I-X: Task Support on the Semantic Web

Stephen Potter, Austin Tate and Jeff Dalton
Artificial Intelligence Applications Institute,
School of Informatics, The University of Edinburgh,
Crichton Street, Edinburgh, EH8 9LE, UK
{s.potter,a.tate,j.dalton}@ed.ac.uk

1 Introduction

The coordination of resource and activity to achieve some common objective is a key task within modern virtual organisations. The Semantic Web initiative promises to increase the number of knowledge and information resources available, presenting more (and more varied) opportunities for interaction. However, as the number and complexity of these interactions increases, so too does the need for task support tools. This extended abstract describes our research into support for mixed-initiative (that is, involving both human and computer agents) collaborative tasks in distributed environments. At the heart of this approach is the *I-X* technology. This is introduced in section 2, while section 3 illustrates the nature of the task support it offers through the description of two applications. Section 4 outlines some future directions that this work will pursue and the final section provides a summary and some conclusions.

2 I-X: A Task Support Architecture

The I-X[1] technology is intended to provide a well-founded approach to allow humans and computer systems to cooperate in the creation or modification of some product, be it a document, plan, design or physical entity [Tate *et al.*, 2003]. The I-X tools support users in selecting and performing processes and creating or modifying process products. A set of *issues* is associated with the process or product, representing unsatisfied requirements, problems arising from critique and so on. Both processes and process products are considered, in the abstract, to consist of (perhaps hierarchically composed) *nodes*: these correspond to activities in the process or parts of the product. The relationships between nodes are defined by a set of *constraints*. Finally, *annotations* can be associated with these elements to capture other, perhaps less formal, information surrounding the collaboration. Together, these elements constitute the <I-N-C-A> (<Issues-Nodes-Constraints-Annotations>) model and provide a unifying framework that allows the communication — using an XML encoding — of elements from one agent to another.

[1]The 'I' of I-X is meant to convey all of 'intelligent', 'intelligible', 'integrated' and 'issue-based', with the 'X' being the uninstantiated variable. See i-x.info for more about I-X.

2.1 The I-X Tool Suite

The principal interface to these tools, the *I-P*[2] (I-X Process Panel) can be seen, at its simplest, as a 'to-do' list for its user; however, when used in conjunction with other I-X agents, it can become a sophisticated workflow and messaging tool. A panel corresponds to its user's 'view', in <I-N-C-A> terms, of the current activity, and the current state of the collaboration is used to generate dynamically the support options the tool provides. For example, associated with a particular activity node might be suggestions for performing it using known procedural decompositions, for invoking an agent offering a corresponding capability, or for delegating the activity to some other agent.

The other tools in the suite include messaging tools and information viewers and editors, used, for example, to allow the user to specify relationships with other agents in the environment, and to create and publish Standard Operating Procedures (SOPs), generic approaches to archetypal activities. Particularly relevant to this discussion is the *I-Q* (I-Query) tool. I-Q is a generic I-X agent shell which, when embodied with the appropriate mechanisms, presents an interface to a particular Semantic Web information resource, providing seamless integration with other I-X agents.

3 Demonstration Applications

In this section we illustrate the use of I-X to support activity involving Semantic Web resources through the brief description of two demonstrations that have been developed.

3.1 Workshop Organisation

This application involves the following scenario: an official of a UK technology research funding body is charged with organising a workshop concerning some particular area of computer science so as to get an overview of its current state.[2] Accordingly, from a set of published SOPs, she selects *Organise workshop*. Now shown on her I-P[2] are the sub-tasks needed to achieve this goal, involving selecting attendees, choosing a location and date, fixing the agenda, and so on.

Further decomposing the *select attendees* task, the initial sub-task is *identify steering committee* for the workshop. An available I-Q agent is known to be capable of performing this

[2]Developed as part of the AKT Project: see www.aktors.org.

task for topics drawn from the ACM classification of computer science.[3] This agent constructs appropriate RDQL[4] queries and sends them via http to an RDQL interface onto an RDF triple store. This database describes the current state of UK research in (predominantly) computer science through some millions of triples extracted from various sources by various techniques, the triples being described according to a number of published ontologies.[5] The RDQL formed by the I-Q agent refers to these ontologies and implicitly contains knowledge of the contents of the triple store, and the agent 'knows' how to communicate with the store and process its responses. However, this is opaque to the I-P^2 user, who need know nothing about this transaction, and, having selected the appropriate topic from the ACM classification and parameterised her message to the I-Q agent, receives a message naming the suggested steering committee along with their contact details a few seconds later.

This sub-task completed, the other steps in the SOP are performed by the user (assisted by links to relevant tools and information) or delegated accordingly. Finally, to discuss this workshop and confirm its dates, location and content with the steering committee, she initiates a videoconference; an additional SOP, downloaded from a meeting-support website,[6] provides experience-based assistance with conferencing technology set-up.

3.2 Search and Rescue

This application involves more complicated interactions with Semantic Web resources. The scenario surrounds the coordination of resources to rescue and care for a downed aviator.[7]

On being alerted about the emergency, the SAR (Search And Rescue) coordinator, through his I-P^2, selects an appropriate SOP containing a number of sequential steps such as *select hospital* and *select SAR resource*. In this environment, the SAR domain and the infrastructures — including medical facilities — of the countries in the locale are encoded according to DAML-O ontologies, with both ontologies and knowledge bases available as web resources.[8]

A particular I-Q agent in this domain has the ability to access and reason with the appropriate ontologies, and so can extract from the knowledge bases information about hospitals offering specialist care facilities (for example, burns units). So, once the nature of the injuries to the airman has been established, this agent can be invoked to suggest the closest appropriate hospitals.

SAR resources — helicopters, patrol boats, etc — are described as DAML-S services, and advertised to a matchmaking service.[9] For the purposes of selecting amongst these resources, a second I-Q agent is able to construct and send to the matchmaker an appropriate DAML-S request, instantiated with the location of the airman and the location of the selected hospital. When selecting an appropriate resource, then, this agent can be invoked to act as an intermediary to the matchmaker, constructing appropriate requests and parsing the returned results.

[3]See `www.acm.org/class/1998/overview.html`.

[4]RDQL is an SQL-like query language for RDF; see: `www.hpl.hp.com/semweb/rdql.htm`.

[5]For more about the triple store see, see `triple-store.aktors.org`.

[6]`i-me.info/resources/coakting`.

[7]Developed in the course of the CoSAR-TS project: see `www.aiai.ed.ac.uk/project/cosar-ts`.

[8]See, for example, the infrastructure ontology at: `www.daml.org/experiment/ontology/infrastructure-elements-ont`, and the knowledge base about a (fictitious) country at: `sonat.daml.org/DAMLdemo/instances/enp/nc-BINNI.daml`.

[9]`www-2.cs.cmu.edu/˜softagents/daml_Mmaker/daml-s_matchmaker.htm`.

4 Future Directions

With particular reference to operating on the Semantic Web, there are a number of areas of work that would enhance the I-X support environment and encourage interoperability, and which we hope to address in the near future. For instance, publishing <I-N-C-A> information according to OWL ontologies would make resources such as SOPs more readily available to a wider community, while describing the capabilities of I-X agents using OWL-S would make these more visible externally, and position I-X more centrally within the developing ideas of web service description and invocation. More generally, some consideration of the whole notion of task support within the Semantic Web is needed: What sort of tasks will be performed? What sort of support is necessary/possible? How might this support best be delivered?

5 Summary and Conclusions

The intention of this extended abstract has been to describe the I-X environment for collaborative task support, with particular reference to placing this in the context of the Semantic Web and its emerging standards, concepts and resources. The potential benefits are mutual: on the one hand, I-X task support is greatly enhanced by exploiting Semantic Web information resources, as illustrated by the applications described above; on the other hand, as the Semantic Web moves towards its goal of empowering users to achieve more than information browsing, the need for integrated intelligent task support of the sort provided by I-X becomes more evident.

Acknowledgments

The work described in this extended abstract is supported, in part, by the AKT IRC and by the DARPA DAML program.

The University of Edinburgh and research sponsors are authorised 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.

References

[Tate *et al.*, 2003] A. Tate, J. Levine, J. Dalton, and A. Nixon. Task Achieving Agents on the World Wide Web In *Spinning the Semantic Web*, Fensel, D., Hendler, J., Liebermann, H. and Wahlster, W. (eds.), MIT Press, 2003.