Presence & Placement: Exploring the Benefits of Multiple Shared Displays on an Intellective Sensemaking Task

Christopher Plaue and John Stasko
School of Interactive Computing
Georgia Institute of Technology
85 5th Street NW
Atlanta, GA 30332-0760
+1 404 894-5617
{plaue, stasko} @cc.gatech.edu

ABSTRACT
Relatively little is known about *how* the presence and location of multiple shared displays changes the performance and dynamics of teams collaborating. We conducted a case study evaluating several shared display configurations with groups collaborating on a data-intensive, sense-making task. Teams completed the same task using either a single display, side-by-side dual, or opposing dual shared displays. The location of the second shared display significantly impacted the ability for teams to make logical connections amongst the data. Users were also significantly more satisfied with the collaboration process using the side-by-side dual display condition than those using a single display.

Categories and Subject Descriptors
H5.2. **Information interfaces and presentation** (e.g., HCI): User Interfaces, graphical user interfaces (GUI), miscellaneous.

General Terms
Measurement, Performance, Experimentation, Human Factors

Keywords
Multi-display environments, meeting spaces, large displays

1. INTRODUCTION
Recent trends in business data indicate that more people, especially managers, are spending even more of their workday in meetings [6]. In conjunction, current trends within many companies indicate a shift away from traditional employee hierarchies to low, flat hierarchies which use self-managing teams [19]. As a result, team members must often make decisions together rather than having one made for them. It is logical to assume that individuals are spending time away from their individual offices working in collaborative spaces.

GROUP '09, May 10—13, 2009, Sanibel Island, FL, USA.

Currently, meeting or conference rooms are commonly equipped with a single large display such as a projector or flatscreen monitor along with a whiteboard and table. However, advances in technology and connectivity now allow workers to ubiquitously access information in collaborative spaces. Computing devices such as laptops or smart phones may be used to share information, presentations, or charts. Thus, meeting spaces seem likely candidates for deploying more than one large display to support collaborative information sharing. Several research lab applications have incorporated multiple large shared displays into their environments [1, 16, 23], but multiple shared displays do not appear to have yet gained widespread deployment in everyday meeting rooms.

While research has shown the benefits of multiple displays for users working on individual workstations [11], we wonder if multiple large, shared displays offer similar benefits to groups of people collaborating within meeting environments. Or, do such displays provide an additional drain on limited cognitive resources? Existing research within the multiple-display environment (MDE) community regarding collaborative spaces typically examines the software infrastructure framework and interaction techniques used to share information. Furthermore, evaluation of MDEs often uses performance-based tasks, such as tracking tasks, or focuses on single users. In contrast, we examine the impact of the *presence of* and *location of* shared displays on a meeting process involving multiple participants. In our controlled user study, we evaluate participants performing a sensemaking task to simulate one type of activity that occurs within meetings. We seek to understand how the presence of shared displays can impact the group sense-making process. Finally, we also discuss user attitudes, constraints, and limitations of such environments.

2. RELATED WORK
Meeting spaces, due to the interaction between individuals, technology, and information, represent an intriguing research domain not only for technology researchers, but also to the management, business, and psychology communities. In this section, we present relevant related work under three themes: research on group decision support systems (GDSS) within the management information science community, collaborative tasks, and multiple display environment (MDE) research.

2.1 Group Decision Support Systems (GDSS)

The management information systems (MIS) community extensively explored group decision support systems (GDSSs), where a centralized computer system facilitates the exchange of ideas, opinions, and preferences within a group [10]. These systems may provide modeling and mathematical techniques or apply rules to the meeting process. However, some systems simply support communication amongst the group.

In an overview of GDSS research, Fjermestad and Hitz report conflicting evidence whether these systems impact the amount of time it takes for a group to reach a decision [8]. Gallupe and DeSanctis [10] found that GDSS usage did not affect the decision-time for complex tasks, but Bui and Sivasankaran [3] found that groups performing a low-complexity task using a GDSS actually required more time than individuals not using the system. It is not clear, for example, which components of a GDSS offer advantages to users. In this paper, we examine one particular aspect of a GDSS—that of shared displays—and its impact on the meeting process.

Furthermore, a survey of research within the MIS community indicates that most *evaluation studies* use fewer than five individuals per testing session, even though prior research shows GDSS system usage is more beneficial for larger groups [8]. We note within the HCI and CSCW community, empirical studies of collaborative settings often investigate smaller groups. For example, the study evaluating Roomware [25] used groups of four participants. In the study described in this paper, we conducted testing sessions with five or six individuals.

2.2 Collaborative Tasks

The psychology and MIS communities have extensively researched the types of tasks people perform while collaborating. Poole *et al* suggest that the type of task may account for up to half of the variance in group performance [20], and furthermore, it is not uncommon for many groups to collaborate in different ways. McGrath developed a classification scheme for the types of tasks groups typically perform, consisting of eight types within four categories [17]. Surprisingly, a survey of research conducted within the community indicates most studies clustered heavily within two of the eight genres; Fjermestad and Hiltz report that 52% of studies focused on *decision-making tasks*, where there is no definitive answer or objective measure of quality [8]. The second-most frequently used task in studies was *creativity* (i.e. brainstorming.)

In this study, we chose to use an *intellective* task, where a solution to a problem exists. Our particular task, described in more detail later in this article, provides groups with a large data set in which solutions to a problem were embedded. We sought to simulate the types of activities we observed occurring in meeting spaces where multiple people come into a room with information to share with other individuals. In addition, this type of task simulated certain aspects of sensemaking activities occurring in visual analytics [27], an emerging field combining information visualization with data analysis techniques. Visual analytics involves making sense of large data sets, where some data may be conflicting, in order to locate expected patterns as well as discover new correlations. However, to control variables in our study, we did not provide participants with any specific visualization tools.

2.3 Multiple Display Environments (MDE):

Research on MDEs has focused on both individual and shared environments. Czerwinski *et al* showed task performance benefits of using multiple displays for individual users [5]. Other research explored how individuals manage information while using multiple displays, finding that users often organize content in terms of primary and secondary importance [11, 13]. For example, content that does not have active focus is often relevant to a user. In this study, we explore how groups use multiple displays in collaborative environments; do they also divide content up in terms of primary and secondary importance, for example?

Several visions of futuristic meeting spaces explored large, shared display usage, such as the Interactive Workspaces project [16] and iLAND/RoomWare [23, 24]. More recently, researchers developed the IMPROMPTU framework to assist users in sharing information across displays using off-the-shelf products, supporting opportunistic, short-lived collaborative moments [2]. Much of this body of research focuses on the development of specialized software architectures facilitating information sharing across multiple devices and displays. Note, however, that many of these systems satisfy the definition of a GDSS since they support the exchange of ideas, opinions, and preferences in a group.

Evaluation of systems supporting MDEs typically focuses on interaction techniques [e.g. 12] and interface representations [e.g. 1]. Much of this body of research focuses on the development specialized software architectures allowing information sharing across multiple devices and displays.

Other MDE research evaluated display positioning. In [28], researchers provide guidelines for positioning displays with respect to the surface used to provide input to the system. Their study involved participants using a tablet to perform a docking task, using a stylus to select and move objects on a screen. This type of task is considered *performance and psycho-motor* in nature, using McGrath's framework. Similarly, Su and Bailey's work on creating guidelines for positioning large displays in interactive workspaces also used a target-task during evaluation and also focused on single users [26]. In this study, we evaluate display presence and positioning via the routine and ubiquitous act of sharing content from a laptop via shared display [18] while individuals perform a sense-making *intellective* task. Furthermore, building upon this existing work, evaluation occurs using multiple individuals for each session.

Finally, several studies also explore the effects of technology use on group work. For example, Streitz *et al*. conducted an empirical study comparing different Roomware configurations' impact on the collaboration process, finding individuals with access to both computer workstations and a Liveboard produced higher-quality work during a brainstorming exercise [25].

3. STUDY

The goal of this study was to evaluate the effects of *presence of* and *location of* multiple shared displays on a simulated meeting environment. Specifically, do dual-shared displays in meeting spaces increase the amount of insights discovered by group members during a sense-making task? Does a second shared display increase the potential for parallel work during the collaboration process? We acknowledge there are a plethora of different configurations and variables that could be tested and each of these variables may influence how displays are used. For

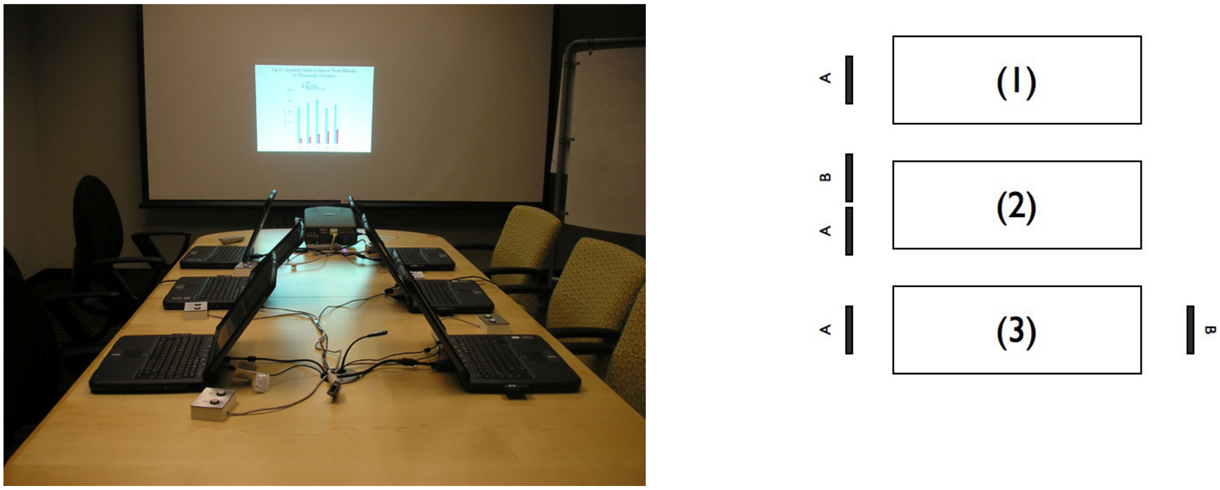

Figure 1. The experimental testing site with six laptop computers (left) and the three shared display configurations: single (1), side-by-side (2), and opposing (3).

the intents of this study, we sought to limit external variability as much as possible, focusing specifically on the three display configurations. We discuss specific design decisions in the following sections.

3.1 Participants

We recruited 105 individuals (26 female) to participate in the study. All but three individuals (a restaurant server, research scientist, and user experience engineer) were students at a technical university with an average age of 22.8 years ($s = 3.53$). Sixty-six participants were in a technology-related major, such as computer science, computational media, or human-computer interaction, and other majors included industrial systems engineering (10), industrial design (4), biology (3), biomedical engineering (3), and psychology (3). Student participants received course credit for their attendance.

We assigned individuals randomly to a testing condition. For this study, it was impractical to find and use existing work groups for evaluation purposes. However, the relatively homogenous student population did remove a source of uncontrolled variance. Student participants were not likely to have adopted a particular style or meeting role, as might be common within established corporate or academic environments. Since participants did not know a majority of their group members, special measures were taken in the design of the experiment to ensure that all group members were familiar with the role of their teammates (as described later in this section).

3.2 Materials

This study simulated an environment in which participants bring in laptop computers containing information relevant to the meeting topic that is likely to be shared. This act of sharing is a ubiquitous and routine practice in both industry and academic meeting rooms [18]. Our goal was not to replace this interaction of "showing data" via sophisticated technologies, but to explore how multiple shared displays impact collaborations. Three different configurations of shared displays were manipulated as the independent variable: single display, side-by-side dual displays, and opposing dual displays, as shown in Figure 1. We selected these configurations to investigate not only the effects of adding an additional shared display to the room, but also to begin to explore the influence of location.

The shared displays consisted of portable XGA projectors running at a resolution of 1024x768. For the multiple display conditions, each projected image was calibrated to be identical in physical size (approximately 1.2 meters diagonally). Due to the logistics of coordinating these sessions, we omitted a condition without any shared display, noting that many existing meeting spaces contain one shared display.

We provided each session with six laptop computers running only a fresh installation of Windows XP and Microsoft Office. Each laptop was preloaded with a unique Microsoft PowerPoint presentation containing information required to solve the task (described in detail in "primary task") when used in conjunction with information from other laptops. Furthermore, each laptop connected to an Altinex programmable video matrix switch.

We placed a control box (see Figure 2) next to each laptop to control the video switch. These control boxes consisted of a physical on/off button to enable mirroring content from a laptop onto a shared display. Pressing a button resulted in the button illuminating and the content being shown on the respective shared

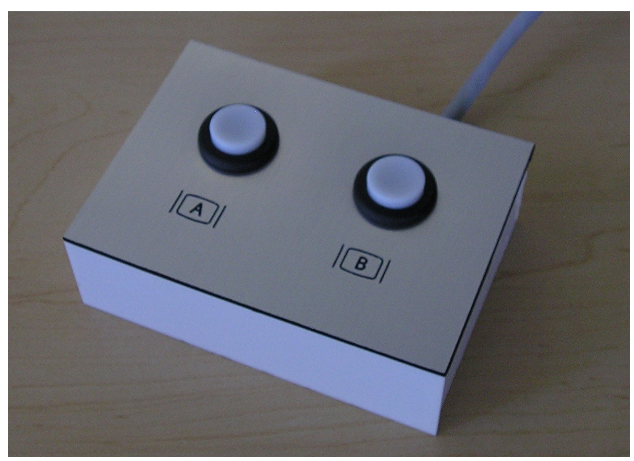

Figure 2. Control box used to share content on the large display in the meeting environment.

display (i.e. left button controls the left shared display). In the single display condition, a sleeve was placed over the second control button. A shared display could be blanked by a participant pressing the respective control button a second time. This switching system allowed individuals to quickly connect and share content onto one of the shared displays.

We acknowledge that several software-based solutions [e.g. 18] have been built by researchers, however none of these systems are widely used outside of research lab applications. Therefore, we opted to use this off-the-shelf video switching solution. Also, this physical interface required minimal training and was easily learned by participants during pilot testing. Again, we sought to simply provide a quick method for individuals to mirror the content of their laptops on the large shared displays.

We furnished the meeting space with a large conference table, six chairs (three on each side of the table), a portable white board, and two-level fluorescent lighting. The conference table seated a maximum of three individuals comfortably on each side of the table allowing for six participants per session. We kept lighting at mid-illumination to reduce glare and improve readability of the projector and provided participants with whiteboard markers, eraser, pens, and paper.

3.3 Primary Task

We provided a scenario based upon "The Bonanza Business Forms Company Case" [15], used in evaluating a GDSS [10]. The scenario posed to participants follows:

Bonanza Business Forms Company sells paper forms for three markets: small business, hospitals, and financial institutions. During the previous three quarters, Bonanza's profits were steadily decreasing while total sales were increasing. Bonanza's management cannot determine the cause of the declining profits, so they decided to bring you in as outside investigators. The goal of the outside investigation team is to determine the cause of the company's problem using a series of reports, charts, and data to identify why the problem is occurring.

We defined six investigator roles: domain researcher, industry trend analyst, sales force consultant, financial analyst, marketing consultant, and advertising trend consultant. Within an experiment session, participants were randomly assigned a role.

For this study, we modified the original business case to create PowerPoint presentations for each investigative role. Information was provided in the form of bulleted information, charts, and graphs. To provide a sense of realism, we formatted the graphs to appear in the default Microsoft Excel font and color scheme. Each role's slides were unique, thus it was not possible to accurately determine solutions to why Bonanza's profits were decreasing amidst increasing sales without individuals sharing information with each other.

Pilot testing indicated the scenario was solvable, yet not trivial, and did not require any specific management or business training. Pilot testing also indicated a potential problem resulting from the unfamiliarity of participants with every member of the group. This had the side effect of participants losing track of which role each person had. A person might, for example, want to see information pertaining to the advertising budget but forget which participant had the role of advertising consultant. To alleviate this problem, a colored band was incorporated at the top of each slide, coding and labeling each investigator role. Corresponding colored placards were placed on top of each participant's laptop.

Procedure: Experiment sessions consisted of six participants (full team), however, all three experimental conditions had one team in which a participant failed to show up. Rather than dismissing the remaining five participants, we chose to run those respective studies with five participants. We discuss these sessions in more detail later.

At the beginning of a session, we randomly assigned participants to a laptop computer and administered informed consent and a demographic survey. The experimenter read an introductory script, outlining that the ultimate goal for the team was to determine why Bonanza's profits were decreasing while sales were increasing. We instructed the group to solve the problem using any strategy they wished and introduced them to the whiteboard, paper, and video-sharing technologies. Furthermore, participants were told that all the information needed to solve the problem was provided. We also allowed participants to use any of the other built-in software applications if they wished, such as the calculator or notepad programs.

The experimenter instructed each group to hand in a list with their answer or answers to the business case dilemma. To encourage groups to be thorough *and* efficient, we offered a $20 per person incentive to the group that solved the scenario most correctly in the shortest period of time.

Finally, after reaching group consensus, each participant individually completed a closing survey. The purpose of this survey was to obtain attitudes towards the meeting process as well as agreement with the group decision.

3.4 Data Collection

We collected data via a combination of surveys, interviews, and video analysis. Each experiment session was videotaped and later coded at one-second intervals to determine 1) which slide was displayed on the shared display(s), 2) who was speaking, 3) who was gesturing towards a display, and 4) who was writing on the whiteboard. Furthermore, we transcribed the dialog from each testing session.

3.5 Insight-Based Evaluation Technique

We sought to evaluate group collaboration performance in the various display conditions using the metrics of *performance*, *collaboration*, and *satisfaction.* Aspects of collaboration and satisfaction were probed through surveys and interviewing. To evaluate performance, we used an insight-based methodology to evaluate each group's collaboration, based on [22]. An *insight* refers to a direct observation from the information provided that is relevant to solving the dilemma posed in the primary task. Examining the number of insights discovered by the teams offered a richer way of exploring collaboration habits than simply measuring time-to-completion. The main challenge was establishing objective metrics over what constitutes an insight.

We consulted two individuals with formal business education training: a 29-year-old project manager with a degree in business administration, and a 30-year-old Masters in Business Administration student. Each individual read through the case study and generated a list of insights presented in the task that lead to the decrease in profits at the fictional company. These raters assigned weighted values to these insights. Since these ratings were highly correlated, we combined them to establish an objective grading rubric to judge accuracy, completeness, and thoroughness of each group's collaboration session.

Furthermore, we defined five important *inferential links*, where groups correctly join two particular insights (i.e. the number of sales for X remain steady [*insight 1*] while the total dollar amount of sales for X is increasing [*insight 2*] → the amount of each sale is increasing [*inferential link*]).

4. RESULTS

As stated earlier, we sought to evaluate group collaboration performance in the various display conditions using metrics of *performance*, *collaboration*, and *satisfaction.* For performance, we investigated trends in discoveries of insights. For collaboration metrics, we considered both quantitative measures, such as whiteboard usage or deictic gesture [4] rates, as well as qualitative findings emerging from transcript analysis. We probed for satisfaction of the meeting process via surveys and interviews.

Each experimental condition consisted of five full teams of six participants at a time. All three experiment conditions had one team in which one participant failed to show up. Since each laptop contained only six slides of information and the participants would be assigned course credit regardless, we chose to run those sessions with five participants. We instructed these groups to double-up on laptop computers as the group deemed fit. Our analyses showed that performance metrics for these groups did not differ substantially from the teams of six; results were within one standard deviation from the mean. Therefore, we opt to include data from the sessions with five participants in the study analysis, noting that a majority of studies evaluating various group support systems used fewer than 5 participants [8].

Since the number of groups per condition is limited, we use a general comparison of trends as an analysis. However, we offer inferential statistical analysis when appropriate to also provide useful indicators of trends.

4.1 Performance

To explore the effect of display condition on performance, we examined the rate of insight and inferential link discoveries amongst the groups. We counted the total number of insights and inferential links for each group, as shown in Figure 3a and Figure 3c. However, the time for each group to reach consensus varied considerably per group, with a mean time to completion of 33 minutes across all groups. For example, one group in the side-by-side display condition reached consensus in just over 8 minutes while another took over 45. This wide variation is perhaps the greatest challenge in performing laboratory studies involving groups and collaboration; the way individuals interact with each other as well as establish a rapport may vary considerably.

To account for this variation in time-to-completion, we calculated an insight rate metric by dividing the number of insights discovered by the amount of time it took to reach consensus to obtain an insight rate. As Figure 3b illustrates, the average insight rate is highest for groups in the side-by-side display condition. However, this condition also has the greatest variance. The rate of insight is lowest for individuals in the single shared display condition.

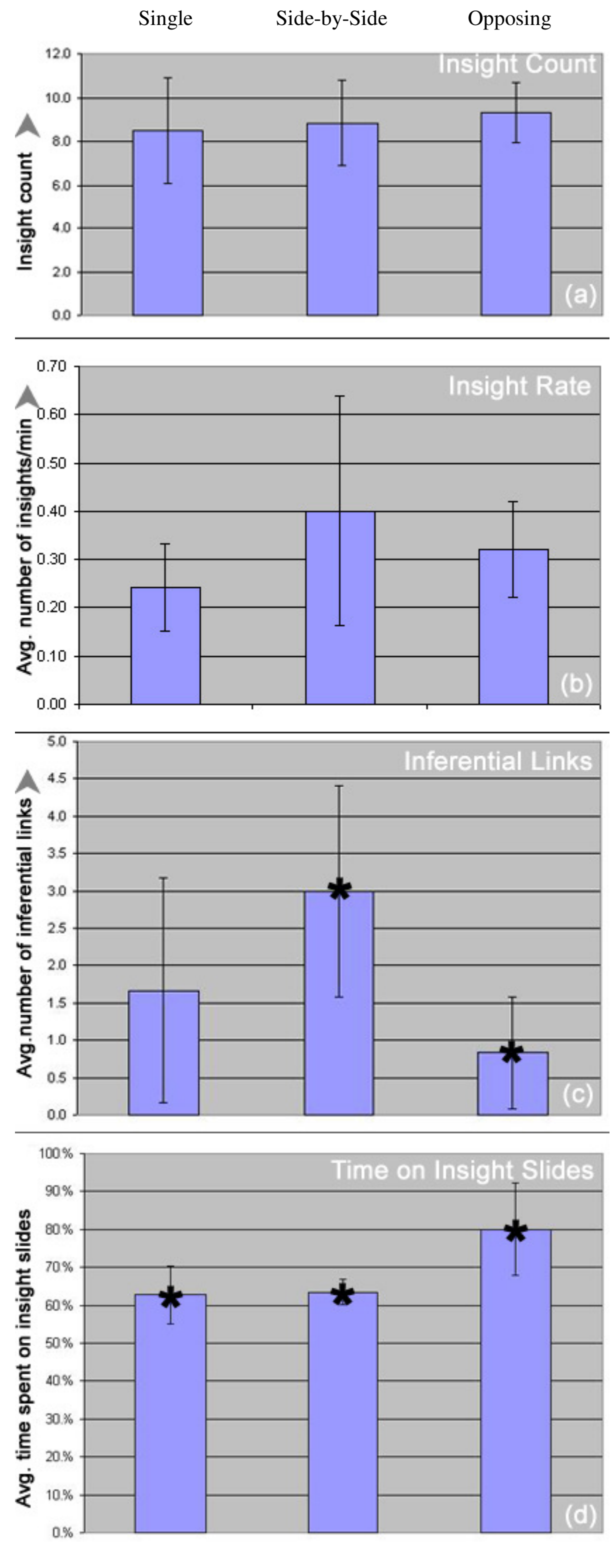

Figure 3. Count of insights (a), insight rate (b), and count of inferential links (c), and time spent on slides with insight content (d). (* indicates a significant difference in performance, and average number of slides viewed per minute. Y-axis arrows indicate the direction of better performance)

Table 1. Count of groups where whiteboard was used, across display conditions.

	Count	Total Duration at Whiteboard (*Average Time Per Encounter*)
Single	4	01:32(*00:46*), 38:45(*12:55*), 06:48(*01:21*), 05:11 (*01:43*)
Side-by-Side	1	06:07 (*00:33*)
Opposing	2	01:57(*00:58*), 04:11(*04:11*)

Table 2. Deictic gesture rate across display conditions.

	Gestures Per Min.	Std. Dev.	Average Duration (s)
Single	0.3	0.35	2.8
Side-by-side	1.1	0.94	2.2
Opposing	1.7	0.50	2.7

There was a significant effect ($F(2,17)=4.448$, $p=.030$) on the number of inferential links made per display condition. A post hoc analysis indicates that participants in the side-by-side condition correctly identified significantly more inferential links than those in the opposing display condition. Although the average number of inferential links was lower under the single display versus the side-by-side condition, this difference was not statistically significant.

We also analyzed the amount of time groups spent displaying slides on the shared displays. An ANOVA analysis yielded a significant difference in this factor ($F(2,17)=8.099$, $p=.004$). Groups in the opposing dual display condition spent a significantly larger percentage of their time showing slides containing insights on a shared display than those in the single display condition or side-by-side display condition.

4.2 Collaboration

To evaluate collaboration, we coded the videotape of each session for gesturing and whiteboard usage, each an act that involves interaction with other individuals in the testing process. Research has shown that individuals often use pointing gestures during collaboration to clarify or enhance a message [9], as shown in Figure 4. We counted each instance of deictic gesture towards the display via coding of the videotape of each session. Likewise, our video analysis also noted when a group member stood and actively wrote on the board.

4.2.1 Quantitative Measures

Groups in the single shared display condition used the whiteboard more frequently than in any other condition (Table 1). Whiteboard usage served primarily as a way to make note of themes or insights. They were also used as persistent information displays while the group performed an analysis of the scenario. The group that used the whiteboard for the longest amount of time in the single display condition did so to analyze trends in the PowerPoint Slides and combine graph trends.

Deictic gesturing towards the shared display was observed across all experiment conditions. To account for the wide range of completion times, we calculate a gesture rate, as shown in Table 2. An ANOVA analysis yielded statistical significance for gesture rate across the conditions ($F(2,17)=6.364$, $p=.010$). Groups in both multiple display configurations had higher gesture rates than those in the single display configuration. However, the average length of a gesture was similar across conditions: 2.8s, 2.2s, and 2.6s for the single, side-by-side, and opposing display conditions, respectively.

4.2.2 Qualitative Measures

To further probe issues of collaboration, we examined the video and transcripts of each collaboration session using inductive techniques to understand commonalities and themes in the data. In conjunction, the independent raters also inspected the transcripts. In particular, themes revolving around how the second shared display impacted collaboration within the meetings emerged.

4.2.3 Secondary Display Impact on Collaboration

One of the independent raters noted commonalities amongst of shared displays usage: "*There was an understanding of a shared display as a group resource which someone had to command. If a second display was present, it looked like it could be employed in various ways*." We explore these themes within this section.

***Use of multiple displays for exploration*:** In this style of collaboration, a group member would notice a particular trend on one of the slides on his or her laptop computer and call the groups' attention to this trend by showing the slide on one of the shared displays. Other group members would then consult their own information and determine if they had charts, data, or information pertaining to the original insight and show this information on the adjacent display. One participant articulated this exploration strategy early on in his group's analysis: "*Let's try to piece together a little bit of what our problem is and then we can put on information from different sources for resolution. To start off, we are looking for a reason why profits have steadily decreased.*" This style of exploration appeared in the both side-to-side dual display condition and the opposing dual-display condition.

Another style of exploration used the second display to actively cycle through another group member's slides to search for correlations as a group. This exploration occurred more frequently in the side-by-side dual display condition than the opposing dual-display condition.

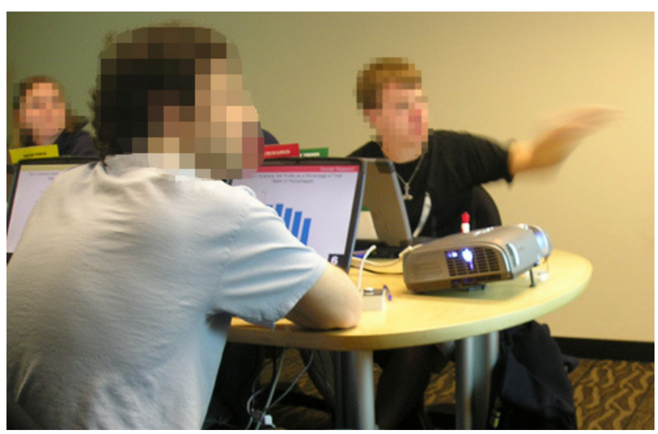

Figure 4. A group member calling attention to the shared display via deictic gesture.

Use of multiple displays to engage: Groups frequently consisted of very different personalities, some people being more outgoing than others. Groups with multiple displays were observed to use the second display to engage a participant who had not been very vocal. For instance, in one group, a participant asked another, "*Is there anything in sales that can help us* [explain this insight]." Another group member asked another, "*You're financial. What do you have regarding healthcare?"* specifically inviting individuals to use the second display.

Using multiple displays to transition group discussion. The multiple display conditions also offered participants an opportunity to segue the dialog to a new topic or insight. On several occasions, if one display was blank or its contents stale (not updated in a while and not the focus of the current dialog), an individual would simply share content on that display, other individuals would take note, and then a new exploration would begin.

Opposing displays used to segregate content. We observed group members using opposing displays to host two different insights. One participant remarked that it was useful to keep different insight themes physically dislocated. In Figure 5, the participant in the foreground and an obscured participant have their attention focused on one display, while other group members have their attention towards the other display.

4.2.4 Collaborative style of meetings

Several trends emerged regarding the collaborative style of the meetings under the different display conditions. One rater noted that there was a sense of high-level discussion occurring within the side-by-side groups with periodic delving into more details.

Discussion within single display groups was described by one rater as often being "ploddy." Individuals would interrupt a current conversation thread to request to see a different chart, an occurrence not as common in the multiple-display conditions. For example, one individual interrupted his own train-of-thought: "*and if you look at sales—can you go back to the same graph?"*

There were also many explicit requests for information to come up, resulting in frequent switches of ownership of the display. On several occasions, individuals attempted to verbally describe some data to be interrupted by "*put it up [onto the shared display*]". Although this occurred within the other display conditions, it appeared much more common under the single-display condition. In the multiple display conditions, group members often seemed to be pro-active in mirroring information onto one shared display without being told to or asking permission to take control.

Several interesting collaboration activities occurred under the single-display condition. In one instance, a group subdivided into different groups on each side of the table and processed information in parallel before collaborating. In another group, individuals broke down into small groups of similar roles to share information before collaborating with the rest of the group. Furthermore, as noted earlier in this section, most of the whiteboard usage occurred within the single-display condition.

4.2.5 Satisfaction

We explored satisfaction with the collaboration process via surveys and interviewing. Participants in the side-by-side condition significantly ranked their satisfaction with the meeting process higher on a Likert scale than those in the single-display condition ($F(2,103)=3.610$, $p=.031$). (Note, one participant failed to answer that survey item).

We also interviewed each group about subjective factors regarding the shared displays. Across all conditions, individuals enjoyed the ability to share information in a quick and lightweight manner. For example, one remarked "*The projector allowed open discussion on the available information"* Another stated, *"It was nice, very fluid. It allowed sharing with zero overhead."* Another individual stated that the shared displays facilitated "*rapidly direct[ing] the group to slides for comparison."* Six negative comments centered primarily upon an occasional lag in the switching device.

Single-display users noted a desire to have a second display to facilitate comparisons amongst the data. When asked of what additional resource would have improved the group decision-making process, groups wanted "side-by-side displays." A few individuals went beyond dual displays, desiring "a *shared desktop so I could 'edit' things on someone else's screen*" or a "*communal whiteboard that could be edited from each person's workstation*." Several participants in groups that did not use the whiteboard also remarked that using the whiteboard would have been useful in organizing thoughts.

The side-by-side display was the only condition that had participants actively remark during the group analysis that they wish they had additional displays. "*We need more than two screens [to explore this]….we need like six."* The desire for more displays was echoed in the post-experiment surveys as a way to further assist groups: "*Two screens is not enough; 4-6 might be better."* One individual went even further, stating that "*only being able to compare two people's slides at a time*" inhibited the generation of ideas within the group. Others remarked that the side-by-side configuration had the benefit of "*show[ing] correlation between two ideas and give supporting evidence towards similar ideas."* However, one participant did note that drawback of the side-by-side configuration was potential ambiguity: "*You were not sure which display someone was talking about."*

The interviews for the opposing display participants illustrated utility in having multiple displays, but there was an overwhelming

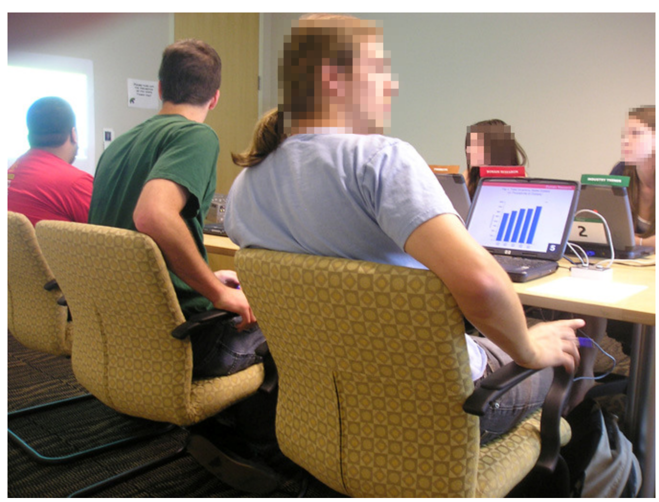

Figure 5. Example of split attention in the opposing display condition.

desire to have the two displays on the same plane, supporting Su & Bailey's guideline that displays should not be orthogonal to each other. One participant remarked that, "*what was distracting was having the displays on opposite sides of the table, making it harder to compare data,"* taking away from the meeting experience. Participants remarked that they did not like looking back-and-forth, however, two participants noted that they were able to determine which display people were talking about simply by seeing which location people were looking at.

5. DISCUSSION

The presence of multiple displays influenced how groups completed the sense-making task in our study, supporting Su & Bailey's findings that different configurations of large displays impact users differently on their tasks. We note, however, that due to the controlled nature of this empirical study, results cannot be easily generalized to other situations and future work needs to further explore these findings—for example, how would performance results change if the groups consisted of 10 individuals? However, in conjunction with other work regarding display placement, such as Wigdor *et al* [28] and Su and Bailey [26], designers can use our results when considering placement of shared displays in collaborative spaces. We offer four themes supported by the findings in this paper when considering supporting an intellective, sensemaking task:

1. Placing multiple shared displays side-by-side offers performance benefits as opposed to placing displays at opposing ends of the table.

With respect to display condition, our results show that groups in the side-by-side shared display configurations were able to make more inferential links between insights than those groups in the opposing shared display configuration. Examining the videotapes and transcripts yielded more exploratory comparisons of information when the dual displays are side-by-side. Interestingly, groups under the opposing-display configuration spent significantly more time displaying slides containing insights than groups in the other two conditions, yet were not able to make these inferential links as often.

Su and Bailey note that displays should be positioned at a 45-degree or lower angle relative to each other, noting that an orthogonal configuration of two displays correlated to lower performance on task times. Similarly, in this paper, the opposing display configuration tended to have lower performance. However, as noted in the previous section, the opposed configuration allowed groups to leverage the physical disparity to focus on different issues. A potential benefit of this arrangement that merits future exploration is to determine whether opposing displays facilitate group members in making more eye contact with each other, potentially impacting the social component of the meeting process.

2. *Side-by-side shared displays offer similar cognitive benefits to using a single shared display with a whiteboard, suggesting analog technologies still merit inclusion in meeting space design.*

Our results showed the ability to share information side-by-side via multiple shared displays is perceived beneficial for groups performing an intellective task compared to opposing display conditions. However, we note that the ubiquitous meeting room configuration of a single shared display and whiteboard offered similar performance characteristics in our study, suggesting that meeting space users will adapt collaborative strategies to the technologies, both analog and digital, provided within a space. This is noteworthy for technology designers as a reminder to not dismiss properties of existing spaces that serve as useful tools.

At first, we were surprised by the lack of a significant difference between the side-by-side and single display groups regarding inferential links. We hypothesized that the side-by-side display configuration would explicitly support comparing and exploring data. Our analysis of the session transcripts showed evidence of comparisons and explorations in multiple shared display groups and also indicated that groups with the side-by-side displays tended to have high-level discourse while groups using the single display were typically more abrupt. Therefore, it seems logical that the smoother collaboration styles combined with easy comparison of slides should have resulted in a performance boost.

However, note in Table 1 that groups under the single display condition used the whiteboard more frequently than under the other two display conditions. These groups used the space to organize thoughts, trends, and themes emerging from their analyses. Interpreting the whiteboard usage through distributed cognition [14], and in particular Salomon's classification scheme [21], we argue that the white board serves as a physical object off-loading the cognitive efforts of the team exploring the data set. This phenomena is shown in Figure 6. For those in the side-by-side display condition, members appear to be performing an act of shared cognition by viewing side-by-side comparisons of information, data, and slides. Groups under the opposing-display condition generally did not use the whiteboard, and due to the disjoint location of displays, were not able to share in cognition as well as their counterparts under the side-by-side condition.

Such a distributed cognition interpretation also provides evidence supporting Streitz *et al*'s point that technology designers need to not only consider the software and technology aspects of meeting spaces, but the physical aspects themselves [25]. Within meeting spaces, physical whiteboards are examples of "informational devices" [25] since they facilitate creating, editing, and displaying information. Our research provides evidence suggesting why existing meeting spaces have not fundamentally changed over time; the basic "tools" provided, both digital and analog, can be appropriated as necessary during collaboration. Furthermore, it is

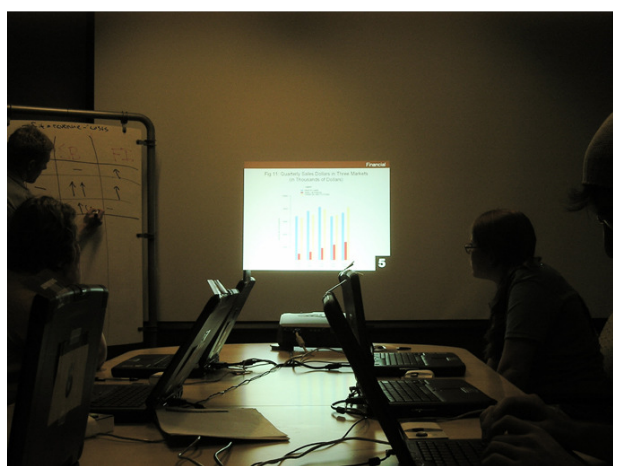

Figure 6. Single display condition group using the whiteboard to record trends.

possible to interpret another finding using a distributed cognition approach. The opposing dual-display group members spent significantly more time showing slides containing insights on the shared display than the other two group conditions support this theory, yet did not demonstrate any substantial performance gains. These groups generally did not use the whiteboard to offload the cognitive efforts of drawing comparisons and correlations amongst the data set. They also did not have the added benefit of shared cognition for side-by-side analysis. Therefore, it seems reasonable to assume that these groups spent more time on these slides continually recalling information and previously discovered insights rather than drawing logical connections between them.

Future work seeks to explore the notion of shared cognition with multiple shared displays in meeting spaces. For this study, we included a whiteboard to replicate this common amenity found within conference rooms. Further studies can remove the whiteboard from the space to examine how shared cognition changes under various multiple display configurations.

3. Multiple-shared displays offer opportunities for individuals to engage team members.

Multiple shared displays impacted the social protocol of the meetings by allowing for new ways for individuals to interject or segue the group discourse. Generally speaking, groups in the side-by-side shared display conditions flowed and interacted more smoothly than groups under the other two conditions. In particular, groups under the single-display condition had an erratic flow where individuals would abruptly talk about their data instead of sharing it visually. It appears that this was also partly due to social factors—did an individual deem their insight was important enough to requisition control of the large display? Many meeting situations have implicit or explicit power relationships that may impact how likely someone is to request "ownership" of a shared display.

Evidence of social norms affecting the flow of collaboration also appeared in the multiple display conditions. When an individual shared content on one display in a multiple display setting, another individual could display content on the second idle/stale display without usurping conversational focus. This also allows shyer group members a lightweight method to introduce their content into the group analysis. Conversely, having multiple displays also prevents one individual from dominating the group conversation. Future work will further explore how multiple shared displays can influence the social dynamics of other meeting situations—for example, can such displays be deployed to engage more team members.

We also note that the presence of multiple displays had a significant impact increasing the number of pointing gestures towards the displays. Therefore, this notable increase in gestures for the multiple display conditions may be interpreted as a method to assist group members in resolving the potential ambiguity of which display's content is in focus. Gesturing is also used to point out comparisons between content shown on both displays.

There was some indication in the opposing display condition that these displays can be used to support parallel discussion, perhaps by physical proximity to each display, even though this did not explicitly occur in the experimentation. However, two participants remarked that one benefit of the opposing display configuration was to use the physical separation to keep analyses between two different insights separate. Interestingly, the only group that showed evidence of breaking down into smaller groups for parallel work occurred in the single display condition.

The opposing display configuration did offer several unique characteristics. The location of these displays gave participants an additional cue to determine which display another group member was looking at (i.e., Figure 5), alleviating some of the ambiguity. However, this did not seem to impact the rate of pointing gestures between the two dual display conditions. This display configuration, however, seemed to promote more eye contact with individuals as they changed attention focus from one side of the room to another. Our camera angles, however, did not allow for us to easily or accurately code for this.

4. An insight-based evaluation offers a useful way to evaluate team collaboration under controlled circumstances.

Finally, we comment on the methodology we used for evaluating displays in meeting spaces. Collaborative settings offer inherent challenges for technology evaluation. Existing research investigating display placement under controlled settings such as [26, 28] used performance tasks (i.e. target-based) in their evaluations. This research offers insight into one specific type of collaborative task, and performance tasks are also fairly easy to implement and evaluate. However, many other complex or information-centric tasks occur during collaboration [17]. Our work contributes a novel way to analyze how multiple individuals come together to make sense of a large amount of information that may be overwhelming at first, a task that frequently occurs in meetings [e.g. 27]. However, as noted in Saraiya *et al,* an insight-based analysis is quite labor intensive and requires domain experts [22], and it is still difficult to control for variance in how well randomly assigned individuals will work together, still representing a challenge to the research community.

6. CONLUSIONS

Our empirical study showed that presence and location of multiple shared displays influenced how groups performed a collaborative analysis on an analytical, sense-making task. Groups using the side-by-side dual shared display condition were able to identify significantly more inferential links amongst insights while problem-solving than those in the opposing dual-display condition. This difference occurred even though groups in the opposing dual-display condition spent significantly more time showing the slides containing the insights on the shared displays. In addition, people in the side-by-side dual display condition indicated a significantly higher rating for satisfaction with the meeting process. We also provide qualitative evidence for previously undocumented ways groups use multiple shared displays to collaborate, engage others, or organize information content, and evidence that the presence of multiple displays changed social protocols in the meeting.

This paper contributes to the HCI community as an empirical investigation of the impact of multiple shared displays within a meeting environment. We also used an insight-based evaluation method where multiple group members performed an information-intensive intellective task and provided evidence that the shared displays themselves play an important role in influencing the collaboration that occurs in such spaces.

7. ACKNOWLEDGMENTS

This work was funded in part by a grant from the National Science Foundation, IIS-0414667. We also extend thanks to

Terence D. West, Joel Stanfield, and Mark Baloga of Steelcase Inc. and to Erika Shehan Poole of Georgia Tech for discussions on collaboration and equipment support. We also thank Blair MacIntyre for loaning equipment for the study, and Don Schoner for his assistance in scheduling facilities. Finally, we thank our two domain experts, Robert Amar and Derek Myers.

8. REFERENCES

1. Biehl, J.T. and Bailey, B.P. Improving interfaces for managing applications in multiple-device environments. In *Proc. of AVI*, ACM Press (2006), 35-42.
2. Biehl, J.T., Baker, W.T., Bailey, B.P., Tan, D.S., Inkpen, K.M, and Czerwinski, M. IMPROMPTU: a new interaction framework for supporting collaboration in multiple display environments and its field evaluation for co-located software development. In *Proc. CHI 08*, ACM Press (2008), 939-948
3. Bui, T., Sivasankaran, T.R. Relation between GDSS use and group task complexity: an experimental study. *IEEE HICSS,* (1990) *69-78*
4. Cassell, J. A framework for gesture generation and interpretation. *Computer Vision in Human-Machine Interaction*, Cambridge University Press (1998), 191-215.
5. Czerwinski, M., Smith, G., Regan, T., Meyers, B., Robertson, G., and Starkweather, G. Toward characterizing the productivity benefits of very large displays. In *Proc. INTERACT 03*, IOS Press (2003), 9-16.
6. Davenport, T.H. Thinking for a living: how to get better performance and results from knowledge workers. HBS Press Book, 2005.
7. DeSanctis, G. and Gallupe, R.B. A foundation for the study of group decision support systems. *Management Science*, 33(5), 1987, 589-609.
8. Fjermestad, J., and Hiltz, S.R. *An assessment of group support systems experimental research: methodology and results.* Journal of Management Information Systems, 15(3), 1998, 7-149.
9. Fussell, S., Setlock, L.D., Yang, J., Ou, J., Mauer, E., & Kramer, A.D. Gestures over video streams to support remote collaboration on physical tasks. *Human-Computer Interaction*, 19, 2004, 273-309.
10. Gallupe, R.B., DeSanctis, G., Dickson, G.W. Computer-based support for group problem-finding: an experimental investigation. *MIS Quarterly*, 12(2), 1988, 276-297.
11. Grudin, J. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. In *Proc. CHI 2001*, ACM Press (2001), 458-465.
12. Guimbretiere, F., Stone, M., & Winograd, T. Fluid interaction with high-resolution wall-sized displays. In *Proc. of UIST '01*, Publisher (2001), 21-30.
13. Hutchings, D.R., Smith, G., Meyers, B., Czerwinski, M., and Robertson, G. Display space usage and window management operation comparisons between single monitor and multiple monitor users. In *Proc. AVI*, ACM Press (2004), 32-39.
14. Hutchins, E. *Cognition in the wild*. Cambridge, MA: MIT Press (1995).
15. Jarvenpaa, S.L. and Dickson, G.W. Bonanza Business Forms Company: a mystery in declining profits. In *Marketing Decision Making*, Ruppel, A.C, O'Dell, W., Trent, R., Kehoe, W (eds), 1988.
16. Johanson, B., Fox, A., Winograd, T. The Interactive Workspaces Project: experiences with ubiquitous computing rooms. *IEEE Pervasive Computing,* 1(2): 67-74.
17. McGrath, J.E. *Groups: interaction and performance*. Prentice-Hall, 1984.
18. Newman, M.W., Ducheneaut, N., Edwards, W.k., Sedivy, J.Z., Smith, T.F. Supporting the unremarkable:experiences with the obje Display Mirror. *Personal Ubiquitious Computing* (2006), Springer-Verlag, 523-536
19. Owens, D.A. Structure and status in design teams: implications for design management. *Design Management Journal,* 1(1): 55-94.
20. Poole, M.S., Siebold, D.R., and McPhee, R.D. Group decision-making as a structurational process. *Quarterly Journal of Speech*, 71 (1985), 74-102.
21. Salomon, G. *Distributed cognitions: Psychological and educational considerations*. New York: Cabridge University Press (1993).
22. Saraiya, P., North, C., and Duca, K. An insight-based methodology for evaluating bioinformatics visualization. *IEEE Visualization and Computer Graphics*, 11(4), 2005, 443-456.
23. Streitz, N.A., Giessler, J., Holmer, T., Konomi, S., Culler-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., and Steinmetz, R. I-Land: an interactive landscape for creativity and innovation. In *Proc. CHI '99*, ACM Press (1999), 120-127.
24. Streitz, N., Prante, T., Muller-Tomfelde, C., Tandler, P., and Magerkurth, C. Roomware: the second generation. In *Proc CHI Extended Abstracts*. ACM Press (2002), 506-507.
25. Streitz, N.A., Rexroth, P., Holmer, T. Does Roomware matter? Investigating the role of personal and public information devices and their combination in meeting room collaboration. In *Proc. ECSCW*, Kluwer Academic Publishers (1997), 297-212.
26. Su, R.E., and Bailey, B.P. Put them where? Towards guidelines for positioning large displays in interactive workspaces. In *Proc. of Interact* '05, Springer-Verlag (2005), 337-349.
27. Thomas, J.J., and Cook, K.A. *Illuminating the path: the research and development agenda for visual analytics.* National Visualization and Analytics Center, 2005.
28. Wigdor, D., Shen, C., Forlines, C., Balakrishnan, R. Effects of display position and control space orientation on user preference and performance. In *Proc. CHI '06*, ACM Press (2006), 309-318.