

Mixed Presence Collaboration using Scalable Visualizations in Heterogeneous Display Spaces

Thomas Marrinan
Argonne National
Laboratory
Lemont, IL, USA
tmarrinan@anl.gov

Jason Leigh
University of Hawaii
at Mānoa
Honolulu, HI, USA
leighj@hawaii.edu

**Luc Renambot, Angus Forbes,
Steve Jones, Andrew Johnson**
University of Illinois at Chicago
Chicago, IL, USA
{renambot, aforbes, sjones, ajohnson}@uic.edu

ABSTRACT

Mixed presence collaboration involves remote collaboration between multiple collocated groups. This paper presents the design and results of a user study that focused on mixed presence collaboration using large-scale tiled display walls. The research was conducted in order to compare data synchronization schemes for multi-user visualization applications. Our study compared three techniques for sharing data between display spaces with varying constraints and affordances. The results provide empirical evidence that using data sharing techniques with continuous synchronization between the sites lead to improved collaboration for a search and analysis task between remotely located groups. We have also identified aspects of synchronized sessions that result in increased remote collaborator awareness and parallel task coordination. It is believed that this research will lead to better utilization of large-scale tiled display walls for distributed group work.

Author Keywords

Mixed presence collaboration; data-conferencing; multi-user interaction; large-scale displays.

ACM Classification Keywords

H.5.3. [Information Interfaces and Presentation]: Group and Organization Interfaces: Computer-supported cooperative work, synchronous interaction, evaluation/methodology.

INTRODUCTION

The increased volume and distributed nature of data that is being generated, collected, and stored are forcing distributed groups to pool resources and knowledge for proper analysis. This has led to the need for better mixed presence collaboration frameworks for sharing and interacting with data-intensive visualization content [12,13]. Successful mixed presence collaboration would allow groups of domain experts to seamlessly communicate and interact with each

other without the need to travel or transfer large volumes of data. Such collaborations would ideally take place in existing workspaces rather than requiring dedicated conference rooms that solely support mixed presence collaboration.

Existing research on mixed presence collaboration primarily focuses on networked tabletop displays [3,13,19]. However, when we surveyed the emerging landscape and studied prevalent technologies used for real-world applications, tabletop displays were only used in specialized circumstances. Vertical displays, on the other hand, are ubiquitous and large-scale tiled display walls are becoming more common – currently found in industries such as financial institutions and government agencies in addition to high-end visualization laboratories. Since each workplace or discipline has unique display technologies that best suit its needs, it is of increasing importance to facilitate effective remote collaboration between heterogeneous display spaces.

For our study, we used the SAGE2 framework (the Scalable Amplified Group Environment) [15,18] to evaluate mixed presence collaboration between large-scale tiled display walls. SAGE2 is a middleware that creates a scalable virtual desktop on screens ranging from a single monitor to a cluster-driven tiled display wall. Multiple users can view and simultaneously interact with an assortment of data-intensive information on a shared display. SAGE2 also has basic remote collaboration support – allowing users to send local content to a distal SAGE2 site. This *data-pushing* technique sends a copy of an application to a remote site where each group interacts independently without affecting the remote instance (similar to sending an email with a document attached that the receiver opens immediately). Figure 1 shows a SAGE2 session and illustrates how content can be pushed to remote sites.

In order to enhance remote collaborator awareness and parallel task coordination, we explored two new techniques for sharing data between display sites: *data-duplication* and *advanced data-synchronization options*. Using the *data-duplication* technique, one section of the large-scale display contains local unsynchronized versions of private data-conferencing content and a second portion of the display contains fully synchronized copies of public data-conferencing content. This technique partitions the large-scale display into two spaces – public content that would be

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of the United States government. As such, the United States Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

CSCW '17, February 25-March 01, 2017, Portland, OR, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-4335-0/17/03...\$15.00

DOI: <http://dx.doi.org/10.1145/2998181.2998346>




Figure 1. SAGE2 multi-user windowing environment running on a large-scale tiled display wall. Left – documents, interactive applications, and shared laptop screens displayed on the shared display. Right – tabs representing remote SAGE2 sites where content can be sent (dragging and dropping any application to a remote site’s tab will *push* the content to their display).

perfectly mirrored with remote collaborators and private content that would be viewed and controlled independently.

Using the *advanced data-synchronization options* technique, collaborators choose which aspects of each shared application are synchronized and which are controlled independently at each site. This technique allows for independent window management and the potential for partial synchronization (in a manner similar to Google Docs [9] where text is synchronized, but each collaborator’s view of the document is independent).

In this paper, we provide the methodology and results of a comprehensive user study to examine distinct remote collaboration strategies between heterogeneous display spaces. We show that continuous synchronization, using *advanced data-synchronization options* or *data-duplication*, is superior to *data-pushing* for complex visual analytics tasks. We also extract specific generalizable features from each of the continuous synchronization techniques that contribute towards improved collaborator awareness and parallel task coordination in a mixed presence setting.

RELATED WORK

Aspects of Collaboration

In general, groupware can be designed to facilitate collaboration in three distinct ways: through communication, cooperation, or coordination [7,8]. Communication is an unstructured exchange of information, such as phone calls or instant messaging chats. Cooperation represents multiple users interacting while working on a common goal, such as conducting a brainstorming session. Coordination refers to separate but interdependent work between multiple users, such as developing a project timeline and budget.

While the three C’s of collaboration interrelate and feed into one another, coordination is the most beneficial for task execution, followed closely by cooperation [14]. Communication, as well as non-collaborative independent work, are both necessary but have been found less beneficial for efficiently completing group tasks. Therefore groupware systems for accomplishing tasks should aim to facilitate reducing independent work and one-on-one communication,

while maximizing the potential for group cooperation and parallel task coordination.

These prior works helped form the first main goal of our research – improving parallel task coordination for mixed presence collaboration.

Collaborator Awareness

Another important aspect of collaboration deals with collaborator awareness. Gutwin and Greenberg highlight the importance of workspace awareness (where collaborators are and what they are doing) in a groupware system [10]. Their framework is designed to help understand the importance of how awareness can be incorporated into groupware systems. The authors emphasize knowledge of the people users are working with, when actions are taking place, and what collaborators are doing. The authors believe that these elements form the core of any groupware system and are considered important to capture and distribute while developing groupware.

Achieving uniform workspace awareness for all collaborators in a mixed presence setting faces the additional challenges of presence disparity and display disparity [19]. Presence disparity occurs because users are naturally more aware of their collocated collaborators than of their remote collaborators. Display disparity can cause users at each site to best utilize their local technology without regard for the technology at the remote sites.

Ocker et al. conducted a study on how to work effectively in a “partially distributed team” (PDT) – a form of mixed presence collaboration where the entire group is working on the same task [17]. Their study identified six key factors for successful long-term collaboration of a PDT. Two of these factors were identified as procedural aspects of team management – awareness and coordination – and therefore considered imperative for successful groupware to address.

McEwan et al. present an awareness model for mixed presence collaboration based on prior work in both collocated collaboration and remote collaboration [16]. Their model consists of relationships between four entities in mixed presence collaboration – people, artifacts, workspace, and sites. They identified the workspace-site relationship as

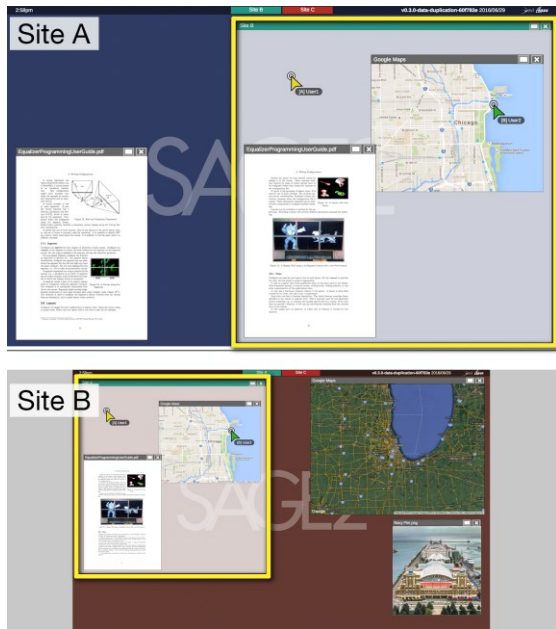


Figure 2. Synchronized content sharing between two sites using the *data-duplication* technique. A perfectly mirrored portal is displayed on each site's SAGE2 wall (highlighted in yellow). Each site can independently position and scale the portal as a whole, but relative size and position of content inside remains the same. Private content, located outside the portal, can be shared and made public by dragging and dropping the window inside the portal.

problematic primarily due to display disparity and suggest providing feedback about remote technologies to reduce conflict.

Due to the critical nature of providing and maintaining awareness during a collaborative session, the second main goal of our research was formed – to improve collaborator awareness in the context of heterogeneous display spaces.

Large-scale Displays for Collaboration

Finally, we looked at prior research in visualization and human-computer interaction to determine an appropriate technological medium for fostering collaborative task execution. There is now conclusive evidence that large display environments significantly amplify the way users make sense of large-scale, complex data [1,2,4].

Additionally, large-scale displays provide affordances for group interaction that can enhance collocated collaboration. Biehl et al. developed and tested a dashboard visualization system on a large-scale shared display that provided information about coworkers and project resources [5]. They discovered the large-scale dashboard enhanced collaborator awareness and increased communication.

Birnholtz et al. conducted a study relating to input device configuration for collocated collaboration on a shared display [6]. Their study involved three people completing a negotiation task in two scenarios – sharing a single mouse pointer or using one pointer per person. Their results found

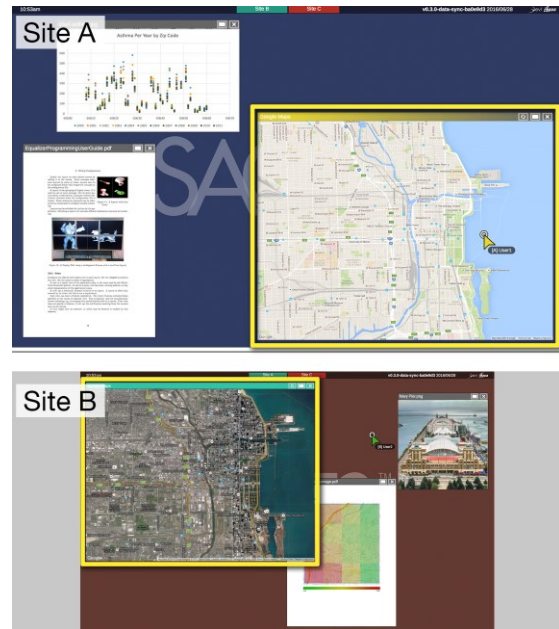


Figure 3. Synchronized content sharing between two sites using the *advanced data-synchronization options* technique. The shared mapping application (highlighted in yellow) has its map type unsynchronized, but the pan and zoom still synchronized. Therefore, both sites are looking the same location, but Site A is viewing roadmaps whereas Site B is viewing satellite images.

that having multiple pointers increased parallel work, but led to less group communication. The authors suggest carefully considering input devices for collaborative activities depending on the desired group behavior.

Since analysis tasks involve complex data, mixed presence collaboration involves collocated groups, and large-scale tiled display walls are being adopted at growing rates, we decided to focus our mixed presence collaboration research around groups using large-scale shared displays.

DATA SHARING AND SYNCHRONIZATION FOR HETEROGENEOUS DISPLAY SPACES

We extended SAGE2's remote collaboration abilities in an attempt to improve parallel task coordination and collaborator awareness by continuously synchronizing shared applications on large-scale tiled display walls. We developed two new data-conferencing techniques to handle shared applications across remote SAGE2 sites – *data-duplication* and *advanced data-synchronization options*. The *data-duplication* technique creates a fully synchronized partition that perfectly mirrors content and interaction icons on each site's display. Figure 2 shows two sites collaborating using *data-duplication*, with shared content fully synchronized and private content kept independent.

Using the *advanced data-synchronization options* technique enables partially synchronized shared applications, where users from either site can control which aspects of an application are continuously synchronized and which aspects are controlled independently. Figure 3 shows how

collaborators could use a partially synchronized shared application using *advanced data-synchronization options*.

We compared these two techniques to SAGE2's existing remote collaboration technique of *data-pushing* – sending a copy of an application and its current data to a remote site without any further synchronization.

Hypothesis

Due to the previous research on collaborator awareness with shared applications, parallel task coordination in groupware systems, and the use of large-scale displays, we formed three hypotheses surrounding the data-conferencing techniques we had developed.

- H1. Both *data-duplication* and using *advanced data-synchronization options* would result in significantly improved collaboration compared to *data-pushing*.
- H2. Effective data sharing strategies would mitigate communication issues and disparity between groups involved in mixed presence collaboration to more closely resemble collocated collaboration.
- H3. Users of a smaller shared display would prefer and perform tasks more effectively using *advanced data-synchronization options* compared to *data-duplication*.

User Study Evaluation

In order to evaluate the effectiveness of these new remote collaboration techniques, we designed a formal user study that would require mixed presence collaboration between two groups. In our study, one group used a larger shared display with an ultra wide aspect ratio (9.17m x 1.72m with a 12294 x 2304 resolution) and the other group used a smaller shared display with a standard widescreen aspect ratio (1.86m x 1.05m with a 3840 x 2160 resolution). Each trial consisted of four participants (two per group). The two groups of participants represented teams from different domains that had a different set of knowledge; each team was required to achieve a common search and analysis goal.

In order to communicate between groups, two video feeds and one audio feed were shared from each site. One video feed came from a camera pointing at the participants so that remote collaborators could see each other. The second video feed came from a camera pointing at the shared display so that remote collaborators could gain awareness about the remote site's technology and visual layout of data.

The fictional problem participants were asked to solve was to find an ideal location to open a new coffee shop within a given city. The problem was split into two tasks – first, to come up with 2-4 potential locations based on a separate set of constraints for each group, and second, to determine the best location from the original selections based on additional information. Each team received separate “prior knowledge” in the form of data printed on a sheet of paper. This served as information known by one group that could not be directly shared digitally using SAGE2 in order to better represent authentic distance collaboration between experts in various

domains. Each group repeated the task three times, once using each data sharing technique.

The task of determining locations for a new coffee shop was chosen since it would not require extensive training for participants to grasp the underlying concept. Recruited participants were required to be regular desktop or laptop users, and have familiarity with mapping software. A total of 44 users participated in this study (33 male, 11 female; ages 18-59) and no compensation was given to volunteers.

User Study Data

In order to evaluate the effectiveness of collaboration, we decided to monitor participants' perception of collaboration by conducting post-use surveys, and measure the quality of collaboration by cataloguing each session with audio/video recordings and user interaction logs. We also measured completion time and task accuracy, but detailed results are omitted from this paper since they do not necessarily correlate with level of collaboration.

The surveys asked participants to score three different aspects relating to the group collaboration on a scale of 1 to 10. First, participants were asked how easy it was to use each data sharing technique and the local collaboration features. This question measured the level of intuitiveness of each data sharing technique. Second, participants were asked how successful each data sharing technique, as well as the local collaboration features, was at facilitating collaboration. This question measured participants' perceptions of the quality of collaboration enabled by each data sharing technique. Finally, participants were asked how much they liked the remote collaboration features. This question was designed to gain insight on how users felt most comfortable collaborating with a remote group. In addition to scoring the three aspects of the data sharing techniques, participants were asked to rank the three data sharing techniques from best to worst.

Videos were recorded so that they could be coded into collaboration modes that each group was engaged in. Collaboration was broken up into four distinct modes: not collaborating, communicating, conferencing, or coordinating. We slightly modified the three C's of collaboration, replacing cooperating with conferencing, to better capture collaboration between the mixed presence groups. By using conferencing as our measure, cooperative tasks only between collocated teammates would not count as higher quality collaboration. Coding the videos provided insight into what percentage of time the groups spent in each collaboration mode, as well as how often the two groups were in the same collaboration mode as each other (which served as a presence disparity indicator).

All participant interaction with SAGE2 was logged in order to create a timeline and look for recurring interaction patterns. The logs consisted of data for each participant on when they moved/resized a window, interacted inside an application, or shared an application with the remote site.

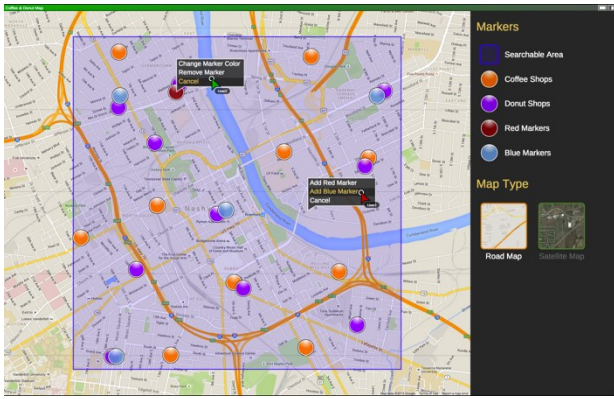


Figure 4. Mapping application used in our user study. Multiple users could simultaneously interact to manipulate the map or add/remove markers.

User Study Methods

Four participants were used per trial, split into two teams of two. Team A was assigned to the larger shared display and had constraints relating to other coffee shops, main streets, and highways. Team B was assigned to the smaller shared display and had constraints relating to donut shops, parking lots, and building roof color. These constraints were chosen since they are easily visible on a mapping application in either roadmap view or satellite view, not due to authentic relevance in determining an ideal location for a coffee shop.

Ten fictional coffee shop and ten fictional donut shop locations were chosen within an approximately 16 km² area of a city in the United States of America. Real features on the map were used for data about roads, parking lots, and building roofs. Four pieces of fictional data were also generated to act as the “prior knowledge” for each team. These pieces of data related to each teams’ unique constraints and modified what was depicted on the map, such as a parking lot being demolished or a new coffee shop opening. The locations of the coffee shops and donut shops were chosen such that there would be five locations in the city that would satisfy all constraints for both teams.

Fictional data was generated for the additional information required to complete the second task of choosing one final location to open a new coffee shop. Team A was given data about crime and storm damages in 16 areas of the city and Team B was given data about family income and business profits in 16 areas of the city. The values for these four variables were chosen such that there would be one optimal area in the city for the final coffee shop.

Data was generated for a total of six cities that were chosen since they were mid-sized urban areas that were not at the location where the study took place.

In order to complete both tasks, the participants were given a custom multi-user mapping application to use within SAGE2. The map could be toggled between two different views – roadmaps and satellite imagery. The map could also be panned and zoomed interactively. A semi-transparent



Figure 5. Teams of two engaged in mixed presence collaboration, working on finding a location to open a new coffee shop.

blue rectangle depicted the approximately 16 km² searchable area within the city. Orange circles were used to denote the locations of existing coffee shops. Purple circles were used to denote the locations of existing donut shops. Participants were able to interactively add and remove red and blue markers. There was no inherent meaning to the different colored markers; each group had to come to a consensus on how they wanted to encode data. A context menu could be used to add, remove, or change the color of a custom marker. Also, a panel on the side of the application could be used to toggle the visibility of any feature. Figure 4 depicts two users interacting with the custom mapping application used in this study. Each participant had his or her own mouse pointer, which was capable of interacting with the mapping application.

Each site had a facilitator who was an expert with the software to perform high-level tasks (e.g. open another map, share a PDF with the other site, etc.) upon request by the participants. Prior to starting the study, each group was allowed to interact with the mapping application, which had sample data using a separate city. This session lasted approximately 5 minutes, and enabled participants to get familiar with the controls of the mapping application. To complete the study, each group of subjects repeated the tasks three times, once for each data sharing technique. A brief training session on how to use the collaborative features of the upcoming data sharing technique preceded each set of tasks. Each time the type of data was the same, but in a different city with different locations of existing coffee shops and donut shops.

The order in which teams used each data sharing technique was counter balanced across test groups as well as which city was assigned to which technique. A within-subjects design was implemented for the user study using three test cases. The entire duration of the study for one group lasted approximately 2 hours. Figure 5 depicts the setup of teams and displays at both sites.




Figure 6. Overall results of the participant survey. Answers were scored on a scale of 1-10, with 1 being worst and 10 being best.

Video Coding and Verification

Audio/video was recorded by using a screen recording software in order to capture the video conferencing window from the camera pointed at the shared display, which also captured the backs of the participants. Videos from both teams were superimposed in the same frame, and audio from each team was isolated with one going to the left channel and the other going to the right channel. This made it easy to view the entire group while also being able to distinguish conversations from each team independently. Videos were coded for collaboration mode at each location – ‘not collaborating’, ‘communicating’, ‘conferencing’, or ‘coordinating’. Additionally, due to technical errors with the audio/video recording, there were a few stretches of time where the collaboration type was unknown.

In order to determine collaboration mode, strict definitions were given for each collaboration mode. A coder would mark down the time and mode of collaboration at any point in the video that the mode changed. Each team was coded separately, since they were not necessarily engaged in the same mode of collaboration the entire time. The definitions for the collaboration modes were as follows:

- *Unknown* – audio/video not available. (
- *Not collaborating* – participants on one side are not (collaborating (participants are silently waiting, talking to each other about something off topic, or working independently without regard for the group). (
- *Communicating* – participants are communicating about the task one-on-one. (
- *Conferencing* – three or more participants are engaged in a discussion (either talking or actively listening). (
- *Coordinating* – the two teams are working in parallel after they have agreed to split the task.

After coding all the videos from the user study for collaboration modes, a second investigator coded one trial (all three runs for both the team using the larger shared display and the team using the smaller shared display) for inter-coder reliability verification. The trial chosen for verification was selected due to the fact that it contained no errors leading to unknown collaboration types, and that the participants engaged in all modes of collaboration at some point during the trial. The coded collaboration modes were discretized into each second of the trial. The overall percent agreement between coders for the entire trial was 93.0%.

Krippendorff's Alpha [11] was also calculated as a more stringent measurement of inter-coder reliability, with values $\alpha \geq 0.800$ corresponding to a reliable coding, values $0.800 > \alpha \geq 0.600$ corresponding to a tentatively reliable coding, and values $\alpha < 0.600$ corresponding to unreliable coding. The overall value of Krippendorff's Alpha for the coding of collaboration modes for the entire trial was 0.897.

EFFECTS OF VARYING DATA SHARING TECHNIQUE BETWEEN REMOTE COLLABORATORS

This section presents the results of the user study. First, the overall results of all trials will be covered, which include participant surveys and audio/video analysis for collaboration mode along with the user interaction logs. Next, we break down our results by display size in order to determine its effect on mixed presence collaboration.

Overall Survey Results

The survey for the formal user study consisted of twelve questions. For the first eleven questions, participants were asked to score certain aspects of collaboration on a scale from 1 to 10 (1 being worst and 10 being best). These questions were about the ease of use, successfulness of collaboration, and how much the participants liked the collaborative features. The twelfth question asked participants to rank the data sharing techniques from best to worst. Results of question 1-11, depicted in Figure 6, show user preference for continuous synchronization, with *advanced data-synchronization options* scoring highest.

A one-way ANOVA test was used to determine significance for survey responses comparing the three data sharing techniques with each other as well as local collaboration. The significance level, or p-value threshold, was set to 0.05. Responses to all three questions yielded significant differences, with ease of use, successfulness of collaboration, and how much the participants liked the collaborative features having p-values of 0.0453, 0.0002, and 0.0052 respectively. These results mean that not only were *data-duplication* and *advanced data-synchronization options* consistently scored higher than *data-pushing*, but that the differences between them were significant. Additionally, a two-tailed t-test was used to compare local collaboration to *advanced data-synchronization options* (the top scorer for remote collaboration techniques). This showed a significant difference in ease of use (p-value of 0.0030), but no significant difference in the successfulness of collaboration (p-value of 1.0000).

Users also ranked the data sharing techniques, with 1 being the best, and 3 being the worst. Results, depicted in Table 1, show that *advanced data-synchronization options* was ranked the best, closely followed by *data-duplication*. A one-way ANOVA test was used to compare the ranking of the data sharing techniques with each other and resulted in a p-value less than 0.0001, well below our 0.05 threshold.

Data Sharing Technique	Rank (1-best, 3-worst)
<i>Data-pushing</i>	2.50 ± 0.09
<i>Data-duplication</i>	1.80 ± 0.12
<i>Advanced data-synchronization options</i>	1.68 ± 0.12

Table 1. Average overall rank for the three data sharing techniques, with 1 being the best and 3 being the worst.

Overall Audio/Video and User Interaction Results

The videos containing the teams from both the larger shared display and smaller shared display were analyzed and coded into sections of time based on the mode of collaboration that each side was partaking in. Each team was coded separately, since the entire group was not always engaged in the same mode of collaboration at the same time. Collaboration modes for each technique were normalized based on the length it took to complete the task, excluding any time there were technical errors causing an unknown collaboration mode, so that each trial carried an equal weight rather than teams who took longer having a larger impact on the averages. In general, using *advanced data-synchronization options* resulted in the most time spent in higher modes of collaboration. Overall collaboration mode results are depicted in Figure 7.

A one-way ANOVA test was used to determine significance for time spent in each collaboration mode. As with the surveys, the significance level was set to 0.05. When using continuous synchronization techniques, the decreased time spent ‘not collaborating’ and increased time spent ‘conferencing’ were significant, with p-values of 0.0498 and 0.0020 respectively. However, the differences in ‘communication’ and ‘coordination’ were not found to be significant, with p-values of 0.2722 and 0.2983.

In order to further analyze collaboration modes, we investigated how often the two teams were in the same collaboration mode as each other. For this analysis, it did not matter which collaboration mode the teams were in, but only if they were in the same mode as each other or not. When the two teams were in the same collaboration mode more often, it signaled a decrease in presence disparity and therefore improved awareness of remote collaborators. The amount of time spent in the same collaboration mode for each trial was normalized based on the length it took to complete the task.

When using the *data-pushing* technique, the two teams were in the same collaboration mode as each other 69.0% (±4.6%) of the time. When using the *data-duplication* technique, the

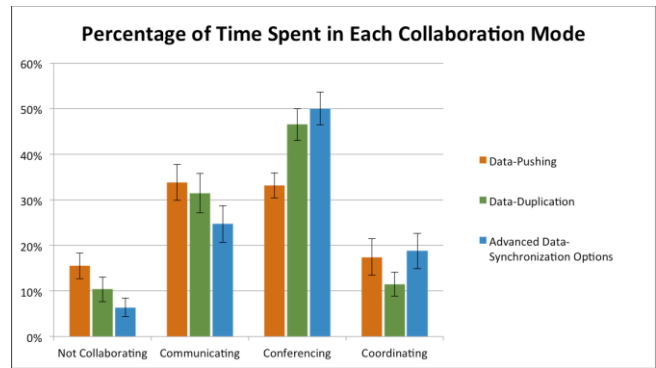


Figure 7. Average overall percentage of time spent in each collaboration mode based on data sharing technique.

two teams were in the same collaboration mode as each other 77.4% (±4.5%) of the time. When using the *advanced data-synchronization options* technique, the two teams were in the same collaboration mode as each other 87.7% (±3.6%) of the time. Using continuous synchronization, especially with the *advanced data-synchronization options* technique, resulted in both teams being in the same collaboration mode for a greater percentage of time, and therefore having a decreased presence disparity.

A one-way ANOVA test was used to determine significance for collaboration mode similarity between the two sites. Overall differences for when both sites were in the same collaboration mode had a p-value of 0.0158 and were therefore determined to be significantly different.

Finally, we combined the collaboration mode coding from the videos with the user interaction logs in order to create a visualization for viewing the timeline data from each trial and to investigate for patterns of interaction. In order to create a timeline view visualization, we took the data from both teams and stacked them on top of each other. First, there is the user interaction log data, where a tick mark is drawn for each type of interaction by each user at the corresponding time in the visualization. The width of the tick mark does not correspond to any data, since it only represents a single point in time. Second, there is the collaboration mode data, where colored blocks represent chunks of time the team spent in each mode. The result shows how each team interacted while in each mode of collaboration, as well as how they collaborated with each other. These visualizations were created using normalized time (each timeline is the same width regardless of task length). This elucidated a couple of interesting patterns for mixed presence collaboration.

First, some groups depicted a turn-taking pattern, where only one of the teams worked at a time and data was sent to the other team when the turn changed. An example of this is shown in Panel A of Figure 8. This trial started with the team using the larger shared display interacting with an application while the team using the smaller shared display was simply talking but not interacting with any application. After a period of time, the roles reversed, with the team using the large display no longer interacting with any application,

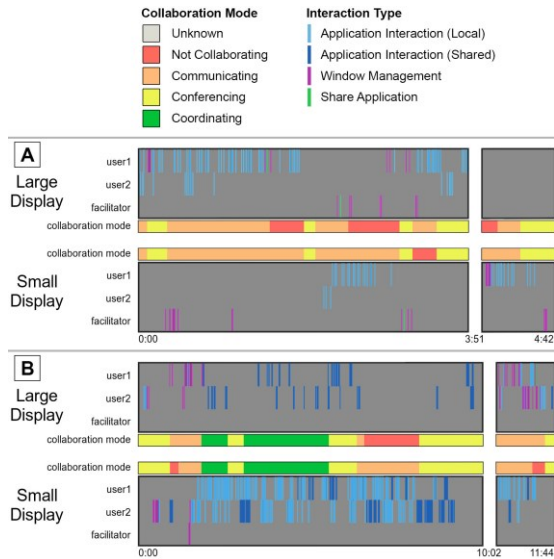


Figure 8. User interaction log and collaboration modes coded from the audio/video analysis. The split time and break in the timeline visualization denotes switching from working on task 1 to working on task 2. Panel A depicts a turn-taking pattern by the two teams. Panel B depicts a finish and wait pattern.

and the team using the small display interacting with an application. In this instance, the roles switched two more times in total.

A second pattern that occurred was having one team finish early and wait for the other team to catch up. An example is shown in Panel B of Figure 8. After both teams were working for a while, the team using the larger shared display finished their work and waited for the team using the smaller shared display. This becomes apparent when evaluating the fact that the team using the larger shared display was no longer collaborating or interacting with any applications. Once the team using the smaller shared display finished, they started a group conference and the team using the larger shared display started working again. This happened again at the very end of task 2, but with the team using the small shared display finishing first and no longer collaborating or interacting with any applications.

Effects of Display Size

In order to analyze the effects of display size, we separately analyzed the data from the team using the larger shared display and the team using the smaller shared display. The larger shared display was approximately 28 Mpixels and was designed to have more than enough screen real estate to show all applications simultaneously without overlap. The smaller

shared display was approximately 8 Mpixels and was designed to have a limited amount of space; requiring users to organize multiple applications, potentially resulting in overlap.

When analyzing the survey results from the participants who used the larger shared display and comparing them to the participants who used the smaller shared display, some interesting differences emerged. The first interesting pattern is that the participants using the smaller shared display gave higher scores on all answers than the participants using the larger shared display. However, when using a two-tailed t-test to compare the differences, the only answers that were significantly different between the two teams were related to the ease of use and how much they like the *data-duplication* technique, with p-values of 0.0174 and 0.0339 respectively. Average scores and their standard errors for the display size dependent survey results are depicted in Figure 9.

There were also slight discrepancies in how participants ranked the three data sharing techniques based on which size display they had used. Participants using the larger shared display ranked *advanced data-synchronization options* best, followed closely by *data-duplication*. Participants using the smaller shared display ranked *data-duplication* best, followed closely by *advanced data-synchronization options*. These results are shown in Table 2. However, when using a two-tailed t-test to compare the differences, there were no significant differences in ranking based on display size, with all p-values > 0.05.

Data Sharing Technique	Rank (Large Display)	Rank (Small Display)
<i>Data-pushing</i>	2.45 ± 0.13	2.55 ± 0.14
<i>Data-duplication</i>	1.95 ± 0.17	1.64 ± 0.17
<i>Advanced data-synchronization options</i>	1.59 ± 0.18	1.77 ± 0.16

Table 2. Display size dependent rank for the three data sharing techniques, 1 being the best and 3 being the worst.

When analyzing the coded collaboration modes from the audio/video recordings, one interesting difference stands out between the participants who used the larger shared display and the participants who used the smaller shared display. Teams using the larger shared display were in a ‘not collaborating’ state more often and ‘communicating’ or ‘coordinating’ less often. This difference exists with all data sharing techniques. However, when using a two-tailed t-test to compare the difference, the only significant difference is




Figure 9. Display size dependent survey results. Answers were scored on a scale of 1-10, with 1 being worst and 10 being best.




Figure 10. Display size dependent percentage of time spent in each collaboration mode based on data sharing technique.

with the ‘not collaborating’ mode when using the *advanced data-synchronization options*, with a p-value of 0.0397. Display size dependent collaboration mode results are depicted in Figure 10.

CONCLUSIONS AND DISCUSSION

This research has provided an understanding on how collaborators work together in a mixed presence scenario, and how the synchronization of data in shared applications affects the quality of the collaboration. We have conducted an experimental study, which provides empirical evidence that synchronized multi-user applications can improve the quality of mixed presence collaboration through increased parallel task coordination and improved remote collaborator awareness. Our study also elucidated interesting impacts that display size has on mixed presence collaboration.

Conclusions

Results from the coded videos and user surveys showed that actual experiences and user perception support H1 – using continuous synchronization improves mixed presence collaboration. Teams spent more time ‘conferencing’ or ‘coordinating’ and reported higher scores for ease of use, successfulness of collaboration, and how much the participants liked the collaborative features when comparing both continuous synchronization techniques to the previously implemented *data-pushing* system. Of the two continuous synchronization techniques, *advanced data-synchronization options* yielded the best results in terms of improving coordination and decreasing presence disparity.

The results of using *advanced data-synchronization options* also supported H2. User survey results showed that local collaboration was primarily reported as superior to remote collaboration, regardless of which data sharing technique was used. However, when comparing *advanced data-synchronization options* with local collaboration, the differences weren’t always significant. Also, teams were in the same collaboration mode as each other nearly 88% of the time when using *advanced data-synchronization options* (compared to 69% of the time when using *data-pushing*). These results provide evidence that data sharing techniques can be used to mitigate the disparities between remote groups to more closely resemble collocated collaboration.

Results were inconclusive in regards to H3, which predicted that teams using a smaller display would prefer *advanced data-synchronization options*. Survey results contradicted our hypothesis – teams using the larger shared display ranked *advanced data-synchronization options* best, whereas teams using the smaller shared display ranked *data-duplication* the best. However, results of the audio/video analysis support our hypothesis – participants using the smaller shared display engaged in higher modes of collaboration most frequently when using *advanced data-synchronization options*.

Another interesting finding about display disparity was that teams using a larger shared display were ‘not collaborating’ more frequently than the teams using a smaller shared display. While at first glance this may appear to favor using a smaller shared display, this result actually appeared due to some of the common patterns of collaboration that groups used, such as turn taking and finish and wait. The teams using a larger shared display were able to finish their piece of a task faster than the teams using a smaller shared display. Therefore, teams using a larger shared display had to wait longer on average, thus leaving them in a ‘not collaborating’ state more often. This observation falls in line with the prior work cited in the related works section about large displays amplifying sense-making and enhancing collocated collaboration.

Discussion

The *data-duplication* technique allowed users to see interaction icons in addition to shared content, which likely was a major factor for increasing remote collaborator awareness. The *advanced data-synchronization options* technique enabled partial synchronization, which allowed users to work in parallel on separate but related pieces of the overall problem. We therefore suggest showing remote interaction icons and providing user controlled partial synchronization as general features for future groupware systems wishing to support mixed presence collaboration.

This research also demonstrated that large-scale displays provide benefits when performing data analytic tasks in mixed presence collaboration. The use of a larger display allowed groups to complete sub-tasks more efficiently and often led to waiting on the groups using a smaller display. Therefore, large-scale displays that are primarily used for tasks involving collocated groups can now be effectively leveraged for the wide array of tasks that require remote collaboration in addition to collocated collaboration.

While our study focused on a visualization-based data analytics task, it is our belief that the lessons learned could broadly apply to other fields. Synchronized applications on large-scale displays could be used in education to co-teach a class between two campuses, conduct a remote business meeting, or even produce a collaborative work of art. Groups in nearly all fields are pooling resources and forming global collaborations to collect and analyze data. The research presented in this paper provides evidence of how these

groups can effectively collaborate across the distance and accomplish complex tasks.

ACKNOWLEDGMENTS

We would like to thank the participants of our user study. This work was supported in part by the National Science Foundation (NSF), awards ACI-1339772, OCI-0943559 and CNS-0959053.

REFERENCES

1. Christopher Andrews, Alex Endert, and Chris North. 2010. Space to think: large high-resolution displays for sensemaking. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '10), pp. 55-64.
2. Robert Ball and Chris North. 2005. Effects of tiled high-resolution display on basic visualization and navigation tasks. In *CHI '05 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '05), pp. 1196-1199.
3. Anastasia Bezerianos and Gregor McEwan. 2008. Presence disparity in mixed presence collaboration. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems* (CHI EA '08), pp. 3285-3290.
4. Xiaojun Bi and Ravin Balakrishnan. 2009. Comparing usage of a large high-resolution display to single or dual desktop displays for daily work. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '09), pp. 1005-1014.
5. Jacob T. Biehl, Mary Czerwinski, Greg Smith, and George G. Robertson. 2007. FASTDash: a visual dashboard for fostering awareness in software teams. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '07), pp. 1313-1322.
6. Jeremy P. Birnholtz, Tovi Grossman, Clarissa Mak, and Ravin Balakrishnan. 2007. An exploratory study of input configuration and group process in a negotiation task using a large display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '07), pp. 91-100.
7. Clarence A. Ellis, Simon J. Gibbs, and Gail Rein. 1991. Groupware: some issues and experiences. In *Communications of the ACM*, vol. 34, pp. 39-58.
8. Hugo Fuks, Alberto Raposo, Marco Gerosa, and Carlos Lucena. 2005. Applying the 3C Model to Groupware Development. In the *International Journal of Cooperative Information Systems* (IJCIS), vol. 14, pp. 299-328.
9. Google Docs. 2016. Retrieved May 12, 2016 from <https://www.google.com/docs/about/>
10. Carl Gutwin and Saul Greenberg. 2002. A Descriptive Framework of Workspace Awareness for Real-Time Groupware. In *Computer Supported Cooperative Work*, vol. 11, pp. 411-446.
11. Andrew Hayes and Klaus Krippendorff. 2007. Answering the Call for a Standard Reliability Measure for Coding Data. In *Communication Methods and Models*, pp. 77-89.
12. Petra Isenberg, Anthony Tang, and Sheelagh Cpendale. 2008. An exploratory study of visual information analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI '08), pp. 1217-1226.
13. KyungTae Kim, Waqas Javed, Cary Williams, Niklas Elmqvist, and Pourang Irani. 2010. Hugin: a framework for awareness and coordination in mixed-presence collaborative information visualization. In *ACM International Conference on Interactive Tabletops and Surfaces* (ITS '10), pp. 231-240.
14. Lotus Development Corporation. 1995. Groupware – Communication, Collaboration, and Coordination.
15. Thomas Marrinan, Jillian Aurisano, Arthur Nishimoto, Krishna Bharadwaj, Victor Mateevitsi, Luc Renambot, Lance Long, Andrew Johnson, and Jason Leigh. SAGE2: A New Approach for Data Intensive Collaboration Using Scalable Resolution Shared Displays. 2014. In *Proceedings of the IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing* (CollaborateCom '14), pp. 177-186.
16. Gregor McEwan, Markus Rittenbruch, and Tim Mansfield. 2007. Understanding awareness in mixed presence collaboration. In *Proceedings of the 19th Australasian conference on Computer-Human Interaction: Entertaining User Interfaces* (OZCHI '07), pp. 171-174.
17. Rosalie Ocker, Mary Beth Rosson, Dana Kracaw, and S. Roxanne Hiltz. 2009. Training Students to Work Effectively in Partially Distributed Teams. In *ACM Transactions on Computing Education*, vol. 9, no. 1, article 6.
18. Luc Renambot, Thomas Marrinan, Jillian Aurisano, Arthur Nishimoto, Victor Mateevitsi, Krishna Bharadwaj, Lance Long, Andrew Johnson, Maxine Brown, and Jason Leigh. 2016. SAGE2: A Collaboration Portal for Scalable Resolution Displays. In *Future Generation Computer Systems*, vol. 54, pp. 296-305.
19. Anthony Tang, Michael Boyle, and Saul Greenberg. 2004. Display and presence disparity in Mixed Presence Groupware. In *Proceedings of the Fifth Conference on Australasian User Interface* (AUIC '04), vol. 28, pp. 73-82.