An Interactive Table for Supporting Participation Balance in Face-to-Face Collaboration

Khaled Bachour, Frédéric Kaplan and Pierre Dillenbourg
Swiss Federal Institute of Technology
1015 Lausanne, Switzerland
{khaled.bachour, frederic.kaplan, pierre.dillenbourg}@epfl.ch

ABSTRACT

In this paper we describe a novel interactive table designed for supporting group collaboration. The table, *Reflect*, addresses the issue of unbalanced participation during group discussions. By displaying on its surface a shared visualization of member participation, *Reflect* is meant to encourage participants to avoid the extremes of over- and under-participation. We report on a user study that validates some of our hypotheses on the effect the table would have on its users. Namely we show that *Reflect* leads to more balanced collaboration, but only under certain conditions. We also show different effects the table has on over- and under-participators.

Author Keywords

Computer-supported Collaborative Work, Computer-Supported Collaborative Learning, Interactive Furniture, Ubiquitous Computing

ACM Classification Keywords

C.3.h Special-Purpose and Application-Based Systems: Ubiquitous Computing; K.3.1.a Computers and Education: Collaborative Learning; K.4.3.b Computers and Society: Computer-Supported Collaborative Work

INTRODUCTION

In situations of face-to-face collaboration, unbalanced participation often leads to undesirable results. These results take the form of suboptimal decisions taken because of poor information sharing or lower learning outcomes for members of a group that did not participate in the group process [9, 13, 14, 15]. One way to overcome this effect is by encouraging members of a group to participate in a more balanced manner. We attempt to achieve this by indicating to individual members their level of participation on a shared display. We embed this display in an interactive table, seen in Figure. 1 that allows users to interact with each other in as natural a manner as possible while giving them feedback on their behavior. This semi-ambient display has the properties of both being in the background of the collaboration process while at the same time remaining visible in a central position of the shared workspace.

Ubicomp 2009, Sep 30 – Oct 3, 2009, Orlando, Florida, USA.

Figure 1. The current design of *Reflect* with color-coded circles around each speaker position indicating how much the person has spoken.

RELATED WORK

Providing feedback to speakers about their participation patterns is not new. Most prominently, DiMicco et al. have studied extensively the effect of such visualizations on speaker behavior. [12, 11] They have studied both the effects of having this information displayed in real-time as the conversation takes place and of having this information displayed between meetings as a replay tool. Their system, *Second Messenger*, showed promising results for this kind of display. The replay tool had a significant effect on speaker behavior after it was displayed. Over-participators spoke less and under-participators spoke more. This desired effect was not completely achieved when only the real-time tool was used. By displaying information in real-time, *Second Messenger* pushed over-participators to reduce their levels of participation but the effect was not as strong for under-participators.

Other researchers have also studied the effects of these visualizations. Bergstrom and Karahalios implemented two systems, the *Conversation Clock* [4, 6] and *Conversation Votes* [5]. In both these systems, a visualization representing the current conversation is projected onto some shared surface. The *Clock* shows which member of the group spoke at each time and allows the users to get a snapshot of the conversa-

tion history every time they look at the surface. *Conversation Votes* goes further and allows members of the group to anonymously "vote" indicating to the table whether or not they agree with what is being said. This information is visualized onto the table along with the speaking patterns of the users. The authors reported varying reactions to the visualizations especially in terms of reactions to long-term and short-term history, as well as changes in behavior among above and below average speakers.

Our work follows a similar approach of displaying information about speaker participation to group members. Our originality comes from achieving similar benefits in terms of balanced speakers while retaining as much of the natural behavior of group members as possible. The display becomes embedded into everyday furniture, and with the use of directional microphone arrays, we eliminate the need for lapel microphones or headsets. The result is a regular table, augmented with a semi-ambient real-time feedback of a conversation taking place around it.

This notion of embedding computing functionality in real-world objects is developing as a research trend on its own. The name "roomware" has been given to this type of device and has been described as an "umbrella" framework for four fields: ubiquitous computing, computer-supported collaborative work, augmented reality, and architecture [19]. Since then, countless new devices have been developed that satisfy the criterion of roomware: real-world objects with embedded computing. Lamps, clocks, tables, walls and floors have been augmented with computational functionality ranging from simple single-purpose devices such as a clock to improve location-awareness for family members [8] to elaborate multi-purpose table surfaces for supporting collaboration [10].

Our own work thus falls within the realm of *roomware* with the specific purpose of augmenting collaborative spaces by embedding within the physical table a tool that helps increase awareness of member participation.

MOTIVATION

Most participation in face-to-face meetings takes the form of verbal communication in that members who are silent are seen as not participating or not contributing to the meeting. In decision-making situations, verbal communication is used to offer opinions, present arguments and share relevant information, all of which are fundamental for good decisions to be reached. In the context of learning, most mechanisms that explain the effects of collaborative learning emphasize the role of verbalization: elaborating explanations [20], building argumentation [3], mutual regulation [7] and conflict resolution [21].

In this section we discuss how unbalanced participation in group meetings has a tendency to reduce the effectiveness of the meeting, either in terms of the quality of the decision made, or in terms of the learning outcomes for individual participants.

Participation Balance and Group Learning

Whether or not they are required to do so, students often find themselves working together in groups. Empirical research has shown that collaborative learning can be more effective than individual learning [17]. However, this is not always the case.

Cohen [9] describes some criteria for group productivity, without which group learners might benefit less than individual learners. Among these, lack of equity in participation is presented as an obstacle to effective learning in a group. Cohen argues that participation is a predictor of learning gains such that the more individual members participate within a group, the more they learn.

Cohen also suggests that the difference in participation is not necessarily related to participants' abilities or their expertise, but rather it is related to their perceived status which can come from any number of stimuli including age, gender, or race of the participant. In some cases, perceived popularity or attractiveness of individuals can lead to more active participation on their part, which in turn leads to lower learning gains for their partners [9, 15].

Unbalanced participation in group learning can thus be seen as a deterrent for effective learning. There is a need then to encourage members to participate in a more balanced manner.

Participation Balance and Decision Making

When decisions are made in group meetings, there is often a substantial risk that one or more participants who hold critical information are unable to effectively share this information [13]. Proper information sharing is thus a crucial aspect of effective decision making. In reality, however, the variety and number of participants who do in fact contribute to the decision-making process is often less than is deemed appropriate by post-hoc analysis [14]. As a result, decisions are made with some relevant and potentially critical information missing, leading to suboptimal results. This could be avoided if group memberss were encouraged to participate in a more balanced manner, permitting all members to contribute. Balanced participation, however, does not guarantee that the information is better shared, as members who would otherwise remain silent, might not use their participation to provide meaningful information. It is in that sense a necessary but not sufficient condition.

Pilot Study

To further illustrate the need for balancing participation, we present a small case study conducted with eight subjects divided into two groups of four. We gave the subjects a task in which they were asked to rank, individually at first and then in group, a list of fifteen objects in order of their importance for survival in the desert. They were given 10 minutes to complete the task individually. They were then asked to discuss the problem for 30 minutes and come up with a single ranking that they all agree upon. This type of task, known as a choice shift task, is used to determine the influence each member of the group had on the group's final decision, by

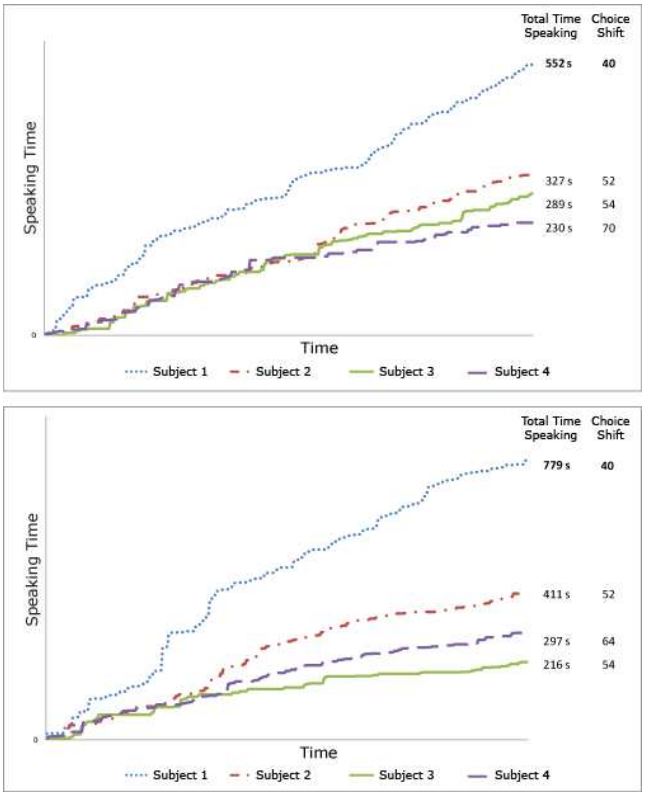

Figure 2. Participation of group members in the choice shift task and their influence on the final group decision.

comparing the group decision with the initial individual decisions of each member. The choice shift for each member is the distance between their initial decision and the final group decision. The most influential member is the one who has the smallest choice shift. In this task, the choice shift was computed as the sum of the differences between the rank given by the individual and that of the group for each object. It ranges from 0 (identical to group decision) to 112 (opposite to group decision).

We measured the individual members' participation in terms of their total talking time during the group discussion phase. We compared that to the individual choice shift. In both groups, one member clearly dominated the discussion, and as can be seen from the choice shift measure in Figure 2, that member had the strongest influence on the final group decision. In other words, the final group decision was closest to that of the member who did the most talking. It is important to note that in both cases, the final group decision was a relatively poor one, according to experts, when compared to some of the original decisions made by individual members. Interestingly, in both situations most participants, including both of the dominating members, were not aware that the conversations they had were not balanced. Moreover, when asked, they were not able to determine which member did in fact dominate the meeting.

We drew two conclusions from the pilot study. The first is a confirmation that group decisions are sometimes strongly influenced by members of a group who participate more than others. The second is that it is not always obvious for members of a group who spoke more than the others, even when one speaker dominated the conversation significantly.

DESCRIPTION OF THE SYSTEM

Reflect is an interactive table designed to address the issue of unbalanced participation. We describe first the conceptual design in relation to its objective. We then detail the physical design which we will motivate with some constraints we imposed on the system.

Design Objectives

The original aim of *Reflect* was to function as a mirroring tool for collaborative groups. The term mirroring tool refers to the informative, rather than normative, nature of the system. Mirrors do not tell their users what they are doing right and what they are doing wrong. A bathroom mirror does not tell a user if their hair looks good or not. It simply shows them a reflection of what their hair looks like, and leaves it for the users themselves to decide what, if anything, needs to be changed. In the same manner, *Reflect* is not meant to judge the quality of the interaction, nor is it meant to actively pursue a more balanced collaboration on the part of its users. Its role in that respect is to inform the users of the current state of the conversation, and it is up to the users to decide what needs to be done. There are instances where one speaker is expected or even required to participate more than the others, for example if that speaker is the expert on the subject of discussion. Our system will thus remain neutral in terms of its judgment of the situation and its role will be strictly informative rather than normative. We cannot deny however that by making available information on participation levels we are potentially inducing an implicit norm among at least some members of the group that participation levels need to be monitored and therefore controlled.

Design Requirements

We required our interactive table to abide by certain principles that we found important for a system meant to follow the *disappearing computer* paradigm of ubiquitous computing. These requirements can be summarized as follows:

1. **Natural**: Regardless of its embedded functionality, the table was required to retain its initial purpose, namely serving as a table before being a display. Having a conversation or a work meeting around the table should involve minimal behavioral change from the natural use of a regular table. In other words, users should be able use the table in the same way they would a regular table, without having to worry about attaching peripherals or other accessories to their bodies. In addition, the table's surface should remain a working surface. Users should be able to place their notes, laptops and even their coffee mugs on the surface of the table.

2. **Unintrusive**: The table should not take too much attention away from the task the users are performing. The information it provides to its users is meant to be minimal and require very little cognitive effort to understand. It is

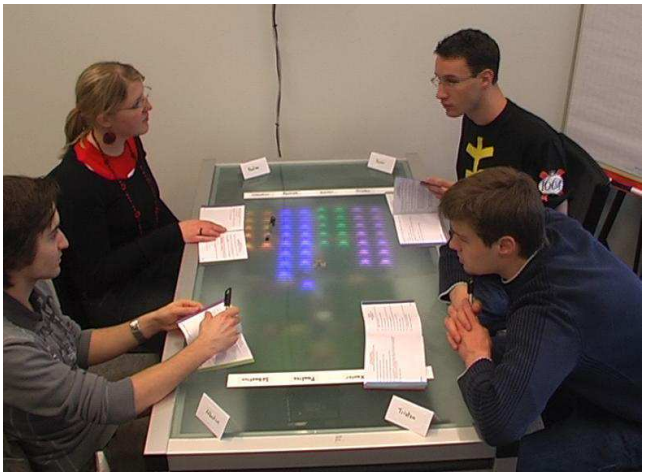

Figure 3. Four subjects taking part in the experiment in the speaker-based condition. Four labeled columns of LEDs can be seen on the table indicating to the users their participation levels.

thus important that the table not draw a lot of attention to itself and away from the real task at hand.

3. **Visible**: Despite the *unintrusive* criterion, the table must nonetheless be noticeable and should not be so discreet that it is ignored completely. The information should be prominently displayed, in a shared location and should be within the peripheral vision of the users, i.e. the part of the users' field of vision that is not the focus of their attention, but of which they have at least a minimal awareness.

Together the *unintrusive* and *visible* criteria form what we refer to as the *semi-ambient* nature of the table.

Physical Design

With the requirements above in mind, *Reflect* was designed as an interactive table for four people. In its center, three microphones, forming what is referred to as a microphone array, allow the system to detect which participant is speaking at each point. This is done by selectively filtering the sounds coming from different directions around the array and converting them into separate channels that can be listened to individually. This process, performed by a special purpose system developed by Illusonic [1], is called beamforming. It permits the table to determine the direction the sound is coming from, and hence the current speaker, reliably and without requiring the users to carry any wearable artifacts such as microphones or other sensors. Users can thus simply sit at the table and begin their collaboration as required by the *natural* criterion. A sturdy glass surface permits the table top to be used as a regular working surface.

The display of the table is a matrix of 8×16 multi-color Light-Emitting Diodes (LEDs) that lay beneath a frosted glass surface. The LEDs are individually addressable and form a very low-resolution screen. This choice of display is mainly motivated by the *unintrusive* criterion. The information displayed on the surface of the table should not be complex enough to require significant attention from the users. Having the display at the center of the table and covering most of its surface, would make the information difficult to miss. The bright light of the LEDs also helps the information retain its *visibility* even in well-lit rooms.

Though it is easy to see why this design satisfies the *natural* criterion, it is less obvious that the resulting table would be *unintrusive* and *visible*. We will refer to this question later in this paper when we describe the results of the user study.

Visualization

Given the input of the beam-forming microphone array and with the LED matrix as output, we were free to design a wide range of visualizations. Notably, the territorial display, seen in Figure 1 visualizes the conversation with four "territories" of lit LEDs, one around each speaker. The territories have different colors for different speakers, and they grow in size according to the speakers' levels of participation.

Another visualization that was implemented is seen in Figure 3. We refer to it as the column visualization, and it shows the participation levels of speakers as columns of LEDs, colored differently for each user. The more a user speaks, the more LEDs in his or her column light up. The result is a simple visualization that makes it very easy, and may even encourage users, to compare their participation levels.

Though our initial favorite was the territorial display, the column visualization was the one chosen for the user study for reasons that we will make clear later on.

Design Limitations

We describe *Reflect* as being a tool for balancing participation. Participation however is not solely measured as the amount of noise produced by each person, which is ultimately what *Reflect* is measuring. One can produce much noise without really saying anything, and active listening is a legitimate form of participation. However, with the current state-of-the-art we are far from being able to judge the quality of a person's contribution automatically. We therefore chose to offload this judgment from our system to the only known system currently capable of this type of judgment, the user. We recall thus that *Reflect* is an informative system and the information displayed on its surface should be clear to the user: it shows how much noise each person is making. The users would thus determine for themselves whether or not the levels of speech exhibited by each speaker is a result of lower or increased real participation on the part of some, or if it is simply due to the increased or decreased verbosity of some members.

USER STUDY

In order to evaluate the effect *Reflect* has on collaborative work, we conducted a user study with the aim of validating two hypotheses:

- **H1. Individuals are more aware of their own and their partners' levels of participation when using *Reflect*.** By validating this hypothesis, we would be able to conclude that the information displayed on the table is seen and as-

similated into the user's mental model of the conversation taking place.

- **H2. Groups that are shown their levels of participation on *Reflect* are more balanced than those that are not.** By validating this hypothesis, we would conclude that the information displayed on the table is actually perceived and used by the participants as a tool to avoid over- or under-participation.

Description of the Experiment

Groups of four students were randomly selected from a pool of bachelor students that had volunteered for the experiments. Subjects were paid 50 Swiss Francs (around 44 US Dollars) for their two-hour involvement in the experiment. The groups were asked to solve a murder mystery task offered to us by Stasser and Stewart [18]. The task materials were translated into French and adapted for groups of four. In this task, each subject was given a copy of investigation logs that included maps, interviews and a snippet of a news article. They were asked to accuse one of three suspects of having committed the murder. Each individual version of the investigation logs contained certain important pieces of information that are not available in others. This ensured that all subjects were required to participate in the discussion in order to gather all the necessary information. This type of task, referred to as a hidden profile task, is often used in experiments involving group decision-making and information pooling.

Experimental Conditions

We used two experimental conditions that were identical except for the content of the information displayed on the surface of the table. In the first condition, the students were shown their levels of participation i.e. how much time each student talked. This condition will be referred to as speaker-based condition, in reference to the speaker visualization that is displayed. In the second, they were shown the focus of the discussion, i.e. how much time was spent discussing the case of each of the three suspects in the murder mystery. This condition will be referred to as the topic-based condition.

In both conditions, the columns visualization was used. In fact, the choice of visualization was motivated by the need for a single visualization that can be used for both conditions. Although the territorial display may have been more suitable for displaying speaker levels, it is not at all suited for displaying the time spent on each topic since, unlike the speakers, the different topics do not have a meaningful spatial position that would justify the location of their corresponding territories. This was not a problem for the column visualization as columns were spatially neutral. By labeling the columns, with white stickers posted on both ends of the table, we were able to attribute any kind of information to what each column represents. Both conditions were thus made as similar as possible to one another, with the exception of what information is displayed on the surface of the table.

Participation levels were detected automatically by the table. The subject of discussion was determined using the "Wizard of Oz" technique i.e. with a human listening to the conversation as it took place and remotely signaling that information to the table.

A third neutral condition, in which no information is displayed on the table, was not included in the design of the study as it would have been quite costly and the benefits of having such a condition were not compelling enough.

Experimental Procedure

The students were first asked to read the investigation logs individually for 30 minutes, during which the table was used as a simple timer that kept the students informed of the time remaining. The students were allowed to annotate their copies of the logs and were told that they would keep the copies with them during the discussion. At that point, the students were not yet informed that their copies of the investigation logs contain information that was not available to others.

The students were then given 60 minutes to reach consensus on a suspect. In order to start the discussion, the students were asked to come up with possible means, motive and opportunity for committing the crime for each suspect. They were informed that, in order to accuse a suspect, they must be convinced that he had all of these three elements against him and that the other two suspects were missing at least one of the elements. The students were then made aware that they may possess unique information that is not available to others. In addition, they were told that they were not permitted to give their copy of the investigation logs to another participant and that each participant was only allowed to read from his or her own copy. Finally, the visualizations were explained to the students, but no mention was made of the theoretical benefit of a balanced discussion either in terms of participation or subject focus.

Data Collection

During their discussion, the students were filmed and their voices were recorded using the built-in microphones of the table. Logs of participation levels and of the time spent discussing each suspect were generated and saved. At the end of each experiment, the subjects were asked to fill in a post-experiment questionnaire that contained 19 questions mostly about the experience they had during the experiment and included four open questions. The questionnaire also asked the users to estimate the amount of time each group member spoke as well as the amount of time they spent discussing each suspect.

RESULTS

Given the logs recorded during the experiments, the user estimations of their participation levels as well as the other questionnaire responses, we were able to make some conclusions about the effect *Reflect* had on the collaboration. One group was excluded from the analysis of logs because of an unintentional error that led to loss of recordings and logs for that group, but not the questionnaires which were included in analysis related purely to questionnaire responses. In this section we report the various interesting observations we made.

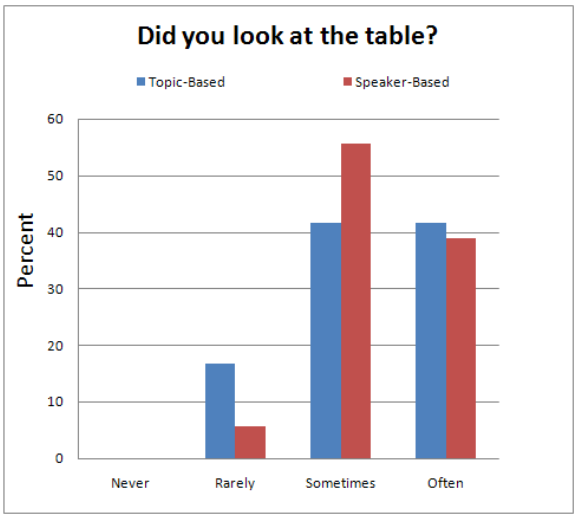

Figure 4. Responses to the question "Did you look at the table?" by condition.

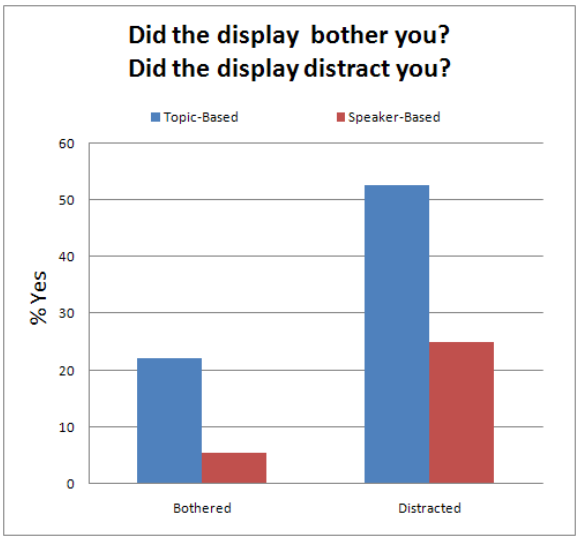

Figure 5. Percentage of subjects who answered "yes" to the quesions "Did the display on the table bother you?" and "Did the display on the table distract you?" across conditions.

The *Visible* and *Unintrusive* Criteria

Recall from the description of the design of the table that compliance with two of its design requirements, the *visible* and *unintrusive* criteria, remained to be verified. We address this issue here.

The post-experiment questionnaire included some questions meant to get a sense of how subjects perceived the table. Some of the questions and their answers will shed some light on this issue. When asked "Did you look at the table?" the vast majority of the subjects in both conditions said they looked at the table either "sometimes" or "often" as seen in Figure 4.

In terms of the intrusiveness of the table, 86% of participants said they were not bothered by the table and 60% said they were not distracted by it. These responses vary across conditions as shown in Figure 5. Note the in the speaker-based condition, which is the condition of primary interest to the study, only 25% reported being distracted by the display.

A small minority of 15% reported feeling "uncomfortable with seeing their participation levels displayed for all to see." Finally, when asked if they would like to use such a table for other meetings, 66% answered "yes" in the speaker-based condition whereas only 25% answered "yes" in the topic-based condition.

We can thus conclude that the table design seemed to satisfy its *visibility* criterion, in that its visualization was looked at most of the time. The subjects also seemed comfortable with the table showing their levels of participation, enough to want to use it in the future. Few reported being bothered by it but a quarter of the users were distracted. These results indicate the table satisfies its *unintrisive* criteria to a large extent, but there is nonetheless room for improvement.

General Effect on Balancing Participation

When looking at the effect the table had on balancing participation levels, we compared how balanced groups were in the speaker-based condition versus how balanced they were in the topic-based condition. We measured balance as the difference between perfect balance (25%) and each user's participation level.

In the post-experiment questionnaire, the subjects were asked the question: "Do you think it is important for members of the group to participate in a more-or-less balanced manner?" When looking at the general effect of the table on the group members' ability to balance their participation, we excluded participants in both conditions that answered "no" to this question (36% of the participants in the study). As we mentioned earlier, *Reflect* is not designed as a tool for enforcing group balance, but rather for supporting it by improving participant awareness. The intention to participate in a balanced manner must thus come from the users themselves, and when this intention is absent, any balancing behavior the user exhibits would likely be coincidental. This could have been avoided by informing the participants before-hand that it is important that they participate equally. We have chosen not to do so because we wanted to keep the collaboration as natural as possible and because informing participants of the benefit of balanced participation could potentially lead to a strong bias in the results, as this information is not usually provided in natural settings.

With the remaining participants (46 subjects), i.e. those who claim to believe that balance in participation is important, we used an independent samples t-test to compare the means of their participation levels across two conditions and obtained a statistically significant difference between the two conditions ($t = 2.18$, $p = 0.036$). In other words, participants who had their participation levels shown to them during the task were statistically more balanced than those who had information about topic focus displayed. This result can be seen

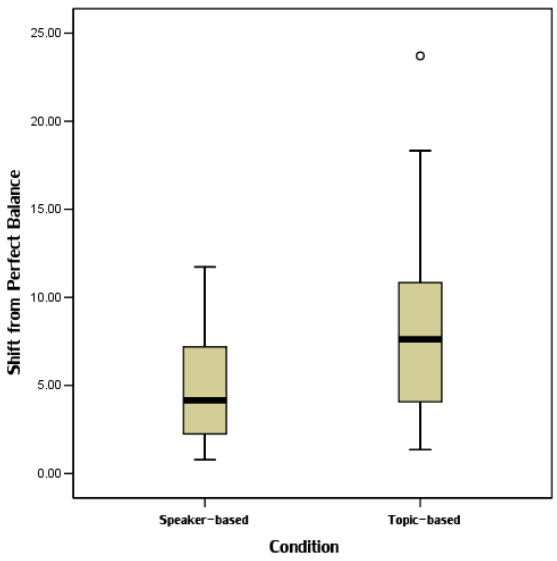

Figure 6. Boxplot showing difference between balance in participation across the two conditions.

in Figure 6.

Effect on Over- and Under-Participators

We studied the effect of the different visualizations on a specific subgroup of participants, namely the extreme participators: those who over-participated and those who under-participated. We were interested in seeing how, over time, these extreme participators modify their behavior. The objective here is to see if spending time around the table would eventually lead to change in behavior. For that, we divided the 60 minute logs into two equal parts of 30 minutes each. We computed the relative participation of each participant during each of the 30 minutes. We then determined those participants who were extreme participators during the first half-hour, and examined how their participation level changes during the second half-hour.

In line with the method used by DiMicco et al. to determine extreme participators [12], we defined over-participators as those who spoke more than the mean participation level (25%) plus the standard deviation of participation levels among all participants. A similar definition was used for under-participators. We ended up with ten over-participators and ten under-participators, divided equally across the conditions.

We observed that, on average, during the first half-hour over-participators in the speaker-based condition spoke less than over-participators in the topic-based condition, though the effect was not significant. More interestingly, in the second half hour, over-participators in the speaker-based condition spoke less than they did during the first half hour while in the topic-based condition, they spoke even more. When comparing the second half-hour participation levels of over-participators across conditions, we observed that there is a significant correlation between participation levels and the condition (t=-3.97, p = 0.004). The effect is similar

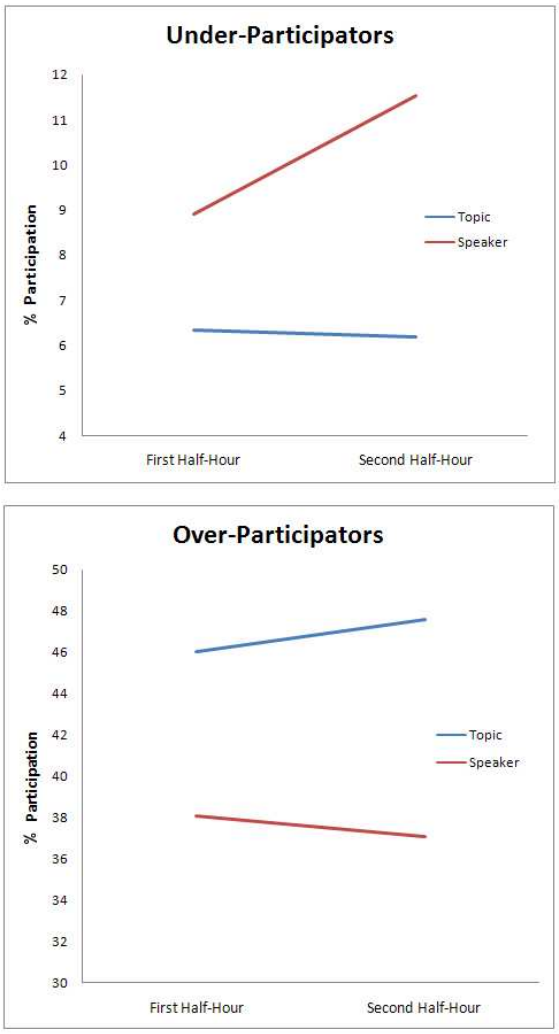

Figure 7. Change in participation levels of extreme participators in both conditions.

when looking at under-participators. During the first half hour, under-participators spoke more in the speaker-based condition than they did in the topic-based condition, and in the second half hour, they increased their participation in the speaker-based condition and reduced it even more in the topic-based condition. However, when comparing the second half-hour participation levels across conditions, the difference is not significant (t=1.304, p=0.22). These results, illustrated in Figure 7, are similar to the findings of DiMicco et al. [12].

Though some of these results do not show a statistically significant effect, which is possibly related to the small number of extreme participators, they do show a trend indicating that the table has the desired effect on participation levels.

Effect on Subject Awareness

We measured the effect the table has on the subjects' ability to estimate both speaker levels for all participants as well as time spent on each topic of discussion (i.e. the suspects). We

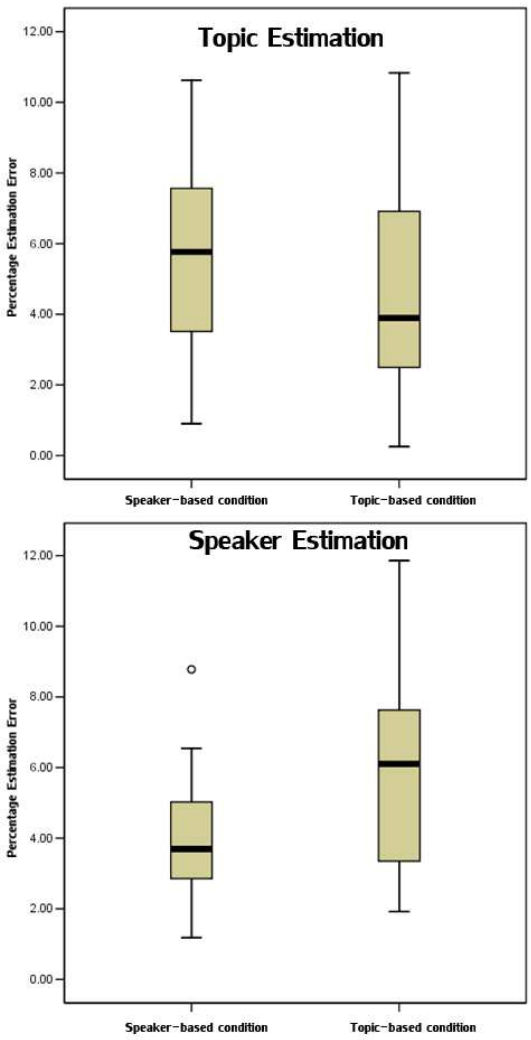

Figure 8. Error levels of while estimating speaker levels and time spent on different suspects across both conditions.

wanted to evaluate how much users are aware of the information displayed on the surface of the table. The subjects where thus asked, as part of the post-experiment questionnaire, to estimate for each member of the group, including themselves, the relative level of participation. Note that the visualization on the table was switched off just *before* the participants were informed that the task is over, and the questionnaire was handed out about a minute afterwards.

For all estimations made, the participants were significantly better at estimating the information in the condition where that information was displayed to them. In other words, when estimating speaker levels, the average error made by the users was significantly lower in the speaker-based condition than in the topic-based condition (t=-3.3, p=0.002), and when estimating time spent on each suspect, the average error was significantly lower in the topic-based condition than in the speaker-based condition (t=2.4, p = 0.02). These results are summarized in Figure 8.

Effect on Topic Balance

In addition to the effect of the table on group balance in terms of participation levels, we also investigated the effect on balance in topic discussion for the topic-based condition. There are of course some conceptual differences between topic balance and participation balance. Unlike participation levels where each member of the group is primarily responsible for his or her own level of participation, no single member is responsible for how much time is spent on each topic. In addition, changes in topic occur much less frequently than changes in speaker, especially near the beginning of the discussion. When the group begins discussing one suspect, they tend to stick to that suspect for a long time before moving to a next one. Finally, the nature of the task does not necessitate that suspects are discussed equally. Some details of the murder mystery require more in-depth discussion than others.

That said, we report that no significant difference occured in terms of topic balance across the conditions (t = -0.24, p = 0.81). The time spent on individual suspects in the experiments varied greatly among groups. Not surprisingly, a large number of participants (70%) felt that it is not "important to spend more-or-less the same amount of time discussing the case of each suspect." In the case of participation levels, we were able to put aside subjects who felt that speaker balance is unimportant. Howver, we cannot do so here since, as stated before, topic balance is not determined by individual users, but by the group as a whole.

Case Study

In order to better understand the effect of the table on our subjects, we present here a brief summary of a case study done with one of the groups that took part in our experiment. A more detailed breakdown of this case can be found in our previous work[2]. These subjects solved the task in the speaker-based condition. We referred to this group in particular because of the different effect the table had on its individual members.

For our analysis we considered the subjects responses to two of the open questions in the post-experiment questionnaire:

1. Can you indicate one or more occasions where the visual display influenced your behavior?

2. Can you indicate one or more occasions where the visual display had a negative impact on the collaboration?

Figure 9 shows the *rate of participation* of each member in this group over time.

Some observations were made about this group discussion.

1. Participant C, whose rate of participation started low but increased to match that of B and D, responded to the second question by saying that when she noticed that her LEDs weren't lit, she got "frustrated."

2. Participant D also exhibits balancing behavior by reducing her level of participation to match two of her group mem-

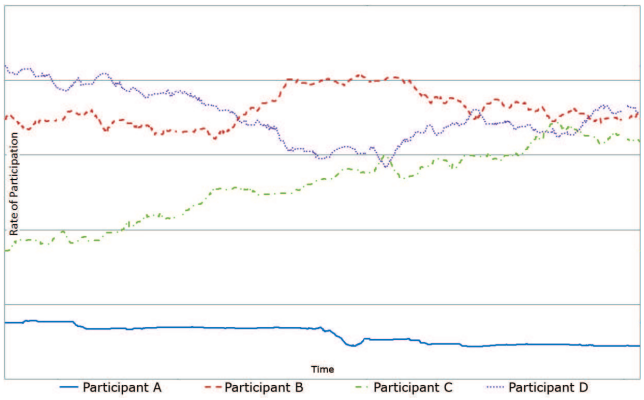

Figure 9. ***Rate of participation*** **of members of one group is the amount of speech produced by each member over a certain amount of time. Three of the speakers' participation rates can be clearly seen to converge, whereas one speaker remains virtually silent.**

bers. In her questionnaire, she explicitly noted that she "tried not to surpass the speaking time of [Participant B]" and that sometimes she "refrained from talking to avoid having a lot more lights than the others."

3. Participant A on the other hand participated very little at the start, and even less in the second half of the discussion. He reported that he *rarely* looked at the table and that he did not feel it is important for members of the group to participate equally. Note that the three other participants reported that they looked at the table either *sometimes* or *often*, and that all three felt that it was important for members of the group to participate equally.

This case study provides further insight into the potential balancing effect this table can have on group discussion as well as the lack of effect it can have on some individuals. It also highlights the informative and not normative role the table has in this kind of setting.

DISCUSSION AND LIMITATIONS

The results of the experiment allow us to draw some conclusions about the effect of a device such as *Reflect* on group behavior. We summarize the main findings here.

Validation of Hypotheses

Our first hypothesis is validated: users are more aware of their participation levels when using the table in speaker-based mode. The significant difference we found when comparing errors in estimating participation levels indicates that the use of the table increased user awareness of these levels. This, of course, does not imply that the users directly used the display of table to learn these levels. It is also possible that by simply knowing that this information was displayed, the users became more conscious of how much they and others were participating. On the other hand, with 88% of the users reporting that they looked at least sometimes on the table (96% in the speaker-based condition), it seems safe to make the claim that the information displayed on the table did indeed increase awareness on participation levels among the members of the group.

The second hypothesis is only partially validated: users who were shown their participation levels are more balanced than those who are not. Though this turned out to be true in general, it is only statistically significant when considering users who claimed to believe it is important to participate in a balanced manner. Given the informative, rather than normative, nature of the table, this is not surprising. The table does not raise a red flag when a participant speaks too much or too little, thus prompting them to balance their behavior. If a user speaks too much and believes it is acceptable to speak too much for whatever reason, being made aware of their over-participation will not push them to reduce their levels of speech.

Our results also showed a significant difference among the second half-hour balance between over-participators across conditions. Under-participators also increased their participation in the speaker-based condition and decreased it further in the topic-based condition, though the difference was not statistically significant. In both cases however, the trend is clear: extreme participators are pushed in the right direction by having the participation levels displayed. However, given the small number of extreme participators, this result is only partially conclusive, and further investigation is needed to establish whether the effect is truly present or not.

Limitiations of the Study

As a first study, this experiment tried to understand the effect *Reflect* has on small groups. Due to the laboratory nature of this study, which was needed for quantitative analysis to be possible, the subjects used the table for short periods of time, and once only. They were working with people they did not know beforehand and will likely never meet afterwards. This limits our ability to generalize the results to possible real-world uses of the table. For example, if a group of four people who work together on a daily basis, have regular meetings on such a table, what will the effect be? Will they eventually lose interest in the feedback provided by the table and start ignoring it? Or will they learn to build a sense of trust with the table as an objective observer and rely on it for guidance? These questions cannot be answered by our one-hour experiments. In the concluding section of this paper, we will describe another study currently being prepared that will address these questions.

The study also did not address the question of group performance either in terms of learning benefits or in terms of making better decisions. However, at this point, we are still trying to study the behavioral effect of the table, not the cognitive one. Balancing participation is seen as the direct objective of *Reflect*, and as such it was the focus of our first study. The actual cognitive effect the table has still needs to be addressed in further studies.

CONCLUSIONS AND FUTURE WORK

In this paper we presented an interactive table, *Reflect*, that is designed to support collaboration between small groups. *Reflect* listens to the conversation taking place around it and displays information on its surface about the levels of participation of the speakers. We conducted a study that shows that

the table does indeed increase awareness of group members about their participation levels. It also, under certain conditions, leads group members to participate in a more balanced manner. We observed a stronger effect of over-participators reducing their participation than of under-parti-cipators increasing theirs.

To further understand the effect of the table, we will soon conduct an real-world study where four prototypes of the table will be placed in four different workplaces for a period of several months. We will observe the effect the table has on groups of people after long-term regular use.

Our ultimate goal will be to address the question of *how* group members are participating, instead of simply *how much*. We are currently exploring the use of pitch and other prosodic features of the voice in order to attribute to each speaker not only a participation level, but also a manner of participation and possibly even a role. Current state-of-the-art indicates that a lot can be told about the outcome of an interaction by simply observing basic vocal features [16]. We aim to incorporate this type of vocal analysis in future versions of *Reflect*.

REFERENCES

1. http://www.illusonic.com.
2. K. Bachour, F. Kaplan, and P. Dillenbourg. Reflect : An interactive table for regulating face-to-face collaborative learning. Lecture Notes in Computer Science. Springer, Berlin / Heidelberg, 2008.
3. M. J. Baker. Argumentation and constructive interaction. In P. Coirier and J. Andriessen, editors, *Studies in Writing: Foundations of Argumentative Text Processing*, volume 5, pages 179–202. University of Amsterdam Press, 1999.
4. T. Bergstrom and K. Karahalios. Conversation clock: Visualizing audio patterns in co-located groups. In *HICSS*, page 78, 2007.
5. T. Bergstrom and K. Karahalios. Conversation votes: enabling anonymous cues. In *CHI '07: CHI '07 extended abstracts on Human factors in computing systems*, pages 2279–2284, New York, NY, USA, 2007. ACM.
6. T. Bergstrom and K. Karahalios. Seeing more: Visualizing audio cues. In *Interact*, 2007.
7. A. Blaye. *Confrontation socio-cognitive et resolution de problemes*. PhD thesis, Centre de Recherce en Psychologie Cognitive, Universit de Provence, France, 1988.
8. B. A. T. Brown, A. S. Taylor, S. Izadi, A. Sellen, J. Kaye, and R. Eardley. Locating family values: A field trial of the whereabouts clock. In J. Krumm, G. D. Abowd, A. Seneviratne, and T. Strang, editors, *Ubicomp*, volume 4717 of *Lecture Notes in Computer Science*, pages 354–371. Springer, 2007.
9. E. G. Cohen. Restructuring the classroom: conditions for productive small groups. *Review of Educational Research*, 64(1):1–35, Spring 1994.
10. P. Dietz and D. Leigh. Diamondtouch: a multi-user touch technology. In *UIST '01: Proceedings of the 14th annual ACM symposium on User interface software and technology*, pages 219–226, New York, NY, USA, 2001. ACM.
11. J. M. DiMicco and W. Bender. Group reactions to visual feedback tools. In *PERSUASIVE*, pages 132–143, 2007.
12. J. M. DiMicco, K. J. Hollenbach, A. Pandolfo, and W. Bender. The impact of increased awareness while face-to-face. *Hum.-Comput. Interact.*, 22(1):47–96, 2007.
13. J. M. DiMicco, A. Pandolfo, and W. Bender. Influencing group participation with a shared display. In *CSCW '04: Proceedings of the 2004 ACM conference on Computer supported cooperative work*, pages 614–623, New York, NY, USA, 2004. ACM Press.
14. G. P. Huber. A theory of the effects of advanced information technologies on organizational design, intelligence, and decision making. *The academy of management review*, 15(1):47–71, Januray 1990.
15. M. Webster Jr. and J. E. Driskel Jr. Beauty as status. *The American Journal of Sociology*, 89(1):140–165, July 1983.
16. A. S. Pentland. *Honest Signals: How They Shape Our World*. The MIT Press, 2008.
17. R. E. Slavin. *Cooperative Learning*. Longman, New York, 1983.
18. G. Stasser and D. Stewart. Discovery of hidden profiles by decision-making groups: Solving a problem versus making a judgment. *Journal of Personality and Social Psychology*, 63(3):426–434, September 1992.
19. N. A. Streitz, P. Tandler, and C. Muller-tomfelde. Roomware: Towards the next generation of human-computer interaction based on an integrated design of real and virtual worlds. pages 553–578. Addison Wesley, 2001.
20. N. M. Webb. Task-related verbal interaction and mathematics learning in small groups. *Journal for Research in Mathematics Education*, 22(5):36–389+, 1991.
21. A.-N. P.-C. Willem Doise, Gabriel Mugny. Social interaction and cognitive development: Further evidence. *European Journal of Social Psychology*, 6(2):245–247, 1976.