KNOWLEDGE MEDIA KMi INSTITUTE

Memetic: An Infrastructure for Meeting Memory

**Simon Buckingham Shum, Roger Slack, Michael Daw, Ben Juby,
Andrew Rowley, Michelle Bachler, Clara Mancini,
Danius Michaelides, Rob Procter, David De Roure,
Tim Chown, Terry Hewitt**

The Open University

Memetic:
An Infrastructure for Meeting Memory

Simon BUCKINGHAM SHUM [a1], Roger SLACK [b], Michael DAW [c], Ben JUBY [d],
Andrew ROWLEY [c], Michelle BACHLER [a], Clara MANCINI [a],
Danius MICHAELIDES [d], Rob PROCTER [b], David DE ROURE [d], Tim CHOWN [d],
Terry HEWITT [c]

[a] *Knowledge Media Institute, The Open University, UK*
[b] *School of Informatics, University of Edinburgh, UK*
[c] *Access Grid Support Centre, Manchester Computing, University of Manchester, UK*
[d] *Intelligence, Agents, Multimedia Group, School of Electronics and Computer Science*
University of Southampton, UK

Abstract. This paper introduces the Memetic toolkit for recording the normally ephemeral interactions conducted via internet video conferencing, and making these navigable and manipulable in linear and non-linear ways. We introduce two complementary interaction visualizations: argumentation-based concept maps to elucidate the conceptual structure of the discourse using a visual language, and interactive event timelines generated from the meeting metadata. We discuss in detail the affordances of Memetic's tools, in particular the *Compendium* hypermedia mapping tool, and the *Meeting Replay* tool that renders the semantic navigation indices into the videoconference replays. Additionally, with respect to methodology and evaluation, we describe how we are engaging diverse end-user communities in the process of designing and deploying these tools.

1. Introduction

'Meetings are where organizations come together. (They) remain the essential mechanism through which organizations create and maintain the practical activity of organizing. They are, in other words, the interaction order of management, the occasioned expression of management-in-action, that very social action through which institutions produce and reproduce themselves.'[1]

---

[1] Corresponding Author: Simon Buckingham Shum, Knowledge Media Institute, The Open University, Walton Hall, Milton Keynes, MK12 5AY, UK; E-mail: sbs@acm.org

The meeting is a pervasive feature of everyday work life and, not surprisingly, there have been numerous attempts to support meeting activity with technology. Examples of previous work in this area include: Bush's [2] Memex – with its 'associative indexing' of texts and artifacts; Engelbart's [3] NLS /AUGMENT[2], which enabled navigation through 'complex information structures' and conceptual mapping; Stefik et al's [4] Colab, which focused on collaborative problem solving and documentation; and the myriad of video whiteboard systems. For the Memetic[3] project, the core issue is getting a sense of what has been achieved in the meeting, how decisions have been made and by whom. This is not meant in a normative sense, rather the aim is to provide a 'natural history' of the meeting and the decisions made therein and to make these visible to participants in a manner which written minutes often occlude. Centrally, the aim is to use the tools developed within Memetic to realize a shared understanding of meetings' work and to map the dimensions of issues raised therein.

Many teams now use video conferencing as an indispensable part of their daily work. It is fair to say that the 'teething troubles' of early internet video conferencing have been largely resolved by improved bandwidth and video conferencing technologies such as the Access Grid (although even state of the art systems can be perceived at times as unstable, so sensitive are people to disruptions in face-to-face discussion).

In contrast, relatively little progress has been made on delivering robust, accessible tools for creating and flexibly navigating records of videoconferences. Whilst not considered useful or even desirable in some contexts (e.g., for reasons of privacy, litigation, intrusiveness, etc.), in the many situations where it would be useful, there is a need for functionality that goes beyond simply replaying/skimming a digital movie.

This paper introduces a toolkit for transforming the normally ephemeral character of video conferencing interactions into persistent records which can be navigated in linear and non-linear ways, and which, as interactions spanning multiple meetings, can be traced and manipulated.

The contributions which this paper seeks to make are twofold. First, with respect to meeting tools, we will introduce two complementary interaction visualizations: argumentation-based concept maps to elucidate the conceptual structure of the discourse using a particular interactional language, and multiple event timelines generated from the meeting metadata. These become semantic navigation indices into a digital video rendered in an advanced form of 'movie player' called Meeting Replay, the design of which is discussed. Second, with respect to methodology and evaluation, we describe how we are engaging diverse end-user communities in the process of designing and deploying these tools.

Section 2 gives an overview of the tools we are developing and these are described in more detail in Section 3. In Section 4, we discuss our methodology for designing, developing and refining the tools. In Section 5, we review related work. We conclude by considering in Section 6 the lessons we have learnt to date and how this work might be developed in the future.

---

2. Overview of Approach

The goal of the Memetic project is to create a toolkit for meeting support. Specifically, this involves:

- Developing tools to record all or selected video streams in a meeting;
- Investigating the scope for *automatically* indexing the video stream timeline with potentially significant events (such as slide changes, visits to websites, progression through agenda items, or changes in speaker), whilst *manually* indexing other significant events which are too complex for automated system detection (such as the raising of arguments and making of decisions).

The critical bottleneck in creating such accounts of meetings is, of course, the overhead of manual indexing. However, the effort required is minimized in Memetic through the use of a concept mapping approach to capturing discussions, mediated via a hypermedia tool called Compendium. While requiring someone to map ideas, these maps are displayed for all to see and validate, and so justify the effort by adding value to the meeting by virtue of the shared focus of attention to the unfolding of decisions that they bring.

Our approach to creating the Memetic toolkit is to take a number of existing tools for meeting support and to refine and integrate them. We now introduce the Memetic tools.

2.1. Access Grid Video Conferencing

The Access Grid (AG) is an open collaboration and resource management architecture for video conferencing based on the metaphor of persistent virtual venues.
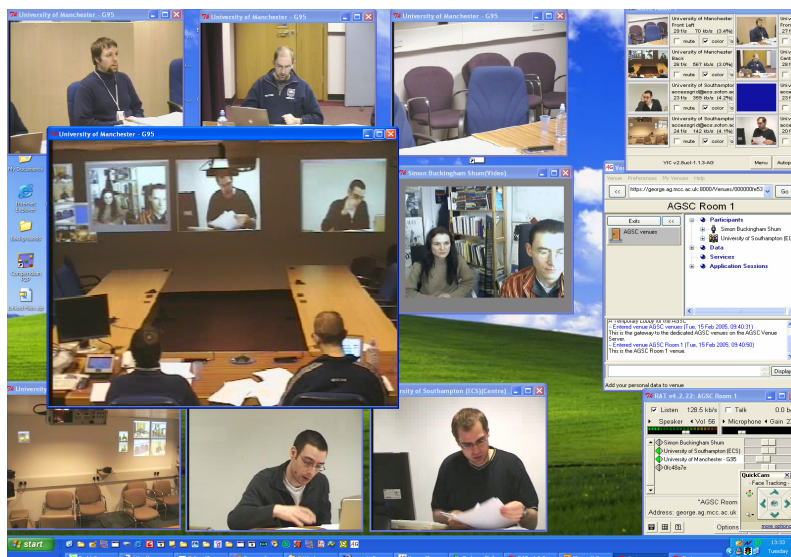

Figure 1. Participating in an Access Grid (AG) videoconference from a personal computer. The enlarged central video window shows participants at a 'venue' in a full AG 'Node' (see text).

A team of researchers collaborating in, for example, a laboratory would expect to find there a set of tools available to help their work; so in a virtual venue, as well as video and audio feeds of all participants, applications and services to aid a specific virtual organization to work together remotely can also be accessed. The philosophy underlying AG is that each group of collaborators has their own virtual venue in which they can store shared objects such as documents and data, together with shared applications, perhaps to aid access to a physical resource such as a radio telescope or electron microscope.

An AG meeting can be attended via a single personal computer (Figure 1) or by going to a full AG 'Node', a designed space with very large display, multiple video cameras and high quality audio system. An intermediate solution is to have a desk-based AG node, consisting of three aligned PC displays and echo-canceling microphones, which provides a richer environment than can be achieved using a single PC desktop with webcam and microphone headset.

AG supports the recording of meetings that can be played and stopped as digital video streams. Our task in the Memetic project is to implement and evaluate extensions to this replay by improving the video replay functionality, and indexing it using Compendium and the Meeting Replay tool.

2.2. Compendium

Compendium is hypermedia software tool for authoring and publishing issue-based Dialogue Maps: concept networks that structure Issues, Ideas and Arguments in a discussion, linked as required to supporting and background multimedia documents and internet resources. Compendium is best thought of as a knowledge management environment for supporting personal/group deliberations and memory, combining hypermedia, modelling and mapping [5].

Compendium uses an approach called Conversational Modeling [10], which extends the technique developed by Conklin [7] termed Dialogue Mapping, which in turn derives from the formative public policy planning work of Horst Rittel [9]. Rittel characterized the concept of "wicked problems", which can only be solved by all stakeholders striving to define the problem and being willing to explore issues dialogically, in what he termed argumentative design. He proposed the Issue-Based Information System (IBIS) as a notation to mediate discourse. Software-supported Dialogue Mapping has been under development since the late 1980s, resulting in a large body of knowledge about the craft, process skills [7], and since 1993, the skills of Conversational Modeling are being analyzed (e.g., [10]). We return to Compendium's affordances as a tool for recording interactions within meetings below.

Figure 2 shows an extract from a Dialogue Map created over several meetings, both face-to-face and virtual. As part of a large scale NASA Mars exploration field trial [11], which deployed and evaluated the tools now being extended in the current project, co-located field geologists at a desert site (a Mars simulation) arranged rock sample photos for analysis. Colleagues (simulating a remote science team back on Earth) reviewed this on the internet and raised queries, linking them into the map as new 🔵 Questions, 🟡 Ideas and ➕/➖ Arguments. The Mars crew then responded (highlighted nodes). In other maps, discussions include links to voice annotations and web datasets. Compendium provided a shared visual focus on the contributions as they were made (particularly useful in the absence of other shared visual referents in virtual meetings),

and a group-validated memory of how contributions connected. The Dialogue Map became the group's evolving, shared picture of their problem .
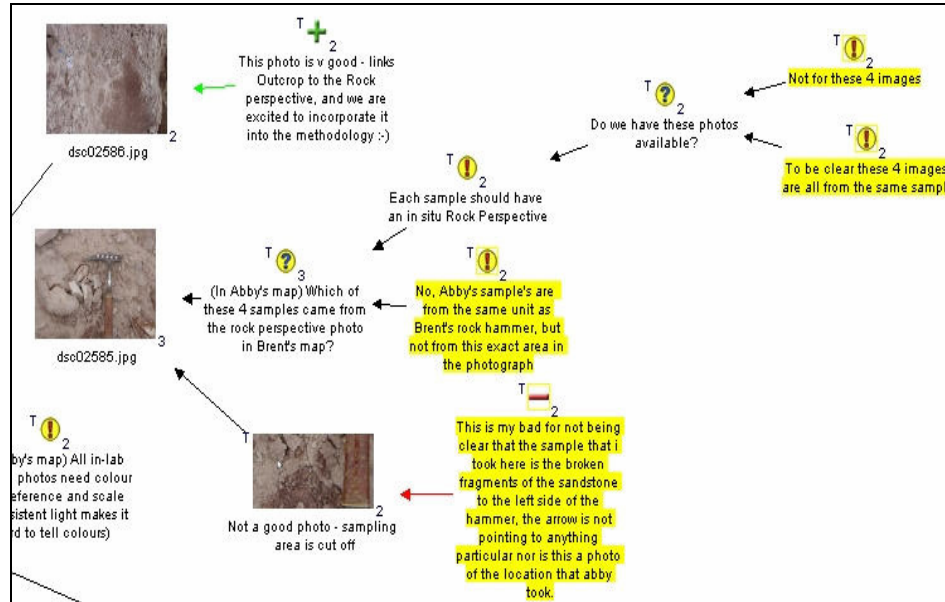

Figure 2. A Dialogue Map created in the Compendium software tool, illustrating its capabilities for integrating media resources with analysis and argumentation from different stakeholders (in both co-present and virtual meetings).

The content of maps may be driven entirely by issues raised by participants, or at the other extreme, discussion can be driven by working through predefined Issue Templates that specify the issues to be tackled, and possibly the options available and the criteria by which they should be judged. The approach can be particularly powerful by blending freeform and predefined maps. In all of the above cases, maps are created by people as an aid to thinking. However, maps can also be automatically generated and read by a software agent (e.g., [11]). The maps then provide hypertext functionality for navigating and linking data elements, and can be combined with any of the above modes of use.

As a semantic, visual hypertext system, Compendium provides several ways to manage the connections between ideas: drawing optionally labeled graphical links between nodes (connections in a given context); transclusion (tracking occurrence of the same node across different contexts); metadata tagging (enabling harvesting of nodes with common attributes across different contexts); and catalogues (managing libraries of nodes and template structures). See [5] for details.

Several significantly-sized case studies have documented the value of rendering real-time interactions as visual maps, whether co-present or online meetings [5] [12][13]. The approach has also been used to model and interpret the key issues and arguments in an extended, asynchronous discourse, rendering a corpus of documents around a controversy such as the Iraq debate as interactive IBIS maps on the Web [14].

2.3. Meeting Replay

The Meeting Replay interface (Figure 3) integrates the Access Grid videos, Compendium database and other indices into the meeting. The meeting can be navigated via the interactive event timelines shown in bottom frame, or from any node in a Compendium client (e.g., to play the video at the point when a particular argument was made). The richness of the video record thus compensates for the terseness of the Compendium maps; in turn, the maps and event timelines provide hyperlinked indices into the video.
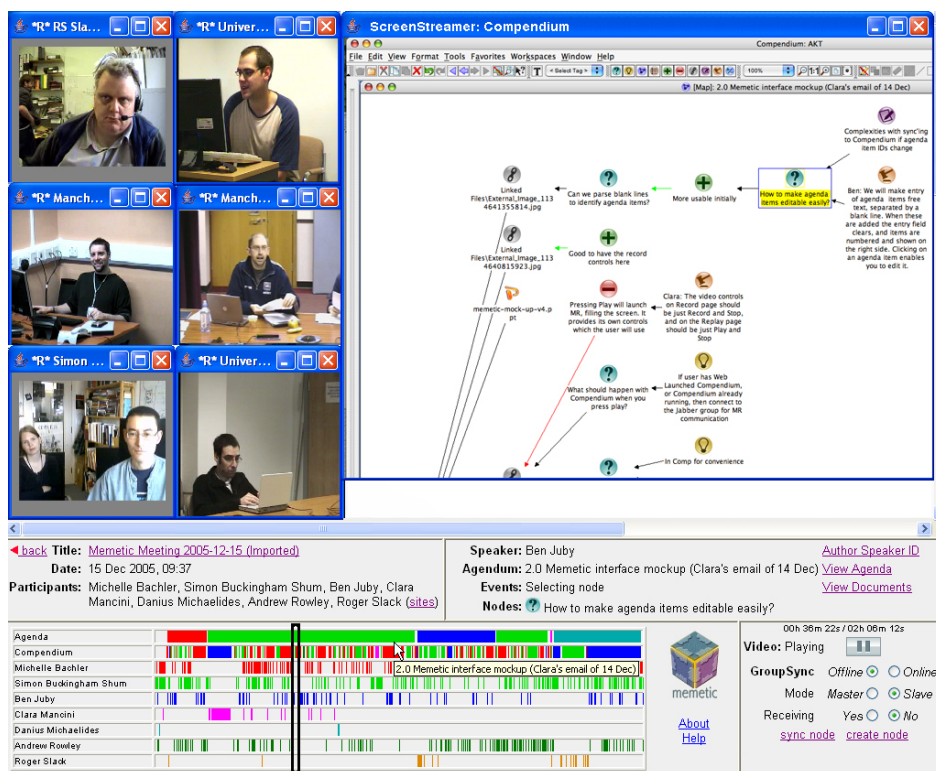

Figure 3. Example of the Meeting Replay web interface

All of the above are integrated within, and launched from, a Meeting Manager website which enables users to book, record, replay and annotate meetings. We have not detailed the technical implementation of the tools, since this is not of primary interest for this paper. Our concern is to focus on the affordances of these tools as a specific form of manipulable record of meeting interactions.

3. Affordances of Memetic tools

What do we know about the affordances of the tools as a resource for recovering and working with past discourse?

3.1. Access Grid video replay

Prior to the Memetic project, the video streams from an AG meeting could be simply played or stopped. Our work improves on this firstly, with a tool called Arena, which adds conventional video player controls to pause, fast forward and rewind. This is further augmented by another new tool called ScreenStreamer which adds a video stream from participants who choose to share the video output from their computer screens. The affordances of video replay are as one would expect of this familiar metaphor: participants know that what they say and do is being recorded, and adjust their behavior accordingly. If they want their presentations to be recorded then they are no longer restricted to AG-aware software, but can opportunistically share anything via ScreenStreamer. With Arena alone, however, when seeking to replay a point in the meeting users must recall or guess the offset from the start time, or simply browse to find the desired moment. Arena has subsequently been embedded in the Meeting Replay tool's upper frame, which as discussed below, expands the medium's affordances considerably.

3.2. Compendium Maps

The key affordances of such maps are summarized in Figure 4. Notational affordances are to some extent intrinsic to the approach, available to anyone who is 'literate' in reading Compendium maps (e.g., Figure 2) and interpreting the cues highlighted in Figure 4 below. But the way in which the notation is used in an interaction can determine its success. The key affordances of the Dialogue Mapping/Conversational Modeling dynamic can be summarized as follows.
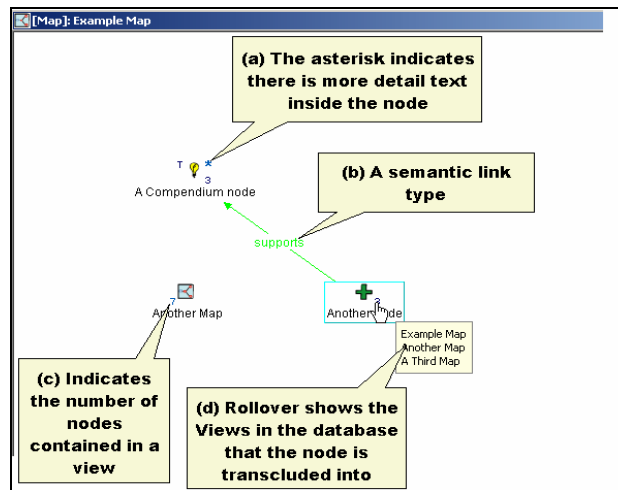

Figure 4. 'Reading' a Compendium discussion map.

The Compendium mapper functions as "technographer", actively crafting structures on a shared display screen that both capture the meanings and ideas of the group and reflect back to it the larger implications of their thinking [16].

3.3. Advantages of Compendium Mapping

Based on the studies cited earlier, and over a decade's experience with the approach in real organizational contexts, we are able to identify a number of affordances of the Compendium's issue-based, highly visual, real time mapping in meetings, physical and virtual (for details see Conklin [7]):

- *Beneficial slowing down.* A complaint sometimes heard when argument/discourse mapping is first introduced to meetings, is that it disrupts the flow of the meeting. When done appropriately, however, we find that it can be extremely beneficial to focus attention on a feature of the hypertext map: potential sources of conflict in the ongoing interaction can be topicalized and conflict can be defused. After a while, people become noticeably unhappy when their contributions are not mapped, but if these are captured on screen, they know that their view has been heard, correctly recorded, and will be harder to ignore when the map is assessed at decision time.

- *A visual trace of the analysis.* The graphical IBIS notation has a number of properties which can 'talk back' [15] to participants to show them when the quality of an analysis can be strengthened. These cues can be learnt implicitly through use (rather like learning how to format a document through experience), or explicitly through training. A workshop can teach participants how to raise effective issues, spot weak or irrelevant positions, and question implicit assumptions. While these are skills offered by a good meeting chair or facilitator, newcomers can use the structure of the IBIS map as a cognitive scaffold to develop these skills.

- *Depersonalization of conflict.* When ideas and concerns are mediated via a shared display, challenges to positions assume a more neutral, less personal tone. We cannot of course claim that this technique resolves all conflicts, but in ill-structured situations where there are competing agendas, it helps participants to clarify the nature of their disagreement (e.g., the definition of 'the problem'; understanding different criteria of 'success'). We have seen Compendium defuse meetings which otherwise looked to be polarized, and bring together parties who were unable to work together.

- *Flexible rhythmic review.* To a surprising degree, collaborative knowledge work can be characterized as 'group list processing'. Whether the list is a set of requirements, budget items, or action items, a common activity is group review of a list of potentially complex elements. While some items draw little comment, others can lead into deep discussions and even debate. With Compendium, the facilitator and the group establish a 'call and response' rhythm during these exercises, creating a sense of shared purpose and momentum. When occasional elements lead to intense discussions about meaning, or spark disagreement among group members, the Compendium practitioner can open a new map and keep facilitating, mapping or modeling the new interaction. With the new issues captured in the shared display, the group can return to the previous review task without losing momentum.

- *Incremental mediation of interaction.* We are discovering a variety of strategies for introducing Compendium to a new group, strongly determined by the context. Some of us simply start to capture the normal discussion in a meeting and, at appropriate points, use it to reflect back to the meeting

personal insights gleaned. Curiosity about what one is doing often leads to natural opportunities to introduce the tool. If the shared display is used from the start, again, there is a spectrum of how strongly discourse is mediated via this display (cf. the continuum in Buckingham Shum and Hammond, [17]). It may be used to punctuate discussion to reflect on progress, but at its most powerful, the discussion and the map shape each other.

It may be argued that the need to be skilled in the use of Compendium is a fundamental weakness of the approach: someone has to become fluent with it, before beginning to take advantage of its affordances. In contrast, our view is that like any other tool that takes time to learn, Compendium yields greater benefits with practice. Mapping ideas in IBIS during a meeting is unquestionably an acquired ability, but one which can be learnt (there is an international Compendium user community, and project members use it to map meetings). We discuss the important issue of cognitive overhead, and strategies for overcoming it, in [5], while other work is beginning to articulate the skills displayed by experienced knowledge mappers [10].

3.4. Meeting Replay

Whereas Compendium is a relatively mature approach to interaction capture and visualization, semantically indexed Meeting Replay of videoconferences is new territory. However, it is possible to begin articulating the affordances of this representation for 'reading' an interaction from a videoconference.

3.4.1. Event Stream Visualization

The key additions to Arena's video replay functionality that the Meeting Replay tool brings are:

- *interactive event timelines* provide a visual index to get an overview of the video, and navigate around it by clicking on an event (Figure 5);
- *integration with Compendium* enables the user to click on a node in a Compendium map to navigate the Meeting Replay to the point in the meeting just before that node was created.
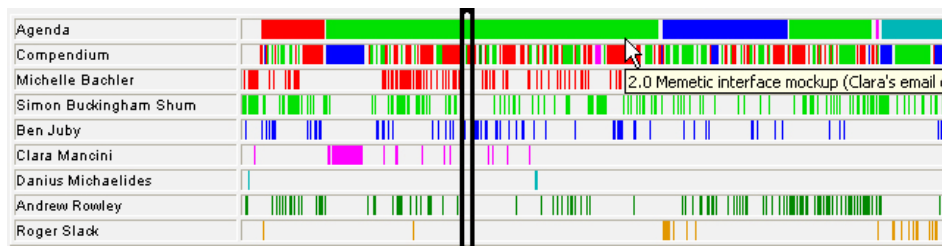

Figure 5. Meeting Replay's interactive event timelines.

The key question is what kinds of information this display can, or should, afford. Taking firstly the event timeline visualization, the following information not normally accessible from a video can be read from the display:

- When an agenda item was discussed (e.g. Figure 5 shows that the second item, in green, was returned to after item 3). Details for a given event are displayed on a mouse rollover.
- Who spoke when, and about which agenda items
- Who spoke a little or a lot.
- Who was speaking when a given Compendium node was created, highlighted, tagged, or a hyperlink followed to an external application or website; this node might be an Issue, Idea or Argument, or a Reference node to an external document such as a spreadsheet, website, photo or slide (see Figure 2).
- What the distribution of Compendium node types is (again, they are color coded by type).
- Combining the above, for instance, one can see at a glance which agenda items or Compendium nodes provoked a lot of discussion, amongst whom, and with an approximate indication of whether there was much argumentation (presence or Pro, Con and Argument nodes)

There is an additional cue provided during navigation around the replay, e.g., on clicking a timeline event or Compendium node. When one jumps into the middle of a video recording, for someone who was not at the meeting or who has forgotten the details, there is an orientation phase while one establishes the context. The Meeting Replay interface offers the user an often reliable cue to the context by displaying the 'current' (most recently created) Compendium node (e.g. see Figure 3: *"How to make agenda items editable easily?"* is displayed as the current node, and can be highlighted in the Compendium map on request, as shown). Thus, although a slide or photo may be on the main display as a context cue to the subject of discussion, the user is also provided with a cue to the particular Issue, Idea, or Argument. The user is, however, left to disambiguate whether this motivated inspection of the slide/photo, arose from it, or indeed, pertains to another topic altogether.

Finally, Meeting Replays can be further annotated in Compendium by anyone in the project, to add missing material that might be useful, or to construct completely new navigational maps around the video, an affordance that we are now investigating to support distributed video data analysis.

3.4.2. Navigating Interactions Spanning Multiple Meetings

An affordance that we have yet to implement in Meeting Replay, but which we are beginning to consider, is navigation of interactions spanning multiple meetings. This is already possible in Compendium, whereby maps from discussions going back years can be retrieved (based on keyword, date, node type, author or metadata), pasted into a current discussion, or even actively cued by the interface by providing auto-completions of a new node's label as the user types it, based on matches to existing nodes (which might come from years back). Once these nodes are linked into a Meeting Replay archive, it will be possible for the user to select from multiple possible Meeting Replays in which a given node has arisen. Similarly, a search on the Compendium database will in effect be a search across multiple videoconferences.

Finally, to represent the contents of meetings, the Meeting Replay semantic web architecture uses an RDF triplestore, from which the Meeting Replay interface is generated. This opens up further possibilities for reasoning over multiple interactions and providing meeting memory services that mine, or act upon, the memory traces.

4. End-User Participatory Design

Meetings are central to the conduct of organizations often with the participation of persons over a number of sites via video conferencing. However, in part because meetings have become part of the fabric of everyday life, there is a sense in which their organization is, to use Garfinkel's [19] term, 'seen but unnoticed'. This obviously has implications for the design of a set of technologies and tools that support records of meetings. The work of participants in making meetings run is often not explicitly formulated or examined but is, nevertheless, a vital resource in the design of technologies for recording interactions. How, then, to get at this stock of practical methods for making meetings work in order to use it as a resource for the design of tools to enhance and possibly transform them?

Following the pioneering work of the Participatory Design (PD) community (Greenbaum and Kyng [20]), the involvement of end-users in the design of technologies and tools has become accepted practice within the world of Information Technology. There are numerous variations on PD and the one that we have followed in the Memetic project goes by the name of 'co-realization'. The aim of co-realization is to develop technologies and tools in co-operation with those who will use them, and to do so over time. Following Trigg, Blomberg, and Suchman [21], co-realization strives to create a situation where "... co-development of CSCW [Computer Supported Collaborative Work] technologies ... means more than engaging prospective users in the design of new computer systems to support their work. It requires that we as designers engage in the unfolding performance of their work as well, co-developing a complex alignment among organizational concerns, unfolding trajectories of action, and new technological possibilities."

Co-realization's orientation to design and development conceptualizes design as a co-operative and situated practice involving end-users and designers of technology as equal partners (Hartswood et al. [22]; Buscher et al [23]). Design and development work is grounded in the lived experience of end-users as they come to use technologies and to appropriate their functionalities and affordances into their work practices and relations.

4.1. Methodology

Our methodological approach takes a twin track: first, workshops with end-users and developers; second, a series of site visits to observe meetings 'in the wild', linked with an ongoing commitment to observe meetings over time via Access Grid and to discuss issues arising from these meetings with developers and end-users. The initial workshop meetings were fora for developers and end-users to interact, discuss the potential of Memetic technologies, and understand how these might be deployed within each organization in order to afford work. Each technology within the Memetic project was demonstrated to users and a workshop discussion allowed users and developers of each technology to discuss the ways that the technologies might a) be used within end-user organizations; b) the potential for enhancements based on site-specific experiences and needs.

The aims of the user requirements workshops were as follows:
a) To provide familiarity with the potential of the technologies and their uses;
b) Predicated on (a), to develop on the day an emergent 'wish list';
c) Discussion of this 'wish list' in plenary sessions

d) A series of 'core' requirements based on (b) and (c);
e) A series of action items based on (d)[4] .

The second track of our approach involves a series of site visits. These involve Memetic researchers in observing the conduct of meetings via Access Grid. The value of these observations lies in the familiarization of the researchers with the setting and the issues that users within each organization face – both in terms of getting meetings organized and in using Memetic technologies within the fabric of these meetings. Site visits also allow the end-users to develop what we might think of as in vivo requirements, i.e., requirements that occur during the meeting and which may have not been envisaged in the workshop.

An example is useful here. During one site visit, an Access Grid meeting involved the use of a shared PowerPoint presentation; collaborators at a remote site were unable to navigate consistently through the presentation and had to be prompted by the presenting site. Presenters interrupted their presentations with utterances such as "could you click onto the next slide [centre name], please" and "click again [site name] yes, that's it". Such interruptions also led to the partner centre moving too far ahead and having to track back over slides with consequent disruption of the narrative of the presentation. After the meeting, one attendee who had been at our user requirements workshop commented to the Memetic representatives that he had "really thought that, yes, there's a place where they could use Screen Streamer, but I couldn't get [remote centre] to set it up. Next time I'll suggest that to them". Such in vivo realizations as to the ways that Memetic technologies could be used are invaluable in both designing systems that afford the work within the user organization and enable buy-in to the project.

In line with the aims of co-realization, our aim is to build upon these partnerships over time as the project unfolds.

5. Related Work

There is a considerable amount of research on capturing physical meetings, to offer a meeting record, but not much on Internet videoconferences (which ironically, are easier to capture since everything is digitally mediated). Some of these are investigating advanced technologies that are not yet widely available within the mainstream environments we are seeking to augment. Speech recognition technologies are being trialed by Pallotta et al [24] who are also experimenting with IBIS as the basis for a schema to model the output of speech recognition tools. Face recognition cameras are prototyped by Cutler et al [25] who then generate event timelines. Abowd et al [26] and Moran et al [27] have demonstrated how activity around digital whiteboards can be indexed and replayed.

The key difference of our work, is that we are working with an infrastructure that is widely available and supported, and our use of real-time IBIS conceptual mapping provides a notation for capturing a wide variety of discourse moves, and an open ended way to generate meaningful event timelines for different user communities, simply by defining different combinations of nodes and metadata as indices into the video.

---

[4] We used Compendium to analyze the materials from the workshops and thereby to capture requirements. This also functioned as a demonstration of the utility of Compendium for participants.

5.1. Emergent Research Questions

The tools are about to deployed with end-user partners. Ultimately, our question is to what extent participants report that the tools improve their work (following a five-month evaluation phase). Specific research questions can be summarized as follows:

- How far does the technology support naturally occurring interactions: where can it be seen to augment these and where does it just get in the way?
- What is the relationship between naturally occurring references to the technology made by end-users during meetings, and their more 'official' accounts when asked by us as researchers (we are helped here because all meeting interactions are recorded)?
- Which parts of the toolkit that we provide will be used, and in what (possibly unexpected) ways?
- What requirements emerge for 'meeting memory services'? Memetic models semantics in discussions (Compendium maps), within the RDF store (People, Agenda items, Meetings, Venues, etc), and within Meeting Replay (Events, Documents, People, etc); these provide a platform for experimenting with services to support projects.

6. Conclusions

In this paper we have described the toolkit which we are developing and integrating: video of participants and shared presentations, hypermedia Dialogue Maps, and interactive event timelines. Each of these has specific affordances for navigating and 'reading' off information about what happened in the meeting. Future papers will report the results of the participatory design engagements with our end-user partners, and the results of the deployment and evaluation of the tools in use.

Acknowledgements

References

1. Boden, D. (1994). *The Business of Talk.* Cambridge: Polity.
2. Bush, V. (1945) "As We Might Think" *Atlantic Monthly* 176 (1), 101-108
3. Englebart, D. and English, W. K. (1968). "Research Center for Augmenting Human Intellect" *Proc. Fall Joint Computing Conference*, San Francisco, CA. AFIPS: Reston, VA; pp. 395-410.
4. Stefik, M., Foster, G., Bobrow, D. G., Khan, K., Lanning, S. and Suchman, L. (1987). "Beyond the Chalkboard: Computer Support for Collaboration and Problem Solving in Meetings. *Communications of the ACM* 30 (1), 32-47.
5. Simon J. Buckingham Shum, Albert M. Selvin, Maarten Sierhuis, Jeffrey Conklin, Charles B. Haley, Bashar Nuseibeh. (2006) To appear in: *Rationale Management in Software Engineering,* (Eds.) Allen H. Dutoit, Raymond McCall, Ivan Mistrik, and Barbara Paech. Springer-Verlag.

6. Selvin, A. and Sierhuis, M. (1999). Argumentation in Different CSCA Project Types. In *Workshop on Computer-Supported Collaborative Argumentation*, CSCL Conf., Stanford, CA (12-15 Dec).
7. Conklin J (2005) *Dialogue Mapping*. Wiley, Chichester
8. Selvin, A. Supporting Collaborative Analysis and Design with Hypertext Functionality. *Jnl. of Digital Information*, 1, 4, 1999 [http://jodi.ecs.soton.ac.uk/Articles/v01/i04/Selvin/].
9. Rittel, H., and M. Webber. (1973). Dilemmas in a General Theory of Planning., *Policy Sciences*, Vol. 4, , pp 155-169. Elsevier Scientific Publishing Company, Inc., Amsterdam.
10. Selvin, A. (2005) Aesthetic and Ethical Implications of Participatory Hypermedia Practice. *TR KMI-05-17, Knowledge Media Insitute, Open University*: http://kmi.open.ac.uk/publications/pdf/kmi-05-17.pdf
11. Clancey, W.J., Sierhuis, M., Alena, R., Berrios, D., Dowding, J., Graham, J.S., Tyree, K.S., Hirsh, R.L., Garry, W.B., Semple, A., Buckingham Shum, S.J., Shadbolt, N. and Rupert, S. (2005). Automating CapCom Using Mobile Agents and Robotic Assistants. *1st AIAA Space Exploration Conference*, 31 Jan-1 Feb, 2005, Orlando, FL. AKT IRC ePrint 375: http://eprints.aktors.org/375
12. Selvin, A M and Buckingham Shum, S J (2002) Rapid Knowledge Construction: A Case Study in Corporate Contingency Planning Using Collaborative Hypermedia. *Journal of Knowledge and Process Management* 9(2):pp. 119-128. http://eprints.aktors.org/48
13. Vedro, S. (2003) Using World Modeling Interviews to Develop Lists of Shared Problems, Opportunities, Know-how and Assets Shared by Wisconsin Public Television and the U.W. System Learning Innovations Center. *Technical Report, Reforging the Links Project*. http://www.reforginglinks.uwex.edu/CIMIssues.html
14. Buckingham Shum, S. and Okada, A. (2005) Modelling the Iraq Debate: Mapping Argumentation in a Document Corpus. *Proc. Compendium Institute Workshop*, 10-11 Nov, 2005, Washington DC http://www.kmi.open.ac.uk/projects/compendium/iraq/
15. Schon, D. (1983) *The Reflective Practitioner*. New York: Basic Books.
16. DeKoven, B. *Connected Executives.* The Institute for Better Meetings, 1990
17. Buckingham Shum, S. and Hammond, N. Argumentation-Based Design Rationale: What Use at What Cost? *International Journal of Human-Computer Studies*, 40, 4, 1994, 603-652
18. Selvin, A. and Buckingham Shum, S. (2005). Hypermedia as a Productivity Tool for Doctoral Research. *New Review of Hypermedia and Multimedia* (Special Issue on Scholarly Hypermedia), 11 (1), 91-101 PrePrint: http://kmi.open.ac.uk/publications/pdf/kmi-tr-05-8.pdf
19. Garfinkel, H. (1967). *Studies in Ethnomethodology*. Englewood Cliffs, NJ: Prentice-Hall.
20. Greenbaum, J. and Kyng, M. (Eds.) (1991). *Design at work: Cooperative design of computer systems*. Hillsdale, NJ: Lawrence Erlbaum Associates.
21. Trigg, R., Blomberg, J. and Suchman, L. Moving document collections online: The evolution of a shared repository. In *Proceedings of ECSCW'99* (Copenhagen Denmark, 1999) Dordrecht. Kluwer
22. Hartswood, M., Procter, R., Rouchy, P., Rouncefield, M, Slack, R. and Voss, A. (2002). Co-realisation: Towards a Principled Synthesis of Ethnomethodology and Participatory Design. In M. Berg, D. Henriksen, J. Pors and B. Winthereik (Eds.), special issue on Challenging Practice: Reflections on the Appropriateness of Fieldwork as Research Method in Information Systems Research, *Scandinavian Journal of Information Systems*, 14(2), p. 9-30.
23. Buscher, M., Hartswood, M., Mogensen, P., Procter, R, Shapiro, D., Slack, R. and Voss, A. (2002). Promises, Premises and Risks: Sharing Responsibilities, Working Up Trust and Sustaining Commitment in Participatory Design Projects. In T. Binder, J. Gregory, and I. Wagner, (Eds.) *Proceedings of the Participatory Design Conference*, Malmo, June 23rd-26th, p. 183-92.
24. Pallotta V., Ghorbel H., Ballim A., Lisowska A., & Marchand-Maillet S. (2004) - Towards Meeting Information Systems: Meeting Knowledge Management. *Proceedings of ICEIS '04 (6th International Conference on Enterprise Information Systems)*, Porto, Portugal, April 14-17, p.464-469.
25. Ross Cutler, Yong Rui, Anoop Gupta, JJ Cadiz, Ivan Tashev, Li-wei He, Alex Colburn, Zhengyou Zhang, Zicheng Liu, Steve Silverberg. Distributed Meetings: A Meeting Capture and Broadcasting System, *Proc. ACM Multimedia*, 2002
26. Abowd, G., Chris G. Atkeson, Jason A. Brotherton, Tommy Enqvist, Paul Gulley, and Johan Lemon, Investigating the capture, integration and access problem of ubiquitous computing in an educational setting. In the *Proceedings of CHI '98*, pp. 440-447, May, 1998
27. Moran, T., Palen, L., Harrison, S., Chiu, P., Kimber, D., Minneman, S., van Melle, W., Zellweger, W. (1997). "I'll get that off the audio": a case study of salvaging multimedia meeting records. *Proc. CHI Conference on Human Factors in Computing Systems*, Atlanta, 202 – 209.