

Pro-active Meeting Assistants: Attention Please!

Rutger Rienks and Anton Nijholt
University of Twente, Enschede, the Netherlands
{rienks, anijholt}@ewi.utwente.nl

Paulo Barthelmess
Natural Interaction Systems, Portland, OR, U.S.A.
Paulo.Barthelmess@naturalinteraction.com

Abstract

This paper gives an overview of pro-active meeting assistants, what they are and when they can be useful. We explain how to develop such assistants with respect to requirement definitions and elaborate on a set of Wizard of Oz experiments, aiming to find out in which form a meeting assistant should operate to be accepted by participants and whether the meeting effectiveness and efficiency can be improved by an assistant at all.

Keywords: Meetings, Assistants, Agents, Pro-active, Efficiency

1 INTRODUCTION

Meetings are often inefficient [46]. Starting with probably the first meeting ever held by humans, people have looked at techniques and protocols to enhance them. The development of technology to support meetings has therefore long been a subject of research [54].

Meetings can nowadays be assisted by a wide variety of tools and technologies, facilitating interaction, saving money and time, and creating opportunities that would not be possible without technology. The foremost benefit of technology so far is its support for meetings in which participants are distributed. Being able to attend meetings remotely results in substantial savings of time and money that might have been otherwise spent on travel. Tele-conferencing systems augmented with additional advanced services such as instant messaging, file transfer and application sharing are becoming more and more prevalent. In the near future meetings will be possible in virtual worlds where participants will be represented by virtual humans [34].

There is also evidence that technology-enabled processes can positively impact meeting performance. Studies reported by De Vreede et al. [12] and Nunamaker Jr. et al. [35] show a significant reduction in labor cost and overall project duration when Group Support Systems (GSS), or Electronic Meeting Systems (EMS) are used. These systems support alternative, technology-enabled meeting processes that can help participants with the formulation of and search for solutions to ‘problems’ listed on the agenda. A participant generally has a computer terminal connected to a central server at his or her disposal through which several problem resolution tools are available. Typical tools are an electronic brainstorming tool, an idea organizer, a topic commenter and a voting support tool.

Despite the huge savings and proven increased efficiency brought about by GSS and similar technology, its adoption has proven sometimes problematic. There are instances in which the use of these systems has been discontinued due to stakeholders’ objections to the (radical) changes in work practices that are introduced by them [35]. This leads us to investigate alternative means for positively influencing meeting outcomes in ways that would encounter less resistance. In particular, we want to investigate how pro-active meeting assistants can be exploited to reap the benefits of technology-enabled meetings instead of being exposed to its drawbacks. Successful automated

meeting assistants can potentially integrate themselves into their surrounding social environment, offering support that blends more seamlessly into users' work practices.

Technology in the field of meeting support ranges from completely passive objects like micro-phones to pro-active autonomous actors such as virtual meeting participants. In earlier work we defined several dimensions that can be distinguished in this spectrum, with the major ones being the *reasoning ability*, the *acting ability* and the *sensing ability* [44]. In this paper we will focus on *pro-active* meeting assistants that are able to act autonomously. Pro-active meeting assistants are those that (preferably in real-time) support the participants and act autonomously in the process either before, during, or after a meeting. For this type of assistants, its operating dimensions are highly dependent on their functionality.

This functionality or sophistication directly depends on the state of the art of automatic collection of appropriate meeting information (the *sensing*) as well as on the required intelligence to use it (its *reasoning ability*) and the means through which the assistant can influence a meeting (its *acting ability*). To aid in this process, so-called 'smart' meeting rooms appeared. These smart rooms embed all sorts of sensors, providing data about the meeting and hence create the opportunity to collect and learn from this data in order to build models. These models may in turn provide insights into interactions and their contents. The first project presenting ideas to augment meetings with various 'smart' technologies was probably Project Nick [10]. This project discussed the incorporation of screens displaying both the agenda and live meeting statistics to aid the meeting process. From that point onward smart meeting rooms appeared at several institutions where large meeting corpora were recorded. In the last four to five years there has been a surge in interest in meeting support. Many large projects were established, including consortia with partners from all over the globe, working on meeting collection, and research on meeting models and support technology [33, 55, IM2 Website, CALO Website, Nectar Website].

The remainder of this paper will elaborate on the concept of pro-active meeting assistants, in particular software agents that aim to assist the meeting process and thereby facilitate more effective and efficient meetings. As there are a lot of ideas but hardly any implemented systems yet, we will, apart from looking at the existing ideas, show how to get from ideas to a full requirements specification. We also present a Wizard of Oz experiment where we simulate several forms of pro-active meeting assistants designed to streamline the meeting process.

2 MEETING ASSISTANTS

Meeting assistants have been the topic of research in various projects, e.g., the Neem Project [16, 3]. In Neem, a basic premise is that assistance has to be provided along multiple dimensions, including the organizational, but also the social and informational. A good meeting is one in which organizational goals are achieved, but not at the expense of the social well-being of a group. Support in Neem revolves around *tools* and *virtual participants*, both of which are designed to explore aspects along the organizational, social and informational dimensions. *Tools* are artifacts that crystallize certain aspects of an interaction, allowing for participants to become aware of and be able to influence these aspects. (e.g. by being able to manipulate items of discussion within an agenda tool.) *Virtual participants* are anthropomorphic assistants. They are designed to have consistent personalities and well-determined roles. *Kwaku* is a virtual participant that takes care of the organizational aspects of a meeting. Kwaku for instance reacts to discussions that extend over the pre-allocated period of time by reminding participants that they might want to move on to the next agenda item. Kwaku "listens" to the reaction of the group (by examining transcribed speech and text message channels) and will either update the agenda tool, moving it to the next agenda item in case of agreement, or leaving it in the current item if its perception is that the suggestion was overruled by the group. *Kwabena* on the other hand is a social facilitator. Kwabena looks after the participants' social well-being, monitoring the actions a group would want to undertake at each point in time, such as take a break, switch topics, change the level of detail, or pace of the interaction. These wishes are expressed via a 'Moodbar' tool that displays a set of possible actions that participants can select by clicking on corresponding buttons. A mechanism is provided to poll the input from the different participants. Kwabena takes the initiative to suggest

the course of action (e.g. taking a break) expressed by the group. (e.g. by voicing the suggestion via all participants' audio systems.) Conversely, if a particular participant is expressing wishes that disagree with the rest of the group, Kwabena communicates in private with this participant, letting him or her know that the rest of the group seems to think differently. Finally, *Kweisi* is responsible for providing the group with additional information. This can happen upon request of one or more participants, but also autonomously, as *Kweisi* perceives (again by analyzing the content of the speech and typed messages) that a certain topic is under discussion for which additional documents are available.

All these assistants can be realized as embodied pervasive software systems that operate alone or in groups, interact with the users and with other participants, and learn user preferences. Neem illustrates an approach to assistance during the meeting. We will now frame ongoing research in the domain of meeting assistants by making a division into assistants that support activities that take place before, during as well as after the meeting.

2.1 ASSISTANCE FOR MEETING PREPARATION

A first opportunity for assistance takes place at the meeting preparation phase. Opportunities during this phase can be related to the identification of the group of people for whom a meeting's particular topics of discussion are of interest. Once it is assured a meeting will take place, a meeting planner can be used to assist with the creation of the actual agenda and with the negotiation of schedule, time and place for the meeting. There has been some research on agents that schedule meetings, for example Garrido and Sycara [19] and more recently Oh and Smith [37]. Bowring et al. [6] proposes the use of agents to optimize schedules given a set of personalized criteria. This is more or less similar to the personalized time management system described in Berry et al. [5] where a personal assistant is described that has, amongst others, the ability to negotiate with other personal assistants for a suitable time and location given people's constraints or preferences.

Such preference-driven negotiations can lead to flexible scheduling of meetings. An assistant that has enough perception and reasoning capabilities could for instance take advantage of its knowledge of people's whereabouts to schedule a meeting at a time during which it knows the people involved will be in the same building, rather than requiring them to waste time commuting on another less convenient day.

Besides dealing with the instantiation of the meeting's agenda, its time and its location, another type of assistant could for example help with the choice of chairman based on the personalities of the participants (see the SYMLOG agent as described in Wainer and Braga [56]). These sorts of assistants could also advise on the desired format of a meeting depending on the expected attendees, or propose a group size for a meeting based on the topics of the agenda as mentioned in Padilha and Carletta [38]. Once the date and location are settled, assistants might inform participants about possible changes in the schedule and gather the documents to be discussed. Others could prepare the data projector, the light settings and temperature settings of the room and schedule the presentations such as mentioned in Chen et al. [8].

2.2 ASSISTANCE DURING THE MEETING

Antunes and Carrio [1] describe three main aspects that pertain to meetings: the *meeting process*, the *meeting resources* and the *meeting roles*. To optimize the meeting process one could have assistants like Kwabena, Kwaku and Kweisi take care of participants' well-being, the organizational and the informational aspects respectively. Assistants could greet the participants and make them feel at ease [9]. Other assistants could look after the content of the meeting, by analyzing the semantics of the group discourse; it would then be possible to gauge the progress of a discussion (convergence/divergence), signal possible repetitions, or determine the level of agreement or disagreement [18]. With respect to the meeting resources, once a meeting starts, the context of the meeting such as the room and equipment can be regulated (e.g. by closing curtains, starting projectors, etc.) - see Oh et al. [36] for more examples. An assistant could also alert participants when someone is calling them [11], or provide background information about other participants.

A final category of meeting assistants can aid specific meeting roles. In the best case the complete role of the secretary could be performed by an automated meeting assistant, as we could have an assistant that creates meeting minutes and takes over all care of other tasks pertaining to a secretary. The role of the chairman could be similarly assisted and eventually completely replaced by an automated assistant. These are complex roles, and much work remains to be done before enough is achieved in terms of understanding of the meeting dynamics and of the issues related to the integration of such assistants into meetings.

An assistant taking over the role of meeting chairman should at a minimum take care of the activities carried out by a human chairman. In a meeting, the chairman has to manage the meeting process in order to maximize the output of the meeting, stick to the agenda and to maintain a positive meeting atmosphere. Guarding agenda and time constraints is an obvious task: taking care of the decision-making process and trying to exploit the expertise of the meeting participants is much less obvious. All sorts of assistants could gather information that could be useful in this respect. A chairman could, for example, be provided with points of view expected from the participants, based on the known background of a participant or on the companies' viewpoints about a topic. All sorts of participants' behavior that could potentially influence the process might be of relevance for a chairman. Niekrasz and Purver [32] describe the usage of a shared discourse ontology that could serve as common ground for these sorts of assistants. We elaborate on the aspect of (semi-)automatic meeting understanding in Section 2.4, where we describe related research we have carried out. Other relevant abilities for an assistant in the role of a meeting chairman are, for example, described by Jebara et al. [22], where a system is able to provide feedback and ask relevant questions to stimulate further conversation. Nakanishi et al. [31] describe a system in the role of a party host, which tries to find a safe common topic of conversation for participants having trouble communicating. This system is able to generate a topic closely related to the ongoing conversation based on a set of detected keywords and a topic tree prepared beforehand. For more elaborate information about leadership issues and required abilities for a meeting chairman the reader is referred to [50, 29].

2.3 ASSISTANCE AFTER THE MEETING

The preservation of meeting information, also referred to as *group memory* is a problem due to the volatile nature of meetings. Apart from the fact that people might be interested in things not captured in the notes, it might take hours to find answers by digging through piles of hard-copy notes. After the meeting, assistants could remind people of commitments and action items they are responsible for. Other assistants might analyze the interaction and produce documents and artifacts that reflect the content of the discussions. An example of such a system is CALO's Charter [24]; this suite of agents analyzes multi-modal interaction during project planning meetings and automatically produces MS-Project renditions of Gantt Charts sketched by participants on interactive boards, thus avoiding the manual labor to reenter the information that would otherwise be necessary. More recently, this system was extended to support collaborative sketching by multiple, potentially distributed participants [4].

Assistants could also provide selective information about the meeting. Three categories of people can be distinguished that might show interest in (parts of) the content or outcome of a meeting: 1) the actual participants, 2) people who did not attend the meeting interested in aspects such as the contributions of a person, or the arguments in favor or against a specific decision, and 3) analysts who are just willing to gather information about meeting processes in general.

The key issue is to provide access to representations of conveyed information from the meeting as mentioned in Palotta et al. [39]. Once this information is available for access, it is shown that people will adapt their way of working based on what they have available in order to increase efficiency [30]. As it might be hard for people to express their informational needs to an assistant, the interface is of utmost importance. Jaimes et al. [21] describes an implementation of a system that helps users to easily express cues people might recall about a particular meeting. A related area of research is the automatic generation of (multimodal) summaries of a meeting (see e.g. [17]). In fact, a summary can be seen as an answer to a question, where the best summary is

perhaps the one that answers the most frequently asked questions. Ultimately, we would want assistants that would be able to answer all questions in a clear and comprehensible manner. For an overview of approaches providing access to meeting records see [52].

A very important question here is what information should be captured [49], which is tightly related to what people would want to remember from meetings. Lisowska [28] gives an overview of typical queries posed to meeting retrieval systems, obtained through questionnaires, that in the future will be evaluated using a Wizard of Oz experiment. Similar research was conducted by Banerjee et al. [2]. It appears that people are interested in two kinds of information: 1) descriptions of the interactions among participants and 2) things that involve elements from the meeting domain itself. In order to provide information about the interaction amongst participants several techniques have to be developed, able to frame the understanding of what is going on in a meeting. Apart from preserving this data for people interested in it after the fact, this type of information would also be highly relevant to almost any assistant operating during the meeting. The next section will therefore elaborate on, and give pointers to ongoing research about the automatic interpretation of human behaviour in a meeting setting.

2.4 INTERPRETING HUMAN BEHAVIOUR

Once a chairman is appointed and given the authority to manage the meeting process he or she is authorized to perform a set of interventions such as selective turn giving and interrupting. These typical actions are triggered based on the displayed behavior of the participants. The occurrence of unwanted situations such as a rare event with a large disturbing impact, or the repetitive occurrence of events with a smaller disturbing impact are typical examples of situations that could trigger an intervention. Human behavior reveals itself through several modalities over time.

The behaviour of meeting participants is generally evaluated relative to social norms and regulated by various means of social control. These norms generally are unstated and unwritten. Typical forms of social norms one might encounter in meetings are that one should not yell or scream, that one should let people finish talking, that one should not start private conversations, that one should not whisper and that one should not engage in ‘Ad Hominem’ arguments. These social norms or conventions define the shared belief of what is normal and acceptable and hence constrain people’s actions. Other restrictions are described by Tracy and Coupland [51], stating that during a conversation a balance should be maintained between various levels of communication. An example of such a balance for a participant is the one between the urge to immediately achieve one’s agenda or objective (task goal) on the one hand and to act in line with social norms and roles (face goal) on the other hand. The intentions of the exhibited behavior are, amongst other things, a combination of the social constraints, the individual agenda of the participants and the amount of effort they are willing to put into realizing a set of predefined goals. To explore some of these behavioral characteristics, one could, for example, analyze its frequency. Simple (possibly automatic) counting of occurrences could suffice in order to get some first impressions. The problem lies in the automatic detection of an observation. How does one know that a specific observation occurred and which sensors are required?

The Human Media Interaction Group at Twente University (HMI) has a long tradition of work in automatic observations of behavioral aspects. Currently HMI is taking part in a European 6th Framework program called Augmented Multiparty Interaction (AMI), which is concerned with research on multimodal interaction, and, as the name suggests, multimodal interaction in a multiparty context. The AMI project concentrates on multi-party interaction during meetings. The main aims of the AMI project are to develop technologies for the disclosure of meeting content and to provide online support for (possibly remote) meetings. Our work consists, amongst other things, of: automatic body pose estimation [41], automatic dominance detection [42], addressee detection [23], emotion analysis [20] and analysis of argumentation patterns of meeting discussions [43]. All of these areas represent open problems that are far from being solved and completely understood. On the other hand the results that are and will be achieved are without doubt beneficial for all meeting assistants that will come to be developed in the future.

The next section describes part of the process of how to develop an actual meeting assistant.

More specifically, it describes an effort to develop formal requirements for a conflict managing meeting assistant, which could work in conjunction with a meeting chairman.

3 REQUIREMENTS FOR A CONFLICT MANAGING MEETING ASSISTANT

This section describes an excerpt of the work from Kernkamp [25], explaining how a conflict managing meeting assistant (CMMA) can be specified in terms of functional requirements. How should such an assistant work? What information does it need from the meeting, what does it derive from this information and what specific actions should it undertake in order to avoid conflicts?

It is well known that conflicts between participants may disrupt meetings. In order to have an efficient meeting it is therefore usually better to avoid them. We foresee a CMMA as a system able to observe a meeting, for instance through cameras and microphones as in the AMI project, and able to detect emerging conflicts (i.e. escalating debates). The CMMA, should in the end be able to act in order to prevent conflicts from escalating, e.g. through interruption or intervention by means of actuators such as a voice-over.

An important step in any systems' development phase is the process of requirements engineering, which results in a (formal) specification explaining what the system should do. With this document programmers and developers can start to build the software. The specification of requirements is recognized as one of the important and difficult areas of systems development [27]. There has already been some research on user-requirements for a meeting assistant. Tucker et al. [53], for example, proposes to specify user requirements by means of use cases in the context of a remote meeting assistant. The process of requirements engineering has, however, never been tackled in the context of a CMMA.

Our goal here is therefore threefold: We needed to devise an appropriate specification technique, to identify obstructing conflict types and to develop a set of corresponding resolution strategies.

3.1 WHAT ARE MEETING CONFLICTS?

A conflict is an open clash between two or more opposing groups (or individuals) aiming for different objectives, or adhering to different norms. This incompatibility is the root cause of the conflict. These clashes generally escalate beyond the traditional 'debate' and may eventually result in deleterious open antagonism.

Note that not all conflicts are bad and some conflicts should not be avoided [13]. If a conflict takes a constructive course, the conflict can potentially be of considerable personal and social value as it is sometimes necessary to strengthen the relations among team members and to improve productivity when it takes the form of a lively argument. It could prevent stagnation, and might stimulate interest and curiosity. Although conflicts are known to be associated with disruption, violence and civil disorder, the negative connotation is not always appropriate. Hence, not all conflicts within organizational meetings should be suppressed. If, on the other hand, conflict seems to harm cooperation and productivity among the members of a team and the conflict takes a destructive course as in a bitter disagreement or when expected to develop over a long period of time, there is room for our conflict managing meeting assistant to intervene.

Duncan [15] has made an overview of various categories in which conflicts emerge. These categories are depicted in Figure 1. This diagram shows the sources of conflict, their types, and how they relate to the physical/psychological and 'in fact'/'in principle' aspects. A source of conflict could be for example a bad relationship between some of the meeting participants. This is a psychological and subjective conflict. An example of a subjective and 'in principle' conflict is for instance when someone's values or beliefs clash with other participants' values or beliefs.

3.2 TECHNIQUES TO AVOID CONFLICTS.

In order to avoid or to limit conflict, it is convenient to have a set of conflict avoidance rules prepared, ready to be applied whenever a conflict emerges or seems about to emerge. These rules legitimize some kinds of behavior and ascertain as a consequence what is, and what is not allowed

during the meeting. If an assistant that can apply the proper rules at the appropriate time can be created, meetings are likely to become more efficient. A simple example of such a rule is to establish that if off-topic discussions (e.g. because of conflict over resources) are detected, it could be useful to make a remark that will encourage the participants to get back on track. More sophisticated rules could be based on the facilitation techniques described in Paulsen [40], which elaborate on how a facilitator creates a safe group environment where people are free to disagree with each other.

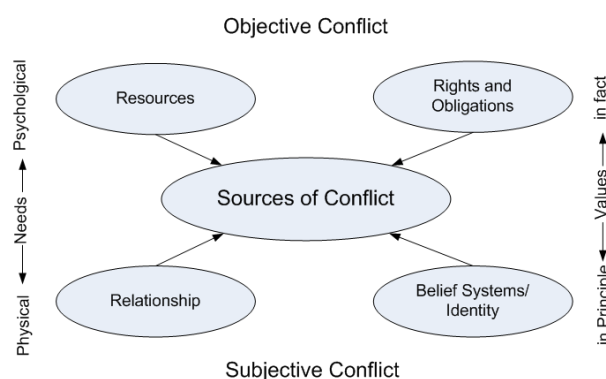


Figure 1: Sources of Conflict

Here we concern ourselves with a set of rules that can be applied to the types elaborated by Sellen et al. [48], hereby excluding rules that physically interfere in the meeting environment. Sellen's model distinguishes the following types of rules:

- Boundary rules (BR); define who is and who is not in the group and can detail the permeability of the group, i.e. whether members can easily enter or exit. It determines the extent to which norms developed within the group can be maintained and shared and the extent to which groups can impose sanctions.
- Aggregation rules (AR); define how a group reaches a collective decision. This can be by majority, unanimity or 'anyone' rule. The unanimity rule is dangerous because if one person objects it can take a lot of time before an acceptable choice is negotiated. The 'anyone' rule means any actor can impose a group choice.
- Position rules (PR); define who can act at any point, so define to a great extent authority. Some positions have higher authority than other (lower) positions.
- Information rules (IR); describe how information is shared and what each actor can know, for instance, whether a member can know what other members have done (or what they are planning to do).

One could reduce conflicts by applying the position rule. This could take the form of the appointment of a chairman with power to take measures against conflicts. Chairmanship is associated with different rights, privileges, and responsibilities that characterize leadership positions within an organizational structure. If these leadership positions or particular rules about positions are in place, conflict is less likely to arise.

Apart from defining a list of possible rules that can be applied, there are several modes or ways in which the rules can be enforced. Ruble and Thomas [47], for instance list the following: *competing, collaborating, compromising, avoiding and accommodating*.

3.3 EXAMINING REAL MEETING CONFLICTS

To get an idea about the most frequently occurring destructive conflicts ten student committee meetings were observed. An overview of the most destructive conflicts is given in Table 1. For all of these, possible conflict resolution strategies have been created using the four types of rules from Sellen presented above. Possible options are listed, together with their type in Table 1. It should be noted that this conflict list is solely based on the ten attended meetings and that there might be other conflicts having a much greater impact in other meetings. A second point to be aware of is the fact that many other possible resolution rules could be thought of. The exact rule as well as its desired mode of execution relate to the eventual ‘personality’ of the CMMA.

#	Conflict	Rule
9	Not all participants agree on the costs of an item.	It should be clear how this choice is to be made and how a decision can be enforced. This could either be done by using the majority, unanimity or anyone rule. In many cases the majority rule should be applied so the decision does not take too much time. When the costs are very high, the choice should be made using the unanimity rule. (AR)
7	Too many off-topic personal matters are discussed, annoying some of the participants.	During the meeting it should be clear which information is to be discussed. If everybody wants to chitchat, the meeting is to be adjourned for 10 minutes before being resumed or the discussion should be stopped through intervention. (IR & PR)
6	Insufficient information is available causing a lot of superfluous discussions.	The information should be available in the next meeting, and the decision is to be postponed. (IR)
6	Not everybody is equally talkative, or some persons are neglected.	If this person is expected to make valuable contributions, he should be invited to give his/her opinion. (BR)
5	People with a relatively high authority and much experience neglect the chairman and keep on arguing with each other.	If this quarrel takes up too much time without being useful, an intervention should assure that all participants stick to the current agenda. The persons should be requested to catch up and told to talk about this outside the meeting, the current status should be summarized and then it is to be explained what is to be discussed. (PR)
4	Personal habits and preferences differ between participants and they argue because of this.	Quarrels between participants about personal preferences and habits should be excluded from the meeting. Sanctions can be considered if the arguers do not comply. (BR)

Table 1: Meeting conflicts and possible corresponding resolution rules

AR = Aggregation Rule, BR = Boundary Rule, PR = Position Rule, IR = Information Rule

3.4 CREATING COMPLETE TASK DESCRIPTIONS

As a method to further specify the resolution strategies into functional requirements, task descriptions were created. Task descriptions are domain level requirements, listing what the user and the system should achieve together [27, 45]. They consist of structured texts that are easy to

understand for the users as well as for the developers. Domain level requirements are useful for defining what is required, without describing how the product is going to look. In this case they focus on the tasks that are to be supported to avoid destructive conflicts.

Table 3 lists the task descriptions for the conflicts emerging when participants do not agree on the costs of an item and when superfluous discussions emerge due to lack of information.

Task	Enforcing a Decision	Controlling a discussion that started while lacking information
Purpose	To make a choice using the unanimity, anyone or majority rule	To shorten the meeting time, people continue trying to find the required information otherwise.
Trigger or Pre-condition	A decision is to be made about something	A discussion is started because of little or no information available
Frequency	Average: 1 decision every 15 minutes	Average: 1 discussion due to insufficient information per meeting
Critical	More than 10 participants in a meeting	More than 3 discussions due to insufficient information
Subtask	- Define the decision. For variant 1a and 1b: - Switch to collaborative handling mode For variant 1c: - Switch to accommodation handling mode- Fill in which choice is made	For variant 1a: - Switch to competing handling mode - Make a list of what information should be available the next meeting - Assign this task to the specific person(s) - Go to next subject For variant 1b: - Switch to compromising handling mode - Talk to the person responsible for collection the specific information or the person with specific knowledge about this topic and ask whats the quickest way to get the information - Assign the person who has to get the information - Continue with the meeting if possible until the information is available
Variants:		
1a	The decision has to be made using the unanimity rule	The subject can be discussed during the next meeting.
1b	The decision has to be made using the majority rule	The subject is to be handled during this meeting.
1c	The decision has to be made using the anyone rule	The topic is important
1d		The topic is not really important

Table 2: Task descriptions for the conflicts arising when *Not all participant agree on the costs of an item* and when *superfluous discussions emerge due to lack of information*.

The first task *Enforcing a Decision* has a subtask that chooses the best conflict handling mode (See e.g. Ruble and Thomas [47]). The collaborative mode is best for a majority or unanimity decision, because when the decision is made, the chairman and the other participants must be heard in order to satisfy both own and other concerns. With the anyone rule, the accommodation mode should be applied in order to satisfy only the experts' concerns. The second task deals with an emerging discussion for which insufficient information is available. Here, the compromise handling mode should be applied if the discussion about the subject is to be concluded during the meeting. If the subject can be talked about during the next meeting the competing mode should

be applied, as one aims for a structured and efficient meeting.

3.5 PUTTING THE CMMA TOGETHER

If properly constructed, the use of a conflict managing meeting assistant sounds promising and has much potential. Although it obviously takes much more to develop a meeting assistant than what has been so far discussed, we have shown a first important step in its production process. Successful implementation will depend upon reliable detection of the required input parameters, as well as efficient implementation of the resolution interventions.

With respect to the CMMA, some specific remarks can be made. First and foremost, all meetings and people are different, resulting in more than just one solution to avoid negative conflicts. The only thing the CMMA is designed for is to suggest its best known solution for a detected conflict. A second point is that in order to make a clearer distinction between what the system and what the user should do, the task description defined above can be split up into use case descriptions. By doing so a clearer view of the role of the user and the role of the system is created, potentially providing more insights in the design phase. A final point we address here is about quality requirements with respect to the usability of the system. Any meeting assistant, including the CMMA should in the end be easy and straightforward to use. Besides, interacting with such an assistant should not take too much time.

4 PUTTING LIVE MEETING ASSISTANCE TO THE TEST: DOES IT WORK?

In work from DiMicco [14] a system called Second Messenger is described that shows real-time text summaries of participants contributions. After increasing the visibility of the less frequently speaking group members, it appeared that these started to speak more frequently than before, whereas the more dominant people started to speak 15% less. This shows that it is possible to build systems that are able to influence the meeting process. This section describes a summary of experiments investigating whether and in what form meeting assistants aiming at improving meeting effectiveness can work in practice. See Kuperus [26] and Broenink [7] for the complete versions. The experiments try to find out how the appearance and associated actions of a meeting assistant can influence the interaction with participants. How should the assistant intervene and act in order to be obeyed and listened to? Another goal of the experiments is to investigate if it is really true that meetings can become more efficient when assistants are applied. The assistant used in all experiments is simulated using a Wizard of Oz technique. This means that the meeting participants are led to believe that they are interacting with an autonomous system, when in fact a human being controls the behavior of the system remotely. This approach is highly suitable, as the implementation of a complete assistant is not easily feasible and it is expected that a good Wizard of Oz experiment will have nearly identical results.

4.1 THE RESEARCH SETTING

The research setting consisted of a meeting ‘control center’ with a computer outside the room where meetings took place. Two monitors were connected to this computer. One of them was inside the room, one outside. The monitor outside the room is used to control the interaction with the meeting and displays the live video footage of the meeting coming from a DV camera located on the ceiling. A microphone is used to capture the audio inside the meeting room. Interaction with the meeting takes place through the monitor and a speaker set inside the meeting room. The ‘control center’ is placed on a moveable cart as meetings take place at various locations. The monitor and a speaker set are able to transmit information to the participants. The consistency of the experiment was guaranteed by the creation of a script which the agent followed.

4.2 THE EXPERIMENTS

As a preliminary investigation to find out which aspects of the meeting were considered useful to influence, questionnaires were issued to 15 different chairmen; 9 were fully completed and returned. Especially off-topic, balance and time indications were pointed out as useful. The chairmen also

expected that information presented on a display would be more beneficial than voiced information. The screen was expected to be less intrusive than the voice-over. (See figure 2(a).)

Using this information, a set of four different systems with varying intrusiveness levels was devised for the experiment. Table 3 shows descriptions of the systems ranked from least to most intrusive according to the perceptions expressed in the questionnaires. Two student committees (of eight and seven members respectively) were subsequently exposed to all four versions of the system over a period of four weeks. Before each meeting we asked the participants to provide the agenda including: an expected time-line; the names of participants; and the chairs they would occupy during the meeting. After each meeting questionnaires were issued in order to discover how the assistant and its actions were received by the meeting participants. Participants were asked, amongst other things, to rate on a seven point scale their perception of the meetings' efficiency, the meeting being off-topic, the meeting being balanced and the system's enjoyableness and intrusiveness. A control group of three more committees was used to correct for 'learning effects' that could occur as people get used to the system.

System	Description
1	Displays messages on a screen when an item is due to be finished in five, two or zero minutes. Also displays messages when something is off topic, a subject takes too long or when a discussion is unbalanced.
2	Similar to system 1, but instead of displaying messages, continuously displaying a clock.
3	Similar to system 1, but instead of displaying messages, voice samples were played.
4	A combination of system 1, 2 and 3. The clock is displayed and messages can be sent either to the screen or played as a voice sample.


Table 3: Description of the systems simulated for the experiment .

4.3 SOME FINDINGS AND RESULTS


To verify whether meeting assisting agents can benefit the meeting process we compared the predefined given agendas with the actual agendas of the meetings of the various systems. The results averaged for the two groups are shown in figure 2(b). It appears that when no system is used at all, the meetings lasted on average 57% longer than what had actually been planned. With System 3 we reached an optimum, shortening the meeting by 27%. Although chairmen might have improved their planning capabilities in the meantime, they were not informed about any of the results.

When we look at the participants' ratings of degree of intrusiveness versus efficiency, figure 2(d)) shows that the added intrusiveness of System 3 pays off in terms of meeting efficiency. Notable is the fact that the perceived efficiency appears to be in-line with the actual efficiency. System 3 also resulted in a slight disturbance increase, whereas its "enjoyability" is rated much lower than Systems 1 and 2 (See figure 2(c)).


After every session the chairmen were asked again to give their opinion about the disturbance and efficiency for both the voice as well as the screen feedback strategies. It appeared that in contrast to the pre-meeting questionnaire results, they now rated them equally for efficiency. Voice messages were still found more intrusive than the text messages, though. An interesting side result was that when the system uses voiced feedback, the participants of the meeting appeared to be much more aware of their own behavior. When they tended to go off topic for example they corrected themselves very quickly, sometimes saying: "off-topic" before continuing with the current item on the agenda. This is probably due to the fact that the system can speak directly to specific participants; participants would therefore try to prevent being corrected by the system. After getting used to a system with voice output, participants did notice and use the information, but did not interrupt their talking. It should be noted that although the above findings speak in favor of a system that assists the meeting process, a lot of additional research is required, for instance by examining a larger number of groups over a longer period of time.




(a) Part of the Chairmen Questionnaire results.



(b) Efficiency of the meetings



(c) Enjoyment vs. Disturbance as rated by the participants



(d) Intrusiveness vs. Efficiency as rated by the participants

Figure 2: Some results of the Wizard of Oz experiment

5 CONCLUSION

We have shown that there is potential for ambient intelligent systems that aid the meeting process. We have discussed a wide variety of possible applications and application areas. A concrete example of how requirements for a conflict management meeting assistant can be developed has been given. We have shown that the results of an experiment utilizing multiple system paradigms of varying degree of intrusiveness; the experiments employed a Wizard of Oz technique. The results show that meeting efficiency can be improved with respect to a baseline in which no meeting assistants are employed.

6 ACKNOWLEDGEMENTS

This work was partly supported by the European Union 6th FWP IST Integrated Project AMI (Augmented Multi-party Interaction, FP6-506811, publication AMI-154). Some material is based upon work supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. NBCHD030010. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the DARPA or the Department of Interior-National Business Center (DOINBC). We thank Gerben Broenink, Johan Kuperus and Reinier Kernkamp for their efforts.

REFERENCES

- [1] Antunes, P. and Carrio, L. (2003). Modeling the information structures of meetingware. In *Proc. of Workshop de Sistemas de Informao Multimidia e Cooperativos (COOP-MEDIA'03)*.
- [2] Banerjee, S., Rose, C., and Rudnicky, A. I. (2005). The necessity of a meeting recording and playback system, and the benefit of topic-level annotations to meeting browsing. In *Proceedings of the Tenth International Conference on Human-Computer Interaction*.
- [3] Barthelmess, P. and Ellis, C. A. (2005). The neem platform: An evolvable framework for perceptual collaborative applications. *Journal of Intelligent Information Systems*, 25(2):207–240.
- [4] Barthelmess, P., Kaiser, E., Huang, X., and Demirdjian, D. (2005). Distributed pointing for multimodal collaboration over sketched diagrams. In *International Conference on Multimodal Interfaces - ICMI*.
- [5] Berry, P., Gervasio, M., Uribe, T., Pollack, M., and Moffitt, M. (2005). A personalized time management assistant: Research directions. In Saphiro, D., editor, *Persistent Assistants: Living and working with AI, workshop at the AAAI Spring Symposium 2005*. AAAI Press.
- [6] Bowring, E., Tambe, M., and Yokoo, M. (2005). Optimize my schedule but keep it flexible: Distributed multi-criteria coordination for personal assistants. In Saphiro, D., editor, *Persistent Assistants: Living and working with AI, workshop at the AAAI Spring Symposium 2005*. AAAI Press.
- [7] Broenink, E. (2006). How a computer actor influences the time-efficiency of a meeting. In *Proceedings of the 4th Twente Student Conference on IT*. Twente University Press.
- [CALO Website] CALO Website. <http://www.ai.sri.com/project/calo/>.
- [8] Chen, H., Finin, T., and Joshi, A. (2004). A context broker for building smart meeting rooms. In *Proceedings of the Knowledge Representation and Ontology for Autonomous Systems Symposium, (AAAI Spring Symposium)*.
- [9] Chen, H. and Perich, F. (2004). Intelligent agents meet semantic web in a smart meeting room. In *Proceedings of the Third International Joint Conference on Autonomous Agents & Multi Agent Systems (AAMAS 2004)*.
- [10] Cook, P., Ellis, C. S., Graf, M., Rein, G., and Smith, T. (1987). Project Nick: meetings augmentation and analysis. *ACM Trans. Inf. Syst.*, 5(2):132–146.
- [11] Danninger, M., Flaherty, G., Bernardin, K., Ekenel, H., Khler, T., Malkin, R., Stiefelhagen, R., and Waibel, A. (2005). The connector - facilitating context-aware communication. In *Proceedings of the International Conference on Multimodal Interfaces*, Trento, Italy.
- [12] De Vreede, G., Vogel, D., Kolfshoten, G., and Wien, J. (2003). Fifteen years of GSS in the field: A comparison across time and national boundaries. In *Proceedings of the 36th Hawaii International Conference on System Sciences*. IEEE Press.
- [13] Deutsch, M. (2003). *International Handbook of Organizational Teamwork and Cooperative Working*, chapter Cooperation and conflict: A personal perspective on the history of the social psychological study of conflict resolution, pages 9–44. Wiley & Sons.
- [14] DiMicco, J. (2004). Designing interfaces that influence group processes. In *Doctoral Consortium Proceedings of the Conference on Human Factors in Computer Systems (CHI 2004)*.
- [15] Duncan, M. (1996). Effective meeting facilitation: Sample forms, tools, and checklists. <http://www.nea.gov/resources/Lessons/>.

- [16] Ellis, C. and Barthelmeß, P. (2003). The neem dream. In *Proceedings of the 2003 conference on Diversity in computing*, pages 23–29. ACM Press.
- [17] Erol, B., Lee, D., and Hull, J. (2003). Multimodal summarization of meeting recordings. In *Proceedings of the IEEE International Conference on Multimedia and Expo (ICME 2003)*.
- [18] Galley, M., McKeown, K., Hirshberg, J., and Shriberg, E. (2004). Identifying agreement and disagreement in conversational speech: Use of bayesian networks to model pragmatic dependencies. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 669–676.
- [19] Garrido, L. and Sycara, K. (1995). Multi-agent meeting scheduling: Preliminary experimental results. In Lesser, V., editor, *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95)*. The MIT Press: Cambridge, MA, USA.
- [20] Heylen, D., Nijholt, A., and Reidsma, D. (2006). Determining what people feel and think when interacting with humans and machines: Notes on corpus collection and annotation. In *Proceedings 1st California Conference on Recent Advances in Engineering Mechanics (to appear), California State University, Fullerton, CA. U.S.A.*
- [IM2 Website] IM2 Website. <http://www.im2.ch/>.
- [21] Jaimes, A., Omura, K., Nagamine, T., and Hirata, K. (2004). Memory cues for meeting video retrieval. In *CARPE'04: Proceedings of the the 1st ACM workshop on Continuous archival and retrieval of personal experiences*, pages 74–85. ACM Press.
- [22] Jebara, T., Ivanov, Y., Rahimi, A., and Pentland, A. (2000). Tracking conversational context for machine mediation of human discourse. In *In AAAI Fall 2000 Symposium - Socially Intelligent Agents - The Human in the Loop*.
- [23] Jovanovic, N., Op den Akker, R., and Nijholt, A. (2005). A corpus for studying addressing behavior in multi-party dialogues. In *Proc. of The sixth SigDial conference on Discourse and Dialogue*.
- [24] Kaiser, E., Demirdjian, D., Gruenstein, A., Li, X., Niekrasz, J., Wesson, M., and Kumar, S. (2004). Demo: A multimodal learning interface for sketch, speak and point creation of a schedule chart. In *Proceedings of the International Conference on Multimodal Interfaces - ICMI'04*, pages 329–330, State College, Pennsylvania, USA.
- [25] Kernkamp, R. (2006). Functional requirements for a conflict managing meeting assistant. In *in Proceedings of the 4th Twente Student Conference on IT*. Twente University Press.
- [26] Kuperus, J. (2006). The effect of agents on meetings. In *Proceedings of the 4th Twente Student Conference on IT*. Twente University Press.
- [27] Lauesen, S. (2002). *Software Requirements - Styles and Techniques*. Addison Wesley.
- [28] Lisowska, A. (2003). Multimodal interface design for the multimodal meeting domain: Preliminary indications from a query analysis study. Technical report, ISSCO/TIM/ETI, Universit de Geneve, Switzerland. IM2.MDM Report 11.
- [29] Misiólek, N. and Heckman, R. (2005). Patterns of emergent leadership in virtual teams. In *Proceedings of the Thirty-Eighth Annual Hawaii International Conference on System Sciences*.
- [30] Moran, T., Palen, L., Harrison, S., Chiu, P., Kimber, D., Minneman, S., Van Melle, W., and Zellweger, P. (1997). I'll get that off the audio : a case study of salvaging multimedia meeting records. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 202–209. ACM Press.

- [31] Nakanishi, H., Ishida, T., Ibister, K., and Nass, C. (2004). *Agent Culture: Human-Agent Interaction in a Multicultural World*, chapter 11, Designing a Social Agent for Virtual Meeting Space, pages 245–266. Lawrence Erlbaum Associates, Publishers.
- [Nectar Website] Nectar Website. <http://www.nectar-research.net>.
- [32] Niekrasz, J. and Purver, M. (2005). A multimodal discourse ontology for meeting understanding. In Bourlard, H. and Bengio, S., editors, *Proceedings of MLMI'05*. LNCS.
- [33] Nijholt, A., op den Akker, H., and Heylen, D. (2004). Meetings and meeting modeling in smart surroundings. In Nijholt, A. and Nishida, T., editors, *Social Intelligence Design. Proceedings third international workshop*, volume 02, pages 145–158. CTIT. ISBN=ISBN 90-75296-12-6.
- [34] Nijholt, A., Zwiers, J., and Peciva, J. (2005). The distributed virtual meetingroom exersize. In Vinciarelli, A. and Odobez, J., editors, *Multimodal Multiparty Meeting Processing, Workshop at the 7th International Conference on Multimodal Interfaces (ICMI)*, pages 93–99.
- [35] Nunamaker Jr., J., Briggs, R., and Mittleman, D. (1995). Electronic meeting systems: Ten years of lessons learned. In Coleman, D. and Khanna, R., editors, *Groupware: Technology and Applications*. Prentice Hall, Englewood Cliffs, NJ.
- [36] Oh, A., Tuchinda, R., and Wu, L. (2001). Meetingmanager: A collaborative tool in the intelligent room. In *Proc. of the MIT Student Oxygen workshop*.
- [37] Oh, J. and Smith, F. (2005). Calendar assistants that learn preferences. In Saphiro, D., editor, *Persistant Assistants: Living and working with AI, workshop at the AAAI Spring Symposium 2005*. AAAI Press.
- [38] Padilha, E. and Carletta, J. (2003). Nonverbal behaviours improving a simulation of small group discussion. In *Proc. 1st Nordic Symp. on Multimodal Comm.*, pages 93–105.
- [39] Palotta, V., Ghorbel, H., Ballim, A., Lisowska, A., and Marchand-Maillet, S. (2004). Towards meeting information systems. In Seruca, E., Filipe, J., Hammoudi, S., and Cordeiro, J., editors, *Proceedings of the ICEIS 2004: 6th International conference on enterprise information systems*, volume 4.
- [40] Paulsen, D. (2004). Leadership essentials: facilitation skills for improving group effectiveness. In *SIGUCCS '04: Proceedings of the 32nd annual ACM SIGUCCS conference on User services*, pages 153–160, New York, NY, USA. ACM Press.
- [41] Poppe, R., Heylen, D., Nijholt, A., and Poel, M. (2005). Towards real-time body pose estimation for presenters in meeting environments. In *Proceedings of the International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision'2005 (WSCG'2005)*, Plzen, Czech Republic. Science Press.
- [42] Rienks, R. and Heylen, D. (2005). Automatic dominance detection in meetings using easily detectable features. In *2nd Joint Workshop on Multimodal Interaction and Related Machine Learning Algorithms*, Endinburgh, Scotland. Springer Verlag. to appear.
- [43] Rienks, R., Heylen, D., and Van der Weijden, E. (2005a). Argument diagramming of meeting conversations. In Vinciarelli, A. and Odobez, J., editors, *Multimodal Multiparty Meeting Processing, Workshop at the 7th International Conference on Multimodal Interfaces*, pages 85–92, Trento, Italy.
- [44] Rienks, R., Nijholt, A., and Reidsma, D. (2005b). *Meetings and Meeting support in ambient intelligence*, chapter 18, in *Ambient Intelligence, Wireless Networking, Ubiquitous Computing*. Artech House, Norwood, MA, USA. In Press.

- [45] Robertson, S. and Robertson, J. (1999). *Mastering the Requirements Process*. Addison Wesley.
- [46] Romano Jr., N. and Nunamaker Jr., J. (2001). Meeting analysis: Findings from research and practice. In *in Proceedings of the 34th Hawaii International Conference on System Sciences*. IEEE Press.
- [47] Ruble, T. and Thomas, K. (1976). Support for a two-dimensional model of conflict behavior. *Organizational Behaviour and Human Performance*, 16:143–155.
- [48] Sellen, J., Lovaglia, M., Mannix, E., Samuelson, C., and Wilson, R. (2004). Investigating conflict, power and status within and among groups. *Small Group Research*, 35(1):44–72.
- [49] Shum, S. (1997). Negotiating the construction and reconstruction of organisational memories. *Journal of Universal Computer Science*, 3(8):899–928.
- [50] Sudweeks, F. and Simoff, S. (2005). Leading conversations: Communication behaviours of emergent leaders in virtual teams. In *Proceedings of the Thirty-Eighth Annual Hawaii International Conference on System Sciences*.
- [51] Tracy, L. and Coupland, N. (1990). Multiple goals in discourse: An overview of issues. *Journal of Language and Social Psychology*, 9:1–13.
- [52] Tucker, S. and Whittaker, S. (2005). Reviewing multimedia meeting recordings: Current approaches. In *in Multimodal Multiparty Meeting Processing, Workshop at the 7th International Conference on Multimodal Interfaces*.
- [53] Tucker, S., Whittaker, S., and Laban, R. (2005). Identifying user requirements for novel interaction capture. In *Symposium 'annotating and measuring meeting behaviour' at Measuring Behaviour*.
- [54] Turoff, M. and Hiltz, S. (1977). Meeting through your computer. *IEEE Spectrum*, (May):58–64.
- [55] Waibel, A., Steusloff, H., and Stiefelwagen, R. (2004). Chil - computers in the human interaction loop. In *NIST ICASSP Meeting Recognition Workshop*, Montreal, Canada.
- [56] Wainer, J. and Braga, D. (2001). Symgroup: applying social agents in a group interaction system. In *Proceedings of the 2001 International ACM SIGGROUP Conference on Supporting Group Work*, pages 224–231. ACM Press.