix

[*Tate et. al., 1998*] Tate, A., Dalton, J. and Levine, J., Generation of Multiple Qualitatively Different Plan Options, Submitted to the Fourth International Conference on AI Planning Systems (AIPS-98), Pittsburgh, PA, USA. Available as O-Plan Technical Paper DARPA-AFRL/O-Plan/TP/44, June 1998.

[*Tate et. al., 1996*] Tate, A., Drabble, B. and Dalton, J., A Knowledge-Based Planner and its Application to Logistics, in Advanced Planning Technology, pp. 259-266, (Tate, A., ed.), AAAI Press, May 1996.

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