

I-Ex : Intelligent Extreme Expedition Support

Stephen POTTER¹, Austin TATE¹ and Rob MILNE²

¹*Artificial Intelligence Applications Institute, University of Edinburgh, UK.*

²*Intelligent Applications Ltd., UK.*

<http://www.aiai.ed.ac.uk/project/ix/>

Abstract. The aim of the I-X research programme is to provide a general framework for performing *mixed-initiative synthesis tasks*, along with a set of tools that supports its use. This framework arises from and builds upon seminal work at the University of Edinburgh in the field of Artificial Intelligence planning. In this paper we describe the framework and tools, before describing the application of I-X to the task of planning and coordinating expeditions to remote locations – such as an attempt on Everest. We call this application *I-Ex*.

Keywords. Synthesis tasks, planning, coordination, intelligible messaging.

1. Introduction

An expedition such as an attempt on Everest obviously requires a lot of coordination and planning, first during the preparatory stages, later during the expedition itself, and finally in concluding the effort. Where previously this would primarily be done ‘manually’, the rise of the internet and the World-Wide Web allows the introduction of greater amounts of computer support and wider access to information from a variety of sources.

The aim of the I-Ex project is to attempt to provide a means to integrate and structure the expedition activity with information from disparate sources and to allow access to intelligent technologies where appropriate. This computer support builds on the more general I-X architecture, which has been developed within the Artificial Intelligence Applications Institute at the University of Edinburgh. Simply put, the purpose of I-X is to provide intelligent computer support to people who are performing some task together. This task may be, say, designing a car, or performing an emergency evacuation, or coordinating an attempt on Everest. The I-X architecture supplies a framework that encourages a methodological approach to the task, based on cycles of issue-raising, handling and resolution. This is underpinned by the ‘intelligible messaging’ of issues, activities and other information among the agents in the system. This allows users to manipulate, transform and transmit information in context-sensitive ways that continually aim to move the process forward. I-X’s foundations in AI planning technologies allow the invocation of automated planners that are able to suggest potential plans for achieving task sub-goals, while its open architecture facilitates the use of external computational services and information.

In this paper we first describe in general terms the I-X approach, and the technology and tools that implement it. Then we discuss the I-Ex project as an

application of I-X, describing the necessary preparations for deployment along with the deployment itself, and the lessons learned from these. Finally we, draw some concluding remarks about the future of task-achieving architectures of this sort.

2. The I-X Approach

I-X is a research programme with a number of different aspects intended to allow humans and computer systems to cooperate in the creation or modification of some product or products such as documents, plans, designs or physical entities – that is, it supports *mixed-initiative synthesis tasks* [1].

The I-X approach involves the use of shared models for task-directed cooperation between agents which are jointly exploring (via some process, which may have been dynamically determined) a range of alternative options for the synthesis of an artifact (termed a *product*) such as a design or a plan. The <I-N-C-A> (Issues-Nodes-Constraints-Annotations) ontology [1] is used to represent a specific product as a set of constraints on the space of all possible products in an application domain. This ontology can be used to describe the product requirements or specification and the emerging description of the product itself. It can also describe the processes involved.

I-X provides a modular systems integration architecture that mirrors the underlying <I-N-C-A> ontology. It provides a “Model-Viewer-Controller” style of architecture. Plug-in components for Issue Handlers, Activity Performers, Constraint Managers, I/O Handlers and Process or Product Viewers allow for specific I-X systems to be tailored for specific tasks. The I-X approach draws on earlier work on Nonlin [2], O-Plan [3-7], Optimum-AIV [8,9], <I-N-OVA> [10,11] and the Enterprise Project [12,13] but seeks to make the framework generic and to clarify terminology, simplify the approach taken, and to increase the re-usability and applicability of the core ideas.

2.1. I-X Process Panels

I-X Process Panels (I-P²) are used to support individual users who are carrying out processes and responding to events in a cooperative working environment – they are the primary means by which useful functionality is delivered to users (Figure 1). A panel supports the tracking of personal or group issues, the planning and execution of activities and the checking of constraints. It can be used to communicate with other panels and any other known services, agents and other co-operative working support tools to form a network of activity and process support in an organisation. This communication is achieved via a range of strategies which vary from simple direct internet ports, through custom name server and brokering systems, to comprehensive, secure, agent communications platforms such as the CoABS Grid [14] and KAoS [15].

I-X Process Panels and their predecessors, the Open Planning Process Panels (O-P³) [7], have been used in a number of prototype and deployed application areas including:

- Air campaign planning process support [5];
- Non-combatant evacuation operations [6];
- US Army small unit operations for hostage rescue [16];
- Coalition and multi-national forces command and control [17,18];
- Search & rescue coordination [19,20];

- Help desks;
- Unmanned autonomous vehicle command and control;
- Cooperative working between e-Scientists [21].

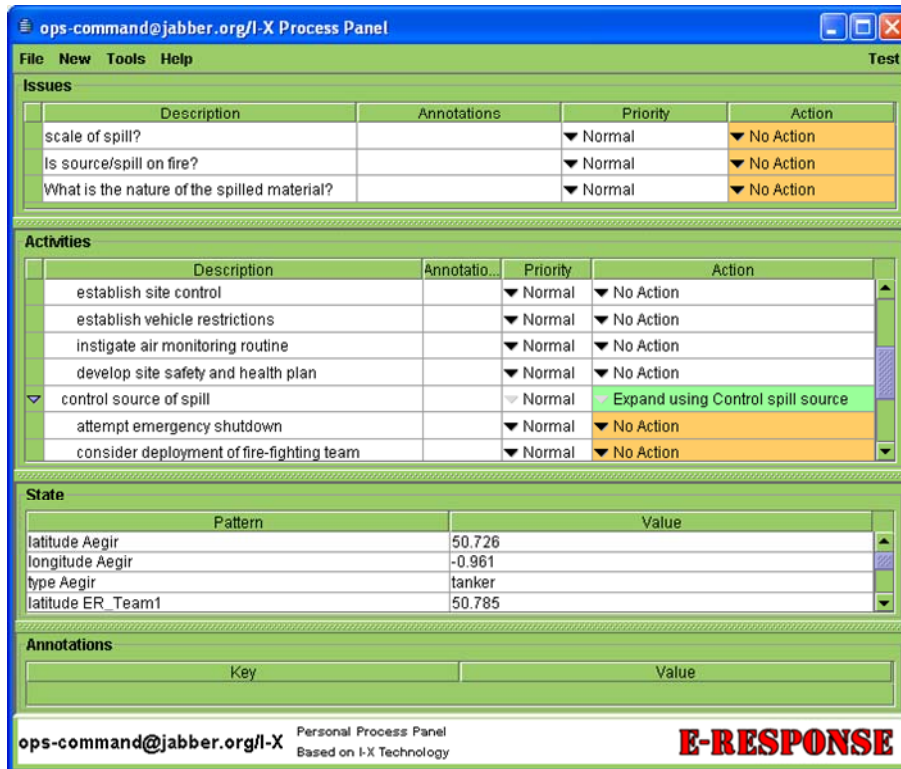


Figure 1. An I-X Process Panel (I-P²), shown here providing an interface on to the response to a (simulated) marine oil spill, with *Issues* raised, the *Activities* currently in hand and the current *State*.

An I-X Process Panel can be seen, at its simplest, as an intelligent ‘to-do’ list for its user (the nature of the ‘intelligence’ of a panel will be discussed below in section 2.2). However, and especially when used in conjunction with other users’ panels, it can become a workflow, reporting and messaging ‘catch all’, allowing the coordination of activity, and hence facilitating more successful and efficient collaborations. A panel corresponds to its user’s ‘view’ onto the current activity, through the presentation of the current items (from the user’s perspective) of each of the four sets of entities comprising the <I-N-C-A> model. The contents of these sets, along with the current context and state of the collaboration, are used to generate dynamically the support options the tool provides. For example, associated with a particular activity might be suggestions for performing it using known procedural expansions, for invoking an agent known to offer a corresponding capability, or for delegating the activity to some subordinate agent in the environment.

To summarise, an I-X Process Panel can accept requests to:

- Handle an issue;
- Perform an activity;

- Add a constraint;
- Note an annotation.

The panel allows its user to resolve issues and perform activities, and note the subsequent change in the state of the collaboration, through:

- Manual (user) activity;
- Invoking local computational capabilities (perform);
- Invoking external capabilities (invoke or query/answer);
- Rerouting or delegating to other panels or agents (pass);
- Planning and executing a composite of these strategies (plan or expand).

The panel receives “progress” or “completion” reports and other event-related messages and, where possible, interprets them to help the user:

- Understand the current status of issues, activities and constraints;
- Understand the current world state, especially status of process products;
- Control the interaction with other agents;
- Annotate the various elements of the model.

An I-X Process Panel can cope with partial knowledge and can operate even where little or no pre-built knowledge of the domain or knowledge of other panels and services is available – effectively becoming a simple to-do list and issue-tracking aid in that case.

The ease and freedom with which people use instant messaging tools in a variety of situations, and the realization that, in activity-oriented situations, many of the messages exchanged refer to items that could be described using the <I-N-C-A> ontology (such as particular issues, activities and various types of preferences and constraints), has led to the adoption of *intelligible messaging* for I-X. By intelligible messages, we mean messages that are structured, where appropriate, using <I-N-C-A> constructs; this provides the recipient (which may be a non-human agent) of any message with additional formal information about the content and encourages users to think about structuring their messages in a more formal, task-oriented manner, and allows access to the intelligence underpinning the I-X system. However, the users retain the ability to send ‘normal’ informal instant messages whenever this seems more appropriate – for example, when discussing things outside the scope of the activity. (An additional advantage of this approach is that by describing I-X Process Panels as providing instant messaging augmented with process, activity and task support, prospective users seem more readily able to grasp what the I-X approach can offer them.)

As I-X Process Panels have been developed and used in more cooperative and human-centric applications (such as in support of scientific meeting and group work [21]), intelligible messaging has come to influence the interfaces and become central to our approach. We have incorporated the use of a Jabber [22] communications strategy, which provides for Instant Messaging using XML content through third-party servers. This allows for simpler ‘out of the box’ deployments of the I-X Process Panels, and on a larger scale.

Running alongside and complementing the functionality offered by the process panels, the I-X tool suite provides a range of support and visualization tools (Figure 2).

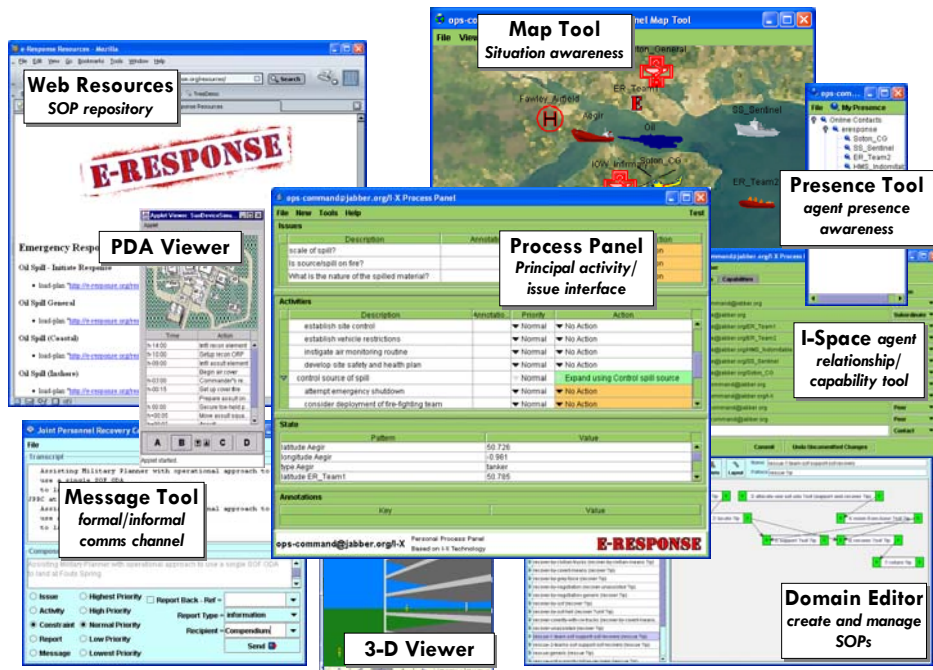


Figure 2. The I-X tool suite.

2.2. I-Plan

The facilities available through the I-X Process Panels include a simple AI planner (I-Plan) used to provide context-sensitive options for the handling of issues, the performance of activities, and the satisfaction of constraints.

For any activity on the panel, an *Action* column shows its current status and, in a drop-down menu, the available options (“activity handlers”) to perform the activity. The status of each activity is indicated by the colour shown in this column (as in Figure 1):

- White indicates that the item is currently not ready for execution (i.e., some temporal ordering, preconditions or other constraints might not be satisfied);
- Orange indicates that the action is ready to be performed and that all preconditions and constraints are met;
- Green indicates that the item is currently being performed;
- Blue indicates successful completion;
- Red indicates a failure for which failure recovery planning steps might be initiated.

The drop-down menu providing the set of options available to perform any item on the panel is dynamically generated and is context-sensitive, reflecting the local knowledge of the capabilities of other panels and available services. It also draws on the in-built planner and other plug-in modules providing specialized handlers. I-Plan can perform hierarchical partial-order composition of plans from a library of single-level plan schemas or *Standard Operating Procedures* (SOPs). This library can be augmented during the planning phase, either through a simple interface allowing the

user to add specific ways to expand a given action (intended for users familiar with the application domain but not AI planning techniques) or with a more comprehensive graphical domain editor. Grammars and lexicons for the domain are built automatically during domain editing to assist the user.

The link to the underlying planner justifies our description of I-X as providing an *intelligent* system. The planner is able to match activity refinements (such as a decomposition of the activity into a number of simpler sub-activities) from the library and allow the user to access these by adding corresponding “expand” handlers to the action menu. At this point, instead of merely being able to tick-off an activity once done, users can exploit the knowledge in the library to break an activity down into sub-activities that, when all performed, accomplish the higher-level task. Of course, sub-activities can themselves be broken down further until a level of primitive actions is reached, at which point the library of procedures no longer contains any refinements that match the activities. One common way of handling an activity is through delegating the task to another (human or computer) agent in the system that is known to be capable of performing it. In this manner, the user collaboratively exploits the knowledge of other agents.

Future developments of I-Plan will provide more assistance with a “How do I do this?” option under the Action menu which will be able to account for other concurrent items on the panel, and account for mutual satisfaction of open variables and other constraints.

2.3. Other I-X Tools

The other tools in the I-X suite include messaging tools and various information viewers (e.g., map, 3D VRML and PDA interfaces) and editors, along with three specific tools: *I-DE*, *I-Q* and *I-Space*.

I-DE (I-X Domain Editor) allows the creation, maintenance and, ultimately, the publication of SOPs, generic approaches to archetypal activities.

I-Q (I-Query) is a generic I-X agent shell which, when embodied with the appropriate mechanisms, provides an agent with the ability to interact with a query service of some kind. It usually responds by adding facts or constraints into the current state of the panel. A typical application, for instance, might be the retrieval of information from some external source such as the Semantic Web.

I-Space is used to maintain organizational relationships with other agents in the environment. The nature of the relationship (for instance, supervisor-supervisee) will influence the nature of the activity-based interactions between these agents: the choices available to an agent will depend (amongst other things) both on its position in the organizational scheme of things and on its awareness of the capabilities and dynamic status (e.g. the current ‘presence’) of other agents. Exchange of agent and organization relationships with tools such as the KAoS Policy Administration Tool (KPAT) is possible [15].

2.4. I-X Message Formats

There are a number of messages that are used within the I-X Process Panels and that can be passed between panels and other services and agents.

- Issues, Activities, Constraints and Annotations;
- Current state information (as world state constraints);

- Plans (composites of Issues, Activities, Constraints and Annotations);
- Reports of progress or completion of nominated activities;
- Text-orientated ‘chat’ messages.

The first three relate to the core underlying ontology on which I-X is based. The other two message types provide status and other contextual information. Reports can be requested when the task of handling a certain activity is passed on to another agent; details of activity *progress* and *completion* (success/failure) are then reported to the original sender of the item. This provides a way to monitor activity progress and receive milestone reports.

2.5. State Information

Information about the current state of the environment can be passed to panels via *world state* constraints. These might come directly from sensors, or through some analysis or reporting system.

A specific type of state information we have found useful is the presence or status information maintained by instant messaging systems, so one can tell if another agent, panel or person is active and available for communications. The Jabber messaging protocol allows users to ‘subscribe’ to the presence of others; in this way, users can construct lists of their contacts, which the Jabber server dynamically augments with their current presence status.

Incoming completion reports and information about the current state sent as constraints can cause other activities to become executable when their pre-conditions or ordering constraints are satisfied.

I-X allows custom state information viewers to be added to expand or replace the default state viewer (a simple tabular representation) that is a constituent of a ‘standard’ I-P² panel. One example of a custom state viewer that has been found to be useful in a number of applications is the BBN OpenMap™ tool [23], which allows the information to be layered on a map (which might be topographical or something more conceptual, as the application demands). Changes to information in any viewer, or coming in via messages from outside of panels are synchronized.

2.6. The <I-N-C-A> Ontology

<I-N-C-A> (Issues-Nodes-Constraints-Annotations) is the basis of the ontology that underpins the I-X approach. It provides the framework for the representation used to describe processes and process products within I-X Process Panels, and gives a structure for the main types of activity-orientated I-X messages. <I-N-C-A> is designed to be a conceptual model that facilitates the expression and communication of task elements for both human and computer agents.

In <I-N-C-A>, both processes and process products are abstractly considered to be made up of a set of *issues*, *nodes* (corresponding to activities in a process, or parts of a physical product), *constraints* of various kinds (which describe relationships between nodes) and *annotations*, which provide rationale and other useful but less formal information. Before discussing in more detail each of the elements which comprise the <I-N-C-A> model, a brief discussion of the origins and development of the model will help give an understanding of its nature.

The forerunner of <I-N-C-A>, <I-N-OVA> [10], when first designed, was intended to act as a bridge to improve dialogue between a number of communities working on formal planning theories, practical planning systems and systems engineering process management methodologies. It was intended to reflect and support the then new work on automatic manipulation of plans, human communication about plans, principled and reliable acquisition of plan information, and formal reasoning about plans. It has since been utilized as the basis for a number of research efforts, practical applications and emerging international standards for plan and process representations. See [24] for a more detailed historical view of the relationships between <I-N-OVA>, the earlier work in AI on plan representations, and the results from the process and design communities and standards bodies.

2.6.1. Issues

Issues state unresolved questions about the task; these can concern unsatisfied objectives, questions raised as a result of analyses, etc. An issue can be thought of as implying potential further constraints to be added into the model in order to address the issue. Until recently, issues in I-X had a task or activity orientation to them, being mostly concerned with actionable items referring to the process underway – that is, with actions in the process space. Lately, however, we have adopted the gIBIS [25] approach of expressing issues as any of a number of specific types of question to be considered [26,27]. The types of questions advocated are:

- Deontic questions – what should we do?
- Instrumental questions – how should we do it?
- Criterial questions – what are the criteria?
- Meaning or conceptual questions – what does X mean?
- Factual questions – what is X? or, is X true?
- Background questions – what is the background to this project?
- Stakeholder questions – who are the stakeholders of this project?
- Miscellaneous questions – to act as a catch-all.

The first five of these are likely to be the most common in our task support environment. This is similar to the Questions-Options-Criteria approach [28], which was used for rationale capture for plans and plan schema libraries in our earlier work [29], and to the mapping approaches used in Compendium [30]. Compendium can in fact exchange its set of issues, activities and some types of constraints and annotations with an I-P² [21,31].

2.6.2. Nodes

The nodes in the specifications describe components that are to be included in the current design. Nodes can themselves be artifacts that can have their own structure with sub-nodes and other <I-N-C-A>-described refinements associated with them. When <I-N-C-A> is being used to describe processes, the nodes are usually the individual activities or their sub-activities. They are usually characterized by a *pattern* composed of an initial verb followed by any number of parameter objects, noun phrases, and qualifiers or filler words describing the activity. An example activity might be:

`(transport package-1 from location-a to location-b)`

Nodes can themselves be thought of as constraints that restrict the space within which an artifact may lie, with the sets of issues and constraints serving to restrict

further the space of solution artifacts. Others have recognized the special nature of the inclusion of nodes (or activities) into a synthesized artifact (or plan) compared to all the other constraints that may be described. In the planning domain, Khambhampati and Srivastava differentiate Plan Modification Operators into *progressive refinements*, which can introduce new actions into the plan, and *non-progressive refinements*, which partition the search space with existing sets of actions in the plan [32]. They call the former genuine planning refinement operators, and think of the latter as providing the scheduling component.

2.6.3. Constraints

The constraints restrict the relationships between the nodes to describe only those artifacts within the design space that meet the requirements. The constraints may be split into *critical constraints* and *auxiliary constraints*; some constraint managers (solvers) can return ‘maybe’ answers to constraints of the latter type, indicating that the constraint being added to the model is satisfactory providing other critical constraints are imposed by other constraint managers. The ‘maybe’ answer is expressed as a disjunction of conjunctions of such critical or shared constraints. More details on the ‘yes/no/maybe’ constraint management approach used in I-X and the earlier O-Plan systems are available in [4].

The choices of which constraints are considered critical and which are considered auxiliary is itself a decision for an application of I-X, and specific decisions will need to be made about how to split the management of constraints within such an application. In other words, it is not pre-determined for all applications. A temporal activity-based planner would normally have object/variable constraints (equality and inequality of objects) and some temporal constraints (maybe just a simple *before* constraint) as the critical constraints. But, in a 3D design or a configuration application object/variable and other critical constraints (possibly spatial constraints) might be chosen. It depends on the nature of what is communicated between constraint managers in the application of the I-X architecture.

2.6.4. Annotations

The annotations add additional human-centric information or design and decision rationale to the information describing the artifact.

3. An Application of I-X: I-Ex

Rob Milne’s planned journey to the Himalayas, his eagerness to exploit fully the opportunities afforded by this trip and his interest in the I-X technology together provided us with a unique chance to apply I-X concepts to a new domain, namely that of coordinating an expedition to some remote place. The name given to this application was *I-Ex*, and it would test both the generality of the I-X and <I-N-C-A> concepts and the robustness of the I-X software under taxing conditions.



Figure 3. Overview of the I-Ex infrastructure.

Figure 3 gives an overview of the various components of I-Ex. The shared task is the expedition itself. In the event of, say, a significant deterioration in weather forecast over the coming hours, the I-X technology might be invoked to re-plan aspects of the climb. This might involve reasoning with available resources (for example, the number of oxygen bottles known to be stored at the South Col), temporal aspects (when will the bad weather pass over? how many hours of daylight are left?) and capabilities (available rescue services, if the situation worsens).

It was envisaged that there would be a number of different human agents involved:

- The mountaineer, carrying I-X on a lightweight Pocket PC/PDA, continually reports his progress and status to the base camp reporting system.
- Progress is monitored by an I-P² user at base camp (which has 'full' internet connectivity), and in turn supplies the mountaineer with next or revised objectives, latest weather updates, etc.
- An expedition support team, based, say, in the UK, and also using an I-P² is kept informed of progress, providing a full reporting system and updating the expedition web site accordingly.

In actuality, time and technology constraints meant that in this instance we would concentrate on the monitoring and communications aspect of the last two of these, effectively meaning that we would have an 'I-X-enabled' laptop at base camp that would communicate progress information to the 'support team' at Edinburgh.

3.1. Tailoring for I-Ex

Before deployment, it was necessary to tailor the general I-X system and tools for this particular application. These preparations concerned communications, task modelling, system 'packaging' and establishing external reporting mechanisms.

3.1.1. Communications in I-Ex

Jabber was chosen as the primary medium for I-X communications between base camp and the support team (with e-mail being used as a secondary channel for images and other non-I-X information). There were two principal reasons for this choice. First, Jabber allows asynchronous communications, in the sense that it is not necessary for the intended recipient of a message to be online when the message is sent: messages can be queued until the next time the recipient logs on. This becomes important when we cannot ensure that agents are online continuously (and where there is a time-difference between agents), and do not wish messages to be lost. However, there are other situations in which synchronous communications are desirable, such as time-critical tasks in which it is important that a particular activity begins almost immediately, and hence, candidates for performing it should exclude any off-line agents. In these cases I-X can prevent any messages being sent to off-line agents.

Secondly, the choice of Jabber meant that we were able to use a third-party server (namely jabber.org), already known from previous applications to be reasonably robust.

3.1.2. 'Packaging' I-Ex

In practice, 'the user' at base camp at any given time might be one of a number of individuals, all having different levels of computer skills and awareness of the I-Ex project objectives. In order to accommodate these different users, it was felt necessary to develop a basic 'reporting' version of I-P², which would be both easy to use and foolproof, as far as possible, in operation. The issues, state constraints and annotations would all be hidden in this reporting panel, leaving just the ascent plan in its current state. When appropriate, the user would simply select 'Done' from the action menu of the current activity; this would generate and send the appropriate reporting messages to the support team panel. In addition, the panel would automatically save the state of the current plan on closing (after having first generated a back-up file of the plan before the latest modifications) and reload it on subsequent start up.

While the aggregate effect of these changes was to limit the autonomy and flexibility of the base camp user and hide much of the underlying functionality, the lack of time and opportunity to familiarise the users meant that this seemed a small price to pay in return for robustness during the experiment.

3.1.3. Task Modelling for I-Ex

With Rob's input, we were able to develop an initial 'ascent plan'. Activities in this plan correspond to achieving particular targets, such as arrival at base camp, acclimatization at camp 1, through to reaching the summit and the subsequent descent to base camp. Some of these activities contain sub-activities; for instance, acclimatization at camp 1 involves the ascent to camp 1 and, after a suitable period, a return to base camp. Figure 4 shows part of this ascent plan (seen here in the minimal reporting panel).

The ascent plan is echoed in a 'monitoring plan', used to control the support team activity. This contains activities such as monitoring the arrival at base camp, which in turn consists of the sub-activities of awaiting a report of arrival at base camp and then updating the external progress reports on the web, to relate this information to a wider audience. Hence, every activity in the ascent plan has a counterpart in the monitoring plan, with the I-X reporting mechanism being used to automatically generate and send

progress messages when the base camp user indicates that the activity is now 'Done'. Figure 5 shows the monitoring plan as seen in a standard I-P².

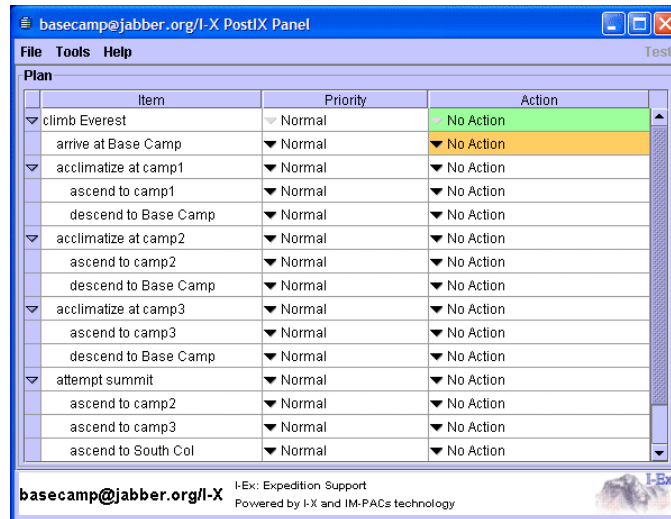


Figure 4. The base camp process panel, showing the initial activities in the ascent plan.

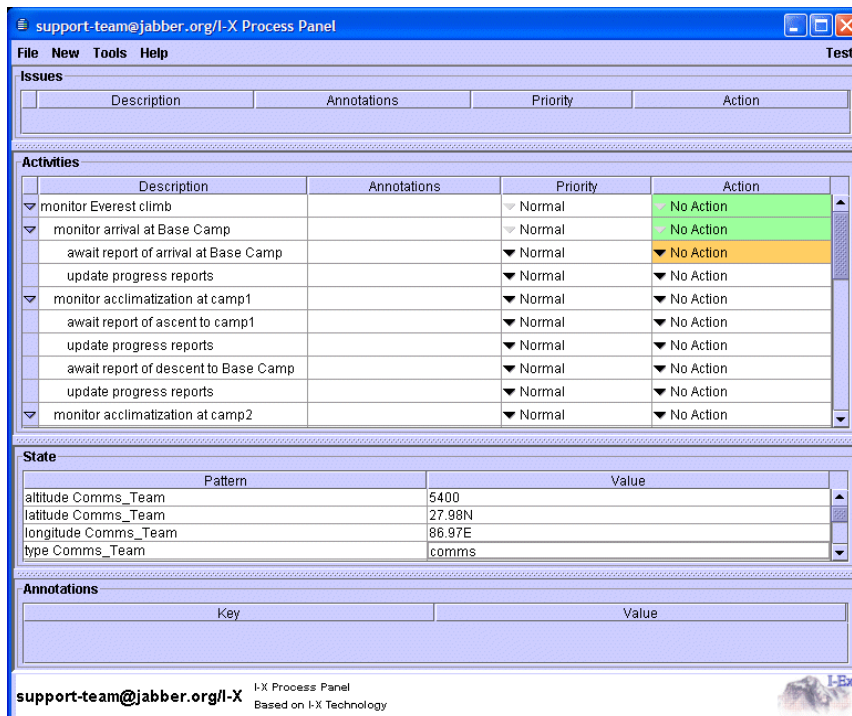


Figure 5. The support team process panel, showing the first steps in the monitoring plan.

3.1.4. External Reporting in I-Ex

A feature of many expeditions and journeys in the internet age is the expedition 'blog', which describes expedition progress and status, as well as photographs, more personal reflections, where appropriate, on the daily routine, surroundings, etc. The blog serves as a central point for broadcasting information intended to keep a general audience informed. Since reporting and monitoring were integral to the I-Ex project, it was suggested by Rob that the provision of a blog would form a natural extension to this, as well as providing him with a useful service. However, in addition to the 'normal' manual approach for updating the blog, we tried to introduce a 'semantic blogging' element, whereby I-X reports could, according to the message type, be automatically converted into the appropriate progress reports and posted on the web site, exploiting the meaning of the underlying <I-N-C-A> messages and the activity progress reporting mechanisms that are built-in to I-X. An autonomous I-X agent was developed for this task, but as it had not been adequately tested before the expedition team set off, it was not deployed. Instead, reports were manually converted into blog entries, and these were supplemented by conventional entries which were sent by e-mail. Figure 6 shows some typical blog entries that were written by Rob alongside a report generated from an I-X progress update.

3.2. I-Ex in Practice: Lessons Learned

The aims of the I-Ex deployment during the Everest expedition were somewhat limited in scope, being restricted to reporting tasks and not exercising the full potential of the I-X environment as a general task-achieving environment. However, within this scope, the deployed system was felt to be a success. The reporting aspects worked well, with the Jabber communications robust enough to enable transmissions when e-mail was unavailable, and the blog and associated web-site attracting much interest.

There are a number of more general lessons that we have been able to draw from the I-Ex application; these might be summarized in two points as follows:

- The importance of appropriate tailoring of the I-X tools: a mixed-initiative task is likely to involve humans of mixed abilities and backgrounds. It is important that the interfaces are modified accordingly to present an appropriate view onto the underlying model, and one which will enable the user to interact appropriately with the models and other participating agents.
- Reporting and monitoring: one of the key roles that technology such as I-X might come to play is in making explicit, in ways previously unimaginable, the current state of the task and its participants through the establishing and use of clear and, where appropriate, formal reporting structures, which may work on different levels: reporting between individual agents, reporting to groups of agents, broadcasting to the community of agents and possibly beyond.

4. Future Directions

Work on I-X and its applications in coalition command and control, search and rescue, help desks, etc. has indicated the value of adopting an 'augmented' style of the instant messaging paradigm which we call activity-orientated 'intelligible messaging'. This

provides a platform for making AI planning technology available in an immediately usable form.

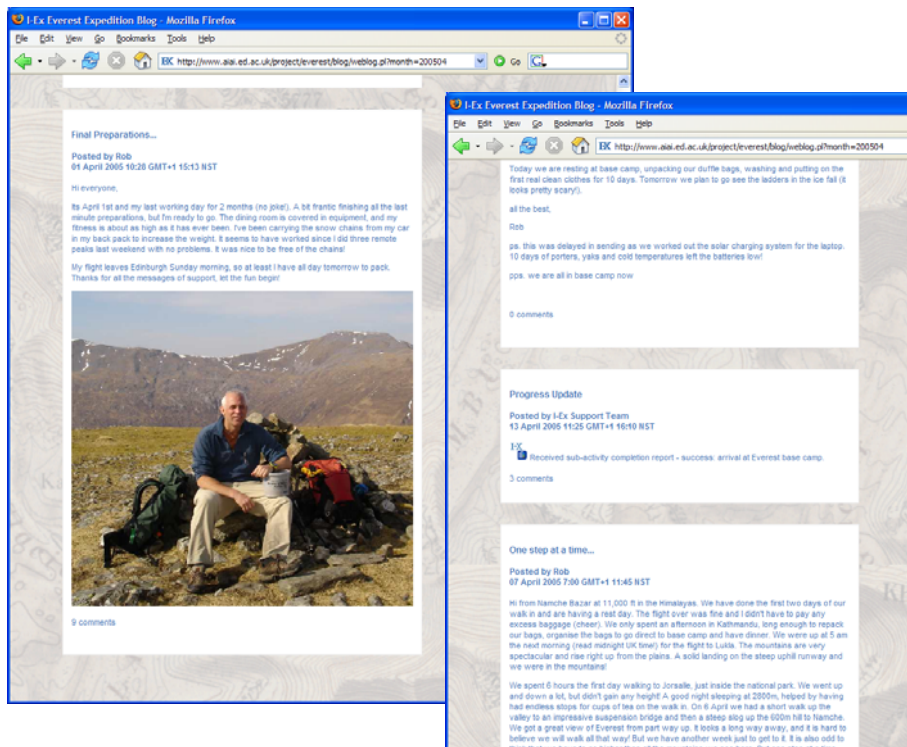


Figure 6. Some of Rob's blog entries. The *Progress Update* entry shown in the center of the right-hand browser was generated from the corresponding I-X activity completion report.

There are many opportunities for extending the initial approach and the technology that implements it. A more comprehensive I-Plan planner is to be incorporated as project work allows, and this will use the O-Plan 'repairing plans on-the-fly' repair technology [33] to recover from failures. The incorporation of more capable constraint managers and optimization algorithms to propose options is possible within the design.

More sophisticated and robust communications strategies are being studied, including logging of message traffic for quality control and audit purposes. We are particularly interested in the changes of process that might be triggered when the status of agents and panels alter (say moving from a synchronous on-line instant messaging mode of communication to an off-line store and forward messaging mode). Ways to describe panels and user/service capabilities in an organization, the roles they play and the authorities they have to act for one another are also being developed.

5. Conclusions

This paper has described the I-X approach to mixed-initiative synthesis tasks, whereby human and computer agents are assisted in their activity-oriented collaborations. The

accompanying suite of tools provides human users with interactive visualizations of the current state of the collaboration from their particular perspectives, underpinned and supported with communications tools, planning aids and organization and domain modellers. As well as what might be thought of in conventional terms as synthesis tasks, the underpinning <I-N-C-A> ontology is general enough to allow the concepts to be applied to more conceptual synthesis tasks such as mounting an expedition: the I-Ex application, though necessarily restricted in scope, is one example of this, with the communications and reporting aspects of the architecture firmly to the forefront.

More technical details on I-X and the current version of the I-P² software (available for demonstration purposes) are available at the project website [34].

6. Acknowledgements

Rob Milne is listed as an author of this paper, as we planned with him for a write up of the joint work on I-Ex. Rob assisted the I-X team in investigating the application of I-X technology to extreme expedition support and used a packaged version of the technology during his Everest expedition. He regularly reported to the Edinburgh team on his use of the system and provided valuable input throughout, and our intention was to jointly write-up the experiences of using I-Ex for academic publication.

This material is based on research within the I-X project and sponsored by a range of sources including the Defense Advanced Research Projects Agency (DARPA) and US Air Force Research Laboratory under agreement number F30602-03-2-0014. Thanks to our co-workers on the I-X project at AIAI (in particular Jeff Dalton, Jussi Stader and Gerhard Wickler) for supporting technologies and valuable discussions relating to this work.

The IM-PACs project is funded by Scottish Enterprise, the European Research Development Fund (ERDF) and AIAI at the University of Edinburgh. The project has received assistance from the Proof of Concept Fund managed by Scottish Enterprise. The Proof of Concept Funds supports the pre-commercialisation of leading-edge technologies emerging from Scotland's Universities, Research Institutes and National Health Service Trusts [Project Reference: 5-SEB-004].

The U.S. Government and the University of Edinburgh are authorized to reproduce and distribute reprints and on-line copies 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 the official policies or endorsements, either expressed or implied, of other parties.

7. References

- [1] Tate, A. (2003) <I-N-C-A>: an Ontology for Mixed-initiative Synthesis Tasks, *Proceedings of the Workshop on Mixed-Initiative Intelligent Systems (MIIS)* at the International Joint Conference on Artificial Intelligence (IJCAI-03), Acapulco, Mexico, August 2003.
- [2] Tate, A. (1977) Generating Project Networks, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-77)*, pp. 888-893, Cambridge, MA, USA, Morgan Kaufmann.
- [3] Currie K. and Tate A. (1991) O-Plan: the Open Planning Architecture, *Artificial Intelligence*, 52, 49-86, 1991.

- [4] Tate A. (1995) Integrating Constraint Management into an AI Planner, *Journal of Artificial Intelligence in Engineering*, 9(3), pp. 221-228, 1995.
- [5] Tate, A., Dalton J. and Levine, J. (1998) Generation of Multiple Qualitatively Different Plan Options, *Proceedings of the Fourth International Conference on Artificial Intelligence Planning Systems (AIPS-98)*, Pittsburgh PA, USA, June 1998.
- [6] Tate, A., Dalton, J. and Levine, J. (2000) O-Plan: a Web-based AI Planning Agent, AAAI-2000 Intelligent Systems Demonstrator, in *Proceedings of the National Conference of the American Association of Artificial Intelligence (AAAI-2000)*, Austin, Texas, USA, August 2000.
- [7] Levine, J., Tate, A. and Dalton, J. (2000) O-P3: Supporting the Planning Process using Open Planning Process Panels, *IEEE Intelligent Systems*, 15(6), November/December 2000.
- [8] Aarup, M., Arentoft, M.M., Parrod, Y., Stokes, I., Vadon, H. and Stader, J. (1994) Optimum-AIV: A Knowledge-Based Planning and Scheduling System for Spacecraft AIV, in *Intelligent Scheduling* (eds. Zweben, M. and Fox, M.S.), pp. 451-469, Morgan Kaufmann.
- [9] Tate, A. (1996a) Responsive Planning and Scheduling Using AI Planning Techniques - Optimum-AIV. In, "Trends & Controversies - AI Planning Systems in the Real World", *IEEE Expert: Intelligent Systems & their Applications*, 11(6), pp. 4-12, December 1996.
- [10] Tate, A. (1996b) Representing Plans as a Set of Constraints – the <I-N-OVA> Model, *Proceedings of the Third International Conference on Artificial Intelligence Planning Systems (AIPS-96)*, Edinburgh, May 1996.
- [11] Tate, A. (2000) <I-N-OVA> and <I-N-CA> - Representing Plans and other Synthesized Artifacts as a Set of Constraints, AAAI-2000 Workshop on Representational Issues for Real-World Planning Systems, *National Conference of the American Association of Artificial Intelligence (AAAI-2000)*, Austin, Texas, USA, August 2000.
- [12] Fraser, J. and Tate, A. (1995) The Enterprise Tool Set – An Open Enterprise Architecture, in *Proceedings of the Workshop on Intelligent Manufacturing Systems, International Joint Conference on Artificial Intelligence (IJCAI-95)*, Montreal, Canada, August 1995.
- [13] Stader, J. (1996) Results of the Enterprise Project, in *Proceedings of Expert Systems '96, the 16th Annual Conference of the British Computer Society Specialist Group on Expert Systems*, Cambridge, UK, December 1996.
- [14] Kahn, M.L. and Della Torre Cicalese, C. (2001) CoABS Grid Scalability Experiments, Autonomous Agents 2001 Conference, *2nd Int'l Workshop on Infrastructure for Agents, MAS, and Scalable MAS*, May 29, 2001.
- [15] Bradshaw, J.M., Duffield, S., Benoit, P. and Woolley, J.D. (1997) KAoS: Toward an industrial-strength generic agent architecture, In J. M. Bradshaw (Ed.), *Software Agents*, 1997, pp. 375-418, Cambridge, MA: AAAI Press/The MIT Press.
- [16] Tate, A., Dalton, J., Jarvis, P. and Levine, J. (2000) Using AI Planning Technology for Army Small Unit Operations, Poster Paper in the *Proceedings of the Artificial Intelligence Planning and Scheduling Systems Conference (AIPS-2000)*, Breckenridge, Colorado, USA, April 2000.
- [17] Allsopp, D.N., Beautement, P., Bradshaw, J.M., Durfee, E.H., Kirton, M., Knoblock, C.A., Suri, N. and Tate, A. and Thompson, C.W. (2002) Coalition Agents Experiment: Multi-Agent Co-operation in an International Coalition Setting, Special Issue on Knowledge Systems for Coalition Operations (KSCO), *IEEE Intelligent Systems*, 17(3), pp. 26-35, May/June 2002.
- [18] Wark, S., Zschorn, A., Perugini, D., Tate, A., Beautement, P., Bradshaw, J.M. and Suri, N. (2003) Dynamic Agent Systems in the CoAX Binni 2002 Experiment, Special Session on Fusion by Distributed Cooperative Agents at the 6th International Conference on Information Fusion (Fusion 2003), Cairns, Australia, July, 2003.
- [19] CoSAR-TS Team (2003) Coalition Search and Rescue Task Support Web Site. <http://www.aiai.ed.ac.uk/project/cosar-ts/>
- [20] Siebra, C. and Tate, A. (2003) I-Rescue: A Coalition Based System to Support Disaster Relief Operations, *The Third International Association of Science and Technology for Development (IASTED) International Conference on Artificial Intelligence and Applications (AIA-2003)*, Benalmadena, Spain, September 2003.
- [21] Buckingham Shum, S., De Roure, D., Eisenstadt, M., Shadbolt, N. and Tate, A. (2002) CoAKTinG: Collaborative Advanced Knowledge Technologies in the Grid, Proceedings of the Second Workshop on Advanced Collaborative Environments, Eleventh IEEE Int. Symposium on High Performance Distributed Computing (HPDC-11), July 24-26, 2002, Edinburgh, Scotland.
- [22] Jabber (2003) Jabber XML Instant Messaging, <http://www.jabber.org>.
- [23] BBN (2003) OpenMap™ Open Systems mapping Technology - <http://openmap.bbn.com>
- [24] Tate, A. (1998) Roots of SPAR – Shared Planning and Activity Representation, *The Knowledge Engineering Review*, 13(1), 1998.

- [25] Conklin, J. and Begeman, M.L. (1988) gIBIS: A Hypertext Tool for Exploratory Policy Discussion. *ACM Trans. Inf. Syst.* 6(4), 303-331, 1988.
- [26] Selvin, A.M. (1999) Supporting Collaborative Analysis and Design with Hypertext Functionality, *Journal of Digital Information* 4, 1999.
- [27] Conklin, J. (2005) *Dialogue Mapping: Building Shared Understanding of Wicked Problems*. Wiley: Chichester.
- [28] MacLean, A., Young, R. M., Bellotti, V., & Moran, T. (1991) Questions, Options, and Criteria: Elements of Design Space Analysis, *Human-Computer Interaction*, 6(3 & 4), 201-250, 1991.
- [29] Polyak, S. and Tate A. (1998) Rationale in Planning: Causality, Dependencies and Decisions, *The Knowledge Engineering Review*, 13(3), September 1998, pp. 247-262.
- [30] Buckingham Shum, S., Selvin, A.M., Sierhuis, M., Conklin, J., Haley, C.B. and Nuseibeh, B. (2006). Hypermedia Support for Argumentation-Based Rationale: 15 Years on from gIBIS and QOC. In: A.H. Dutoit, R. McCall, I. Mistrik, and B. Paech (Eds.), *Rationale Management in Software Engineering* Springer-Verlag/Computer Science Editorial.
- [31] Chen-Burger, Y. and Tate, A. (2003) Concept Mapping Between Compendium and I-X, Informatics Report Series, University of Edinburgh, EDI-INF-RR-0166, May 2003.
- [32] Khambhampati, S. and Srivastava, B. (1996) Unifying Classical Planning Approaches, Arizona State University ASU CSE TR 96-006, July 1996.
- [33] Drabble, B., Dalton, J. and Tate, A. (1997) Repairing Plans on the Fly, in *Proceedings of the NASA Workshop on Planning and Scheduling for Space, Oxnard CA, USA, October 1997*.
- [34] I-X Project Website: <http://www.aiai.ed.ac.uk/project/ix/>.