The Distributed Virtual Meeting Room Exercise

Anton Nijholt
University of Twente
CTIT, PO Box 217
7500 AE Enschede, the Netherlands

anijholt@cs.utwente.nl

Job Zwiers
University of Twente
CTIT, PO Box 217
7500 AE Enschede, the Netherlands

zwiers@cs.utwente.nl

Jan Peciva
Brno University of Technology
Faculty of Information Technology
612 66 Brno, Czech Republic

peciva@fit.vutbr.cz

ABSTRACT

In this paper, we describe our research on distributed virtual meeting rooms. Starting point is our research on multi-party interaction, where the interaction may take place in real, augmented and virtual environments. Moreover, those that interact may be humans, human-controlled avatars, (semi-) autonomous agents, mobile robots, etc. In this paper the emphasis is on connecting meeting environments and transforming perceived and captured meeting activity into multimedia representations of that activity, including re-generation of activity in virtual reality, and making this re-generation perceivable in different, internet-connected, environments. This re-generation is useful for on-line meeting assistance, remote meeting participation and assistance, and for off-line access to meeting information.

Categories and Subject Descriptors

H.5.0 [**Information Systems**]: Information Interfaces and Presentation – General.

General Terms

Algorithms, Design, Experimentation, Human Factors.

Keywords

Smart environments, multiparty interaction, augmented reality, on-line and off-line meeting assistance, embodied agents, virtual reality, inhabited virtual environments.

1. INTRODUCTION

In this paper we introduce our research on a distributed virtual meeting room (DVMR). Starting point of the research are the observations that (1) there is a growing need for teleconferencing, (2) current, commercially available teleconferencing systems are hardly used, (3) current teleconferencing systems are very much biased towards transmitting video and don't consider other ways of transmitting participants' contributions, including manipulating their contributions, let alone, providing means to offer meta-

information about the conference or meeting, (4) current teleconferencing systems assume that all conference participants are remote, rather than assuming that there can be several people in the same location taking part in the conference, and, finally, (5) current teleconferencing systems do not make use or anticipate research results in the areas of image processing, artificial intelligence, animation, virtual reality and information visualization. Obviously, these observations also hold when we look at web casting, remote viewing and audio/web conferencing. Here the emphasis is on offering the viewer advanced viewing facilities (panoramic views, speaker image, whiteboard and sheet views), although also here we see attempts to introduce interactivity. Also automatic camera and microphone control based on speaker localization or viewers' interests (e.g., made explicit by his gaze [1] is considered. However, in general there is poor media richness, interrupted media, delay in media delivery, and, in particular, lack of interactivity. Lack of interactivity means lack of engagement and a poor sense of presence [2].

The situation is slightly different when looking at Computer Supported Collaborative Work (CSCW) systems. The comparison is not completely fair, because here, from the beginning research issues were much more advanced since workers are assumed to collaborate in non-verbal ways (sharing notes, sharing objects), and therefore it is an advantage to have their actions made visible for each other and have a virtual environment designed that supports these activities, while the traditional viewpoint of meetings is that only the verbal interaction is important and needs to be captured. Joint virtual workspaces allowing access from 'remote' places and offering tools for designers and scientists to design and experiment are the future workspaces.

Hence, it is obvious to expect research on smart environments and ambient intelligence earlier to be associated with computer supported collaborative work than with teleconferencing. Research on smart environments and ambient intelligence is about capturing information various kinds of sensors (audio, video, motion and location sensors, wearables, etc.) about activities in an environment, interpreting and enriching this information, and making it available to inhabitants or virtual agents informing and guiding inhabitants. When we speak of inhabitants, we include situations where the environment is virtual and there is only computer-mediated contact between the inhabitants and situations where several people are in the same physical environment and others are allowed to enter this environment, and be virtually present, from remote places. Here, with virtual we do not necessarily mean virtual reality.

EU and DARPA funded research projects on multimodal interaction have been designed to provide the link between smart

environments and ambient intelligence research and meeting or teleconferencing research. EU projects are AMI (Augmented Multi-party Interaction) [3] and CHIL (Computers in the Human Interaction Loop) [4]. The main DARPA funded project on small group meetings is CALO (Cognitive Agent that Learns and Organizes) [5].The research reported in this paper grew out of the AMI project. While in the AMI project the main starting point was the off-line browsing of meeting information, in the research reported here the emphasis is on (1) using AMI technology for real-time visualizing of interpreted meeting information, and (2) using network technology to give real-time access to this information. In our distributed virtual meeting room exercise we have confined ourselves to a real-time representation of meeting events in an environment inhabited by embodied agents representing meeting participants. However, this is only one way to have a real-time mapping and transformation from meeting events to a remotely accessible multimedia representation that allows remote participation, remote experiencing and remote access to meta-information of a meeting.

In the next section we have an overview of technology that has been developed to perform our DVMR exercise. It shows one particular way of connecting smart meeting environments, where each environment can have a number of inhabitants (meeting participants) or just one inhabitant (meeting participant). In section 3 we discuss the exercise as it was demonstrated during a recent conference (MLMI 2005, Edinburgh). Section 4 discusses further ideas based on recently obtained AMI results and more sophisticated ideas about mapping meeting events into multimedia representations that allow real-time remote access and participation.

2. DVMR RELATED OBSERVATIONS

2.1 Viewpoints on AMI Technology

The AMI project is a comprehensive research effort on modeling multi-party interactions in the context of meetings. Multi-party interactions are multimodal of nature; hence, multimodal interactions between meeting participants are subject of research. Group dynamics, group interaction, goals and aims of group members, current verbal and nonverbal meeting interaction, emotions, speech, gestures, poses and facial expressions need to be modeled in order to allow recognition and interpretation. This recognition and interpretation is needed to allow off-line access, but also to allow real-time support of activities by the meeting participants.

In the context of the AMI project we have looked at ways to map meeting activities to multimedia presentations that can be accessed on- or off-line. Obviously, there are many ways to do so. With real-time recognition and interpretation information about the meeting can be made accessible for meeting participants. That is, we have real-time meeting support, but we can of course as well make this information accessible for people not present in the meeting room. When they have real-time access to this information and when they can real-time manipulate this information they in fact take part in the meeting. In an off-line situation more effort can be given to the interpretation of the meeting data and to efforts to allow access to the data in such a way that the meeting in the past can nevertheless be experienced by an off-line 'meeting participant'.

2.2 On Experiencing and Attending Meetings

To get closer to both objectives we have looked at mapping meeting events to representations of these events, possibly enriched with meta-information about the meeting, in 3D virtual reality environments. In previous papers [6,7] we discussed how to obtain enriched virtual reality representations of meeting events from annotated meeting data, where part of the annotations could be obtained automatically and in real-time and where part of the annotation needed to be done manually. Obviously, the real-time obtained meta-information can be made accessible real-time, while the more comprehensive knowledge, obtained by integrating automatically and manually obtained information, can only be made accessible off-line.

2.3 Visualizing Meetings and Meeting Events

Real-time recognition and interpreting multi-party interaction in the context of meetings is the main research issue of the AMI project. Comprehensive interpretation of meeting interaction is far from being possible, it would require comprehensive interpretation and computational modeling of human individual and group behavior. Nevertheless, there is a level of available speech and image processing techniques that allows us to map captured (through microphones and cameras) meeting events (verbal and nonverbal interaction, identifying participants, and tracking of participants in the meeting environment) to multimedia on-line and off-line presentations of these events.

We have looked at transforming meeting events to events in a virtual reality representation of a meeting environment where embodied agents play the role of meeting participants. Given the limitations of real-time speech and image processing techniques, our main interest has been the mapping of the behavior of human meeting participants to the behavior of their representations as embodied agents in a virtual meeting environment. Being able to do this is a prerequisite of further and more intelligent processing of meeting information, including real-time access to meeting data and real-time participation in meetings.

2.4 Capturing Meeting Activity

In our research we have looked at capturing meeting activities from an image processing point of view and at capturing meeting activities from a higher-level point of view, that is, a point of view that allows, among others, observations about dominance, focus of attention, addressee identification, and emotion display. We will return to these issues in forthcoming sections, but here we will look at capturing a limited selection of nonverbal meeting interactions (posture, gestures, and head orientation) only, and we look at possibilities to transform them to a virtual reality representation of a meeting room and its meeting participants.

In order to capture nonverbal activities of meeting participants we studied posture and gesture activity, using our vision software package. Our flock-of-birds software package was used to track head orientation of some of our four party meetings.

The computer vision software processes low resolution, monocular image sequences from a single camera. A silhouette is extracted, shadows are removed and skin color is extracted from the silhouette in order to locate hand and head. Silhouette matching is used to match a projection of a human body model to the extracted silhouette. This allows us to display animated

representations of meeting participants in a (3D) virtual reality environment. The 3D positions of head, elbows and hands can reasonably be calculated [8]. 3D technology based upon portable standards, like VRML/X3D and H-Anim avatars is used. For some meetings to be recorded electromagnetic sensors were mounted on the heads of the participants for tracking their head movements. Especially in meetings this allows us to record and real-time display head orientations of the represented meeting participants. Although there can be differences in head orientation and gaze direction, it nevertheless allows a sufficiently realistic representation of focus of attention behavior (addressing persons, looking at a speaker, looking at notes or looking at the white board in the meeting room).

2.5 A Virtual Meeting Room Representation

The research described in the previous subsections allows us to design a virtual meeting room (VMR) in which the activities of human meeting participants are represented. Virtual reality allows us to view the room from all possible angles, for example the viewpoint of a participant, and by means of a head mounted display we can become immersed.
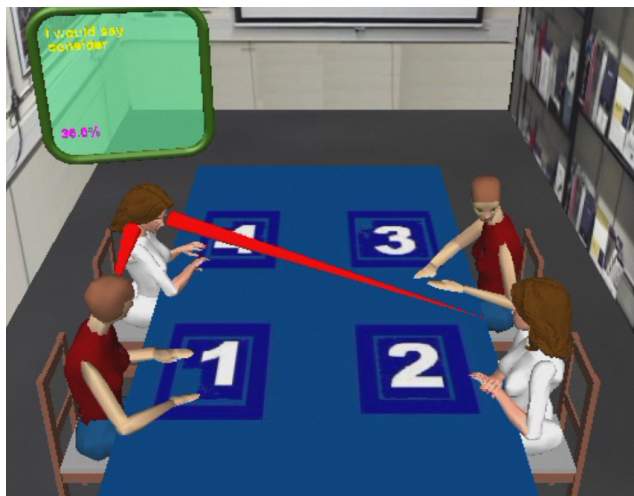
Figure 1. The virtual meeting room showing gestures, head movements, the speech transcript, the addressee(s) of the speaker and the percentage a person has spoken until that moment

Due to our limited number of capturing devices, but also because of imperfect capturing technology and corresponding algorithms, the representation of meeting events is far from perfect. A more perfect representation can be obtained if we are able to use other real-time and automatically obtained annotations of a meeting or, forgetting about real-time constraints, use manually, off-line obtained annotations. Annotations might include results of speech recognition, dialogue structure recognition, talkativity, movements, speaker localization, turn-taking, slide changes, etc. and they can real-time trigger changes (viewpoint changes, adding of metadata, etc.) in the VMR. As such it can play a useful role during a meeting, either for remote viewers or for the meeting participants themselves. We will return to that below. First, we distinguish the following useful applications of our VMR environment (cited from [9]):

• First of all, it allows a 3D presentation and replay of multimedia information obtained from the capturing of a meeting. Depending on the state of the art of speech and image processing (recognition and interpretation) one may think of manual annotation replay, replay based on both manually and automatically obtained annotations and interpretations and replay purely based on fully automatically obtained interpretations. Obviously, when the meeting environment has the intelligence to interpret the events in the meeting environment, it can transform events and present them in other useful ways (summaries, answers to queries, replays offering extra information, visualization of meta-information, etc.);

• Secondly, transforming annotations, whether they are obtained manually or automatically, can be used for the evaluation of annotations and annotation schemes and of the results obtained by, for example, machine learning methods. Current models of verbal and nonverbal interaction, multi-party interaction, social interaction, group interaction and, in particular considering our domain of meeting activities, models of meeting behavior on an individual or on a group level, are not available or only available for describing rather superficial phenomena of group interaction [10]. Our virtual room offers a test-bed for eliciting and validation of models of social interaction, since in this representation we are able to control the display of various independent factors in the interaction between meeting participants (voice, gaze, distance, gestures, facial expressions) and therefore it can be used to study how they influence features of social interaction and social behavior.

• Thirdly, a virtual reality environment can be used to allow real-time and natural remote meeting participation. In order to do so we need to know which elements of multi-party interaction during a meeting need to be presented in a virtual meeting in order to obtain as much naturalness as possible. The test-bed function of a virtual meeting room, as mentioned above, can help to find out which (nonverbal) signals need to be mediated in one or other way.

As mentioned in the first bullet, the VMR allows us to reconstruct a meeting, but when useful we can do it in a different way. Gestures can be exaggerated, pointing can be done such that it is better recognizable, speech can be improved, and we can even have different combinations of modalities than used in the real meeting.

3. PERFORMING THE DVMR EXERCISE

As mentioned, a VMR can be used to allow real-time and natural remote participation. Participation requires real-time interaction with other meeting participants. This section is concerned with real-time use of the VMR.

3.1 VMR for Life Meeting Assistance

Let us first consider the situation where we offer the VMR to the meeting participants inhabiting the physical meeting room while they are interacting. While meeting they can get all kinds of information about the meeting presented in this virtual environment and they can use it as a domain-dependent browser asking questions like: Who is this person, what did he/she say about this topic in a previous meeting, why is this person getting

upset when we talk about this topic, etc.[1] Hence, due to this visualization, meeting participants may feel stimulated to ask questions related to behavior of meeting participants, meta-information displayed in the environment and events taking place (without disturbing the meeting). Clearly, when looking at the VMR from this point of view it serves the role of providing life meeting assistance to the meeting participants present in the real meeting room. The visualization provides the context for the user to interact with the system and it provides the context for the system to interpret and assist the user.

Remote on-line viewing of the VMR is of course no problem. That is, non-meeting participants can get access and see what's going on. This does not require interactivity, although, inherently to virtual reality, any viewpoint can be taken, meaning that e.g. the viewpoint from an empty chair at the table can be taken. This audience is not necessarily visible by the meeting participants. A slight extension should allow visualization of the audience, for example as avatars, and making the meeting participants, still assuming that they use the VMR as life meeting assistant, aware of who is in the audience. We have not done this yet, but it fits in a tradition of multi-user virtual environments, where in this case the multi-user environment can be constrained to a public gallery, not disturbing the meeting. Obviously, many other ideas common in multi-user virtual reality environments and distributed virtual environments, including the various ways of distribution of data and processes can be introduced here [11].

3.2 VMR for Distributed Meeting Assistance

The general objective of the DVMR exercise is that we have different smart locations, each equipped with cameras, microphones and probably other sensors. These smart locations are inhabited by one or more meeting participants that take part in a virtual meeting room in which all locations and their inhabitants join. That is, we connect smart meeting rooms, we can connect individual remote participants to smart meeting rooms and we can connect many individual remote participants to one joint virtual meeting room. From every location we need to capture meeting behavior of the inhabitants and make it available to the joint virtual reality meeting room that can be accessed by every meeting participant from every location. (see Figure 2, where we have connected two smart meeting rooms, and Figure 4, where we apply transformations to captured activities before displaying them in the joint virtual meeting room.

In our exercise, demonstrated at the MLMI 2005 conference in Edinburgh, local constraints and resource limitations did not allow us to demonstrate the full potential of our technology. We confined ourselves to a situation where one remote meeting participant joined a meeting of three embodied agents in the form of an animated embodied agent in a virtual representation of the
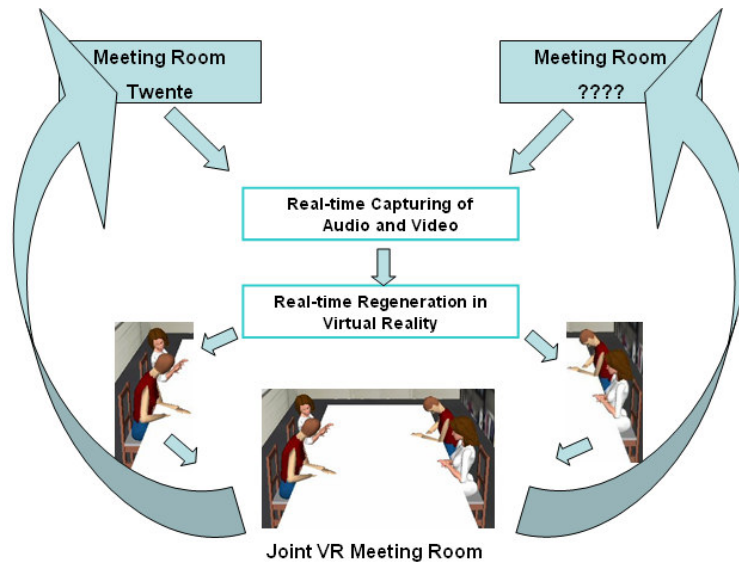

**Figure 2**. Capturing and re-generation of meeting activities from remote locations

IDIAP smart meeting room. Capturing of the remote participant was done using a simple web camera and electromagnetic sensors to measure head orientation. In the near future we may expect that these latter sensors can be replaced by other, less non-obtrusive, sensors (e.g. rfid tags, glasses, headsets, wearables). At the moment we use technology developed in our group (vision software, flock-of-bird sensors, a multi-agent platform, and DVMR clients and server software) for tracking meeting activity in remote and connected environments and transforming it into activities displayed by embodied agents in a joint virtual reality meeting room. A remote participant, in fact every participant, can see the DVMR with avatars representing the meeting participants and can see the meeting activities of those participants.

In Figures 2 we have illustrated the situation where two smart meeting rooms are connected and the captured information is displayed in a virtual reality meeting room.

The technology used within the DVMR experiment differs substantially from normal video conferencing technology. Rather than sending video data as such, this data is transformed in a format that enables analysis and transformation. For the DVMR experiment the focus was on representing poses and gestures, rather than, for example, facial expressions. Poses of the human body are easily represented in the form of skeleton poses, essentially in the same format as being used for applications in the field of virtual reality and computer games. Such skeleton poses are also more appropriate as input data for classification algorithms for gestures. Another advantage for remote meetings, especially when relying on small handheld devices, using wireless connections, is that communicating skeleton data requires substantially less bandwidth than video data. A more abstract representation of human body data is also vital for combining different input channels, possibly using different input modalities. Here we rely on two different input modalities: one for body posture estimation based upon a video camera, and a second input channel using a head tracker device. Although the image recognition data for body postures also makes some estimation of the head position, it turned out that using a separate head tracker was much more reliable in this case. The general conclusion is,

---

[1] In a video of the CHIL project a virtual meeting room is introduced that allows meeting participants to on line identify other meeting participants, allowing the avoidance of embarrassing situations where names and faces are forgotten.
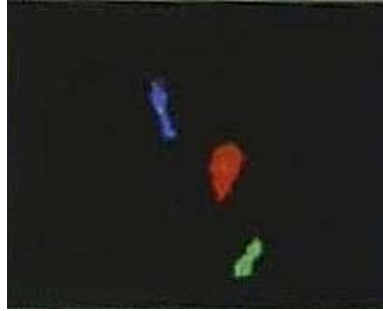
Figure 3 Remote participant, making some gesture, the gesture recognition, and the representation within the VMR, as seen through the eyes of one of the other participants.

not so much that everyone should use a head tracker device, but rather that the setup as a whole should be capable of fusing a wide variety of input modalities. This will allow one to adapt to a lot of different and often difficult situations. In the long run, we expect to see two types of environment for remote meetings: specialized meeting rooms, fully equipped with whatever hardware is needed and available for meetings on the one hand side, and far more basic single user environments based upon equipment that happens to be available. The capability to exploit whatever equipment is available might be an important factor for the acceptance of the technology. In this respect, we expect a lot from improved speech recognition and especially from natural language analysis. The current version of the virtual meeting room requires manual control, using classical input devices like keyboard or mouse, in order to look around, interact with objects etcetera. It seems unlikely that in a more realistic setting people that are participating in a real meeting would like to do that. Simpler interaction, based upon gaze detection but also on speech recognition should replace this situation.

The DVMR-Server transforms its input to an up-to-date distributed virtual meeting room. Objects in the DVMR can be controlled/moved by the DVMR inhabitants. As an example, since many of our recorded meetings are design meetings devoted to the design of a remote control, we designed a remote control and put it in the DVMR as an example of how real and remote meeting participants can discuss and manipulate the properties of this remote control. Clearly, visualizing and manipulating objects that are under discussion, whether they represent physical objects or documents and presentations, is an important issue in advanced meeting technology.

The remote participants have a virtual position at the table, and can watch the meeting from that virtual position or, if they prefer, can watch the meeting from a more global point of view. The same hold true for the other participants: they will see the remote person at his or her virtual position, making the movements and gestures of the real person. The technology is based upon simple consumer web cams, together with image recognition technology that extracts key features, like body position and gestures. This process is illustrated in figures 3. A more global view is presented in figure 4.

Figure 5 shows that there is the possibility to transform meeting activities to other media, modalities and appearances before displaying them to meeting participants. We have chosen to make transformations from and to modalities, since that shows how detailed we can go; obviously, modality changes, changes of

combinations of modalities and replacing human modalities by other media to present activities and information can be considered.

Each computer running the DVMR is transforming inputs from its input devices to its virtual meeting room replicas. Our distributed version of the VMR is using recent developments in database technology based on delayed commits for time-stamped transactions in replicated multi-version databases. Objects in the DVMR can be controlled and moved by the DVMR inhabitants. As an example, since many of our recorded meetings are design meetings devoted to the design of a remote control, we designed a remote control and put it in the DVMR as an example of how real and remote meeting participants can discuss and manipulate the properties of this remote control.

The DVMR replicates meeting data among all participating computers. Three types of replication can be distinguished:

- static data

- primary-copy replication

- delta consistency replication

Static data are those that never changes. Once they are set at creation or loading time, they stay the same. No synchronization is necessary for these data.

Primary-copy data includes, for example, avatars and all other objects that are modified just by one computer. Primary-backup
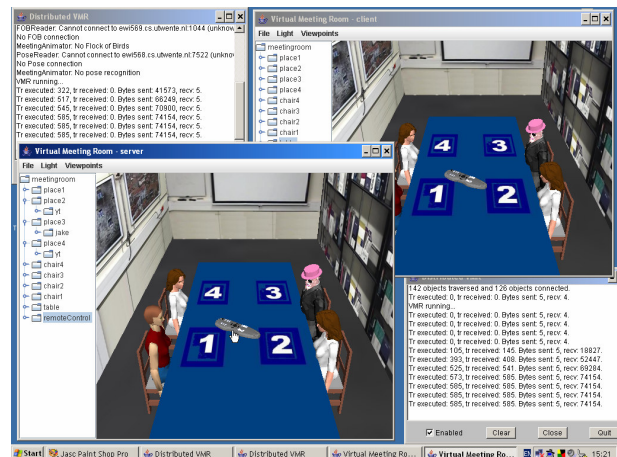
Figure 4. Global view of two synchronized meeting rooms.
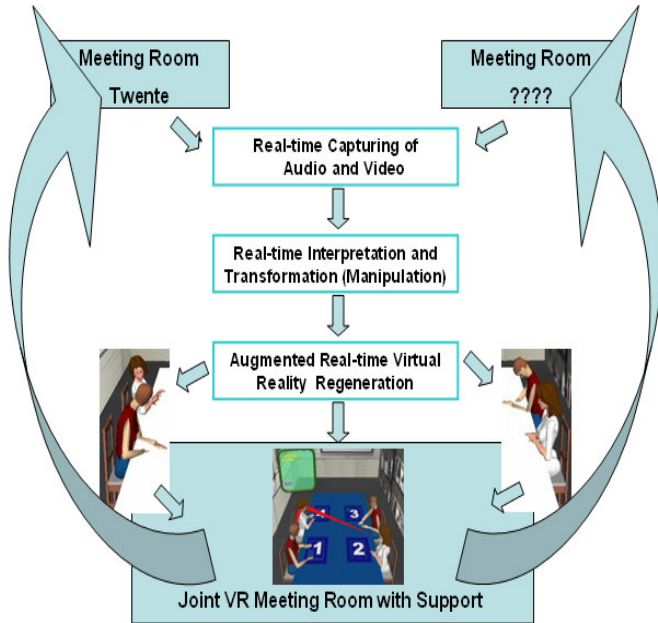
Figure 5. Capturing, manipulation and re-generation of meeting activities from remote locations

replication [12] means that just the computer holding the primary replica can modify its value. All other computers hold "backups" that can not be directly modified without contacting the primary replica.

The most complicated situation is when several participants want to modify one object simultaneously. For example, more people want to manipulate the remote control. Concurrent writes from different computers with different values may break the consistency of the scene, for instance, one participant might see a yellow remote control whereas others see a blue one. Therefore, some rules are necessary for definition and maintenance of consistency. Our solution is based on time-stamped transactions with delayed commits. It can be considered as extended delta consistency [13]. The concurrency rules grant that when concurrent writes happen, the same write will win over another write on all computers resulting in the same virtual scene databases on all computers, thereby ensuring consistency.

The concurrency rules are defined as follows:

- each write is causally dependent on previous write

- when more concurrent writes are found, the earliest of them is accepted

To get the rules working, global order of writes is necessary. Therefore each write is marked by globally unique timestamp. We are using Lamport Timestamps [14] without final acknowledging. Then, the global order of writes is established by sorting them based on their timestamps.

Because the order of writes cannot be established immediately, since it takes some time for updates to propagate through the network, we are using delayed commit. It means that the order of writes is considered temporary or speculative until writes are older than longest network latency between computers. Then, we are sure that no other older write will arrive to change the order of

these old writes. The moment of turning write from "speculative" to "permanent" we call commit.

The concurrent writes are detected by the commit operation. If the write is about to commit and it does not depend on the last committed value, it has to be aborted since it depends on already overwritten value or other aborted write.

Dividing writes into two classes - speculative and permanent – makes it possible to do some optimizations related to network latency. For example, it is possible to change a color of the remote control speculatively and user can see the result immediately without waiting until network communication is done and the value is committed. This optimistic behavior has the advantage of showing an immediate response to the user, while consistency of the scene is guaranteed by existence of "committed" scene. If any write that is still not committed turns out to break the scene consistency, it is aborted and its effect is removed from the "speculative" scene that is shown to the user.

4. CONCLUSIONS

Currently we work on the integration of speech recognition in the DVMR in order to select or manipulate objects or agents in the environment. For speech software we use SpeechPearl XML (ScanSoft). Another issue that is being looked at is capturing and mediating gaze behavior between remote meeting participants. Sharing and manipulating shared objects is another issue that is requiring our attention. Obviously, and also part of our current research efforts, is virtual reality visualization of meeting events, interpretations of these meeting events in order to produce semantic preserving transformations, and presentation of these meeting events using various media sources. We also hope to integrate current research efforts on personalization of embodied agents, facial expressions and emotion display into our efforts to obtain meeting environments that allow for more natural meeting experiences.

5. ACKNOWLEDGMENTS

6. REFERENCES

[1] Smith, D. et al. OverHear: Augmenting attention in remote social gatherings through computer-mediated hearing. *Proceedings CHI 2005*, Portland, Oregon, USA, ACM 1801-1804.

[2] Lombard, M., & Ditton, T. B. At the heart of it all: The concept of presence. *Journal of Computer-Mediated Communication*, 3(2) (1997).

[3] McCowan, I., Gatica-Perez, D., Bengio, S., Moore, D., Bourlard, H. Towards Computer Understanding of Human Interactions. In: *Ambient Intelligence*, E. Aarts et al. (Eds.), Lecture Notes in Computer Science, Springer-Verlag Heidelberg, 235 - 251. See also: http://www.amiproject.org/.

[4] Waibel, A., Steusloff, H., Stiefelhagen, R. CHIL - Computers in the Human Interaction Loop. 5th Intern. Workshop on *Image Analysis for Multimedia Interactive Services*, April, 2004, Lisbon, Portugal. See also http://chil.server.de/.

[5] CALO: http://www.cse.ogi.edu/CHCC/Projects/CALO/

[6] Reidsma, D., Op den Akker, R., Rienks, R., Poppe, R., Nijholt, A., Heylen, D., and Zwiers, J. Virtual Meeting Rooms: From Observation to Simulation. *Proc. Social Intelligence Design 2005*, March 2005, Stanford University, CD Proceedings, R. Fruchter (Ed.).

[7] Rienks, R., Nijholt, A., and Reidsma, D. Meetings and Meeting Support in Ambient Intelligence. Chapter in: *Ambient Intelligence, Wireless Networking, Ubiquitous Computing.* Th. A. Vasilakos & W. Pedrycz (eds.), Artech House, Norwood, MA, USA, 2005, in press.

[8] Poppe, R., Heylen, D., Nijholt, A., Poel, M. Towards real-time body pose estimation for presenters in meeting environments. *Proc. 13th Intern. Conference in Central Europe on Computer Graphics, Visualization and Computer Vision.* V. Skala (Ed.), Plzen, Czech Republic, 2005, 41-44.

[9] Nijholt, A. Meetings in the Virtuality Continuum: Send Your Avatar. In *Proceedings Cyberworlds 2005*, A. Sourin (Ed.), November 2005, Singapore, to appear.

[10] Traum, D., Rickel, J. Embodied agents for multi-party dialogue in immersive virtual worlds. *Proceedings of the first International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002),* Bologna, Italy, July 2002, 766-773.

[11] Saar, K. VIRTUS: A collaborative multi-user platform. *Proceedings of the fourth symposium on Virtual reality modeling language.* Paderborn, Germany (1999) 141-152.

[12] Wiesmann, M., Pedone, F., Schiper, A., Kemme, B., Alonso, G. Understanding replication in databases and distributed systems. *Proceedings of ICDCS'2000*, Taipei, Taiwan, R.O.C., 2000, 264-274.

[13] Singla, A., Ramachandran, U., Hodgins, J. Temporal notions of synchronization and consistency in Beehive. *In Proceedings of the Ninth Annual ACM Symposium on Parallel Algorithms and Architectures (Newport, Rhode Island, United States, June 23 - 25, 1997).* SPAA '97. ACM Press, New York, NY, 211-220.

[14] Lamport, L. Time, clocks and the ordering of events in a distributed system. *Communications of the ACM*, 21(7), July 1978, 558-565.