Collaboration Applications in Smart Spaces

Alice Oh, Howard Shrobe, Andy Chang, Krzysztof Gajos, Tyler Horton, Stephen Peters & Rattapoom Tuchinda

Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

http://www.ai.mit.edu

The Problem: Every day, people are engaged in various types of collaboration tasks in the Intelligent Room [1]. However, no software exists that can support collaboration effectively and smoothly in a smart environment setting.

Motivation: Collaboration is one of the most important activities that can benefit from the software and hardware infrastructure of smart spaces. It is difficult, however, to design and build useful software to assist human-human and human-computer collaboration in and across smart spaces. The application must be flexible to allow many users to connect and share their information, but at the same time be protected from outside intrusions and unwanted data sharing. The user interface must be natural and efficient so that it is easy for the users to take advantage of all of the features. There are also social and psychological implications of collaboration that must be considered. Each of these issues is a difficult problem on its own, and we do not attempt to solve any of them completely in this project, but we will try to put together a set of practical solutions for useful collaboration applications.

There are three types of collaboration we would like to support. The first is for group collaboration within a smart space. This includes face-to-face meetings, informal discussions, brainstorming sessions among others. The second is for collaboration across different spaces. The third type is collaboration between human users and computer agents.

Previous Work: Computer meeting support has been a topic of active research in the Computer-Supported Cooperative Work (CSCW) community. However, their focus has been modifying the meeting process via use of technology, rather than developing systems to better support the meeting process. For example, in [5], all communication is done via typed input from meeting participants, each sitting in front of a workstation. While this may increase participation and improve the group decision making process, it has the disadvantage that the communication style is very different from the natural-occurring meetings, thereby forcing all participants to learn and adapt to the new meeting process. In [6], they take a step further and support different types of meetings, including traditional face-to-face meetings, but the meeting record is again based on the written material produced during the meeting.

A team at Xerox PARC has looked at capturing multimedia meeting records and buidling tools for *salvaging* [4]. This requires the salvager to go through the multimedia meeting record after the meeting and produce an annotated record. A similar approach is taken in [2], but with personal digital assistants (PDAs) and capturing only written notes. These two systems offer flexibility of supporting unstructured meetings, the latter providing support for meetings held anytime, anywhere (assuming the participants have PDAs).

The Interactive Systems Lab has been working on a system for multimodal meeting record, in which audio and video are processed and annotated automatically with rich information such as speaker identity, participants' emotions, etc. [3]. Their system uses sophisticated technologies such as large-vocabulary speaker-independent speech recognition, eye-gaze tracking, and dialogue processing. However, while the paper mentions that it captures the action items when they are reiterated at the end of the meeting, the system itself does not impose and support any structure to a meeting. This makes it more difficult to use the meeting record as a simple reminder of the meeting to the participants.

Approach: We are working on three applications, one for each type of collaboration.

The first is called the MeetingManager. It is a group meeting support application that takes advantage of the natural structure of the meeting that many organizations practice for an efficient meeting process. In some sense, it is a combination of the meeting process-focused work of the CSCW community, and the free-form meeting support

of the HCI community. The MeetingManager takes advantage of minimal structure for a meeting (an agenda with issues and commitments) to produce automatic, real-time annotated multimedia meeting record. The MeetingManager has four components that work together: a planner, a facilitator, a summarizer, and a browser. Before the meeting takes place, the participants use the planner to make and store the agenda. During the meeting, the facilitator takes the stored agenda for real-time meeting assistance. After the meeting, a brief summary of the recorded multimedia is emailed to the participants. Finally, the meeting is stored in a database available for convenient browsing.

The second is a multi-user game application for Boggle. It uses a server-client architecture to enable multiple users to connect to a single game server. This is an instance of collaboration across different locations. The clients access the server through message passing, allowing the server to continue serving the needs of the current clients should one of them be disconnected in the middle of the game. At this time, any user is allowed to connect to the server and share information, but we are working on protocols to ensure that security issues that may arise can be resolved.

Finally, we are working on an application for human-computer collaboration. We plan to build a tourist information browsing and planning application to be used in the Intelligent Room. The knowledge base that the computer has can be delivered to the user via sophisticated multimodal output capabilities of the Room. In turn, the user will be able to talk to the computer through multiple modalities including speech, gesture, head movements. Together the user and the computer will solve task-oriented problems.

Impact: These applications showcase the capabilities of the Intelligent Room. In doing so, we raise important research questions about human-computer interaction, human-human interaction with the help of the computer, and information sharing across locality and time. These collaboration tools will be useful prototypes for further development in our Intelligent Room, and will also serve as good examples to build upon in other smart spaces and ubiquitous computing environments.

Future Work: We plan to develop these applications further to include more sophisticated HCI technologies. We hope to consider the problems of security and access, as well as more traditional AI problems of planning and problem solving.

Research Support: This research is supported by the MIT Project Oxygen.

References:

[1] Michael Coen. Design principles for intelligent environments. *The Fifteenth National Conference on Artificial Intelligence*, 1998.

[2] Richard C. Davis, Jason A. Brotherton, James A. Landay, Morgan N. Price, and Bill N. Schilit. Notepals: Lightweight note sharing by the group, for the group. *Proceedings of Human Factors in Computer Systems: CHI99*, 1999.

[3] Ralph Gross, Michael Bett, Hua Yu, Xiaojin Zhu, Yue Pan, Jie Yang, and Alex Waibel. Towards a multimodal meeting record. *Proceedings of ICME 2000*, 2000.

[4] T. P. Moran, L. Palen, S. Harrison, P. Chiu, D. Kimber, S. Minneman, W. van Melle, and P. Zelweger. I'll get that off the audio: A case study of salvaging multimedia meeting records. *Proceedings of Human Factors in Computing Systems, CHI97*, pages 202–209, March 1997.

[5] J. F. Nunamaker, Alan R. Dennis, Joseph S. Valacich, Douglas R. Vogel, and Joey F. George. Electronic meetings to support group work. *Communications of the ACM*, pages 40–61, March 1991.

[6] Norbert A. Streitz, Jorg Geissler, Jorg M. Haake, and Jeroen Hol. Dolphin: Integrated meeting support across liveboards, local and remote desktop environments. *Proceedings of the Conference on Computer-Supported Cooperative Work, CSCW'94*, 1994.