Graphing Meeting Records - An Approach to Visualize Information in a Multi Meeting Context

Bettina Kirchner [1], Jan Wojdziak [2], Mirko de Almeida Madeira Clemente [2], Hannes Leitner [2], Rainer Groh [2]
[1] GTV – Gesellschaft für Technische Visualistik mbH
[2] Technische Universität Dresden, Department of Computer Science, Chair of Media Design

Structured Abstract

Purpose – Meeting notes are effective records for participants and a source of information for members who were unable to attend. They act as a reference point to decisions made, to plan next steps, and to identify and track action items. Despite the need for a multi meeting solution (Tucker and Whittaker, 05), meetings are often displayed as separated as well as descriptive documents. The aim of this work is to enhance access to overlapping meeting contents and existing coherences beyond a decoupled description. A visual representation of meeting content can lead to meeting records which are more comprehensible and more time efficient. Furthermore, it enables the depiction of knowledge that is often lost in conventional meeting records.

Approach – Our goal was to define a general structure for meeting items, integrating content categories and relations between successive meetings. In this paper, we present a model based approach to visualize meeting content as well as content relations in order to support the preparation, execution and follow-up of meetings. Due to the fact that contents of consecutive meetings refer to each other (Post et al., 04), we consider meetings as a series of events. The resulting model substantiates the transformation of content as well as content relations into a visual form.

Value – The proposed solution focuses on the model that is integrated into an interactive visualization. Thus, a novel approach to explore meeting records is provided. The model was proved to be suitable for meeting contents in various use cases. Examining the content in its visual representation across multiple consecutive meetings enhances the identification of any linked information at a glance over even long periods of time. Hence, important pieces of information will not be disregarded.

Practical implications – The approach of our multi meeting protocol application is realized as a browser-based implementation that displays data from JSON objects. With this interactive visualization, the user can browse, search, and filter meeting content and get a deeper understanding of topics, their life cycle and relations to other topics. This leads to an overall comprehension of project or business progression that

highlights topics that need to be addressed. Thus, the viewer is supported in preparing, executing, and following up meetings successfully and qualified to structure records in order to keep a clean transcript of a meeting.

Keywords – Multi Meeting, Interactive Visualization, Meeting Records, Meeting Content Relations, Topic Evolution, Meeting Model

Paper type – Practical Paper

1 Introduction

Different methods and tools exist in the field of automated meeting capturing and analysis in order to reduce loss of knowledge (Geyer et al., 05; Yu and Nakamura, 10) or structuring the captured contents for easier retrieval (Richter et al., 01). Their objective is to support carrying out tasks on a day to day basis or help making important business decisions. Conventional meeting records often exist in the form of single text documents. They are most likely stored in individual file systems as separate files, one for each meeting, without their contents explicitly or visibly referencing contents from preceding meetings. This makes it difficult for anyone to locate specific information over the course of time – even more so, when meeting participants do not know or remember the respective meetings. In the following, they do not know which files to go through to find objects of conversation or decision.

Surprisingly, only few systems provide access to contents of successive meetings. This stands in contradiction to the fact that meetings refer to each other. They should not be regarded as single, but as a series of events and thus should be viewed jointly (Post et al., 04). Tucker and Whittaker state that available meeting browsers[1] deliver a solution for this problem by providing search mechanisms for a set of meeting recordings. Yet, with the query leading to one or more single meetings that match the search criteria, the records are viewed separately again (Tucker and Whittaker, 05). Even if search mechanisms solve the problem of identifying a single relevant meeting from a series of meetings in parts, it is not yet satisfying. As a consequence, Tucker and Whittaker further express the need for a system that is able to examine decisions and other content items across a series of meetings. They predict that a "multi-meeting analysis tool" would be of great value to meeting participants. Such a system should be able to integrate multiple meetings in one view instead of displaying single meetings separately.

1 Meeting browsers are a form of user interface that supports meeting information retrieval. They can be part of a larger smart meeting environment (Yu and Nakamura, 10).

We intend to redesign the process of retrieving meeting information following these requirements by providing an interactive visualization for meeting contents and their relations across a series of meetings. This paper introduces a protocol-browser that displays multiple meetings and their coherences simultaneously, giving the viewer a powerful tool for analysing, exploring and searching information in its entirety.

State of the art meeting browsers[1] and visualization techniques for text documents are described in section 2. Subsequently, section 3 introduces a model for meeting contents and relations. It factors in user studies that we considered to be relevant for the development. In chapter 4, the concept for an interactive visualization is presented that is based on the model and the previous research. A description of benefits, followed by an outlook on future work and a summary in chapter 7, concludes the paper.

2 Related Work

We reviewed related work in three different areas: Meeting Browsers for meeting content analysis and as a state-of-the-art overview, a topic evolution model for content relation analysis, and text mining visualization techniques.

2.1 Meeting browsers

Beyond conventional meeting records in the form of text documents, meeting browsers are able to store multimodal data. This includes, for example, textual transcripts from recorded conversation, with additional information such as the speaker of a certain piece of information. On top of that, the meeting browsers “Archivus” (Lisowska et al., 05) and “TeamSpace” (Richter et al., 01) provide access to a series of meetings. Both tools store consecutive meetings and implement search mechanisms for a set of recordings as described in section 1. The tools display one meeting record at a time without explicit reference to preceding or successive meetings. However, their overall handling and classification of meeting content provides a solid basis for structuring meeting contents and the derivation of relevant interactional necessities.

2.2 Topic evolution

A wide range of tools compute topic evolution from large time based text collections using mathematical algorithms and probabilistic models (Blei and Lafferty, 06; Liu et al., 12; Gohr et al., 09; Cui et al., 11). On top of that, the text mining tool “TextFlow” (Cui et al., 11) includes a visualization for the gathered data. After extracting topics from a set of documents, their life cycles are traced and their changes are displayed along a timeline. Cui et al. further give an insight into the dynamics of the topic evolution theory. They define four major events in the “life” of a topic: topic birth, topic death, topic merging and topic splitting. The latter events concern more than

one topic at a time. Thus, relations between different pieces of time based content can be established. Their solution for illustrating the topics is based on the stream graph[2] visualization. In addition, they use glyphs to represent the events (topic birth, death, split and merge), integrating them into the stream graph at the point of time they occur.

2.3 Visualizing time-based text collections

Besides TextFlow, "Event River" (Luo et al., 12) and "StoryFlow" (Rose et al., 09) present two alternative visualization techniques that focus on topic evolution in time based text collections. They show topic relations and topic development in different ways, with the joint objective to support information retrieval and understanding. In Event River, topics emerge from real life events. They are represented as coloured bubbles. Their size correlates to their importance at the time and their colour links them to other bubbles with the same colour, forming a story in their entirety. In addition to the visual representation, the authors provide interaction mechanisms to get an overview as well as detailed information. StoryFlow is simpler in its presentation and without interactive elements. Yet, the tool covers another aspect of topic evolution that is not immanent in Event River: topic splitting and merging. This feature includes the ability to draw a link between different topics by use of plain lines, regardless of their affiliation to different stories.

3 Meeting content and relational structure

In order to create an interactive visualization for several successive meetings, the initial step was to define a general model for meeting contents and their relations across multiple meetings. In the following subsections, we describe the underlying model of the multi-meeting-protocol-browser. It consists of a content structure as well as a relational structure.

3.1 Meeting content structure

As part of the long term evaluation of the TeamSpace meeting browser, Lipford and Abowd conducted a user study. Among other things, it revealed that most meetings share a common order of events and comparable goals (Lipford and Abowd, 08). At the same time, they discovered that the contents are similar in their structure for a wide range of meetings. The results of the study further suggest that a meeting is identified primarily by the date it took place on, and that the topics discussed in a meeting are the main indicators for structuring its content. This observation is of great value to our approach, as it enables applying only one model to different kinds of meeting situations and contexts.

2 Another example for a stream graph visualization is ThemeRiver (Havre et al., 00).

Another relevant aspect for building a model was to categorize content items. Therefore, types of meeting content were considered. The review of four studies and concepts (Brandl et al., 10; Ispas et al., 10; Khan, 93; Lisowska et al., 04) that vary in purpose, but hold answers to this question, helped us to identify the following categories as most common regarding different kinds of meetings: tasks, ideas, decisions, dates and general information. The last mentioned category covers all content items that do not fit any of the other categories. In summary, the first part of the meeting structure is developed from common meetings that were subdivided into topics. Each topic contains one or more content items of different categories.

3.2 Content relations

As a next step, a specification for links between those topics and content items is needed. For this purpose, the previously described topic evolution theory (Cui et al., 11) and the general meeting cycle were considered. The latter evolves around the fact that meetings follow an order of preparation, execution, recording and processing of their contents in a possibly infinite loop (Jain et al., 03). This implies the occurrence of dependencies between meetings. The fact that the content of consecutive meetings does relate has already been established by Cook et al. With “Project Nick”, they developed a system that supports the execution of meetings by technical means (Cook et al., 87). Part of their system is based on the idea that information enters a meeting through different channels, with the prior meeting and its recordings being one of them. A piece of information that emerges in one meeting, or information is derived from it later on, can be part of a successive meeting. This means that former information influences current and future pieces of content. Therefore, the theory of content items splitting up or merging with one another as described by Cui et al. (Cui et al., 11) can be sustained.

In order to define specific kinds of relations, we reviewed Marchand-Maillets model for meeting recordings (Marchand-Maillet, 03) and a user study conducted by Lisowska et al. preliminary to the development of the Archivus meeting browser (Lisowska et al., 03). In accordance to our prior definition of a general meeting structure, we extracted the following types of relations:

- topic is part of multiple successive meetings (r1)
- information is mentioned in multiple meetings (r2)
- content item is linked to a topic and the other way around (r3)
- content item is related to another single content item (r4)
- topic leads to a different topic (r5).

Relations (r1) and (r2) differ from the others, as they describe an ongoing link for continuous pieces of information, while relations (r3) to (r5) lead to a different topic or content item, whereby a new link is defined. The complete model including meeting structure and content relations across a series of meetings is shown in Figure 1.

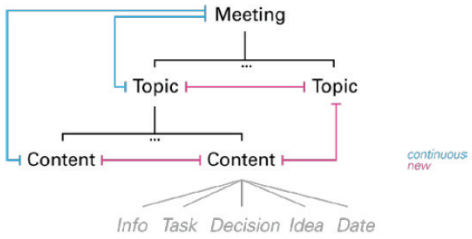

Figure 1: The derived meeting content and relation model, with the abstract relations leading across successive meetings

4 Interactive visualization

Based on the meeting model, the transformation of meeting contents as well as their relations into a visual form is required to transfer knowledge effectively and efficiently. A time and topic based interactive visualization has been developed to ensure traceability. The key factors that build the foundation for the concept emanate from the described model and the reflections of the preceding chapters. The realization as an interactive visualization is explained in the following subsections.

4.1 Visual representation of multiple meeting records

As pointed out earlier, the most important indicator when looking up information from a series of meetings is the date on which the meeting was held (Lipford and Abowd, 08). Therefore, we decided to use a time based visualization for the integrated view of multiple meetings. We arranged the meetings of interest along a horizontal timeline that depicts the respective dates. Since the basic structure of a single meeting is determined by its topics, we clustered the individual content items by the topic they belong to. The blocks that are formed using this approach are stacked vertically below the regarding meeting date at the timeline and tagged with a label that displays the topic's identifier. The content item categories are represented by a different color for each category. The content item is visualized as a rectangular instance within its topic and colored accordingly. The height of the rectangle refers to the amount of contained information (without weighing in its importance or priority). This visual representation of meeting content enables the viewers to identify separate topics as well as the kind and amount of information at a glance. At the same time, users are not restricted to view only one meeting at a time, but enabled to scroll through the whole duration of a project, business related conferences, scientific reviews or any other field of work that includes a set of conversational recordings.

Existing relations between topics and content items are illustrated across meetings as path elements, leading from one item to another. Narrow lines link different topics or content, while broad grey or colored paths illustrate continuous information. The endings of the relations that represent new links convey information as well. It is possible that the linked item is not displayed in the current view and therefore the viewer does not know which type of content or topic the item leads to or is derived from. Due to this, we added colored elements to both ends of the links: circles for a content item and straight lines when the link is connected to a topic. In case of a content item leading to a topic, for example, the beginning of the line would be a grey circle (grey for the topic), while the end would be a straight horizontal line in the color of the content's category (Figure 2).

4.2 Interactive elements

Beyond the overview the visualization provides, the user can gather information from comprehensive meetings by using simple mouse and tap interaction. Each visual element implies some form of interaction, and the visualization offers several methods to explore, search and filter the recordings for general or specific information.

First, there is the timeline that can be dragged in order to scroll back and forth. It can also be resized. This interaction changes the zoom level of the visualization in terms of narrowing or widening the time scale, bringing more meetings into the view or pulling them further apart. The depicted meeting dates themselves include interactivity as well as the topic label and the content item's rectangle. Clicking on any of these elements will show all written conversation they contain – from the transcript of the whole meeting down to a single line of text in a content item.

Lastly, the relation lines and their endings can be selected. On the one hand, a selected relation shows every piece of information that is connected to it, thereby fading out everything else. Selecting its anchor points, on the other hand, leads the user directly to its linked item, whether it is a topic or a content item. The latter interaction only works for new links, since continuous relations are not equipped with an ending due to the fact that the connected elements represent the same content. Additional filter and search components help to find specific information. Selecting a value from a list of topics focuses the selected topic in the view. A range of checkboxes highlight contents according to their category. In addition, a plain search input shows the search results by scrolling to and opening them within the visualization.

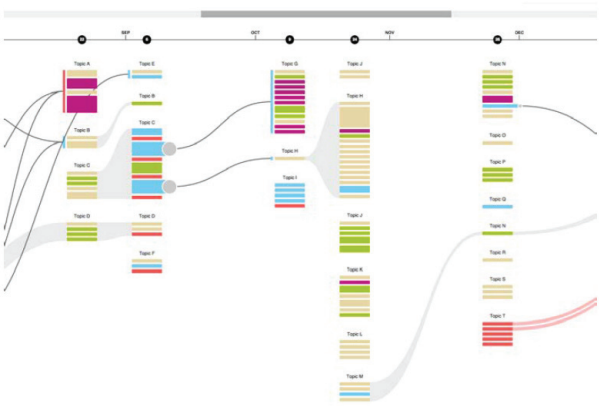

Figure 2: Visual representation of multiple meeting records

4.3 Realization

The design of a multi meeting protocol is realized as a browser-based application. HTML, CSS and JavaScript ensure the basic functionality and layout. D3.js (Data Driven Documents; Bostock, 15) is used to draw the visual elements as scalable vector graphics (SVG). The data is bound to the SVG elements via d3.js and integrated as JSON objects. For application and review of the interactive visualization, we used the transcripts of 21 project meetings that were held over the course of 20 months. After transforming them into JSON, relations and categories were added to the content. The result is shown in Figure 3. Viewing the transcripts in the multi-meeting-protocol-browser, it becomes clear that all contents fit the structure and no content-related aspect is missing in the application, compared to the textual documents. The visible relations clarify that some topics influence the project from the beginning to the last recorded meeting. This would likely have remained undiscovered without the visualization.

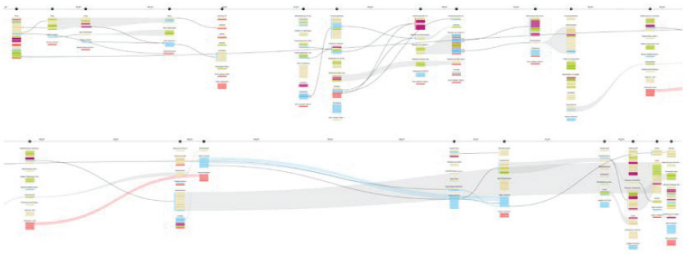

Figure 3: Meeting transcripts visualized by the multi-meeting-protocol-browser

5 Benefits

The specified requirements are met by the implementation. Most importantly, the defined structure is visually represented by the timeline, the topic blocks and the coloured content items. Relations are graspable in form of lines as well as paths and therefore cannot be missed by the viewer. This leaves a reinsuring sense of completeness, while knowing that no important piece of information will be overlooked.

In terms of information retrieval, the protocol browser supports both the need to look up specific contents by providing search and filtering functionalities, as well as the ability to get a general overview over dialogues. The user can easily learn which topics are very large or complex. By means of the interactive elements, the user can switch between overview and detail, fading out information that does not bare relevance for the task at hand or fading in a particular piece of information. The timeline transfers the content in chronological order, thus disclosing dependencies. This enables the viewer to trace the dynamic evolution of a topic for the duration of its presence – the time the topic emerges, the context it was mentioned in first, when it ceases to exist and the outcome if it leads to new topics or content items. Thereby, the visualization is able to display all of the four events in the life of a topic as defined by Cui et al. (Cui et al., 11). All in all, the multi-meeting-protocol-browser creates a deeper understanding of the information at hand and reduces the complexity when looking at a set of meeting recordings simultaneously.

By offering a general structure to the meeting participants, the transcription process can be facilitated in terms of organizing the content. As to whether the approach supports the preparation and planning of upcoming meetings, the tools does not tell the viewer which steps need to be tackled next. It does display the state of a topic implicitly, though. This is achieved by the thorough composition of contents, based on meeting transcripts and built on the presented content and relation model. The state of the topic would be defined by tasks issued, decisions made or ideas mentioned. Decisions mostly represent closed matters, whereas ideas and tasks suggest that there is work to be done and things to be discussed in the future.

6 Future Work

The paper presents a tool as a standalone solution, missing import for meeting recordings. The tool does not yet include any form of information capture – the content needs to be integrated via structured JSON. In a further step, we will develop a designated meeting protocol editor that includes tagging and linking of content and provides export and import functionality for standard text document files.

Except for our own intense utilization, the interactive visualization has not yet been evaluated widely. While the theoretical aspects of the model could be demonstrated suitable, an evaluation of the interactive elements and tangibility of the visualization needs to be conducted in order to improve and extend its functionality and presentation.

On another note, the approach to describe different types of content items using only five fixed categories is an attempt to fit the concept to a wide range of meetings. When it comes to a tool that supports individual workflows, the categories should be kept open for expansion. Therefore, our visualization is not limited to these five categories. Adding more categories would simply result in extending colours and category names, without altering the core of the model or the visualization in general.

7 Conclusions

In this paper, the authors present a novel approach to explore meeting records. The proposed solution focuses on a model for multiple meeting contents and their relations as well as its application in an interactive visualization. The model has proved to be suitable for multi meeting contents. Examining the content in its visual representation across a series of meetings, it is possible to identify any linked information at a glance, no matter how much time has passed between the items. Hence, important pieces of information will not be overlooked. Furthermore, the user can access all information available to a series of meetings, using mouse or tap interaction combined with filter and search mechanisms provided by the multi-meeting-protocol-browser. This helps accessing information quickly and changes the way meeting recordings are perceived – not as a mere collection of text files, but possibly as a "multi meeting analysis tool" (Tucker and Whittaker, 05).

Our main goal in developing this tool was to aid the preparation, execution and follow-up of consecutive meetings. At the moment, one key item is missing – an editor for meeting protocols. However, the visual representation and interactive elements enhance meeting information retrieval greatly, compared to the current usage of conventional meeting recordings.

References

Blei, David M. and Lafferty, Joohn D., (2006) "Dynamic Topic Models", Proceedings of the 23rd international Conference on Machine Learning, pp. 113–120.

Bostock, M., (2015) https://d3js.org/

Brandl, P., Richter, C. and Haller, M., (2010) "NiCEBook – Supporting Natural Note Taking", Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 10), ACM, Atlanta, Georgia, USA, pp. 599–608.

Cook, P. et al., (1987) “Project Nick: Meetings Augmentation and Analysis”, ACM Transactions on Office Information Systems (TOIS), Vol. 5, No. 2, pp. 132–146.

Cui, W. et al., (2011) “TextFlow: towards better understanding of evolving topics in text.”, IEEE transactions on visualization and computer graphics, Vol. 17, No. 12, pp. 2412–2421.

Geyer, W., Richter, H. and Abowd, G.D., (2005) “Towards a Smarter Meeting Record – Capture and Access of Meetings Revisited”, Multimedia Tools and Applications, Vol. 27, No. 3, pp. 393–410.

Gohr, André et al., (2009) “Topic Evolution in a Stream of Documents”, Proceedings of the 2009 SIAM International Conference on Data Mining, pp. 859–872.

Havre, S., Hetzler, B. and Nowell, L., (2000) “ThemeRiver: visualizing theme changes over time”, IEEE Symposium on Information Visualization, pp. 115–123.

Ispas, A., Signer, B. and Norrie, M.C. (2010) “A Study of Incidental Notetaking to Inform Digital Pen and Paper Solutions”, Proceedings of the 24th BCS Interaction Specialist Group Conference. British Computer Society, Dundee, United Kingdom, pp. 374–383.

Jain, R., Kim, P. and Li, Z. (2003) “Experiential Meeting System”, Proceedings of the 2003 ACM SIGMM workshop on Experiential telepresence - ETP ’03, ACM Press, New York, USA, pp. 1–12.

Khan, F., (1993) A Survey of Note-Taking Practices, Hewlett-Packard Laboratories.

Lipford, H.R. and Abowd, G.D., (2008) “Reviewing Meetings in TeamSpace”, Human-Computer Interaction, Vol. 23, No. 4, pp. 406–432.

Lisowska, A., (2003) “Multimodal Interface Design for the Multimodal Meeting Domain: Preliminary Indications from a Query Analysis Study”, Report IM2. MDM-11.

Lisowska, A., Popescu-belis, A. and Armstrong, S., (2004) “User Query Analysis for the Specification and Evaluation of a Dialogue Processing and Retrieval System”, Proceedings of the 4th International Conference on Language Resources and Evaluation, pp. 993–996.

Lisowska, A., Rajman, M. and Bui, T.H., (2005) Archivus: A System for Accessing the Content of Recorded Multimodal Meetings, In Machine Learning for Multimodal Interaction, pub. S. Bengio and H. Bourlard, Springer Berlin Heidelberg, pp. 291–304.

Liu, Shixia et al., (2012) “TIARA: Interactive, Topic-Based Visual Text Summarization and Analysis”, ACM Transactions on Intelligent Systems and Technology, Vol. 3, Issue 2, pp. 1–28.

Luo, D. et al., (2012) “EventRiver: visually exploring text collections with temporal references.”, IEEE transactions on visualization and computer graphics, Vol. 18, No. 1, pp. 93–105.

Marchand-Maillet, S., (2003) “Meeting Record Modelling for Enhanced Browsing”, Technical Report 03.01, Computer Vision and Multimedia Laboratory, Computing Centre, University of Geneva, Switzerland.

Post, W.M., Cremers, A.H.M. and Henkemans, O.B., (2004) “A research environment for meeting behaviour”, Proceedings of the 3rd workshop on social intelligence design, ed. A. Nijholt & T. Nishida, pp. 159–165.

Richter, H. et al., (2001) “Integrating Meeting Capture within a Collaborative Team Environment”, Ubicomp 2001: Ubiquitous Computing, Springer, pp. 123–138.

Rose, S. et al., (2009) “Describing story evolution from dynamic information streams”, IEEE Symposium on Visual Analytics Science and Technology, pp. 99–106.

Tucker, S. and Whittaker, S., (2005) “Reviewing Multimedia Meeting Records: Current Approaches”, Proceedings of the 2005 (ICMI) International Workshop on Multimodal Multiparty Meeting Processing, pp. 33–38.

Yu, Z. and Nakamura, Y., (2010) “Smart meeting systems: A Survey of the State-of-the-Art and Open Issues”, ACM Computing Surveys, Vol. 42, No. 2, pp. 1–20.