

Tools for Distributed Facilitation

Michael J. McQuaid,¹ Robert O. Briggs,² Duffy Gillman,¹ Rosie Hauck,¹ Chienting Lin,¹
Daniel D. Mittleman,³ Jay F. Nunamaker, Jr.,¹ Marshall Ramsey,¹ Nicholas Romano,⁴

¹University of Arizona
Tucson, Arizona

²Ventana Corporation
Tucson, Arizona

³DePaul University
Chicago, Illinois

⁴University of Tulsa
Tulsa, Oklahoma

Abstract

We have extensive experience with distributed facilitation of synchronous virtual meetings since 1992. More recently, we have gained experience in asynchronous virtual meetings and have learned lessons from which best practices have been distilled. Best practices for synchronous and asynchronous distributed facilitation give insight into requirements for toolsets. Our experience with toolsets, including (1) widely available unstructured toolsets, (2) virtual reality offices, and (3) persistent visualization, provide direction for further tool development.

1. Distributed Facilitation

Virtual meetings, meetings between geographically dispersed participants, are being supported by more and more software systems, as catalogued by Woolley [10]. Yet documented experience with virtual meetings remains limited, as exemplified by Beise et al. [2], who found that only 15 of the 34 facilitators investigated had ever participated in virtual meetings. Facilitators in this study had expectations about virtual meetings, but lacked experience to inform differences between virtual meetings and face-to-face meetings. Without experience, these facilitators don't know what, if any, new toolsets will be required for virtual meetings.

Some early studies have classified virtual meetings as *synchronous*, where participants are geographically dispersed but interacting at the same time, and *asynchronous*, where participants may join and leave the meeting and return at different times and any two participants may not be involved at the same time, but nevertheless interact with each other. Another way to classify tools is by the communication channels: text alone or with the addition of audio or video. Table 1 shows six modes of virtual meetings and

how challenges *accumulate* (i.e., each cell inherits the problems of the cells above and to the left) from mode to mode as we add communication channels and relax synchronicity.

Activities of virtual meetings can be grouped into two kinds, content and process. Content refers to the matter of the meeting, the issues, objectives, and organizational outcomes. Process refers to the techniques, such as brainstorming, converging, summarizing, and voting, used to achieve outcomes, as well as to the organization of those techniques.

Both process and content must be discussed explicitly. The term "back channel" refers to a conduit for communication about other than the main content. Typically, "back channel" communications are communications about the processes used to address content and the mechanisms used to support processes. For instance, the content may be that the group needs to agree on (a) measurable goals for improvement. The process is to (b) brainstorm possible goals through (c) an electronic brainstorming tool, (d) converge on best goals through (e) multi-voting, and (f) summarizing specific measures for these goals through (g) group writing. In this example, (a) is content, (b), (d), and (f) are processes, and (c), (e), and (g) are process support mechanisms. Group facilitators typically need to be able to explicitly separate and discuss these issues.

We have had extensive experience in synchronous virtual meetings since 1992. For instance, Mittleman et al. [7] reported on findings and lessons learned from a synchronous distributed group facilitation study. These findings illuminate needs arising in virtual meetings. How can these needs be met?

Our experience with synchronous meetings has given us an advantage in asynchronous, but we have found difficulties carrying over expertise from one to the other. Mittleman et al. [6] developed a set of best practices from the prior study and a related study of asynchronous distributed group facilitation. The characteristics of these practices point the way to the requirements for tools to assist in facilitation of virtual meetings.

	$\overrightarrow{\text{less synchronous}}$	<i>Synchronous</i>	<i>Asynchronous</i>
<div style="display: flex; align-items: center;"> <div style="margin-right: 5px;">↓</div> <div style="text-align: center;"> <i>more channels</i> </div> </div>	<i>Text</i>	Focus and context: Which window are participants looking at? Why aren't participants typing in text right now?	Establishing timeframe: Will people return to the same window tomorrow? How can I contribute without immediate feedback?
	<i>Audio</i>	Give and take of focus: I can only listen to one voice at a time. When should I stop talking if I can't see whether participants understand me?	Labeling, storing, retrieving, sorting, processing clips: How can I choose the most interesting, relevant clip to listen to?
	<i>Video</i>	Framerate, connectivity: Am I understood? Where do I restart after a technical interruption? How do I know one site has lost video reception and another has lost audio transmission?	Magnification of audio problems: Can I store as much video as participants want to transmit? How do video clips fit into the task completion? What's the minimum video I can view to get the gist?

Table 1: Six modes of virtual meetings

2. Best Practices Findings

The best practices developed by Mittleman et al. [6] are summarized in Table 2. They offer some insight into research questions and tool development for distributed facilitation. The best practices can be grouped by synchronous and asynchronous virtual meetings. In synchronous meetings, participants are geographically dispersed, but meet at the same time. Asynchronous meetings, by contrast, involve participants dispersed by time as well as geography. Global work teams, for example, must be able to operate asynchronously because so many time zones may be involved. For synchronous meetings, some important patterns emerge from the given best practices. These patterns share a common basis: signals we take for granted are missing from virtual meetings. We can't see puzzled faces, hand gestures, body language, ebbs and flows of physical energy and involvement. We have to be more explicit in the absence of feedback signalling understanding or puzzlement, recognition or bafflement. Consequently, the best practices have the following character. Many practices call for more advance planning for virtual meetings as opposed to face-to-face meetings. Many practices call for verbal interjection during virtual meetings. Some practices call for separate channels for content and process communication. It is typical of the practices described that they can be measured after the fact to assess the degree of success. One practice suggests a team dictionary to solve the problem of a shared vocabulary not evolving during the virtual meeting process.

Asynchronous meetings suggest some additional patterns as well as reiterating those for synchronous meetings. One practice suggests a chat window as a back channel for communication about process rather than content. Another practice mandates a meeting scoreboard. Several practices emphasize the importance of explicit instructions.

One practice insists on reminding participants about who is present at any given time. Practices relating to focus suggest visual images of data as well as a wide view of participants.

The above best practices led the authors to seven steps for asynchronous meetings, summarized in Table 3.

3. Impact of Best Practices Findings

Implementation of many of the best practices mentioned above could be distilled to a variant of the ever-unpopular dictum: work harder, accomplish more! Worse still, many implementations could be distilled to the demoralizing dictum: exchange and digest more information!

Simon [9] identifies (as in many of his works) attention as the scarce resource to manage in an information intensive environment. In Simon's view, any addition to information exchange and digest must viewed as a challenge for the management of human attention. In this spirit, the best practices mentioned above may be viewed as challenges to develop implementations to permit participants to accomplish more without relying solely on working harder and challenges to manage attention to alleviate the burden of increased information exchange and digestion.

Are there mechanisms to ease the pain implied by the above best practices? There may be, but further research is needed in this new area. Our aim here is to identify research questions and tools based on prior findings.

The study by Mittleman et al. [7] rests squarely on a link between cognitive load for three kinds of processes and group productivity. The three kinds of processes are (1) communication, (2) deliberation, and (3) information access. The best practices seem to increase the load for (1) communication the most, followed by (3) information access. The theory, called Focus Theory of Group Productivity, is described by Briggs [3].

<i>Lesson</i>	<i>Best Practices</i>
1. Converging from a distance entails different requirements: People can't see restlessness, agreement, sub-group coalitions.	<ul style="list-style-type: none"> ● Explicitly structure the convergence process. ● Hold frequent process checks. ● Use ad-hoc teams to negotiate compromise solutions.
2. Different techniques build a team over a distance: Building loyalty across sites requires bonding across sites. Different time zones have different lunch hours.	<ul style="list-style-type: none"> ● Engage in distributed breaks. ● Create multi-site sub-teams for breakout work. ● Photos and bios in <i>People</i> window matter more.
3. Following a meeting from a distance requires different signals / planning: cues are missing to signal success, completion, transitions. Completion of a major task needs a gesture or tally to give participants the sense of victory.	<ul style="list-style-type: none"> ● Create a scoreboard ● Focus transitions ● Enunciate interim goals ● Make explicit pre-meeting plan
4. Focus wanders easily when working over a distance. Some tasks require inherently more processing from some participants. Bottlenecks occur when skills of subsets are required.	<ul style="list-style-type: none"> ● Engage vested interest ● Assign parallel tasks ● Default video feed should prefer overview to talking heads ● Default video feed should include data affected by process
5. People forget who is at a distributed meeting. Participants sometimes defer discussion until relevant players can be brought together, forgetting relevant players <i>are</i> together!	<ul style="list-style-type: none"> ● Remind them every 10 minutes ● Use the people window a lot ● Reflect user names when facilitating
6. Remote users feel like second class citizens. If one site is perceived as "main" site, morale droops at other sites, participation crumbles.	<ul style="list-style-type: none"> ● Place facilitator at an independent site ● Employ co-facilitators at multiple sites ● Focus local people on remote users
7. Audio channels need different attention in a distributed meeting: Lack of gestures and eye contact makes source and target of message ambiguous.	<ul style="list-style-type: none"> ● Talk to a person, not the group ● Engage in a dialogue with someone you know ● Frequently check on audio quality (intelligibility)
8. Network connections are unpredictable: focus wanders, participation diminishes, and chances of success dwindle during long waiting periods due to disconnects.	<ul style="list-style-type: none"> ● Pre-establish protocols and test ● Anticipate a learning curve ● Establish a re-bootstrap mechanism ● Consider bandwidth cost with multimedia benefit ● Have on-call tech support at each site ● Have a fallback plan
9. Communication channels need to be specified: How should I send a message about a tool malfunction? Which window should I look at to see if we're switching to a new task?	<ul style="list-style-type: none"> ● Separate task and process channels ● Use video to support process rather than document talking heads ● Use process support tools to focus group attention on specific channels
10. New physical environment issues emerge in distributed sessions: Meeting rooms are often designed to promote face-to-face interaction, deterring tool use, leaving remote users at a disadvantage.	<ul style="list-style-type: none"> ● Design to support online focus within rooms ● Support distributed breaks ● Provide online attendance monitoring
11. Asynchronous meetings differ from synchronous meetings.	<ul style="list-style-type: none"> ● Follow the seven steps for asynchronous meetings (Table 3)

Table 2: Problems to Solve and Best Practices

3.1. Communication

The communication construct includes the devotion of attention to choosing words, behaviors, images, and artifacts and presenting them to other team members through some medium. In the reported study, media included a software interface with multiple windows devoted to content, processes, and remote video, and a voice telephone interface.

There is a problem transitioning from face-to-face problem created shifting from traditional meeting to distributed meeting with technology support in that people are used to processing many more signals and have capacity to do so.

The switch from traditional to virtual meetings changes communication, stripping participants of face-to-face signals and customary signal processing conveniences. Two focusing points about communication are

- Over half our brain's neurons are devoted to vision

processing.

- People have experience in face-to-face environments and already have ways of receiving signals present in face-to-face environments.

The preceding points relate to the notion that people do process colossal numbers of signals in face-to-face environments. They suggest that providing an emulation of a face-to-face environment might alleviate the additional communications burden by accessing available processing power that remains unused by present distributed collaboration software.

For example, one best practice is to frequently remind participants of who is and isn't present. There may be less cognitive load in allowing participants to see avatars (visual representations) of each other and to see those avatars enter and leave a room than to create a new type of presentation to reveal that information.

Another best practice calls for participants to be explicit about directing a communication to another individual. It is certainly explicit for a participant's avatar to walk over to an avatar of another participant and speak directly to that avatar.

One best practice is to focus on transitions. If each process has a representative artifact in a visual representation of the workspace, it may be obvious when participants start and finish using that artifact. For instance, if voting is represented by a ballot box, and avatars carry "ballots," we know when everyone votes.

3.2. Information Access

The information access construct includes the devotion of attention to finding, storing, retrieving, and processing information needed by the group.

One best practice calls for the creation of a team dictionary to store an evolving common language. It may enhance the use of the dictionary to have a visual representation of

1. Make sure participants perceive direct vested interest in the task.
2. Make sure all participants agree that there is no easier way to accomplish the task.
3. Make sure users know that management values the output of the task.
4. Correspond in advance with each participant directly to confirm their participation.
5. When possible, begin asynchronous project with a synchronous meeting.
6. Make the pre-meeting plan much more explicit than would typically be the case for a face-to-face meetings.
7. In every tool you use, create an extra place for team members to engage in back-channel communication to the facilitator.

Table 3: Seven steps for setting up Asynchronous Meetings

the dictionary in a shared representation of a workspace. Further, it may help to have a user's avatar appear to interact with the representation of the dictionary as notification that the dictionary has been updated or perused.

4. Toolset Experience

We discuss three kinds of tools we have used to gain insight into requirements for toolsets for distributed facilitation. We are actively developing a fourth toolset, GroupSystems Online, using some of the experience gained. The first toolset we describe, Vizard, a persistent visualization tool, is under active development. The second, a typical unstructured toolset available for free from the Web, is representative of many such simple tools. We did not develop the unstructured toolset, but simply retrieved it from the vendor website and used it for a semester-long project. Third, we discuss a VRML/Java virtual office prototype we developed for study.

4.1. Persistent Visualization: A Meeting Scoreboard

Suppose you enter a ballpark in the third quarter of a tied football game but notice that the time of possession displayed on the scoreboard is highly lopsided. The typical inference in this situation is that one defensive unit is about to collapse in exhaustion and that the team controlling the ball will win. This is just one example of how a scoreboard summarizes an athletic event in a way that shapes our expectations for the future.

We have long recognized the need for tools to fulfill the role of a scoreboard at an athletic event, summarizing the key events to date. An asynchronous meeting, unfolding over time, may require a different kind of summary than a face-to-face meeting. This problem is well illustrated by the new toolset, GroupSystems Online(tm), a web accessible version of GroupSystems(tm) (www.ventana.com).

Two tools demonstrate the problem, People Window and Participation Meter. People Window, shown in Figure 1, shows icons representing participants, along with their names and optional graphics. The icons can be rearranged, for instance, to reflect the positions of participants in a room, and are often used to check the names of unfamiliar participants. The option to add participant photos is less used than icon rearrangement in the authors' experience, perhaps because photos seem less needed in a face-to-face meeting. A distributed toolset needs different elements. The floorplan metaphor for moving the icons around to correspond to room positions lacks motivation in the distributed context, and the inability to link in a URL to a graphic or bio ignores the WWW environment in which the toolset operates.

A subtler example can be found in the participation meter, shown in Figure 2. This attempt at some sort of scoreboard facility shows raw number of contributions over time, and can be customized on the panel shown in Figure 3. One clear limitation here arises from the major typical use of the tool in a face-to-face meeting. The facilitator uses the tool to determine when participation has dropped off to a low level as a signal to transition to another process. In this context, time intervals like 30 seconds, 1 minute, 3 minutes, and 5 minutes make perfect sense.

On the other hand, consider the example in Figure 2. Here, a tool has been opened during a kick-off face-to-face meeting and is to be left open for contributions during the week leading to the next face-to-face meeting. Our experience suggests that the intensity of participation in an asynchronous meeting is positively correlated with recency of face-to-face contact and impending deadlines. If this is true, we need a more detailed look at participation on the day of the kick-off meeting to determine whether people participated from their desktops after the kick-off meeting, and more detailed look at the day of the subsequent face-to-face meeting. Ideally, Participation Meter would allow us to toggle the granularity of analysis for participation: by tool, by activity, by varying units of time, and by participant (if participants don't regard it as Big Brother).

How can we develop a more appropriate meeting scoreboard for an asynchronous meeting? We base one component of such a scoreboard on the notion of persistent visualization. Persistent visualization is a term we have coined to refer to visualization of the relationships between texts (text work objects) arriving in the workspace over time and at too rapid a rate to permit individual perusal. In particular, a successful persistent visualization diminishes cognitive load by emphasizing changes in relationships between texts over time, and deemphasizing unchanging relationships. This toolset allows the persistent visualization to function like a scoreboard, providing a newcomer with a concrete picture of the current state of work in progress.

We have developed a toolset in Java and C to explore persistent visualization. The key concepts are

- *multidimensional scaling* to optimize distances between text objects on a display,
- *animation of hierarchical clustering* to provide user control of the level of detail to be viewed,
- *prototype effects theory*, offering a theoretical guide to permit only a subset of a display to be redrawn and to emphasize change on the display, and
- *integration* with a collaborative writing tool to capture user interaction and provide new output for user workflow.

4.2. Unstructured Toolset

Enormous numbers of unstructured toolsets are freely available across the Web. It is tempting to see how productive we can be with such software. We report one experience with an unstructured toolset that has left us feeling the need for more process support than the typical collection of unstructured whiteboards, chat windows, voting, and brainstorming can provide.

Researchers at CMI have used various software that tries to better resemble face-to-face interaction by opening more channels of communication. One early experience involved five researchers at three sites attempting to do a software design task using an unstructured toolset offering many kinds of unstructured collaborative aids, including whiteboards, chat windows, electronic brainstorming, and voting. The focus of our activity was to learn about the software design process through design of an actual software tool and to use a variety of tools in a joint effort over two months.

The group also worked with various tools in support of object-oriented design, but the only other tools specifically meant to support distributed collaboration were email and video and phone conferencing, all three of which were used more extensively than the unstructured toolset.

Participants reported that they were not inspired to use the unstructured toolset because (1) they could not easily adjust their workstyles to suit the tools, (2) they felt no pressure to use the tools because whenever they logged in they found that no one else was logged in, (3) because of connection difficulties due to use of a non-tested operating systems (NT 3.51 hadn't been verified by the toolset vendor and it was not possible for this participant to switch to NT4.0, Win95, Macintosh OS, or Unix) and possibly firewalls (the same participant was behind a firewall and had no access to non-firewall systems with which to compare), and (4) because not everyone could use the toolset due to (3).

Notably, these problems can be overcome in a software enhanced face-to-face interaction [8]. Some research has

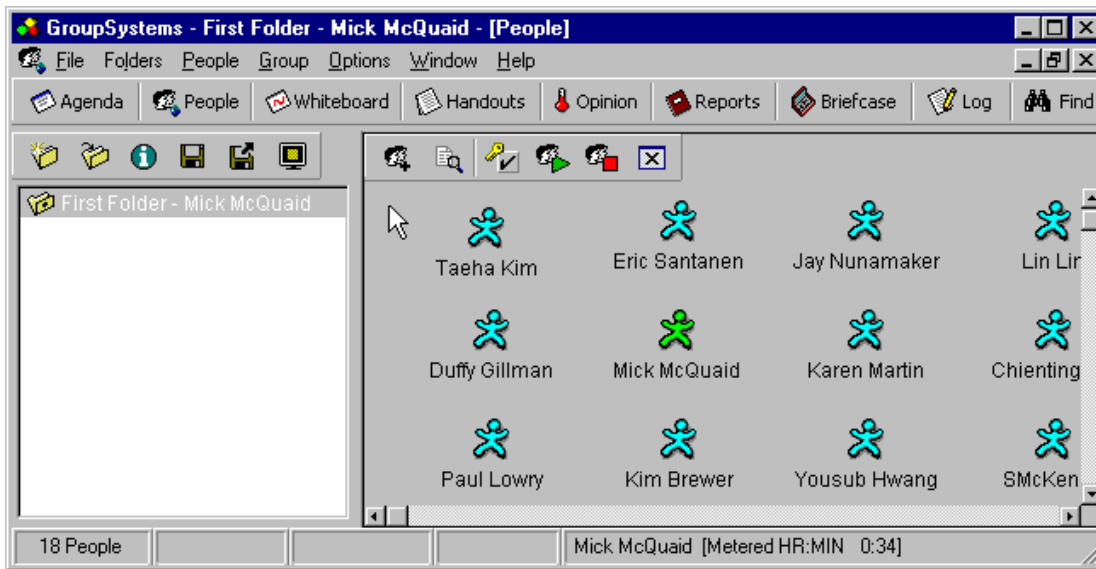


Figure 1. The People Window in GroupSystems Online(tm)




Figure 2. The Participation Meter in GroupSystems Online(tm), set to 1 day intervals

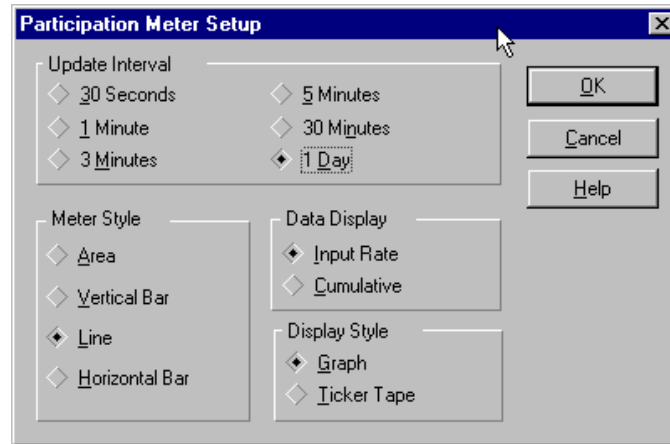


Figure 3. The Participation Meter Settings Window in GroupSystems Online(tm)

explored reason (2) above, lack of peer pressure [1]. Unfortunately, we did not explore the unstructured toolset in enough depth to determine whether the difficulty with adjusting to the tool was due to an inherent process problem or was more an artifact of the tool. At least one participant liked the tool and did not ascribe the difficulties in getting used to using the tool to its design. This participant emphasized the problem of lack of awareness of other participants, their status and focus.

It is also notable that all participants preferred phone conferencing with email of documents as the primary support for distributed collaboration. Comments by participants were typically either about feeling in control of these technologies and feeling that it was possible to obtain necessary contact to determine things that couldn't be spelled out by phone.

One participant comment captivated us more than others and spurred much of the subsequent work. The comment was that interacting with the toolset took too much attention for any attention to be left over to interact with other participants.

We concluded on the basis of this last comment that it may be that the software approach of presenting a rich interface also presents a cognitive load that limits productivity. For our group, doing software design meant learning some new object oriented design tools and the time we spent learning any new tool was precious. We found that each of us felt more stimulated by the presence of others online, but that we often couldn't make the tool do what we wanted quickly enough to engage our teammates while online.

The above experience has led us to investigate the link between emulating face-to-face interaction and increased productivity through reduced cognitive load. We would like to know whether reducing the cognitive load of interacting with the system will lead to increased productivity.

4.3. Virtual Reality Toolset

Virtual reality environments [5] contain objects analogous to some we've already experienced. For instance, a picture of a telephone in a VR environment can be reasonably assumed to offer a link to a two way communications tool. A picture of a mailbox can be reasonably assumed to permit the deposit and retrieval of email. A picture of a room full of people around a conference table can be reasonably assumed to indicate a meeting in progress. The pictures of people in the conference room can be assumed to represent the real people they resemble.

Representations of objects and phenomena, like those described above, use connections already present in our minds (if the objects and phenomena are familiar - the environment may do nothing for a participant from a different culture). By using connections already present, analogy reduces the need to form new connections, reducing cognitive load.

We can easily remember that our boss has the corner office on the top floor and can easily find our way there from the day we start work. The cognitive load to go there is less than to send the boss email with Mime compliant attachments.

A receptionist or uniformed security guard can be assumed to be able to offer guidance about the locations of other participants. There is less cognitive load in tapping one of these official representatives on the shoulder to ask for directions than to find a menu selection for participant location.

In our software design task described above, we found that we had a very limited amount of time to achieve our objectives and spent much of that time focused on the software. We came to believe that we would have been more



Figure 4. Library with open book on floor and avatar

productive if we had not been required to think about the remote collaboration software so much. In any task limited by time, attention we can divert toward the task goal is a precious resource.

4.4. Initial VR Experience

We developed a prototype virtual reality environment and used it, arriving at some conclusions about where to go from here. Now we describe our experience.

We constructed a virtual reality office prototype using VRML and Java and tested it using SGI Cosmo Player and Sony Community Place Browser [4] on several networked computers. We connected three computers using Sony Community Place Bureau to enable interactions between avatars. Community Place has built-in multi-user functionality, including customizable avatars. On-screen avatar controls include facial expressions for happiness, sadness, puzzlement, and anger and related body positioning, as well as other appearance shortcuts.

We constructed offices for a fictional magazine, *Popular Driving*, and included publicly accessible multimedia material from automobile manufacturers, as well as tools for remote participants to choose feature articles. The prototype included several office features:

- A receptionist able to exchange messages and give directions to inner offices. The receptionist could report on previous logins by editorial staff. See figure 6.
- A library containing bookshelves and books which can be activated to obtain access to a document search and browse tool. These books included promotional material about cars under consideration for editorial review. There was also an avatar for a product representative from one manufacturer. See figure 4.
- A meeting room with a conference table, chairs, a gavel, and ballot box. The gavel can be activated to conduct a vote. The ballot box can be activated to participate in a vote. These were attached to a voting tool on another server configured with possible articles to include in the magazine. See figure 5.
- Individual offices containing mailboxes.
- A special room, without analogy in the physical world, for testing visualization tools.

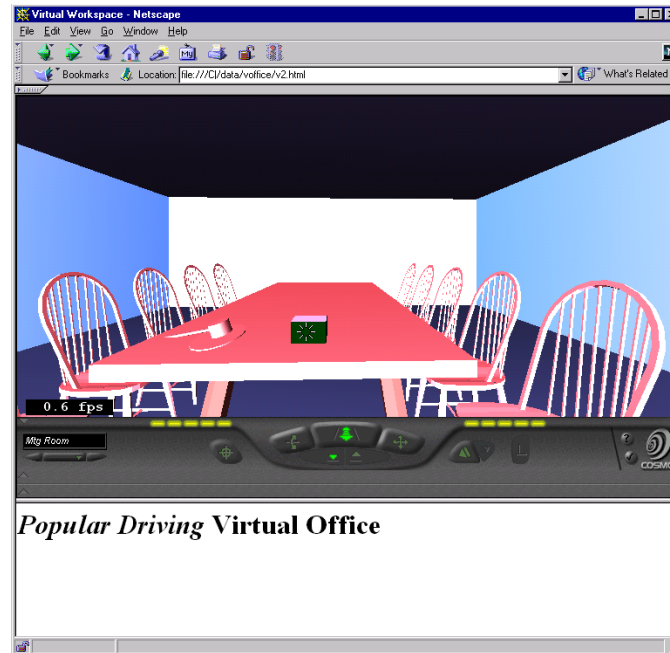


Figure 5. Conference room with voting tools

4.5. VR Participant comments

The prototype was quite unstable for our first investigation. Our initial prototype's instability was a major focus of participant comment. In contrast to the unstructured toolset, we were not able to apply our first prototype to a large task like software design. This instability prompted us to recruit fellow researchers as participants in a cooked task: choosing editorial content for a fictitious magazine. These participants were willing to put up with far greater instability than most users, so we're not ready to generalize our results beyond them.

Our first prototype ran only on machines that provided a framerate of less than 2 frames per second. Participants repeatedly mentioned this slow framerate and the resulting jerky motion. Other than instability, this was the major technical issue participants focused on.

We noticed that participants constantly tested the boundaries of the system. For instance, we set defaults such that participants could not walk through walls, but we did not disable participant ability to change this default. Participants frequently changed defaults to enable them to walk through walls, to alter the lighting that surrounded them, and to permit them to sink partly into the floor so that they could walk with only their heads and upper torsos showing. Some participants even turned off gravity settings so that they could fly. Participants frequently commented on these adjustments, comparing notes on how they could alter the

environment. Many of these comments had a tone of control or power ("... yeah, I could make it let me fly."), as if the participant obtained satisfaction by controlling environmental settings and finding the limits of those settings.

Participants frequently made use of the menu selections for facial expressions and body positions. Participants commented that they could easily determine the mood of fellow participants, and felt that they could easily communicate their own mood. Participants appeared to devote considerable focus to refining their own communication. One on-screen selection in the Community Place browser permits the participant's avatar to jump up and down, extending its arms. Jumping up and down became a common attention-attracting behavior. Participants often quickly moved the mouse between a facial expression and a body position, saying later, for instance, "... I kept jumping up and down and frowning until you noticed."

This last group of comments speaks most directly to the issue of reducing cognitive load, so we conclude from this initial experience that we need to build another prototype that restricts the other areas to focus participants better on their interactions.

We conclude that requirements for avatar and virtual meeting rooms include shortcuts, allowing participants to concentrate on what they do rather than how they do it, shortcuts that leave the avatar to perform an analogical task, while permitting freedom for the represented person. An example would be that when a person wish to address the




Figure 6. Visiting the virtual receptionist

crowd, he selects a megaphone. Meanwhile, his avatar walks to a megaphone, picks it up, and moves it from side to side while addressing other avatars. Having selected the megaphone, the avatar owner need not engage in detailed behaviors to manipulate the avatar. Yet the cognitive load for other participants is reduced by the avatar's behavior.

5. Further Work

Our experience with distributed meetings gives us a basis to further develop and refine our prototypes, both in persistent visualization and virtual reality. The forthcoming release of GroupSystems Online promises to let us test our ideas on a system offering extensive process support and to integrate our "meeting scoreboard." Our long-term goals include creation of a more comfortable environment, eliciting recognition of the workspace as a shared space, with the feel of interacting in a shared physical space, a familiar space, a space you can return to and feel as if you're "there."

References

- [1] M. S. Ackerman and B. Starr. Social activity indicators for groupware. *IEEE Computer*, pages 37–42, June 1996.
- [2] C. M. Beise, F. Niederman, and P. M. Beranek. Group facilitation in a networked world. *Group Facilitation: A Research and Applications Journal*, 1(1):1–2, 1999.
- [3] R. O. Briggs. The focus theory of group productivity and its application to the development and testing of electronic group support systems. Technical report, Management Information Systems, University of Arizona. Tucson, AZ., 1994.
- [4] Y. Honda, K. Matsuda, J. Rekimoto, and R. Lea. Virtual society. In *Proceedings of VRML '95*, pages 109–116. ACM Press, December 1995.
- [5] A. Kamble and R. Arora. Virtual reality - a technology and applications review. *Electronics information and planning*, 25(7):360–382, 1998.
- [6] D. D. Mittleman, R. O. Briggs, and J. F. Nunamaker, Jr. Best practices in facilitating virtual meetings: Some notes from initial experience. *Group Facilitation Journal (forthcoming)*, 2000.
- [7] D. D. Mittleman, R. O. Briggs, J. F. Nunamaker, Jr., and N. Romano, Jr. Lessons learned from synchronous distributed GSS sessions: Action research at the US Navy Third Fleet. *Proceedings of the 10th EuroGDSS Workshop, Copenhagen*, pages 63–79, June 1999.
- [8] J. Nunamaker, Jr., R. Briggs, and D. Mittleman. Electronic meeting systems: Ten years of lessons learned. In D. Coleman and R. Khanna, editors, *Groupware: Technology and Applications*, chapter 7, pages 149–193. Prentice-Hall, Englewood Cliffs, NJ, 1995.
- [9] H. A. Simon. Knowledge and the time to attend to it. *International Conference on High Performance Global Corporations*, April 1995.
- [10] D. R. Woolley. Conferencing software for the web. <http://thinkofit.com/webconf>, 1999.