

Open Research Online

The Open University's repository of research publications and other research outputs

CoAKTinG: Collaborative Advanced Knowledge Technologies in the Grid

Conference or Workshop Item

How to cite:

Buckingham Shum, Simon; De Roure, David; Eisenstadt, Marc; Shadbolt, Nigel and Tate, Austin (2002). CoAKTinG: Collaborative Advanced Knowledge Technologies in the Grid. In: Second Workshop on Advanced Collaborative Environments, Eleventh IEEE Int. Symposium on High Performance Distributed Computing (HPDC 11), 24-26 Jul 2002, Edinburgh, Scotland.

For guidance on citations see [FAQs](#).

© 2002 The University of Edinburgh

Version: Accepted Manuscript

Copyright and Moral Rights for the articles on this site are retained by the individual authors and/or other copyright owners. For more information on Open Research Online's [data policy](#) on reuse of materials please consult the policies page.

oro.open.ac.uk

CoAKTinG: Collaborative Advanced Knowledge Technologies in the Grid

Simon Buckingham Shum[†] David De Roure^{††} Marc Eisenstadt[†] Nigel Shadbolt^{††} Austin Tate^{*}

[†]Knowledge Media Institute
Open University
Milton Keynes
MK7 6AA, UK
sbs@acm.org
m.eisenstadt@open.ac.uk

^{††}Intelligence, Agents, Multimedia Group
University of Southampton
Highfield
Southampton SO17 1BJ, UK
dder@ecs.soton.ac.uk
nrs@ecs.soton.ac.uk

^{*}Artificial Intelligence
Applications Institute
University of Edinburgh
80 South Bridge
Edinburgh EH1 1HN, UK
a.tate@ed.ac.uk

Abstract

Grid infrastructures coupled with semantic web linkage and reasoning open up intriguing new possibilities for scientific collaboration. In this short paper, we outline the research agenda and collaboration technologies under development within the CoAKTinG project: Collaborative Advanced Knowledge Technologies in the Grid. CoAKTinG will provide tools to assist scientific collaboration by integrating intelligent meeting spaces, ontologically annotated media streams from online meetings, decision rationale and group memory capture, meeting facilitation, issue handling, planning and coordination support, constraint satisfaction, and instant messaging/presence. Their integration is illustrated through an extended use scenario.

1 Introduction

The *Advanced Knowledge Technologies* Interdisciplinary Research Collaboration [AKT IRC: www.aktors.org] is a 6 year, \$10M project to develop knowledge management technologies, funded by the UK's Engineering and Physical Sciences Research Council (EPSRC). The related *CoAKTinG* project [www.aktors.org/coaktin], funded as part of the UK's e-Science Initiative on Grid computing, aims to integrate and adapt AKT and related technologies specifically to support distributed scientific collaboration. As part of the AKT project's conception of the convergence of knowledge technologies and grid computing as the Semantic Grid [www.SemanticGrid.org], CoAKTinG will provide tools to assist scientific collaboration by integrating intelligent meeting spaces, ontologically annotated media streams from online meetings, decision rationale and group memory capture, meeting facilitation, issue handling, planning and coordination support, constraint satisfaction, and instant messaging/presence. These approaches are summarised below.

Smart spaces. Scientists may wish to be in a variety of places when they are in communication with remote colleagues: experimental labs, meeting rooms, data analysis suites, or travelling. This component of the project will be combining Access Grid node spaces [www-fp.mcs.anl.gov/fl/accessgrid/ag-spaces.htm] with portable smart devices that support a variety of broad and narrow bandwidth connections to other people/devices. A smart space, as we conceive it, will recognise significant events in a meeting and insert metadata into the AV stream, described next.

Ontologically annotated audio/video streams. Few researchers have the time to sit and watch videos of meetings; an AV record of an online meeting is thus only as useful as its indexing. Moreover, indexing effort must negotiate the cost/benefit tradeoff or it will not be done. Our prior work has developed ways to embed 'continuous metadata' (of which one form is hyperlinks) in media streams [3]. We will now be embedding metadata grounded in one or more ontologies for scientific collaboration. Additionally, decisions and key discussions (as captured in Compendium – see below) can be recovered.

Issue handling, tasking, planning and coordination. We will be building applications using I-X Intelligent Process Panels [2, 6] and their underlying <I-N-CA> (Issues, Nodes, Critical and Auxiliary) constraint-based ontology for processes and products [6] [i-x.info]. The process panels provide a simple interface that acts as an intelligent "to do"

list that is based on the handling of issues, the performance of activity or the addition of constraints. It also supports semantically task directed “augmented” messaging and reporting between panel users. A common ontology of processes and process or collaboration products based on constraints on the collaborative activity or on the alternative products being created via the collaboration is the heart of this research. We envisage the creation of a library of process models to support the issues, options and constraints associated with common types of meeting held by a given scientific group.

Collective sensemaking and group memory capture. Whilst meetings are a pervasive knowledge-based activity in scientific life, they are also one of the hardest to do well. “Meeting technologies” tend either to over-structure meetings (e.g. Group Decision Support Systems), or ignore process altogether, and simply digitize physical media (e.g. whiteboards) for capturing the products of discussion. The *Compendium* approach [www.CompendiumInstitute.org] occupies the hybrid middle-ground – ‘lightweight’ discussion structuring and mediation plus idea capture [5], with import and export to other document types. “Dialogue maps” [CogNexus.org] are created on the fly in meetings providing a visual trace of issues, ideas, arguments and decisions.

Enhanced presence management and visualisation. The concept of *presence* has moved beyond the ‘online/offline/away/busy/do-not-disturb’ set of simple state indicators towards a rich blend of attributes that can be used to characterise an individual's physical and/or spatial location, work trajectory, time frame of reference, mental mood, goals, and intentions. Our challenge is how best to characterise presence, how to make it easy to manage and easy to visualise, and how to remain consistent with the user's own expectations, work habits, and existing patterns of Instant Messaging (IM) and other communication tool usage. Working with the Jabber open source XML-based communications architecture, we will be extending its IM capabilities with ‘ontology of presence’ and ‘knowledge profiles’. A prototype called *BuddySpace* [1] [kmi.open.ac.uk/projects/buddyspace] also adds visual and map-based ‘buddy lists’ to display presence information that is mapped onto visualisations, both geographical (e.g. a map of a building, or a region), and conceptual (e.g. a workflow chart or project plan, a design or experiment). The scale of the map can be altered to reflect anything from global positioning, to school and workplace office layouts, to experimental assemblies. Moreover, not only people have presence states: devices, documents, and indeed any arbitrary resource can have a presence state indicated on our ‘desktop radar’ display.

2 Use Scenario

We now present an extended scenario to illustrate how these tools could be usefully combined to support scientific collaboration. This is one of a series of use scenarios that we are using to drive the integration work. Each of the tools introduced above, and appearing in the scenario, is an implemented system. However, their integration as described in the scenario is fictitious, and some of the user interfaces have been visually augmented to reflect how we envisage user interaction.

A research team is holding a meeting over the internet to discuss the results of a recent series of experiments. Three of the team (Anna, Ben and Clive) are present via high bandwidth Access Grid meeting rooms that their institutions have, providing multi-screen video channels and high quality audio to the CollabClients on their laptops. Daisy’s university hasn’t yet installed an Access Grid room, so she is joining them via a desktop CollabClient in her office. Elli cannot make the meeting, as she is on the road to a conference.

Ben has displayed the agenda on his process panel, and also as a visual Compendium map of open issues, ready to capture their discussion as it unfolds (Figure 1).

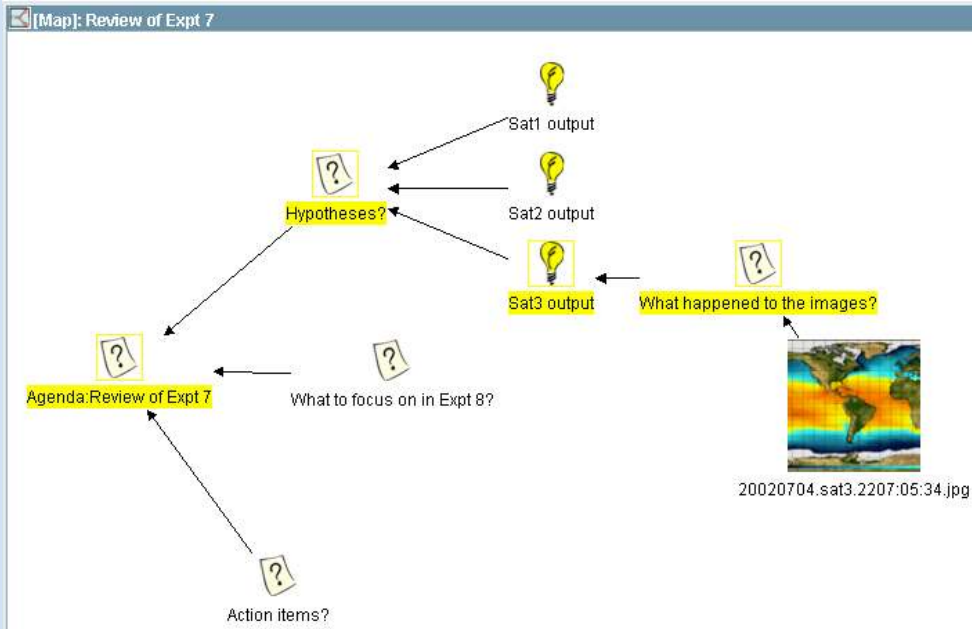


Figure 1: Compendium dialogue map setting the agenda, ready to capture discussion

Anna shows a graph from the last experiment and they discuss this, scribbling a few notes on it. They discuss the three options facing them for the next experimental run, using Ben’s dialogue map to track the pros and cons for each option. One issue re-opens a discussion that they started two months earlier at a workshop. Ben opens up the map for this discussion so they can remind themselves what they covered last time. Several of the ideas generated in that discussion were contributed by Felix, who has moved to another project now. However, on mousing over an *Idea* icon referencing him (Figure 2), they can see that he’s online, but not video-enabled, so they open up a text chat window (Figure 3).

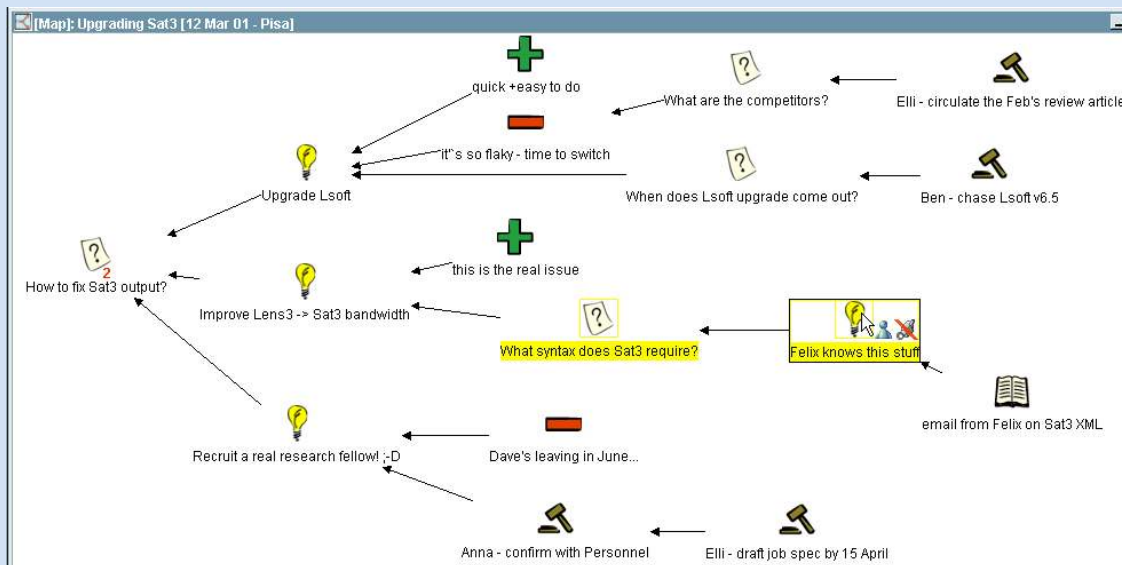


Figure 2: Revisiting the map from a “Sat3” discussion two months earlier: mousing over a contribution from Felix indicates his current availability online.

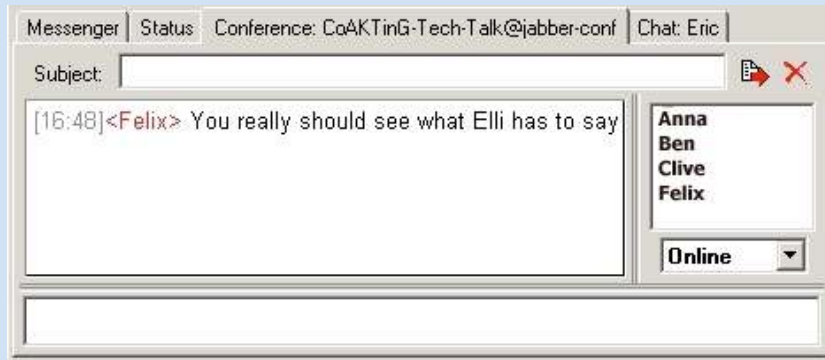


Figure 3: The meeting opens the recommended communication channel (textchat) to Felix

Felix provides some helpful commentary, and then refers them to Elli, who has flagged her own high-priority interest in (and capabilities related to) this topic, as an earlier contributor to supporting arguments.

Although she's on the road right now, Elli's Presence Manager tool (Figure 4) knows that high-priority alerts on subjects that she has flagged as important can be re-routed automatically to her mobile phone as a plain text message alert, so such an alert is sent to Elli's phone.

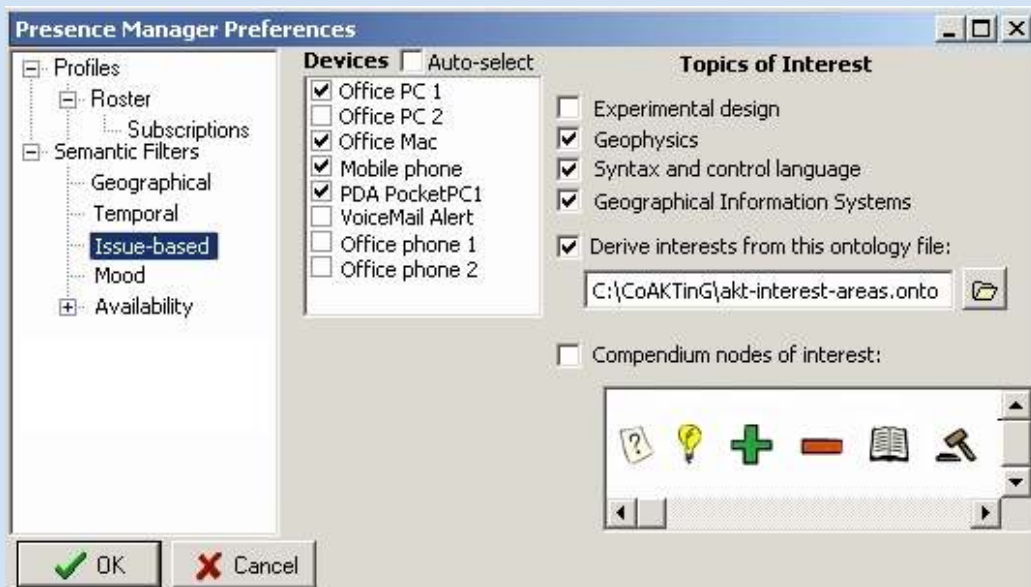


Figure 4: Elli's Presence Manager, giving her control over how she wants to be alerted to different kinds of messages.

Ben marks their agreement on other actions, and from this generates a 'To Do' list. On each of their laptops, a small panel appears listing action items, showing each item's status, waiting to be checked as "Done" (Figure 5). Some actions begin to be dealt with via intelligent agents authorised to perform these autonomously.

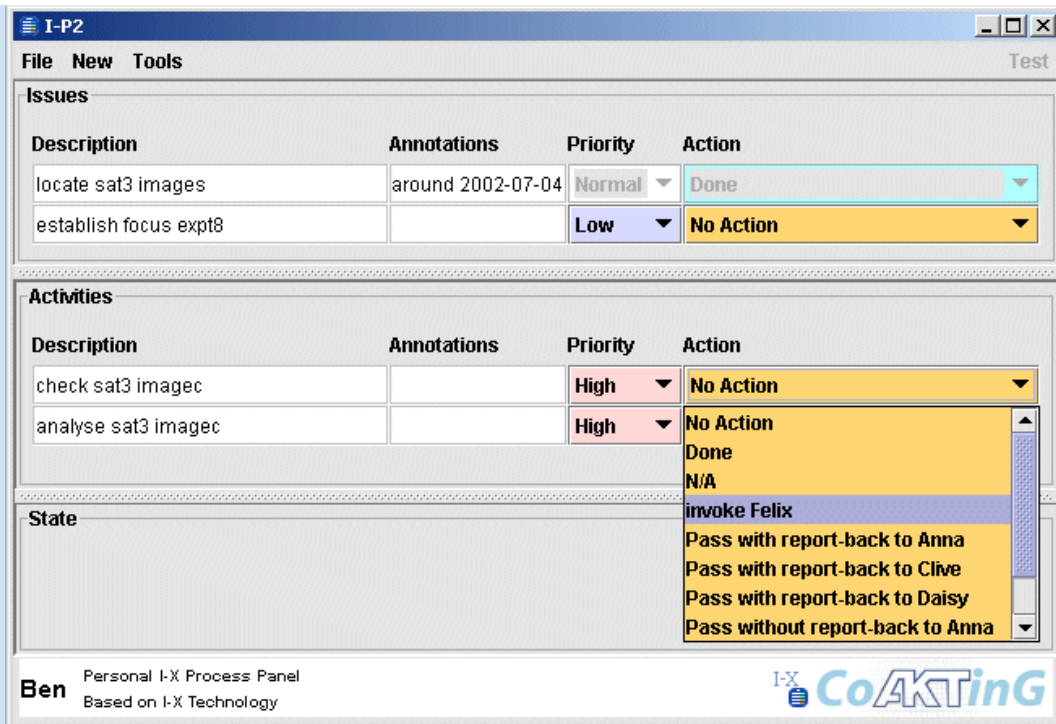


Figure 5: Ben's Process Panel: as issues and action items are agreed in the meeting and captured in the dialogue map, items appear on his active 'To Do' list.

Two hours later when she's checked in at her conference hotel, Elli, already aware of a high-priority query because of her text message alert, boots her laptop and on her desktop sees at a glance that the team are all "Away", and that a project bid still needs one more signature. But flashing in red is a *Question* icon from Compendium: *What syntax does the Sat3 require?* (Figure 6).

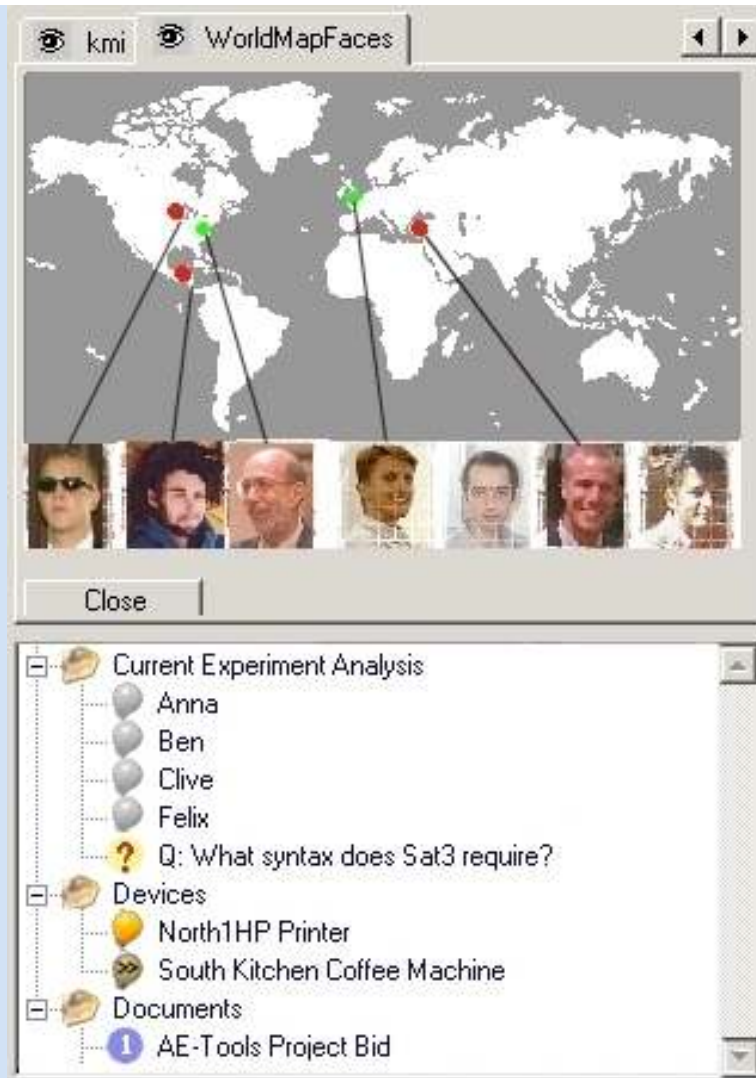


Figure 6: Elli's presence list indicates the state of key people, devices, and documents, and flashing in red, an open issue on which the team needs her input.

Mousing over it, she sees this is both an issue relevant to her own high-priority interest and from the datestamp, an issue on which the team needed her input that afternoon. She double-clicks on it to launch her CollabClient (Figure 7). It displays Clive presenting his slide and the state of the Compendium dialogue map at the time. She replays the key minute of discussion between Anna and Clive when they needed her input. She sees the messaging window appear as they connect to Felix, and sees his reference to consult her.

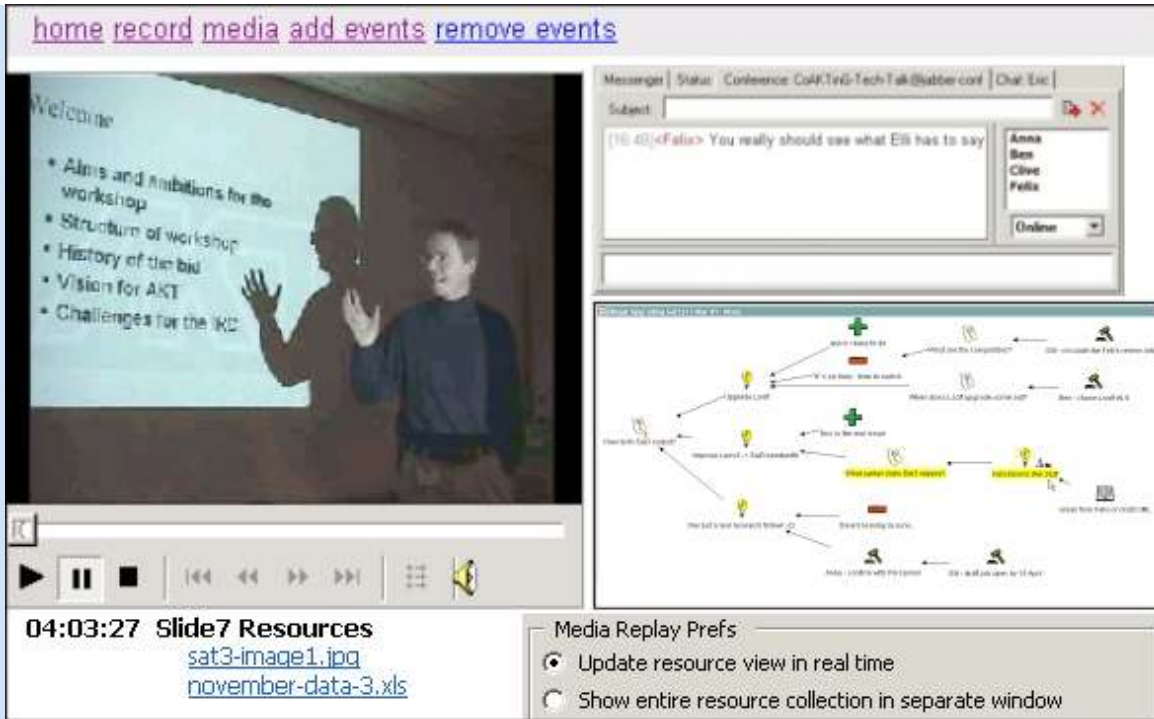


Figure 7: Elli replays the key segment from the meeting that she missed.

She brings the dialogue map to the front, records an audio annotation, and places it as an answer to the open question on the map.

Her *To Do* panel is automatically synchronised as she does this. She briefly scans it, noting with a groan the new items that have appeared, then leaves her agent to do what it is authorised to act upon while she goes to bed. She will see what it has left her that needs her attention in the morning.

The above scenario is generated from considering feasible integrations (within the next two years) between the technologies summarised at the start. Compendium's dialogue maps can show the presence status managed by Jabber of people, documents, and other artifacts. Issues, action items and constraints resulting from discussions can lead to I-X Process Panel entries and its planning and execution aids can semi-autonomously handle some of these. Nodes on dialogue maps can be made into active objects indicating status on desktops. Audio/video can be arbitrarily replayed with significant events (such as relevant application events) embedded as metadata in the media stream. Devices and the software tools that utilise them to communicate with the various users know their limitations, and so do the best they can when someone tries to contact or use them.

3 Acknowledgements

This work is supported partially under the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University. The I-X project is sponsored by the Defense Advanced Research Projects Agency (DARPA) under grant number F30602-99-1-0024.

The authors' employers 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 official policies or endorsements, either express or implied, of other parties.

4 References

[1] Eisenstadt, M. (2002). From Buddy Lists to Buddy Space: Scaleable Experiences of InterPersonal Presence. *Proceedings of the Presence and Interworking Mobility Summit (PIM2002)*, June 11-13, 2002, Helsinki, Finland. [www.pulver.com/pim]

[2] Levine, J., Tate, A. and Dalton, J. (2000). O-P³: Supporting the Planning Process using Open Planning Process Panels. *IEEE Intelligent Systems*, Vol. 15, No. 6, November/December 2000.

[3] De Roure, D.C., Moreau, L. and Hall, W. (2002). HyStream - Applying Open Hypermedia to Multimedia Streams. *Individual Grant Review Report, EPSRC GR/M84077/01* [www.ecs.soton.ac.uk/~dder/hystream/hystream-report.htm]

[4] Selvin, A., Buckingham Shum, S., 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*, March 4-7, 2001, Austin TX. This and other case studies at: [www.CompendiumInstitute.org]

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

[6] Tate, A., Levine, J., Dalton, J. and Nixon, A. (2002). Task Achieving Agents on the World Wide Web, In *Creating the Semantic Web*, Fensel, D., Hendler, J., Liebermann, H. and Wahlster, W. (eds.), 2002, MIT Press.
