

Investigating the Role of a Large, Shared Display in Multi-Display Environments

RUNNING HEAD:

Shared Display Use in MDEs

AUTHORS:

James R. Wallace, Stacey D. Scott, Eugene Lai, Deon Jajalla

AFFILIATION:

University of Waterloo, Waterloo, Ontario, Canada

CORRESPONDENCE:

James R. Wallace
Systems Design Engineering, DC 2583
University of Waterloo
Waterloo ON Canada
N2L 3G1

jrwallac@uwaterloo.ca

519 888-4567 x 36813

ABSTRACT

We conducted an empirical study to investigate the use of personal and shared displays during group work. The collaborative environments under study consisted of personal workspaces, in the form of laptops, and a shared virtual workspace displayed on a nearby wall. Our study compared the use of the large shared display under two different interface content conditions; a status display that provided an overview of the group's current task performance, and a replicated view of the shared workspace that allowed task work to occur on the shared display. The study results suggest that while participants used their personal displays primarily to perform the task, the shared display facilitated several key teamwork mechanisms. In particular, the provided status display best

facilitated monitoring of group progress, whereas the replicated content display best facilitated conversational grounding. Regardless of the shared display content, having a shared, physical reference point also appeared to support synchronization of the group activity via body language and gaze.

Keywords

multi-display environments, evaluation, design, display configuration, input redirection, personalized views, content replication, job shop scheduling task

1. INTRODUCTION

Many Computer-Supported Cooperative Work (CSCW) researchers have explored computing systems that take advantage of the collaborative benefits that a physically shared display can provide in the form of Single Display Groupware (SDG) systems (e.g. Elwart-Keys et al., 1990; Scott et al., 2003; Shoemaker & Inkpen, 2001; Stewart, 1999). However, the increasing availability of mobile, personal devices and wireless networking infrastructure is rapidly expanding the possible design solutions that can be developed to support co-located collaboration, including Multi-Display Environments (MDEs) that enable group interaction through combinations of multiple, heterogeneous personal and shared devices connected by an underlying groupware architecture (e.g. Biehl & Bailey, 2004; Johanson et al., 2002). As these MDE systems continue to be designed, developed and evaluated in the CSCW literature, the list of potential uses for additional displays continues to grow. For example, personal displays can be used to provide scratch space, to access personal data, or to access a personally tailored view of a shared workspace. Alternatively, shared displays may provide an

overview of the group activity, specific views on the group's tasks, or a fully interactive group workspace.

Unfortunately, the literature offers few guidelines that help system developers decide how to incorporate heterogeneous personal and shared devices in order to create an effective collaborative work environment. To address this problem, this research aims to explore the roles that personal and shared displays each play during collaborative work and to discover how these displays can best be leveraged to enable individuals to perform the required *taskwork* (Pinelle et al., 2003), and the group to perform the required *teamwork* (Pinelle et al., 2003), to effectively accomplish a collaborative activity. Taskwork is defined as the work required to accomplish the task, whereas teamwork is defined as the work required to coordinate the group while doing this work.

The study presented in this paper is the second of an ongoing research series designed to explore group process in multi-display environments. In the first study, we investigated the impact of different display configurations on taskwork and teamwork during a collaborative optimization task (the job shop scheduling task described below) (Wallace et al., 2009). This study revealed that providing groups with only a single, shared display promoted group awareness, whereas providing each group member with a personal display with customized views, in addition to a shared display, promoted task accuracy as group members could easily focus on their individual aspects of the task. In addition, the study revealed that in the configuration in which both personal and shared displays were available (i.e. the MDE condition), participants rarely used the shared display, even though the display provided such additional information as the other group

members' mouse cursors (similar to the 'replicated content' display used in the current study).

Our informal observations, however, suggested that when participants did use the shared display, it appeared to play a different role than their personal displays. This display seemed to enable them to mentally "step back" from their individual task interactions to obtain an overview of the group's progress. This apparent difference in the purposes served by the shared and personal displays called into question the utility of replicating the task workspace content on the shared display, rather than using that display to provide other possible content views. If users were mentally "stepping back" when viewing the shared display, for example, perhaps displaying an overview of the task would be more appropriate. To explore this issue further we conducted a follow-up study detailed in this paper, explicitly investigating the impact of providing different types of task-related content on a shared display.

The results of this follow-up study, presented in detail in this paper, indicate that both shared display content configurations provide advantages to the group. Groups provided with a replicated view of the task interface that appeared on their personal display were better able to ground conversation, whereas groups provided with a status display were better able to monitor one another's progress. The results also revealed, however, that regardless of the specific type of task content shown on the shared display, its physical presence facilitated the use of body language and gaze in synchronizing the groups' activity.

Before describing the study in detail we first overview the related work. Then, we describe the experimental task used in our study, the Job Shop Scheduling (JSS) task (Tan et al., 2005; Tan et al., 2008), and the two shared display interfaces used in the study. Next, we present the results of the study, and interpret these results, along with the larger picture provided by both these results and those from our earlier study, within the theoretical context of the CSCW literature. Finally, we discuss the implications of these results for the design of computer-supported collaborative environments and reflect on the value of shared displays for supporting collaboration.

2. ROLE OF DISPLAYS IN COLLABORATIVE ENVIRONMENTS

Moore's law (2005; 1965) has created a world in which laptops, personal digital assistants (PDAs), phones, tablets and other personal devices are significantly more powerful than desktop computers were 10 years ago. These devices are also incredibly mobile and wirelessly connected to the Internet, local area networks, and other nearby devices. As these personal devices have become capable of supporting increasingly complex interactions, their role has moved from that of a stand-alone device to an integral part of the collaborative environment (e.g. Greenberg et al., 1999; Johanson et al., 2002; Myers, 2000; Nacenta et al., 2005; Shen et al., 2003; Sugimoto et al., 2004).

As collaborative environments have continued to evolve, research has focused on the development of new types of interaction and leveraging the combination of personal and shared workspaces to improve on a group's efficiency, outcomes and satisfaction while collaborating. Research into Group Decision Support Systems (GDSS) has explored methods to improve a group's efficiency at multiple levels (DeSanctis & Gallupe, 1987; Fjermestad & Hiltz, 1997; Nunamaker et al., 1996),

such as by shaping the way in which groups perform tasks and processes, much of this research, however, was conducted prior to the mass adoption of mobile devices. Research in the field of CSCW has explored methods by which users can share content between personal and shared workspaces (e.g. Biehl & Bailey, 2004; Collins et al., 2007; Lanir et al., 2008; Wigdor et al., 2009) and interact with shared content via their personal devices (e.g. Berry et al., 2005; Ishii & Kobayashi, 1992; Sugimoto et al., 2004). These developments have lead to the adoption of multi-device and multi-display groupware in a variety of collaborative environments from the demanding and complex setting of a commercial airline cockpit (McKay, 2009) to the more casual setting of media sharing in the living room (Seifried et al., 2009).

As the opportunities for multi-display interaction increase, an important design consideration for developers of collaborative systems is to understand the relative strengths and weaknesses of different display types. In particular, it is important to understand what role certain display types can play in supporting collaboration. In this research, we focus on understanding the role that a *shared display* can play in a collaborative environment.

2.1 Shared Displays

Large, shared displays are often used to provide a shared workspace in the form of an interactive digital table (e.g. Morris et al., 2006; Ryall et al., 2004; Sugimoto et al., 2004), or a nearby wall display (e.g. Biehl & Bailey, 2004; Hailpern et al., 2007; Johanson et al., 2002). These displays are often used to support group interactions; thus, they are uniquely positioned to support group process, or teamwork (Pinelle et al., 2003). This investigation is focused on two types of

shared displays, each with a specific intended purpose in aiding collaboration; *status displays* and *replicated content displays*.

Status displays tend to consist of non-interactive data, and are used to help monitor group activity. This functionality has been referred to as “at-a-glance awareness” (Plaue et al., 2009), and may support awareness of projects, issues or group progress in the form of status update information (Carroll et al., 2003). For example, large status displays are often seen in war room configurations where users are assigned specialized subtasks, and provide a mechanism by which users can monitor the progress of the group. In research, projects such as Notification Collage (Greenberg & Rounding, 2001), group participation displays by DiMicco et al. (2004), and FASTDash (Biehl et al., 2007) have deployed status displays in office settings to successfully support awareness of presence, participation, and activity with shared task resources. Projects such as MERBoard (Huang et al., 2006) have explored providing status displays in the support of more specialized groups, such as NASA’s space operations.

On the other hand, replicated content displays tend to support synchronous, tightly-coupled communication and coordination. Research has shown that such workspaces support collaboration by improving the efficiency with which groups collaborate (Gergle, 2006). For example, shared displays enable non-verbal communication such as gestures (Baker et al., 2002; Gutwin et al., 1996), and provide a shared visual reference that facilitates communication grounding (Clark & Brennan, 1991). Projects such as Pebbles (Myers, 2000), UbiTable (Shen et al., 2003), and Caretta (Sugimoto et al., 2004) have previously explored the use of

shared workspaces to provide shared access to task resources, to facilitate sharing of personal artifacts, and as a space to share personal task work, respectively.

In a previous investigation of collaborative environment design (Wallace et al., 2009), we directly compared groups working in single- and multi-display collaborative environments in a laboratory setting. Our results indicated that multi-display collaborative environments put groups at a disadvantage in terms of awareness, yet provided a “sheltered” workspace in which individuals could focus on more cognitively demanding aspects of the task. We also noticed that in multi-display configurations, individuals spent most of their time performing work on their personal display, and that the large, shared display appeared to be used for “stepping back” and obtaining a different perspective on the task. As a follow-up to that study, we wanted to more fully explore the interface design choices that led to this difference and to clarify the role that shared displays play in supporting collaboration. We now discuss how shared displays have been utilized in multi-display collaborative environments to date, and outline research that suggests how these displays should be arranged to optimally support a group.

2.3 Laboratory Studies of Display Configuration

The status display in NASA’s space operations centre and the replicated controls in a commercial airline cockpit have proven their utility through decades of in-context use, feedback, and iteration from their respective designers and stakeholders. Yet, the lessons learned in deploying these systems have not led to the development of generic guidelines for the design of collaborative environments. Laboratory studies of collaboration provide an opportunity to investigate and directly compare design choices that have been validated,

deployed and utilized in disparate field settings, and to observe generalizable differences between the two display types. In the case of shared displays for supporting group work, laboratory studies that compare alternative display configurations provide an opportunity to more fully explore how these displays support the collaborative process, and to derive guidelines for the development of collaborative environments for supporting new tasks.

Previous empirical laboratory studies explored how the presence and placement of displays can impact collaborative processes, effectiveness, efficiency, and satisfaction (e.g. Forlines et al., 2006; Hawkey et al., 2005; Su & Bailey, 2005). For example, in a comparison of alternative display configurations for supporting a decision-making task, Plaue & Stasko (2009) compared traditional whiteboard use, the use of a single shared display, and the use of two large shared displays within a single study. Comparing these configurations within a controlled setting allowed Plaue & Stasko to explore the impact of display placement on performance, to compare the cognitive benefits of analogue techniques such as whiteboards to their digital counterparts, and to investigate how the inclusion of multiple displays changed the group's social protocols and provided new ways for group members to interact with one another. These direct comparisons would not be as meaningful had the display configurations been studied in separate environments where environmental, social, cultural, and group and individual differences, introduce too many confounding variables to make comparisons of personal and group interactions with technology possible.

While previous studies focused on exploring the impact of the presence and location of large, shared displays, they have not yet addressed how to optimally

design the display's content. Fjermestad & Hiltz (1997), in an overview of GDSS research, recognized this shortcoming and noted that "most experiments seem to (falsely) assume that all GDSS's are a standard 'package' that will have the same effect" (p. 4). Our primary motivation for conducting this work was to leverage the controlled environment of a laboratory study to explore the impact of alternative shared display types on a group's collaborative processes and efficiency. As this aspect of collaborative environment design is currently unexplored, we conducted a mixed-methods study that gathered both quantitative data that could be directly compared to our previous study, and also qualitative data to enable an investigation into how alternative shared display configurations support teamwork and taskwork.

3. STUDY METHOD

A repeated measures study design was used to provide an opportunity to train participants in the use of both status and replicated content displays, to enable repetition between groups, and to help account for variations introduced by individual and group differences. By observing groups as they worked with both status and replicated content displays, we hoped to reduce the known impact of individual and group differences on verbal and deictic communication, and in interactions with the experimental apparatus (e.g. Greenberg, 1991). The controlled laboratory setting of the study also enabled for replication and repetition between groups, which improved our ability to capture a broad spectrum of group activity under the same conditions. Finally, the inclusion of qualitative measures provided an opportunity to explore the rich interactions between group members and technology that occurred during each trial and to investigate how the different experimental conditions impacted teamwork and

taskwork processes. The following sections detail the experimental task, study design, participants, setting, and procedure used in the study.

3.1 Experimental Task

The Job Shop Scheduling (JSS) task (Tan et al., 2008) emulates authentic optimization tasks such as the scheduling of manufacturing apparatus on a plant floor, and has been classified as an intellective task (McGrath, 1984). Tan et al. (2008) note that the task is particularly useful for the study of group work in laboratory settings because when assigned to multiple participants it elicits information sharing behavior and requires coordination between group members. Since information sharing and coordination are two aspects of teamwork that are supported by shared displays, the JSS task is useful in contrasting the differences between technologies used to support highly synchronous, co-located activities. For example, in their study of alternative input and display configurations, Tan et al. found that they were able to observe how users “adapt their communication to the available collaboration tools” (p. 12), a property of the task that we hoped would aid our comparison of shared display configurations. Finally, The JSS task has objective measures of performance, such as solution time, solution quality, and solution efficiency, simplifying comparisons of task outcomes between display configurations and groups.

To complete the task, participants optimize the scheduling of six “jobs”, each composed of six ordered operations (Figure 1, A). These operations are dependent on six resources (Figure 1, B) that can only be used by one operation at a time. A solution is considered valid if no two operations are simultaneously utilizing a shared resource (an *overlap error*), and if no two operations within the same job are scheduled to occur at the same time (an *order error*). Once a candidate solution is found, each group member must agree on a final solution using the “Submit” button on their personal display (Figure 1, C).

A solution “scrubber” was also made available as a component of the personal interface that enabled users to view and load previous task states (Figure 1, D). The scrubber was included to better facilitate exploration within the problem space. In particular, the scrubber was potentially helpful as solutions to optimization tasks are susceptible to local minima and maxima within the problem space, and the scrubber can be used to backtrack to previous solutions, or to




Figure 1. The Job Shop Scheduling Task interface. Each job is indicated by blocks of a particular colour, with six operations (A) per job. Each job has one operation that must be completed by each of the six resources (B). Group member job assignments are indicated on the right side of the JSS interface (F). Participant submit label indicators and the submit button (C) are also provided. In this iteration of the study, a solution scrubber (D) was also incorporated into the interface to facilitate backtracking to a previous solution. The excerpt on the right (E) shows a close-up view of a personal display, where unassigned job components are displayed in a less salient fashion to provide a personal workspace for the task.

restart the puzzle entirely. A similar backtracking tool was found to be particularly beneficial in supporting groups conducting a city planning task in the Caretta project (Sugimoto et al., 2004).

Solutions can be compared between trials and groups using quantitative measures of the task's outcome: *solution quality* and *errors*. Solution Quality is defined as the degree to which a solution is optimal, and is measured as the difference between a solution's completion time as measured by the total scheduled time to complete all jobs and that of an optimal schedule. Errors are defined as the total number of overlap and order errors present in a submitted solution. Similarly, quantitative measures of *job component moves*, *conflicts*, and *utterances* provide a means to compare the taskwork performed by the group, and are measured at both the group and individual levels. Job component moves are defined as the number of times a participant clicks and drags a job component to a new position in the solution space. Conflicts are defined as the number of times that two participants simultaneously click on the same job component. Utterances are defined as the number of times a participant spoke during the trial.

3.2 Design

A 2 (shared display type) x 2 (task structure) design was used, with shared display type as a within-subjects factor, and task structure as a between-subjects factor. Thus, each group completed 2 display configuration trials, in one of two task structure conditions. The two shared display type configurations used in the study included *status display* and *replicated content display* (Figure 2). In all conditions, each personal display acted as a personal workspace that showed only its owner's mouse cursor. Operation components on the personal displays were visually differentiated to increase the salience of resources assigned to each group




Figure 2. Interface differences between display type configurations: status display configuration(left) shows task interface on laptop displays, and an overall status display on the shared display; replicated content configuration (right) shows a replicated copy of the task interface on the laptop displays as well as on the shared display.

member. In contrast, unassigned operations and operations assigned to others were visually de-saturated (Figure 1, E).

In the replicated content display configuration, the shared display also contained a view of the JSS interface that was shown on the personal displays, with two key differences. First, all three participants' cursors were visible on the shared display. Second, job components were visually differentiated (i.e. semi-transparent) only if no participant maintained control over them. The shared and personal display interfaces were carefully aligned to ensure that mouse coordinates were mapped identically between both displays; this mapping was used to facilitate virtual deixis across displays (e.g. pointing or gesturing with the mouse cursor).

The status display condition also displayed all participants' mouse cursors, but provided an alternate content view on the shared display (Figure 3). In this view, participants were able to see task status graphs corresponding to efficiency measures (Figure 3, A), a clock indicating remaining trial time (Figure 3, B), and

an error display which indicated any job components which overlapped (Figure 3, C).

Previous work indicates that group members working on personal displays have reduced awareness of their team member's actions and intentions (Gutwin & Greenberg, 1998; Hawkey et al., 2005). Such reduced awareness may reduce the ability to coordinate the use of shared resources, such as the job operations in the JSS task, thus limiting the number of shared task resources that require group members to coordinate may provide advantages in content replication-based environments that provide limited awareness. In order to investigate this issue, we included two levels of task structure: *shared access* (SA) and *negotiated access* (NA).

In the SA condition, the JSS interface allowed any group member to access any job operation throughout the task session; thus, groups had to coordinate their




Figure 3. The status display interface. Participants using this interface were presented with (A) graphs representing the optimality of their current solution, (B) a clock and (C) an error display, in which job components which were overlapping with other pieces were visually less salient.


interactions with the available job components to avoid conflicts in which multiple group members attempt to move the same piece at the same time. In the NA condition, the JSS interface allowed participants to negotiate job assignments via checkboxes in the task interface (Figure 1, F). Once assigned, participants maintained ownership of job components (i.e. other members were “locked out” from manipulating that job component) until the “owner” de-selected the corresponding checkbox. For example, when an individual takes ownership of the “red” job, they become responsible for positioning all red components in the overall task schedule until they relinquish ownership.

3.3 Participants

Thirty six participants (20 male, 16 female), aged 18 to 27 (average age 21), were recruited from the University of Waterloo community and were organized into twelve groups of three for performing the experiment. Groups were randomly assigned to the between-groups factor, task structure, with an equal amount of groups completing each of the SA and NA task structure conditions. In the SA condition, three groups consisted of participants who knew each other and volunteered together, two groups consisted of pairs who volunteered together matched with individual volunteers, and one group consisted of three randomly-matched individual volunteers. In the NA condition, three groups consisted of participants who knew each other and volunteered together, one group consisted of a pair who volunteered together matched with an individual volunteer, and two groups consisted of three randomly-matched individual volunteers.

All participants had normal or corrected-to-normal vision, and were tested for colour-blindness prior to beginning the study. All participants indicated that they were familiar with laptop use, and indicated that they used a laptop on a weekly or

The study was conducted in an open lab space with participants seated around the three sides of a 2m x 1m table. A shared display was projected on a wall approximately 1.7m away from the table at a resolution of 1024x768 pixels over a 2m x 1.5m area, and each participant was seated such that they were facing or adjacent to the public display, and seating positions were kept constant across all trials (Figure 4).



17

Three Lenovo T61 Thinkpad laptop computers (2x2GHz, 1GB RAM) provided participants with input to the public display via a dedicated 802.11g wireless network and the Swordfish groupware architecture (Wallace et al., 2006). Each laptop had a mouse attached for input, and participants' cursors were displayed on both their personal and shared workspaces.

3.5 Procedure

Participants first received a brief introduction to the study from the experimenter, and then completed an informed consent form, colour-blindness test, and a demographic questionnaire that included questions concerning laptop and large display use. Next, the experimenter presented a 10 minute PowerPoint tutorial describing the JSS task, the task goals, and error cases. Participants then completed one 10-minute practice trial in each display condition to reduce anticipated learning effects (Tan et al., 2008), and to ensure that participants were familiar with all of the interface features before proceeding to the experimental trials.

Next, participants completed two 20 minute experimental trials, one for each of the two display conditions. The order of presentation of the shared display content conditions was counter-balanced across groups. After each trial, participants completed a questionnaire eliciting their opinions on the shared display, and their experienced workload via a NASA-TLX (Hart & Stateland, 1988). Once all trials were completed, each group participated in a semi-structured interview which elicited responses regarding difficulties with the task, features missing from the shared display interfaces, and preference data. Finally, participants were thanked for their time and paid for their participation.

3.6 Data Collection and Analysis

Participants' interactions with the JSS interface were automatically captured into computer logfiles. Their conversations and interactions in the physical workspace were recorded on videotapes. Participants' subjective responses were collected via the post-condition questionnaires and a post-experiment semi-structured interview.

Two-way repeated measures analysis of variance (RM-ANOVA) tests were conducted to investigate possible statistical differences in task performance, as measured by the solution time (faster being better), number of errors (fewer errors being better), and task efficiency (fewer number of job component moves being better). An alpha of 0.05 was used for these parametric statistical tests, and where sphericity assumptions were violated, the Huynh-Feldt method was used for corrections. Tukey tests using the Bonferroni adjustment were used for pairwise post-hoc comparisons. The Likert-scale ratings collected from the questionnaires were analyzed using Wilcoxon paired signed-ranks tests.

The video and interview data were reviewed to identify behavioural or conversational patterns and notable participant comments. The video data were transcribed and a basic conversational analysis was performed in which the number of utterances was compared across conditions, using a two-way RM-ANOVA, to determine whether the independent variables had any impact on communication efficiency.

4. RESULTS

The quantitative data analysis revealed that the shared display type and the task structure factors had minimal impact on the taskwork and teamwork measures included in this study. The results do, however, reveal interesting differences between data collected from this study and our previous study that compared groups performing the JSS task in both a single-display groupware (SDG) setting and a multi-display environment (MDE) (Wallace et al., 2009). Surprisingly though, the analysis of participants' questionnaires revealed that participants' perceived value of the shared display type differed across conditions. The results from both quantitative analyses are detailed below. We then report the results of our in-depth qualitative analysis in order to probe further into the possible impact of the shared display type condition on the groups' taskwork and teamwork processes.

4.1 Task Performance Data

To understand the impact of the shared display content on taskwork, we examined a number of task performance metrics across replicated content display and status display conditions, including number of errors committed, solution quality, conflicts, task time, and number of job components moved. No significant differences were found across these measures. Moreover, these data were similar to the same measures collected for the MDE condition in our previous study (Wallace et al., 2009). Table 1 summarizes these results, along with the results discussed below and the comparison data from our previous study.

On average, however, groups in this study took 14% longer to complete the task, and produced schedules that were 40% shorter (i.e. more optimal) than groups in the MDE condition of our previous study. Though these improvements are well

within the large between-group variation in task times and solution optimality observed in both studies, we hypothesize that the trend of increased performance may be caused by the addition of the “scrubber” feature in this second study (it was not available in the previous study). The ability to rollback solutions encouraged more exploration of the solution space, possibly leading to more time spent performing the task. As groups more fully explored the solution space, they were more likely to come across more optimal solutions to the JSS task. The interaction logs show that five groups took advantage of this capability, and that those groups loaded a previous solution state an average of three times per trial (SD 1.95).

We also examined the impact of the shared display content on teamwork. In particular, we examined the amount of communication that occurred across conditions. No significant differences were found in the number of utterances groups made across shared display types. Similar to the task time results, though, there was an increase in the average number of utterances in both shared display type conditions compared to the MDE condition in our previous study (64% increase in the status display condition and 78% increase in the replicated content condition). Again, we hypothesize that the increase in group communication between the two studies may result from the introduction of the scrubber feature, and the resulting tendency of groups to explore the solution space. This issue is explored further in Section 5.4 below. Overall, the content of the shared display had little effect on task performance or the amount of group communication in this study.

Table 1. Mean values and standard deviations (in parentheses) for task efficiency and solution quality measures, and ANOVA results for comparisons between experimental conditions.

Study	Measures									
	Utterances		Job Component Moves		Task Time (seconds)		Errors per Trial		Solution Optimality	
	NA	SA	NA	SA	NA	SA	NA	SA	NA	SA
Study 1 (MDE Condition Only)	166.17 (70.06)	120 (70.97)	234.83 (104.52)	399.83 (228.19)	811.17 (326.70)	807.83 (291.62)	0.00 (0.00)	0.167 (0.408)	25.00 (21.213)	25.00 (35.355)
Study 2, Status Display	293.50 (191.01)	176.00 (123.70)	240.67 (95.43)	351.00 (109.64)	867 (337.22)	935.83 (257.44)	0.00 (0.00)	0.167 (0.408)	23.33 (22.51)	14.00 (19.494)
Study 2, Replicated Content Display	302.50 (131.56)	207.67 (127.78)	274.83 (71.86)	361.17 (149.86)	903.33 (190.97)	990 (264.89)	0.167 (0.408)	0.167 (0.408)	12.00 (16.432)	10.00 (10.00)
Study 2, Comparison Between Display Conditions	F(1,10)=1.037, P=.332		F(1,10)=.856, P=.377		F(1,10)=.422, P=.531		F(1,10)=2.000, P=.188		F(1,7)=.718, P=.425	
Study 2, Comparison Between Task Structure Conditions	F(1,10)=1.555, P=.241		F(1,10)=2.248, P=.165		F(1,10)=.259, P=.622		F(1,10)=2.000, P=.188		F(1,7)=.021, P=.890	
Study 2 Total (Average over entire study)	298 (156.44)	191.833 (126.75)	257.75 (86.18)	356.08 (131.40)	885.17 (285.58)	962.92 (271.34)	0.083 (0.289)	0.167 (0.389)	18.182 (19.909)	12.00 (14.757)

To understand the impact of task structure on taskwork, we tested the same task performance metrics discussed above across NA and SA task structure conditions. Similarly, no significant differences were found for errors, solution quality, and job component moves between task structure conditions, and conflicts were only possible in the SA condition, so no comparison was made. The data for these metrics were consistent with those found in the MDE condition of our previous study. Finally, we compared the number of utterances across display conditions to understand the impact of task structure on teamwork. No significant differences were found. Therefore, overall task structure did not appear to impact teamwork or taskwork in this study.

4.2 Preference and Subjective Workload Data

Though no significant task performance differences were found across shared display type conditions, the analysis of the questionnaire data revealed that participants perceived the status display condition to be more helpful than the replicated content display, as evidenced by participants agreeing more strongly with the statement, “The shared display helped us solve the puzzle” in the status display condition than in the replicated content display condition ($W=61$, $p=0.0178$). However, no significant differences were found for other preference measures such as “I felt our group worked well together” ($W=19$, $p=0.3472$) or “I felt that it took a lot effort to solve the puzzle” ($W=15$, $p=0.522$). Similarly, subjective workload, as assessed by the NASA-TLX (Hart & Stateland, 1988), was not significantly different across shared display type or task structure conditions for any of the six assessed dimensions.

4.3 Shared Display Use

Throughout the JSS task, all groups would alternate between “strategy” and “activity” phases of work. In the strategy phases, groups would actively discuss job component moves, whether to load a previous solution, or overall strategy in performing the JSS task. After deciding on a course of action, participants moved to an activity phase in which they would focus primarily on their individual laptop displays to complete their task moves. The strategy phases were relatively short, and overall, participants spent most of their time working on their personal laptop displays during the study trials.

An analysis of the coded video data revealed that groups looked at the shared display more frequently in the status display condition than the replicated content display condition ($F(1,10)=14.395$, $P=.004$). However no differences were between structure conditions ($F(1,10)=.481$, $P=.504$). The increased display use in the status display condition suggested that there may be underlying behavioural differences between display conditions.

To investigate these differences, a post-hoc analysis of the collected data was performed in which groups were sorted into three categories according to their observed behaviour using investigator’s field notes, activity graphs created from logged computer data, and transcribed video: *optimizers*, *satisficers*, and *other*. Groups classified as optimizers repeatedly loaded previous solution states, and used nearly the entire allocated time to perform the task. Satisficers rarely loaded previous solution states, and submitted their solutions in approximately half of the allotted trial time. Groups classified as other exhibited both types of behaviour throughout their two sessions, or were judged not to be in one of the other two

categories. In total, three groups were classified as optimizers, four groups were classified as satisficers, and five groups were classified as other. In conducting this analysis of group behaviour, two optimizer and two satisficer groups were analyzed in detail – the observational results presented here are based on this analysis of the four selected groups.

Monitoring

While working on their personal displays, many users continuously monitored the shared display; participants would look up from their personal displays, glance at the shared display, and then return to working on their personal display. Groups looked at the replicated content display once per minute on average, whereas groups looked at the status displays approximately three times per minute on average. We did not specifically ask about this behaviour during the post-condition questionnaire or post-study interview; however it appeared from the video that users were briefly consulting the portion of the display which indicated if there were errors in the solution. By glancing up at the status display's error indicator, participants were able to maintain awareness of the overall solution state without significant effort. Interestingly though, the presence of this behaviour did not impact task performance.

In the replicated content display condition, participants did not exhibit monitoring behavior as frequently; however the replicated content display was used as a “safety net” when unexpected events occurred. For example, in the case of Group 7's replicated content display trial, Participant 1 loaded a previous solution state without realizing that it would do so for the entire group (Figure 5). When Participants 2 and 3 realized that their solution state had changed, they

immediately referred to the replicated content display to identify the source of confusion. The transcript follows:



Figure 5. Participants 2 and 3 refer to the shared display after Participant 1 unexpectedly loads a previous solution state.

(4:13) G7P3: yea, now I move ... [p1 loads solution] ... oh wait, what, what?

(4:15) [p1 checks shared screen]

(4:16) [p3 checks shared screen]

(4:16) [p2 checks shared screen]

(4:17) G7P2: hey why'd you do that

(4:18) G7P1: oh crap it happens for everybody?

(4:20) G7P3: yea man

(4:21) G7P1: I didn't know that

(4:22) G7P2: of course, we're ... [shakes head]

(4:23) G7P1: don't worry we'll go back

(4:24) G7P1: we'll go back in time

[p1 loads previous solution state, and group resumes work]

Communication Grounding

Verbal references to the shared display tended to embody the puzzle or JSS task as a whole or the group state rather than fine-grained references to individual job components. However, in some cases, participants were more active in using the

shared display to explicitly communicate fine-grained task details. Grounding typically occurred when participants were engaged in the “activity” phases of the task, when there was difficulty in gaining the attention of fellow collaborators who were actively engaged with their personal display. One example of grounding occurred when Group 5 was considering the submission of a solution in their status display trial. Participant 1 suggested that they submit the solution. Before submitting, Participant 3 identified an error in the solution, and had to get Participant 2’s attention on the shared display by tapping him on the shoulder repeatedly until he was acknowledged (Figure 6).



Figure 6. Participant 3 uses the status display to explain a problem with the current solution state to Participant 2.

- (30:13) **G5P1:** okay?
- (30:13) **G5P3:** good
- (30:15) **G5P3:** eh?
- (30:15) *[p2 checks shared screen]*
- (30:19) **G5P1:** sure?
- (30:20) **G5P3:** yes
- (30:21) **G5P3:**oouya
- (30:22) **G5P1:** *[Laughs]*

(30:22) *[p2 checks shared screen]*
(30:31) *[p3 checks shared screen]*
(30:35) **G5P3:**wha *[repeatedly taps p2 until p2 looks up] [points at shared screen]*
(30:37) *[p2 checks shared screen]*
(30:38) *[p1 checks shared screen]*
(30:38) **G5P2:** what?
(30:40) **G5P3:** green 'C' 'D'
(30:46) *[p3 checks shared screen]*
[p2 and p3 go back to working on their personal displays]

Instances of grounding were not limited to the shared display; some participants also used their collaborators' personal displays as tools for grounding. For example in Group 7's replicated content display trial, Participant 2 decided to get Participant 1's attention by pointing directly on her (P1s) personal display (Figure 7). By pointing directly on a personal display, users could bypass the "getting attention" phase.




Figure 7. In Group 7's replicated content display trial, Participant 2 explains a series of job component moves by directly referring to Participant 1's personal display.

In addition to these active examples of grounding, participants would refer to the shared display on their own when receiving instructions from a collaborator that

they did not understand. For example, in one case Group 4's Participant 2 asked if a job component could be moved using verbal deixis (e.g. "Move this 'A'"). Since Participant 3 could not see Participant 2's mouse cursor on their personal display, they quickly glanced at the shared display where all cursors were present (Figure 8).

(26:13) **G4P2**: can we move this 'A' forward more?

(26:14) [p3 checks shared screen]

(26:15) **G4P2**: so the rest of this can move back again?

(26:16) **G4P3**: yea

[p2 continues working]



Figure 8. After Participant 2 suggests a job component move, Participant 3 checks the shared display to resolve verbal deixis.

Synchronization

Participants frequently monitored one another's posture, orientation and physical gestures while performing the task, and the awareness gained from this monitoring aided in synchronizing group activity. The most common example of this synchronization activity was observed when groups were nearing the end of the task and were deciding whether or not to submit their current solution. Typically, as participants emerged from activity phases on their personal display, they would look up at the shared display. Other participants would recognize their body

language and would face the shared display as well. Once all three participants were focused on the shared screen, a consensus was reached and a solution was submitted. Figure 9 illustrates a group shifting from an activity phase of work towards submitting a final solution, with participants in varying states of transition.



Figure 9. Towards the end of their trial, Participants 2 and 3 lean back and inspect the shared display. This behaviour typically initiated a “strategy” phase of collaboration or a consensus to submit.

5. DISCUSSION

Consistent with our previous work (Wallace et al., 2009), participants in this study tended to focus on their personal displays while performing taskwork. This tendency appears to arise from a combination of the cognitive nature of the JSS task and the personalized workspace provided on the laptop displays, since participants were better able to focus on the task (Wallace et al., 2009). Participants’ subjective reports indicated that they did not feel that the shared display was necessary to complete the task, and that they felt they would be satisfied with a personal workspace that integrated the shared display’s functionality. While these comments suggest that the shared display was not

overtly perceived to add significant value to the task, observed use suggests it offers an important, if subtle, benefit in fostering teamwork. The benefits of mutual gaze and a shared visual workspace are well established in the CSCW and small group research literature (Argyle & Cook, 1976; Boyle et al., 1994; Fussell et al., 2003; Gergle et al., 2006), however understanding how to apply existing theory and to optimize shared displays for use within MDEs is a relatively new area of study. In studying MDEs, the complexity of multi-user, multi-display interaction means that the validity in applying theories developed for single- and dual-display configurations to MDEs is uncertain.

This work represents a first step towards understanding how existing theory can be applied to these more complex working environments, and in particular, has provided insight into how the alternative MDE relationship models support collaboration. Our observations suggest that the replicated content display better supported grounding, whereas the status display better supported monitoring. In both cases, the physical presence of a shared display appears to have facilitated the synchronization of group activity. After discussing how these results can help to refine existing theory, we will discuss more concrete implications for the design of MDEs.

5.1 Replicated Content Displays & Grounding

Our observations of participants utilizing the replicated content display for grounding support the hypothesis that the replicated content display supports grounding in ways that the status display cannot. Clark and Brennan (1991) provide a theoretical basis for these benefits when they assert that “grounding changes with the medium” (p 140). In MDEs that support content replication,

constraints on visibility and reviewability are removed, and the production, reception, and display costs of referential grounding are alleviated by the presence of shared views. That is, alternative views of task resources in MDEs allow users to simultaneously work in a sheltered workspace while maintaining awareness of shared task resources; by working in such a hybrid environment, users can effectively work as individuals while retaining ties to the group. This hybrid working environment, however, comes at the expense of an increased cost of communication between individuals; in our study, participants would occasionally gesture towards their personal display when talking to peers with the (often mistaken) expectation that their peers would understand their deictic references.

In our study, grounding behaviour was observed on the shared display despite the personal displays providing awareness of much of the group activity. In both shared display conditions, all task components were visible on each of the personal displays (only the salience of components was altered between personal displays), however all participants' cursors were only visible simultaneously on the shared display. Thus, when participants needed to regain awareness of their peers' actions, they often looked to the shared display where cursor information was available for the entire group. This display appeared to provide a type of "openness", similar to the team-optimized tools described by Hutchins' (1990) investigations of naval navigation tools. Such tools provide a visibility of others' taskwork that contributes to task awareness and coordination. The observed use of the replicated content display would indicate that participants were able to work with a cursory awareness of each others' activity most of the time, however the more open shared display was useful for repairs when communication broke down (Clark & Brennan, 1991).

Our observation of these phenomena, primarily in the replicated content display condition, suggests that such shared display content is particularly effective in supporting grounding in group work, even when personalized views may replicate much of the replicated content display. These results are particularly interesting because, while environments that provide access to shared resources via personalized views (e.g. Berry et al., 2005; Sugimoto et al., 2004) have been explored in the literature, the hardware configuration in our study provides a unique perspective on the role of personal and shared displays. For example, Sugimoto et al. (2004) found that in the Caretta environment consisting of a table and handheld computers, personal work was conducted on handhelds while group work and negotiations predominantly occurred on the shared tabletop. Our results suggest that groups working with more powerful personal devices (e.g. laptops) and carefully designed personal user interfaces may rely more heavily on their personal devices, and only use a replicated content display to manage group awareness.

5.2 Status Displays and Monitoring

On the other hand, groups working in the status display condition tended to utilize the shared display for monitoring, rather than grounding. In this condition, only one view of the shared workspace can be seen on participants' laptops, limiting the opportunity for users to be aware of their peers' task interactions, and thereby increasing the amount of effort required for grounding. Despite these shortcomings, the shared display's alternate task view was useful in that participants were able to more seamlessly monitor task progress. This utility was demonstrated not only by the participants repeated use of the display (i.e. looking

at the status display approximately three times as often as the replicated content display), but also through participants' self-reported preference for the status display configuration on the post-condition questionnaires.

As Grudin (2001) discusses for single-user settings, an advantage of multi-display configurations is that the division of tasks amongst multiple displays can reduce the cognitive load associated with transitioning between tasks. In our study, the partitioning of task information appeared to provide participants support for transitioning between group and individual work, as illustrated by our observations of participants shifting their attention from personal to shared displays, and vice-versa, as they transitioned between activity and strategy phases of work. In this sense, the shared display not only acted as a common workplace for the group, but also as a *secondary* display for individual work. Our findings support observations by Biehl et al. (2007) in their evaluation of FASTDash, a tool developed to support workspace awareness in programming teams. In their study, Biehl et. al. found that most programmers preferred to maintain an open copy of FASTDash on a secondary monitor, rather than refer to the large, replicated content display projected nearby. One interpretation of these results is that displays designed to support activity awareness (Carroll et al., 2003) may be viewed by group members as supporting an individual task (i.e. the act of monitoring) rather than group work, and may therefore be best implemented as secondary displays within a personal workspace.

Our implementation of an active status display also contrasts observed use of MDEs in which shared displays are used to display less frequently updated content. Our work, and that of Biehl et al.'s (2007), placed an emphasis on

sustained use of shared displays for the display of real-time information regarding the state of the replicated content display. In other cases, the “at-a-glance” availability of the status display has previously been identified as a strength of shared displays used for monitoring group work (Huang et al., 2006; Plaue et al., 2009). For example, in a field study of conference meeting room use at a global corporation, Plaue et al. (2009) suggest that idle displays be used for peripherally relevant information such as performance metrics.

We interpret the variety of shared content types and monitoring activity to be indicative of the flexibility required in displaying content in MDEs. Even in scenarios that are marginally different from one another, individual preferences may impact the optimal display configuration for a given group, task, or environment. Window management tools should take these preferences into account; some users may prefer to have at-a-glance information available on a nearby wall display, whereas others may wish to have it available on a secondary personal display on their desk.

5.3 Physical Design of MDEs & Synchronization

The results of this study suggest that the physical presence of the shared display appears to have been beneficial in synchronizing group activity regardless of its content. One participant commented that “if [the status display] was on the laptop, then we wouldn’t be communicating as much.” Often, participants glancing at the display would trigger group interactions through a change in body position, and participants concentrated on this phenomenon during the post-study interviews. One participant explained, “sometimes looking at the shared [display, it’s] like

everyone's actually talking like instead of looking at their screens like it's a time to gather around.”

While we did not explicitly test alternate participant seating positions, we hypothesize that the “around the table” configuration employed in our study helped to facilitate the use of body language in collaboration. This hypothesis varies somewhat from Sommer’s (1969) findings in which collaborators preferred to sit side-by-side during cooperative work, but in adjacent corner configurations for conversations, citing the ability to share physical artifacts as motivating the preference for adjacent seating position in collaborative settings. In the case of collaborative environments using content replication, the sharing of physical artifacts is not a major concern. Our participants reported that a seating configuration in which body language is more easily observed promoted group interactions; one participant explained their preference for “around the table” seating configurations by saying “if we sat in a straight row we wouldn’t be discussing in a circle, we’d be talking to a wall.” These results suggest that the face-to-face configuration augmented with content replication utilized in this study may provide a “best of both worlds” setup in regards to Sommer’s reported seating preferences.

The benefits of face-to-face configurations are further clarified by Kendon’s F-formation theory (Kendon, 1990) and its description of the role of gaze and body language in moderating collaboration. Kendon observed that group members’ orientation dictates “transactional space”, or the common workspace utilized by a group, and that a peer’s orientation and gaze relative to the group’s transactional space is often used to communicate intent or motivation in conducting group

work. For example, a group of peers working around a table would define the physical space between them as transactional space (i.e. the table's surface), and group interactions would then be carried out in that shared space. Kendon's theory suggests that gaze and body position relative to this transactional space often moderates collaboration (see also, Cook's description of gaze in moderating conversation (Argyle & Cook, 1976; Cook & Lalljee, 2009)).

An interesting difference between our work and the work motivating Kendon's theory is that F-formations were, like Sommer's (1969) theory, developed in a purely physical domain. Our results help to interpret the theory's application in cases where digital devices are used to support collaboration. Our observations suggest that users maintained a transactional space at the shared display, whereas personal displays were maintained largely as separate, personal workspaces. As participants shifted their gaze and body orientation between the shared and personal displays, transactional space was established, broken, and re-established, marking transitions between group and individual work. Kendon observed similar behaviour in that participants often rapidly shift between group and individual work (e.g. quickly check to see if anyone new is in the room), or establish more long term shifts in gaze (e.g. synchronize with new group members entering the group).

Notably, the physical layout of participants in our study closely resembles a commonly used configuration in today's workplace; collaborators were seated at a table, each with laptops open in front of them, and a projected display on a nearby wall. As Kendon's F-formation theory was first published in 1990, before mobile computing, and in particular laptops, became common, our results provide an

opportunity to elaborate on how it may be applied to more technically-driven environments. In our study, the laptops' vertical displays created a partial visual barrier between participants, effectively dividing what would traditionally be called the group's transactional space. Our results suggest that having users seated in a face-to-face configuration facilitates the use of body language and gaze to synchronize activity between personal and shared displays in MDEs.

However, not all participants agreed that the face-to-face configuration was optimal. One participant, who tended to direct his collaborators by pointing to their personal displays, indicated in the post-study interview that he would have preferred to sit side-by-side, as it would be easier to gesture towards the shared display. The presence of these differing views leads us to believe that in cases where physical sharing is important (e.g. preference for referring to others' laptops), a side-by-side seating configuration may be appropriate, however when virtual sharing is possible (and implemented), a face-to-face (e.g. around the table) seating arrangement is preferable.

Finally, while the presence of a nearby, shared wall display appeared to facilitate the use of body posture and other non-verbal communication in synchronizing collaborators' interactions, it is important to note that in this study both shared display configurations displayed task-related content. We assert that this behavior was observed because the shared display contained task-related content; if the display contained non-task-related content, it would not have played a role in synchronizing collaboration to the extent observed in this study.

5.4 Classifying Collaborative Tasks

McGrath's task circumplex (McGrath, 1984) is widely cited by academic papers while describing collaborative tasks, particularly in experimental settings. When the JSS task was introduced by Tan et al. it was classified as *intellective*, because it has a demonstrably correct solution and participants can objectively compare one solution with another (Tan et al., 2008). During our first investigation of groups completing the JSS task we observed that the strategies groups used to solve the puzzle were similar to those used in a *decision-making* process. That is, while in theory there is always a means to objectively compare two potential solutions and demonstrate that one is better than the other, individuals within the group were not always capable of doing so. Consequently, groups tended to adopt hill-climbing strategies where each step was suggested to the group by an individual and negotiated, and their final solution was largely dependent on the first job components they scheduled.

As mentioned in section 4.1, we added the 'scrubber' feature (Figure 1, D) to the JSS interface with the goal of reducing the cost to groups when backtracking and experimenting with alternative solutions during their trial. Subjectively, we found that the scrubber was effective, and that groups in this study were willing to explore alternative solutions more often than they had in the previous study; however we also did not see any change in the processes that groups used in solving the puzzle, and hill-climbing was the dominant strategy used throughout the study. Our quantitative results indicated that participants spent more time solving the JSS task, and had better solutions on average than in the first JSS study. It is also worth noting that our participants were largely undergraduate students studying Math or Engineering, who would presumably be more inclined

to apply rigorous mathematical discussions to the task than the population at large. We assert that due to its inherently large solution space (6 jobs with 6 components each, and several restrictions on component order and positions) the JSS task is solved by participants as a decision-making task rather than intellectual. As our study focused on group behaviour in relation to technologies provided by the collaborative environment, and not on specific task-related behavior, we believe that the results of this work are supported despite the difference between intended and actual task type.

This difficulty in consistently classifying collaborative tasks is not limited to the JSS task. Fjermestad & Hiltz (1997), in their review of experimental studies of collaboration, also observed that “some tasks had been described as belonging in two or three [of McGrath’s] categories, by different authors” (p 6), indicating that there is at least some disagreement in the community as to how tasks should be classified. One of McGrath’s stated objectives in developing the circumplex was to provide a set of categories that were both mutually exclusive and useful in that they expound differences between and relations among tasks (McGrath, 1984). We suggest that based on our experience with the JSS task, even though theoretically a task may fall into one classification, participants may *choose* to use an alternative method when completing the task. Hackman defines this behavior as *task redefinition* (Hackman, 1969), and notes that it affects both group processes and collaborative outcomes. Thus, it may benefit the research community to revisit McGrath’s task circumplex and explore classification schemes grounded in the potential processes employed by groups while solving a task rather than the task requirements alone.

6. IMPLICATIONS FOR DESIGN

In our previous work, we observed that individuals would often “step back” and assess the state of the problem using the shared display. In this work, the role of the shared display was further examined to reveal that it can facilitate monitoring, communication grounding, and synchronization. As discussed by Birnholtz et al. (2007), one factor that appears to play an important role in collaboration is group strategy; in multi-input collaborative environments, users appear to tackle problems using an egocentric approach. The use of the shared display that we observed in this study, however, was collaborative in nature, and supports lightweight transitions between personal and group work. Plaue and Stasko (2009) also note that shared display configurations influence group process by providing opportunities for group interaction that otherwise would be unavailable to the group. They state that “having multiple displays also prevents one individual from dominating the group conversation” (p 187). Thus, the presence of shared displays in MDEs presents an opportunity to balance the otherwise egocentric processes that can occur in collaborative environments when individuals are provided with personal input devices.

6.1 The Value of Shared Displays in MDEs

In determining a shared display’s value within the MDE it is important to understand and balance the costs at design time with the benefits provided to the group. For example, at the time of writing this article the monetary cost of large displays, either projected or LCD, ranges from \$300 to \$25,000, which can be a prohibitive factor in building MDEs in many contexts. In order to support group interactions with a shared display, the environment must provide physical access to the display, often at the expense of additional seating. Hardware and software

must be designed to accommodate the transfer of artifacts between personal and replicated content displays, introducing complexity into the environment. Our results provide some guidance in designing collaborative technologies to support both collaborative processes and taskwork.

Our results suggest that carefully designed personal displays can be effective in supporting taskwork. Feedback from participants during our previous study (Wallace et al., 2009) helped to fine tune the JSS task interface, and led to the design of customized personal workspaces. The effectiveness of the personal devices was particularly salient when comparing performance differences between display configurations – groups in the MDE conditions outperformed groups using a SDG configuration. Feedback from participants in this study reaffirms this decision. For example one participant in our study noted that while working on the task “you just want to focus on the thing you’re focused on,” suggesting that introducing additional communication tools into the personal workspace may detract from its usefulness.

We interpret the group’s prevalent use of both shared displays, despite our significant efforts to fine tune the personal workspaces, as an indication of the value of the shared display. The conceptual separation of a task-intensive, personal workspace and a replicated content display facilitated individual taskwork while simultaneously supporting grounding. The physical separation of a personal space and a group-centric replicated content display also appears to provide value to the group, particularly in leveraging existing social indicators (body position, gaze) in moderating interactions within the MDE. In supporting tasks where coordination between group members is essential, such as in

commercial airline cockpits, command centers or war rooms, the additional cost of a shared display may be acceptable and the advantages identified here may easily justify additional development costs. In supporting tasks that have less rigorous group awareness and coordination requirements, such as home gaming and multimedia systems, the additional costs may not be justified. In either case, we believe that the results of this work can help designers make better informed decisions regarding when to consider the inclusion of shared displays in a multi-display environment.

Finally, the display conditions utilized in our combined studies were limited to SDG configurations (i.e. a single, shared display) and MDE configurations with three personal laptops and a single shared display. Moreover, they were conducted in a controlled, laboratory setting with artificial groups and an emulated intellective/decision-making task. To fully understand the breadth of design choices available to designers of collaborative environments, and the impact of those choices on teamwork and taskwork, a more complete series of studies contrasting a wider range of display configurations, settings and tasks warrants investigation. For example, in situ field studies of collaborative environments would provide an opportunity to explore long-term use of technologies and to observe how social, cultural, ergonomic, activity and temporal contextual factors influence the use of shared displays outside of the lab (Wallace & Scott, 2008).

7. CONCLUSION

In this work, we have described an empirical study of groups performing the JSS task in a multi-display environment; discussed theoretical implications of the observed monitoring, grounding and synchronization behaviours involving the

shared display provided in the environment; and derived implications for the design of collaborative environments where effective taskwork and awareness are important design considerations. The primary motivation behind this study was to more fully explore observed behaviour in our previous work comparing SDG and MDE display configurations that revealed that participants predominantly used the provided personal displays in the MDE configuration to complete the JSS task, but that the shared display appeared to play a role in moderating group process. This follow-up study helped to clarify which aspects of the collaboration process are supported by the inclusion of the shared display. In particular, we found that participants used the replicated content display to ground conversation, whereas they used task status information on the shared display to monitor the overall group activity. Furthermore, the physical presence of the shared display, and arrangement of participants around it, facilitated group synchronization by providing a shared physical reference for non-verbal communication mechanisms such as body posture and gaze.

Computer supported collaborative environments in practice will consist of both personal and shared devices; understanding how to design interfaces that support collaboration in such environments is an important problem for CSCW researchers. While guidelines and requirements have been developed to aid designers in developing single- and multi-display groupware (e.g. Elwart-Keys et al., 1990; Scott et al., 2003), there are still significant gaps in the literature, particularly in understanding how specific displays support group work and in how to design an MDE to support specific tasks. These questions are unique to MDE research, in that MDEs can be designed with any number, type, and arrangement of displays. The results presented here contribute to such an

understanding by contrasting specific implementations of replicated content and status displays, and in further developing theory based on SDG research for use in the design of MDEs.

ACKNOWLEDGEMENTS

This research is funded by the Natural Science and Engineering Council of Canada (NSERC). The authors thank the members of the Collaborative Systems Lab and Human Factors Research Group at the University of Waterloo for their insights and assistance in conducting this research. The authors also thank Dr. Neil Randall for his assistance while revising drafts of this work.

REFERENCES

- Argyle, M., & Cook, M. (1976). *Gaze and mutual gaze*. Cambridge: Cambridge University Press.
- Baker, K., Greenberg, S., & Gutwin, C. (2002): *Empirical development of a heuristic evaluation methodology for shared workspace groupware*. Proceedings of the ACM conference on Computer Supported Cooperative Work (CSCW), New Orleans, Louisiana, USA: pp. 96 - 105.
- Berry, L., Bartram, L., & Booth, K. S. (2005): *Role-based control of shared application views*. Proceedings of the ACM Symposium on User Interface Software and Technology (UIST 2004), Seattle, WA, USA: ACM, pp. 23 - 32.
- Biehl, J. T., & Bailey, B. P. (2004): *ARIS: an interface for application relocation in an interactive space*. Proceedings of Graphics Interface (GI 2004), London, Ontario, Canada: Canadian Human-Computer Communications Society, pp. 107-116.
- Biehl, J. T., Czerwinski, M., Smith, G., & Robertson, G. G. (2007): *FASTDash: a visual dashboard for fostering awareness in software teams*. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI), San Jose, California, USA: ACM, pp. 1313 - 1322.
- Birnholtz, J. P., Grossman, T., Mak, C., & Balakrishnan, R. (2007): *An exploratory study of input configuration and group process in a negotiation task using a large display*. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI), San Jose, California, USA: ACM, pp. 91 - 100.
- Boyle, E. A., Anderson, A. H., & Newlands, A. (1994). The effects of visibility on dialogue and performance in a cooperative problem solving task. *Language & Speech*, 37(1), 1-20.
- Carroll, J. M., Neale, D. C., Isenhour, P. L., Rosson, M. B., & McCrickard, D. S. (2003). Notification and awareness: synchronizing task-oriented collaborative activity. *International Journal of Human-Computer Studies*, 58(5), 605 - 632.

- Clark, H. H., & Brennan, S. E. (1991). Grounding in Communication. In L. B. Resnick, J. Levine & S. D. Teasley (Eds.), *Perspectives on socially shared cognition*. Washington, DC: Psychological Association, pp. 127-149.
- Collins, A., Apted, T., & Kay, J. (2007): *Tabletop File System Access: Associative and Hierarchical Approaches*. Proceedings of the IEEE International Workshop on Horizontal Interactive Human-Computer Systems (TABLETOP 2007), Newport, RI, USA: IEEE, pp. 113 - 120.
- Cook, M., & Lalljee, M. G. (2009). Verbal Substitutes for Visual Signals in Interaction. *Semiotica*, 6(3), 212-221.
- DeSanctis, G., & Gallupe, R. B. (1987). A Foundation for the Study of Group Decision Support Systems. *Management Science*, 33(5), 589-609.
- DiMicco, J. M., Pandolfo, A., & Bender, W. (2004): *Influencing group participation with a shared display*. ACM Conference on Computer Supported Cooperative Work, Chicago, Illinois, USA: ACM, pp. 614 - 623.
- Elwart-Keys, M., Halonen, D., Horton, M., Kass, R., & Scott, P. (1990): *User interface requirements for face to face groupware*. Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI), Seattle, Washington, United States: ACM, pp. 295 - 301.
- Fjermestad, J., & Hiltz, S. R. (1997): *Experimental Studies of Group Decision Support Systems: An Assessment of Variables Studied and Methodology*. Proceedings of the International Conference on System Sciences, Wailea, HI, USA: IEEE Computer Society, pp. 45-65.
- Forlines, C., Shen, C., Wigdor, D., & Balakrishnan, R. (2006): *Exploring the effects of group size and display configuration on visual search*. In Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW), Banff, Alberta, Canada: ACM Press, pp. 11 - 20.
- Fussell, S. R., Setlock, L. D., & Parker, E. M. (2003): *Where do helpers look?: gaze targets during collaborative physical tasks*. In Extended Abstracts of the ACM Conference on Human Factors in Computing Systems (CHI EA), Ft. Lauderdale, Florida, USA: ACM Press, pp. 768 - 769.
- Gergle, D., Kraut, R. E., & Fussell, S. R. (2006): *The impact of delayed visual feedback on collaborative performance*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI): ACM Press, pp. 1303 - 1312.
- Gergle, D. R. (2006). *The value of shared visual information for task-oriented collaboration*. Carnegie Mellon University, Pittsburgh, PA, USA.
- Greenberg, S. (1991): *Personalizable Groupware: Accommodating Individual Roles and Group Differences*. Proceedings of the ECSCW '91 European Conference of Computer Supported Cooperative Work Amsterdam, NL: Kluwer Academic Press, pp. 17 - 32.
- Greenberg, S., Boyle, M., & Laberge, J. (1999). PDAs and shared public displays: Making personal information public, and public information personal. *Personal and Ubiquitous Computing*, 3(1), 54-64.
- Greenberg, S., & Rounding, M. (2001): *The notification collage: posting information to public and personal displays*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI), Seattle, Washington, United States: ACM Press, pp. 514 - 521.
- Grudin, J. (2001): *Partitioning digital worlds: focal and peripheral awareness in multiple monitor use*. Proceedings of the SIGCHI Conference on Human

- Factors in Computing Systems (CHI), Seattle, Washington, United States: ACM Press, pp. 458 - 465.
- Gutwin, C., & Greenberg, S. (1998): *Design for individuals, design for groups: tradeoffs between power and workspace awareness*. Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work (CSCW), Seattle, Washington, United States: ACM, pp. 207-216.
- Gutwin, C., Greenberg, S., & Roseman, M. (1996). Workspace Awareness in Real-Time Distributed Groupware: Framework, Widgets, and Evaluation. In A. Sasse, J. Cunningham & R. L. Winder (Eds.), *Proceedings of HCI on People and Computers XI*. London: Springer-Verlag, pp. 281-298.
- Hackman, J. R. (1969). Toward Understanding the Role of Tasks in Behavioral Research. *Acta Psychologica*, 31, 97-128.
- Hailpern, J., Hinterbichler, E., Leppert, C., Cook, D., & Bailey, B. P. (2007): *TEAM STORM: demonstrating an interaction model for working with multiple ideas during creative group work*. Proceedings of the 6th ACM Conference on Creativity & Cognition (C&C '07), Washington, DC, USA: ACM Press, pp. 193 - 202.
- Hart, S. G., & Stateland, L. E. (1988). Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Human Mental Workload*, 139-183.
- Hawkey, K., Kellar, M., Reilley, D., Whalen, T., & Inkpen, K. M. (2005): *The proximity factor: impact of distance on co-located collaboration*. Proceedings of the ACM International Conference on Supporting Group Work (GROUP 2005), Sanibel Island, Florida, USA: ACM Press, pp. 31 - 40.
- Huang, E., Mynatt, E., & Trimble, J. (2006). Displays in the Wild: Understanding the Dynamics and Evolution of a Display Ecology. In *Lecture Notes in Computer Science: Pervasive Computing*. Berlin: Springer, pp. 321-336.
- Hutchins, E. (1990). The Technology of Team Navigation. In J. Galegher, R. E. Kraut & C. Egido (Eds.), *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*. Hillsdale, NJ, USA: Erlbaum Associates, pp. 191-220.
- Ishii, H., & Kobayashi, M. (1992): *ClearBoard: a seamless medium for shared drawing and conversation with eye contact*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Monterey, California, United States: ACM Press, pp. 525 - 532.
- Johanson, B., Hutchins, G., Winograd, T., & Stone, M. (2002): *PointRight: experience with flexible input redirection in interactive workspaces*. Proceedings of the ACM Symposium on User Interface Software and Technology (UIST 2002), Paris, France: ACM, pp. 227 - 234.
- Kendon, A. (1990). *Conducting Interaction: Patterns of Behavior in Focused Encounters*. New York: Cambridge University Press.
- Lanir, J., Booth, K. S., & Tang, A. (2008): *MultiPresenter: a presentation system for (very) large display surfaces*. Proceedings of the 16th ACM international conference on Multimedia, Vancouver, British Columbia, Canada: ACM, pp. 519-528.
- McGrath, J. (1984). *Groups: Interaction and Performance*. Englewood, NJ: Prentice-Hall.
- McKay, P. (2009). *Design of Collaborative Systems for Modern Cockpits*. University of Waterloo, Waterloo, Ontario, Canada.

- Moore, G. (2005). Excerpts from A Conversation with Gordon Moore: Moore's Law. Video Transcript. Intel Corporation. http://download.intel.com/museum/Moores_Law/Video-transcripts/Excerpts_A_Conversation_with_Gordon_Moore.pdf. Retrieved Oct 29th, 2010.
- Moore, G. E. (1965). Cramming more components onto integrated circuits. *Electronics*, 38(8).
- Morris, M. R., Paepcke, A., Winograd, T., & Stamberger, J. (2006): *TeamTag: exploring centralized versus replicated controls for co-located tabletop groupware*. Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Montréal, Québec, Canada: ACM, pp. 1273 - 1282.
- Myers, B. A. (2000): *The Pebbles Project: using PCs and hand-held computers together*. Extended Abstracts on Human Factors in Computing Systems (CHI EA), The Hague, The Netherlands: ACM, pp. 14 - 15
- Nacenta, M. A., Aliakseyeu, D., Subramanian, S., & Gutwin, C. (2005): *A comparison of techniques for multi-display reaching*. Proceedings of the SIGCHI conference on Human Factors in Computing Systems (CHI), Portland, Oregon, USA: ACM, pp. 371 - 380.
- Nunamaker, J. F., Briggs, R. O., Mittleman, D. D., Vogel, D. R., & Balthazard, P. A. (1996). Lessons from a dozen years of group support systems research: a discussion of lab and field findings. *Journal of Management Information Systems Special issue: Information technology and its organizational impact*, 13(3), 163 - 207
- Pinelle, D., Gutwin, C., & Greenberg, S. (2003). Task analysis for groupware usability evaluation: Modeling shared-workspace tasks with the mechanics of collaboration. *ACM Trans. Comput.-Hum. Interact.*, 10(4), 281-311.
- Plaue, C., & Stasko, J. (2009): *Presence & Plancement: Exploring the Benefits of Multiple Shared Displays on an Interactive Sensemaking Task*. Proceedings of the ACM 2009 International Conference on Supporting Group Work (GROUP '09), Sanibel Island, FL, USA: ACM Press, pp. 179-188.
- Plaue, C., Stasko, J., & Baloga, M. (2009). *The conference room as a toolbox: technological and social routines in corporate meeting spaces*. Paper presented at the Proceedings of the Fourth International Conference on Communities and Technologies (C&T '09).
- Ryall, K., Forlines, C., Shen, C., & Morris, M. R. (2004): *Exploring the effects of group size and table size on interactions with tabletop shared-display groupware*. Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work (CSCW), Chicago, Illinois, USA: ACM, pp. 284 - 293
- Scott, S. D., Grant, K. D., & Mandryk, R. L. (2003): *System Guidelines for Co-located, Collaborative Work on a Tabletop Display*. European Conference on Computer Supported Cooperative Work (ECSCW), 14-18 September 2003, Helsinki, Finland: Springer, pp. 159 - 178.
- Seifried, T., Haller, M., Scott, S. D., Perteneder, F., Rendl, C., Sakamoto, D., & Inami, M. (2009): *CRISTAL: a collaborative home media and device controller based on a multi-touch display*. Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces (TABLETOP 2009), Banff, Alberta, Canada: ACM, pp. 33-40.
- Shen, C., Everitt, K. M., & Ryall, K. (2003): *UbiTable: Impromptu Face-to-Face Collaboration on Horizontal Interactive Surfaces*. Proceedings of

- UbiComp'03: 5th Int'l Conference on Ubiquitous Computing, October 12-15, 2003, Seattle, WA: pp. pp. 281-288.
- Shoemaker, G. B. D., & Inkpen, K. M. (2001): *Single display privacyware: augmenting public displays with private information*. Proceedings of the SIGCHI conference on Human Factors in Computing Systems (CHI), Seattle, Washington, United States: ACM, pp. 522 - 529.
- Sommer, R. (1969). *Personal space: The behavior basis of design*. Cliff, NJ: Prentice-Hall.
- Stewart, J. E. (1999): *Single display groupware: a model for co-present collaboration*. Proceedings of the SIGCHI conference on Human Factors in Computing Systems (CHI), Pittsburgh, Pennsylvania, United States: ACM, pp. 286 - 293.
- Su, R. E., & Bailey, B. P. (2005). Put Them Where? Towards Guidelines for Positioning Large Displays in Interactive Workspaces. In M. F. Costabile & F. Paternò (Eds.), *Human-Computer Interaction - INTERACT 2005*. Heidelberg: Springer Berlin, pp. 337-349.
- Sugimoto, M., Hosoi, K., & Hashizume, H. (2004): *Caretta: a system for supporting face-to-face collaboration by integrating personal and shared spaces*. Proceedings of the SIGCHI conference on Human Factors in Computing Systems (CHI 04), Vienna, Austria: ACM, pp. 41-48.
- Tan, D. S., Gergle, D., & Czerwinski, M. (2005). *A Job-Shop Scheduling Task for Evaluating Coordination during Computer Supported Collaboration*: Microsoft Research, Redmond, WA, USA.
- Tan, D. S., Gergle, D., Mandryk, R., Inkpen, K., Kellar, M., Hawkey, K., & Czerwinski, M. (2008). Using job-shop scheduling tasks for evaluating collocated collaboration. *Personal Ubiquitous Computing*, 12(3), 255-267.
- Wallace, J., Ziola, R., Ha, V., & Inkpen, K. (2006): *Swordfish: user tailored workspaces in multi-display environments*. Extended Abstracts on Human Factors in Computing Systems (CHI EA), Montréal, Québec, Canada: ACM, pp. 1487 - 1492.
- Wallace, J. R., & Scott, S. D. (2008): *Contextual Design Considerations for Co-located, Collaborative Tables*. Proceedings of the IEEE Conference on Tabletops and Interactive Surfaces (TABLETOP 2008), 1-3 Oct. 2008, Amsterdam, Netherlands: IEEE, pp. 57 - 64.
- Wallace, J. R., Scott, S. D., Stutz, T., Enns, T., & Inkpen, K. (2009). Investigating teamwork and taskwork in single- and multi-display groupware systems. *Personal and Ubiquitous Computing*, 18(8), 569-581.
- Wigdor, D., Jiang, H., Forlines, C., Borkin, M., & Shen, C. (2009): *WeSpace: the design development and deployment of a walk-up and share multi-surface visual collaboration system*. Proceedings of the 27th International Conference on Human Factors in Computing Systems (CHI), Boston, MA, USA: ACM, pp. 1237 - 1246.