An Ego-Centric and Tangible Approach to Meeting Indexing and Browsing

Denis Lalanne, Florian Evequoz, Maurizio Rigamonti, Bruno Dumas, and Rolf Ingold

DIVA Group, Department of Informatics, University of Fribourg, Bd. de Pérolles 90
CH-1700 Fribourg, Switzerland
firstname.lastname@unifr.ch
http://diuf.unifr.ch/diva

Abstract. This article presents an ego-centric approach for indexing and browsing meetings. The method considers two concepts: meetings' data alignment with personal information to enable ego-centric browsing and live intentional annotation of meetings through tangible actions to enable ego-centric indexing. The article first motivates and introduces these concepts and further presents brief states-of-the-art of the domain of tangible user interaction, of document-centric multimedia browsing, a traditional tangible object to transport information, and of personal information management. The article then presents our approach in the context of meeting and details our methods to bridge the gap between meeting data and personal information. Finally the article reports the progress of the integration of this approach within Fribourg's meeting room.

1 Introduction

With the constant growth of information a person owns and handles, it is particularly important to find ways to support information organization as well as personalized access to information. Information is dematerializing in our daily and professional life and thus, people are often experiencing the "lost-in-infospace" effect, i.e. overloaded with information and tasks to do. Our documents are multiplying in very large file hierarchies, meetings attendance is increasing, emails are no longer organized due to lack of time, our pictures are no longer stored in photo-albums, our CDs are taking the form of mp3 files, etc. What we often miss in our daily-life and professional life are personal access to information, either tangible or digital, like used to be books in our shelves or the librarian who knew our interests.

Google and Microsoft recently tried to solve the "lost-in-infospace" issue by providing, respectively, a desktop search engine and a powerful email search engine, in order to minimize the effort made by people to organize their documents and access them later by browsing. However, in order to find a file, one has to remember a set of keywords or at least remember its "virtual" existence. If one does not remember to have a certain document, browsing could still be helpful. Browsing can reveal related keywords and documents that help you remember, since the process of browsing works by association, like our human memory does [15][20][28]. For this reason, information is generally easier to retrieve, and to "remember" if it is associated to **personal information**, either in a **digital** or **physical** form.

A. Popescu-Belis, S. Renals, and H. Bourlard (Eds.): MLMI 2007, LNCS 4892, pp. 84–95, 2007.

Meetings are central in our professional lives, not only formal meetings but also meetings at the coffee machines. Numerous recent works have tried to support the meeting process with recorders, analyzers and browsers, however most of those projects try to automate and generalize the whole process for every user, with the central long term goal to automate meeting minutes authoring [5, 6, 9, 16, 22, 26]. Moreover, all these projects try to hide technological support during the meeting, instead of exploring new ways to improve human-human communication with technology or new human-machine interaction to augment meeting rooms or enable intentional live annotations. Our claim is that replaying and browsing a meeting is personal in that it depends of each individual person's interests and that it is hard to find an agreement on what is interesting for everybody during a meeting. This claim is sustained by a recent survey we performed on 118 users [3]. This survey clearly shows that people have different needs from the meeting recordings and often personal needs, for in-stance what are the tasks they need to do or what do they need to prepare for the following meeting. Another interesting finding is that people use often their personal emails in order to access meeting information (reminders, outcomes, retrieve files).

For the above reasons, we propose in this article an ego-centric approach, complementary to other projects, which consider (a) meeting participants as potential live annotators and (b) persons who consult meetings afterwards as individuals, and thus with their own interests, derived from their personal information. The proof of concept of this tangible and personal access to information is currently assessed through two major applications:

- The control of our daily life information through the design, implementation and evaluation of tangible interfaces, called MeModules [1], allowing access, control, personalization and materialization of personal multimedia in-formation;
- The control of our professional information, meetings, emails and documents. This application, in the context of the NCCR on Interactive Multimodal Information Management (IM2), aims at recording meeting with audio/video devices, and further at analyzing the captured data in order to create indexes to retrieve interesting parts of meeting recordings.

Although these applications share common type of multimedia information, either personal or professional, the way information is accessed can not be done in a unique manner because users do not share the same roles and context of interaction in the two applications. Our assumption is that an organization of meetings centered on users could bridge the gap between these two applications and thus re-enforce user experience in meeting browsing. For this reason, this article discusses mainly the second application.

A brief state-of-the-art of tangible user interfaces is depicted in the next section. Further, the section presents a state-of-the-art on past and current document-centric browsers. The third section presents our ego-centric approach for indexing and browsing meetings, with tangible live annotations first, and secondly through the alignment with personal information structure. Finally, the last section presents our meeting room and the advancement of the integration of this ego-centric vision. The conclusion finally wrap-up the article and presents the future works.

2 State of the Art

We believe that tangible interfaces can be useful to establish links between our memory and information. They roost abstract information in the real world through tangible reminders, i.e. tiny tagged physical objects containing a link towards information sources that can be accessed by several devices. In the context of meetings, physical documents are well suited to play the role of such tangible hypermarks, that could provide entry points in the digital meeting minutes, or enable live personal annotation of meeting parts.

Brief states-of-the-art on tangible user interactions and document-centric tangible interfaces are presented in this section. Moreover, as our goal is to bridge the gap between the personal and professional information spaces by a unified ego-centric approach of both personal and professional fields, the domain of personal information management is introduced at the end of the section.

2.1 Tangible User Interfaces

Over the last couple of years, it has been demonstrated that Graspable, or Tangible User Interfaces (TUIs), a genre of human-computer interaction that uses physical objects as representations and controls for digital information [24], make up a promising alternative for the 'traditional' omnipresent graphical user interface (GUI). TUIs have also shown a high potential for supporting social interaction and collaboration. The majority of existing systems have targeted learning or playful learning, office environments, or collaborative planning tasks [25]. While quite a few systems have demonstrated the technical feasibility of associating digital media with tangible objects [12,24], these have often remained stand-alone proof of concept prototypes or applications of limited functionality.

Two existing TUIs are particularly close to our approach: MediaBlocks [23] and Phenom [12]. MediaBlocks consist of generic symbolic tangibles which act as containers for information. The system enables people to use tagged physical blocks to "copy and paste" (or move) digital media between specialized devices and computers and to physically compose and edit this content. The Phenom souvenirs are personal iconic tangibles embedded with digital ID tags to online content (an URL).

2.2 Document-Centric Tangible Interfaces

Several works explored tangible interface using printed or printable documents. The DigitalDesk [27] is centered on the co-existence of physical papers with digital information. It is built around an ordinary physical desk. It identifies and tracks documents on the table, and identifies user's pointing. Thanks to a computer-driven electronic projector above the desk, the system augments the real world with electronic objects onto the surface and onto real paper documents.

In fact, documents have properties that justify their users' acceptance. These properties are presented in this subsection, along with projects taking benefit of them. Unlike other augmented media, printed and written documents preserve their information. For instance, the RASA [17] project uses augmented maps and Post-it™ notes for military operations. In normal conditions, the system support users with speech

and gesture. In case of blackout, militaries can continue their works in the classical way. People are able to associates pages topology and indexes to document message. Books with Voice [14] allows to hear oral histories thanks to books tagged with barcodes. This project was founded on observation that historians disposing of both printed document and audiotapes preferred to retrieve information in the transcript, because it is rich in visual information such as page numbers, indexes, etc. Moreover, printed documents are not intrusive. Palette [18] for example allows man-aging slideshows presentation thanks to printed index cards, identified by RFID chips, and its evaluation showed that users concentrated exclusively on the communication task. In addition, paper is a portable interface. Take for instance the interactive paper maps for tourists that have been studied in [19] and which allow to get multi-media information by pointing on interesting areas. Finally, printed document are not exclusively associated to working or unpleasant tasks. Thus, Listen Reader [2] allows consulting a book, while the atmosphere is adapted with music in respect to the text focused by readers.

2.3 Personal Information Management

Personal information management (PIM) supports the daily life activities people perform in order to manage information of all kind. Though an old concept, it has been receiving a growing interest in the recent years due to the increasing amount of information we face, from disciplines as diverse as information science, cognitive psychology, database management or human-computer interaction [21].

Among personal information tools, two main approaches have been considered: searching versus semantic ordering. The first category contains systems such as MyLifeBits [10] or Stuff I've Seen [7]. They are improved query-based desktop search engines allowing querying all types of personal information in a unified way. Stuff I've Seen presents the query results as lists than can be sorted chronologically or by relevance. MyLifeBits also provides more flexible time-based visualizations of query results. Keys for retrieving personal information in these systems include the content of documents and other metadata than can be directly extracted (time, type of document, author). MyLifeBits also allows user-created collections of documents, opening the door to semantic annotations of personal information. Systems falling in the second category focus on the semantic aspect of personal information. Haystack [13] is a platform allowing the user to define annotations and collections over its own personal information to enable semantic browsing. Alternatively, SEMEX [29] tries to generate semantically meaningful associations automatically.

Several studies have been conducted to understand users' habits and needs in terms personal information. A particularly interesting one was conducted by Boardman [4]. Among other conclusions, he notes the preference of users for browsing over searching and the potential of integration between emails and files that often appear to have strong similarities. We leverage those conclusions in our approach of personal information management.

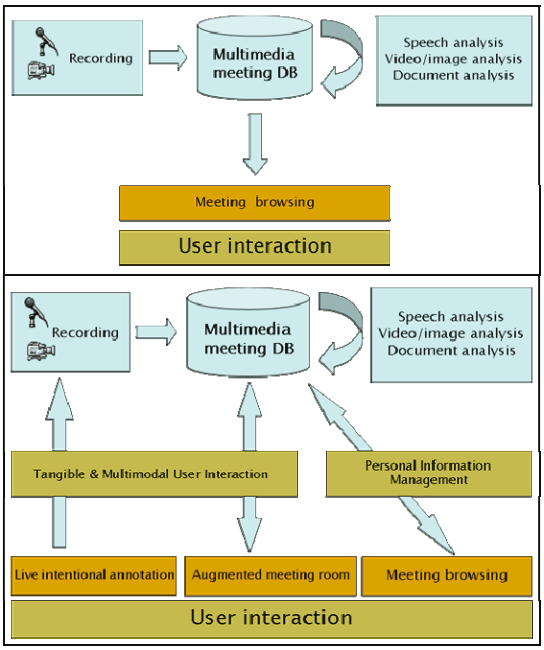

Fig. 1. Standard versus ego-centric meeting recoring, indexing and browsing

3 Ego-Centric Indexing and Browsing of Meetings

This section presents our ego-centric approach to index and browse meetings. Figure 1 represents our approach, at the applicative level, in perspective with the standard approach generally followed for recording, analyzing, indexing and browsing meetings. On the top part, the standard approach contains three sequential steps: (1) first synchronized audio/video recording of meeting; (2) in a post-production phase, analysis of multimedia meeting records is performed on raw-data to extract search-able indexes, for instance speech to text, facial expressions recognition, document analysis, etc.; (3) and finally, users can browse on meetings using the previously extracted annotations as cues to access the searched parts. In our approach, at the bottom, our goal is to enrich this process with personal annotations and with personal in-formation. These two aspects are reflected in the following two tasks of our ego-centric approach:

1. Personal live annotation of meetings using intentional tangible interaction techniques during the recording phase;
2. Browsing multimedia meeting archives via personal information structure and tangible shortcuts as specified by the user during the recording phase.

3.1 Indexing and Browsing Meetings Through Personal Tangible Anchors

Figure 2 illustrates how meetings and personal information can be intentionally linked thanks to tangible user interactions. It shows that meetings are limited in time, whereas personal information evolves during an entire person's life. Consequently, we could represent all the relationships between these applications from meeting time point of view, which is in general decomposed in pre-production, in-meeting and post-production phases.

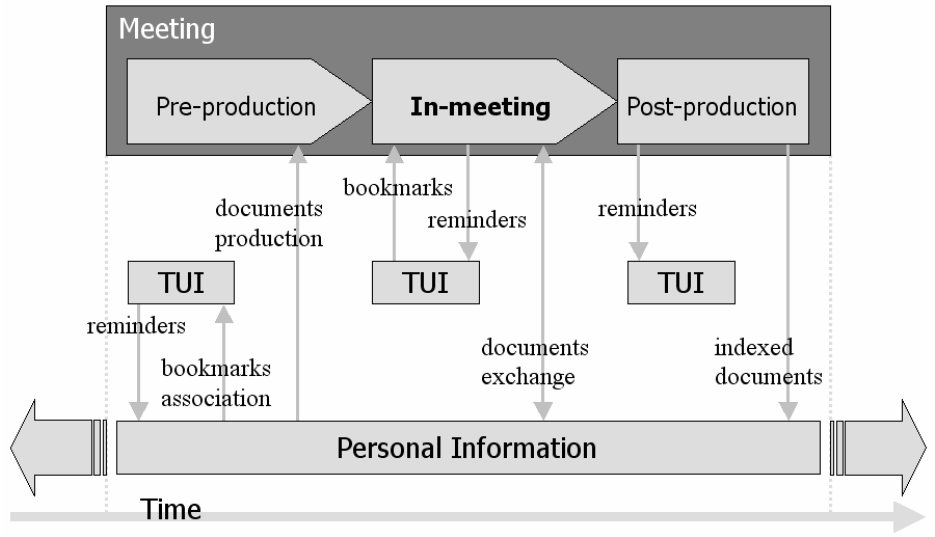

Fig. 2. The schema synthesizes the interactions binding personal information to meetings in function of time

Pre-production consists in preparing all the material the person aims to present during the meeting. In classical systems dealing with meeting, the participants prepare slideshows, printed documents, etc. At this stage, in our approach, tangible user interfaces can be used to create tangible shortcuts to the prepared documents, i.e. they allow further access to stored multimedia information during the meeting. Moreover, a participant could define a set of tangible annotators to use in the in-meeting phase, in order to bookmark interesting information according to freely defined categories that match their personal information structure.

At **in-meeting** phase, personal tangible shortcuts created in the previous step can be used to present documents to participants or share information. Furthermore, participants can bookmark some meeting part by intentionally putting an object, representing a personal category (e.g. project's name, person's name, etc.) extracted from their personal information structure, or a personal annotation (interesting, to-do, etc.), on the RFID reader. Tagged printed documents can also be exchanged between participants. This possibility has the great advantage of 1) presenting new documents

not uploaded in the pre-production phase and 2) allowing to annotate paper in a human-natural way (pen strokes, hand gesture, etc.).

The **post-production** phase includes all the analysis applied to recorded data. In our model, the participant can not only replay (the summary of) the meeting, but also access to some specific parts thanks to his personal tangible bookmarks, categories or annotations. Moreover, "category" bookmarks used during the meeting allow to automatically cluster meeting's parts, organizing the information in respect of person needs and experience.

The main contribution of the model presented in this section consists in proposing an alternative for accessing and browsing the meeting in a user-centered manner. This solution can also solve another interesting problem, which has been identified in [22] and taken into account only in [20]: to consider meetings as sequences of correlated and evolving events, instead of single and isolated experiences.

In addition, we believe that physical documents, i.e. printed documents, can help achieve the same goal than tangible personal object. TDoc is a first prototype extension of the FriDoc meeting browser [15], providing exclusively tangible interaction. By means of printed versions of the documents discussed or presented during a meeting, identified thanks to RFIDs, the user can access the vocal content recorded during that meeting, related to the paragraph he/she selected in the printed document using simple Phidgets such as sliders and IR sensors [11]. This work is a first step towards the implication of physical documents, and more generally tangible interactions, at several stages of meeting recording, analysis and browsing. Similarly to tangible objects, printed documents could indeed be engaged in three different tasks: (1) Aid "live" annotation of the meeting by providing tangible and collaborative mechanisms, like for instance voting cards (agreement/disagreement, request for a turn taking, etc.) that could be used for later browsing; (2) Serve as triggers to access services provided by the meeting room; (3) Help browsing a recorded and fully annotated meeting.

3.2 Indexing and Browsing Meetings Through Personal Information Structure

As we have recalled in section 2.3, our personal information (PI) contains valuable signs of the different roles and interests we have in life. Therefore, an abstract structure put on top of our PI, gathering pieces of meta-information about us, could be helpful in a professional context, by proposing an egocentric vision of professional document archives or events, and particularly meetings.

The extraction of an abstract structure from raw PI data is not an easy task, though. Extending the previously mentioned conclusions of Boardman's study [4], our approach is based primarily on the emails, which form a rich subset of PI, obviously linked to other pieces of PI (agenda, documents, visited websites, etc.). Indeed, a standard mailbox typically contains clues about the activities, topics of interest, friends and colleagues of the owner, as well as relevant temporal features associated to activities or relationships. Therefore, it is a rich entry point into the personal information space. Three dimensions of personal information are particularly well represented in emails: the thematic, social and temporal dimensions. We extract features pertaining to each dimension from the raw email archive. Thematic features can be extracted from the email subjects and bodies using traditional text mining methods. Social features stem from the social network built from the email archive considering

the co-occurrences of addresses in emails headers as social links between people. Temporal features consist in emails timestamps. On top of the features, we use clustering methods to help structuring the email archive according to the different dimensions.

Once a structure has been elicited from emails, the remaining personal information can be linked to it. For this purpose, multi-media information mining techniques shall be used on personal information in order to extract thematic, temporal and social features. Further, cross-media alignment techniques shall be used to link those to the elicited email archive structure. Obviously, professional documents (e.g. related to meetings) can be introduced in the process as well and thus become integrated into the whole structure.

Furthermore, the personal information structure acts as a lens for visualizing PI and browsing through it. Visual clusters, filtering mechanisms, as well as views related to different dimensions of PI can be implemented more easily thanks to the PI structure extracted from emails.

In the course of our research on personal information management, a system aiming to extract the first-level PI structure from emails is currently being implemented. The data on which it works consists of the mailbox of one individual containing nearly 6000 emails and 3500 addresses. The social and thematic dimensions of emails have already been explored to some extent: (1) a social network has been built using similarity measures between people's email addresses based on the frequency of exchanged emails between people, and this network can be visualized as a graph using the "spring" layout method; (2) exploiting the statistical similarity based on the co-occurrences of words in the subjects and contents of emails, an agglomerative hierarchical clustering has been performed, which aims at finding a thematic organization of emails. The result of this clustering has been fed into a treemap visualization. However, the study conducted so far tends to show that no dimension alone (whether thematic, social or temporal) can fully grasp the complexity of one's mailbox. Therefore, our plan is to combine and link several visualization techniques applied on each dimension to help the user browse through his personal email archive.

Plans for future works mainly include the reinforcement of the PI structure extraction from emails and the alignment of PI and professional information with this structure. As new dataset, the AMI meeting corpus, which notably includes emails, will be

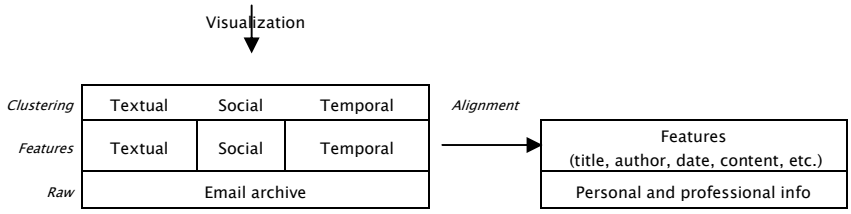

Fig. 3. (1) Extraction of the personal information structure from email, (2) alignment with the remaining personal and professional information at the feature level and finally (3) browsing in the personal and professional information through the email visual clusters

used in order to lay the foundations of an egocentric meeting browser, taking profit of the PI structure's metadata to guide the user towards the desired piece of information in meetings.

4 Towards an Ego-Centric Meeting Room

The Meeting Room in Fribourg has been built in the context of the NCCR on Interactive Multimodal Information Management (IM2) [15] as show on figure 4. This application aims at recording meeting with audio/video devices, and further at analyzing the captured data in order to create indexes to retrieve interesting parts of meeting recordings. In this context the University of Fribourg was in charge of analyzing the static documents and to align them, i.e. to synchronize them, with the other meeting media.

Roughly 30 hours of meeting have been recorded at the University of Fribourg thanks to a room equipped with 10 cameras (8 close-ups, one per participant, 2 overviews), 8 microphones, a video projector, a camera for the projection screen capture and several cameras for capturing documents on the table. Camera and microphone pairs' synchronization is guarantied and they are all plugged to a single capture PC, thanks to three IVC-4300 cards. A meeting capture application pilots the capture. It has a user-friendly interface to start, pause and stop recording, to control post-processing operations such as compression (for streaming and archiving) and to control file transfers to a server. This application is part of a more general Organizer tool for specifying the cameras and microphones to be used, the participants' position, camera's frame-rate, etc. The Organizer tool also assists users in the preparation, management and archiving of a meeting. This includes services for registering meeting participants, gathering documents and related information.

At the time of writing, RFID readers have been integrated in our meeting room, one for each meeting participant. First of all, they enable participants to register to the meeting room, thus automatically entering the metadata related to the recordings, and also identifying the tagged documents. Secondly, thanks to the synchronization of

Fig. 4. Fribourg Smart Meeting Room environment

RFID readers with the audio/video recording, users can intentionally bookmark interesting meeting instants.

In our current environment, meetings data and personal information have not been fully linked. The AMI meeting corpus, which notably includes emails, is currently used in order to lay the foundations of our egocentric meeting browser, taking profit of the PI structure's metadata to guide the user towards the desired piece of information in meetings.

Further, we are currently working on ways to augment live meetings with multimodal interaction techniques such as voice or gesture. There are three aspects we plan to handle:

- Controlling meeting room services (printing, projection, etc.): how multimodal interaction can help interacting naturally with the meeting room to project documents, exchange documents, control a slideshow, etc.;
- Annotating the meeting records live (personal bookmarks): how tangible interaction combined with multimodal interaction can enable book-marking or annotation of moments of interest in a meeting for future replay or browsing (for instance using voice to label the bookmark);
- Augmenting human/human communication and collaboration capabilities: the goal is to build an ambient room able to reflect and enhance the productivity of a meeting, for instance by projecting on the table a visualization of the dialog structure or the speakers' time of intervention.

Towards this end, a prototype allowing the control of a multimedia presentation via voice commands, gestures and tangible objects has been developed; this prototype is built upon HephaisTK, a toolkit allowing rapid prototyping of multimodal inter-faces. This toolkit is based on a multi-agents middleware, using meaning frames fed by a central blackboard as a preliminary fusion mechanism. HephaisTK allows developers to rapidly prototype multimodal human-machine dialogs using SMUIML (Synchronized Multimodal User Interaction Markup Language) [8]. HephaisTK and its SMUIML will be gradually extended and applied to handle the three applications described above.

5 Conclusion

This article presents an ego-centric and tangible approach to meeting recording, indexing and browsing. The approach proposed takes benefit (1) of the alignment of personal information with meeting archives to enable ego-centric browsing and (2) of tangible interactions during the meeting to add personal annotation in real time onto meeting data, linking meeting data with personal information.

The article presents a preliminary solution and implementation for managing personal information through emails mining and clustering, that we believe is the core of personal information. Based on this central information structure we expect to build a personal access to meetings' archives, thanks to emails/meetings alignment. The article further explains how tangible interaction, as well as printed documents, can be another way to bridge the gap between meetings and personal information, and presents in detail its involvement in pre- and post-production as well as during meetings.

The article finally presents Fribourg smart meeting room and how it implements this ego-centric vision, and the future plans to augment it with multimodal interactions.

Last but not least, the various stage of the vision presented in this article should be carefully evaluated in the future with users in order to measure not only the benefits gained following this ego-centric approach at the retrieval/browsing stage, but also how it modifies the meeting structure itself at the recording stage.

References

1. Abou Khaled, O., Lalanne, D., Ingold, R.: MEMODULES as tangible shortcuts to multimedia information. Project specification (2005)
2. Back, M., Cohen, J., Gold, R., Harrison, S., Minneman, S.: Listen Reader: An Electronically Augmented Paper-Based Book. In: proceedings of CHI 01, pp. 23–29. ACM Press, New York
3. Bertini, E., Lalanne, D.: Total Recall Survey Report., Technical Report, Fribourg (2007)
4. Boardman, R., Sasse, M.A.: stuff goes into the computer and doesn't come out: A cross-tool study of personal information management. In: CHI 2004. Proceedings of the SIGCHI conference on Human factors in computing systems, pp. 583–590. ACM Press, New York, NY, USA (2004)
5. Chiu, P., Boreczky, J., Girgensohn, A., Kimber, D.: LiteMinutes: An Internet-Based System for Multimedia Meeting Minutes. In: proceedings of the 10th international conference on World Wide Web, Hong Kong, pp. 140–149 (2001)
6. Cutler, R., Rui, Y., Gupta, A., Cadiz, J.J., Tashev, I.: He, L.-w., Colburn, A., Zhang, Z., Liu, Z., Silverberg, S.: Distributed meetings: A meeting capture and broadcasting system. In: proceedings of Multimedia 02, Juan-Les-Pins, France, 503-512 (2002)
7. Dumais, S., Cutrell, E., Cadiz, J.J., Jancke, G., Sarin, R., Robbins, D.: Stuff I've seen: A sys-tem for personal information retrieval and re-use. In: SIGIR 2003. Proceedings of the 26th annual international ACM SIGIR conference on Research and development in information retrieval, pp. 72–79. ACM Press, New York (2003)
8. Dumas, B., Lalanne, D., Ingold, R.: Creating Multimodal Interfaces with HephaisTK, a Programming Toolkit using SMUIML Markup Language. In: Poster at the 4th Joint Workshop on Multimodal Interaction and Related Machine Learning Algorithms, Brno, Czech Repub-lic (June 2007)
9. Erol, B., Li, Y.: An Overview of Technologies for E-Meeting and E-Lecture. In: proceedings of ICME 2005 (2005)
10. Gemmell, J., Bell, G., Lueder, R., Drucker, S., Wong, C.: Mylifebits: fulfilling the memex vision. In: MULTIMEDIA 2002. Proceedings of the 10th ACM international conference on Multimedia, pp. 235–238. ACM Press, New York (2002)
11. Greenberg, S., Fitchett, C.: Phidgets: Easy Development of Physical Interfaces through Physical Widgets. In: proceedings of the ACM UIST 2001, pp. 209–218. ACM Press, Orlando, USA (2001)
12. van den Hoven, E.: Graspable Cues for Everyday Recollecting, PhD Thesis, Department of Industrial Design, Eindhoven University of Technology, The Netherlands (2004), ISBN 90-386-1958-8
13. Karger, D., Bakshi, K., Huynh, D., Quan, D., Sinha, V.: Haystack: A general-purpose information management tool for end users based on semistructured data. In: CIDR, pp. 13–26 (2005)

14. Klemmer, S.R., Graham, J., Wolff, G.J., Landay, J.A.: Books with Voices: Paper Transcripts as a Tangible Interface to Oral Histories. In: proceedings of CHI 2003, pp. 89–96. ACM Press, Ft. Lauderdale, USA (2003)
15. Lalanne, D., Ingold, R., von Rotz, D., Behera, A., Mekhaldi, D., Popescu-Belis, A.: Using static documents as structured and thematic interfaces to multimedia meeting archives. In: Bengio, S., Bourlard, H. (eds.) MLMI 2004. LNCS, vol. 3361, pp. 87–100. Springer, Heidelberg (2005)
16. Lalanne, D., Lisowska, A., Bruno, E., Flynn, M., Georgescul, M., Guillemot, M., Janvier, B., Marchand-Maillet, S., Melichar, M., Moenne-Loccoz, N., Popescu-Belis, A., Rajman, M., Rigamonti, M., von Rotz, D., Wellner, P.: The IM2 Multimodal Meeting Browser Family, IM2 technical report (2005)
17. McGee, D.R., Cohen, P.R., Wu, L.: Something from nothing: Augmenting a paper-based work practice with multimodal interaction. In: proceedings of DARE 2000, Helsingor, Den-mark, pp. 71–80 (2000)
18. Nelson, L., Ichimura, S., Pedersen, E.R., Adams, L.: Palette: A Paper Interface for Giving Presentations. In: proceedings of CHI 1999, pp. 354–361. ACM Press, Pittsburgh, USA (1999)
19. Norrie, M.C., Signer, B.: Overlaying Paper Maps with Digital Information Services for Tourists. In: Proceedings of ENTER 2005. 12th International Conference on Information Technology and Travel & Tourism, Innsbruck, Austria, pp. 23–33 (2005)
20. Rigamonti, M., Lalanne, D., Evéquoz, F., Ingold, R.: Browsing Multimedia Archives Through Intra- and Multimodal Cross-Documents Links. In: Renals, S., Bengio, S. (eds.) MLMI 2005. LNCS, vol. 3869, pp. 114–125. Springer, Heidelberg (2006)
21. Teevan, J., Jones, W., Bederson, B.: Introduction. In communication ACM 49(1), 40–43 (2006)
22. Tucker, S., Whittaker, S.: Reviewing Multimedia Meeting Records: Current ApproacMultimodal multiparty meeting processing workshop. In: ICMI 2005. International Conference on Multimodal Interfaces (2005)
23. Ullmer, B., Ishii, H.: MediaBlocks: Tangible Interfaces for Online Media. In: Proceedings of CHI 1999, pp. 31–32 (1999)
24. Ullmer, B., Ishii, H.: Emerging frameworks for tangible user interfaces. IBM Systems Journal 39(3-4), 915–931 (2000)
25. Ullmer, B.: Tangible Interfaces for Manipulating Aggregates of Digital Information. Doctoral dissertation, MIT Media Lab (2002)
26. Waibel, A., Schultz, T., Bett, M., Denecke, M., Malkin, R., Rogina, I., Stiefelhagen, R., Yang, J.: SMaRT: The Smart Meeting Room Task at ISL. In: proceedings of ICASSP 2003, vol. 4, pp. 752–755 (2003)
27. Wellner, P.: Interacting with Paper on the Digital Desk. Commun. ACM 36(7), 86–96 (1993)
28. Whittaker, S., Bellotti, V., Gwizdka, J.: Email in personal information management. Commun. ACM 49(1), 68–73 (2006)
29. Dong, X.L., Halevy, A.Y.: A platform for personal information management and integration. In: CIDR, pp. 119–130 (2005)