

Augmenting Design Deliberation with Compendium: The Case of Collaborative Ontology Design

Simon Buckingham Shum, Enrico Motta, John Domingue

Knowledge Media Institute, Open University, UK

sbs@acm.org

1 Introduction

In this position paper, I describe ongoing work using the [Compendium](#) approach to hypertext-augmented collaborative modelling (HACM). Compendium is being used to support a team specifying, designing and implementing an ontology-based information portal.

The [Advanced Knowledge Technologies](#) (AKT) consortium is a six year research project between five UK universities. One of its main goals is to build significant demonstrations of knowledge management tools, some of which are based on ontologies, that is, explicit conceptual models of concepts, their attributes and relationships. The [Mifflin](#) hypertext tool for Compendium is itself a tool that will be developed as part of a new AKT-related project ([CoAKTinG](#)), and so is being evaluated and refined through a bootstrapping approach in which we use it to support our own work.

Mifflin is being used by the author to support AKT members in designing a prototype called the AKTive Portal (designed to be a next generation portal infrastructure that supports the capture, indexing, dissemination and querying of information). The first application of the AKTive Portal is for the AKT Project itself, based on an underlying ontology for scientific knowledge management which describes the world of AKT research, such as researchers of different levels and their affiliations, documents of various status, research events, AKT technologies, methods and research themes, etc. The structure of the ontology had to be designed from three existing ontologies that partially described research processes and products, and Mifflin was used over a series of meetings not only as an integrative medium to capture general discussions associated with the AKTive Portal (e.g. responsibilities and deadlines), but in particular, for visualizing the structure of the ontology, in order to agree on decisions, and capture the rationale. The resulting ontology is now being deployed and refined, and Mifflin will be used to support this maintenance and revision phase of the lifecycle.

In the next section I illustrate the way in which Mifflin was used to support the design process, including its integration with other technologies already in use, through data import and export. I then reflect on the lessons learnt to date.

2 Compendium-supported collaborative ontology design

Let us start with the physical setting for this work. Figure 1 shows a room set up for a Compendium-assisted collaborative design session. The key addition to a conventional meeting room with a data projector is the second projector and screen on the right, driven by the Compendium facilitator. All participants can see both screens without difficulty, so that attention can be switched without effort to either screen.

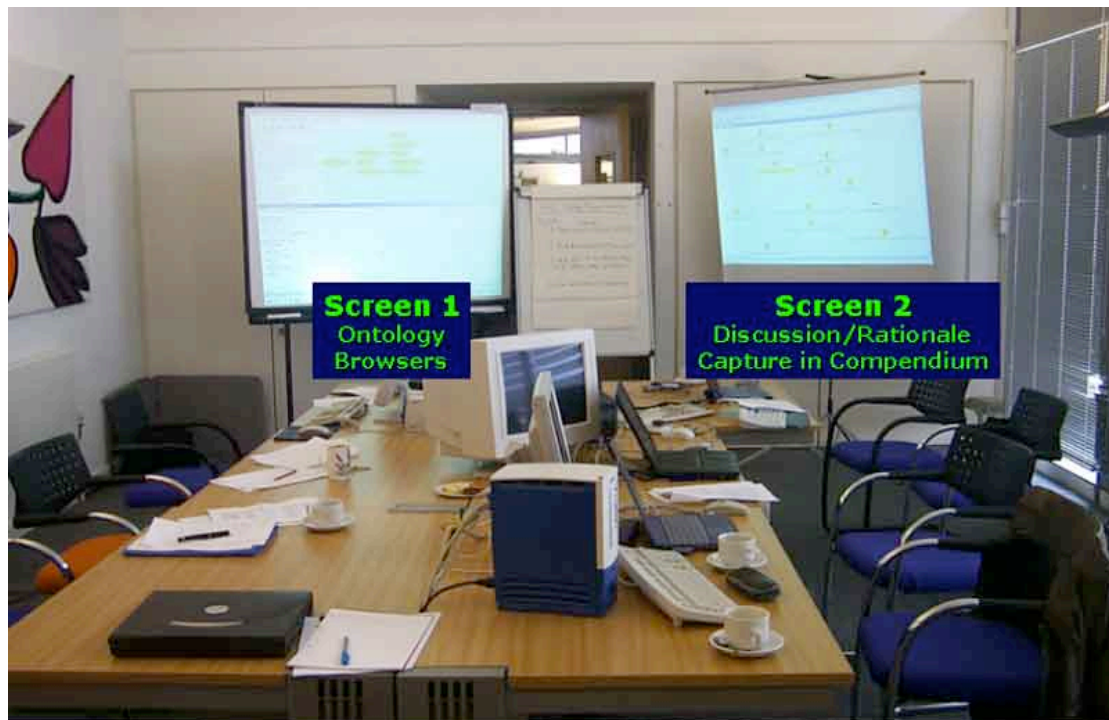


Figure 1: Room setup for a Compendium-assisted collaborative design session. The key addition to a conventional meeting room is the second data projector and screen on the right, driven by the Compendium facilitator.

A workshop was held to specify the requirements for the AKTive Portal, with Mifflin being used to structure and capture the discussion. For instance, Figure 2 shows a number of scenarios for querying the Portal, generated in this workshop, a map that was returned to several times in the subsequent design sessions to verify that the ontology could support the scenarios.

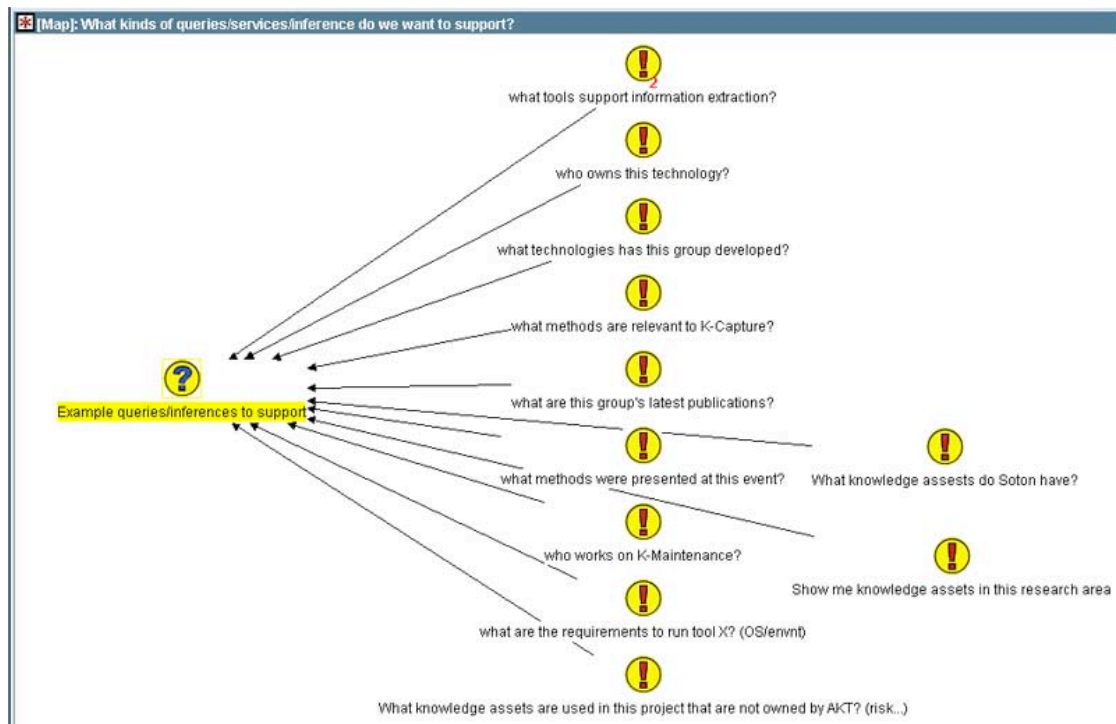


Figure 2: Capturing a discussion on end-user scenarios that the system should support. This served as a reference point to verify the design as it emerged in subsequent sessions.

The results of this first workshop were exported to HTML and published to the project as a web discussion document using the [D3E](#) toolkit (Figure 3).

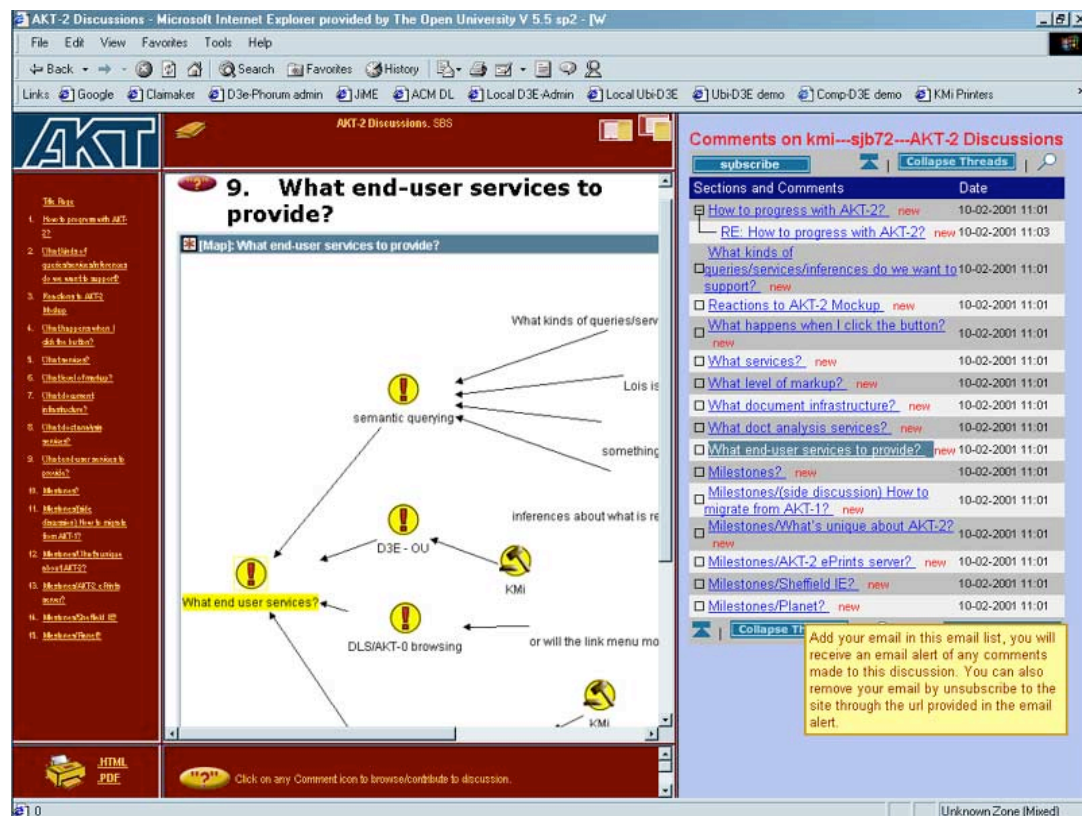


Figure 3: Conversion of a Mifflin HTML export to a D3E discussion document.

A member of the project then wrote a document proposing a way to merge three existing ontologies for scientific knowledge management (plus elements from ‘upper level ontologies’ of general concepts). This became the point of departure for discussions in a subsequent workshop. The document was exported from Microsoft Word into Mifflin (each paragraph becomes a hypertext node, and the document’s tree structure is graphed explicitly: Figure 4). The nodes for the key paragraphs specifying the primary classes required for the ontology were then transcluded into new maps in order to discuss them.

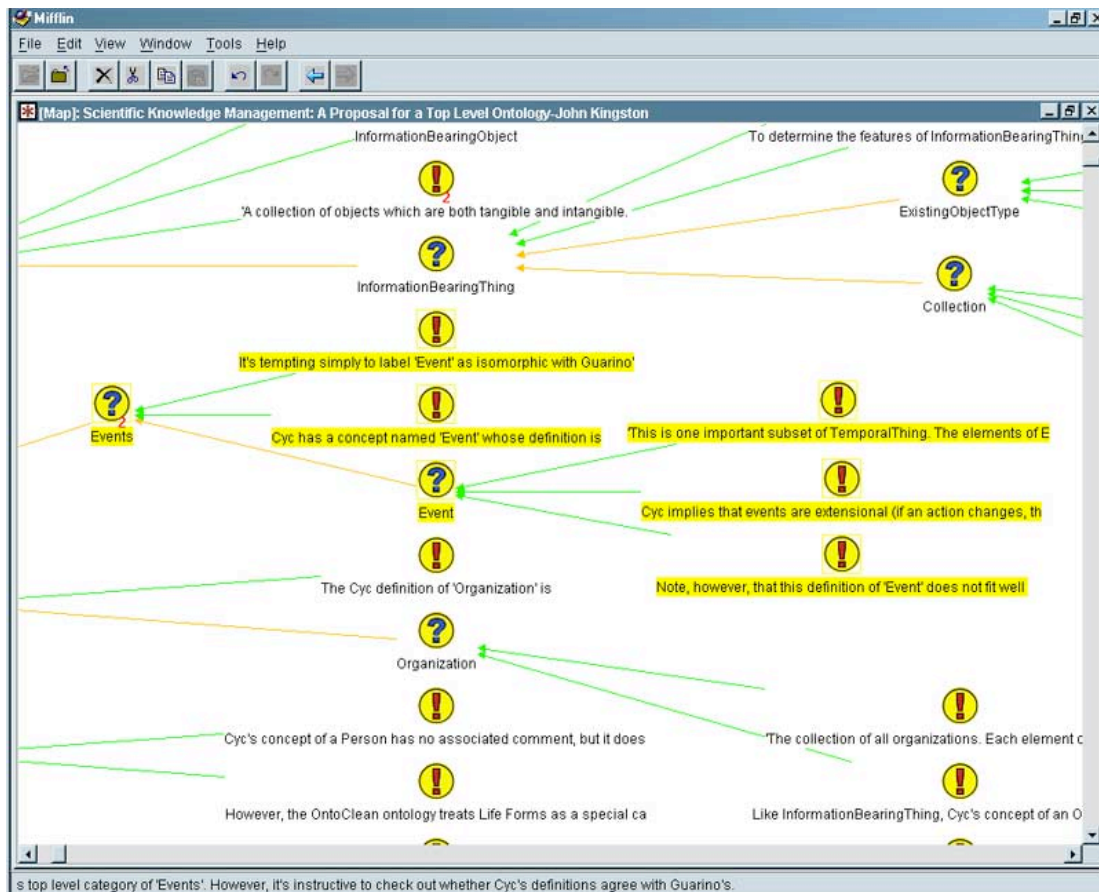


Figure 4: Import from Microsoft Word of the draft specification document. The key paragraphs for debate (highlighted) are transcluded into new maps, providing the anchor points for discussion.

Not surprisingly, ontology merging requires easy access to, and switching between, the source ontologies. All three of the existing ontologies were accessible through ontology browsers (on various laptops in the room that could be displayed), and the upper level ontologies were on the web. These were all displayed on Screen 1 (see Figure 1) as required.

Mifflin's 'agenda map' set a visual compass for the material to be covered over the day-long workshop (Figure 5).



Figure 5: The agenda map in Mifflin specifying the ontology classes that have been agreed, and now need to be designed.

On Screen 2 (see Figure 1), free-flowing discussion about a class was mapped and displayed in Mifflin (e.g. Figure 6 on *Information Bearing Object*).




Figure 6: Dialog Map of free-flowing discussion about the ontology.

Mifflin was used not only to capture/mediate free-flowing discussion as in [Dialog Mapping](#), but following the Compendium approach of embodying a formal methodology within Mifflin's simple Question/Answer notation, each unspecified element in the ontology became

a specific type of *Question* node (e.g. *Subclass?*, *Slots?*), answers to which were constrained to legitimate ontology entities (*classes*; *slotnames*), at least, as far as those present could tell at that moment (some changes were still required when it came to implementation). The associated ontology structure was thus mapped, either linked directly to the discussion, placed underneath it, or in a separate map, depending on size (e.g. Figure 7 shows the class structure for *People*).

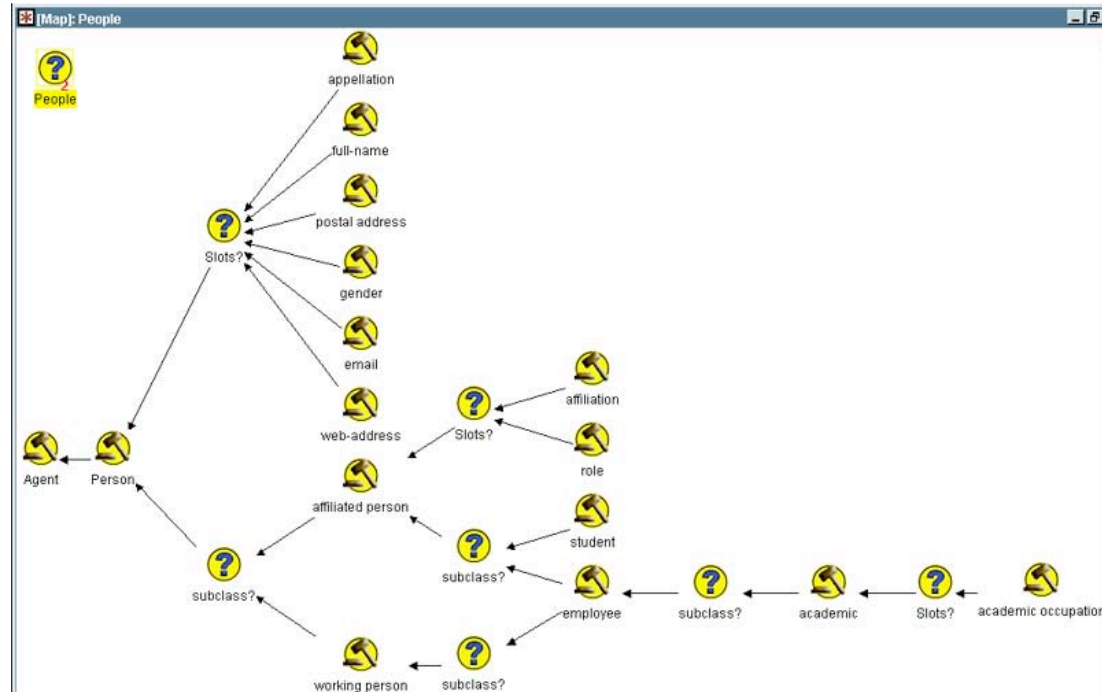


Figure 7: Modelling the structure of the *People* class within the ontology, systematically using *Question* nodes (*Subclass?*, *Slots?*) to open and complete discussion about each unspecified element.

As this structure was mapped in Mifflin and agreed, an ontology engineer also began to check this ‘on the fly’ in an ontology engineering environment on his laptop ([WebOnto](#) [1], periodically displayed on Screen 1). This more detailed work helped to verify answers to questions that arose, and in some cases highlight issues that had not been addressed (Figure 8).

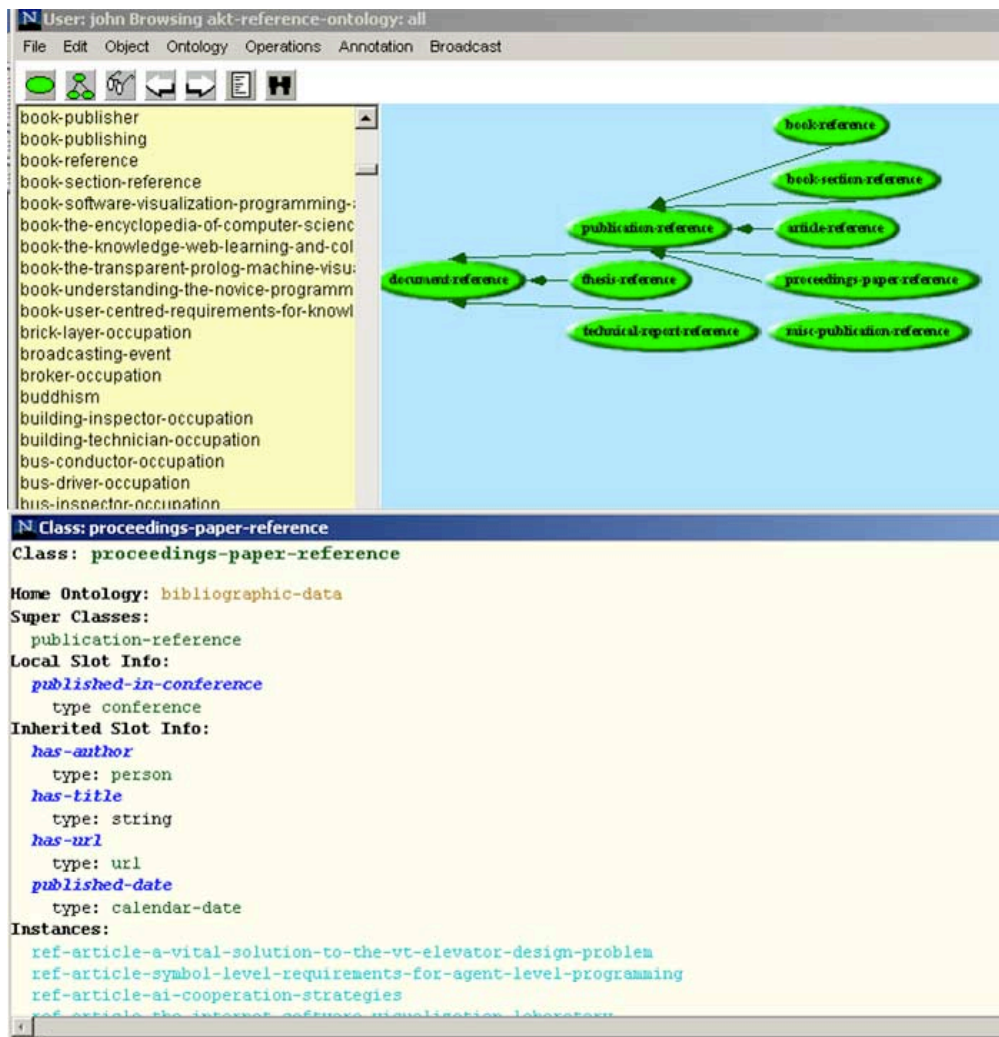


Figure 8: Checking the new ontology in a formal ontology engineering environment. The workshop concluded with agreement on the next steps to be taken (Figure 9).




Figure 9: Wrapping up the workshop by agreeing on actions.

Following the workshop, the agreed specification for the ontology was implemented (requiring further ‘cleaning’) in a formal ontology modelling language ([OCML](#)), and version 1 of the [AKT Reference Ontology](#) was released. An HTML export from the OCML environment was converted using D3E into a private project discussion document (Figure 10), which is the locus of all feedback from partners as they begin to build technologies on the ontology, and new issues arise.

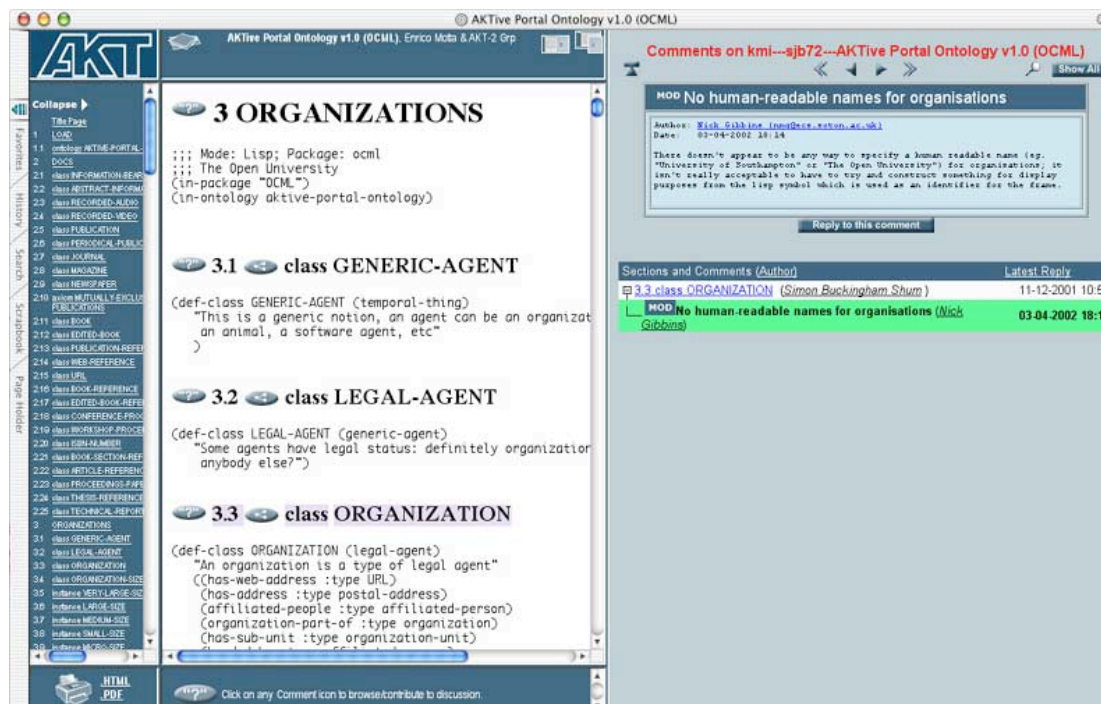


Figure 10: The ontology from the workshop is published for feedback and further discussion as it is deployed.

3 Reflections on Compendium

Although still at a relatively early stage, this work is starting to provide a non-trivial Mifflin database in support of a real system development project. Mifflin’s main functions in the project to date have been to provide:

- structure for face-to-face collaborative sensemaking: what needs to be tackled, what has been tackled, what are the arguments, what are we agreeing?

- the rationale for a (local) ontology engineer when implementing the agreed specification

- a memory aid for the group in meetings (facilitator brings up maps of previous discussions), and between meetings (facilitator circulates Mifflin maps of agreed action items and responsibilities)

- a memory aid for the facilitator to track decisions after meetings in order to coordinate the project

The following reflections on the process of using Compendium summarise where we have got to in this case study, and are offered as contributions to the workshop’s discussions.

Dialog Mapping interventions. Mifflin was used in this case firstly for what Jeff Conklin calls *Dialog Mapping*: facilitated, IBIS-based capture, display and reflection back to participants of their deliberations. As a facilitator who knew the participants reasonably well, and who was also coordinating the development of the AKTive Portal, I felt confident to chair the meeting, push for clarification of contributions, and push for commitments to action items, all captured in Mifflin. In contrast, in other case studies Compendium practitioners

have been brought in as outsiders who do not know the participants, and with little understanding of the specific problem under discussion. The process-centric nature of the approach is evident in that they report still being able to make valuable contributions to such meetings (although not always).

Compendium interventions. The specifically *Compendium*-oriented aspect of this use of Mifflin was the additional constraining of the maps for modelling ontology class structures using *Subclass?* and *Slots?* Question nodes. This was understood and accepted by the participants without explanation. The resulting structures were not directly exported into the OCML/WebOnto ontology engineering environments, but transcribed. However, it is not clear at this point if much additional value would have been gained from some form of Mifflin-to-OCML export, since there remains significant modelling work to properly specify an ontology that goes beyond what was captured in Mifflin in the meeting.

‘Mifflin literacy’ and scepticism. The participants in the meetings described were researchers in knowledge-based technologies, and as such, extremely literate with formal knowledge representations (Mifflin would be considered semi-formal by this community). Interestingly, some members in this group had already expressed some scepticism about Compendium, similar to ‘more normal, less formal’ newcomers in other Compendium case studies who are not used to having their discussions mediated and structured. By the end of the workshops reported, the scepticism had changed to enthusiasm, to the extent that participants came to expect their contributions to be captured, and (half-jokingly) expressed concern when I was not in the room.

Placement of the Mifflin display screen. Although an apparently trivial point, in previous uses of Mifflin where the screen was positioned too far from the main screen, it felt more awkward to direct attention to Mifflin. What counts as ‘too far’ depends on the orientation of the participants. Screen 2 (in Figure 1) could in fact be moved right next to Screen 1 if the flipchart between them was not in use.

Compatibility with existing software tools. The importance of HACM tools being compatible with existing tools has already been proven in a previous case study [1], in which Microsoft Word and Visio diagrams were generated from constrained Compendium maps driven by specific questions. Mifflin’s Word import and HTML export has been used in this case, although it is fair to say that whilst very useful, these have not proven *critical* to this project so far. The Word import provides paragraph-level granularity tracing back to the imported source document (although we could of course implement this in other ways). The HTML export has enabled us to rapidly publish Mifflin maps back to Word, and to the web, in particular, for discussion in D3E. This is useful for the simple reason that other participants are not yet running Mifflin, and so depend on Word/HTML for their record. We do not know if they have used Mifflin records themselves in this form, but think it unlikely to date (see above for Mifflin’s main role so far).

We anticipate that the D3E sites will be extremely important as we move into longer term deployment and revision of the ontology. It is our intention to implement a Mifflin import filter that can convert the threaded web discussion back into a hypertext map, in order to ‘complete the loop’ and support iterations from synchronous work (Mifflin), to asynchronous (D3E), and back again, providing a more complete group memory trace of all deliberations. This would be an example of what Engelbart has proposed as a [Dynamic Knowledge Repository](#).

Issues concerning the ‘art and craft’ of Compendium practice. The workshop call highlighted some specific questions concerning the difficulties that beginners often have in learning to ‘do Compendium’ in meetings. “Nothing worthwhile comes for free” goes the adage, so what is the nature of the costs?

Literacy. I have written about this elsewhere using a [literacy metaphor](#): we need to practise reading and writing in new notations and software environments, just as we

have practised reading and writing natural language prose from an early age. To develop this, once we have learned our ‘ABC’, we learn for the rest of our lives how to write effectively in different genres, and in different computer-supported environments. I also find compelling the metaphors proposed by Conklin and Selvin, of mastering a musical instrument, or improvising in a band.

Cognitive overheads. My [empirical studies of designers](#) using the Xerox QOC notation (an IBIS derivative) with pen and paper demonstrated the cognitive ergonomic challenge of chunking, naming, classifying and linking ideas, in real time, and in support of different modes of designing.

Personal confidence. One then needs to add to this the pressure of doing it in front of potentially sceptical peers, and in real meetings, and the obstacles can be intimidating to the newcomer. Confidence needs to be built in safe spaces such as training workshops, and in one’s own private space for personal work, or on a non-displayed laptop in meetings, or post-hoc after meetings.

From white to black belt. There’s no doubt that an intuitive ability to sense what is going on a meeting, and make appropriate interventions, is required to work as a ‘black belt’ practitioner. It is an interesting question to reflect on whether the IBIS notation itself provides leverage in this: that is, can someone learn to make good interventions more quickly when using IBIS, than without? There are certainly visual affordances from dialogue maps that would suggest that this can indeed be possible (e.g. significant gaps). However, lower levels of practice are still possible simply by driving a meeting from templates, which relieve the practitioner of much of the cognitive overhead by placing in external, visual memory much of representational scaffolding required: what questions to ask, how to structure them, how to code them with metadata, how to transclude them. Referring back to the workshop call, the road to “jazz musician” may therefore involve going via “automobile repairman” (or perhaps “instrument repairman”), and far from being a waste of time, this journey is crucial to really understanding your tools.

4 Acknowledgements

This work was funded in part by 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. We acknowledge the contributions of our colleagues to the AKT Reference Ontology, which will be detailed in future documents.

5 References

- [1] Domingue, J. (1998) Tadzebao and WebOnto: Discussing, Browsing, and Editing Ontologies on the Web. *11th Knowledge Acquisition for Knowledge-Based Systems Workshop*, April 18th-23rd. Banff, Canada.
 - [2] Motta E. (1999). *Reusable Components for Knowledge Modelling*. IOS Press, Amsterdam, The Netherlands. ISBN: 1 58603 003 5. <<http://kmi.open.ac.uk/people/motta/book/>>
 - [3] Selvin, A. M. and Buckingham Shum, S. J. (2002, in press). Rapid Knowledge Construction: A Case Study in Corporate Contingency Planning Using Collaborative Hypermedia. *Knowledge and Process Management*, Vol. 9, (Issue 2) (April/June, 2002). <<http://www3.interscience.wiley.com/cgi-bin/jtoc?ID=6242>> PrePrint: <<http://kmi.open.ac.uk/tr/abstracts/kmi-tr-92.html>>
-