

Final Technical Report

Collaborative Operations for Personnel Recovery

**Austin Tate, Jeff Dalton, Jussi Stader, and Gerhard Wickler –
AIAI, Edinburgh, UK**

Jeffrey Hansberger – USJFCOM/J9, Norfolk, VA

Artificial Intelligence Applications Institute
The University of Edinburgh
Appleton Tower, Crichton Street, Edinburgh EH8 9LE, UK

Principal Investigator: Prof. Austin Tate. Tel: +44 131 650 2732

Contract No. F30602-03-2-0014
Effective Date of Contract: February 18, 2003
Reporting Period: January 1, 2005 – May 31, 2007
Contract Expiration Date: May 31, 2007

Report Date: August 31, 2007 (revised January 30, 2008)

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency, US Air Force Research Laboratory, or the U.S. Government.

Acknowledgements

This material is based on research sponsored by the Defense Advanced Research Projects Agency (DARPA) under agreement number F30602-03-2-0014. The University of Edinburgh, the U.S. government and research sponsors 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.

Thanks to the other members of the I-X project team: Stuart Aitken, Jessica Chen-Burger, Thomas French, John Levine, Natasha Lino, Stephen Potter, Claurton Siebra and Eleanor Sim.

We are grateful for support from the DARPA program managers, John Allen and Alexander Kott, and for the guidance and assistance with access to subject area resources from the project manager, Mark Gorniak at Air Force Research Lab, Rome.

Note that information, displays and photos in this report are from an unclassified element of a Command Post Exercise (CPX) at the Personnel Recovery Education and Training Center (PRETC) in Fredericksburg, Virginia. Thanks to the Director and his staff for allowing the Co-OPR project team to attend this exercise and gather information for use on the project. All documents and displays used in the exercise were unclassified, contain fictional training data and do not refer to real people.

Contents

1.	Summary	1
2.	Introduction	3
2.1	Aims of the Project	3
2.2	Overview of the Final Report	3
3.	I-X Technology	5
3.1	Issues.....	6
3.2	Nodes	7
3.3	Constraints	8
3.4	Annotations	10
4.	The Co-OPR Problem	13
4.1	The Personnel Recovery Domain	13
4.2	Command Post Exercises.....	15
4.3	Requirements for I-X	16
5.	Developing I-X and the Co-OPR Application	21
5.1	Core I-X Components	23
5.2	The Domain Editor	23
5.3	White Cell Support: I-Sim	25
5.3.1	Discrete Event Simulation	26
5.3.2	Process-level Simulation and Time	27
5.4	Other Tools and Viewers	28
5.4.1	I-Space	29
5.4.2	Agent Capabilities.....	29
5.4.3	Handling Activities through Task Distribution.....	31
5.4.4	Structured Instant Messaging.....	31
5.4.5	Documentation	31
5.5	The Co-OPR Application.....	32
6.	Experiments and Evaluation	34
6.1	Experimental Set-Up.....	34
6.2	Execution of the Experiments	35
6.3	Results.....	37
6.3.1	Task Analysis.....	37
6.3.2	Benefits: Cognitive Attributes	38
7.	Conclusions and Future Work	45
8.	References	47
Appendix A:	Screen Shots from the Final Experiment	52
Appendix B:	List of Publications	85
Appendix C:	List of Software and On-line Documentation Available	89
Appendix D:	Project Web Sites	90
Abbreviations	91

List of Figures

Figure 1: The JPRC director with checklists on his desks.....	14
Figure 2: Generic Scenario Map.....	15
Figure 3: Organization of Agents in the Scenario	16
Figure 4: Incident board towards the end of a CPX.....	18
Figure 5: Asset board used during CPX	19
Figure 6: Map used during CPX.....	20
Figure 7: Busy White Cell during CPX.....	20
Figure 8: Initial state of the JPRC director's I-X Process Panel.....	21
Figure 9: Initial state of the JPRC director's I-X Process Panel.....	24
Figure 10: The I-Sim Tool during the simulation (at 15:27 simulated time).....	25
Figure 11: The I-Space Tool: Relations.....	29
Figure 12: The I-Space Tool: Capabilities.....	30
Figure 13: Fragment of Co-OPR task analysis	37
Figure 14: Hierarchical Task Analysis of a Personnel Recovery Center Director. Lightly highlighted tasks represent areas which Co-OPR is hypothesized to positively influence. .	40
Figure 15: The I-X Name Server with all relevant agents registered.....	52
Figure 16: The I-X for the White Cell (trainers).....	52
Figure 17: The basic parameters to start the scenario simulation.....	53
Figure 18: The I-Sim Tool for controlling the events that will occur.....	53
Figure 19: The I-Sim Clock showing Simulated Time (at the White Cell).....	53
Figure 20: The I-Sim Tool during the simulation (at 15:27 simulated time).....	54
Figure 21: Initial state of the JPRC director's I-X Process Panel.....	54
Figure 22: JPRC director's panel with the first task (from the JTFC).....	55
Figure 23: JPRC director's panel with the first task expanded.....	55
Figure 24: JPRC director's panel – JPRC setup in progress.....	56
Figure 25: JPRC director's panel – JPRC setup completed.....	56
Figure 26: JPRC director's panel – second phase: distribution of the ATO.....	57
Figure 27: JPRC director's panel – third phase: first incident in progress	57
Figure 28: JPRC director's panel – third phase: second incident report just arrived	58
Figure 29: JPRC director's panel – information gathering for two parallel incidents in progress	59
Figure 30: The Watch Supervisor's Screen (1).....	60
Figure 31: The Watch Supervisor's Panel (2).....	61
Figure 32: The Watch Supervisor's Panel (3).....	61
Figure 33: The Watch Supervisor's Screen (4).....	62
Figure 34: The Watch Supervisor's Panel (5).....	63
Figure 35: The Watch Supervisor's Panel (6).....	63
Figure 36: The Watch Supervisor's Screen (7).....	64
Figure 37: Initial state of the JPRC controller's I-X Process Panel.....	65
Figure 38: JPRC controller's Panel with first task – to bring up the CSAR asset board.....	65
Figure 39: JPRC controller's Panel with next task – to review the ATO	66
Figure 40: The ATO document – served from I-Serve.....	66
Figure 41: JPRC controller's Panel – review of the ATO completed.....	67
Figure 42: Asset board displaying the CSAR resources and their state	68

Figure 43: Task to plan and execute mission for first incident arrived	69
Figure 44: Planning a rescue mission for the second incident.....	69
Figure 45: Using the I-Plan tool to automatically generate options	70
Figure 46: Using the I-Plan Option tool to compare options.....	71
Figure 47: Selecting an option in the current space.....	72
Figure 48: The Virtual Operations Center (as a local HTML resource).....	73
Figure 49: Shared display (1).....	74
Figure 50: Shared display (2).....	75
Figure 51: Shared display (3).....	76
Figure 52: Shared display (4).....	77
Figure 53: Shared display (5).....	78
Figure 54: Shared display (6).....	79
Figure 55: Shared display (7).....	80
Figure 56: The SPINs (document served by I-Serve).....	81
Figure 57: The I-X Process Panel for the JTFC.....	82
Figure 58: The I-X Process Panel for the JTFC after completion of the first 2 phases	82
Figure 59: The I-X Process Panel for the JTFC at the end of the experiment.....	83
Figure 60: The I-X Process Panels for the other RCCs	84

1. Summary

The Collaborative Operations for Personnel Recovery (Co-OPR) project sought to provide collaborative task support for a Search and Rescue coordination center. The project aimed to create a prototype “Personnel Recovery (Experimental) Pack” (PREP) and to demonstrate its use.

A number of requirements capture, knowledge gathering and transition workshops and meetings were held. This included an initial requirements setting workshop early in 2005, meetings at the USJFCOM Joint Personnel Recovery Agency’s (JPRA) Personnel Recovery Education and Training Center (PRETC) in Fredericksburg, Virginia in June 2005, a review meeting in Edinburgh in August 2005, and attendance at a Command Post Exercise (CPX) at the PRETC in November 2005. These initially established the potential areas for use of Co-OPR and I-X tools in training exercises. In the second project year such tools were developed and tested using a training exercise as held at the PRETC, using observers from PRETC and USJFCOM and an evaluation expert from USJFCOM/J9.

The project was provided with a rich set of urban and rural scenarios by JPRA/PRETC which together are unclassified versions of scenarios used within the PRETC training courses and Command Post Exercises. At the time, these stretched the capabilities of the current and envisaged technologies within Co-OPR/I-X. Refinement of the scenarios alongside PRETC, and knowledge engineering to capture information on standard operating procedures and responses were a key part of making the work relevant to the potential real use of Co-OPR/I-X for Personnel Recovery.

The core I-X technology was packaged into a number of checkpoint releases to make available the features required to meet the application needs. I-X version 4.3 released in November 2005 to checkpoint the results achieved on the first 12 months of work with the PRETC. It formed a basis for the work on really using the technology at the PRETC. New “white cell” aids for training were made available in an initial version. A new I-Sim simulation capability, and advanced option exploration tools have all been improved significantly to make them more usable, including the features of the I-Plan AI planner and its capability for plan repair after failures. The features of the Domain Editor (I-DE) and its ability to browse and update or augment standard operating procedure knowledge dynamically during missions were enhanced. The final release of I-X that includes all the new developments achieved in the Co-OPR project is version 4.5.

The results of the work were packaged, along with Personnel Recovery domain specific models, as a web site and/or CD which could be considered as a prototype “Personnel Recovery (Experimental) Pack” of tools to assist a Joint Personnel Recovery Center (JPRC) and associated operational staff in performing their operations. The versions of PREP produced were used in one workshop or Command Post Exercise at the PRETC under guidance from Dr. Jeff Hansberger at training related workshops already organized by USJFCOM/J9 Expt. and Fred Kleibacker, the (now former) Director of the PRETC in Fredericksburg, Virginia. Co-OPR team members were engaged with these workshops to

show the tools in realistic settings, to assist with training where possible, and to gather experimental feedback.

Realistic use of tools for Personnel Recovery requires that the systems can work with emerging technology for geo-positioning, survival radios, evasion aids, robotic or semi-automated rescue aids or robots, and doctrine or tactics, techniques and procedures for Personnel Recovery. A number of short studies of these “complementary technologies” were made which explored deployment and inter-working aspects of these with the Co-OPR/I-X technology.

2. Introduction

The aim of this report is to describe the work and results of the Co-OPR project from January 2005 to May 2007. In this section we will begin by describing the aims as they were defined at the beginning of the project. An overview of the different sections will follow.

2.1 Aims of the Project

Personnel recovery teams must operate under intense pressure, taking into account not only hard logistics, but ‘messy’ factors such as the social or political implications of a decision. The *Collaborative Operations for Personnel Recovery (Co-OPR)* project has developed decision-support for sensemaking in such scenarios, seeking to exploit the complementary strengths of human and machine reasoning.

In the first phase of the project in 2004 (Tate *et al.*, 2006) Co-OPR integrated the *Compendium* sensemaking-support tool for real time information and argument mapping with the *I-X* artificial intelligence planning and execution framework to support group activity and collaboration. Both share a common model for dealing with issues, the refinement of options for the activities to be performed, handling constraints and recording other information. The tools span the spectrum from being very flexible with few constraints on terminology and content, to knowledge-based relying on rich domain models and formal conceptual models (ontologies). In a personnel recovery experimental simulation of an UN peacekeeping operation, with roles played by military planning staff, the Co-OPR tools were judged by external evaluators to have been very effective.

In the second phase of the project, which is described in this report, a closer cooperation with the Personnel Recovery Education and Training Centre (PRETC) in Fredericksburg, VA was used to identify additional requirements from a group of people who could be end users of the envisaged application. These requirements were used to further drive the development of the generic *I-X* framework as well as develop an application based on the framework that is aimed at supporting personnel recovery tasks. In addition to this originally envisaged aim, the co-operation from PRETC made it possible to extend the application to support the kind of personnel recovery training undertaken at PRETC and extend *I-X* with some generic tools that support general training applications. The resulting application has been tested in a series of experiments that continue the experiments conducted in the first phase of the project, starting from a view that corresponds to the *I-X* way of supporting a task towards an application that mirrors the training scenarios developed by PRETC.

2.2 Overview of the Final Report

The remainder of this report is organized as follows: Firstly, we will describe the *I-X* framework that has been used as a basis for the development undertaken. More specifically we shall describe the ontology underlying the whole system and approach.

This is necessary for understanding the philosophy behind I-X Process Panels, the principal user interface of an I-X application. This ontology describes a plan in terms of 4 components: issues, nodes, constraints and annotations. These components and their representation will be described in detail.

Next, we will describe the problem that has been addressed in the Co-OPR project: personnel recovery. After an overview of the domain, we will look more closely at the training exercises that the PRETC conducts to prepare US military personnel for the task of running a Joint Personnel Recovery Centre and the various component centers. Observations made by the project team during one of these exercises lead to a number of requirements that will be described in this report.

With the basic I-X technology and the problem described, we will continue to show how this technology was applied to develop a tool that can be used to support collaborative personnel recovery as observed at PRETC. The I-X framework already came with a number of tools that were expected to be useful in this domain, and these will be described here. For example, the intelligence in the I-X Process Panels is achieved using a library of standard operating procedures, an approach based on HTN (Hierarchical Task Network) planning (Sacerdoti, 1975; Tate 1977). The HTN planning system built into I-X is seamlessly integrated into the system. I-X is not meant to only support single agents in responding to an incident, but it also provides mechanisms for connecting a number of I-X Process Panels and supporting a coordinated multi-agent response. The key here is a simple agent capability model that automatically matches tasks to known capabilities for dealing with these tasks. In addition to the existing tool set, a number of new tools that were added to the generic framework will be described.

When the I-X framework is instantiated with a domain-specific model, we refer to it as an I-X application. Such an application has been developed during the Co-OPR project for the task of personnel recovery and personnel recovery training. A brief description of this application will conclude the next section.

The application was evaluated through a series of experiments that were conducted at AIAI in Edinburgh (using remote observation) and later at USJFCOM/J9 facilities in Norfolk, VA. We will describe the experimental set-up that was used to evaluate the Co-OPR application and how the execution of these experiments was performed.

Finally the results of these experiments will be presented and analyzed, showing that the I-X framework provides a number of very useful features that can be exploited for general task-supporting applications and personnel recovery in particular.

3. I-X Technology

There are a number of tools available that help people organize their work. One of these is provided with virtually every organizer, be it electronic or paper-based: the “to-do” list. This is because people are not good at remembering long lists of potentially unrelated tasks. Writing these tasks down and ticking them off when they have been done is a simple means of ensuring that everything that needs to be done does get done, or at least, that a quick overview of unaccomplished tasks is available. In responding to an emergency this is vital, and the larger the emergency, the more tasks need to be managed.

I-X is a framework that can be used to create an application in which multiple agents adopt a task-centric view of a situation, and which supports the necessary coordination of their activities to respond to that situation. The I-X Process Panel provides the functionality of a to-do list and thus, it is a useful tool when it comes to organizing the response to an emergency. The idea of using a to-do list as a basis for a distributed task manager is not new (Kreifelts, Hinrichs and Woetzel, 1993). However, I-X goes well beyond this metaphor and provides a number of useful extensions that facilitate the finding and adaptation of a complete and efficient course of action.

I-X Process Panels constitute the primary user interface to an I-X application. A panel more or less directly reflects the ontology underlying the whole I-X system, the <I-N-C-A> ontology (Tate, 2003), which is a generic description of a synthesis task (such as design, planning or configuration), dividing it into four major components: Issues, Nodes, Constraints, and Annotations. When used to describe processes, nodes are the activities that need to be performed in a course of action, thus functioning as the items in an intelligent to-do list. The other elements contain issues as questions remaining for a given course of action, information about the constraints involved and the current state of the world, and notes such as reports or the rationale behind items in the plan.

In <I-N-C-A>, both processes and process products are abstractly considered to be made up of a set of Issues which are associated with the processes or process products to represent potential requirements, questions raised as a result of analysis or critiquing, etc. They also contain Nodes (activities in a process, or parts of a physical product) which may themselves have sub-nodes making up a hierarchical description of the process or product. The nodes are related by a set of detailed Constraints of various kinds. Finally there can be Annotations related to the processes or products, which provide rationale, information and other useful descriptions.

<I-N-C-A> models are intended to support a number of different uses:

- for automatic and mixed-initiative generation and manipulation of plans and other synthesized artifacts and to act as an ontology to underpin such use;
- as a common basis for human and system communication about plans and other synthesized artifacts;

- as a target for principled and reliable acquisition of knowledge about synthesized artifacts such as plans, process models and process product information;
- to support formal reasoning about plans and other synthesized artifacts.

These cover both formal and practical requirements and encompass the requirements for use by both human and computer-based planning and design systems.

3.1 Issues

The issues in the representation may give the outstanding questions to be handled and can represent decisions yet to be taken on objectives to be satisfied, ways in which to satisfy them, questions raised as a result of analysis, etc. Initially, an <I-N-C-A> artifact may just be described by a set of issues to be addressed (stating the requirements or objectives). The issues can be thought of as implying potential further nodes or constraints that may have to be added into the specification of the artifact in future in order to address the outstanding issues.

Until recently, the issues used in I-X work had a task- or activity-orientation to them, being mostly concerned with actionable items referring to the process underway – i.e., actions in the process space. This has caused confusion with uses of I-X for planning tasks, where activities also appear as nodes. This is now not felt to be appropriate, and as an experiment we are adopting the gIBIS orientation of expressing issues as questions to be considered (Conklin, 2003; Selvin, 1999). This is advocated by the Questions–Options–Criteria approach (MacLean *et al.*, 1991) – itself used for rationale capture for plans and plan schema libraries in earlier work (Polyak and Tate, 1998) and similar to the conceptual mapping approaches used in Compendium (Selvin *et al.*, 2001).

For example, in the personnel recovery domain, the question “what is the location of the isolated person?” is an issue that needs to be addressed in order to develop the final recovery plan.

More formally, I is the set of unresolved issues in the current plan. An issue is represented by a syntactic expression of the form $l:M(O_1,\dots,O_n)$, where:

- l is a unique label for this issue,
- M is a symbol denoting a primitive plan modification activity, and
- O_1,\dots,O_n are plan-space objects, i.e. they are issues, nodes, constraints or annotations. The number of such objects, n , and the interpretation of each object in the context of the issue, will depend on the particular primitive plan modification activity represented by this issue.

Issues can be seen as primitive meta-level activities, i.e. things that need to be done to the plan before it becomes a solution to a given planning problem. This approach is inherited from O-Plan (Currie and Tate, 1991; Tate *et al.*, 2000a) and is also seen in planners such as OPIS (Smith, 1994). The most commonly found primitive meta-level activities carried

out by planners, but usually only implicit in their underlying implementation or internal plan representation, are:

- Achieving a goal (in classical planners): Let p be a world-state proposition and τ be a time point, then the primitive meta-level activity of achieving p at τ can be represented as the issue:

$$l_1:\text{achieve}(p,\tau)$$

- Accomplishing a complex activity (in HTN planners): Let $a \in N$ be a complex activity. Then the primitive meta-level activity of accomplishing a can be represented as the issue:

$$l_2:\text{refine}(a)$$

Here, `achieve` and `refine` are examples of symbols denoting primitive plan modification activities. Note that these symbols are not domain specific but specific to the planning process by which these types of issue are handled.

Issues can be either ‘negative’, in which case they can be thought of as flaws in the plan, or they can be ‘positive’, presenting opportunities.

3.2 Nodes

The nodes in the representation describe components that are to be included in the design. Nodes can themselves be artifacts that can have their own structure with sub-nodes and other <I-N-C-A> refinements associated with them. The node constraints (which are of the form “include node”) in the <I-N-C-A> model set the space within which an artifact may be further constrained. The “I” (issues) and “C” (constraints) restrict the artifacts within that space which are of interest. In the case where the design corresponds to a plan, nodes will represent activities that need to be performed.

For example, “locate the isolated person using beacon” is an activity in a rescue plan that could be introduced to address the example issue given above.

More formally, N is the set of activities (nodes) to be performed in the current plan, i.e., in this <I-N-C-A> object. An activity is a syntactic expression of the form $l:\alpha(o_1,\dots,o_n)$, where:

- l is a unique label for this activity,
- α is a symbol denoting an activity name, and
- o_1,\dots,o_n are object-level terms, i.e. they are either constant symbols describing objects in the domain, or they are as yet uninstantiated variables standing for such objects.

Time points constitute a special class of domain objects that are found as parameters of an activity. Specifically, two time points, one representing the beginning and the other the end of an activity, are often used as parameters.

In the context of I-X, nodes represent the object-level activities in the plan, i.e., things that need to be performed by some agent to execute the plan. As mentioned above, activities can be of two types from the perspective of the planner:

- Primitive activities: primitive activities can be carried out directly by an agent executing the plan. For example, in a search and rescue domain, the primitive activity of flying the aircraft `ac1` from location `loc1` to location `loc2` may be represented as:

$$l_3:\text{fly}(ac1,loc1,loc2)$$

- Complex activities: complex activities cannot be accomplished directly by the agent executing the plan but need to be refined into primitive activities. For example, the complex activity of rescuing an isolated person `ip` may be represented as:

$$l_4:\text{rescue}(ip)$$

In this example, `fly` is a primitive activity symbol and `rescue` is a complex activity symbol in their domain (note that activity symbols are always domain specific). It follows that there has to be an activity schema defined for the domain that has the name `fly` and describes when this activity schema is applicable and how it will change the world when applied, and there has to be a refinement defined in the domain that accomplishes a complex activity with the name `rescue` and describes how exactly it can be broken down into detail and accomplished.

Note that the set N of activities in the plan may contain both complex activities and the primitive activities that have been chosen to implement them.

3.3 Constraints

The constraints restrict the relationships between the nodes to describe only those artifacts within the design space that meet the objectives. The constraints may be split into “critical constraints” and “auxiliary constraints” depending on whether some constraint managers (solvers) can return them as “maybe” answers to indicate that the constraint being added to the model is okay so long as other critical constraints are imposed by other constraint managers. The “maybe” answer is expressed as a disjunction of conjunctions (i.e. using an and/or tree) 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 (Tate, 1995).

The choices of which constraints are considered critical and which are considered as auxiliary are decisions for an application of I-X, as are specific decisions on how to split the management of constraints within such an application. It is not pre-determined for all applications. A temporal, activity-based planner would normally have object/variable constraints (e.g. equality and inequality of objects) and some temporal constraints (maybe just the simple `before{time-point1, time-point2}` constraint) as the critical constraints. But, for example in a 3-D design or a configuration application, object/variable and some other critical constraints (possibly spatial constraints) might be chosen. Which constraints

are used depends on the nature of what is communicated between constraint managers in the application of the I-X architecture.

For example, constraints on the execution of a helicopter recovery plan could be that the “location of the isolated person is within range of the helicopter” and “the weather is safe for flying”.

More formally, C is the set of constraints that must be satisfied by the current plan (<I-N-C-A> object). A constraint is a syntactic expression of the form $l:c(v_1, \dots, v_n)$, where:

- l is a unique label for this constraint,
- c is a symbol denoting a constraint relation, and
- v_1, \dots, v_n are constraint variables, i.e., they can represent domain objects (including time points), variables in activities (which may have binding constraints attached).

Constraints represent the relations that must hold between the different objects related in the constraints for the plan to be executable. In the context of planning, the most commonly used constraints are of the following types:

- Ordering constraints: Let v_1, v_2 be variables in the plan representing time points. Then the constraint that v_1 has to be before v_2 can be represented as:

$$l_5:\text{before}(v_1, v_2)$$

- World-state constraints: Let p be a world-state proposition and v a variable representing a time point in the plan. Then the fact that p is a condition that has to hold at the time point represented by v , or the fact that p is an effect of an activity that holds at time point v can be represented respectively as:

$$l_6:\text{cond}(p, v)$$

$$l_7:\text{effect}(p, v)$$

- Variable binding constraints: Let v be a variable mentioned in some activity $a \in N$ and s be a constant symbol in the planning domain. Then the fact that v must take the value s can be represented as:

$$l_8:\text{value}(v, s)$$

These are just some of the constraint types that can be defined. The objects related to each other can be of different types. This is reflected by the domains of the constraint variables representing them. They can be world-state propositions as in conditions and effects, or they can be variables used in activities representing time points or other domain objects in the plan as in ordering and variable binding constraints.

3.4 Annotations

The annotations add additional, often human-centric information or design and decision rationale to the description of the artifact. They may be informal or they may adhere to some detailed syntax (which is not specified as part of <I-N-C-A>). They are normally expressed as “keyword = value” annotations. This can be of assistance in making use of products such as designs or plans created using this approach by helping guide the choice of alternatives should changes be required.

For example, the fact that the activity “locate the isolated person using beacon” was added to the plan to address the issue represented by the question “what is the location of the isolated person” is used to annotate the plan with some rationale information.

Annotations can be used to record arbitrary information about the plan (and the annotations form a part of this plan – hence the plan becomes, in some sense, self-descriptive). However, we want to discuss their specific use of annotating plans with one particular type of rationale, namely the rationale information that can be recorded by the planner during the planning process. More formally, let A be the set of annotations attached to the current plan. In the specific case, an annotation will be a syntactic expression of the form $l_a:r(l_p:O, l_m:M, O_1, \dots, O_n)$, where:

- l_a is a unique label for this annotation,
- r is a rationale predicate relating a plan-space object to other plan-space objects,
- $l_p:O$ is a labeled plan-space object that is part of the current plan, i.e., it is an issue, an activity, a constraint or an annotation,
- $l_m:M$ is an issue that was formerly in the plan and has since been resolved, i.e., it is a primitive meta-level activity that has been performed by the planner, and
- O_1, \dots, O_n are plan-space objects that may or may not be labeled.

An annotation of this type represents the fact that the plan-space object O was introduced into the plan as part of performing the plan modification activity M , and possibly involving other plan-space objects O_1, \dots, O_n . The rationale predicate r denotes the relationship between these objects and describes the justification for including O . Thus, the interpretation of such an annotation depends on the rationale predicate r used. The different labels are necessary to specify the exact object that is being referred to. This is necessary as there might be two activities in the plan which are identical except for the label. The following examples illustrate the use of rationale annotations of this form.

- Let $l_m:\text{achieve}(p, \tau)$ be an issue in the current plan and let $\alpha(o_1, \dots, o_n)$ be an activity schema defined in the domain that has an effect that unifies with p under the substitution σ . Suppose the planner introduces a new activity $l_p:\sigma(\alpha(o_1, \dots, o_n))$ into the plan to address the issue $l_m:\text{achieve}(p, \tau)$. Then the following annotation can be added to the plan to record the rationale for adding $l_p:\sigma(\alpha(o_1, \dots, o_n))$:

$$\text{naap} (l_p:\sigma(\alpha(o_1,\dots,o_n)), l_m:\text{achieve} (p, \tau) , p)$$

In this case *naap* is a rationale predicate that expresses that a new activity, the first argument, was introduced into the plan to address the issue of achieving some proposition (the second and third arguments respectively). Thus, the argument types for this particular rationale predicate are an activity $a \in N$, an issue $M \in I$ in which the plan modification activity symbol is *achieve*, and a world-state proposition. Furthermore, the last argument, the proposition p , must be the same as the one to be achieved in the plan modification activity, and it must be unifiable with one of the effects of the activity $a \in N$.

In this case, a second rationale annotation could be introduced in a similar fashion to express the fact that $l_p:\sigma(\alpha(o_1,\dots,o_n))$ has to be performed before the time point τ .

- Let $l_m:\text{refine} (a)$ be an issue in the current plan and let there be a refinement Δ defined in the domain that can be used to accomplish a under the substitution σ by refining it into, amongst other things, activities $\sigma(\alpha_1(o_1,\dots,o_n))\dots\sigma(\alpha_k(o_1,\dots,o_n))$. Note that the elements into which a is refined can together be seen as an <I-N-C-A> object, i.e. they can be issues, nodes, constraints and annotations. Suppose the planner uses Δ to refine a and this adds new activities $l_{p1}:\sigma(\alpha_1(o_1,\dots,o_n))\dots l_{pk}:\sigma(\alpha_k(o_1,\dots,o_n))$ to N to address the issue $l_m:\text{refine} (a)$. Then, the following annotation can be added to the plan to record the rationale for adding each $l_{pi}:\sigma(\alpha_i(o_1,\dots,o_n))$, $1 \leq i \leq k$:

$$\text{nadi} (l_{pi}:\sigma(\alpha_i(o_1,\dots,o_n)), l_m:\text{refine} (a) , \Delta)$$

(One such annotation would be added for each new activity α_i .) In this case *nadi* is a rationale predicate that expresses that a new activity, the first argument, was introduced into the plan to address the issue of refining some proposition in accordance with some particular refinement in the domain (the second and third arguments respectively). Thus, the argument types for this rationale predicate must be an activity $a \in N$, an issue $m \in I$, where the plan modification activity symbol has to be *refine*, and a refinement. Furthermore, the last argument, the refinement Δ , must be defined as accomplishing a complex activity that can be unified with a .

Similarly, if appropriate, analogous rationale annotations could be introduced to express the fact that other <I-N-C-A> elements of the refinement – such as issues or constraints – were also introduced as part of this refinement.

Rationale predicates of this type are usually specific to a type of issue. Hence, *naap* rationale will always relate to an *achieve* issue, and *nadi* rationale will always relate to a *refine* issue. However, there may be multiple rationale predicates that may be used with the same issue – which one is used will depend on how the planner actually resolved the issue. For example, achieving a proposition at some time point can be done by introducing a new activity before the time point or by maintaining the truth of the

proposition if it was true at another, previous time point. Thus, the relation between rationale predicates and issues is not one-to-one: issues need not always be resolved in the same manner.

Note too that this type of rationale, recording justifications for the inclusion of objects into the plan, is only one type of rationale that may be recorded in a plan. For example, we may want to record why a specific way of refining a plan was chosen among the various available options. While this type of information would be very useful to record, we believe that this will best be approached by use of a separate decision structure. It is in general not possible to extract useful knowledge of this kind from a search-based planning algorithm that tries out many possibilities and backtracks upon failure. At any choice point, there may be a large number of reasons why all the leaf nodes that are in the search space under the choice point represent failures in the search, and it may be hard to abstract these into meaningful rationale. However, there are also choice points in a search space where a decision is forced or made via user selection from open alternatives and it may be most useful to record this as part of the rationale for the plan.

4. The Co-OPR Problem

Personnel Recovery (PR) is the sum of military, diplomatic and civil efforts to achieve the recovery and reintegration of isolated personnel. During any military operation Joint Force Commanders and Staff are responsible for, and must be prepared to accomplish, the PR tasks throughout a specified operational area or else determine and accept the risk of not doing so (Joint Publication 3-50, 2005). In order to be prepared, the USJFCOM/JPRA Personnel Recovery Education and Training Center (PRETC) in Fredericksburg, VA, trains US military personnel in the execution of PR tasks. This training consists of classroom sessions in which the necessary knowledge is taught, and it consists of Command Post Exercises (CPX) in which the students have to perform PR tasks in a simulated, fictitious military operation.

4.1 The Personnel Recovery Domain

The running of a personnel recovery center can be divided into four main tasks:

- setting up;
- processing the air tasking order (ATO);
- information collection and management of rescue operations; and
- shift handover.

Setting up a personnel recovery center is a one-off task, performed to ensure that the center is ready and prepared to deal with incidents. This involves simple tasks like checking all the communication channels like telephones, fax and Internet chat are working properly, making sure the information about weather and current code words is up to date, etc. At the end of this process the JPRC has to report to the Joint Task Force Commander that it is ready and all set up. This process is performed according to a checklist that the director of each center has to complete (see Figure 1).



Figure 1: The JPRC director with checklists on his desks

The next task is the processing of the ATO. This usually involves getting a general overview of the planned operation and extracting the assets that are assigned to PR missions and their locations. Occasionally, assets will be relocated immediately, i.e. before any incidents occur, if it is felt that they are more likely to be useful in other locations. Note that these tasks are not necessarily sequential as the ATO may be released at any time.

The main task of the PR center is, of course, personnel recovery. To perform this task the PR center collects incoming reports about incidents as they occur. Reports may come in through various channels and from various sources. The first step thus consists of accurately recording the incident reports as they come in. This may be particularly challenging if the information comes in over noisy channels, e.g. a bad radio connection.

With an accurate record of the report, incident information has to be assigned to an existing incident, or a new incident has to be established. This sensemaking is often very difficult as reports tend to be incomplete when they come in and different sources observing the same incident often have very different views of what exactly they have seen.

In addition to incident reports coming in, the PR center has to be pro-active in collecting information, and collect information about the personnel to recover from their unit and other sources, including the “Isopreps”.

The PR center has to track all the information relating to the various incidents that are being reported, and once sufficient information is available, they have to plan, launch and monitor a rescue mission. The planning may also lead to no mission being launched. The

different options for recovery that should be considered are described in (Joint Publication 3-50, 2005). Again, checklists are available to ensure all the necessary actions are taken by the PR center. Note that the ATO processing tasks need to be performed by all rescue centers, which adds to the complexity of the problem.

The final task is the shift handover. Incidents often cannot be resolved in a single shift and when the next shift comes in a briefing needs to be prepared that informs the new shift about all the ongoing and completed rescue operations.

4.2 Command Post Exercises

Command Post Exercises (CPX) are performed at the Personnel Recovery and Training Center (PRETC) as part of the Personnel Recovery course. The course consists of classroom teaching sessions and the CPX in which students are divided into groups, each group playing the role of a rescue center that has to respond to some incidents that are emulated by the trainers.




Figure 2: Generic Scenario Map

In training, the context for the incidents and rescue missions that need to be launched is a generic military operation, which is set in an area corresponding to the generic map shown in Figure 2. In the figure, Country-1 represents the country that is being assisted and that is in conflict with its immediate neighbors. A shared coastline makes the involvement of the Navy possible. Country-1 also has rural as well as urban areas that

make for an interesting variety of potential incidents. Finally, a neutral country provides an overseas base that may play a role.

The students are divided into four groups and placed in different rooms where they act out the activities performed by the different Rescue Coordination Centers (RCCs). In the CPX the Joint Personnel Recovery Center (JPRC) is co-located with the Air Force RCC. All other agents are role-played by the trainers at the PRETC. An overview of the organizational relationships between the different agents is given in Figure 3.




Figure 3: Organization of Agents in the Scenario

4.3 Requirements for I-X

In observing the Command Post Exercises at the PRETC, we have identified a number of ways in which I-X technology and the user interfaces or tools we can provide may be able to support those involved in search and rescue. I-X uses in a JPRC/RCC could include:

- Communications
 - Simple Chat
 - Structured chat
 - Information sharing
- Task Support
 - Checklists

- To do list
 - Progress reporting
 - Plan option aids
- Whiteboards
 - Incident (see Figure 4)
 - Weather/Codes/Info
 - Assets (see Figure 5)
- Mapboards
 - Terrain and GIS features (see Figure 6)
 - Routes, Restricted Operations Zones, etc.
 - Town and road plans
 - Sketch maps
- Web Resources
 - Fact Book
 - Phone List
 - Codes
- Mission Folders
 - Attachments



Figure 4: Incident board towards the end of a CPX

Many of these features were already supported in the I-X framework generically. However, the JPRC and RCCs make heavy use of wall mounted whiteboards, maps, overlays on maps, and pin board material such as codes, phone lists, etc. We have implemented whiteboard and map orientated "viewers" that can all simultaneously share the same state in a single panel for display and sharing. We also explored ways in which the state underlying specific views can easily be shared with other users and I-X panels, and ways in which variances between the incoming and current believed state on any panel can be highlighted, such that the changes can initiate issues, activities, constraints or notes that need to be incorporated into the local plan.




Figure 5: Asset board used during CPX

We have also created a "white cell" support panel to assist the trainers in a CPX (see Figure 7). This will allow:

- Driving a simulation of the world in which the training takes place, including starting and stopping moving assets such as fuel tankers, trucks, planes and ships.
- Setting the world clock as seen by all other I-X panels and users to a simulated time.
- Allowing master scenario event lists (MSELs) to be input and assisting in driving the simulation
- Assisting in logging, noting training issues for reporting back, etc.

All these features are now part of the I-X framework and can be included in any I-X application. The first application to use them is the Co-OPR application described next.



Figure 6: Map used during CPX




Figure 7: Busy White Cell during CPX

5. Developing I-X and the Co-OPR Application

The user interface to the I-X system, the I-X Process Panel, shows four main parts that reflect the four components of the <I-N-C-A> ontology described above. They are labeled “Issues”, “Activities”, “State”, and “Annotations”, as shown in Figure 8.




Figure 8: Initial state of the JPRC director’s I-X Process Panel

In the case of the artifact to be synthesized being a course of action, the nodes that will eventually make up the artifact are activities, and these play the central role in the view of an I-X panel as an intelligent to-do list. Users can add an informal or formal description of a task to be accomplished to the activities section of the panel where it will appear as the description of that activity. Each activity consists of four parts listed in the four columns of the activities part of the panel:

- Description: This can be an informal description of a task such as “do this” or it can be a more formal pattern consisting of an activity name (verb) followed by a list of parameters such as
(deploy ?team-type)
where the words preceded by a question mark are variables that need to be bound before the task can be dealt with.

- Annotation: This can be used to add arbitrary pieces of information to a specific activity.
- Priority: This defines the priority of the activity. Possible values are Highest, High, Normal, Low, or Lowest.
- Action: This field contains a menu that gives the various options that are available to deal with the activity and is the focus of intelligent task synthesis in I-X Process Panels.

The Action field allows the user to mark the task as “Done”, which corresponds to ticking off an item in a to-do list. Other options that are always available for this field are “No action”, the default value until the task has been dealt with, or “N/A” if the activity does not make sense and is “not applicable” in the current context.

The entries in the Action menu related to an activity are determined by activity handlers. These handlers are modules that can be plugged into the I-X system and define ways in which activities can be dealt with. If an activity handler matches an activity it can add one or more entries to the activity’s action menu. The most commonly used activity handler in the context of HTN planning adds “Refine” items to this menu, and this is the point where the to-do list becomes intelligent.

Instead of just being able to tick off an activity, users can use the knowledge in a library of standard operating procedures 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. This mechanism supports the user in two ways:

- The library of standard operating procedures may contain a number of different refinements that all match the present activity. All of the applicable procedures are added to the action menu by the activity handler, thus giving the user a comprehensive and quick overview of all the known standard procedures available to deal with this task.
- When a refinement for an activity is chosen, the I-X Process Panel shows all the sub-activities as new items in the to-do list. This ensures that users do not forget to include sub-activities, a common problem especially for infrequently applied procedures.

Both of these problems become more severe when the user is under time pressure and lives depend on the decisions taken.

Note that the intelligence of the to-do list comes in through the underlying HTN planner that finds applicable refinements in the library and, on demand, can complete a plan to perform a given task automatically, propagating all constraints as it does so. Equally important, however, is the knowledge contained in the library of standard operating

procedures. From the perspective of the user this means that I-X can actively suggest ways of performing an activity on the to-do list, or I-X can allow the user to explore the set of options currently available.

5.1 Core I-X Components

Standard operating procedures describe the knowledge underlying the intelligent to-do list. The formalism is based on refinements used in HTN planning and will be explained next. However, users are not expected to learn this formalism. Instead, they can use a domain editor and its graphical user interface to define the library of procedures.

What are known as standard operating procedures to domain experts are called refinements or methods in HTN planning (Ghallab *et al.*, 2004). Methods formally describe how a task can be broken down into sub-tasks. The definition of a method consists of four main parts:

- Task pattern: an expression describing the task that can be accomplished with this method;
- Name: the name of this method (there may be several for the same task);
- Constraints: a set of constraints (e.g. on the world state) that must hold for this method to be applicable; and
- Sub-task network: a description of the sub-tasks into which this method refines the given task.

The task pattern of a method is used for matching methods to items in the activity list. If the task pattern matches the activity the method will appear in the action menu of the activity in the panel as a possible expansion. This is also where the name of the method will be used: the menu displays an entry “Refine using *name*” where *name* is the name of the method. In this way, the user can easily distinguish the different options available. The constraints are used to decide whether the method is applicable in the current world state. If the constraints are satisfied, the method can be selected in the action menu, otherwise the unsatisfied constraints can be seen as issues, namely sub-goals that need to be achieved in some way. Finally, the network contains the list of sub-tasks that will be added as activities to the panel when the method is selected. The ordering constraints between sub-tasks are used to show in the interface those sub-tasks that are ready for tackling at any given time.

5.2 The Domain Editor

Figure 9 shows an example of the I-X Domain Editor for defining standard operating procedures (SOPs). The definitions of SOPs form part of a coherent “domain model” which describes information relevant for a particular application. The panel on the left lists all the currently defined procedures by name, and the task pattern they match. One, called “rescue-generic”, is shown as being edited. There are a number of views

available to edit a refinement. The one shown is the graphical view, which shows all the direct sub-tasks with their begin- and end-time points. Arrows between these activities indicate temporal ordering constraints, for example, the activity “locate ?ip” cannot be started before “report ?ip” has been completed. However, the activities “support ?ip” and “recover ?ip” can then be performed in parallel. Other views show the conditions and effects that can be defined for refinements.




Figure 9: Initial state of the JPRC director’s I-X Process Panel

The I-X Domain Editor, I-DE, already included intuitive editing facilities for activities of SOPs and their refinement into detail. During the project, it has had an overhaul and a few features added:

- There is an object modeling facility, which lets the modeler specify classes (or types) of objects that are related to the activities in the process models. These classes are organized into a hierarchy and they can have properties (or attributes);
- The object modeling facility is integrated with the activity modeling support. This helps the modeler to keep track of modeling decisions and to keep world-state constraints consistent;
- The form-based panels now have facilities for adjusting the size of components, including a hide/show option which will alert the user if hidden components contain information in the current model;

- The graphical view of process models has been improved and now has an additional layout option that is more suitable for viewing models with many, sparsely related nodes;
- There is a search facility for names of refinements (activity specifications) and names of object classes;
- The ordering of the list of all defined refinements has been improved - they can be viewed in alphabetical order or in the order in which they were entered;
- The grammar panel, which shows terms and formal patterns used, has been improved. In addition to showing which patterns are used in nodes and issues, it now also shows which constraint types are used in the domain.

The I-X Domain Editor is tightly integrated with the I-X framework. SOPs can be defined in advance, or they can be defined while rescue operations are in progress. New SOPs are available to process panel users as soon as they are “published” in the SOP library, adding to the options available for dealing with tasks.

5.3 White Cell Support: I-Sim

The I-Sim tool integrates the functionality of a discrete event simulator and a number of process-level simulators into the I-X framework. Usually there will be only one I-X Process Panel that has access to this tool, e.g. the panel representing the instructors in a training scenario. In such a set-up the I-Sim tool gives the instructors the necessary control over the simulation, while other panels that do not have access to the tool may only see the result of the simulation. As a result, control over the development of the scenario is centralized, whereas I-X panels supporting responders are expected to be distributed. An example I-Sim tool is shown in Figure 10, displaying events taken from the Co-OPR application during experiment E (described in section 6).

Status	Time	Thread	Event
COMPLETED	07/06/00 15:00	BUCKET-33	Activity[send-activity E2C-ZEN Navy-RCC (note-incident-report "BUCKET 31 REPORTS BUCKET 33 TOOK A SAM
COMPLETED	07/06/00 15:02	BUCKET-33	Activity[send-activity E2C-ZEN Navy-RCC (note-incident-report "BUCKET 33 AND 34 ARE NORTHBOUND TRYIN
COMPLETED	07/06/00 15:04	GOPHER-63	Activity[send-activity E2C-ZEN Navy-RCC (note-incident-report "GOPHER 64 REPORTS GOPHER 63 IS SEAHAW
COMPLETED	07/06/00 15:08	BUCKET-33	Activity[send-activity E2C-ZEN Navy-RCC (note-incident-report "BUCKET 33 HAS EJECTED AT
COMPLETED	07/06/00 15:10	GOPHER-63	Activity[send-activity E2C-ZEN Navy-RCC (note-incident-report "GOPHER 63 IS ON GROUND WITH INJURED AR
COMPLETED	07/06/00 15:16	BUCKET-33	Activity[send-activity E2C-ZEN Navy-RCC (note-incident-report "BUCKET 34 IS OVERHEAD BUCKET 33 WITH VIS
COMPLETED	07/06/00 15:25	BUCKET-33	Activity[send-activity E2C-ZEN Navy-RCC (note-incident-report "BUCKET 34 IS DEPARTING FOR AR")]
WAITING	07/06/00 15:30	BUCKET-33	Activity[send-activity E2C-ZEN Navy-RCC (note-incident-report "FELIX 43 REPORTS CONTACT WITH BUCKET 33
WAITING	07/06/00 15:34	GOPHER-63	Activity[send-activity E2C-ZEN Navy-RCC (note-incident-report "KNIGHT 13 REPORTS THAT GOPHER 63 IS IN H
WAITING	07/06/00 15:40	BUCKET-33	Activity[send-activity E2C-ZEN Navy-RCC (note-incident-report "FELIX 43 PASSES COORDINATES FROM BUCKE
WAITING	07/06/00 16:04	GOPHER-63	Activity[send-activity E2C-ZEN Navy-RCC (note-incident-report "KNIGHT 13 PASSES THAT GOPHER 63 REPORT

Figure 10: The I-Sim Tool during the simulation (at 15:27 simulated time)

The tight integration of the I-Sim tool with the I-X framework is achieved through use of a shared model of activity, <I-N-C-A> (see section 3), that is maintained in the panel to which the tool belongs. Thus, I-Sim can extract all required information to start the simulation directly and without any user interaction. Furthermore, updated state

information can be written back to this shared model, thus allowing all the attached viewers to display this information in a consistent way, i.e. shared state information is visible in all other tools. Finally, state can also be shared with other panels allowing students selective access to simulation results. The impact of these design choices is that shared models do not give rise to undesired inconsistencies that could disrupt a training exercise and increase maintenance overheads.

For a training session the scenarios have to be scripted in advance, with each incident to be addressed corresponding to a sequence of events. The I-Sim tool allows the user to load one or more such incident scripts at a time, thus allowing the combination of different incidents. Events are associated with the incident they were loaded from and the tool displays this as a ‘thread’ running through the scenario, information that is not available to the trainees. Note that this means that multiple incidents and their simulation can take place in parallel. Each event is ‘timed’ and the time represents the start time of the event relative to the start of the incident. The ‘actual’ start time of the simulation has to be defined when an incident is loaded. Additional incidents that are added into the scenario after the start of the simulation will begin at the time of loading. This gives instructors the flexibility to extend the scenario quickly and according to desired training outcomes.

Another aspect of the simulation that can be controlled through the I-Sim tool is simulated time. At the beginning of the simulation the time acceleration factor can be set to determine how fast the simulation will progress. This can be any positive number, including numbers less than one (meaning simulation time passes slower than real time). The time acceleration factor currently can only be set at the beginning of the simulation. However, the simulation can be paused and resumed at any point. Jumping forward and backwards is not possible as this has technical implications for the process-level simulators used, i.e. they too need to support such jumping, and this may not always be feasible. It is possible to skip forward to the next event, but this means that this event will be injected into the scenario now rather than at its scheduled time. The intended result of these time-management facilities is a reduced cognitive load on the instructors.

The current status of events is always shown in the I-Sim tool window, thus giving instructors a quick overview.

The aim of the above controls is to give the instructors flexibility. This is achieved by giving instructors the option to create a scenario on the fly from pre-defined components and add more incidents at a later point if required. Furthermore, instructors have some limited control over simulated time by setting the time acceleration factor and pausing and resuming the simulation to adapt to the trainee’s progress.

5.3.1 Discrete Event Simulation

The discrete event simulator provided by the I-Sim tool is tightly integrated with the I-X framework. The simulation of actions is based on the same representation in the domain model used by all the components and tools in the framework and can be edited using the integrated domain editor. This domain model is based on a hierarchical task network

(HTN) planning (Ghallab et al., 2004) representation and consists mainly of refinements and actions. Actions are primitives in the sense that the planner has no further knowledge about the internal structure of these activities; it only looks at actions at the level where they are described in terms of preconditions and effects. The choice of an HTN model for activity representation allows for a description of unfolding incidents at various levels of abstraction, e.g. strategic, tactical and operational event and action sequences. HTN planning also has been used successfully in a number of realistic domains.

When I-Sim simulates an action, its default behavior is to use only the preconditions and effects of the action. First, the preconditions are evaluated against the current state of the world as represented in the agent's current <I-N-C-A> model. For the simulation to be able to start, the preconditions have to be uniquely satisfiable, that is, there must be exactly one way to bind all the remaining variables in the preconditions. If there is no way to bind the variables consistently, the action cannot be applied; if there is more than one way, this is still a choice that the simulator cannot make. If an action can be simulated, the effects of the action are asserted just before the simulation is terminated. This ensures that the simulation is consistent with the domain model to which instructors and students have access during the exercise.

Actions may be instantaneous, e.g. sending a message to another panel notifying it that an event has been observed, or they may take time as specified in the action description, e.g. flying a rescue helicopter from one location to another. If a time is specified, the default behavior of the simulator is to wait for the specified amount of time before the effects are asserted. This does not result in any change of world state during the execution however, which is simply due to the fact that there is no further information available to the discrete event simulator.

5.3.2 Process-level Simulation and Time

Discrete events are often sufficient from a controller's perspective, but finer grained models are required to simulate the development of a situation, e.g. a fire, an oil spill, or the spreading of a virus. As opposed to discrete event simulations, process-level simulators, often based on mathematical models, emulate a world that appears to be changing continuously.

In I-Sim, process-level simulators can be attached to actions in the domain model. When asked to simulate an action, I-Sim first verifies the preconditions as described above. Then, instead of simply waiting, it executes a process-level simulation which may (periodically) assert facts into the simulator's state to simulate the continuing change of the world. Finally, I-Sim asserts the effects of the action as for the discrete event simulation. As a result the simulation progresses naturally and continuously. Furthermore, the use of mathematical models means that events and resulting world states may be more realistic and contain more detail than a human could generate. However, I-Sim has no control over the quality of these process-level simulations.

Note that I-Sim allows for multiple simulations to be run in parallel. Interference or interaction between different simulators is coordinated via the shared <I-N-C-A> model.

That is, constraints asserted by one simulator are broadcasted to all other simulators, allowing them to immediately react to the changed state.

5.4 Other Tools and Viewers

As activities are the nodes that make up a course of action, it is only natural that the activity part of the I-X Process Panel forms the centre of attention for our view of I-X as an intelligent to-do list. We have implemented a cut-down interface called Post-IX which shows only this part of the panel (and so provides a minimal or ‘entry level’ interface to the system). We shall now briefly describe the other parts of a panel and how they are used.

World state constraints are used to describe the current state of the world. Essentially, these are a state-variable representation of the form “*pattern = value*” allowing the user to describe arbitrary features of the world state. Usually, these features describe aspects of objects related to the activities to be performed. The world state constraints are displayed in the I-X Process Panel in the constraints section. However, it is not expected that users will find this list of facts about the world style representation very useful. Thus, I-X allows for the registration of world state viewers that can be plugged into the system. For example, BBN Openmap has been used in a number of applications to provide a 2-D world map with various features, showing – for example – locations of relevant objects. 3-D virtual reality viewers have also been explored. Most importantly, such world state viewers can be automatically synchronized with the world state constraints such that icons in the map always represent current positions of the entities they represent. Constraints are propagated and evaluated by constraint managers that are plugged into the I-X system.

Issues can be seen as a meta to-do list: instead of listing items that need to be done to deal with an emergency in the real world, they list the questions or outstanding items that need to be dealt with to make the current course of action complete and consistent. Often, these will be flaws in the current plan, but they can also be opportunities that present themselves, or simply facts that need to be verified to ensure a plan is viable. Issues can be either formal, in which case registered issue handlers can be used to deal with them just like activity handlers deal with activities, or they can be informal.

Annotations are used for descriptive elements, such as comments about the course of action as a whole, and are stored as “*keyword = value*” patterns.

So far we have described I-X as a tool for assisting a single person in organizing and executing the response to an emergency. However, I-X can also support the coordination of the response of multiple agents. I-Space is a tool in which users can register the capabilities of other agents. These capabilities can then be used from an I-X panel through inter-panel communication. Augmented instant messaging can be used to directly communicate with other responders via their panels.

5.4.1 I-Space

Every I-X panel can be connected to a number of other I-X agents. Each I-X agent represents an agent that can potentially contribute to the course of action taken to respond in an emergency. The I-Space holds the model of the other agents and can be managed with a simple tool as shown in Figure 11.




Figure 11: The I-Space Tool: Relations

Associated with each agent are one or more communication strategies, which define how messages can be sent to this agent. By default, a built-in communication strategy simply sends XML-formatted messages to a given IP-address and socket. Alternatively, a Jabber strategy (Jabber, 2003; 2006) is available for using an instant messaging mechanism for communication. New communication strategies can be added to communicate with agents implemented using different frameworks.

Usually, users will be less concerned with the question of how communication takes place as long as the system can find a way, but more with the relationships between the different agents in the I-Space. Within an organization a hierarchical structure is common, so collaborating agents are usually either superiors or subordinates. They can also be modeled as peers, which is also how agents from other organizations can be described. If the agent to be integrated into the virtual organization is a software agent it is described as a (web) service. Finally, a generic relation “contact” is available, but it does not specify what exactly the relationship to this agent is.

5.4.2 Agent Capabilities

At present there is only a relatively simple capability model implemented in I-X. The idea behind this model is that activities are described by verbs in natural language and thus, a task name can be used as a capability description. Parameter values are currently not used to evaluate a capability. Each agent is associated with a number of capabilities that can be called upon.

Agent	Verb Capabilities
JPRC-Controller	note-incident-report, display, note, plan-and-exe...
Navy-RCC	check-ready
JPRC-Watch-Supervisor	note-incident-report, display, update-incident-inf...
SOF-RCC	check-ready
JTFC	do-nothing, establish, ensure, receive
white-cell	do-nothing
Army-RCC	check-ready

Figure 12: The I-Space Tool: Capabilities

In the future it will be possible to use a much more sophisticated model. The problem with more complex representations is often that matching capabilities to tasks can be computationally expensive, and when the number of known capabilities becomes large, this can be a problem, which is why the current model is so simple. On the other hand, capabilities can often only be distinguished by a detailed description. One approach to this trade-off is to provide a representation that is flexible, allowing for a more powerful representation where required, but retaining efficiency if the capability description is simple (Wickler, 1999).

Conceptually, the description of a capability is similar to that of an action, which is not surprising as a capability is simply an action that can be performed by some agent. A capability description essentially consists of six components:

- **Name:** The name of a capability corresponds to the verb that expresses a human-understandable description of the capability.
- **Inputs:** These are the objects that are given as parameters to the capability. This may be information needed to perform the capability, such as the location of a person to be recovered, objects to be manipulated by the capability, such as paper to be used in a printing process, or resources needed to perform the capability.
- **Outputs:** These are objects created by the capability. Again, this can be information such as references to hospitals that may have been sought, or they can be new objects if the capability manufactures these.
- **Input constraints:** These are effectively preconditions, consisting of world state constraints that must be true in the state of the world just before the capability can be applied. Usually, they will consist of required relations between the inputs.
- **Output constraints:** These are similar to effects, consisting of world state constraints that are guaranteed to be satisfied immediately after the capability has been applied. Usually, they will consist of provided relations between the outputs.
- **I-O constraints:** These cross-constraints link up the inputs with the outputs. For example, a prioritization capability might order a given list of options according to some set of criteria. A cross-constraint, referring to both the situation before

and after the capability has been applied, is necessary to say that the given list of options and the prioritized list contain the same elements.

This capability model can be used to describe the abilities of real-world agents that ultimately must be deployed to do things, or for software agents that provide information that can be used to guide the activity in the physical world.

5.4.3 Handling Activities through Task Distribution

From a user's perspective, task distribution is integrated into the user interface through the "action" menu in the activities part of the panel as just another option available to deal with an activity. The agent relationship is used to determine in which way the activity can be passed to another agent, for example, if the other agent is a subordinate the activity can simply be delegated to the agent.

The capability model is used to filter the options that are listed in the action menu. Currently there is the option of specifying no capabilities for an agent in which case the agent will always be listed. If there is a list of capabilities associated with an agent then these options will only be listed if there is an exact match of the verb capability.

5.4.4 Structured Instant Messaging

Another tool that is widely used for the coordination of efforts in response to an emergency is instant messaging. Like a to-do list, it is very simple and intuitive, but it lacks the formal structure that is needed when the scale of the event that needs to be addressed increases. As for the to-do list, I-X builds on the concept of instant messaging, extending it with the <I-N-C-A> ontology, but also retaining the possibility of simple and informal messages. Thus, users can use structured messaging when this is appropriate, or continue to use unstructured messaging when this is felt to be more useful.

The structured version can be activated by selecting a message type: issue, activity, constraint or annotation, rather than a simple chat message. An <I-N-C-A> object with the content of the message will then be created and sent to the receiving I-X agent. Since all messages between agents are <I-N-C-A> objects, the receiving agent will treat the instant messenger generated message just like any other message from an I-X panel, e.g. the message generated when a task is delegated to a subordinate agent. In this way, structured instant messaging can be seamlessly integrated into the I-X framework without losing the advantages of informal communications.

5.4.5 Documentation

To make the I-X framework more usable, the I-X documentation has had a major overhaul. There is now a set of guide documents that supports developers of I-X applications, users of process panels (i.e. users of I-X applications), and modelers who specify standard operating procedures:

- The I-X Process Panels User Guide, describing the basics of I-X process panels and giving an overview of the tools associated with them;

- The I-X “Configurer's Guide”, describing methods for building I-X applications using step-by-step examples;
- The I-X Domain Editor Guide, describing I-DE and how I-X models are created;

There are two instructive demonstrators, I-Demo-Basic and I-Demo-Cooperation, which illustrate the components and use of I-X applications, and there is a set of web pages associated with the I-X releases.

5.5 The Co-OPR Application

Personnel recovery teams must operate under intense pressure, taking into account not only hard logistics, but "messy" factors such as the social or political implications of a decision. The Collaborative Operations for Personnel Recovery (Co-OPR) project has developed decision-support for sensemaking in such scenarios, seeking to exploit the complementary strengths of human and machine reasoning (Buckingham Shum *et al.*, 2006; Tate *et al.*, 2002). Co-OPR integrates the Compendium sensemaking-support tool for real time information and argument mapping, with the I-X artificial intelligence planning and execution framework to support group activity and collaboration. Both share a common model for dealing with issues, the refinement of options for the activities to be performed, handling constraints and recording other information. The tools span the spectrum from being very flexible with few constraints on terminology and content, to knowledge-based relying on rich domain models and formal conceptual models (ontologies). In a personnel recovery experimental simulation of an UN peacekeeping operation, with roles played by military planning staff, the Co-OPR tools were judged by external evaluators to have been very effective.

The first step in developing an I-X application consists of deciding which agents to support. For the Co-OPR application it was clear that the most important agent is the JPRC which coordinates the efforts of the different RCCs. Three roles in the JPRC of particular importance are that of the director, who has to manage the centre and make sure everything that needs to be done gets done, the watch supervisor, who is in charge of sensemaking and maintaining the information related to the various incidents on shared displays (white boards in a CPX), and the controller who manages the recovery assets and has to come up with plans for individual recovery missions. Three I-X Process Panels were used to support these roles. Only the controller's panel had the I-X option management facility enabled (not described here) which can be used to explore possible courses of action and compare different recovery plans (see figures in Appendix A). Other RCCs were supported by a single panel only.

Another agent that plays an important role in the training scenario is the “white cell” that drives the scenarios and simulates the events that lead to the incidents the JPRC has to deal with. An I-X Process Panel was used to support this role by allowing for an additional communication channel with the other agents supported by panels. Finally, some other agents that play only minor roles in the different scenarios were included, e.g. the Joint Task Force Commander (JTFC) that has to give authorization for certain missions. Thus, the organization of the agents in the application is as shown in Figure 3.

To implement the task support it was necessary to model a set of standard operating procedures that could be used as refinements in the I-X Process Panel as described above. The refinements used were derived from two sources. Firstly, the U.S. manual for Personnel Recovery 0 was used as a base for knowledge engineering. Secondly, the checklists used by the PRETC during a CPX were imported into I-X using an experimental import facility. However, the resulting model still required some knowledge engineering, in this case using the I-X Domain Editor.

The application so far can be considered as a simple customization of I-X for the task at hand. However, during the real CPX a number of other tools were used to support the JPRC and other RCCs. It was felt that these were needed for the I-X application too, and corresponding extensions to I-X were implemented.

Whiteboards: The JPRC and RCCs make heavy use of wall mounted whiteboards, maps, overlays on maps, and “pin-board” material such as codes, phone lists, etc. We have implemented whiteboard and map orientated “viewers” that can all simultaneously share the same state in a single panel for display and sharing. We are now exploring ways in which the state underlying specific views can easily be shared with other users and I-X panels, and ways in which variances between the incoming and currently believed state on any panel can be highlighted, such that the changes can initiate issues, activities, constraints or notes that need to be incorporated into the local plan.

White-Cell Support: We have created a white cell support panel to assist the trainers in a CPX. This will allow:

- Driving a simulation of the world in which the training takes place, including starting and stopping moving assets such as fuel tankers, trucks, planes and ships.
- Assisting in logging, noting training issues for report back, etc.

6. Experiments and Evaluation

The objective of experiments C thru E was to demonstrate that features known to be useful in personnel recovery operations, as identified by task analyses performed by USJFCOM, could be provided by I-X in general and by the Co-OPR application specifically. Furthermore, potential benefits for more flexible operations could result. An additional aim was to provide support that could benefit the training staff involved in personnel recovery.

The aim of each of these experiments was to emulate one half-day round of a CPX usually held at the PRETC. Experiments A and B were held during an initial phase of the Co-OPR project. CPX exercises were observed by the project team and researchers in October 2005 (see figures in section 4), and materials were provided to enable research and experimentation. The experimentation was designed to demonstrate and stress the I-X technology components in response to various individual events in sample missions and events provided by JPRA/PRETC. Following a number of progressively more realistic trials held in AIAI's experimental Emergency Response Coordination Center (e-RCC) during April and May 2006, the initial Co-OPR experiment C was conducted on 1st June 2006 following trials of the experimental setup and Internet collaboration software on 30th May 2006. This was followed by experiment D on 9th October 2006 and finally, experiment E on 27th April 2007. This evaluation section presents the results of these experiments.

6.1 Experimental Set-Up

The experiments all concentrate on a number of personnel recovery incidents that arise during a military operation called Operation Able Sword, which nominally takes place in Country-1 (see Figure 2) on some given dates in June/July 2000. Each experiment covers setting up a JPRC which is co-located with an Air Force RCC and checks with associated RCCs for the Navy, Army and Special Operation Forces (SOF) that they are ready for operations, prior to declaring to the JTFC that the JPRC is active. Incidents of various kinds are dealt with, and a final operation is to prepare a shift change briefing. The aim of the experiments was to allow for an evaluation of the I-X technology as a support tool for both trainers and trainees. At this stage the evaluation was performed with an observer from USJFCOM/J9.

Figure 49 to Figure 55 illustrate the progress during the experiment from the point of view of the JPRC. The double-screen setup was projected in the room such that all members of the JPRC could see the shared information displays, e.g. the electronic whiteboards. Internet application sharing technology was used to let observers remotely view the operations.

6.2 Execution of the Experiments

The evaluation focused on the cognitive tasks that the JPRC director, JPRC watch supervisor, and JPRC controller performed when working in tandem to respond to the incidents that came into the JPRC as an emergency response coordination centre. This evaluation was necessarily limited in that, without a corresponding analysis of the performance with the current in-situ systems and (manual) processes, a comparative assessment of the influence and value of the I-X system is not possible. However, an analysis of the results provides some interesting insights.

The evaluation methodology was straightforward. The director, watch supervisor and controller roles were played by members of the I-X development team. In addition to being familiar with the use of I-X systems and with its deployment for this particular domain, the participants had gained a basic competence in the objectives, approaches and working practices of the JPRC through observation and completion of basic training courses. An independent observer, a non-participant in the exercise (and also a member of the I-X team), was to observe their behavior (aided and augmented by self-reporting by the subjects), determine the nature of the task that was currently being performed and the time at which the task began and ended, plus any additional comments or observations. In addition, the exercise was being video-taped, which would allow a retrospective analysis, perhaps with the assistance of the ‘director’, ‘watch supervisor’ and ‘controller’, of any points during the exercise where the precise nature of the immediate task in hand was not clear. Importantly, the experiment was also observed by a member of the sponsoring organization familiar with personnel recovery and with systems evaluation. During experiment C this was done remotely using Internet collaboration and desktop sharing tools including video teleconferencing.

Once this was done, in an attempt to generalize the various tasks that had been performed, where appropriate each task was classified into one of several course-grained ‘cognitive categories’, namely:

- information-gathering: these tasks involved searching for information that was required before the overall activity of the JPRC could be moved forward. In certain cases, this may involve looking up information in on-line databases, or paper-based manuals, or it may involve (simulated) phone-calls to appropriate colleagues.
- sense-making: these tasks involved an analysis and interpretation of information with the aim of understanding the problem, enumerating the different options that were available, listing the pros and cons of possible courses of action, and so on.
- decision-making: these tasks involved the subject making a clear choice among competing possible activities that would serve to achieve the objectives of the JPRC by effecting activity in other agents and then enacting this activity. So, deciding to send a rescue helicopter to a particular destination and issuing the appropriate orders would be an example of a decision point, whereas deciding to

look at a map would not, since it has no effect on other agents (and, instead, would probably be an instance of information-gathering).

- housekeeping: these tasks involved the initial set-up of the JPRC environment, documentation of decisions, logging of calls, etc.

The first three of these categories (the housekeeping category being an artifact arising from the need to manage the JPRC and the ‘paperwork’ it generates) emerge from consideration of several different ‘best practice’ approaches to command and control and decision-making in general. For instance, Boyd’s well-known OODA loop (Osinga, 2006) – Observe, Orient, Decide, Act – can be seen to correspond with these three tasks: observe is essentially synonymous in this context with information-gathering and orient is synonymous with sense-making, and since enacting most of the decisions that are taken by the JPRC staff is done by issuing commands to others (i.e., in I-X terms, sending an activity to another agent) and this is done on the click of a mouse button, for our analysis we do not attempt to differentiate the decide and act activities, but instead we conflate these two OODA tasks into the single decision-making category. Similarly, Wohl’s SHORe (Stimulus, Hypothesis, Option, Response) framework (Wohl, 1981) can be seen to be analogous to our categories, with stimulus (Wohl’s shorthand term for the information correlation and fusion phase) corresponding to information-gathering, hypothesis (Wohl’s situation analysis phase) corresponding to sense-making, and the option and response phases being conflated into the single decision-making task (for the same reason outlined above).

The correspondence between these different models is summarized in Table 1. The fundamental concept underlying all of these models is that a methodical approach to each cycle of the command and control ‘loop’, based on assembling information, interpreting that information, appraising possible courses of action and making and enacting decisions should lead to clear, consistent, and – ultimately – correct behavior in situations where the pressure is great and time is short. Our empirical hypothesis here is that the use of the I-X system can encourage its users to adopt such a methodical approach to their task.

Table 1. Comparison of different Command-and-Control frameworks as they apply in this context; only part of the *act* (OODA) and *response* (SHORe) activities occurs within the context of the JPRC.

Phase	OODA	SHORe	“JPRC Experiment C” Analysis
1	<i>observe</i>	<i>stimulus</i>	<i>information-gathering</i>
2	<i>orient</i>	<i>hypothesis</i>	<i>sense-making</i>
3	<i>decide</i>	<i>option</i>	<i>decision-making</i>
4	<i>act</i>	<i>response</i>	

6.3 Results

6.3.1 Task Analysis

A fragment of the task analysis performed on the activities observed during experiment C can be seen in Figure 13. A similar analysis was performed for experiment E highlighting the progress made with the Co-OPR application.

37	0:51:15		decision-making	request launch authority
38	0:53:00	JD	decision-making	mission establish
39	0:55:00	JD	decision-making	send plan to whitecell
40	0:55:45			[wait for confirmation from whitecell]
41	0:57:02	JD	housekeeping	updates state assuming mission complete
42	1:00:00	GW	housekeeping	rescue figgy, note reporting agency, establ
43	1:01:00	GW	information-gathering	request info - incident location
44	1:02:13	GW	information-gathering	receive location
45	1:02:36	GW	information-gathering	validate incident (WoD)
46	1:03:30	GW	information-gathering	awaiting message
47	1:04:00	GW	housekeeping	broadcast WoD update to all subordinates
48	1:04:46	GW	information-gathering	ISOPREP - VOC
49	1:06:58	GW	sense-making	identify threats - I-X mesgs
50	1:09:47	GW	sense-making	ensure conditions, ensure OSC
51	1:11:10			[interrupt] duress word
52	1:13:00	GW	sense-making	reviews plan (threats) so far

Figure 13: Fragment of Co-OPR task analysis

Notwithstanding the provisos noted above about the inability at the time of writing to perform a full comparative evaluation, the analysis is encouraging for the use of the I-X in this task. In general, the use of SOPs encouraged a methodical approach to the overall JPRC activity: instances of information-gathering were followed by instances of sense-making which led to decision-making episodes, with no instances of, for example, a decision-making activity being interrupted or abandoned due to the lack of a crucial piece of information. In addition, at several times during the exercise, important messages arrived which interrupted the current activity and diverted the cognitive attention of the director or controller. Such interruptions can serve to disrupt the flow of the Center, but in the majority of cases, the framework provided by the SOPs allowed a quick resumption of activity once the message had been dealt with.

The analysis also highlighted some areas where further support might prove helpful. In addition to dealing with interruptions, the arrival of new information which demands that the decisions made earlier in the process are to be re-appraised (and, in one case during the experiment, wholly abandoned, with rescue resources 'recalled') is currently difficult to handle using the SOP framework (and would seem to require something akin to 'exception-handling' procedures). Successfully dealing with such situations seems to rely too much on the experience and initiative of the human in question. This would seem to be a general problem with any SOP-based system rather than with I-X per se, but technology that can offer more support would obviously be of great benefit.

The time devoted during the experiment to each of the task categories is also interesting. While roughly the same amount of time was spent in information-gathering, sense-making and decision-making during the exercise, a surprisingly large amount of time was

spent housekeeping – twice as long as the time spent for any of the other categories. This is due, in part, to the time required to initialize the JPRC and check that its procedures and communications are in place, and then later to produce a report summarizing the session activities for the next duty officer. Providing automated assistance for these tasks may reduce the workload of the humans involved while also ensuring a more rapid and efficient establishment of the Center and hand-over of duty.

Aside from an analysis of the cognitive tasks performed by the system users, the experimentation also highlighted a number of open issues with the current prototype. Firstly, support for the white cell was rather limited at this stage. Only the structured messaging feature was a real advantage provided by I-X. However, the way the scenario was driven was adapted to this way of delegating tasks, which does not correspond well to the way the real CPX works. This in effect removes a large part of the sense-making task from the problem and shifts the focus onto the planning activities, an area in which I-X is strong. Secondly, the two panels used by the director and the controller are equipped with independent <I-N-C-A> models which may lead to inconsistent world state representations within the JPRC. While this did not occur during the experiment, it is a potential problem that was noted. Finally, a few problems with the user interface need to be addressed for future versions, e.g. the lack of a mechanism to draw the user's attention immediately to new, incoming activities.

6.3.2 Benefits: Cognitive Attributes

The I-X technology in the Co-OPR program was observed and evaluated in each of the Co-OPR experiments by a Human Factors expert, Dr. Jeff Hansberger from the Army Research Laboratory representing the field element at the US Joint Forces Command (USJFCOM). Based on an assessment of the Co-OPR capabilities, observations of these Co-OPR capabilities in the context of a Personnel Recovery (PR) situation, and a task analysis of the PR tasks, the I-X technology is expected to provide multiple benefits and enhancements to the Personnel Recovery Center Director and staff. These potential benefits will be assessed within a distributed cognitive system framework that will focus on the interaction between the PR user and Co-OPR.

Distributed cognition (Hutchins, 1995) is a theoretical framework that explains cognitive activities as embodied and situated within the work setting and the artifacts used in the environment. Distributed cognition (D-Cog) emphasizes the distributed nature of cognitive phenomena across individuals, tools/technologies, and internal/external representations. The unit of analysis goes beyond the cognitions of a single individual and focuses on the functional system as a whole to examine the relation between individuals, the task environment, and artifacts used for task completion. Such functional systems have 6 basic distributed cognitive attributes: 1) Coordination across agents 2) situation assessment, 3) mental models, 4) memory demands, 5) attentional control, and 6) workload management.

Among the 6 D-Cog attributes, the I-X technology is expected to improve upon 5 of them.

- 1. Coordination across agents**
- 2. Situation assessment**
3. Memory demands
4. Attentional control
5. Workload management

Among those 5 attributes, the attributes of coordination across agents and situation assessment are the most likely areas to be enhanced by Co-OPR and can be directly associated with components of a task analysis of the Personnel Recovery (PR) actions done by USJFCOM (Bolstad, Cuevas, & Costello, 2005). As part of the assessment of the I-X technology to PR, components of the task analysis were highlighted to represent areas most likely enhanced by the I-X technology (see Figure 14). The potential benefit to each D-Cog attribute will be addressed in turn along with the specific Co-OPR capability that supports that task.




Figure 14: Hierarchical Task Analysis of a Personnel Recovery Center Director. Lightly highlighted tasks represent areas which Co-OPR is hypothesized to positively influence.

Coordination across Agents

The coordination of information and actions across the Joint Personnel Recovery Center (JPRC) and the service Rescue Coordination Centers (RCCs) is a critical part of any PR mission. To illustrate its importance, coordination activities comprise 3 of the 5 major PR tasks in the task analysis (Figure 14). Co-OPR was evaluated to have a positive influence on all these 3 major PR coordination tasks: 1) Coordinate assets for assignment, 2) Coordinate the recovery plan, and 3) Provide effective communications.

Coordinate assets for assignment: In order to coordinate assets for a PR mission, the staff needs to have knowledge of the currently available assets (Figure 14). Not only does the staff need to know what assets are available, it needs to be updated quickly and efficiently on any change in availability among its current assets. Both types of information are presented to the Co-OPR user through the I-Plan process panel when assets are involved. Only relevant assets are presented according to the constraints of the mission (e.g., ground assets are not presented for a water recovery). In addition to intelligently constraining the available assets, relevant assets that are currently not available due to its use in other actions or other reasons are also made known to the user.

Access to this information in the context of specific actions within the PR mission provide the clear and current information needed to coordinate a PR plan while minimizing re-planning due to outdated or unknown asset constraints.

Coordinate the Recovery Plan: During execution of a planned PR mission, it is important to accurately assess the plan progress. The monitoring of the mission progress sub-task (Figure 14) allows for the appropriate coordination with all the involved parties and indicates any needed deviations and adaptations to the original plan. Co-OPR should enhance the ability to monitor the status of the mission through its I-Plan process panel that provides directed feedback on the current status of actions and requests. This information is both presented in a format that facilitates a “quick look” capability with a scan across many actions to see what tasks are still to be accomplished as well as a format that enables greater detail to be drilled down to (e.g., who the task is assigned to or was completed by).

The I-Plan panel also enables a potentially very effective means to update and conduct a handover between Commander/staff shifts or changes. Through its visualization of what tasks have been completed for a mission and which have not, it quickly updates a Commander on the progress of the mission and potentially what the next step(s) should be. It also presents a dynamic historical account of progress within the current and other missions if the Commander wishes to explore more information and context.

Provide Effective Communications: Co-OPR should contribute to both communication sub-tasks of communication of mission critical information to higher command and to other recovery centers (Figure 14). The 3 components that should have the largest impact on getting mission critical information to the JPRC and the supporting RCCs is the I-Plan panel, electronic white

boards, and the structured instant messaging. As mentioned for the coordination of available assets, the I-Plan aids in communicating the available and relevant assets for the mission. This information is communicated to higher command within the context of the mission and the tasks required to complete the mission.

One of the strengths of Co-OPR is its ability to structure and share communications across the JPRC and RCCs through its I-Plan panel, electronic white boards, and the structured instant messaging. The facilitation of this shared information dissemination provides greater common ground (Clark & Brennan, 1991) among the users and should enable greater shared awareness and understanding of the environment. Each of the 3 Co-OPR components mentioned above provides a way to share information or a virtual space in which information can be shared. The I-Plan panel shares information on specific tasks and task assignments. The virtual white board creates a shared virtual space for mission relevant information and could reduce errors by providing a single information source for mission information compared to each RCC independently maintaining the same information. Finally, the structured instant messaging system works in conjunction with the I-Plan and provides additional context to the information being sent (e.g., by linking it with an action defined in the I-Plan panel).

Situation Assessment

The accurate and timely assessment of the situation is critical for any PR mission. Co-OPR has several component capabilities that could provide an enhanced ability to analyze the PR event and develop situation awareness tasks identified in the PR task analysis (Figure 14).

Analyze the Isolated Individual Event: The first action in the analysis of a new PR event is to validate the information from the event to deconflict, validate, and obtain background information on the event. The semi-structured nature of tasks and communications in Co-OPR, based on the <I-N-C-A> ontology, proceduralizes many actions and naturally begins to identify and work out any conflicts within the available information. Co-OPR also supports courses of action (COA) development through its COA tool that could be used to help recommend appropriate support or recovery plans.

Develop Situation Awareness: Co-OPR capabilities that are hypothesized to improve situation awareness have already been identified in other sections (e.g., “Provide Effective Communications”). It is expected that the primary means by which Co-OPR facilitates the development of situation awareness and the assessment of the friendly situation in particular is through the structure that <I-N-C-A> provides for communications and actions. This structure facilitates quick inspection of actions and communications, which can improve shared situation awareness among distributed RCCs.

Memory Demands

The I-Plan panel’s intelligent to-do list capability aids in the memory demands on the user by acting as an external memory device and therefore distributing the memory demands of the task across the user and Co-OPR. This joint memory system is referred to as a transactive memory

system (Wegner, 1987) where another person (e.g., spouse remembers billing schedule or phone numbers) or object (e.g., telephone speed-dial, PDA address book, etc...) can be used to encode knowledge external to their internal memory system. Past research has shown that transactive memory systems enable better utilization of knowledge and allow higher levels of performance to be reached (Moreland & Argote, 2003).

The Co-OPR intelligent to-do list possesses the steps and procedures needed to accomplish a PR mission and therefore doesn't require the user to recall this information from their own memory. This proceduralization of tasks can significantly reduce errors of omission (Reason, 1990). The list also acts as a constant reminder of tasks that are awaiting completion, ones that have been completed, and ones that are in progress.

Attentional Control

The I-Plan panel's intelligent to-do list capability also directs the attention of the user to needed and unaccomplished tasks. The to-do list capability does not, however, force the user into actions without options and the flexibility to customize actions to the specific situation.

Workload Management

The framework of the NASA TLX will be used to discuss the implications of the Co-OPR system on workload. The NASA TLX is a subjective workload assessment tool that allows users to perform subjective workload assessments on operator(s) working with various human-machine systems. NASA TLX is a multi-dimensional rating procedure that derives an overall workload score based on a weighted average of ratings on six subscales.

1. 'Mental demand' refers to how much mental and perceptual activity was required (thinking, deciding, calculating, remembering, looking, searching, etc.) during the task. The respondent should consider whether the task was easy or demanding, simple or complex, or exacting or forgiving.
2. 'Physical demand' measures the experienced required physical activity in relation to whether the task was easy or demanding, slow or brisk, slack or strenuous, or restful or laborious.
3. The amount of time pressure experienced is measured by the 'temporal demand' subscale. It addresses issues such as whether the pace of interaction was slow and leisurely or rapid and frantic.
4. 'Performance' refers to how successful respondents think they were in accomplishing the goals of the task set by the experimenter, and how satisfied they were with their performance in accomplishing these goals.
5. The criteria of 'effort' requests the respondents to assess how hard they had to work (mentally and physically) to accomplish the level of performance they achieved.
6. Finally, evaluation of the 'frustration' level measures how insecure, discouraged, irritated, stressed and annoyed versus secure, confident, relaxed, and complacent subjects felt during the task.

Due to the memory demand reduction and the use of Co-OPR as a transactive memory device discussed above, lower mental demand should be evident for a user of Co-OPR. The temporal and effort workload dimensions should be improved due to the intelligent agent support that Co-OPR provides. Finally, an improvement in the performance dimension is expected due to all the benefits of the Co-OPR system mentioned in this evaluation.

Summary

The capabilities of the Co-OPR system have the potential to support the Commander and staff of a JPRC & supporting RCCs across many cognitive dimensions. The cognitive attributes it shows direct support for when linked to the PR hierarchical task analysis is the coordination across agents and situation assessment. Other capabilities of Co-OPR show promise in supporting and improving the memory demands, attentional control, and workload management of the user when interacting with Co-OPR to complete a PR mission. Future Co-OPR work would test these observations against human behavioral modeling results (e.g., Hansberger & Barnette, 2005) and additional experimental testing with emphasis on distributed cognitive data collection and analysis.

7. Conclusions and Future Work

In this report we have described the I-X framework and I-X Process Panels, which can be seen as providing a distributed and intelligent to-do list for agent coordination in emergency response and, more specifically, in personnel recovery as performed by a JPRC. The to-do list analogy provides users with a familiar metaphor that should make an I-X application easy to understand. However, I-X extends this concept in two important ways.

Firstly, items on the to-do list can be expanded using pre-defined standard operating procedures. Such procedures are available in many scenarios but usually only in the form of books or manuals that, even if they are to hand, are often too cumbersome to use in a real emergency. The encoding of such standard operating procedures in I-X is supported by a graphical domain editor. The intention is, of course, that this takes place before an emergency occurs. As a result, this knowledge is at hand and can be used when it is most needed. The HTN planner that is available in I-X uses the library of standard operating procedures to update the Process Panel, showing the user the various ways in which an item on the to-do list can be dealt with. Thus, the apparent intelligence of the panel is the knowledge encoded by a domain expert before an emergency occurs, but it is adapted to the context and can also be composed dynamically if the context required.

The second extension provided by I-X is the capability model. This allows for a number of panels to be linked to respond in related ways to an emergency. For the user this means that the panel can suggest other agents that may be able to deal with an item if they choose to advertise a matching capability. Furthermore, the panel provides support for the management of such task distribution by sending activities with their parameters and keeping track of reports relating to that activity as they come back.

Both these extensions are integrated into the panel in a seamless way. Together these technologies are used to effectively support emergency responders in organizing a collaborative response quickly and efficiently.

Of the I-X applications currently under development at AIAI, the Co-OPR application was chosen as a test case and a series of experiments were performed in which the Co-OPR application was used to support the task of personnel recovery training. This shows that I-X can indeed be used to build applications that support task-centric activities in this domain, and that the two features focused on in this report, namely intelligence through integrated standard operation procedures, and coordination support through linked panels, are useful in supporting the overall activity of a JPRC. More specifically, an analysis of the experiments shows that the hierarchical structure of the tasks in the to-do list helps users to focus their efforts and avoid distractions, and if interrupted, it helps them to quickly continue with important decision making without having to repeat information-gathering or sense-making activities that have already been completed.

Thus, the I-X framework is useful for developing task-supporting applications. Specifically, it proved useful in its initial form showing that it is indeed a generic framework. With the additional viewers now in place the framework is even more powerful and should be able to support an even wider range of task-centric applications.

8. References

- Allsopp, D., Beautement, P., Bradshaw, J.M., Carson, J., Kirton, M., Suri, N. and Tate, A. (2001) Software Agents as Facilitators of Coherent Coalition Operations, 6th International Command and Control Research and Technology Symposium, US Naval Academy, Annapolis, Maryland, USA, 19-21 June 2001.
- 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 Cooperation in an International Coalition Setting", Proceedings of the Second International Conference on Knowledge Systems for Coalition Operations (KSCO-2002), Toulouse, France.
- Allsopp, D., Beautement, P., Kirton, M., Tate, A., Bradshaw, J.M., Suri, N. and Burstein, M. (2003) The Coalition Agents Experiment: Network-Enabled Coalition Operations, Special Issue on Network-enabled Capabilities, Journal of Defence Science, Vol. 8, No. 3, pp. 130-141, September 2003.
- BBN (2003) OpenMap™ Open Systems mapping Technology - <http://openmap.bbn.com>
- Bolstad, C.A., Cuevas, H.M., & Costello, A. (2005, March). Final Report for the Personnel Recovery Education and Training Center (Cooperative Agreement DAAD19-01-2-0009). Marietta, GA: SA Technologies.
- Bradshaw, J. M., Beautement, M. Breedy, L. Bunch, S. Drakunov, P. Feltovich, P., Raj, A., Johnson, M., Kulkarni, S., Suri, N. & A. Uszok (2004). Making agents acceptable to people. In N. Zhong & J. Liu (Ed.), Intelligent Technologies for Information Analysis: Advances in Agents, Data Mining, and Statistical Learning. (pp. 361-400). Berlin: Springer Verlag.
- Bradshaw, J. M., Uszok, A., Jeffers, R., Suri, N., Hayes, P., Burstein, M. H., Acquisti, A., Benyo, B., Breedy, M. R., Carvalho, M., Diller, D., Johnson, M., Kulkarni, S., Lott, J., Sierhuis, M., & Van Hoof, R. (2003). Representation and reasoning for DAML-based policy and domain services in KAoS and Nomads. Proceedings of the Autonomous Agents and Multi-Agent Systems Conference (AAMAS 2003). Melbourne, Australia, New York, NY: ACM Press,
- 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.
- Buckingham Shum, S., Selvin, A., Sierhuis, M., Conklin, J., Haley, C. and Nuseibeh, B. (2006). Hypermedia Support for Argumentation-Based Rationale: 15 Years on from gIBIS and QOC. In: Rationale Management in Software Engineering (Eds.) A.H. Dutoit, R. McCall, I. Mistrik, and B. Paech. Springer-Verlag: Berlin

- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine & S. D. Teasley (Eds.), *Perspectives on socially shared cognition*. Washington: D.C.: American Psychological Association.
- Conklin, J. and Begeman, M. L. (1988). gIBIS: A hypertext tool for exploratory policy discussion. *ACM Transactions on Office Information Systems*, 4(6), 303-331.
- Conklin, J. (2003) Dialog Mapping: Reflections on an Industrial Strength Case Study. In: P.A. Kirschner, S.J. Buckingham Shum and C.S. Carr (Eds.), *Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making*. Springer-Verlag: London, ISBN 1-85233-6641-1, pp. 117-136.
- 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.
- Currie K. and Tate A.. (1991) O-Plan: the Open Planning Architecture. *Artificial Intelligence*, Vol. 52, pp 49-86.
- Damianou, N., Dulay, N., Lupu, E. C., & Sloman, M. S. (2000). Ponder: A Language for Specifying Security and Management Policies for Distributed Systems, Version 2.3. Imperial College of Science, Technology and Medicine, Department of Computing, 20 October 2000.
- Ghallab, M., Nau, D., and Traverso, P. (2004) *Automated Planning – Theory and Practice*, chapter 11. Elsevier/Morgan Kaufmann.
- Hansberger, J.T. and Barnette, B.D. (2005). Human performance modeling for operational command, control, and communication. *Proceedings of the 49th Annual meeting of the Human Factors & Ergonomics Society*, Orlando, FL.
- Hutchins E. (1995). How A Cockpit Remembers its Speeds. *Cognitive Science*, 19, 265-288.
- Jabber (2003) Jabber XML Instant Messaging, <http://www.jabber.org>.
- Jabber (2006) Jabber: Open Instant Messaging and a Whole Lot More, Powered by XMPP. <http://www.jabber.org/>
- Joint Publication 3-50 (2005) Joint Doctrine for Personnel Recovery, 2nd Draft, March 2005.
- Khambhampati, S. and Srivastava, B. (1996) Unifying Classical Planning Approaches, Arizona State University ASU CSE TR 96-006, July 1996.
- Kreifelts Th., Hinrichs E., and Woetzel G. (1993) Sharing To-Do Lists with a Distributed Task Manager. In: de Michelis G. and Simone C. (eds.) *Proceedings of the 3rd European Conference on Computer Supported Cooperative Work*, pp 31-46, Milano, 13-17 September 1993, Kluwer, Dordrecht.

- MacLean, A., Young, R., Bellotti, V. and Moran, T. (1991) Design space analysis: Bridging from theory to practice via design rationale. In Proceedings of Esprit '91, pages 720-730, Brussels, November 1991.
- McIlraith, S. A., Son, T. C., & Zeng, H. (2001). Semantic Web Services. IEEE Intelligent Systems, 46-53.
- Moreland, R. L., & Argote, L. (2003). Transactive memory in dynamic organizations. In R. Peterson & E. Mannix (Eds.), *Leading and managing people in the dynamic organization* (pp. 135-162). Mahwah, NJ: Erlbaum.
- Osinga, F. (2006) *Science Strategy and War, The Strategic Theory of John Boyd*, Abingdon, UK: Routledge, ISBN 0-415-37103-1.
- Polyak, S. and Tate, A. (1998) Rationale in Planning: Causality, Dependencies and Decisions. *Knowledge Engineering Review*, Vol.13(3), pp 247-262.
- Polyak, S. and Tate, A. (1999) A Common Process Ontology for Process-Centred Organisations, *Knowledge based Systems*, 2000. Earlier version by S. Polyak published as University of Edinburgh Department of Artificial Intelligence Research paper 930, 1998. See <http://www.aiai.ed.ac.uk/project/oplan/documents/1999/99-sebpc-cpm.pdf>
- Rathmell, R.A. (1999) "A Coalition Force Scenario 'Binni — Gateway to the Golden Bowl of Africa'", Proceedings of the International Workshop on Knowledge-Based Planning for Coalition Forces, (ed. Tate, A.) pp. 115-125, Edinburgh, Scotland, 10th-11th May 1999.
- Reason, J. (1990). *Human Error*. Cambridge University Press.
- Sacerdoti E. (1975) The Nonlinear Nature of Plans. In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), pp 206-214.
- Schlenoff, C., Gruninger M., Tissot, F., Valois, J., Lubell, J., Lee, J. (2000). The Process Specification Language (PSL): Overview and Version 1.0 Specification," NISTIR 6459, National Institute of Standards and Technology, Gaithersburg, MD.
- Selvin, A.M. (1999) Supporting Collaborative Analysis and Design with Hypertext Functionality, *Journal of Digital information*, Volume 1 Issue 4.
- Selvin, A.M, Buckingham Shum, S.J., Sierhuis, M., Conklin, J., Zimmermann, B., Palus, C., Drath, W., Horth, D., Domingue, J., Motta, E. and Li, G. (2001) *Compendium: Making Meetings into Knowledge Events*, Knowledge Technologies 2001, Austin TX, USA, March 4-7, 2001.
- Smith, S.F. (1994) OPIS: A methodology and architecture for reactive scheduling. In Zweben, M. and Fox, M.S. (eds), *Intelligent Scheduling*, chapter 8, pages 29--66. Morgan Kaufmann, San Francisco, CA, USA.

Tate A. (1977) Generating Project Networks. . In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), pp 888-893.

Tate, A. (1995) Integrating Constraint Management into an AI Planner. Journal of Artificial Intelligence in Engineering, Vol. 9, No.3, pp 221-228.

Tate, A. (1996) The <I-N-OVA> Constraint Model of Plans, Proceedings of the Third International Conference on Artificial Intelligence Planning Systems, (ed. Drabble, B.), pp. 221-228, Edinburgh, UK, May 1996, AAAI Press.

Tate, A. (1998) Roots of SPAR, in Special Issue on Ontologies, Knowledge Engineering Review, Vol. 13(1), March, 1998, Cambridge University Press.

Tate, A. (2000a) “<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, at the National Conference of the American Association of Artificial Intelligence (AAAI-2000), Austin, Texas, USA, August 2000.

Tate, A. (2000b) “Intelligible AI Planning”, in Proceedings of the Twentieth British Computer Society Special Group on Expert Systems International Conference on Knowledge Based Systems and Applied Artificial Intelligence, Cambridge, UK, December 2000.

Tate, A., Dalton, J. and Levine, J. (2000a) “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.

Tate, A., Levine, J., Dalton, J. and Nixon, A. (2001) “Task Achieving Agents on the World Wide Web”, in “Creating the Semantic Web”, Fensel, D., Hendler, J., Liebermann, H. and Wahlster, W. (eds.), MIT Press, 2001.

Tate, A., Dalton, J. and Stader, J. (2002) I-P2- Intelligent Process Panels to Support Coalition Operations. In Proceedings of the Second International Conference on Knowledge Systems for Coalition Operations (KSCO-2002). Toulouse, France, April 2002.

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, pp 125-130.

Tate, A., Buckingham Shum, S.J., Dalton, J, Mancini, C. and Selvin, A.M. (2006) Co-OPR: Design and Evaluation of Collaborative Sensemaking and Planning Tools for Personnel Recovery, Open University Knowledge Media Institute, Technical Report KMI-06-07, March 2006.

Uszok, A., Bradshaw, J. M., Jeffers, R., Suri, N., Hayes, P., Breedy, M. R., Bunch, L., Johnson, M., Kulkarni, S., & Lott, J. (2003). KAoS policy and domain services: Toward a description-

logic approach to policy representation, deconfliction, and enforcement. Proceedings of Policy 2003. Como, Italy.

Wark, S., Zschorn, A., Perugini, D., Tate, A., Beutement, 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.

Wegner, D. M. (1986). Transactive memory: A contemporary analysis of the group mind. In B. Mullen & G. R. Goethals (Eds.), *Theories of group behavior* (pp. 185-208). New York: Springer-Verlag.

Wickler, G. (1999) Using Expressive and Flexible Action Representations to Reason about Capabilities for Intelligent Agent Cooperation. PhD thesis, University of Edinburgh.

Wohl, J.G. (1981) Force Management Decision Requirements for Air Force Tactical Command and Control. *IEEE Transactions on Systems, Man, and Cybernetics*, 1981, 11.

Appendix A: Screen Shots from the Final Experiment

A.1 Infrastructure

The Name-Server is a registry that knows about the addresses of all the agents in the scenario. The screen shot in Figure 15 shows that all the agents have been registered on their respective machines, some with fixed ports assigned to them, others at dynamic ports. All agents in the scenario can now send messages to each other and the Name-Server window plays no further role in the experiment.




Figure 15: The I-X Name Sever with all relevant agents registered

A.2 White Cell

The white cell is the first agent to be started after the Name-Server.




Figure 16: The I-X for the White Cell (trainers)

Figure 16 shows the white cell's I-X Process Panel in its initial state. In this experiment, the panel is used to get access to the I-Sim tool which drives the scenario. The panel could also be used to visualize the current state for the trainers, but no use of it in that way is made in this experiment.



Figure 17: The basic parameters to start the scenario simulation

To start the simulation the white cell must specify the initial simulated time point, which is June 7, 2000 at 15:00 as shown in Figure 17. This is one of the actual training scenarios used by the PRETC.




Figure 18: The I-Sim Tool for controlling the events that will occur

The I-Sim control window shown in Figure 18 shows all the events that are found in the MSELs for the scenario. Initially, all time points are relative to the start of the simulation. Events are listed by the thread they belong to, which can be freely recombined by the trainers to dynamically modify the scenario.



Figure 19: The I-Sim Clock showing Simulated Time (at the White Cell)

The I-Sim clock shown in Figure 19 is a small window that shows the current simulated time. This tool is available to all agents.




Figure 20: The I-Sim Tool during the simulation (at 15:27 simulated time)

Figure 20 shows the I-Sim control tool after a number of incident reports have been created. These are marked as completed in the tool. Time points that were initially relative have now been replaced with absolute (simulated) times.

A.3 JPRC

There are three agents in the JPRC that are supported by I-X Process Panels: the JPRC director, the watch supervisor, and the controller.

The JPRC Director




Figure 21: Initial state of the JPRC director's I-X Process Panel

Figure 21 shows the initial state of the JPRC director's panel. The only information showing in the panel at this stage is the current world state.




Figure 22: JPRC director's panel with the first task (from the JTFC)

Since annotations are not used in this experiment and state information is best viewed through the various state viewers provided, the director has minimized the state and activities parts of the panel in Figure 22. This figure also shows the arrival of the first task, to set up the JPRC, which is coming from the JTFC. The action menu related to this task shows how it can be addressed and the director is about to select a refinement, which is a common way of dealing with tasks.




Figure 23: JPRC director's panel with the first task expanded

After the refinement has been selected the panel shows the various sub-activities into which the overall task has been broken down. In Figure 23 the sub-tasks are listed and the color-coding suggests that 4 of the 8 sub-tasks shown can be done immediately.




Figure 24: JPRC director's panel – JPRC setup in progress

Figure 24 shows the JPRC director's panel still during the setup of the JPRC with a number of tasks now completed. The displaying of the I-X information displays, the shared displays described later, has been delegated to the watch supervisor and is in progress.




Figure 25: JPRC director's panel – JPRC setup completed

Figure 25 shows the director's panel after the first phase, the setup of the JPRC, has been completed.




Figure 26: JPRC director’s panel – second phase: distribution of the ATO

The second phase of the experiment begins when the JTFC sends out the ATO. This is initially received by the director and Figure 26 shows the director looking through the action menu for ways of dealing with this task.




Figure 27: JPRC director’s panel – third phase: first incident in progress

The third phase of the experiment constitutes the heart of the work performed by the JPRC. Figure 27 shows the director's panel with the first incident report received and the first steps to deal with the incident already completed, one delegated to the watch supervisor where it is in progress, three more ready to be dealt with, and one final task that cannot be done yet.




Figure 28: JPRC director's panel – third phase: second incident report just arrived

While the first incident is still being dealt with, a second report related to a different incident reaches the director. Note that this new activity is shown in pink in Figure 28 indicating that it has not been looked at. New activities from other agents are always shown in this color to draw the user's attention to them.




Figure 29: JPRC director's panel – information gathering for two parallel incidents in progress

Figure 29 shows the JPRC director's panel with both incidents still in progress. If the shift handover were to take place now, the panel could be read line by line as a status report, e.g. as follows:

- The setup of the JPRC and the reviewing of the ATO have been completed.
- The first incident is not completed but in progress.
 - An incident folder has been created and information from this report has been used to update it.
 - The collecting of necessary incident data has been delegated to the watch supervisor and is in progress.
 - The ensuring of conditions for the incident has been started but none of its sub-tasks have been started.
 - Mission planning is not yet possible.
- The second incident is not completed but in progress.
 - An incident folder has been created ... etc.

The Watch Supervisor

The watch supervisor's main responsibility is to collect the information relating to the various incidents that need to be dealt with. This information is displayed on a shared display which can be seen by everybody in the JPRC. The screen shots from the shared display will be explained below. Here we will look at the screen shots taken from the watch supervisor's own computer that is only meant to support the watch supervisor. Each screen shot is labeled with a number here and the screen shots from the shared display that were taken at the same time show the same number in their caption.




Figure 30: The Watch Supervisor's Screen (1)

Figure 30 shows the screen of the watch supervisor after the completion of the first phase, the setup of the JPRC. All activities delegated to the watch supervisor have been completed, which means that all the shared information displays are now set up and visible. In addition, the watch supervisor has decided to display the Virtual Operations Center on their screen, which is essentially a set of locally available HTML pages displaying static information. Availability of this information had to be verified as part of the setup of the JPRC.




Figure 31: The Watch Supervisor's Panel (2)

Incidents reports usually come to the watch supervisor's panel and Figure 31 shows some reports being dealt with while others are still waiting to be looked at.




Figure 32: The Watch Supervisor's Panel (3)

Later, the watch supervisor has caught up with some of the reports that are coming in. Figure 32 shows only one activity in pink, which is a new incident report that has been received.




Figure 33: The Watch Supervisor’s Screen (4)

Part of the process of collecting incident data involves retrieving the Isoperp, which is a data sheet held by the unit of the isolated personnel. In Figure 33 the watch supervisor has used the I-Serve agent to obtain this information. It is currently displayed as a Word document on the screen. The important information extracted from this document, e.g. the number of personnel that are isolated, will be put on the shared display, whereas the Isoperp goes into file related to the incident.




Figure 34: The Watch Supervisor's Panel (5)

Figure 34 shows the watch supervisor binding variables – a tool is used that suggests possible values that are extracted from the current context of the panel.




Figure 35: The Watch Supervisor's Panel (6)

Towards the end of the experiment the watch supervisor's panel shows almost all incident reports dealt with completely. Only the first for each incident is still in progress. This is because the first incoming message is used to inform the director that there is a new incident. The director has to deal with all the tasks that are to do with this new incident, which includes the rescue. Subsequent reports relating to the same incident are dealt with locally by the watch supervisor, which is why all of them are completed in the panel in Figure 35. When the rescue has been completed the respective tasks on both, the director's and the watch supervisor's panel will show the activity as completed.




Figure 36: The Watch Supervisor's Screen (7)

Finally, Figure 36 shows the watch supervisor's screen at the end of the experiment. One of the rescues has now been completed; the other one is still in progress. Also, the phone book which is part of the VOC is visible. Presumably, the watch supervisor has used the phone to report the success.

The Controller

The task of the controller is mostly to manage the rescue of isolated personnel once the information about an incident is sufficient. Thus, this role becomes active relatively late in the experiment, apart from the analysis of the ATO to extract the available CSAR resources.




Figure 37: Initial state of the JPRC controller’s I-X Process Panel

Figure 37 shows the initial state of the JPRC controller’s panel.




Figure 38: JPRC controller’s Panel with first task – to bring up the CSAR asset board

During the setup of the JPRC, the only task for the controller is to display the asset board. This task has just arrived on the controller’s panel in Figure 38. Note that, initially there will be no information on this board as the CSAR resources only become known when the ATO is distributed in phase 2 of the experiment.

In the experiment described here, the ATO is a document that is stored in Word format on the I-Serve server and can be retrieved by the controller. Figure 40 shows a sample ATO as used by the PRETC for the June 7 scenario used in the experiment.




Figure 41: JPRC controller’s Panel – review of the ATO completed

Figure 41 shows the controller’s panel with the review of the ATO completed. The different sub-tasks are still expanded.

The screenshot shows a window titled "JPRC Controller Panel Assets View". Inside the window, there is a menu bar with "File" and "View". Below the menu bar is a section labeled "Assets" containing a table. The table has eight columns: "Object", "type", "callsign", "ac-type", "latitude", "longitude", "alert-time", and "note". The rows list various assets such as robin71, verbiage31, woodpecker61, and wren54. At the bottom of the window, there are two buttons: "Commit changes" and "New Row".

Object	type	callsign	ac-type	latitude	longitude	alert-time	note
robin71		ROBIN71					
robin72		ROBIN72					
verbiage31		VERBIAGE31					
verbiage32		VERBIAGE32					
woodpecker61		WOODPECKER61					
woodpecker62		WOODPECKER62					
woodpecker63		WOODPECKER63					
woodpecker64		WOODPECKER64					
wren54		WREN54					
wren55		WREN55					
wren56		WREN56					
wren57		WREN57					

Figure 42: Assert board displaying the CSAR resources and their state

The corresponding asset board listing all the available CSAR assets is shown in Figure 42. This corresponds more or less directly to the asset board used during a CPX but is linked in with the panel's state that is shared with all JPRC panels.




Figure 43: Task to plan and execute mission for first incident arrived

After the review of the ATO the controller remains inactive in the experiment described here until incidents have progressed far enough to plan and execute a rescue mission. The first such task has just appeared on the controller’s panel shown in Figure 43.




Figure 44: Planning a rescue mission for the second incident

Planning a rescue mission is done by first creating the options available. This is done with the Option tools only available in the controller panel – the current option is always displayed at the top of the panel. By default, this is the “Base” case option (see Figure 43). The base option is

used for activities that the agent is committed to by convention. In Figure 44 the controller has created a new option for the incident, named after the call sign of the lost aircraft: gopher. This is a hypothetical list of activities that is only being considered. For a rescue task this is initialized with the rescue activity as the only entry; the state copied from the base option. The controller may then use the panel in the usual way to create a rescue plan, i.e. by working through the activities using the different ways of dealing with them offered in the action menu.




Figure 45: Using the I-Plan tool to automatically generate options

Instead of creating plans manually the controller may decide to use the planning tool which is part of I-X to flesh out a partial plan in which only the top-level strategy has been decided manually. Figure 45 shows the I-Plan tool on top of the controller's normal panel. Five options have been created so far as can be seen from the name of the option: "gopher-5".




Figure 46: Using the I-Plan Option tool to compare options

When a number of options have been generated another tool can be used to compare the different options: the I-Plan Option tool. This tool is shown in Figure 46. This tool lists all the available options in the Option Tree on the left. Currently shown are the base option and the gopher option under which 5 different, automatically generated, complete plans are shown. On the right a comparison matrix lists various features that characterize the different options. The “run” feature at the bottom can be used to run a simulation of each option individually, which may show even more information about the option. To inspect each option in detail the controller can select each option in the tree and the corresponding plan will be displayed in the panel.




Figure 47: Selecting an option in the current space

Options can be selected from the panel using the Options menu as shown in Figure 47. Once an option has been adopted it needs to be executed which should result in the actual rescue if everything goes according to plan.

The procedure for the second incident is similar to the above and therefore no further screen shots are given here.

Shared Information Displays and Resources

The shared information displays are usually projected onto a large screen within the JPRC so that they can be seen by all members of the JPRC. The only exception is the Virtual Operations Center shown in Figure 48 that is available to everybody individually, but still displaying the same information.




Figure 48: The Virtual Operations Center (as a local HTML resource)

The following screen shots all show the shared display that is controlled by the watch supervisor. The screen shots are numbered and the corresponding screen shots from the controller's personal screen are labeled accordingly and were explained above.




Figure 49: Shared display (1)

Figure 49 shows the basic windows that are shared at an early state of the experiment. The largest window displays an overview map of the area with icons indicating the positions of different forces and relevant objects.




Figure 50: Shared display (2)

Figure 50 shows the shared display after the incidents have occurred. The incident board in the bottom left now lists both incidents and the information relating to them.




Figure 51: Shared display (3)

Simulated time is now 15:23 as shown by the I-Sim Clock in Figure 51. Multiple maps are still displayed – the main one giving an overview and two smaller ones focusing in onto the two incidents to give more detailed information.




Figure 52: Shared display (4)

7 minutes later in simulated time not much has changed. The small white board next to the map shows that the first letter of the word of the day has now been used for authentication (and may be compromised).




Figure 53: Shared display (5)

Two more simulated minutes later, Figure 53 shows that the first incident is close to the coast at a point where there is a lot of action and the JPRC needs to keep an eye on that.




Figure 54: Shared display (6)

The window at the bottom right in Figure 54 shows the current weather, which appears to be stable. New information about the second incident now shows the isolated personnel on land.




Figure 55: Shared display (7)

Finally, the rescue forces have arrived at the scene of the first incident as can be seen on the small map in the top right corner – the overview map shows the same information, but at a scale that does not show enough detail in the area of the incident.




Figure 56: The SPINs (document served by I-Serve)

Another shared information resource is constituted by the documents available from the I-Serve agent. For example, Figure 56 shows the current SPINs. These documents are shared but available to each agent in the scenario individually, similar to the Virtual Operation Center shown in Figure 48.

A.4 Other Agents in the Scenario

Normally the experiment described involves more agents than described so far. The JTFC is usually role-played by the white cell, but we have decided to support it with its own panel. Similarly, the Army, Navy and SOF RCCs would have a set of director, watch supervisor and controller panels available to them. To simplify, each of these RCCs is represented by a single panel here and these are hardly used in this experiment.




Figure 57: The I-X Process Panel for the JTFC

Figure 57 shows the initial state of the JTFC panel. Note that this panel has an activity at start-up, which has been delegated to the JTFC director here.




Figure 58: The I-X Process Panel for the JTFC after completion of the first 2 phases

In Figure 58 two tasks on the JTFC director's panel are shown as completed, which means the experiment must now be in the third phase – dealing with incidents.




Figure 59: The I-X Process Panel for the JTFC at the end of the experiment

The JTFC also has to authorize a number of activities undertaken by the JPRC. These authorization tasks can be passed to the JTFC which can then deal with them in the usual way. Figure 59 shows a number of such request, all of which have been dealt with.




Figure 60: The I-X Process Panels for the other RCCs

As mentioned before, the remaining RCCs do not play a significant role in this experiment. Their panels are shown in Figure 60. In a much larger experiment these could be replaced by more elaborate versions with multiple panels supporting multiple agents.

Appendix B: List of Publications

All available via <http://www.aiai.ed.ac.uk/project/ix/documents/> or <http://i-x.info/documents/>

2005

Dyer, D., Cross, S., Knoblock, C.A., Minton, S. and Tate, A. (2005) Planning with Templates, Guest Editor's Introduction, Special Issue on Planning with Templates, IEEE Intelligent Systems, March/April 2005, pp. 13-15, IEEE.

Nareyek, A., Freuder, E.C., Fourer, R., Giunchiglia, E., Goldman, R.P., Kautz, H., Rintanen, J. and Tate, A. (2005) Constraints and AI (Planning, Notes from the Workshop on Constraints and AI Planning at AAAI-2000, Austin, Texas, USA, August 2000) IEEE Intelligent Systems, March/April 2005, 20(2), pp. 62-72, IEEE.

Berry, D., Usmani, A., Torero, J., Tate, A., McLaughlin, S., Trew, A., Baxter, R., Bull, M. and Atkinson, M. (2005) FireGrid: Integrated emergency response and fire safety engineering for the future built environment, invited talk, Workshop on Ubiquitous Computing and e-Research, National eScience Centre, Edinburgh, UK 18-19 May 2005.

Chen-Burger, Y-H. and Lin, F-P. (2005) A Semantic-based Workflow Choreography for Integrated Sensing and Processing, Proceedings of the 9th IEEE International Workshop on Cellular Neural Networks and their Applications, Hsin-chu, Taiwan, May 28-30, 2005.

Lino, N., Tate, A., and Chen-Burger, Y-H. (2005) Semantic Support for Visualisation in Collaborative AI Planning, in Proceedings of the Workshop on The Role of Ontologies in Planning and Scheduling, at the International Conference on Automated Planning & Scheduling (ICAPS), June 2005, Monterey, California, USA.

Siebra, C. (2005) Planning Requirements for Hierarchical Coalitions in Disaster Relief Domains. Selected Papers from AI-2003/4 Poster Session, Expert Update Vol. 8, No. 1, pp. 20-24, Summer 2005, The Specialist Group on Artificial Intelligence, British Computer Society (BCS-SGAI). BCS SGAI Expert Update Vol. 8, No. 1, Summer 2005, ISSN 1465-4091.

Chen-Burger, Y-H., Hui, K., Preece, A.D., Gray, P.M.D. and Tate, A. (2005) Workflow Collaboration with Constraint Solving Capabilities. Selected Papers from AI-2003/4 Poster Session, Expert Update Vol. 8, No. 1, pp. 48-60, Summer 2005, The Specialist Group on Artificial Intelligence, British Computer Society (BCS-SGAI). BCS SGAI Expert Update Vol. 8, No. 1, Summer 2005, ISSN 1465-4091.

Siebra, C. and Tate, A. (2005) Integrating Collaboration and Activity-Oriented Planning for Coalition Operations Support. Proceedings of the 9th International Symposium on RoboCup 2005, 13-19 July 2005, Osaka, Japan.

Page, K.R., Michaelides, D.T., Buckingham Shum, S.J., Chen-Burger, Y-H, Dalton, J., De Roure, D.C., Eisenstadt, M., Potter, S., Shadbolt, N.R., Tate, A., Bachler, M. and Komzak, J. (2005) Collaboration in the Semantic Grid: A Basis for e-Learning. *Applied Artificial Intelligence*. 19, (9-10), 881-904 Taylor & Francis.

Berry, D., Usmani, A., Terero, J., Tate, A., McLaughlin, S., Potter, S., Trew, A., Baxter, R., Bull, M. and Atkinson, M. (2005) FireGrid: Integrated Emergency Response and Fire Safety Engineering for the Future Built Environment, UK e-Science Programme All Hands Meeting (AHM-2005), 19-22 September 2005, Nottingham, UK.

Siebra, C. (2005) A Unified Approach to Planning Support in Hierarchical Coalitions, Doctor of Philosophy Thesis, School of Informatics, University of Edinburgh, October, 2005.

2006

Siebra, C. and Tate, A. (2006) Integrating Collaboration and Activity-Oriented Planning for Coalition Operations Support. *Springer Lecture Notes on Artificial Intelligence*, Volume 4020, pp.561-568, (eds. Bredendfeld, A., et al.), RoboCup 2005 - Selected Papers from the RoboCup International Symposium, 13-19 July 2005, Osaka, Japan. [Reprint of Siebra and Tate (2005)]

Tate, A., Buckingham Shum, S.J., Dalton, J., Mancini, C. and Selvin, A.M. (2006) Co-OPR: Design and Evaluation of Collaborative Sensemaking and Planning Tools for Personnel Recovery, Open University Knowledge Media Institute, Technical Report KMI-06-07, March 2006.

Wickler, G., Tate, A. and Potter, S. (2006) Using the <I-N-C-A> Constraint Model as a Shared Representation of Intentions for Emergency Response, Proceedings of the First International Workshop on Agent Technology for Disaster Management (ATDM), at the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006), Future University, Hakodate, Japan, May 8-12, 2006.

Potter, S., Tate, A. and Wickler, G. (2006) Using I-X Process Panels as Intelligent To-Do Lists for Agent Coordination in Emergency Response, Proceedings of the Information Systems for Crisis Response and Management 2006 (ISCRAM2006), Special Session on "Multiagent Systems for Disaster Management and Response", Newark, New Jersey, USA, May 15-17, 2006.

Tate, A. (2006) The Helpful Environment: Geographically Dispersed Intelligent Agents That Collaborate, Special Issue on "The Future of AI", *IEEE Intelligent Systems*, May-June 2006, Vol. 27, No. 3, pp 57-61. IEEE Computer Society. [DOI Bookmark: - <http://doi.ieeecomputersociety.org/10.1109/MIS.2002.1005628>]

Wickler, G., Potter, S. and Tate, A. (2006) Recording Rationale in <I-N-C-A> for Plan Analysis, Workshop on Plan Analysis and Management, International Conference on Automated Planning and Scheduling (ICAPS-06), 6 June 2006, Lake District, England.

Siebra, C. and Tate, A. (2006) An Investigation into the Use of Collaborative Concepts for Planning in Disaster Response Coalitions, in Proceedings of the 2006 IEEE Workshop on Distributed Intelligent Systems (DIS-2006), (eds. Marik, V., Gruver, W.A., Pechoucek, M. and Preucil, L.) pp. 253-258, Prague, Czech Republic, June 2006, IEEE Computer Society.

Tate, A., Dalton, J., Bradshaw, J.M. and Uszok, A. (2006) Coalition Search and Rescue - Task Support: Intelligent Task Achieving Agents on the Semantic Web. Interim Technical Report (Final DAML Program Technical Report), January 2003-December 2004, AIAI Contract No. F-30602-03-2-0014 (DARPA Order No. P105/00), IHMC Contract No. F-30602-00-2-0577. Published as Air Force Research Laboratory Technical Report AFRL-IF-RS-TR-2006-91, via DTIC (<http://www.dtic.mil>), Accession Number ADA449044, July 2006.

Vaccari, L., Marchese, M., Giunchiglia, F., McNeill, F., Potter, S. and Tate, A. (2006) OpenKnowledge Deliverable 6.5: Emergency response in an open information systems environment, OpenKnowledge Project, July 2006, <http://openk.org>.

Potter, S., Tate, A. and Milne, R. (2006) I-Ex: Extreme Expedition Support. In "Robert Milne: A Tribute to a Pioneering AI Scientist, Entrepreneur and Mountaineer" (ed. Bundy, A.), IOS Press.

Tate, A. (2006) The Helpful Environment: distributed agents and services which cooperate, in "Cooperative Information Agents X" (eds. Klusch, M, Rovatsos, N. and Payner, T.R.), Lecture Notes in Artificial Intelligence (LNAI 4149), pp. 23-32, Springer. Invited talk to Tenth International Workshop on Cooperative Information Agents (CIA-2006), Edinburgh, 11-13 September 2006. <http://www.dfki.de/~klusch/cia2006/>

Alexander, D. (2006) Linking an HTN Planner to a Universal Robot Controller for High-Level Activity Control, Master of Science Dissertation, School of Informatics, University of Edinburgh, September, 2006.

Fragkakis, M. (2006) Modelling and Brokering Emergency Requests-For-Assistance, Master of Science Dissertation, School of Informatics, University of Edinburgh, September, 2006.

Wickler, G., Potter, S., Tate, A. (2007) Using I-X Process Panels as Intelligent To-Do Lists for Agent Coordination in Emergency Response, International Journal of Intelligent Control and Systems (IJICS), Special Issue on Emergency Management Systems, Vol. 11, No. 4, Dec. 2006. (<http://www.ijics.org/>)

Tate, A., Dalton, J. and Potter, S. (2006) Intelligible Messaging - Activity-oriented Instant Messaging.

2007

Wickler, G., Tate, A. and Hansberger, J. (2007) Supporting Collaborative Operations within a Coalition Personnel Recovery Center, Proceedings of the Fourth International Conference on Knowledge Systems for Coalition Operations (KSCO-2007) (Lawton, J., Patel, J. and Tate. A.

eds.) pp. 14-19, Waltham, MA, USA, 1-2 May 2007. In Proceedings of the International Conference on Integration of Knowledge Intensive Multi-Agent Systems Modeling, Evolution and Engineering (KIMAS '07).

Wickler, G., Tate, A. and Potter, S. (2007) Integrating Discrete Event and Process-Level Simulation for Training in the I-X Framework, Proceedings of the 4th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2007), B. Van de Walle, P. Burghardt and C. Nieuwenhuis (eds.), Delft, the Netherlands, May 2007.

Potter, S., Kalfoglou, Y., Alani, H., Bachler, M., Buckingham Shum, S., Carvalho, R., Chakravarthy, A., Chalmers, S., Chapman, S., Hu, B., Preece, A., Shadbolt, N., Tate, A. and Tuffield, M. (2007) The Application of Advanced Knowledge Technologies for Emergency Response, Proceedings of the 4th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2007), B. Van de Walle, P. Burghardt and C. Nieuwenhuis (eds.), Delft, the Netherlands, May 2007.

Tate, A. (2007) Planning and Doing Things, AISB Quarterly, Spring 2007, The Society for the Study of Artificial Intelligence and Simulation of Behaviour.

Lino, N.C.Q. (2007) Semantic Based Support for Visualisation in Complex Collaborative Planning Environments, Doctor of Philosophy Thesis, School of Informatics, University of Edinburgh, June, 2007.

Wickler, G., Potter S., Tate, A., Pěchouček, M. and Semsch, E. (2007) Planning and Choosing: Augmenting HTN-Based Agents with Mental Attitudes. International Conference on Intelligent Agent Technology (IAT 2007), Silicon Valey, 2-5 November 2007. IEEE Computer Society, Web Intelligence Consortium and Association for Computer Machinery. (to appear)

Appendix C: List of Software and On-line Documentation Available

I-X Latest Version Download

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

I-X Version 4.5 Download

<http://www.aiai.ed.ac.uk/project/ix/release/4.5/>

Appendix D: Project Web Sites

Co-OPR project web site

<http://www.aiai.ed.ac.uk/project/co-opr/>

I-X project web site

<http://www.aiai.ed.ac.uk/project/ix/>
<http://i-x.info>

I-X Co-OPR application resources

Available on specific request to project team

Abbreviations

The following abbreviations and acronyms are used within this report. They are collected together here to act as a reminder wherever the context is not clear.

AIAI	Artificial Intelligence Applications Institute
ATO	Air Tasking Order
CISA	Centre for Intelligent Systems and their Applications
COA	Course of Action
Co-OPR	Collaborative Operations for Personnel Recovery
CPX	Command Post Exercise
CSAR	Combat Search And Rescue
DARPA	Defense Advanced Research Projects Agency
D-Cog	Distributed Cognition
GIS	Geographical Information System
HTML	Hypertext Markup Language
HTN	Hierarchical Task Network
<I-N-C-A>	Issues – Nodes – Constraints – Annotations Ontology
I-DE	I-X Domain Editor
IP	Internet Protocol
I-P2	I-X Process Panel
I-Plan	I-X Planning System
I-Sim	Intelligent Simulator
I-X	Intelligent Technology Research Program
JPRA	Joint Personnel Recovery Agency
JPRC	Joint Personnel Recovery Center
JTFC	Joint Task Force Commander
MSEL	Master Scenario Event List
OODA	Observe – Orient – Decide – Act
O-Plan	Open Planning Architecture
PR	Personnel Recovery
PREP	Personnel Recovery (Experimental) Pack
PRETC	Personnel Recovery Education and Training Center
RCC	Resource Coordination Center
ROZ	Restricted Operations Zone
SHORe	Stimulus – Hypothesis – Option – Response
SOP	Standard Operating Procedure
SPIN	Special Instructions
TLX	Task Load Index
USJFCOM	US Joint Forces Command
VOC	Virtual Operations Center
XML	Extensible Markup Language