

Seeing More: Visualizing Audio Cues

Tony Bergstrom and Karrie Karahalios

Department of Computer Science
University of Illinois Urbana-Champaign
Urbana, IL 61801
{abergst2, kkarahal}@cs.uiuc.edu

Abstract. Using audio visualization, we seek to demonstrate how natural interaction is augmented with the addition of interaction history. Our *Conversation Clock* visualization captures and represents audio in a persistent and meaningful representation to provide social cues not available in an otherwise ephemeral conversation. In this paper we present user study evaluation of the *Conversation Clock* as utilized by familiar groups and demonstrate how individuals use the salient cues to evaluate their own interaction.

1 Introduction

Having spent thousands of years of socializing, communication and speech are natural interactions for most people. It is a skill so honed that most people do not consider how complicated communication is. More than just speaking, face to face communication can be broken into 7% words, 38% tone of voice and 55% body language [1]. These social cues are sensed and utilized as necessary during conversation; however, details of interaction are promptly forgotten [2][3] or the cue is interpreted and reciprocated unconsciously [4].

Social interaction contains observable aspects of conversational dominance, interruption, length of turns, turn taking, interruption, overlapping of speech, the rhythm and flow of conversation, etc. These aspects can be detected by reviewing or recording interaction, but are difficult to utilize as cues during a conversation. With the *Conversation Clock*, we use these aspects to visually augment collocated conversation with new conversational cues.

The *Conversation Clock* makes salient aspects of conversation that would be otherwise unnoticed. In this paper, we demonstrate the effects of augmenting collocated conversation to gauge how people perceive themselves and others. As results of our user study have shown, participants in conversation become more aware of their own and other participants' interaction and react to cues differently creating a more balanced conversation.

2 Related Work

The bulk of our work is situated in audio visualization. Before discussing our own project, we will briefly present previous work in visualizing sound and conversation.

DiMicco et al. describe groups of four and how their participation might be affected by a shared display[5][6]. The display was projected onto two walls facing each other. Participants in this system must significantly divert their gaze to view and interpret the displayed information. The visualization consisted of a labeled histogram, each bar depicting the corresponding participant's contribution to discussion. Though the contribution represented aggregated interaction data, participants could not directly examine the conversation history. Based on the measured contribution, the histogram explicitly labeled participants as "over-participating," "participating" and "under-participating" on the shared screen. The goal of this system was to aid group decision making by balancing interaction during discussion. During a study of group problem solving situations, DiMicco measured the change of participation levels. Observational and experimental data indicated that the over-participators were more likely to back off than the under-participators were to speak up.

In later work, DiMicco compiled and examined Second Messenger, a suite of visualizations, for asynchronous review of a meeting [5][7]. Having shown that live displays influence group participation, the suite explored more expressive visualizations. These visualizations are examined post-meeting, at times when participants have opportunity to contemplate and interpret the meaning of the visuals. Each visualization focused on different aspects of interaction including relative participation, turn-taking, overlapping speech, and history.

Previous work with Visiphone examined the effect of visualizing audio in remotely mediated conversation between two locations[8][9]. Visiphone presented a domed projection surface, allowing participants to gather around and focus their gaze on the device while talking. A spiraling and cascading stream of circles on the surface continuously show speech and aural activity between the two spaces. The visualization conveyed the degree of aural contribution by adjusting the diameter of the circles as they spiraled down the sides. Simultaneous activity on both ends of the Visiphone connection presents overlapping circles allowing the device to reveal patterns of conversational dominance, turn taking, and interruption. The visualization allows you to "see things you know, but do not realize you know." Visiphone directly influenced the interaction between two parties without explicitly providing direction. Participants in the local spaces adjusted the volume of their own speech to visually match the volume of the remote speakers rather than by ear. Additionally, conversational dominance became obvious as the dome's color becomes visually dominated by a single hue. Due to the salience of balance in communication, Visiphone was recommended for use in areas such as marriage counseling. We extend this investigation to explore how this interaction changes when there are no remote individuals and whether visualization provides the same influence when participants are in the same space.

As a museum installation, Vigas presented Artifacts of the Presence Era to demonstrate history of a space over time [10]. This visualization used a geological metaphor to aggregate time. Pictures of the space were regularly taken, portions of which were layered on the already existing history. Audio input was sampled

to determine the shape of the next layer. The audio samples, based on volume, revealed a larger portion of the image during periods of activity. Older layers were compressed as the new images were added to facilitate a geological metaphor. The result is a visualization of the installation space that highlights the change of activity and gives a sense of a space's history.

A work by Karahalios, Telemurals, connected two remote spaces by abstractly rendering and projecting video [11]. Individuals were obscured but became more visible and clear as both spaces interacted. Speech recognition provided a channel of interaction as it was displayed along with the abstract depiction of the other space. Telemurals, effectively increased conversation to at least five times its normal level.

Numerous artistic installations have visualized live conversation and sound [12][13]. Work by Levin and Lieberman included *Hidden Worlds*, a visualization of speech as three dimensional bubbles that float over the table as participants wear special glasses to see these bubbles. Another piece, *RE:MARK* creates shapes symbolic of spoken phonemes that escape from the head of a person's shadow. Many other installations were also on display. These installations have been successful at encouraging both individual and group interaction by allowing users to manipulate a visual depiction of their aural input.

Another branch of audio visualization is music visualization. Most computer based music players utilize visualizations such as G-Force, Geiss, Advanced Visual Studio (AVS), or something similar [14][15][16]. To generate visuals, these visualizations incorporate physical characteristics of audio input such as amplitude, frequency analysis, beat detection, etc. Used for performance and entertainment, these visualizations focus on the physical characteristics of sound and not the interaction between multiple sound sources.

Aside from audio visualization, text-based chat systems have been augmented to include additional visual cues for remote participants. Systems like Babble and Chat Circles have sought to create a new interaction environments to show conversational activity, history and status [17][18].

In the following sections, we describe our own work with audio and visualization in order to provide social mirrors to collocated groups. We present a user study and demonstrate how the *Conversation Clock* makes interaction patterns more salient to participants.

3 Conversation Clock

As we describe in our previous work [19][20], the *Conversation Clock* visualizes the interaction visualizes the interaction patterns of up to four individual and provides participants with a communal social mirror. Representing each participant with an associated color, the *Conversation Clock* builds a visual history of ephemeral audio (Figure 1). Using a circular structure to represent the passage




Fig. 1. Participants are seated around the *Conversation Clock* table (left). Lapel microphones monitor conversation while the visualization of history is projected in the center. The *Conversation Clock* provides a visual history of interaction and communication (right). Each contribution displays bars colored to indicate the speakers' identities. The lengths of these bars indicate the degree of participation, measured by volume. As a conversation progresses, a history is built with concentric rings reminiscent of the rings on a tree.




Fig. 2. Above we see a close up of aural contribution represented in the *Conversation Clock*. The length of individual bars correspond to the aural activity of speakers; longer bars indicate louder speech. Individual speech is seen as solid colored bars while layered bars demonstrate simultaneous speech. The dots at the right indicate no aural contribution. As time progresses clockwise (left to right in this image), we see the changing dynamics of conversation contribution.

of time, the *Conversation Clock* emphasizes recent interaction in the outer ring while showing past interaction in the interior.

Aural contribution is gauged by sampling audio from a microphone assigned to each participant. We chose to use a calibrated amplitude as the main measure of determining speakers and contribution. Individual samples are translated into discrete marks on the table; the amplitude of the sample proportionally determines the length of the projected mark (Figure 2). These marks can be examined for current speaker, shown as the longest mark on the table; and multiple speakers, shown with overlapping marks.

The *Conversation Clock* displays a wide variety of conversational cues during interaction: turn taking, interruption, conversational dominance, silence, agreement,

aural back-channels, mimicry, time spans, rhythm and flow. The cues persist on the table allowing participants to view the visualization when convenient.

4 Conversational Context

The *Conversation Clock* does not explicitly represent the context of a conversation. Rather than telling people how to interpret their conversation, we allow participants at the table to imbue meaning into the visualization and visual cues provided by the mirror. In this section we explain the intricacy of context in conversation and how it has influenced our design.

Conversation and language are flexible. Understanding spoken words is not the equivalent of understanding meaning. Similarly, observing characteristics of interaction is not the same as understanding their significance. It is easy to note occurrences of overlapping speech, changes in speaker, and periods of silence; examining recordings and sensor data reveals these events. Showing interruptions as opposed to a change of turn requires context.

Overlapping speech is a prime example where context must be used to understand its meaning. In the course of normal conversation, speech overlaps more often than people perceive [2]. Overlapping speech might be a hostile interruption, a back channel comment, an emphasis of a point, a change of turns or a misjudged transition. Looking at patterns without context, one cannot judge the true intent of the speakers.

The context of a conversation includes location, time, speaker, topic, word choice, inflection, etc. These are relevant to interpreting the meaning of the conversation. To understand interaction, one must examine cultural expectations, relationships of participants, gender, conversational style, etc. Altering the context of an individual significantly affects conversational patterns and style [21].

Context determines how interaction should be interpreted. Consider the simple case of remaining silent in response to a question. The silence could be an indication the questioned doubts his or her response, doesn't like to speak, or wasn't paying attention. When both parties are fully aware of the question, the questioner might prolong the silence, in order to apply pressure for a response. In another context, silence might be used to build curiosity by playfully choosing not to respond and building curiosity. Additionally, some people think silently when asked a question. These are all hypothetical situations; there are many more that can be arrived at by altering the context of the situation. Knowing that context is important to the interpretation of silence [2].

To be able to correctly characterize interactions such as interruptions, overpowering, conflict, and agreement we must first detect context. Our system is not sophisticated enough to analyze and interpret interaction adequately. To cope with context, we chose to provide no explicit interpretation in our visuals. Participants are presented with a visualization of their interactions over time, a social mirror of their conversation. Interpreting the context is left to the participants, making the *Conversation Clock* a meaningful and personal object to each group and conversation.

5 User Study

In order to gauge the influence of the *Conversation Clock* a user study was conducted over the course of a month. Participants were asked to commit to multiple sessions during that month under different experimental conditions. From these multiple sessions we gain better insight into how the *Conversation Clock* influences both groups and individuals.

5.1 Configuration

The user study took place in the HCI/Graphics lab in the UIUC Computer Science building. Four lapel microphones were wired to the bottom of our table. Tags on each individual microphone indicated the color associated with that microphone. Participants were informed that this color visualizes their activity in the *Conversation Clock*. A DV-camcorder monitored interaction from a nearby tripod during all sessions. During conversation, this video recorded gestures and physical references for later analysis.

5.2 Scenario

Participants were solicited as groups. Each member was expected to be familiar with all other members of the group and have a predefined purpose for meeting. We chose to use familiar groups in order to ensure a more comfortable and natural interaction environment. We felt that using random groups over the course of three separate sessions would alter interaction patterns as individuals became more familiar with the other participants in their group. We were also wary of mixing groups as each set of individuals would likely have different group dynamics due to the makeup of the group [21]. The focus of group conversation ranged from weekly meetings for a class project to casual conversation about movies and music.

Prior to each meeting, groups were asked to plan for a 30 minute session. However, groups were allowed to finish early as this would capture the natural break up of a session and would not force interaction and conversation.

5.3 Sessions

Each group met for three sessions. Groups first met in a session with no visualization present. Their interaction was observed and recorded. This session offered a baseline for later comparisons. In the second session, we introduced participants to the *Conversation Clock*. The final session was a repeat of the conditions in first session, allowing us to capture the natural variation inherent in different conversations amongst the same group.

In all sessions, participants wore microphones and underwent a calibration process. The *Conversation Clock* was generated for each session, though only projected for the second session. The purpose of running the visualization was

to generate comparable data between the two conditions for analysis. The *Conversation Clock* stores a log of data that can be used to count turns taken and measure the amount spoken during the sessions.

When the visualization was first projected onto the table, participants were reminded of the color associated with their microphone. Speaking was encouraged to better understand how the visualization worked. This period allowed participants to become familiar with the visualization before the session began.

Within the visualization condition, two settings were examined. Each group would use the *Conversation Clock* with a visualization history of 5 minutes and a history of 20 minutes. To account for possible ordering effects, half of the groups were randomly chosen to receive the short history first. Participants were informed that they would be interrupted after 10 or 20 minutes (appropriately). A brief survey was administered and the *Conversation Clock* was cleared and set to use the remaining history setting before proceeding.

6 Study Results

During the second session, participants completed a brief questionnaire and survey. The surveys followed both the 10-minute condition and the 20-minute condition. In both cases participants reported their interaction to have been moderately altered by the table. Additionally, using a Wilcoxon Signed Rank Test, we found that the 20-minute session with the longer history had a significantly greater perceived effect upon participants (Table 1).

Table 1. Results of the Likert Scale surveys in the *Conversation Clock* User study. Using a Wilcoxon Signed Rank Test, we show that the longer history was perceived to have a greater effect on participation than the short history.

Measurement	10-minute	20-minute	Wilcoxon Z	Significance
<i>Distracting</i>	2.56	2.80	-1.364	$p < .2$
<i>Helpful</i>	2.46	2.36	-.966	$p < .4$
<i>Altered You</i>	2.48	3.00	-2.871	$p < .004$
<i>Altered Others</i>	2.65	3.02	-2.066	$p < .04$
<i>Satisfaction</i>	3.24	3.04	-1.475	$p < .2$

Level of satisfaction with the *Conversation Clock* varied greatly (ranging from 1 to 5) depending on the individual, group, and task. In the end, participants were neither particularly satisfied nor greatly dissatisfied. Participants reported the *Conversation Clock* to be somewhat distracting and somewhat helpful.

In addition to self-reported measures of interaction, coded video and audio recordings provided additional measures. The two sources of data we utilized were the audio logs and video.

Video coding and a repeated measures ANOVA revealed participants significantly changing their gaze patterns ($F(3,3)=15.56$ $p=0.03$) by averting their eyes more often from the conversation and to the *Conversation Clock* (Figure 3).




Fig. 3. The *Conversation Clock* contributed significantly to the increase of gaze movement on top of the gaze movement already present

A notable change in the table gestures was observed, though our sample could not show significance ($F(3,3)=5.98$ $p=0.09$).

Further examining the change in gaze patterns induced by the *Conversation Clock*, the 10-minute and 20-minute sessions were compared via paired t-test ($t(6)=0.94$ $p=0.38$). No significant change was observed, indicating gaze patterns were not affected by the length of visualized history.

Examining the audio data, we split participants into two categories. Based upon the initial non-visualization session, participants were categorized as above average and below average participants. As shown in Dimicco’s earlier work, we expected the different categorizations of individuals to alter their participation differently [6]. This categorization is made strictly for analysis; participants were never informed of their categorization.

Running a repeated measures ANOVA, we found significant alteration in some areas (Table 2). Specifically, we noted reductions in turn length of the above average participants, and the increase in the number of turns taken by below average speakers. Having seen significance overall, we investigated the two visualization conditions. Once again, we found no indication of notable differences between the short and long history.

The *Conversation Clock’s* observed data was also sorted to examine how interaction changed over time. As no significant differences were observed between

Table 2. Audio data, automatically logged and analyzed, was available for each participant. Expecting different changes to occur in each, participants were categorized as above average or below average participants based upon the initial session. We ran a repeated measures ANOVA and found above average participants change the length of there turns while below average participants changed the number of turns.

Participation	Metric	F value	Significance
Above Average	Leading	$F(3, 10) = 1.29$	$p = 0.3$
	Turns	$F(3, 8) = 1.19$	$p = 0.4$
	Length	$F(3, 5) = 9.22$	$p = 0.02$
Below Average	Leading	$F(3, 7) = 0.52$	$p = 0.7$
	Turns	$F(3, 9) = 3.89$	$p = 0.05$
	Length	$F(3, 12) = 3.32$	$p = 0.06$

Table 3. Re-sorting to indicate the first condition seen and the second condition seen, it becomes apparent there were learning or novelty effects. Participants in both the above and below average participation categories significantly altered the amount of time spent leading. Examining the other notable statistics, the above average category seems to have altered the length of their turns while the below average participants seemed to change the number of turns to accomplish this feat.

Participation	Metric	N	Condition 1	Condition 2	Paired t	Significance
<i>Above Average</i>	Leading	15	20.24	16.85	$t(14) = 2.66$	$p < 0.02$
	Turns	12	1.79	1.67	$t(11) = 0.839$	$p = 0.4$
	Length	11	14.87	12.21	$t(10) = 2.16$	$p = 0.06$
<i>Below Average</i>	Leading	8	6.86	11.20	$t(7) = -3.56$	$p < 0.01$
	Turns	11	0.98	1.19	$t(10) = -2.06$	$p = .07$
	Length	12	7.69	9.30	$t(11) = -1.13$	$p = .3$

the 10 and 20 minute sessions, we became interested in how interaction changed with continued exposure. We see in Table 3, the most affected audio feature was the amount of time spent leading the conversation. This is true for both the above average and below average categorizations.

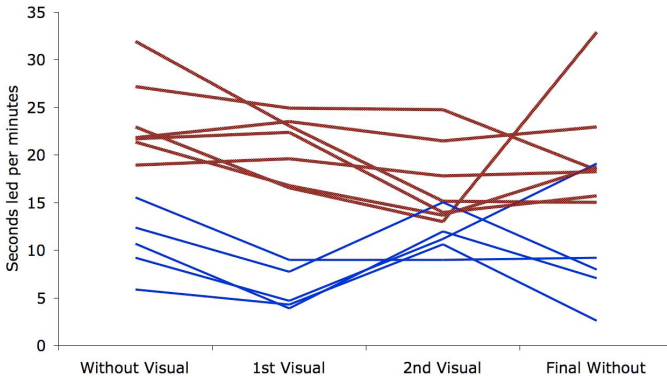


Fig. 4. Filtering participants based on their survey responses, we chose the individuals who reported a 4 or 5 when asked if their interaction was altered. Viewed chronologically, we see the same pattern in the individuals, below average participants (upper lines) lead more while above average (lower lines) lead less in the second visualization session.

Limiting our scope to the people reporting to be most affected, reporting a 4 or 5 on the written survey, we see a similarly striking pattern (Figure 4). After an initial break in period, the participants' conversation converges and reduces variance between speakers.

Data coded from the video is also emphasized by a chronological examination. Both the gaze patterns ($F(3,4)=24.5$ $p=0.005$) and gestures to the table




Fig. 5. When viewed chronologically, we see the first visualization session attracted a large number of glances away from conversation (left). Individual groups are indicated with black lines and the average with a thicker blue line. Overall, the *Conversation Clock* attracted looks and glances from participants in addition to other glances away from the table (Figure 3). Conversely, the *Conversation Clock* reduced the number of gestures made by participants during conversation (right).

($F(3,4)=10.20$ $p=0.02$) are significantly altered, according to a repeated measures ANOVA. Examination of the actual data shown the *Conversation Clock* still increasing the number of gazes averted from conversation and objects referenced in conversation and the decrease in the number of gestures (Figure 5).

6.1 Qualitative

The first session was perceived by most participants to be a normal interaction. As we purposely chose to use groups that were already formed and familiar, participants were in a familiar scenario. We consider the lab environment to be a small change, as many groups we studied do not meet at a regular location. A few minor comments were made when questioned, but nothing more than:

“After a while it was [normal], once we got on topic... That and we had to sit at a table.”

Observing the interaction of the groups during the sessions, we noticed that many of these initial sessions had skewed participation levels. This included three groups that had a combined four participants with less than 10% of contributed audio. Only one group reported an unbalanced conversation. Another stated it was very typical to what they had seen in past meetings.

The second session displayed our social mirror, the *Conversation Clock*. Participants seemed more aware of their own and others’ participation.

“I realized that I could monitor my speech patterns by watching the colors. It was interesting to train myself not to say ‘umm’ as much or pause.”

“It’s easy to judge who is driving conversation.”

“I was trying to look at the circle to see whether we were balanced.”

“It’s more salient, what’s going on.”

The awareness of conversation affected both speakers and listeners. Participants reported that when speaking they would desire to finish quickly and pass the conversation over to another; however, when listening, participants were aware of their silence and were compelled to speak.

During one group's discussion, a collaborative decision on what movies to watch at a weekly meeting, participants felt there should be a balance between themselves. However upon seeing the visualization they realized how heavily dominated the conversation was:

"I noticed when you're the one talking, you want to stop. But if you're mid topic you couldn't stop, because you had to finish your topic. But as soon as you finished your topic, you'd shut up."

When the same group was asked if their interaction had been a balanced conversation the heaviest contributor (about 40% of the leads and 40% of the turns) responded, *"Not after looking at the clock."* This group also indicated an expectation of equal participation in this group, whereas a similarly 'unbalanced' viewed the contributions differently,

"Project managers communicate more than testers or developers, thus it is reasonable for it to be unbalanced."

Another participant making a heavy contribution to conversation commented on the completed ring as a milestone, noting,

"By the time the ring had returned to the starting point I was like 'That's too much' and a stopped talking because I feel like I completed that ring. That's one unit that's mine."

Whereas participants who had been silent for a while could be prompted or compelled to speak by the *Conversation Clock*.

"Based upon more of the history after a while, when I did glance at the clock and then it like reminded me to prompt [Blue] or direct questions to him."

"It became all red, should green or yellow speak next?" - Yellow

Very little conversation during the sessions directly referenced the *Conversation Clock*. Over the course of conversation a group would generally only mention the *Conversation Clock* one time if at all. During the experiment, the focus was on topic for their group. However, during the break between conditions and after the experiment was over, participants focused direct attention to the clock. Some just spoke to watch their contribution, others would yell and try to make their own color dominate, another group began to whistle. Groups reported it was *"fun"*, *"interesting"* and *"amusing"*.

The *Conversation Clock* encountered criticism in groups with a task based purpose or a more defined schedule. Groups with more structure and predefined interaction had less flexibility. One participant explained his lack of contribution,

“Due to our group roles each member was expected to talk about a certain project anyway, so I had little input on something I didn’t know.”

Some participants also found the visualization distracted them from the meeting at hand.

“I found it a bit distracting, I don’t like things in my visual space.”

Though *Conversation Clock* was not visible in the third session, participants reported thinking about it during their interaction. One participant stated:

“The first [session] I wasn’t aware of it ... the second [session] I looked at the clock and tried to make it balanced. Now, I don’t have a measure ... but I still tried to use the impression from last time. Because he always talked more [last time], I tried to consciously talk more [this time].”

Another participant who had been the heaviest participant in his group during session two commented:

“I think I consciously tried not to talk. Then I realized I wasn’t talking at all.”

7 Discussion

Our research illustrates that live social mirrors, like the *Conversation Clock*, influence the dynamics of interaction among collocated groups. By providing visual cues generated from the ephemeral audio conversation, we have created a visualization that allows participants to evaluate their own interaction and the history of conversation during that conversation. In this section we will summarize and discuss our findings.

- *People perceive a notable difference between short and long histories.* According to self reporting, longer histories are more effective at altering one’s own interaction. However, no noticeable change in gaze, gestures made, time spent leading, turns taken, or length of turns ever corroborated this observation. We speculate if this perceived change does exist, it was something which we did not measure or the difference in what we did measure may be more subtle. Additionally, due to fallibility of self reporting, there may be no real difference in interaction [3].
- *The Conversation Clock encourages participants to remove their gaze from the conversation.* People looked away from the conversation more often. Glances to the *Conversation Clock* did not replace other glances that occur naturally during a conversation, and participants did not report a loss of quality interaction. As supported by the qualitative data, the social mirror led individuals to become curious of their representation. Some glances to the table are a result of inspecting one’s own interaction.

- *Above average speakers alter the length of their turns while below average speakers alter the number of turns.* These observations illustrate how the common social mirror can be individually interpreted. The different interpretation could be seen in the their verbal adjustments and interview responses. Below average participants commented on domination of the entire table, noting it was mostly one color. Whereas above average participants mentioned completing rings as a milestone. This difference in observation and perception of interaction demonstrate the principle of small multiples [22]. In the larger picture made visible by history, one can see how they have contributed over the course of conversation. In the outer ring, detail is exposed for the recent past, distinctly affecting the current speaker.
- *Participants showed particular interest in their own interactions.* Qualitative feedback indicated that people were most aware of their own interaction and monitored their own contributions. They felt they were speaking too much or too little based upon what they saw. Additionally, some participants evaluated his own contribution and usage of filler words like “umm.”
- *Ordering demonstrates a notable break in period for the Conversation Clock.* Many patterns of change in interaction become apparent when viewed chronologically. Nearly all our measurements indicated significant change. This strong change motivates a break-in period. Future studies are needed to investigate the continued, both regular and prolonged, use of the *Conversation Clock* to fully understand it’s effect. This point is underscored by one group’s comment almost 30 minutes into their second session when the realized, “*This IS useful.*”

8 Conclusion

Audio visualization provides an informative medium to convey social, interactive, and meaningful cues. We have used varied lengths of audio history to show that live visualization of audio through social mirrors can provide influential cues for individual participation in conversation. Participants alter themselves in order to equalize the contribution of individuals.

References

1. Mehrabian, A.: Silent messages. Wadsworth, Belmont, California (1971)
2. Tannen, D.: Gender and Discourse. Oxford University Press, Oxford (1994)
3. Bernard, H.R., Killworth, P., Kronenfeld, D., Sailer, L.: The problem of informant accuracy: The validity of retrospective data. *Annual Reviews Anthropology* 13 (1984)
4. Whyte, W.H.: The Social Life of Small Urban Spaces. Municipal Art Society of New York (1988)
5. DiMicco, J.M.: Changing Small Group Interaction through Visual Reflections of Social Behavior. PhD thesis, Massachusetts Institute of Technology (2005)
6. DiMicco, J.M., Pandolfo, A., Bender, W.: Influencing group participation with a shared display. In: CSCW '04. Proceedings of the 2004 ACM conference on Computer supported cooperative work, pp. 614–623. ACM Press, New York (2004)

7. DiMicco, J.M., Hollenbach, K.J., Bender, W.: Using visualizations to review a group's interaction dynamics. In: CHI '06 extended abstracts on Human factors in computing systems, New York, USA, pp. 706–711. ACM Press, New York (2006)
8. Donath, J., Karahalios, K., Viégas, F.: Visualizing conversation. *Journal of Computer Mediated Communication* 4(4) (1999)
9. Donath, J., Karahalios, K., Viégas, F.: Visiphone. In: ICAD 2000: International Conference on Auditory Display (2000)
10. Viégas, F.B., Perry, E., Howe, E., Donath, J.: Artifacts of the presence era: Using information visualization to create an evocative souvenir. In: INFOVIS '04. Proceedings of the IEEE Symposium on Information Visualization (INFOVIS'04), Washington, DC, USA, pp. 105–111. IEEE Computer Society Press, Los Alamitos (2004)
11. Karahalios, K., Donath, J.: Telemurals: linking remote spaces with social catalysts. In: CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems, pp. 615–622. ACM Press, New York (2004)
12. Levin, G., Lieberman, Z.: *Messa di Voce*. Audiovisual Performance and Installation (2003)
13. Levin, G., Lieberman, Z.: In-situ speech visualization in real-time interactive installation and performance. In: NPAR '04: Proceedings of the 3rd international symposium on Non-photorealistic animation and rendering, pp. 7–14. ACM Press, New York (2004)
14. Geiss, R.: Geiss, <http://www.geisswerks.com/geiss/index.html>
15. Frankel, J.: Advanced visualization studio. (Nullsoft)
16. O'Meara, A.: G-force. <http://www.soundspectrum.com/g-force/>
17. Erickson, T., Smith, D.N., Kellogg, W.A., Laff, M., Richards, J.T., Bradner, E.: Socially translucent systems: social proxies, persistent conversation, and the design of 'babble'. In: CHI '99: Proceedings of the SIGCHI conference on Human factors in computing systems, New York, NY, USA, pp. 72–79. ACM Press, New York (1999)
18. Donath, J., Viégas, F.B.: The chat circles series: explorations in designing abstract graphical communication interfaces. In: DIS '02. Proceedings of the conference on Designing interactive systems, pp. 359–369. ACM Press, New York (2002)
19. Bergstrom, T., Karahalios, K.: Visualizing audio patterns in co-located groups. In: HICSS '07: Proceedings of the Fourtieth Annual Hawaii International Conference on System Sciences (2007)
20. Karahalios, K., Bergstrom, T.: Visualizing audio in group table conversation. In: TABLETOP '06: Proceedings of the First IEEE International Workshop on Horizontal Interactive Human-Computer Systems, Washington, DC, USA, pp. 131–134. IEEE Computer Society Press, Los Alamitos (2006)
21. Tannen, D.: Teachers classroom strategies should recognize that men and women use language differently. *The Chronicle of Higher Education* (1991)
22. Tufte, E.R.: *The Visual Display of Quantitative Information*. Graphics Press (1983)