

Memetic: Semantic Meeting Memory

Danius Michaelides¹, Simon Buckingham Shum², Ben Juby¹, Clara Mancini², Roger Slack³,
Michelle Bachler², Rob Proctor³, Michael Daw⁴, Andrew Rowley⁴,
Tim Chown¹, David De Roure¹, Terry Hewitt⁴

¹Electronics and Computer Science,
University of Southampton, UK

²Knowledge Media Insititute,
The Open University, UK

³School of Informatics,
University of Edinburgh, UK

⁴Manchester Computing,
University of Manchester, UK

Abstract

This paper introduces the Memetic toolkit for recording meetings held over Internet-based video conferencing technologies, and making these navigable in linear and non-linear ways. We introduce the tools and technologies that form the toolkit and discuss the semantics of the information they capture.

1 Introduction

How do people remember meetings? Consider the following stories, typical of the ways we recount meeting events to others:

- Sam arrived late, at least 15 minutes in I think. By then Kim had already got agreement for her trip
- Anne announced she was quitting her job 5 minutes before we had to move back into the main session
- Joe's slide triggered negative vibes from the network guys
- It was a shame Lin had to go just before Sandra's killer demo
- We looked at that spreadsheet several times, and no-one complained once
- All Steve did was disagree with everything the marketing people suggested
- Millie was in full flow and then guess what her connection went down
- Bill is a lousy chairman we only covered half the agenda and even then we jumped all over the place

- There was a massive debate for about 20 minutes when everyone pitched in
- We made this decision quite early on, even though Jo said towards the end that it was doomed if we didn't take on board her reports findings
- The mood completely changed when Daisy arrived

How can we provide support for recovering such critical incidents from a recorded meeting? How would we record it, and how would we model and index it? This paper reports on the Memetic project, which has developed an integrated toolkit for meeting capture supporting real time and post hoc annotation. This generates a "semantic video player" supporting hypertextual navigation around a video replay, and semantic search across a meeting database. Section 2 gives an overview of the Technologies in the Memetic toolkit. Section 3 looks at the approach we have taken to model the semantics of discourse. Section 4 explores some of the facilities we intend to explore with the toolkit.

2 Technologies

In this section we review the tools and technologies in the Memetic toolkit and look at how they integrate at both a system level and information level.

2.1 Meeting Manager

The Meeting Manager is the main entry point for using the Memetic environment. It is a web-based application that provides users with a front-end to the meeting data, allowing them to create, edit and review meetings. Users would typically engage in the following lifecycle:

book meeting Users enter data about the meeting such as when, who will be attending, agenda items, associated documents etc. Much of this information can be authored later, but the primary (internal) effect of booking the meeting is that the meeting is known to the system by a generated URI.

record meeting At the start of the meeting, the user instructs the Meeting Manager to record the meeting. Tools such as Compendium and ScreenStreamer can be launched, associating themselves with the meeting. Access Grid video conferencing tools are also launched.

refine meeting data At any point information about the meeting can be added or edited. Post-meeting is a good time to update participants and documents used during the meeting. Compendium maps can be uploaded.

review meeting The meeting can be reviewed with Meeting Replay. Compendium can be used to map the meeting post-hoc. Annotations can be associated with the meeting (e.g. speaker identification).

2.2 Access Grid Tools

The Access Grid[5, 1] is an open collaboration and resource management architecture for video conferencing. It is characterised by being able to support large scale distributed meetings and uses a virtual venue metaphor to provide persistent meeting points in which users can collaborate and share documents and data.

Meetings are usually held in Access Grid “Nodes”, which are designed spaces with large display areas, multiple cameras and effective room-based audio systems. Users can also attend sessions from their desktop, using a single PC based solution. Laboratories, specialist equipment and output from complex visualization systems are often included in sessions to provide a rich environment for researchers to collaborate in.

Memetic has developed two additional tools for use within Access Grid sessions. Arena is a server component that provides support for recording and playback of video from Access Grid sessions. ScreenStreamer is a tool for sharing computer screens that integrates into the Access Grid framework. The distinction between ScreenStreamer and other screen sharing technologies is that it uses the same network protocol (RTP) as the other Access Grid streams. As a result it can also be recorded by Arena

2.3 Compendium

Compendium is a hypermedia 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 resources on the Internet. Compendium is best thought of as a knowledge management environment for supporting personal/group deliberations and memory, combining hypermedia, modelling and mapping[4].

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).

Several significantly-sized case studies have documented the value of rendering real-time interactions of visual maps, whether co-present or online meetings [4, 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.

2.4 Meeting Replay

The Meeting Replay interface (Figure 1) is a web-based interface which integrates the meeting metadata, video and indices into a single view. The upper half of the interface shows the video streams of the meeting. As there may be any number of different streams, the user is free to select which streams are most relevant and arrange them accordingly. In this case, six video streams are shown alongside a larger view of the Compendium screen capture which was created with ScreenStreamer.

The lower half of the interface consists of three panes presenting data about the meeting, and a fourth control pane. The north-west pane comprises of general information about the meeting and is entirely static. The north-east pane, conversely, is a dynamic view of information relevant to the current instant in time; current speaker, current agenda item and current activity in Compendium.

The south-west pane shows the complete meeting in a timeline view with a number of lines depicting specific types of annotations. A slider, which indicates the current time in the meeting, can be dragged to navigate to a specific point. Each line in the timeline groups a number segments which depict the annotations. Each segment is clickable (causing the replay to jump to the start time of the annotation) and has a tooltip to provide further details of the annotation. Segments are colour-coded based on a scheme which is specific to the type of annotation. Colour schemes typically indicate different instances of an annotation. In the Agenda line, for example, each item in the agenda has

a specific colour, allowing the user to differentiate between the different topics of the meeting. The colour scheme for the Compendium annotations is slightly different in that it indicates the type of the Compendium node involved in the annotation. For example, this allows the user to quickly pick out all the Compendium action items created during the meeting.

This visual index of the meeting provides a powerful mechanism for users to quickly navigate around meeting. With time being the major dimension, the structure and other aspects of the meeting become readily apparent: the duration of discussion for each agenda item; who spoke a little or a lot; which agenda items promoted discussion; what was the nature of the discussions (from the types Compendium nodes, e.g. lots of exploratory discussion vs lots of decision making). It also highlights that annotations have varying levels of detail and varying levels of information structure.

2.5 Systems Architecture

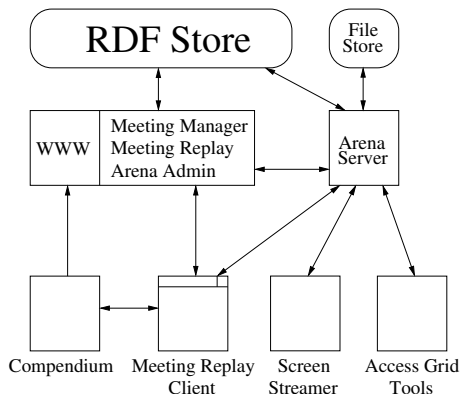


Figure 2. Architecture of Memetic components

Figure 2 shows the architecture of the Memetic system. The lower four tools are client side tools Compendium, Meeting Replay, ScreenStreamer and the standard Access Grid tools. The server components consist of a Web server, the Area server, and an RDF store. The Web server provides web-based interfaces for the Meeting Manager, Meeting Replay and an Arena administration facility. These are implemented as JSP pages. The Arena server provides the record and playback of media streams over the RTP network protocol, using RTSP as the control mechanism. The RDF store is Jena based with the Derby backend and is used through to store all meeting metadata.

Meeting replays are generated dynamically from the RDF triples in the datastore, by taking the meeting URI as

a start point and performing a number of RDQL queries. These results are used to expand an HTML template that forms the skeleton of the Meeting Replay interface. The dynamic panes in the interface then load their data via Javascript encoded pages, generated from servlets making RDQL queries. The video streams are displayed by a Java applet which communicates with the Arena server. The rest of the interface is a combination of HTML and Javascript. An additional applet is used to provide synchronization amongst a number of Meeting Replay clients, via the Jabber communications protocol. It also enables synchronization with Compendium: i) users can click on a node in Compendium and Meeting Replay will jump to the relevant point in the meeting and ii) users can map a discussion in Compendium as they replay the meeting.

2.6 Ontologies

All the meeting metadata is specified in RDF and uses a number of layered ontologies: the Memetic ontology, the CoAKTiNG ontology[7], AKT's portal and support ontologies[2, 3] and Dublin Core. These are described in OWL.

AKT's portal and support ontologies model academic life, defining concepts such as people, publications and papers. Whilst it has the concept of a meeting, CoAKTiNG developed this further to include support for agendas, multiple sites and time-based annotations (in order to model speaking events and Compendium activity). This was further refined in Memetic, and extended to model the more complex media streams and management aspects found with a move to Access Grid and wider deployment.

3 Modelling the semantics of discourse

The Compendium tool is specifically designed to enable someone in a meeting to map contributions in real time, organized around a simple ontology called IBIS (Issue-Based Information System)[12]. Compendium comes pre-loaded with node and link types for IBIS, for connecting key issues, possible responses to these, and relevant arguments.

Figure 3 shows a design rationale extract from a project meeting, in which an issue is raised, two options explored, and one justified. Figure 4 shows the use of Compendium simply to record decisions (about metadata). While these might simply have been recorded in a word processor or slide tool, such tools do not support (i) the possibility of capturing important discussion/rationale if it arises, or (ii) the reuse of a decision in subsequent other contexts see the links on the bottom node to its other appearances in the database. Users can also define their own custom modeling language, by building their own palettes of icons (called

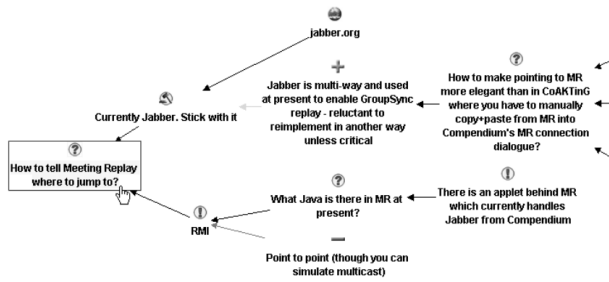


Figure 3. Extract from a software design meeting, in which Compendium is used to map issues, options, arguments, the decision, and a relevant website. (This meeting was an Internet video conference, with Compendium viewed by participants via a desktop sharing application.)

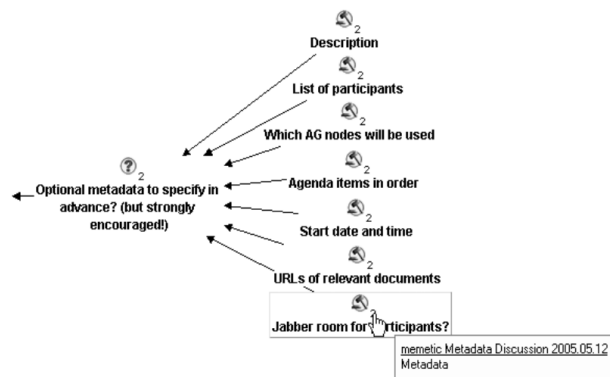


Figure 4. Recording decisions (in this case without any significant rationale) in Compendium. Rolling the mouse over the digit on a node displays a link menu to other maps which contain the node.

Stencils) and relational types (Linksets). This is not currently a full meta-modeling tool, however, in that constraints cannot be specified between nodes and links: any two nodes can be linked using any linktype.

Compendium maps are not flat drawings, but views onto a relational database that can be rendered in multiple formats. A given node (e.g. representing an idea, argument, entity, or document) can appear and be updated in multiple views. Since any application document or URL can be dragged and dropped into a map as a Reference node, so an external document can be linked into one or more discussions and tracked that is, given one or more meaningful contexts where it plays a role. Corrections or updates to a node are immediately updated in every context in which it appears. This provides precisely the representational capability needed to build semi-structured models in which a particular object is systematically reused (e.g. an idea, plan, person, system, location).

4 Future Work

With a basic infrastructure to capture and replay meetings in place, there are many enhancements to the Memetic tools that can be explored.

End users may wish to improve the knowledge and understanding of the meeting by making their own annotations. These annotations could be private or public, and of standard types or domain-specific. Simple tagging mechanisms can be provided within Meeting Replay, mimicking in part the behaviour already provided by Compendium. Enabling the use of domain-specific annotations, however, presents issues of how to allow users to integrate their annotations into the Memetic system, such as providing an

ontology for their annotations, how to represent them in the Meeting replay and the design of associated search mechanisms.

With a corpus of meetings at a user's disposal, effective search mechanisms are vital. Compendium's own search facility allows tracking repeated instances of an issue across its maps, and hence all meetings. Such a search would result in a set of nodes, transcluded from the original meetings, from which Meeting Replay can be launched at the relevant point in the meeting(s). Similar search mechanisms could be provided within Meeting Manager.

Services could be built to provide users with information post-meeting, such as sending them a reminder of their action items prior to the next meeting.

The existing Meeting Replays provide a complete record of meetings held, which may well be too much for users to digest post meeting. Traditional meeting minutes already demonstrate the desire for condensed, edited summaries of the meeting. An open research question would be whether automatically generated highlights can be created by applying inference rules to the wealth of data accumulated about a meeting. Summaries may take the form of visualizations that are more focused on a specific attribute or property than Meeting Replay's timeline.

Other inferences could be drawn, such as building up a profile of a person's expertise/interests based on who was speaking when a specific subject occurred in a meeting. The subject of a meeting can be inferred from the creation of Compendium nodes, or other activity such as the viewing of documents or the following of web links. The types of Compendium nodes created may also influence the profile - if someone seems to consistently be speaking as an answer to a given question, then they presumably know something

about the question and idea nodes' labels/details.

The combination of speaker annotations and Compendium activity could also be used to characterise the nature of the interaction

5 Related Work

There have been numerous attempts to support meeting activity with technology, such as Bush's Memex with its "associative indexing" of texts and artifacts; Engelbart's NLS/AUGMENT, which enabled navigation through "complex information structures" and conceptual mapping. Other more recent systems such as NoteLook [6] and Distributed Meetings [9] have provided support for capturing and indexing meetings, but have very limited semantics associated with the data limiting the amount of automated processing and the types of queries that may be performed on that data.

Our ontology combined with IBIS has been used in more basic research which seeks to extract meeting events from unstructured audio and video records. Pallotta et al [10] compare the work of the IM2 and Calo projects, both of which model discourse moves using IBIS (plus other argumentation schemes), and in one case, also use the CoAKTinG ontology that we developed prior to Memetic. This is ambitious, long term work which if successful, has the advantage of not requiring the explicit mapping of IBIS structures in a meeting. However, IBIS structures inferred from naturalistic dialogue will always be of poorer quality than human-constructed representations. Moreover, the fundamental basis of the Compendium Dialogue Mapping approach is that the construction of the map is not merely to serve as an index to support possible replay or rationale recovery, but adds value to the meeting as it proceeds (see Conklin[8], who has documented in detail the principles and hands-on information mapping skills by which this operates).

Other research is investigating ways in which to model discourse moves in meetings, and also draws on our work. Rienks et al [11] report on a representational scheme which combines our work on the use of IBIS for mapping meetings, together with other argumentation models, to derive an argument diagram of a meeting transcript. However, they do not describe how this would support a practical meeting memory tool such as that described here.

6 Conclusions

In this paper we have described the ongoing development of the Memetic toolkit for recording and playing back Access Grid meetings. We believe our approach and choice of tools enables us to create a rich source of semantic information about the contents of meeting and the collaboration

between participants. Our playback environment demonstrates how this information can be used to create a novel interface that enables users to navigate to quickly to the points in meetings that interest them. We exploit the system further by providing search mechanisms, visualization and other services that are useful to meeting participants.

Acknowledgments

The Memetic project is funded by the UK's Joint Information System Committee's Virtual Research Environments programme.

References

- [1] Access Grid website. <http://www.accessgrid.org/>.
- [2] AKT Project. AKT Portal Ontology. <http://www.aktors.org/ontology/portal>.
- [3] AKT Project. AKT Support Ontology. <http://www.aktors.org/ontology/support>.
- [4] S. J. Buckingham Shum, A. M. Selvin, M. Sierhuis, J. Conklin, C. B. Haley, and B. Nuseibeh. Hypermedia Support for Argumentation-Based Rationale: 15 Years on from gIBIS and QOC. In A. H. Dutoit, R. McCall, I. Mistrik, and B. Paech., editors, *Rationale Management in Software Engineering*. Springer-Verlag, 2006.
- [5] L. Childers, T. Disz, R. Olson, M. E. Papka, R. Stevens, and T. Udeshi. Access grid: Immersive group-to-group collaborative visualization. In *4th International Immersive Projection Technology Workshop*, July 2000.
- [6] P. Chiu, A. Kapuskar, S. Reitzeier, and L. Wilcox. Meeting capture in a media enriched conference room. In *Second International Workshop on Cooperative Buildings (CoBuild'99)*, volume 1670 of *Lecture Notes in Computer Science*, pages 79–88. Springer-Verlag, 1999.
- [7] CoAKTinG Project. CoAKTinG Meeting Ontology. <http://www.aktors.org/coacting/resources/ontology/meeting>.
- [8] J. Conklin. *Dialogue Mapping*. Wiley, Chichester.
- [9] R. Cutler, Y. Rui, A. Gupta, J. Cadiz, I. Tashev, L. He, A. Colburn, Z. Zhang, Z. Liu, and S. Silverberg. Distributed meetings: A meeting capture and broadcasting system. In *10th ACM Conference on Multimedia*, 2002.
- [10] V. Pallotta, J. Niekrasz, and M. Purver. Collaborative and argumentative models of meeting discussions. In *Workshop on Computational Models of Natural Argument, IJCAI*, Edinburgh, 2005.
- [11] R. Rienks, D. Heylen, and E. van der Weijden. Argument diagramming of meeting conversations. In *Seventh International Conference on Multimodal Interfaces*, Trento, Italy, October 2005.
- [12] H. Rittel and M. Webber. Dilemmas in a general theory of planning. *Policy Sciences*, 4:155–169, 1973.
- [13] A. M. Selvin and S. J. Buckingham Shum. Rapid knowledge construction: A case study in corporate contingency planning using collaborative hypermedia. *Journal of Knowledge and Process Management*, 9(2):119–128, 2002.

The screenshot displays the Meeting Replay web interface. On the left, a grid of six video windows shows participants: *R* RS Sla..., *R* Univer..., *R* Manch..., *R* Manch..., *R* Simon..., and *R* Univer... The main window, titled 'ScreenStreamer: Compendium', shows a diagram titled '[Map]: 2.0 Memetic interface mockup (Clara's email of 14 Dec)'. The diagram consists of several nodes connected by arrows, with annotations such as 'Can we parse blank lines to identify agenda items?', 'More usable initially', 'How to make agenda items editable easily?', 'Ben: We will make entry of agenda items free text, separated by a blank line...', 'Clara: The video controls on Record page should be just Record and Stop, and on the Replay page should be just Play and Stop', and 'What should happen with Compendium when you press play?'. Below the diagram is a control panel with fields for 'back Title: Memetic Meeting 2005-12-15 (Imported)', 'Date: 15 Dec 2005, 09:37', 'Participants: Michelle Bachler, Simon Buckingham Shum, Ben Juby, Clara Mancini, Danius Michaelides, Andrew Rowley, Roger Slack (sites)', 'Speaker: Ben Juby', 'Agendum: 2.0 Memetic interface mockup (Clara's email of 14 Dec)', 'Events: Selecting node', and 'Nodes: ? How to make agenda items editable easily?'. At the bottom left is an agenda table, and at the bottom right are video controls including 'Video: Playing', 'GroupSync' (Offline/Online), 'Mode' (Master/Slave), and 'Receiving' (Yes/No).

Agenda	Color
Compendium	Red
Michelle Bachler	Green
Simon Buckingham Shum	Blue
Ben Juby	Yellow
Clara Mancini	Purple
Danius Michaelides	Orange
Andrew Rowley	Light Blue
Roger Slack	Light Green

Figure 1. The Meeting Replay web interface