Up Close and Personal: Collaborative Work on a High-Resolution Multitouch Wall Display

MIKKEL R. JAKOBSEN and KASPER HORNBÆK, University of Copenhagen

Multitouch wall-sized displays afford new forms of collaboration: They can be used up close by several users simultaneously, offer high resolution, and provide sufficient space for intertwining individual and joint work. The difference to displays without these capabilities is not well understood. To better understand the collaboration of groups around high-resolution multitouch wall displays, we conducted an exploratory study. Pairs collaborated on a problem-solving task using a 2.8m × 1.2m multitouch display with 24.8 megapixels. The study examines how participants collaborate; navigate relative to the display and to each other; and interact with and share the display. Participants physically navigated among different parts of the display, switched fluidly between parallel and joint work, and shared the display evenly. The results contrast earlier research that suggests difficulties in sharing and collaborating around wall displays. The study suggests that multitouch wall displays can support different collaboration styles and fluid transitions in group work.

Categories and Subject Descriptors: H.5.2 [**User Interfaces (D.2.2, H.1.2, I.3.6)**]: Evaluation/Methodology, Graphical User Interfaces (GUI); H.5.3 [**Group and Organization Interfaces**]: Computer-Supported Cooperative Work

General Terms: Design, Human Factors

Additional Key Words and Phrases: Large high-resolution display, wall-display, multi-touch, user study, group work, colocated collaboration, user tracking, proxemics, territoriality

ACM Reference Format:
Mikkel R. Jakobsen and Kasper Hornbæk. 2014. Up close and personal: collaborative work on a high-resolution multitouch wall display. ACM Trans. Comput.-Hum. Interact. 21, 2, Article 11 (February 2014), 34 pages.
DOI: http://dx.doi.org/10.1145/2576099

1. INTRODUCTION

Large interactive displays provide shared spaces around which groups can collaborate. Large displays have been widely used to support groups in face-to-face meetings [Elrod et al. 1992], office work [Streitz et al. 1999], military command [Dudfield et al. 2001], and high schools [Brignull et al. 2004]. Use of large displays has been investigated not only for groups but also for individuals and in public settings such as for public exhibitions [Jacucci et al. 2010]. Extensive research has been conducted on large interactive displays, covering both vertical and horizontal display orientations (for overviews see Robertson et al. [2005], Huang et al. [2006], and Müller-Tomfelde [201 0]). For instance, large displays have been studied with respect to input devices [Birnholtz et al. 2007; Hornecker et al. 2008]; use of gestures, body orientation and position for interacting

---

This work has been supported in part by the Danish Council for Strategic Research under grant 10-092316.
Authors' address: Mikkel R. Jakobsen and Kasper Hornbæk, Department of Computer Science, University of Copenhagen, DK-2300, Copenhagen S, Denmark; email: mikkelrj@di.ku.dk.

DOI: http://dx.doi.org/10.1145/2576099

Fig. 1. Users working on a problem-solving task around a high-resolution multitouch wall display.

with large displays [Vogel and Balakrishnan 2004; Shoemaker et al. 2010]; group dynamics in collaboration [Scott et al. 2004; Hawkey et al. 2005; Inkpen et al. 2005; Tang et al. 2006; Isenberg et al. 2012]; and sensemaking [Andrews et al. 2010].

Compared to the physical size and resolution of most large displays investigated in previous research, emerging technology allows for larger, wall-sized displays with high resolutions [Ni et al. 2006]. Increasing the physical size and resolution of displays has several benefits. For example, more display space allows multiple windows to be displayed simultaneously, which can result in increased performance and user satisfaction [Czerwinski et al. 2003; Bi and Balakrishnan 2009]; it also allows users to navigate more effectively through physical movement [Ball and North 2007]; and it gives room for multiple users to collaborate [Vogt et al. 2011].

We investigate high-resolution wall displays that support interaction up close (see Figure 1) and their benefits for collaborative work of colocated groups. Compared to most of the display systems that have been studied thus far, high-resolution multitouch wall displays (a) support multiple simultaneous touch points to allow multiple users to interact at the same time, (b) have sufficient resolution to allow reading and other work to be carried out up close, and (c) have sufficient space to allow users to mix parallel and joint work at the display.

Each of these characteristics may change interaction and collaboration. For instance, the type and configuration of input devices have been found to affect group work on large displays (e.g., Birnholtz et al. [2007]). Thus, allowing multiple users to interact by touch instead of, say, a single mouse may change interaction and group dynamics. With pixels-per-inch above 60, reading directly from the screen may happen up close [Ashdown et al. 2010]. This, too, may change work dynamics compared to what is suggested by research that has investigated users interacting with the display at a distance (e.g., 1.5m in Birnholtz et al. [2007], 1.5m to 3.0m in Bi and Balakrishnan [2009], 1m to 2m in Isenberg et al. [2009]) or where resolution is too low to comfortably read up close (e.g., 32 pixels per inch in Birnholtz et al. [2007], 2.7 pixels per inch in Haller et al. [2010]). Finally, earlier research suggests that the amount of display space influences collaboration (e.g., Inkpen et al. [2005]). On tabletops, for instance, the relation between private and public parts of the table shape the interaction and performance of a group [Scott et al. 2004]. In addition, it is not clear how the benefits of high-resolution wall displays that have been described for individual users, such as supporting spatial organization of information and physical navigation [Andrews et al. 2010; Ball and North 2007], may appear for group work.

This article presents an exploratory study of how pairs collaborate around a high-resolution multitouch wall display (Figure 1). We use a complex problem-solving task involving a large document collection that provides for varied types of collaboration and that allows us to compare to earlier work [Andrews et al. 2010; Isenberg et al. 2012; Vogt et al. 2011]. The main contribution is to complement research on collaboration around tabletop displays (e.g., Scott et al. [2004], Tang et al. [2006], and Isenberg et al. [2012]) and seated interaction with vertical high-resolution displays (e.g., Birnholtz et al. [2007], Isenberg et al. [2009], and Vogt et al. [2011]) by providing empirical data for wall-sized displays that support up-close interaction. Our study extends earlier findings about collaboration styles and proximity, territoriality, and physical navigation. However, we also note contrasting findings, and we discuss the specific conditions under which different findings may hold. Further, we discuss possible implications of our findings for the design of interactive wall display applications that support colocated work.

2. RELATED WORK

Large high-resolution displays may benefit collaborative work in several ways: They can show more information at a time, allow users to physically navigate the display, and help multiple persons share the display and interact simultaneously. In the following text, we first review research investigating how large displays may change interaction and increase performance on single-user tasks. We then review research that has looked at collaboration around large displays.

2.1. Benefits of Large Displays

Increasing display size has been shown to improve performance and user satisfaction for many tasks, including personal desktop computing [Czerwinski et al. 2003; Bi and Balakrishnan 2009; Ball and North 2005a], sensemaking [Andrews et al. 2010], map navigation [Ball and North 2005b], and 3D navigation [Tan et al. 2006].

A key advantage of increasing display size is that the user can view multiple windows at the same time with reduced navigation. This has been shown in several studies: Czerwinski et al. [2003] found that users performed multiwindow office tasks better and spent less effort managing windows for a 42” wide display compared with a 15” display; Ball and North [2005a] observed and interviewed users of a display consisting of 3×3 LCD panels, reporting that it helped switching between tasks and enhanced the user’s awareness for secondary tasks; and Bi and Balakrishnan [2009] studied users working for 5 days with a 5m wide display that benefited multiwindow and rich-information tasks, and improved awareness of peripheral applications. Increasing display size helps users manage multiple windows by partitioning them in focal and peripheral regions [Grudin 2001; Ball and North 2005a; Bi and Balakrishnan 2009]. Studies suggest that given a large tiled display, users tend to place windows used for their primary task in the center region of the display [Ball and North 2005a; Bi and Balakrishnan 2009]. For sensemaking, large displays can provide a form of external memory that allows users to spatially reference information and to structure information in models that support their analysis task [Andrews et al. 2010].

Many studies including those just described have had users seated in front of the display. Bi and Balakrishnan [2009] allowed users to place their chair in front of a 5m wide display; a distance of 1.5m to 2.5m was the optimal range for most users in that they could clearly perceive the content on the large display and their visual field covered sufficient screen real estate. When users can freely move, large displays may further prompt physical navigation and reduce the need for virtual navigation. For instance, Ball and North [2005b] compared simple navigation tasks using a zoom+pan interface with one, four, and nine tiled monitors, using data that fit the largest display. Participants performed tasks faster and felt less frustration with nine monitors than with one monitor; they also engaged in more physical navigation and less virtual navigation. Later studies have provided further evidence for the importance of physical movement for solving map tasks on a 100 megapixel display [Ball and North 2008].

The benefits of increased display size for managing multiple windows motivated our research. Whereas the work mentioned earlier has investigated how individuals use display space, we investigate how display space is used for organizing windows by multiple users. We extend previous research to understanding how multiple users physically navigate information on a high-resolution wall display.

2.2. Colocated Synchronous Collaboration around Shared Displays

Early research on shared workspaces for colocated collaborative problem solving includes Liveboard [Elrod et al. 1992] in Xerox PARC's Colab and DynaWall in i-Land [Streitz et al. 1999]. This research has given general observations about interaction with large displays embedded in office environments and has described how group activities can be flexibly organized in such environments. Huang et al. [2006] reviewed large-display groupware to identify common factors that influence their adoption and use. The large-display systems they examined were designed for casual ad hoc use and for informal communication and awareness. Some findings from synchronous use were reported. For instance, initial field studies of BlueBoard found mostly one person driving the use of the display, whereby others often stepped back to form an audience [Russell et al. 2002]. However, the groups that were observed using BlueBoard were also found to fluidly take turns. One reason for the turn-taking might be that the display, like in most research around this time, only supported a single touch point at a time.

We study collaboration of colocated groups around high-resolution wall displays that support interaction up close. Probably the first instance of a high-resolution display for up-close interaction is that of Guimbretière et al. [2001], which was evaluated for brainstorming sessions in a pilot study with five groups of designers. Users reacted very positively to the wall-interaction metaphor, but the study did not describe how groups worked together around the display. Descriptions of groups collaborating around large high-resolution vertical displays that support up-close interaction are generally lacking. However, much research has investigated collaboration around tabletop displays and around smaller or lower-resolution vertical displays used at a distance. In the following text, we describe key findings from such research and their relation to the present study.

Coupling Styles. Members of a group may work closely together or in parallel depending on their task. For instance, for mixed-focus tasks (which we later study) group members shift between individual and shared work. Olson and Olson [2000] "use the concept of coupling to refer to the extent and kind of communication required by the work." They described tightly coupled work as nonroutine, ambiguous, typically requiring frequent communication among group members. Loosely coupled work, in contrast, has fewer dependencies or is more routine. Coupling can also be described as the frequency with which members need to interact relative to the amount of work that needs to get done [Salvador et al. 1996]. Mutual awareness between group members, Gutwin and Greenberg [2002] argued, is important for groups to be able to switch between loosely coupled and tightly coupled work. They note that “people keep track of others’ activities when they are working in a loosely coupled manner, for the express purpose of determining appropriate times to initiate closer coupling.”

For interaction around large displays, coupling seems to be related to the physical arrangement of the members in a group. Tang et al. [2006] conducted two observational studies of two-person groups standing around a tabletop display. They identified six styles of collaborative coupling and found these styles to be related to different physical arrangements of the group members around the table. They concluded that “groups use tighter coupling styles when working together closely, preferring common, global views.” Participants stood physically closer when they worked closely together, and they stood further apart when they worked independently. Isenberg et al. [2012] studied 15 two-person groups working on a visual analytics task while seated across from each other around a tabletop display. They identified eight collaboration styles, extending the code set used by Tang et al. [2006]. Eleven of the 15 groups spent more than half of their time (70%, on average) in close collaboration, while the remaining four groups spent 60% of their time working in parallel. Isenberg et al. [2012] highlighted that the face-to-face arrangement was useful for discussions.

Wall displays do not allow for the same physical arrangements as the tabletop displays researched in the earlier discussion. The present work extends this observational research to wall displays, for which studies of collaborative coupling have been lacking.

Location and Territoriality. When people work around a shared tabletop, they sometimes partition the space into territories. This has been described both for work with pen and paper [Tang 1991; Scott et al. 2004] and for interactive tabletop displays [Kruger et al. 2004; Tang et al. 2006]. People have been found to use the area of the table immediately in front of them to form a personal space [Tang 1991; Kruger et al. 2004]. Scott et al. [2004] extended these findings in a study of three small groups working with pen and paper while seated around a tabletop. Results from a spatial analysis showed that people partitioned the space into personal, group, and storage territories. Scott et al. [2004] argued that tabletop territories appear to help groups coordinate their interactions. Given limited display space, group members may not switch effectively between group work and individual work or may require more explicit coordination [Scott et al. 2004]. However, limited space may work well for tightly coupled collaboration. In a study by Tang et al. [2006], users were nonterritorial: Because they were working together all the time, the entire workspace was group territory. In another study, they found that physical positioning appeared to be related to territorial behavior in that participants tended to avoid interacting with areas physically close to their partner. Isenberg et al. [2012] reported that “all pairs negotiated a left-right orientation to the screen, reading documents mostly on one side, storing them on the other.” This was in part attributed to the limited display space and the participants’ fixed positions.

Territoriality in collaboration around shared displays has mainly been found in studies of relatively small tabletops (e.g., 93cm diameter in Scott et al. [2004]). Many questions, therefore, remain about how groups share larger surfaces such as multitouch wall displays. First, size influences territoriality [Scott et al. 2004], but it is not clear how. Second, territories seem less prominent when people move [Tang et al. 2006] instead of remaining seated around a table—territories may be even less significant when people move in front of a wall display. Third, the critical role of orientation in territoriality on tabletops [Kruger et al. 2004] seems less significant for vertical displays, for which groups share the same orientation.

Proximity. Proximity has been found to be important in collaboration around large displays. Tang et al. [2006] found that proximity was related to coupling: Participants were close when using tighter coupling styles. Users may also find it more effective and enjoyable to work closely together [Hawkey et al. 2005]. In a study of 12 pairs collaborating on route-planning tasks around a large display, Hawkey et al. [2005] varied proximity between participants and proximity to the shared display so that participants were both near the display, were both far from the display, or one was near while the other was far from the display. Their results suggest that participants enjoyed interacting directly with the display through touch compared to interacting at a distance using a pressure-sensitive tablet.

The present study extends the observational study of collaboration of Tang et al. [2006] to describe the relation of proximity and coupling for pairs working around wall displays. Although we do not experimentally control for proximity as Hawkey et al. [2005] did, we examine users' proximity to the display as well as to each other. A notable contrast to these studies is the use of mixed-focus tasks that promote parallel as well as joint work.

Input. The type and configuration of input devices has been shown to influence collaboration around shared displays. Birnholtz et al. [2007] compared two input device configurations, a single shared mouse and one mouse per group member, for a mixed-motive negotiation task. Twelve groups of three worked on a newspaper layout task while seated in front of a large display. Multiple mice allowed groups to work in parallel but also lowered the perceived quality of discussion. The single mouse sometimes caused frustration for those not controlling it and allowed one group member to dominate the task. Hornecker et al. [2008] compared multitouch and multiple mice for 13 groups of three on a collaborative design task at a tabletop display. They found higher levels of awareness and a higher incidence of verbal shadowing with touch. Participants were more likely to work in parallel in the touch condition. They also found that touch caused more actions that interfere with each other, but that interactions were more fluid and interferences were resolved more quickly. Furthermore, in a within-subjects comparison of single-mouse, multiple-mice, single-touch, and multitouch, Marshall et al. [2008] measured equity of interaction, verbal equity, and perceived equity for each three-person group. They concluded that "interactive participation is more equal with touch input and multiple entry points than with mice or single input, but verbal participation is not."

These experimental results overall seem to support that multitouch capabilities benefit pairs in the type of mixed-focus task that we study—allowing for parallel work and equal interactive participation.

Display Orientation. Several researchers have argued that horizontal displays are better at supporting collaboration than vertical displays. Rogers and Lindley [2004] compared how eight three-person groups worked together on a tourist itinerary planning task for three conditions: a horizontal table display, a vertical wall display, and a PC monitor. They concluded that the horizontal display condition facilitated the most collaborative and fluid interactions: Participants rarely switched places in front of the vertical display. However, participants were seated around a coffee table in the vertical condition; thus, more effort was required to stand up and move toward the display. Also, because a single-user input device was used in all conditions, interaction was largely based on turn-taking. Inkpen et al. [2005] found that participants made more pointing gestures with a horizontal display than with a vertical display in a collaborative task. Participants moved more and liked the freedom of movement with the vertical display, but several participants were worried about the physical effort of standing for longer activities. In contrast, participants found the horizontal display comfortable to use and a more natural surface to collaborate around.

A notable limitation of earlier research comparing orientation is the modestly sized displays used. The background for the present work is that wall-sized displays are well suited for collaboration; therefore, it is natural to discuss these previous findings. Our observational study does not aim to compare vertical and horizontal displays, however.

Display Size. Increasing display size has several benefits, as described earlier, but few have studied the influence of display size on group work. Ryall et al. [2004] varied both display size and group size in a study of groups assembling a poem from words distributed on a tabletop display. Although participants found a large display better for the task, no performance effects were found for display size. Also, larger groups performed better. The authors noted that the task largely involved searching for words, which can be performed faster in parallel but probably does not gain from more display space.

Compared to the aforementioned study, we use a task that benefits both from increased display space and from multiple persons working in parallel. Although we describe how groups utilize the space of a wall-sized display, the present research is observational and not designed to investigate the relation between display size and performance.

2.3. Summary

In the introduction, we proposed three characteristics of high-resolution multitouch wall displays: (1) simultaneous touch input from multiple users, (2) reading and other work can be carried out up close, and (3) users can mix parallel and joint work at the display. There are several open questions about how people can work together around such displays. First, research has described how collaboration style relates to the physical arrangement around tabletop displays, but wall displays allow for other physical arrangements. What role does proximity and physical arrangement play in collaboration when space allow users to work both physically close and further apart?

Second, descriptions of how groups use territories for coordinating collaboration around tabletop surfaces may not apply as well to wall displays. Some research suggests that use of territories is less pronounced when people move around the shared surface. If people physically navigate a wall display, how will they use territories?

Third, research has shown how people benefit from increased display size and from physically navigating a wall display, but it is not clear how these benefits materialize for colocated groups. How will groups physically navigate when they are collaborating on a task?

3. USER STUDY

We conducted a study to investigate how pairs work together around a high-resolution multitouch wall display. The study can be characterized as descriptive research [Rosenthal and Rosnow 1991]. The decision for doing descriptive (rather than experimental) research was motivated by the many questions that emerged from reviewing the literature. The behavior of groups working around large high-resolution multitouch displays needs to be more carefully described. Mapping out the behavior of groups using this type of display may help identify relationships and generate hypotheses for further research.

3.1. Interface

Participants in this study used a touch-based interface for exploring a document collection. The interface allows users to search a collection of news articles and open individual articles for reading, to view PDFs and images, and to take notes. The searches,

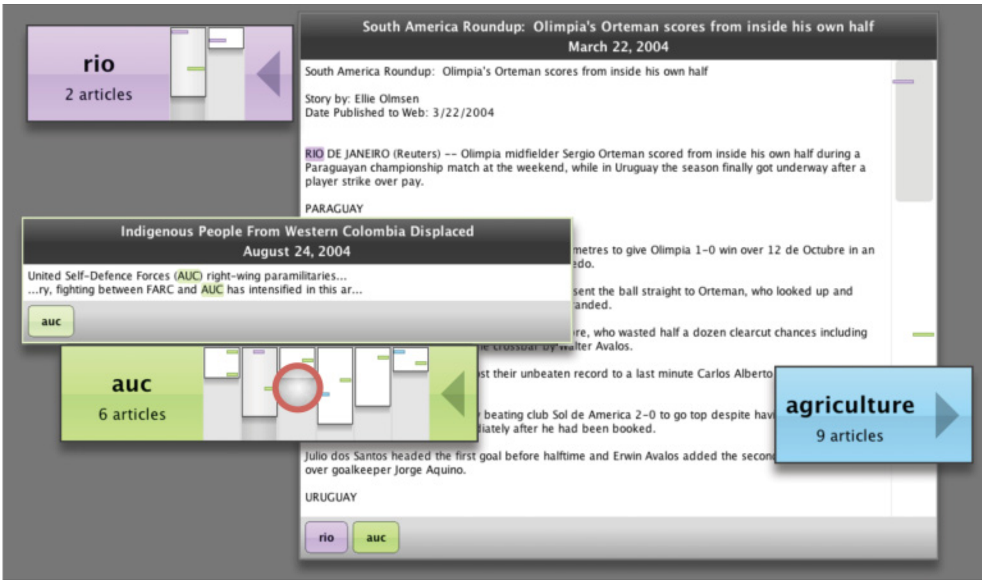

Fig. 2. The interface for searching the news article collection: Two search bars show tilebars for documents containing the terms “auc” and “rio”; one search bar for “agriculture,” which has been minimized; a window showing an article containing the search terms “auc” and “rio”; and a preview of an article containing the term “auc,” which is shown while holding a finger on its tilebar (the red circle indicates the touch point).

articles, PDFs, images, and notes can be freely moved within the interface by touching and dragging with a finger. The interface resembles Cambiera [Isenberg et al. 2012], except for several simplifications and changes made when adapting the interface to a high-resolution wall display. We used Cambiera as an inspiration because it provided a nice baseline for touch-based interaction by multiple users.

Searching. The interface contains four search buttons located near the bottom of the display. Each button has a different hue. When a search is initiated with a button, all interface elements associated with that search use a color variant of that hue. Tapping a search button opens a search bar and an on-screen keyboard. Entering a search term and hitting Enter searches for news articles that contain at least one occurrence of the search term. The search bar expands to show each of the resulting articles as a tilebar, a modified version of TileBars [Hearst 1995]. The height of the tilebar is proportional to the length of the article. Colored rectangles indicate sections of the article that contain a search term (indicating the term’s frequency); if an article is found by more than one search, its tilebar contains a column for each of the matched search terms. Figure 2 shows an example. Search bars can be minimized to hide the tile bars (e.g., in Figure 2, the search for “agriculture”) by tapping the right-side button. Also, search bars can be closed by holding a finger on the bar and tapping the icon that appears after 1 second.

Browsing Search Results. Touching a tilebar in a search bar opens a preview of the article in a window attached to the search bar. The preview includes the article’s title and publication date, a preview of the lines containing the search term, and all the search terms that are found in the document. Figure 2 shows the preview for one of the articles containing the term “auc.” Search results can be browsed quickly by dragging the finger sideways across the tilebars: The preview updates to show the article of the currently touched tilebar.

Article Windows. Articles can be opened for reading by dragging the finger up or down from a tilebar, whereby the preview window detaches from the search bar and then follows the touch; when releasing the preview window, an article window is opened. Articles can be scrolled using a scrollbar at the right edge of the window. Search terms are highlighted with their individual colors. Also, all the search terms contained in the article are listed at the bottom of the window. Article windows can be closed by holding a finger on the window and touching the icon that appears after 1 second. An article can be opened in multiple windows (e.g., by different users).

Highlighting of Opened or Selected Articles. When a user selects a article, by touching its tilebar or article window, all other representations of the article in the interface are highlighted with a colored border. Moreover, the tilebars for articles that have been opened in a window are shown in a darker shade (see Figure 2).

PDF Documents and Images. In addition to the collection of news articles that can be explored through searching, the interface provides views of PDF documents and images. For PDF documents that contain several pages, a thumbnail panel is shown. The users can select a page by touching its thumbnail. In contrast to searches and article windows, the views of PDF documents and images cannot be closed.

Notes. The interface also has a simple feature for making notes. Holding a finger anywhere on the background for 1 second brings up a yellow note and an on-screen keyboard for entering text into the note. A note can be closed by tapping the icon that appears after holding a finger on the note for 1 second.

3.2. Apparatus

Participants used a vertical multitouch display containing 24.8 megapixels (7,680×3,240). The active display area measures 2.8×1.2m, with the bottom edge 89cm above the floor. The display is back-projected and consists of 12 projectors, each with 1,920×1,080 pixels. Projector images are arranged as tiles in a 4×3 grid. The display has a resolution around 68 pixels per inch, which is appropriate for reading [Ashdown et al. 2010]. The display is operated by a single computer equipped with two Radeon 5870 graphics cards.

For detecting touches on the display surface, we used camera-based tracking with diffused surface illumination. Six cameras, capturing 640×480 pixels at 30 frames per second, are connected by Firewire to a computer. The computer runs Community Core Vision to process the input for tracking touch points. The touch points detected within each camera image are multiplexed by a custom program written in Java.

In order to analyze the proximity of participants to the display and to each other, we tracked the position of participants using a camera mounted in the ceiling. The camera was mounted with a wide angle lens in order to cover the room, and captured 640×480 pixels at 15 frames per second. Participants wore colored baseball caps (one red, one yellow) that were detected by a custom program written in Python using the OpenCV library.

3.3. Participants

We recruited 30 participants (8 female) to take part in the study as pairs—in the following, individual pairs are referred to as P1 through P15, and participants in a pair as Red and Yellow. Participants were 18 to 41 years old ($M = 26.3$, $SD = 6.4$). Participants were between 164cm and 194cm tall. Each pair of participants knew each other as students (20), friends (8), coworkers (4), roommates (2), or as a couple (2). Participants knew each other well ($M = 2.3$, $SD = 1.1$, on a scale from 1: "Very familiar" to 7: "Not familiar at all"), and had known each other for 2.8 years, on average ($SD = 3.0$). Participants were given an amount equal to about $40 for joining the study.

3.4. Task

Participants worked on the "Stegosaurus" scenario from the interactive session of the VAST 2006 challenge [Grinstein et al. 2006]. The task scenario involves finding a hidden plot. The dataset consists of a background document explaining the scenario, 230 news articles, three images, one map, one spreadsheet, and three reference documents. The background document gives some starting points for the investigation. Participants must search for and read relevant documents to make hypotheses and gather evidence, filter out irrelevant information, and connect the data. Participants were given a maximum of 1.5 hours to work on the task. We interviewed the participants afterward to learn what they had found out. The dataset was constructed and thus has a known ground truth, which helps determining the progress made by participants. This was for instance useful by Andrews, Endert, and North [2010] for their quantitative comparison of user performance with two displays of different sizes. We were mainly interested in describing *how* groups worked around this type of display, rather than in the outcome of their work, but we do describe how groups progressed on the task.

We chose this task because it provides for varied types of collaboration around the wall display. Group members may find it useful to search or read in parallel, but they also need to share information or work jointly to plan their investigation or to connect evidence. This type of mixed-focus collaboration, in which group members shift frequently between individual and shared activities, requires support for both individual and group needs [Gutwin and Greenberg 2002]. We argue that the wall display provides sufficient space both for working in parallel (e.g., exploring different data in separate parts of the display) and for joint work (i.e., on shared views of the data).

3.5. Procedure

Participants were first given an introduction to the wall display and the interface: We described the functions of the interface and gave participants 10 minutes to try the interface with a sample document collection and to ask questions about the interface. After the introduction, each participant was handed a baseball cap and the head tracking system was calibrated. Participants were then briefed on the task with reference to the background document shown in the center of the display. As part of the briefing, we described the different types of data that were available to them, including the PDF documents and images that were arranged in piles on the display. We gave participants a break after around 45 minutes so that they could sit down if they needed to. After having worked on the task for 1.5 hours, participants were asked to summarize their findings, explaining their hypotheses and how they were backed by the data. Last, participants were administered a questionnaire asking about their age, sex, and height, their relation to their group partner, and questions about their collaboration on the task. Sessions lasted 2:12 hours, on average.

3.6. Data Collection and Analysis

Several types of data were analyzed. First, each session was recorded with video and audio. The recordings comprised 24:27 hours in all, excluding time spent on introductions and filling in the questionnaires. In an initial pass, we determined periods in which participants were working on the task. Excluding time spent on "nonwork"—briefing, breaks, debriefing, and interruptions—19:27 hours (80%) of video were used in the following analysis. Furthermore, some periods could not be coded for visual behavior because a participant was out of sight. On average, pairs spent 79 minutes

Table I. Codes Describing the Level of Verbal Communication between the Participants
Codes are mutually exclusive and exhaustive.

Category	Examples
Silence. None of the participants are talking. Participants may make low inaudible sounds (e.g., sighs, grunts).	Participants are reading the background document.
One is talking. A is talking, while B is silent or is making brief verbal acknowledgements (yeah, okay).	One participant is sharing information from an article he is reading, while the other is reading another document.
Both are talking. A and B are both taking, typically taking turns. Participants are either engaged in "conversational sequences" in which an utterance by A depends on an utterance made by B (and vice versa), or "parallel sequences" in which an utterance of A does not depend on an utterance by B.	One participant asks about a person she has read about in a document, and participants start discussing a possible family relation between persons they each have read about.

(SD = 9.7 minutes) on the task. Because of technical problems, two sessions, P2 and P4, finished before the participants had spent 1.5 hours on the task.

Second, we tracked the position of the participants to quantitatively describe how they moved in front of the display. The tracking data comprise 19:46 hours, excluding briefing, breaks, and debriefing. Because participants left the area of tracking and because of occasional tracking errors, we did not receive data for 2.7 minutes (SD = 5.3), on average, per participant. Our analyses relate the tracking data to coding of the video for verbal and visual behavior. However, because the tracking data cover periods that were not coded for behavior (i.e., participants out of sight), and vice versa (i.e., tracking errors), our analyses use only those periods where we have tracking data together with visual and verbal coding (18:26 hours and 18:48 hours, respectively).

Third, we instrumented the system to automatically collect data describing participants' interaction with the system, including locations of all the windows so that we could determine how the display space was used throughout the task.

Finally, we collected participants' answers to the questionnaires and took note of their comments during debriefing about their use of the display.

3.6.1. Coding of Verbal Communication and Visual Attention. For describing participants' collaboration, we coded the verbal communication and visual attention of participants. Instead of using one code set as Tang et al. [2006] and Isenberg et al. [2012] did, the coding of verbal communication has been separated from the coding of visual attention. This allows us to define more *physically based codes*, which are easier to identify during coding, compared to more *socially based codes*, which rely on abstractions and require inference [Bakeman and Quera 2011]. Also, because the two code sets represent different dimensions, we can analyze contingencies [Bakeman and Quera 2011].

In one pass, sessions were segmented into mutually exclusive and exhaustive states that characterize whether participants are talking, and if so, whether one or both are talking (see Table I). We decided against further categorizing the verbal communication. When both were talking, there was usually some connection between what they were saying, but attempting to infer the degree of coupling in participants' work from how they were talking about the same sources of information, persons, hypotheses, and so on, proved to be unreliable. In another pass, sessions were segmented into mutually exclusive and exhaustive states that characterize participants' visual attention to each other and to the display using the categories in Table II. For describing the visual attention states of participants, we used the five eye gaze patterns shown in Figure 3: One is looking *away* from the display; both are looking at the display, but at *different* areas; both are looking at the *same* area of the display; one is looking at the *other*, who is looking at the display; both are looking at each other, often with *eye* contact. A visual

Table II. Codes Describing the Visual Attention of the Participants
Codes are mutually exclusive and exhaustive.

Category	Examples
Same area. A and B are predominantly looking at the same area of the display. They may switch between areas of the display, but their gaze follow to the same area.	Participants are reading the background document.
Look. A is looking at B, who is looking at the display.	A is looking in B's direction while B is reading aloud and explaining something from a document he is looking at.
Eye contact. A and B are predominantly looking at each other. Participants typically do not sustain eye contact but frequently look also at the display. However, both participant are predominantly looking at the other participant.	Participants look at each other while discussing possible reasons for an incident.
Divided or mutual attention. There is no dominant gaze pattern; the gaze of one or both participants switches frequently, which indicates that their attention is divided. This state is often accompanied by conversation and by pointing at the display.	Participants have explored different news articles and have started discussing connections between three articles and possible connections to other documents; A is making sure B is typing a name correctly, while checking the spelling in an article.
Different areas. A and B are predominantly looking at different areas of the display. Sometimes, participants are not looking at any particular area but switch between different areas of the display with no clear focus.	A is reading a document on dangerous chemicals while B is searching for articles about a person; participants have finished discussing a lead and are scanning the display to pick up new information to work with.
Disengaged. A is predominantly looking away, engaged neither neither with the display nor with B. B is looking at the display.	One participant is taking a glass of water; one participant seems unmotivated and is not involved in the task.

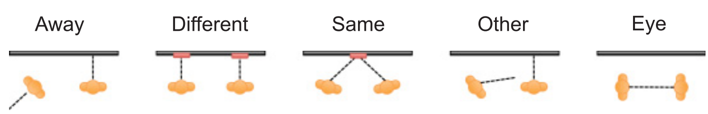

Fig. 3. Patterns of the participants' gaze at an instant in time.

attention state is determined as a period of at least five seconds that is predominantly characterized by one of these gaze patterns. We added an additional category, *divided or mutual attention*, which characterizes periods where none of these gaze patterns are predominant.

The verbal communication and visual attention codes together with the proximity data indicate whether group members are working tightly together: Participants standing close and talking together while looking at the same view are tightly coupled, while participants standing far from each other and not talking while looking at different areas of the display are loosely coupled. These codes in combination can be used to describe the coupling styles coded in earlier research [Tang et al. 2006; Isenberg et al. 2012]. For instance, in *active discussion* [Isenberg et al. 2012], participants are typically looking at each other while conversing (*eye contact* or *mutual attention* together with *both are talking* in our schema); *sharing of the same view* [Isenberg et al. 2012] and *same problem same area* [Tang et al. 2006] is coded as visual attention on the *same view* in our schema, indiscriminate of whether participants are silent or talking.

Verbal communication and visual attention were coded by two observers. Two samples of the video (40 minutes and 80 minutes) were coded by both observers as well as the first author. All three observers discussed and refined the coding scheme until a reasonable interobserver agreement was reached: $\kappa = 0.85$, both for verbal communication and for visual attention, which indicates over 90% observer accuracy [Bakeman

Table III. Participants' Answers to Questions About How They Worked Together During the Study

Question	M	SD
How closely did you work together with your partner during the study?		
We worked (1) together / (7) independently all the time	2.8	0.8
How closely did you monitor the work of your partner during the study?		
I (1) was fully aware / (7) had no idea what my partner was doing	3.4	1.5
Did you divide up tasks between you? "Yes": 13 "No": 2		
Did you and your partner work effectively as a team to solve the scenario?		
The group work was (1) very effective / (7) very ineffective	2.4	0.9
How much did you contribute to solve the tasks?		
I contributed (1) the most / (7) the least to solve the scenario	3.7	0.7
How satisfied are you with your work in the team to solve the scenario?		
(1) Very satisfied / (7) Very dissatisfied	2.5	1.2

and Quera 2011]. We calculated time-unit Kappa with tolerance [Bakeman and Quera 2011], using a 2-second tolerance as recommended by Mudford et al. [2009].

3.6.2. Participant Tracking. Based on head tracking data, we analyzed participants' location in front of the display. The precision of the tracking data was measured for eight points on the floor. For each point, we measured the distance to the four corners of the room and to the center point of the display, and we compared the physically measured location with the location determined from the web camera image. On average, the camera-determined distance deviated from the physically measured distance by 5cm ($SD = 2.8$cm). To measure how much participants moved based on the tracking data, we used a modified version of the Douglas-Peucker algorithm [Douglas and Peucker 1973] to compensate for jitter from the tracking system. Also, to determine tracking precision, we walked a path connecting eight points five times. We measured the physical distance of the path, which was compared to the distance determined from the tracking data. We found an average error of $M = 5.3$cm.

3.6.3. Activity and Window Logging. In all, 8,219 interaction events (e.g., starting a search or dragging a window, but excluding keyboard taps) were logged. For each interaction event, we identified the participant touching the display so that we could determine which areas of the display that each participant interacted with. We automatically identified the participant touching the display for events where only one participant was within 76cm of the display location (69% of the events). For the remaining events, an observer was given the time, location, and interaction type of each event to help identify the participant from the video recordings using a commercial analysis tool. The observer also coded a sample of 10% of the events where only one participant was within 76cm of the display location so as to validate the automatic identification. Interobserver agreement was excellent (Cohen's $\kappa = 0.93$). Another observer coded a sample of 10% of the events for each pair, taken at random times in the video. The interobserver agreement was excellent (Cohen's $\kappa = 0.85$).

3.6.4. Questionnaire. In addition to questions about participants' age, sex, and how well they knew each other, the questionnaire contained questions adapted from the study of Isenberg et al. [2012]. The answers to those questions (summarized in Table III) describe participants' perception of how they collaborated during the task.

Participants commented in free text on questions about how they worked on the task. Nine pairs commented on how closely they worked together (see first question

Table IV. Verbal Communication, Summary Statistics for Codes

	Frequency	Relative duration	Mean duration	Median duration
Both are talking	1,139	28.9%	17.8s	12.7s
One is talking	1,505	25.9%	12.1s	9.0s
Silence	1,619	45.2%	19.6s	10.6s

Table V. Visual Attention, Summary Statistics for Codes

	Frequency	Relative duration	Mean duration	Median duration
Same area	633	26.7%	28.9s	19.2s
Look	125	2.4%	13.2s	11.7s
Eye contact	67	1.1%	11.3s	9.3s
Divided or mutual attention	424	14.4%	23.3s	18.7s
Different areas	836	54.8%	45.0s	26.6s
Disengaged	20	0.6%	19.7s	14.3s

in Table III). Most comments said something similar to "reading was individual, but we communicated often," a few to the effect that they started out individually and later put their clues together. Seven pairs provided comments to the question about how closely they monitored each other: Many were positive on how they communicated their findings along the way, but a few comments indicated problems—one suggested making regular "time outs" to give each other updates. All but two answered that they divided tasks between them. Of those, 11 pairs wrote that they divided articles to read; 11 divided searching; a few specified how they divided topics or areas of their investigation.

4. RESULTS

4.1. Collaboration

Participants switched fluidly between working together and in parallel around the high-resolution multitouch wall display. Proximity was found to play a clear role in how tightly coupled participants worked. Also, the display afforded flexibility in collaboration styles, and great variation was found in how the pairs collaborated. This is described in the following text in terms of the level of verbal communication, visual attention, and proximity between participants.

4.1.1. Verbal Communication and Visual Attention. The codes describing the verbal communication of participants are summarized in Table IV. Overall, participants talked 54.8% of the time and were silent 45.2% of the time. Periods of silence were often interrupted by talking, as indicated by the mean duration of silences ($M = 20$s). There was an overweight of shorter periods as indicated by the median of 11s, and a tail of silent periods that are noticeably longer (up to 6:24 minutes). The codes for verbal communication thus show that participants frequently talked: We coded *both are talking* 29% of the time, *one is talking* 26% of the time. By itself, *silence* might indicate loosely coupled work, whereas *both are talking* would typically indicate tightly coupled work. However, to better understand coupling in participants' work, we related verbal communication to visual attention and the proxemics data.

The codes describing the visual attention of participants are summarized in Table V. Overall, participants spent most of the time looking predominantly at *different areas* (54.8%), which might indicate that they worked in parallel. Participants kept visually attending to different areas for relatively long periods at a time ($M = 45$s). The other visual attention states (except *disengaged*, 0.6%) indicate tightly coupled work by the way that participants are visually paying attention to the same area or to each other.

Table VI. Verbal Communication and Visual Attention, Adjusted Residuals (z_{rc})

	Both are talking	One is talking	Silence
Same area	13.81	19.53	−29.93
Look	13.00	11.98	−22.51
Eye Contact	18.22	6.86	−22.76
Divided or mutual attention	83.67	−4.39	−72.70
Different areas	−77.80	−18.13	87.27
Disengaged	−6.93	−8.24	13.64

Participants were predominantly looking at the *same area* 26.7% of the time, and 14.4% of the time their visual attention was *divided*. Participants spent much less time looking at each other (*look*: 2.4%, *eye contact*: 1.1%) and for much shorter periods at a time ($M = 13.2$s and $M = 11.3$s, respectively) compared to when they looked at the *same area* ($M = 28.9$s) or *different areas* ($M = 45$s). Participants were also observed to make quick glances at each other. This behavior is not described earlier because only periods with a predominant gaze pattern lasting more than 5s were coded. These short occasional glances we think are important for participants in maintaining awareness of what each other are doing and for coordinating their work (e.g., identifying opportunities for switching to joint work).

The codes for the visual attention and verbal communication of participants can be combined to better describe coupling in their work. Across sessions, we find an association between the two sets of codes, Pearson's $\chi^2(10, N = 67{,}965) = 12{,}620.72$, $p < 0.001$. Tests of association were performed on a tally of 1-second units. As a measure of strength of the association, we calculated Cramer's $V = 0.30$, which indicates a medium effect size according to Cohen [1988, pp. 222–224]. Table VI shows the association between the individual codes in a contingency table. Each cell in the table shows the adjusted residual, a standardized measure of how the observed frequency of the pair of codes differ from the expected frequency (i.e., assuming the codes are not related). A positive residual indicates that the observed joint frequency of the pair of codes is greater than expected; a negative residual indicates that their frequency is less than expected. Also, the adjusted residuals can be compared to a 1.96 criterion for indicating significance at the 0.05 level [Bakeman and Quera 2011], but multiple tests increases the risk of type I errors. We focus only on associations with large residuals. Appendix A explains how the adjusted residuals are calculated from the observed and expected frequencies. Among the joint frequencies of the codes (i.e., the cells in the table), the most clear association is found for *different areas*, for which participants were observed more often to work in *silence* than would be expected ($z_{rc} = 87.28$) and less often talking with one another ($z_{rc} = -77.8$). This means that participants working in parallel at different areas of the display often remained silent. For the other codes—those we take as indication of tight coupling—participants were observed more often talking and less often silent than would be expected (positive residuals for *both are talking* and *one is talking*, negative residuals for *silence*; see Table VI). This is the most clear for *divided or mutual attention*, which is observed more often together with *both are talking* ($z_{rc} = 83.67$) and less often with *silence* ($z_{rc} = -72.70$) than would be expected. Here, participants often were talking while referencing different areas of the display in order to build a shared understanding of data, to discuss hypotheses, and so on.

The summary of the codings might suggest that participants spent more time working in parallel (*different areas*: 54.8% and *silence*: 45.2%) than working together. However, from the questionnaire answers, most of the participants felt they worked more

Table VII. Proximity, Summary Statistics for Codes Derived from the Tracking Data

		Frequency	Relative duration	Mean duration	Median duration
Intimate	0–46 cm	434	18.4%	29.2s	18.0s
Personal, close	46–76 cm	905	35.0%	26.7s	16.0s
Personal, far	76–120 cm	711	26.1%	25.3s	15.0s
Social	120–370 cm	333	20.5%	42.5s	21.0s

Table VIII. Verbal Communication and Proximity, Adjusted Residuals (z_{rc})

	Intimate	Personal close	Personal far	Social
Both are talking	7.03	17.56	−5.21	−21.84
One is talking	4.15	9.23	−7.10	−7.17
Silence	−10.08	−24.18	11.01	26.28

together than independently ($M = 2.8$ on a scale from 1: together to 7: independently). Only one pair answered that they worked individually for the most part. We note that participants switched frequently between coupling styles, as indicated by the mean durations of codes that range between 11s and 45s. As for working in parallel, all but two pairs answered that they divided up tasks between them. Participants felt they worked effectively as a group ($M = 2.4$, $SD = 0.9$), and they were all satisfied with their work ($M = 2.5$, $SD = 1.2$).

4.1.2. Proximity and Coupling in Group Work. Table VII summarizes the time participants spent in the different proxemic distance zones defined by Hall [1963].[1] Proxemics was introduced by Hall for studying how people interpret and use space in relations with other people. We use the physical distances originally given by Hall, but it should be noted that they depend on culture, physiology, and other factors. It should also be noted that Hall distinguished between close and far intimate zones and close and far social zones. However, because participants spent less than 1% of the time at intimate close distance and at social far distance, we summarized times for the intimate zone and the social zone. On average, participants spent about half of the time at less than one arm's length from each other: 35% of the time at close personal distance (defined by [Hall 1966] as 46–76cm), and 18.4% of the time at intimate distance (0–46cm). At these distances, referred to as *physically close* in the following, both participants can typically reach or point to a particular area of the display, which may be important for tightly-coupled work. The other half of the time participants were not physically close: 26.1% of the time at personal far distance (76–120cm) and 20.5% of the time at a social distance (120–370cm). Participants typically remained in a proxemic zone less than 30s at a time, except at social distance where they remained longer at a time ($M = 42.5$s).

Proximity between participants played a role in how tightly coupled they worked. This can be described by relating proximity to the codings of verbal communication and visual attention. First, verbal communication differed by proximity, Pearson's $\chi^2(6, N = 67{,}690) = 1{,}169.72$, $p < 0.001$. However, Cramer's $V = 0.09$ indicates a small effect size [Cohen 1988]. Table VIII shows the association between the categories as adjusted residuals. There seems to be a systematic association between verbal communication and proximity: Participants were observed more frequently talking (positive z_{rc}) and less frequently silent (negative z_{rc}) when participants were physically close

[1]An hysteresis tolerance (10% of the region width) was used in determining transitions between zones so as to avoid frequent switching caused by small movements near thresholds between regions.

Table IX. Visual Attention and Proximity, Adjusted Residuals (z_{rc})

	Intimate	Personal close	Personal far	Social
Same area	114.14	16.22	−60.67	−63.39
Look	−11.12	−11.19	6.61	16.81
Eye Contact	−0.80	0.41	1.81	−1.69
Divided or mutual attention	−6.78	30.13	−2.65	−26.30
Different areas	−92.95	−31.79	52.18	70.64
Disengaged	−1.81	−3.53	8.79	−3.64

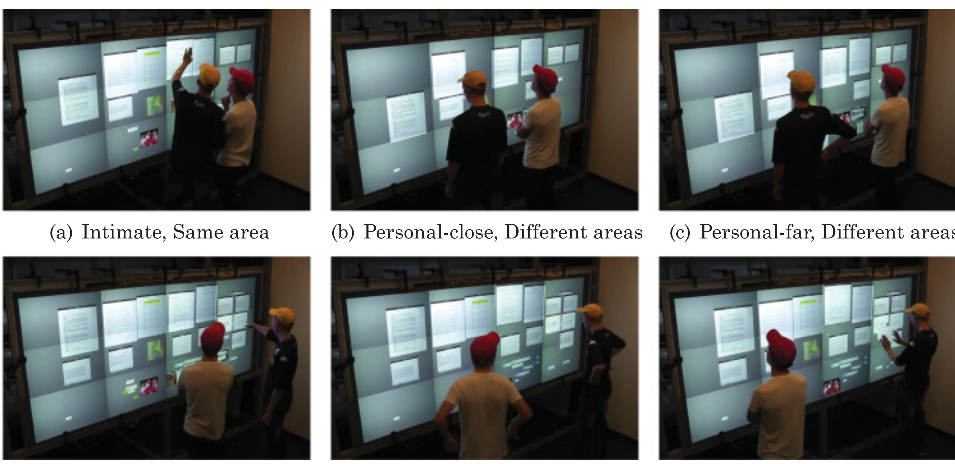

(a) Intimate, Same area (b) Personal-close, Different areas (c) Personal-far, Different areas

(d) Personal-far, Same areas (e) Social, Look (f) Social, Different areas

Fig. 4. Video snapshots of P3 for varying proximity and visual attention between participants.

than would be expected. Conversely, *silence* was more frequently observed, and talk less frequently, when they were not physically close than would be expected. Although there is a significant association—participants were physically close more often while talking together—the association is not strong. It is clear that participants also did benefit from talking while being apart. For instance, during one instance of P6 working physically close together (11:00–12:00, both are talking), while Red started searching, Yellow moved to another part of the display so as to provide information for Red's search from a document located there.

We also found that visual attention differed by proximity, Pearson's $\chi^2(15, N = 66{,}377) = 19{,}918.36$, $p < 0.001$. Cramer's $V = 0.32$ indicates a large effect size [Cohen 1988]. Table IX shows the association between the categories. Figure 4, which shows representative video snapshots for varying proximity and visual attention between participants, may help understand the associations. The most obvious associations are seen for *same area* and *different areas*. Participants were more frequently observed looking at the same area when physically close (positive z_{rc}), less frequently when not physically close (negative z_{rc}); they were more frequently observed looking at different areas when not physically close, less frequently when physically close. This is not surprising because it becomes more difficult for participants to see the same area when the distance increases (compare Figures 4(a) and 4(e)). This also explains why the *look* state is more frequently observed when participants are far than close. A participant might often look when the other calls for their attention (as was the case in Figure 4(e)). This is a possible reason why look states are relatively short lived. Particularly *same area* is frequently observed at intimate distance ($z_{rc} = 114.14$): Looking closely at the

(a) (b) (c)

Fig. 5. Video snapshots of P6 showing transition to tight coupling. (a) Participants work in separate areas, (b) then Yellow turns his head and leans toward Red when he calls for his attention, and (c) Yellow moves closer to read together with Red.

same area is a good reason for standing this close, which might otherwise be socially inappropriate and uncomfortable for pairs that are not intimately related [Hall 1966]. To compensate for such intimate closeness, participants may limit their movements and remain facing toward the display. Last, we note that *divided or mutual attention* was more frequently observed together with conversation ($z_{rc} = 83.67$) and at close personal distance ($z_{rc} = 30.13$). This is different than when participants look at the same area, which was observed more frequently together with one talking ($z_{rc} = 19.53$) than with conversation ($z_{rc} = 13.81$), and more frequently with intimate ($z_{rc} = 114.14$) than with personal close ($z_{rc} = 16.22$) distance.

Several points can be made from the associations just described between the verbal communication, visual attention, and proximity of participants. First, the associations indicate that participants are more often tightly coupled when physically close, more often loosely coupled (i.e., they are more often silent and looking at different areas) when not physically close. Second, periods of loosely coupled work are typically longer. Third, the display seemed flexible in supporting different working arrangements. For instance, participants were often tightly coupled at intimate and close personal distances, although they could of course be working in parallel while physically close. Figure 4(b) shows one such instance for P3 where participants are reading different documents that happen to be placed closely together on the display. Also, participants could quickly shift from loosely to tightly coupled work through minimal physical movement. For instance, a participant could simply turn his head or move closer to the other if called for to share information on the display. Figure 5 shows such an instance, where Red is calling attention to something in a document he is reading; Yellow turns his head and, because they are at social distance, moves closer to share the same view.

4.1.3. Variation in Collaboration Style and Proximity. There was great variation across pairs in collaboration among participants. This can be seen in Figure 6, which summarizes the time that participants in each pair spent in different modes of verbal communication and visual attention, and at different distances from each other. First, the level of verbal communication differed across pairs: time spent in *silence* varied between 17% (P14) and 75% (P12), in conversation between 13% (P12) and 49% (P7). Second, the time spent in different modes of visual attention varied: Most importantly, time spent looking at the *different areas* varied between 32% (P5) and 72% (P7). Third, the time pairs spent at different distances varied: P1 spent only 3% of their time at *intimate* distance, whereas P5 spent 55% of their time at this distance.

Together, the data indicate that pairs used diverse collaboration styles. One striking example is P7: This pair spent the most time looking at different areas (72%), which by itself would indicate loosely coupled work. However, they also spent the most time in conversation (49%), which would indicate tightly coupled work. The coding data showed that they were talking together while looking at different areas for more than 20% of the time (compare to $M = 6\%$ across pairs). They also were among the pairs

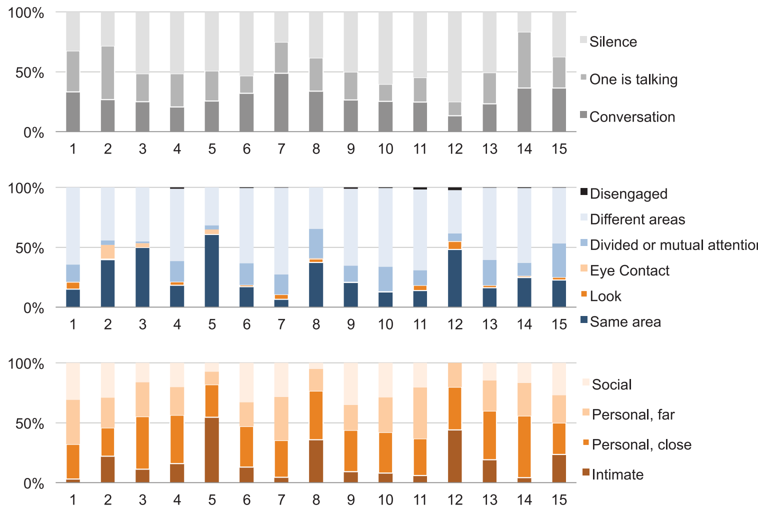

Fig. 6. The relative amount of time participants in each pair (1–15) spent: (top) in different modes of verbal communication; (middle) in different modes of visual attention; (bottom) at different distances from each other.

that spent the least amount of time (35%) being physically close to each other. Another pair, P12, looked at the same area almost half of the time but were silent 75% of the time. Based on the difference in collaboration styles, we emphasize that care must be taken in interpreting the codings for a single dimension. The combination of codings on multiple dimensions of collaboration provide a richness in description of the group work that is missed in the analysis of the individual codings.

4.2. Display Sharing

We analyzed participants' movement relative to the display to understand how they shared the display. A key result was that most participants shared the display quite evenly. We also found great variation in how participants used the display space.

First, we quantified how pairs shared the display over time by calculating an index of display partitioning (IDP_{time}). The index is calculated by partitioning the area in front of the display into N regions, then determining how much longer each region is occupied by one participant relative to the other:

$$IDP_{time} = \frac{\sum_{r=1}^{N} |t_{Rr} - t_{Yr}|}{\sum_{r=1}^{N} t_{Rr} + t_{Yr}},$$

where t_{pr} is the time spent by participant $p \in \{R, Y\}$ (wearing the Red or the Yellow hat) in region r. We calculated IDP_{time} for $N = 10$, that is, 28cm wide regions, which corresponds to the width of a document on the display. Two participants occupying separate regions for the entire task would give $IDP_{time} = 1$. The median $IDP_{time} = 0.35$ suggests that most pairs shared the display quite evenly. The number of regions influences IDP_{time}: For instance, using larger regions results in a lower median $IDP_{time} = 0.24$, $N = 4$. However, this is highly dynamic and variations among pairs depend on how they moved relative to the display throughout the session.

Whereas most pairs shared the display quite evenly in terms of the time each participant spent in different regions, pairs might have partitioned the display in terms of how often they interacted with different regions. To investigate this, we also calculated IDP_{int} (similar to IDP_{time}) determining how many of participants' touch

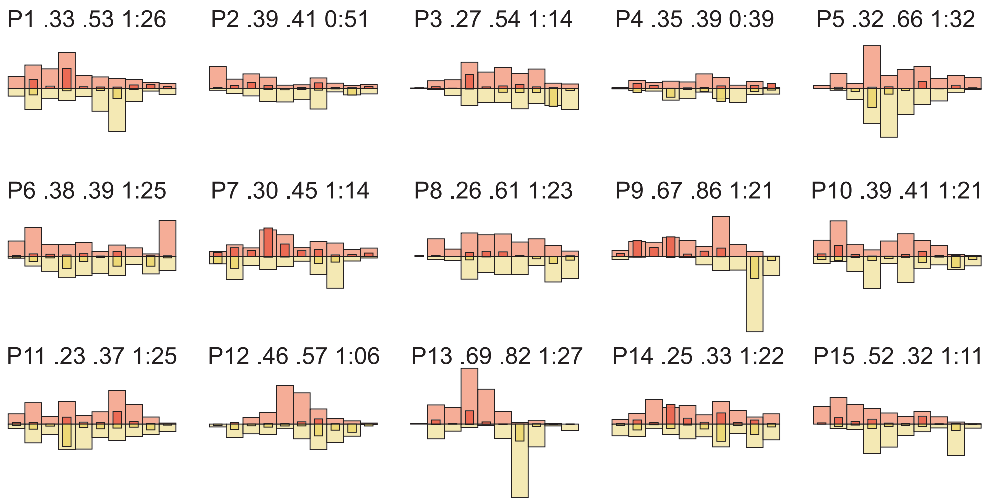

Fig. 7. The time participants spent in front of the different regions of the display (large bars), and their interactions with the different regions of the display (interior bars). Each of the ten bars shows the sum of time spent (and periods of interaction, for interior bars) by participant in the region, divided among Red (above the x-axis) and Yellow (below the x-axis). Above each chart is the pair number (e.g., P1), the index of partitioning IDP_{time} (e.g., 0.33) and IDP_{int} (e.g., 0.53), and the total time each pair was tracked (e.g., 1 hour and 26 minutes). Note that P2 and P4 were interrupted due to technical problems.

interactions with each display region were made by one participant relative to the other. This index of partitioning has a somewhat higher median $IDP_{int} = 0.45$ compared to $IDP_{time} = 0.35$. Because IDP_{int} is based only on periods where participants interact with the display, it is more sensitive to partitioning in activities involving interaction. One such activity for which participants might partition the display more clearly is searching. The interface provided several entry points for starting a search. However, once a participant started searching in a particular place, they might do many searches in that particular place. Reading, in contrast, does not require much interaction (except for scrolling or dragging a window to a better reading location). Because participants often read the same documents to share information, this impacts the index of partitioning.

Despite the higher IDP_{int}, both measures of display partitioning are low, which indicates that relatively few pairs partitioned the display clearly between them. However, we found great variation in how pairs shared the display and some clearly did partition the display. To describe differences across the pairs, we visualized the time participants occupied—and the number of interactions with—each region of the display. This is shown for each pair in Figure 7, which also lists the two indexes of partitioning, IDP_{time} and IDP_{int}. First, we see that IDP_{time} varies greatly across pairs. The lowest $IDP_{time} = 0.23$ is found for P11, in which the participants occupied different regions for a comparable amount of time. In contrast, participants in P13, which had the highest $IDP_{time} = 0.69$, mostly spent time in front of their respective parts of the display. Similar variation is seen for participants' interaction with the display: IDP_{int} ranges from 0.32 for P15, who interacted the most equally with different regions, to 0.86 for P9, who seemed to partition the display quite clearly between them. Yellow of P9 spent more than half of the time in only a single region in front of the display—the second region from the right. The visualizations in Figure 7 show how some participants were located mostly in front of particular regions of the display or interacted the most with the display in particular regions. For instance, in P13, Red spent the most time and interacted the most with the display in region 4, while Yellow spent the most time and interacted the most in region 7. The other pair showing the most clear partitioning of the display is P9 ($IDP_{time} = 0.67$ and $IDP_{int} = 0.86$), in which Red occupied and interacted almost exclusively with the entire left part of the display, while Yellow remained in a small area concentrated around region 9 in the right side of the display. Variations in how pairs partitioned the display can be seen. For instance, in some pairs both members worked in separate, spatially concentrated regions (e.g., P1, P10, P13); in other pairs, one member working in a concentrated region while the other used more regions (e.g., pairs P9, P12, P15). P3 spent close to an equal amount of time in different regions, but their amount of interaction varied across regions.

Another way of looking at how participants partitioned the display is whether participants stayed on one side of their partner. Half of the pairs spent 69% or more time with their partner on one side. Some spent almost an equal amount of time on the left and right sides of their partner, whereas, for instance, in P13, Red spent 83% of the time to the left of Yellow.

There seems to be no straightforward relation between the way pairs worked in different collaboration styles and how they shared the display (compare Figures 6 and 7). For instance, P11 had the lowest IDP_{time}, but spent much of the time working on different areas. P7 similarly had a low IDP_{time} and spent much time working on different areas. Pairs that worked much of the time in parallel might be expected to partition the display more clearly, but this was not the case for these pairs. Group members may have switched places and worked for an equal amount of time in different regions, but at different times. In contrast, P5 had higher-than-average IDP_{int} but worked most of the time looking at the same area and being physically close.

4.3. Distance and Physical Navigation Relative to the Display

We analyzed participants' movement relative to the display to understand how large multitouch displays support up-close interaction and to describe the role of physical movement for navigation.

4.3.1. Proximity to the Display. Most of the time, participants were close to the display: 60% of the time ($M = 47.5$ minutes, $SD = 13.0$ minutes) was spent within a comfortable distance for interacting with the display (<46cm, "upper arm or elbow distance," [Hall 1963]); 31% ($M = 24.2$ minutes, $SD = 15.9$ minutes) within arms reach (<76cm). Most activity including reading documents is often best done at a close distance. Around 6% of the time, participants were at a distance where they could not reach the display ($M = 4.5$ minutes, $SD = 2.9$ minutes). On average, participants stepped back from the display (and stayed there for at least 5 seconds) 43 times ($SD = 20.8$). They stepped out of touching distance for several reasons: to read more comfortably, to gain an overview of the display, or to move around their partner to access a different part of the display. Participants' proximity to the display varied across participants. One participant, who spent 10% of the time away from the display, said that he liked walking around while thinking; others worked almost exclusively up close throughout the task.

4.3.2. Movement. Participants moved 329m ($SD = 82$m), on average. Moreover, participants moved past their partner, for instance to access another part of the display, 53 times ($SD = 33.7$), on average. This indicates that although participants worked more in some areas of the display or more to one side of their partner, they often found it useful to move in front of the display.

4.4. Use of the Display

When participants were asked to summarize their findings, an average of 48% of the display was covered by windows. However, the percentage of display space used

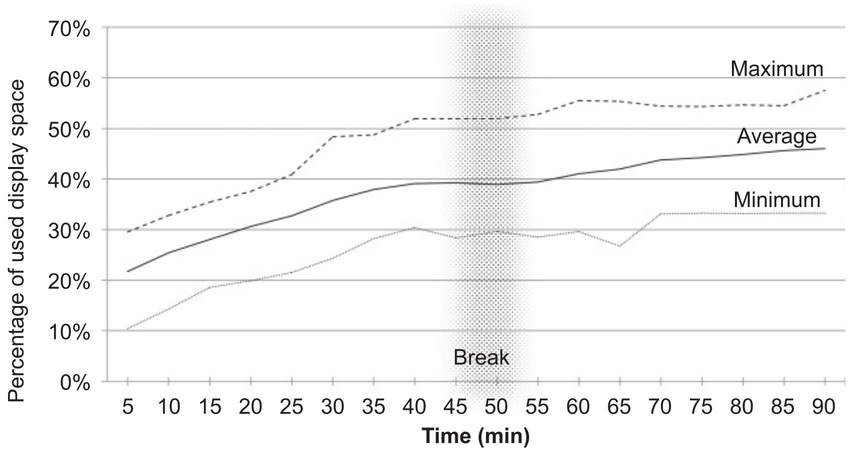

Fig. 8. Percentage of the display space that is covered by windows at different times since participants started working on the task: shown as the average for all pairs, and for the pairs that had the minimum and maximum percentage of usage. Most pairs had a break in the task around 45 to 55 minutes into the session, which explains why the average percentage of used display space does not increase around that time.

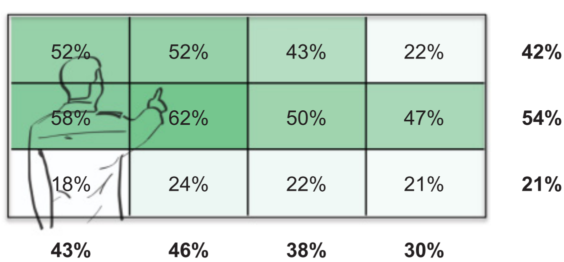

Fig. 9. Percentage of the display space that is covered by windows at the time when participants were asked to summarize their findings, summarized for the 12 areas covered by each projector.

varied between 33% and 58% across pairs ($SD = 7$). Figure 8 shows that the amount of space that was used increased over time, which suggests that by the end of the task, participants were still putting more of the available display space in use for the task. A general interpretation is that participants kept searching and producing article windows throughout the task, arranging them on the display and making use the available space. However, participants left a relatively large amount of the space unused. For comparison, Hutchings et al. [2004] studied 39 users' computer activity for 3 weeks and found that less than one fifth of the display was empty for 89.9% of the time for single-monitor users and more than 71% of the time for dual-monitor users.

One reason why participants used a relatively low percentage of the available display space compared to what other research shows is that they used some areas of the display more than others. To describe this, we divided the display into 12 areas, three rows and four columns, and determined the amount of display space used in each area. This is shown in Figure 9. The display area of the center row was covered the most by windows (54%, on average, at the end of the task). The area of the bottom row was covered the least (21%). The center row of the display provided the most appropriate height for reading, given the height of the participants ($M = 180$cm). The top row was also accessible to most participants; however, the bottom row required most participants to bend down to be able to read. Thus, participants often moved article windows to the

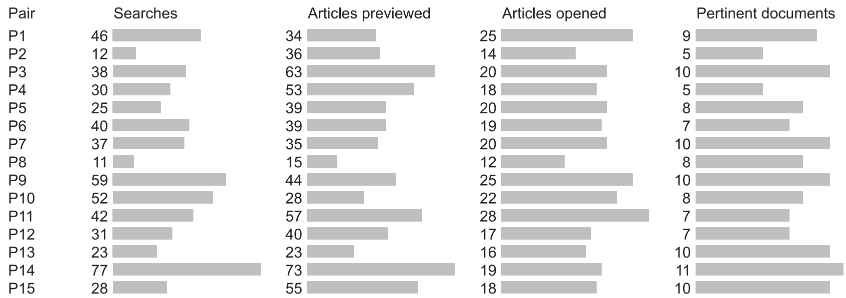

Fig. 10. The number of searches performed, articles previewed and opened, and pertinent documents shown in the display at the end of the session.

center of the display for reading. The lower part of the display was mainly used for storing information.

4.5. Task Progress

We determined how much progress each pair had made on the task by analyzing their interaction with the system. Figure 10 summarizes pairs' searches for relevant information in the article collection. The last column shows the number of documents they had opened that are pertinent to the scenario; nine news articles and two PDF documents contained pertinent information (we reviewed the display to determine if pairs had specific pages with pertinent information in the two PDF documents shown on the display). The number of pertinent documents is an important task performance metric according to Kang et al. [2009].

On average, pairs made 36.7 searches ($SD = 17.4$), and previewed 42.3 news articles ($SD = 15.5$) by browsing through the search results. Pairs opened 19.5 news articles on average ($SD = 4.2$). On average, pairs found 8.3 ($SD = 1.9$) of the 11 documents containing information pertinent to the hidden plot. The search efforts of the pairs varied much: P14 did the most searches (77), previewed the most articles, and uncovered all the pertinent documents. In comparison, P13 did only a one-third as many searches as P14 did but missed one relevant document. P8 searched only 11 times, previewed 12 articles, and found 8 documents pertinent to the plot.

Only one pair, P14, had all information necessary to uncover the plot shown on the display by the end of the session. Other pairs missed between one and four pieces of the plot, except the two groups (P2 and P4) that were interrupted before the time limit because of technical problems. There was no relation between the number of pertinent documents found and how satisfied participants were with the work on solving the scenario.

We considered whether the number of pertinent documents found by each pair would be related to how they collaborated (verbal communication, visual attention, proximity, *IDP*, and use of display space). No patterns emerged except that the number of pertinent documents seems related to how much participants talked. The more time participants spent talking, the more pertinent documents they found, Pearson's $r = 0.67$, excluding the two pairs (P2 and P4) that were interrupted. One interpretation is that talking helped pairs better coordinate their search for relevant information.

5. DISCUSSION

The aim of this work has been to describe how people work on a shared task given a high-resolution multitouch wall display that allows simultaneous interaction by multiple users. We have sought to better understand the benefits of such displays for group

Table X. Main Observations, Their Relation to Earlier Research on Large Displays, and Open Questions

Observations from this study	Previous research	Open questions
Participants interacted simultaneously with touch display, both in joint and parallel work.	Wall displays: Turn-taking with single-user input in ad-hoc tasks [Russell et al. 2002]; more parallel work with multiple mice in mixed-motive tasks, but lower discussion quality [Birnholtz et al. 2007]; awkward turn-taking around vertical display with pen-based input [Rogers and Lindley 2004].	For multitouch wall displays, compare to other (e.g., mixed-motive) tasks; three or more simultaneous users.
Participants moved but stayed within touching distance 91% of the time.	Wall displays: Users mostly preferred to sit 1.5–2.5m away from wall display [Bi and Balakrishnan 2009]; many other studies have users' seated at a distance [Birnholtz et al. 2007; Isenberg et al. 2009].	Understand benefits of backing away; relative impact of direct touch or resolution on proximity.
Participants frequently switched between parallel and joint work.	Tabletop displays: Similar findings for the same task [Isenberg et al. 2012].	
Participants worked side by side, looking at the display and rarely at each other for longer periods of time.	Tabletop displays: Users often sit face-to-face across tabletop displays [Tang et al. 2006; Scott et al. 2004].	Understand how physical arrangement impacts awareness in group work.
Proximity associated with collaborative coupling: Participants physically navigated relative to each other while switching between parallel and joint work.	Tabletop displays: Similar findings where users frequently moved to gain a shared perspective [Tang et al. 2006].	
Users shared display evenly with no explicit negotiation for space. Larger display than used in related research.	Tabletop displays: Similar to "transient territories" when users move [Tang et al. 2006]; in contrast, clear use of territories by seated users [Scott et al. 2004].	Understand transient use of territories when users move; compare to larger groups and different types of task.
Users moved 300m, on average, navigating among different parts of the display.	Wall displays: Users preferred physical navigation to virtual navigation [Yost et al. 2007]; spatial organization on display may support external memory [Andrews et al. 2010].	Understand interplay of spatial cognition and movement in group tasks.

work and to answer questions that emerged from our review of previous research. Main findings are summarized and compared to earlier research in Table X, which also indicates areas where further research is needed. Next, we discuss both our findings and areas for future research.

5.1. Characteristics of Multitouch High-Resolution Wall Displays

Based on our findings, we revisit the characteristics of high-resolution multitouch wall-displays described earlier. The first characteristic of these displays is that they allow simultaneous input by multiple users. This seems important particularly for loosely coupled work, when group members are working in parallel. We saw several instances of participants interacting simultaneously with the display, for instance, when searching for different subjects at the same time as well as for working on a shared task of organizing information in the workspace (e.g., one group member writing notes, the other arranging windows). It seems that less coordination is required for task switching than with single-user input. Russell et al. [2002] found users quickly forming a turn-taking practice around a wall display, where they can easily coordinate their joint actions, but still the single-user input and the limited size seem to restrict parallel activity compared to larger multi-input displays. Birnholtz et al. [2007] found this to hold for mouse-based input: Multiple mice allowed groups to work in parallel, whereas a single mouse sometimes caused frustration to those not controlling the mouse. Scott et al. [2003] found support for recommending simultaneous user actions around tabletop displays, arguing that concurrent interaction would free users from having to monitor collaborators to know when the system is available and instead focus on the task at hand. We echo this recommendation for wall displays.

The second characteristic of these displays is that their resolution is sufficient for reading and carrying out other activities up close. Participants in the study worked within reach of the display 91% of the time. The interface did not allow participants to resize information on the display. Still, most text in the interface could be read from farther away. The ability to interact directly with the surface may induce physical proximity to the wall display. Judge et al. [2008] observed a similar need for close physical proximity to the wall in a study of affinity diagramming, which also involves direct interaction with the wall, albeit through paper notes. Display resolutions sufficient for reading up close may also invite physical proximity. For comparison, Bi and Balakrishnan [2009] reported that users preferred to sit between 1.5m and 2.5m away from a 5m wide display containing 32 pixels per inch—the seating position must trade off proximity to one part of the display with a view of the entire display. Most studies of collaboration around wall-sized displays have had users seated several meters from the display (e.g., Isenberg et al. [2009] and Birnholtz et al. [2007]). Seated group members have asked for additional visual aids for drawing attention to mouse cursors [Isenberg et al. 2009]. We got no such feedback from participants. Joint visual attention may be eased by close proximity to the wall display where group members can easily point and gesture.

The third characteristic of these displays is that they provide ample space for groups to work both jointly and in parallel. The data describing verbal communication, visual attention, and proximity of participants suggest that physical closeness—participants were often within touching distance of each other—is associated with joint work, while participants more often were further away from each other when working in parallel, on separate parts of the display. Perhaps more important, participants seemed to switch frequently and fluidly between parallel and joint work. This was indicated by the short overall duration of the codes for different modes of communication, attention, and proximity between participants. One challenge in collaborative work is that participants need to monitor each others' work to maintain awareness [Gutwin and Greenberg 2002]. Our video coding showed that participants frequently talked, participants commented in the questionnaire that they talked to keep each other updated, and we observed what Isenberg et al. [2009] describe as "running commentary." However, a few participants in our study suggested problems with monitoring each others' work. We saw participants making frequent glances at each other, and we wonder if group monitoring is more difficult when standing side by side up close to a wall display than when members can all view the entire display [Birnholtz et al. 2007].

5.2. Proxemics in Tightly and Loosely Coupled collaboration

Previous research has described how coupling style relates to the proximity and physical arrangement of users around displays and has suggested how different types of display support collaboration differently. Based on our results, we discuss whether these observations hold also for wall-sized multitouch displays, which allow for other physical arrangements. We also discuss how different findings might be attributed to differences in the display setups used in research.

The analysis of collaboration in groups based on the proxemics data showed a significant association between participants' proximity to each other and how tightly coupled they were working. When they were within arms reach of each other, at *intimate* or *personal close* proxemic distance, they were more likely to visually attend to the same area of the display or to each other and more likely to talk together. In contrast, when they were farther from each other, at *personal far* or *social* distance, they were more often silent and more often attending to different areas of the display. This supports the finding of Tang et al. [2006] that group members stood physically close when working closely together and further apart when working independently around a tabletop display. The present study extends previous work in showing that users' proximity to each other seems related to their coupling style regardless of the orientation of the display.

In addition to proximity, the role of users' physical arrangement and orientation toward each other has also been discussed in previous work. Clearly, tabletop displays support different physical arrangements that wall displays—users can sit facing each other on opposite sides of a tabletop, but this is not possible with a wall-display—and some physical arrangements may be beneficial for particular tasks. A face-to-face arrangement was common in the study of Tang et al. [2006], which they suggested was because it gave a good position to work on the same problem while providing face-to-face communication. However, Tang et al. [2006] also suggested that side-by-side may reduce visual distraction when group members work independently. Isenberg et al. [2012] highlighted two benefits of a face-to-face arrangement. One benefit was artifact-centered information sharing: Team members could point to relevant information on the display that their partner could see. The other benefit was direct communication: Team members did not need to turn away from the table to communicate. Indeed, participants in our study sometimes had to turn their head or move closer in order to view the same document. However, we question whether this is a consequence of the larger size rather than the vertical orientation of the wall display. The size of the tabletop display used by Isenberg et al. [2012] allowed both participants to see the entire display at all times. This size of display requires less effort from participants to share information, but this comes at the cost of less display space for individual work and for spatially organizing information for the task, as was also noted by Isenberg et al. [2012]. Also, whether groups benefit from particular physical arrangements may depend on the information being shared on the display. For users sitting across from each other around a tabletop, text documents, for instance, probably need to be reoriented for reading by different users; participants in our study were often observed to move closer to read the same document on the wall-display. Also, Tang et al. [2006] found that groups in one study were "highly mobile, with individuals frequently moving around the table to gain a shared perspective of the area of interest." In summary, the relative benefits of different physical arrangements and the resulting tradeoff between different display orientations that support such physical arrangements are not straightforward and are likely task dependent. Observations from the present study suggest that wall-sized displays provide more space for group members to work in parallel side by side. This possibly reduces visual distraction [Tang et al. 2006], while allowing group members to fluidly share information by turning their head or moving closer.

Concerning physical arrangement, previous research has also suggested problems with collaborating around vertical displays that did not emerge in the present study. Rogers and Lindley [2004] concluded that horizontal surfaces encouraged groups to work together in more cohesive ways than with a vertical display and they suggested several possible reasons: Groups shared a single Mimio pen for input that was found awkward to hand to each other, and only one group member was standing by the display at a time while the others had to stand back or be seated. Based on their findings, they suggested designing "vertical displays [...] that enable all group members to view the whole display and interact with it at the same time, without having to stand to the side, or feel uncomfortable switching between tasks or socially awkward." The type of wall display we studied might overcome some of the problems observed by Rogers and Lindley [2004]: Group members could view and interact with the display simultaneously through multitouch without having to stand to the side (given the much larger display of 280cm×120cm compared to 59cm×43cm used in Rogers and Lindley [2004]). Our data show that participants worked up close for 91% of the time, that they reported working effectively as a team, and that they frequently switched between individual and joint work. This suggests clear benefits of wall-sized multitouch displays for colocated groups. However, the present findings need to be extended to larger groups (Rogers and Lindley [2004] studied three-person groups) to understand the limitations in physical arrangements around this type of wall display.

5.3. Display Sharing and Territoriality

We question how groups use territories for coordinating their work around a wall display that they can physically navigate. Vogel and Balakrishnan [2004] claimed that explicit space partitioning is required for sharing a large vertical display. Participants in our study, however, seemed to work in parallel without explicitly negotiating for space. Many of the pairs shared the display evenly. Different factors may be at play. Size has been suggested to influence territoriality: If space is insufficient, groups may require more explicit coordination in switching between group work and individual work [Scott et al. 2004]. The display sharing behavior seen in our study seems to support the idea that territories are less important for coordinating group work for this type of display, at least for the mixed-focus task studied here.

Movement may also influence territoriality. In studies of groups seated in fixed positions around a shared display, users have tended to partition space into personal and other territories [Scott et al. 2004; Kruger et al. 2004; Isenberg et al. 2012]. Such space partitioning was less prominent in our study. Also, Tang et al. [2006] found that users that moved displayed less territoriality. As users moved, "others were no longer restricted from operating in those areas" [Tang et al. 2006].

Tang et al. [2006] also found that individuals avoided interacting with areas physically close to their partner, but such territorial behavior was not clear in our study. We did not see participants physically push or grab each other to gain access to areas of the display as Tang et al. [2006] saw. Situations where one participant reached forward to interact with the area in front of the other participant were not uncommon. These instances occurred during tightly coupled work where, as Tang et al. [2006] suggested, the goals and intentions of each individual are often known to all. Participants' interactions did not seem to interfere with each other's work. We did not systematically ask participants about interference, however.

Finally, the type of mixed-focus task we studied may invite groups to divide the work in ways that do not require them to partition the display into territories. All but two pairs answered in the questionnaire that they divided tasks between them. Almost all participants said they divided search or reading among them. The working area needed for carrying out these individual activities is likely easy for two persons to find on the wall display. None of the pairs in the present study seemed to run out of display space during the task. We are not sure why or under which circumstances users might run out of space. The availability of space might influence how users share the display in general and the benefit of ample space needs to be explored further. Important future work, therefore, is to investigate use of wall displays in the long term, for other tasks (e.g., involving large datasets), and for larger group sizes.

5.4. Physical Navigation in Group Work

Last, we were interested in understanding the role of physical navigation in group work around a wall display. When people are allowed to move freely, they may benefit from physically navigating a wall display [Ball and North 2007]. One difference from the present work to studies of seated interaction with large high-resolution displays is that participants could effectively navigate many documents through physical movement: They could turn their head, lean, or walk to access a particular document. Participants thus moved more than 300m, on average, often passing each other to access different parts of the display or to switch between coupling styles.

Physical navigation around a wall display may also benefit groups working on a shared task because they have a common spatial reference to information on the display. This may be particularly so for the type of task that participants collaborated on in this study. Andrews et al. [2010] used the same task. They suggested that the display may provide a form of external memory that allows the user to spatially reference information. Several users working together can build a shared spatially organized workspace in which information can be spatially referenced, which may support communication and coordination in the group. For instance, group members' arrangement in front of particular parts of the workspace provide context for their conversation; physical navigation of one member of a pair to another part of the display help the other member identify an opportunity for joint work. To understand these potential benefits, more empirical studies are needed that investigate the interplay of spatial organization and physical navigation in group work.

Participants' movement in front of the display may have been influenced by the design and implementation of the wall display interface. For instance, four search buttons were placed in fixed positions on the display. This provided multiple entry points, allowing both participants to search at the same time. However, the fixed positions restrict participants to particular locations for initiating a search. Moreover, the interface provided no functionality for grouping and organizing information on the display. Individual windows had to be moved manually. Organizing documents by moving them around through touch interaction may have been cumbersome.

5.5. Design of Wall Display Support for Colcated Work

Our study findings may have implications for the design of wall display systems that support colocated work. First, we found participants to collaborate in a variety of ways. Some groups in our study worked tightly together in close proximity to each other much of the time, and other groups worked most of the time on different areas of the display while talking. Wall display systems need to accommodate this variety of work styles. The implications of our findings are similar to research on tabletop collaboration: Groups were found to use a variety of coupling styles [Tang et al. 2006], different work styles and collaboration strategies [Isenberg et al. 2012], which tabletops should be designed to flexibly support. Additionally, our results show that groups frequently switched between parallel work and joint work. The type of mixed-focus tasks that we studied requires support for fluid transitions, as also suggested by Tang et al. [2006].

One way of flexibly supporting colocated collaboration is by using proxemics data as input [Ballendat et al. 2010]. For instance, we found an association between collaboration style and proximity, which can be utilized for supporting different interactions depending on proxemics. Systems may utilize more fine-grained tracking data (e.g., changes in orientation) than what we collected in this study. Such data can be used to infer when groups switch from parallel to joint work and thereby support these transitions.

The study also suggested that the wall display provided sufficient space for members to work effectively in parallel on separate areas of the display and, therefore, with less risk of distraction. Participants worked in parallel for longer periods at a time. Participants said they communicated to keep each other updated. However, some participants' comments suggest a need for helping loosely coupled groups to gain an awareness of each others work. Persistent information about interactions of group members with the display (e.g., who last touched or moved a document) could support colocated work on wall displays as well as remote collaboration [Hajizadeh et al. 2013]. Such information may further help the transition from loosely coupled to tightly coupled work. Finding that tightly coupled groups were more successful on the task, Isenberg et al. [2012] suggested showing clear visual connections between the information that group members work on to encourage coupling.

Last, the interface used allows free organization of searches and documents, and allows users to create their own instances of a document. Thus, it may support different styles of working on the task that we studied. Isenberg et al. [2012] that used a similar interface for the same task conclude so. We also think that having multiple entry points for searching, as mentioned earlier, may have supported parallel activity. However, similar to Isenberg et al. [2012], participants in our study missed features for grouping information and for linking associated documents shown on the display.

6. CONCLUSION

Wall-sized displays with higher resolution and support for multitouch offer new ways of supporting colocated collaboration. This article contributes empirical research to help understand the characteristics of such displays: They can be used up close by several users at a time, they offer high resolution for working up close, and they provide sufficient space for varied collaboration styles. The study examined pairs working on a problem-solving task involving a collection of documents. Video recordings and data from tracking of participants were analyzed to describe the communication and attention of group members, their navigation relative to each other and to the display, and their sharing of the display.

The results extend findings from studies of people collaborating around other types of shared display. The wall display used in the study seemed to accommodate different working styles (tabletop displays have similarly been found to support varied collaboration styles [Tang et al. 2006; Isenberg et al. 2012]); proximity between group members was associated with how tightly coupled they were working (similar to findings of Tang et al. [2006]); and groups switched fluidly between parallel and joint work (as they did in the study of Isenberg et al. [2012]).

The results also show differences to findings from earlier research on collaboration around large displays. For instance, several studies have found groups to use territories for managing space on tabletop displays [Scott et al. 2004; Kruger et al. 2004]. The groups in the present study shared the display quite evenly, which is similar to the finding that "'territories' were transient" when users moved [Tang et al. 2006]. Groups likely need to negotiate for space only when it is scarce. The space provided by the wall display seemed ample for coordinating group work on the task that we studied. However, more studies are needed to understand the needs for space on wall displays and transient use of territories when users move.

More work is needed to understand how groups can benefit from high-resolution multitouch wall displays. In particular, the present study needs to be extended to larger groups and to other types of task that may bring out other group dynamics around this type of display. Several pertinent questions need answering: How will

Table XI. Observed and Expected (in Parentheses) Joint Frequencies of Visual Attention and Verbal Communication (Seconds)

	Conversation	One is talking	Silence
Same area	6034	5748	6381
	(5314)	(4770)	(8079)
Look	707	660	281
	(482)	(433)	(733)
Eye Contact	442	277	24
	(217)	(195)	(330)
Divided or mutual attention	6334	2372	1018
	(2845)	(2554)	(4325)
Different areas	6325	8768	22231
	(10921)	(9802)	(16601)
Disengaged	52	31	307
	(114)	(102)	(173)

Table XII. Observed and Expected (in Parentheses) Joint Frequencies of Verbal Communication and Proximity (Seconds)

	Intimate	Personal close	Personal far	Social
Conversation	3959	7908	4889	3003
	(3637)	(6917)	(5160)	(4046)
One is talking	3410	6640	4223	3260
	(3227)	(6138)	(4578)	(3590)
Silence	5089	9149	8564	7596
	(5595)	(10642)	(7938)	(6224)

wall displays support concurrent interaction of larger groups, how will large groups be able to monitor each other, and what territorial behavior will emerge? How will larger groups interweave individual and joint work (for instance, compared to displays that enforce turn-taking)? Other questions have emerged in the discussion of our results and of previous research. These questions may help guide future research on colocated collaboration around wall displays.

APPENDIX A: JOINT FREQUENCIES AND ADJUSTED RESIDUALS

Results reported in Section 4.1 are based on the observed joint frequencies (o_{rc}) and expected frequencies (e_{rc}) of individual pairs of codes of visual attention, verbal communication, and proximity. Table XI shows the joint frequencies that are used for calculating the adjusted residuals shown in Table VI, Table XII shows the joint frequencies that are used for calculating the adjusted residuals shown in Table VIII, and Table XIII shows the joint frequencies that are used for calculating the adjusted residuals shown in Table IX. The expected frequencies are based on the assumption that the two categories are not related.

The adjusted residual z_{rc} for the pair of codes at row r and column c is calculated as

$$z_{rc} = \frac{(o_{rc} - e_{rc})}{\sqrt{e_{rc}(1 - p_c)(1 - p_r)}}, \quad (1)$$

where p_c and p_r is the probability for the column c and row r, respectively. "The adjusted residual indicates the extent to which an observed joint frequency differs from chance: It is positive if the observed is greater than chance and negative if the observed is less than chance" [Bakeman and Quera 2011, p. 110]. The magnitude of adjusted residuals can be compared across pairs of codes within the same contingency table. According to Bakeman and Quera [2011], the adjusted residuals can be compared to a 1.96 criterion

Table XIII. Observed and Expected (in Parentheses) Joint Frequencies of Visual Attention and Proximity (Seconds)

	Intimate	Personal close	Personal far	Social
Same area	8326	7080	1576	696
	(3287)	(6203)	(4597)	(3591)
Look	120	376	525	579
	(297)	(561)	(416)	(325)
Eye Contact	129	265	214	132
	(138)	(260)	(192)	(150)
Divided or mutual attention	1568	4611	2354	971
	(1767)	(3335)	(2472)	(1931)
Different areas	2153	10879	12441	11070
	(6794)	(12822)	(9503)	(7424)
Disengaged	50	88	159	42
	(63)	(119)	(88)	(69)

for indicating significance at the 0.05-level, or typically a higher criterion value around 3 because of the increased risk of type I errors with multiple tests.

ACKNOWLEDGMENTS

We thank Niels Dalum Hansen for help with developing the tracking system. Niels also assisted in conducting the study, as did Søren Sander—we thank them both. We also thank Cæcilie Kolling-Wedel and Johan Lynnerup for helping with the analysis of video recordings, and Jørgen Bansler for providing useful comments on an earlier draft of the article. Finally, we thank the reviewers for their thorough comments and help on improving this article.

REFERENCES

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*. ACM, New York, NY, 55–64.

Mark Ashdown, Philip Tuddenham, and Peter Robinson. 2010. High-resolution interactive displays. In *Tabletops - Horizontal Interactive Displays*, Christian Muller-Tomfelde (Ed.). Springer Verlag, 71–100.

Roger Bakeman and Vicenc Quera. 2011. *Sequential Analysis and Observational Methods for the Behavioral Sciences*. Cambridge University Press, Cambridge, UK.

Robert Ball and Chris North. 2005a. Analysis of user behavior on high-resolution tiled displays. In *Proceedings of INTERACT: the IFIP TC13 International Conference on Human-Computer Interaction*. Springer, 350–363.

Robert Ball and Chris North. 2005b. Effects of tiled high-resolution display on basic visualization and navigation tasks. In *Extended Abstracts of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 1196–1199. DOI:http://dx.doi.org/10.1145/1056808.1056875

Robert Ball and Chris North. 2007. Visual Analytics: Realizing embodied interaction for visual analytics through large displays. *Comput. Graph.* 31, 3 (2007), 380–400.

Robert Ball and Chris North. 2008. The effects of peripheral vision and physical navigation on large scale visualization. In *Proceedings of Graphics Interface*. Canadian Information Processing Society, Toronto, Ontario, Canada, 9–16.

Till Ballendat, Nicolai Marquardt, and Saul Greenberg. 2010. Proxemic interaction: designing for a proximity and orientation-aware environment. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (ITS'10)*. ACM, New York, NY, 121–130. DOI:http://dx.doi.org/10.1145/1936652.1936676

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*. ACM, New York, NY, 1005–1014.

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*. ACM, New York, NY, 91–100.

Harry Brignull, Shahram Izadi, Geraldine Fitzpatrick, Yvonne Rogers, and Tom Rodden. 2004. The introduction of a shared interactive surface into a communal space. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*. ACM, New York, NY, 49–58.

Jacob Cohen. 1988. *Statistical Power Analysis for the Behavioral Sciences*. L. Erlbaum Associates.

Mary Czerwinski, Greg Smith, Tim Regan, Brian Meyers, George Robertson, and Gary Starkweather. 2003. Toward characterizing the productivity benefits of very large displays. In *Proceedings of INTERACT: the IFIP TC13 International Conference on Human-Computer Interaction*. 9–16.

David Douglas and Thomas Peucker. 1973. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *Cartographica* 10, 2 (Oct. 1973), 112–122.

H. J. Dudfield, C. Macklin, R. Fearnley, A. Simpson, and P. Hall. 2001. Big is better? Human factors issues of large screen displays with military command teams. In *Proceedings of the 2nd International Conference on Human Interfaces in Control Rooms, Cockpits and Command Centres. People in Control. (IEE Conf. Publ. No. 481)*. IEE, 304–309. DOI:http://dx.doi.org/10.1049/cp:20010480

Scott Elrod, Richard Bruce, Rich Gold, David Goldberg, Frank Halasz, William Janssen, et al. 1992. Liveboard: a large interactive display supporting group meetings, presentations, and remote collaboration. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 599–607.

Georges Grinstein, Theresa O'Connell, Sharon J. Laskowski, Catherine Plaisant, Jean Scholtz, and Mark Whiting. 2006. The VAST 2006 Contest: A tale of Alderwood. In *Proceedings of IEEE Symposium on Visual Analytics Science and Technology*. IEEE, 215–216.

Jonathan Grudin. 2001. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 458–465. DOI:http://dx.doi.org/10.1145/365024.365312

François Guimbretière, Maureen Stone, and Terry Winograd. 2001. Fluid interaction with high-resolution wall-size displays. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, 21–30.

Carl Gutwin and Saul Greenberg. 2002. A descriptive Framework of Workspace Awareness for Real-Time Groupware. *Computer Supported Cooperative Work* 11 (November 2002), 411–446.

Amir H. Hajizadeh, Melanie Tory, and Rock Leung. 2013. Supporting awareness through collaborative brushing and linking of tabular data. *IEEE Trans. Visual. Comput. Graph.* 19, 12 (2013), 2189–2197.

Edward T. Hall. 1963. A system for the notation of proxemic behaviour. *American Anthropologist* 65, 5 (1963), 1003–1026. DOI:http://dx.doi.org/doi:10.1525/aa.1963.65.5.02a00020

Edward T. Hall. 1966. *The Hidden Dimension*. Doubleday, Garden City, NY.

Michael Haller, Jakob Leitner, Thomas Seifried, James R. Wallace, Stacey D. Scott, Christoph Richter, Peter Brandl, Adam Gokcezade, and Seth Hunter. 2010. The NiCE Discussion Room: integrating paper and digital media to support co-located group meetings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 609–618.

Kirstie Hawkey, Melanie Kellar, Derek Reilly, Tara Whalen, and Kori M. Inkpen. 2005. The proximity factor: impact of distance on co-located collaboration. In *Proceedings of the ACM SIGGROUP Conference on Supporting Group Work*. ACM, New York, NY, 31–40.

Marti A. Hearst. 1995. TileBars: visualization of term distribution information in full text information access. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM Press, New York, NY, 59–66. DOI:http://dx.doi.org/10.1145/223904.223912

Eva Hornecker, Paul Marshall, Nick Sheep Dalton, and Yvonne Rogers. 2008. Collaboration and interference: awareness with mice or touch input. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*. ACM, New York, NY, 167–176.

Elaine M. Huang, Elizabeth D. Mynatt, Daniel M. Russell, and Alison E. Sue. 2006. Secrets to success and fatal flaws: the design of large-display groupware. *IEEE Comput. Graph. Appl.* 26, 1 (2006), 37–45.

Dugald Ralph Hutchings, Greg Smith, Brian Meyers, Mary Czerwinski, and George Robertson. 2004. Display space usage and window management operation comparisons between single monitor and multiple monitor users. In *Proceedings of the Working Conference on Advanced Visual Interfaces*. ACM Press, New York, NY, 32–39. DOI:http://dx.doi.org/10.1145/989863.989867

Kori Inkpen, Kirstie Hawkey, Melanie Kellar, Regan Mandryk, Karen Parker, Derek Reilly, Stacey Scott, and Tara Whalen. 2005. Exploring display factors that influence co-located collaboration: angle, size, number, and user arrangement. In *Proceedings of HCI International*.

Petra Isenberg, Anastasia Bezerianos, Nathalie Henry, Sheelagh Carpendale, and Jean-Daniel Fekete. 2009. CoCoNutTrix: collaborative retrofitting for Information Visualization. *Computer Graphics and Applications: Special Issue on Collaborative Visualization* 29, 5 (September/October 2009), 44–57.

Petra Isenberg, Danyel Fisher, Sharoda A. Paul, Meredith Ringel Morris, Kori Inkpen, and Mary Czerwinski. 2012. Co-located collaborative visual analytics around a tabletop display. *IEEE Trans. Visualiz. Compu. Graph.* 18, 5 (2012), 689–702.

Giulio Jacucci, Ann Morrison, Gabriela T Richard, Jari Kleimola, Peter Peltonen, Lorenza Parisi, and Toni Laitinen. 2010. Worlds of information: designing for engagement at a public multi-touch display. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 2267–2276.

Tejinder K. Judge, Pardha S. Pyla, D. Scott McCrickard, Steve Harrison, and H. Rex Hartson. 2008. *Studying Group Decision Making in Affinity Diagramming.* http://eprints.cs.vt.edu/archive/00001043/. (Jan. 2008). http://eprints.cs.vt.edu/archive/00001043/

Youn-ah Kang, C Gorg, and John Stasko. 2009. Evaluating visual analytics systems for investigative analysis: Deriving design principles from a case study. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology (VAST)*. IEEE, 139–146.

Russell Kruger, Sheelagh Carpendale, Stacey D. Scott, and Saul Greenberg. 2004. Roles of orientation in tabletop collaboration: comprehension, coordination and communication. *Comput. Supported Coop. Work* 13, 5–6 (2004), 501–537.

Paul Marshall, Eva Hornecker, Richard Morris, Sheep Dalton, and Yvonne Rogers. 2008. When the fingers do the talking: A study of group participation for different kinds of shareable surfaces. In *Proceedings of TABLETOP*. IEEE Computer Society, 37–44.

O. C. Mudford, N. T. Martin, J. K. Hui, and S. A. Taylor. 2009. Assessing observer accuracy in continuous recording of rate and duration: three algorithms compared. *J. Appl Behav Anal.* 42, 3 (2009), 527–539.

Christian Müller-Tomfelde. 2010. *Tabletops - Horizontal Interactive Displays* (1st ed.). Springer.

T. Ni, G. Schmidt, O. Staadt, M. Livingston, R. Ball, and R. May. 2006. A survey of large high-resolution display technologies, techniques, and applications. In *IEEE Virtual Reality*. Alexandria, Virginia, 223–234.

Gary M. Olson and Judith S. Olson. 2000. Distance matters. *Hum.-Comput. Interact.* 15 (September 2000), 139–178. Issue 2. DOI:http://dx.doi.org/10.1207/S15327051HCI1523_4

George Robertson, Mary Czerwinski, Patrick Baudisch, Brian Meyers, Daniel Robbins, Greg Smith, and Desney Tan. 2005. The large-display user experience. *IEEE Comput. Graph. Appl.* 25, 4 (2005), 44–51. DOI:http://dx.doi.org/10.1109/MCG.2005.88

Y. Rogers and S. Lindley. 2004. Collaborating around vertical and horizontal large interactive displays: which way is best? *Interact. Comput.* 16, 6 (Dec. 2004), 1133–1152.

Robert Rosenthal and Ralph L. Rosnow. 1991. *Essentials of Behavioral Research: Methods and Data Analysis* (second ed.). McGraw-Hill.

Daniel M. Russell, Clemens Drews, and Alison Sue. 2002. Social aspects of using large public interactive displays for collaboration. In *Proceedings of the International Conference on Ubiquitous Computing*. Springer-Verlag, London, UK, 229–236.

Kathy Ryall, Clifton Forlines, Chia Shen, and Meredith Ringel Morris. 2004. Exploring the effects of group size and table size on interactions with tabletop shared-display groupware. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*. ACM, New York, NY, USA, 284–293.

Tony Salvador, Jean Scholtz, and James Larson. 1996. The Denver model for groupware design. *SIGCHI Bull.* 28, 1 (Jan. 1996), 5258. DOI:http://dx.doi.org/10.1145/249170.249185

Stacey D. Scott, Karen D. Grant, and Regan L. Mandryk. 2003. System guidelines for co-located, collaborative work on a tabletop display. In *Proceedings of the European Conference on Computer Supported Cooperative Work*. Kluwer Academic, Norwell, MA, 159–178.

Stacey D. Scott, M. Sheelagh, T. Carpendale, and Kori M. Inkpen. 2004. Territoriality in collaborative tabletop workspaces. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*. ACM, New York, NY, 294–303.

Garth Shoemaker, Takayuki Tsukitani, Yoshifumi Kitamura, and Kellogg S. Booth. 2010. Body-centric interaction techniques for very large wall displays. In *Proceedings of the Nordic Conference on Human-Computer Interaction*. ACM, New York, NY, 463–472.

Norbert A. Streitz, Jörg Geissler, Torsten Holmer, Shin'ichi Konomi, Christian Müller-Tomfelde, Wolfgang Reischl, Petra Rexroth, Peter Seitz, and Ralf Steinmetz. 1999. i-LAND: an interactive landscape for creativity and innovation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 120–127.

Desney S. Tan, Darren Gergle, Peter Scupelli, and Randy Pausch. 2006. Physically large displays improve performance on spatial tasks. *ACM Trans. Comput.-Hum. Interact.* 13, 1 (2006), 71–99. DOI:http://dx.doi.org/10.1145/1143518.1143521

Anthony Tang, Melanie Tory, Barry Po, Petra Neumann, and Sheelagh Carpendale. 2006. Collaborative coupling over tabletop displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 1181–1190.

John C. Tang. 1991. Findings from observational studies of collaborative work. *Int. J. Man-Mach. Stud.* 34 (February 1991), 143–160.

Daniel Vogel and Ravin Balakrishnan. 2004. Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. In *Proceedings of the ACM Symposium on User Interface Software and Technology*. ACM, New York, NY, 137–146.

Katherine Vogt, Lauren Bradel, Christopher Andrews, Chris North, Alex Endert, and Duke Hutchings. 2011. Co-located collaborative sensemaking on a Large High-Resolution Display with Multiple Input Devices. In *Proceedings of INTERACT: the IFIP TC13 International Conference on Human-Computer Interaction*, Vol. 6947. 589–604.

Beth Yost, Yonca Haciahmetoglu, and Chris North. 2007. Beyond visual acuity: the perceptual scalability of information visualizations for large displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, New York, NY, 101–110. DOI:http://dx.doi.org/10.1145/1240624.1240639

Received March 2013; revised November 2013; accepted January 2014