Avoiding Interference: How People Use Spatial Separation and Partitioning in SDG Workspaces

Edward Tse[1], Jonathan Histon[2], Stacey Scott[1] and Saul Greenberg[1]
[1]Department of Computer Science, [2]Department of Psychology
University of Calgary
Calgary, Alberta, Canada, T2N 1N4
+1 403 220 6087

ABSTRACT
Single Display Groupware (SDG) lets multiple co-located people, each with their own input device, interact simultaneously over a single communal display. While SDG is beneficial, there is risk of *interference*: when two people are interacting in close proximity, one person can raise an interface component (such as a menu, dialog box, or movable palette) over another person's working area, thus obscuring and hindering the other's actions. Consequently, researchers have developed special purpose interaction components to mitigate interference techniques. Yet is interference common in practice? If not, then SDG versions of conventional interface components could prove more suitable. We hypothesize that collaborators spatially separate their activities to the extent that they partition their workspace into distinct areas when working on particular tasks, thus reducing the potential for interference. We tested this hypothesis by observing co-located people performing a set of collaborative drawing exercises in an SDG workspace, where we paid particular attention to the locations of their simultaneous interactions. We saw that spatial separation and partitioning occurred consistently and naturally across all participants, rarely requiring any verbal negotiation. Particular divisions of the space varied, influenced by seating position and task semantics. These results suggest that people naturally avoid interfering with one another by spatially separating their actions. This has design implications for SDG interaction techniques, especially in how conventional widgets can be adapted to an SDG setting.

Categories and Subject Descriptors
H5.3. Group and organization interfaces: Computer supported cooperative work.

Keywords
Single display groupware (SDG), co-located collaboration.

1. SINGLE DISPLAY GROUPWARE (SDG)
The vast majority of work meetings, whether scheduled or casual, are between co-located people. When information is located on a single computer, these co-located people often share and manipulate the information on the display even though traditional computers know nothing about groups [12]. People 'get by' through over-the-shoulder and side-by-side viewing of computer monitors, through projected displays, and by sharing input devices through turn-taking.

Yet applications that are aware of the presence of multiple users can significantly enhance how people work together. They can allow simultaneous interaction [19][20]; they can remember personal information and modes; they can allow individuals to bring their private information into the public arena [6]. Consequently, researchers in CSCW are now paying considerable attention to the design of single display groupware (SDG) i.e., applications that support the work of co-located groups over a physically shared display [18]. What distinguishes SDG from conventional applications is that each participant has his or her own input device, that the SDG application distinguishes individuals by these devices, and that it allows everyone to interact simultaneously with the common display.

While SDG research has evolved somewhat slowly over the last decade, it has received a recent explosion of interest: the availability of multiple monitors, high resolution displays, large projected displays, multi-touch displays (e.g., [5][17]), and electronic table-tops naturally afford multi-person interaction. As a result, researchers are now looking at the nuances of SDG interaction. Our own particular interest is in the risk and mitigation of *interference*.

1.1 Interference in SDG
In a physical workspace, people's use of space is constrained by the location of their arms and bodies. While there is much potential for one person to interfere with others if they are working over a small space, they naturally mediate their actions through social protocol and unspoken actions [13]. For example, people standing in front of a whiteboard often turn-take, or choose a side to work on, or wait for others to complete an action in an area before moving into it to mitigate against their bodies, arms and actions getting in the way of each other [20].

Physical occurrences of interference disappear when collaborators interact in an SDG workspace with indirect input devices, such as mice or laser pointers. Because people now have 'virtual hands' (the cursors), this environment enables people to easily reach any part of the workspace and even to work atop each other. While this ability to simultaneously work in the same area of the workspace increases interaction capabilities, it also introduces a different possibility of interference, where one person's virtual actions potentially disturb the productivity of others. For example, a person could raise an interface component, such as a floating palette, a menu or a dialog box, thus obscuring a nearby person's working area underneath the raised menu or dialog [23]. As a result some researchers have developed tools and widgets that mitigate interference effects (see next section).

Yet to our knowledge, no one has carefully examined whether interference really occurs. That is, do people actually work concurrently and within close proximity in a SDG workspace, thus increasing the likelihood of interference? Or do they naturally *spatially separate* their activities from one another, perhaps to the point where they restrict their activities to separate areas, or *partitions,* of the workspace? If they do, these minimize the potential for interference. This paper addresses these questions through an investigation of the natural behavior of co-located collaborators sharing an SDG workspace. These answers are important: if interference proves rare, then the design of SDG interaction techniques would shift away from inference-avoiding techniques and towards SDG tools, objects and applications that recognize people's natural use of space.

To set the scene, we first review existing interaction techniques for mitigating interference effects, and what is known about spatial separation and partitioning of workspaces. The subsequent sections detail our study and its results. We then show the relevance of our results by providing examples of what they mean to the design of SDG interaction techniques.

2. RELATED WORK

2.1 Techniques that mitigate interference

Several notable interaction techniques have been proposed to help reduce the potential for interference in an SDG workspace.

Bederson et al. [2] proposed the use of *Local Tools* within a children's application, where all application functionality is accessed through tools represented as separate icons scattered across the shared workspace. Because all tools are visible in the workspace, there is no need for participants to use transient windows to display them, e.g., menus or floating palettes. However, this solution is limited to applications with relatively few tools, as those with large numbers would overwhelm the interface.

Zanella and Greenberg [23] investigated the effects of using *transparent interface components* in SDG workspaces. Using a popup menu as an example, the idea is that even if someone raised a menu over another person's work area, that person working in the space could see through the menu and continue to work underneath it as if it were not there (the inset shows how one person continues to draw underneath a menu raised by the other person). Their study involved people playing a game where one person would raise a menu over another who was drawing on the surface. They found that transparency helped reduce – but not eliminate – inefficiencies caused by interference.

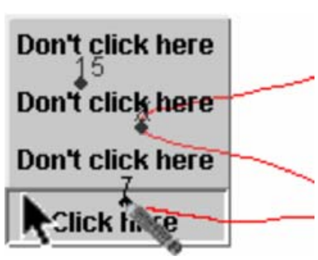

Shoemaker et. al. [16] reduce interference by exploiting personal shutter glasses, where two participants sharing a common display also see their own 'private' information within it. The first and second person's shutter glasses show only frames displayed during odd or even refresh periods respectively. While public information is shown across both display frames, private information particular to one person is shown only in the odd (or even) frame. For example, while both may see a common drawing surface and the marks made on it, person 1 also sees information that person 2 cannot, e.g., a menu raised by person 1 is not visible by person 2. Although this approach could eliminate interference by letting only the appropriate person see their transient window, it requires the use of special glasses that can be ergonomically and socially awkward. It may also make some actions mysterious: a transient window's invisibility means the other person may not know what the first person is doing [9][13].

Finally, several researchers have implemented PDAs as a secondary input / output device to augment the SDG display [6][11][14]. The idea is that individual PDAs can display person-specific information (e.g., a palette of operations) and that each person can use the PDA to control what appears on the main display. For example, Greenberg et al. [6] explain how a PDA shows a list of personal notes created by its owner, and how the owner can use the PDA to select particular notes to make them visible on the common SDG display. Similarly, Rekimoto [14] shows how a PDA can implement a pick and drop metaphor to move information between the common display and a PDA, and between PDAs. One disadvantage is the need and cost of PDAs. Another is that because participants cannot see what the others are doing on their PDAs, they lose information that could otherwise benefit fine-grained collaboration [9][13].

2.2 Spatial separation and partitioning

The approaches discussed above assume that people using SDG may often work atop each other, and as a consequence it is important to mitigate possible interference effects. But is this assumption correct?

One counter-argument is that people will spatially separate their actions through social protocol. Because they avoid one another, interference would be rare. To explain, Pinelle et. al. [13] argues that teamwork within a shared visual workspace is derived from the affordances and constraints of the shared space where collaboration occurs. Furthermore, they argue that there is a set of basic, core actions that happen in shared workspace collaboration. Called the *mechanics of collaboration*, these are the small-scale actions and interactions that group members must carry out in order to perform collaborative work while sharing space. One aspect of these mechanics involves the *coordination of shared access* to tools, to objects within the space, and to the space itself. People coordinate through mechanical actions such as: obtaining the resource (thus excluding others from using it), reserving the resource (by moving closer to it and explicitly or implicitly notifying others of their intentions), and protecting their work (by monitoring other's actions in the area and notifying others when problems are anticipated). They also transfer resources by handing off objects (through verbal or physical give and take) and by placing objects in particular locations and notifying others about the handoff. All the above serve to spatially separate actions, and to coordinate those moments of close interaction.

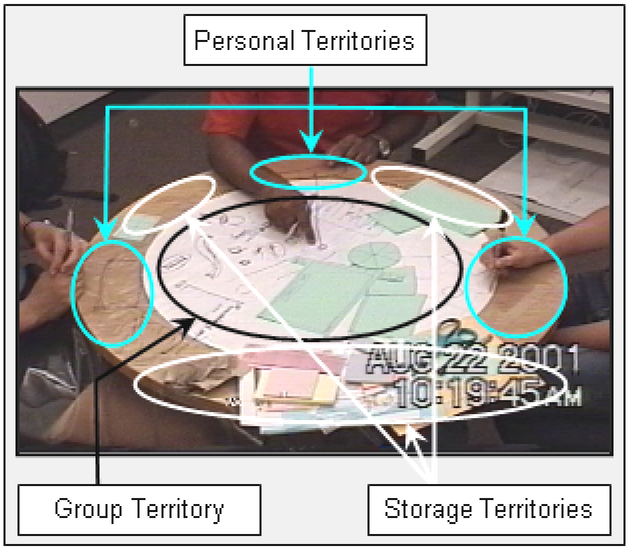

Figure 1. Territories in collaborative work, from [15].

Another view considers that spatial separation happens because people naturally partition the work space, and that individuals constrain their actions to disjoint partitions. Because of this partitioning of actions, interference would be rare. We do know that partitioning of our physical space is a natural human behavior, both in the large and in the small. For example, partitioning of the white-collar workplace is a very familiar concept. They are typically divided into interior offices, meeting rooms, coffee/lunch rooms and cubicles, and there are strong social norms about who does what in these spaces. In the broader context of our social environment, such partitioning is referred to as human *territoriality*. On the most basic level, territoriality serves an important role in organizing our interpersonal and group interactions to facilitate social order [1][21].

We also know that territorial behavior extends to physically shared workspaces. Several observational studies of tabletop collaboration have shown that people using traditional media (e.g., pen and paper) on a shared tabletop surface partition their workspace into several territories. For example, both Tang [19][20] and Kruger et. al., [10] found that people use the area immediately in front of them to define a personal space. More recently, Scott [15] further differentiated people's partitioning of the table surface into personal, group, and storage territories; Figure 1 illustrates the distribution of the areas created by 3 people working on a task over a round table. Scott argues that these territories help group members organize their collaborative activities. Because collaborators in physical spaces typically spend much of their time in different and disjoint areas, perhaps the actual potential for interference is small.

The problem is that collaborating in a virtual workspace with indirect input devices, such as mice, removes constraints and cues available in the physical workspace. This potentially increases the opportunity for interference. Yet perhaps the social and behavioural norms defining personal spaces and how they are negotiated may still be in effect, which would mitigate interference effects.

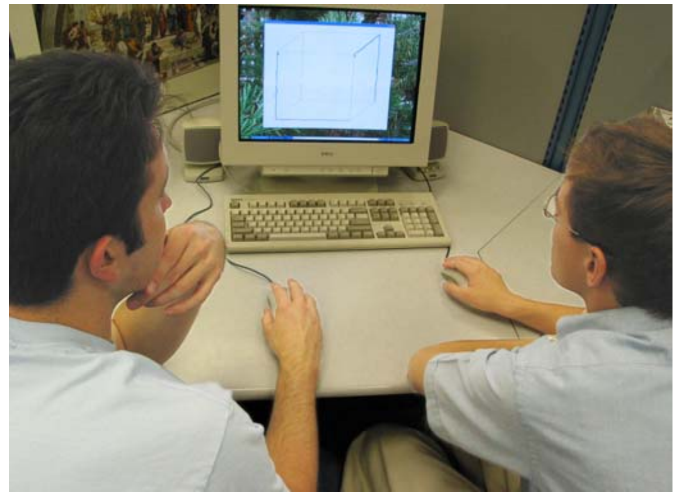

Figure 2. Our experimental set up.

3. A STUDY OF PARTITIONING IN SDG

We believe that understanding spatial separation and partitioning – and thus the realistic likelihood of interference in a shared virtual workspace – will provide valuable insights into the design of SDG interfaces. In order to determine whether spatial separation and partitioning emerges in a shared virtual environment, we investigated pairs completing a shared drawing activity in an SDG desktop environment, as described below.

Hypotheses. Based on our initial understanding of collaboration in SDG environments and results from a number of exploratory pilot studies, our experiment was designed to investigate four working hypotheses.

1. Participant pairs will move and draw simultaneously i.e., they will not exclusively turn-take as they perform their task.
2. Each participant will spatially separate his or her actions away from where the other person is working.
3. Each participant will constrain his or her work to a partition i.e., a portion of the workspace that has minimal overlap with the other participant's working area.
4. The way participants divide the workspace will depend primarily on the image semantics (the underlying structure suggested by the diagram), and secondarily on the participant's seating arrangement with respect to the drawing.

Participants. Forty-eight university students (41 males and 7 females), all paid volunteers, worked in self-selected pairs to complete a series of collaborative tracing and drawing exercises. Almost all partners knew each other well; most reported meeting each other several times a week.

Equipment. Two mice and a single keyboard were attached to a standard computer with a single upright 19 inch CRT display running at 1280 x 1024 resolution. Each person had their own mouse and a clear view of the display, as pictured in Figure 2. A custom single display groupware application, created atop the SDGToolkit [22], presented either a pre-defined image or an empty canvas to participants within an 1100 x 900 pixel workspace. The software let participants simultaneously draw atop these images with their mice (multiple cursors are visible), and let them advance to the next image once they had completed a trial. For example, Figure 3 shows an example of participants in the process of tracing a ‘cupboard’; their drawn lines are thick and blue, whereas the source image consisted of thin black lines.

Procedure. We first administered a pre-test questionnaire to collect demographic information and experiences with groupware. Next, participants were instructed to perform thirteen sequential drawing tasks with their partners. The pair was told that:

- the first eleven tasks would consist of tracing over an image of a line drawing that appeared on the screen with their mouse (the first three were practice trials),
- their goal was to completely trace these images as quickly and as accurately as possible,
- they could each draw on the image at the same time,
- only one trace was required, i.e., if one person traced a line, the other participant would not have to retrace it,
- when they judged the tracing task to be complete, they could advance to the next exercise, and
- the final two trials would consist of creating a free-form drawing (a windowed house and a car) on a blank drawing canvas.

The images used in the three practice trials are illustrated in Figure 4. These trials let participants become familiar with the task and with the simultaneous drawing capabilities of the system. The next eight tasks presented the images in Figure 5 (top two rows), where the order of presentation of the images was counter-balanced across the subject pairs. The final two trials presented a blank image for free form drawing. After completing these trials, participants were asked to fill out a post-test questionnaire asking them about their awareness of spatial separation and partitioning that occurred, and of any particular partitioning strategies used.

Data collection. During the trials the application software logged all participants’ mouse actions at a fine grain. Every 0.166 seconds, we recorded the state of each mouse: its ID (each mouse generated a unique ID), what button (if any) was pressed, the coordinates of its cursor, and a time stamp. With this data, we could reconstruct how people moved their mice, what they had drawn, and what concurrent actions were performed. We also captured all participants’ verbal comments through audio recordings, and took field notes of any salient events.

Categories of image semantics. During pilot explorations preceding this study, we observed that collaborators frequently divided the drawing space into non-overlapping areas of work. We noticed that in certain images, partitioning occurred due to a natural split presented by the image being traced. This split could be based, for example, by the image’s spatial layout, its component objects, or by line orientation. For example, the sofa image (Figure 5 top left) was drawn in several pilot trials using an *upper/lower* division, likely because it had many long continuous horizontal lines that would require more time to complete if each person only drew half of a line.

To validate this phenomenon, we created four categories of images to see how it would affect peoples’ partitioning behaviours (Hypothesis 4). Each column in Figure 5 shows a category, with two images per category.

1. ***Upper/Lower.*** The image naturally divides into two areas located above and below each other.
2. ***Left/Right.*** The image naturally divides into two areas located left and right of each other.
3. ***Inside/Outside.*** The image naturally divides into two areas, where one closed area is located inside the other.
4. ***Unknown.*** The division, if any, is uncertain.

The final free-form drawing tasks have no *a priori* division semantics. However, we wanted to see whether people would still divide the workspace, and if so, if those divisions reflected the semantics of the drawing as it was being created, and if participants verbally negotiated their partitioning.

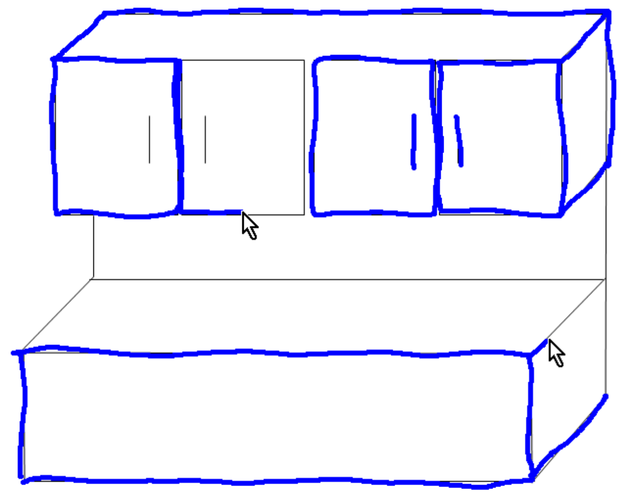

Figure 3. A typical tracing exercise

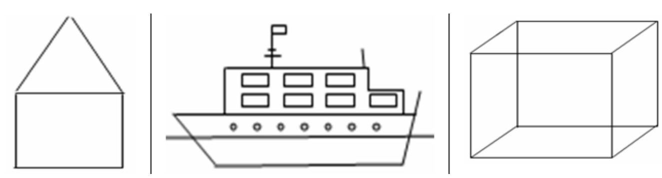

Figure 4. The trial images

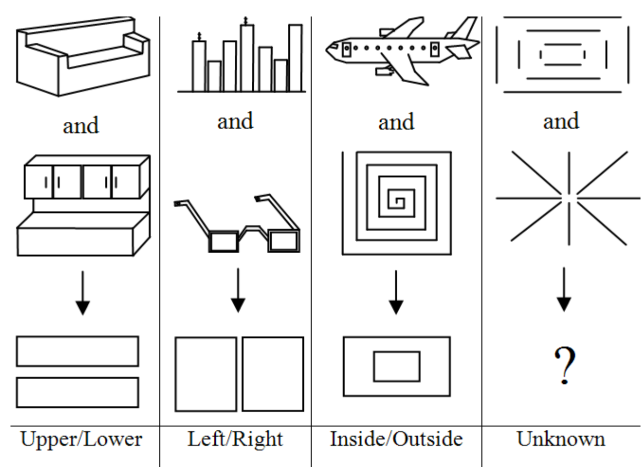

Figure 5. Image categorization.

Design justification. Two aspects of our experimental design warrant some discussion and justification.

First, we chose an SDG setting comprising multiple mice and a conventional monitor because we believe this presents a 'worst case' environment for interference. Indirect input (in contrast to a stylus or multi-touch surface) means that people will lack many of their normal shared workspace cues [13]. Similarly, the small size of the monitor (in contrast to a large projected display) means that people are more likely to compete for space. If we can show that partitioning occurs within this setting, then it almost certainly will occur in the other settings.

Second, we chose tracing and free-form drawing because they have properties common to many shared workspace activities. Tracing represents a simple well-constrained spatially-oriented task that lends itself to divide and conquer strategies while still allowing considerable degrees of freedom of how people would go about doing it. Free-form drawing represents a completely unconstrained creative activity that is representative of how people use a shared workspace for brainstorming ideas, for collaborative designs, and for other types of emergent activities whose outcome is unknown [4][19][20].

4. RESULTS AND DISCUSSION

4.1 Do participants interact simultaneously?

Our first hypothesis suggests that people will move and draw simultaneously i.e., they will not exclusively turn-take as they draw.

Data analysis. Recall that all mouse movements, button states and drawing coordinates are logged at .166 second intervals. By using this data, we can easily categorize people's actions and count the time when one or both mice are *still* (coordinates are not changing), when they are *moving* (coordinates are changing but the drawing button is not pressed), and when they are *drawing* (the mouse button is pressed). Table 1 expresses these categories and counts as an averaged percentage of the task time during the tracing, drawing tasks as well as an overall combined average, including standard deviations.

Results. Depending on how one interprets the data, there is a moderate to very high amount of simultaneous activity. The strictest interpretation of simultaneous action considers only the time when both partners are drawing, which is an average of 43% and 14% for the tracing and drawing task respectively (Table 1). However, a broader interpretation considers simultaneous actions as being the sum of the times when both partners are either drawing or moving (i.e., both draw + one draw/one move + both move). This gives an averaged time for simultaneous interaction of 87% and 63% for the tracing and drawing task respectively.

Our observations of how people performed the tasks informed our understanding of the differences between the tracing and drawing task. First, the 'both still' category reflects dead times just before and after each trial, as well as times when people are just talking to one another during the task. We see that times are much lower for the tracing task (3% *vs.* 11%), which makes sense because people doing free-form drawing need to do a bit more planning and coordination than when they trace. This matches our observations of what actually happened. Second, the times when one moves or draws while the other is still are also lower in the tracing task (10% *vs.* 25%). We explain this by our observations that a leader/follower role emerged during the drawing task: at the beginning of the trial one person would wait for the other to start moving and drawing, and then join in when they had a sense of what was going on. Third, perhaps the most interesting difference is that there are far greater moments when both are drawing when we compare the tracing to drawing task (43% *vs.* 14%), but fewer moments when both are moving (13% *vs.* 20%). We ascribe this difference to increased gesturing in the free-form drawing task, i.e., we saw one or both people gesture around the workspace as they discussed their task. This finding accords with other studies of how people frequently gesture over physical workspaces [19].

Table 1. Average portion of time where partners are still, moving or drawing (as a percentage of each trial)

	Tracing	Drawing	Overall
Both still	3% ± 3	11% ± 12	4% ± 7
One move, one still	6% ± 3	16% ± 7	8% ± 6
Both move	13% ± 6	20% ± 6	14% ± 6
One draw, one still	4% ± 3	9% ± 5	5% ± 4
One draw, one move	31% ± 12	29% ± 9	31% ± 11
Both draw	43% ± 15	14% ± 8	38% ± 18

In general, collaborators took advantage of the simultaneous multiple input capabilities provided in the SDG environment, where one person would move or draw at the same time as the other. Because the tracing task required less explicit planning, we saw quite a bit of simultaneous drawing. Because the drawing task required somewhat more planning and coordination, we saw less simultaneous drawing but more instances where one person was gesturing around the workspace as the other was also gesturing or drawing. Both cases substantiate that simultaneous interaction is a common and frequent event. Thus the potential for interference is high.

4.2 Do participants spatially separate actions?

The second hypothesis suggests that a participant will spatially separate his or her actions away from where the other person is working.

Data analysis. We measured special separation by calculating the cursor separation distance (in pixels) between participants during each trial.

Results. Table 2 summarizes the mean separation between cursors by task type, both in pixels and as a proportion of the 1100x900 pixel workspace this distance represents. Figure 6 plots the actual separation distribution over time as a histogram, where the total pixel separation across all trials and groups are calculated as a percentage of the total time and then split into ten equal bins. For additional detail, the inset in Figure 6 further partitions the 0-127 pixel bin into 5 bins. Both Figure 6 and Table 2 clearly show that participants' were typically interacting at quite a distance from each other; on average, 37.5% of the workspace away from each other (Mean=413 pixels, SD=124). While the average separation of their interactions was higher in the tracing trials compared to the free form drawing task, this difference is not significant (p=.22). From the inset in Figure 6, we see that on average cases of near misses is very small i.e., only ~3% of partner's actions are within 50 pixels of one another.

Table 2. Mean and proportional separation distance by task type (including standard deviations)

	Tracing	Drawing	Overall
Separation Distance (pixels)	446 ± 65	380 ± 119	413 ± 124
Proportion of Workspace (%)	40.5 %	34.5 %	37.5 %

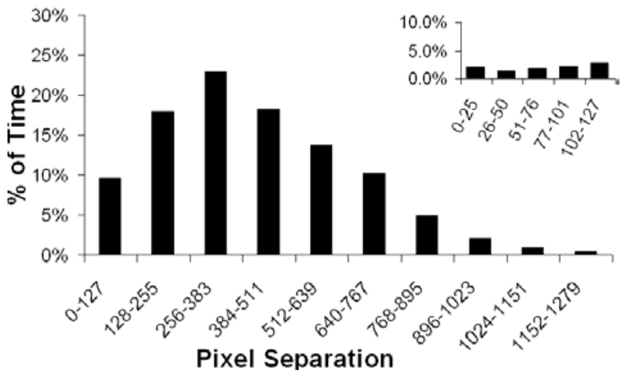

Figure 6. Pixel separation across all trials and all groups

4.3 Do participants partition the workspace?

The third hypothesis suggests that each participant will constrain his or her work to a partition i.e., a portion of the workspace that has minimal overlap with the other participant's working area.

Data analysis. In order to determine whether participants partitioned the workspace, we created visualizations of each participant's drawing activity from their logged data. These visualizations separated the mouse actions of each participant by reproducing each in a different colour and shade. Figure 7 depicts the visualization for a typical drawing, where we show drawing actions by the person seated on the left as a thick line, and by the person seated on the right as a thin line. Three people then independently judged whether the participants had partitioned the workspace into two discrete areas. Their decisions were recorded on coding sheets that identified whether partitioning occurred, categorized by image type (Table 3). For example, coders evaluated Figure 7 as showing clear partitioning.

We also categorized participant responses in the post-test questionnaire to see if they were aware of their partitioning behaviours and strategies (Table 4).

Results. The coding of partitioning within the visualizations are detailed in Table 3[1]. Results reveal that across all of the images, regardless of the pair or task type, participants overwhelmingly partitioned their drawing of the image (i.e., 222/233 = 95%). Inter-coder reliability was 100%; all three coders uniformly agreed on the presence or absence of partitioning.

Participants' post-test questionnaire responses confirmed their tendency to work in separate areas of the workspace. When asked to articulate how their drawing actions were coordinated, the

[1] Due to technical problems, coding could not be performed on 7 of the free form drawings, thus the total number adds up to 233 instead of 240.

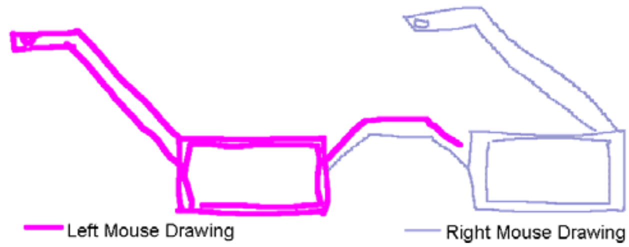

Figure 7. A drawing visualization from one of the trial tasks

Table 3. Partitioning occurrences for each image type

	Tracing									Drawing		
	Cupboard	Sofa	Glasses	Cityline	Spiral	Airplane	Star	Broken Boxes	**Total**	Car	House	**Total**
Partitioned	23	24	23	23	23	23	23	22	**184**	18	20	**38**
Unpartitioned	1	0	1	1	1	1	1	2	**8**	2	1	**3**

Table 4. Coordination strategies given by participants

Coordination Strategy	# of Remarks	Example Remark
Sides of the Screen	18	"We generally kept to our respective sides."
Opposite areas of the drawing	14	"Each of us would start on opposite sides of the picture."
Partner avoidance	8	"I drew where my partner wasn't drawing."
Coordinated but no explanation of strategy	5	"We collaborated well…" and "After the first few drawings we started working together quite well."
No effective Strategy	2	"We did not coordinate well."

majority of the responses could be grouped into the three main categories shown at the top of Table 4: (1) sides of the screen, (2) opposite areas of the drawing, and (3) partner avoidance. The first category refers to one person working on one particular side of the screen and his or her partner working on the other side, each side generally matching their seating position. The second category refers to people working on opposite sides of the drawing. The third category refers to people trying to work somewhere other than where their partner was currently working i.e., a simpler form of spatial separation that does not necessarily require partitioning. For the remaining responses, the fourth row in the Table 4 refers to remarks that mentioned that the pair coordinated their actions, although the partners did not (or could not) articulate the particular strategy they used. Finally, the fifth row refers to those few remarks that suggested that coordination did not occur.

The audio recordings and our field notes revealed that participants appeared to coordinate their drawing activities quite effortlessly, often with little to no verbal negotiation. Participants' responses reflected this observation with remarks such as: "I found it interesting that we became coordinated without any explicit attempts to do so, i.e., you do this, I'll do that."

4.4 What factors influence partitioning?

The fourth hypothesis suggests that the way participants partition the workspace will depend both on the image semantics and on the participant's seating arrangement with respect to the drawing.

Data Analysis. When coders identified visualizations as being partitioned, they also coded the dimension that partitioning occurred along i.e., left/right, upper/lower, etc. For example, coders all agreed on the left/right split in Figure 7. Again, inter-coding reliability was high (91%), where coders completely agreed on how 202 of the 233 images were partitioned[2]. We used this coding to count how often collaborators partitioned their interactions to match the semantic structure of the image (Table 5). We also analyzed field notes and responses to the post-test questionnaire to determine whether partners consciously used this underlying structure to help them coordinate their actions.

In order to investigate the influence of seating position, we combined the logged data with our field notes to determine where participants were seated in relation to their on-screen cursor actions. This information was then combined with the image coding to determine the correlations between partitioning and the physical location of each participant.

Results. As mentioned in the experimental methodology, the images used in the tracing task were based on the four image categories: *upper/lower*, *left/right*, *inside/outside*, and *unknown*. As expected, coders identified examples of partitioning corresponding to each of these categories. However, they also identified a *diagonal* category that describes a division of activity along a diagonal axis of the image, such as the upper-left/lower-right split shown in Figure 8. Consequently, *diagonal* was added as a new category in Table 5.

The coding results are shown by image type in Table 5. For the tracing task, we see that most pairs partitioned a given image in a similar manner. For example, 15/21 cupboard images followed an upper/lower division (as evident in Figure 3), while 22/23 cityline images followed a left/right division. However, the ambiguous structure of the *star* and *broken boxes* diagrams is also reflected in the greater variations in how pairs partitioned those images. We also see that the partitioning strategy adopted by most groups on an image did not always match the *a priori* categorization of that image. For example, 13/20 pairs partitioned the sofa left/right, even though the sofa is classified as upper/lower.

We believe these coding results show that images with clear underlying structures will strongly affect how the tracing task was divided. As shown in Table 5, the clear upper/lower structure of the cupboard, the *left/right* structure of the *glasses* and *cityline* images, and *inside/outside* structure of the *spiral* image were frequently leveraged by participants to help coordinate their

Table 5. Number of groups that performed each partitioning type

Task:		**Tracing**								**Drawing**	
Image Category:		**U/L**		**L/R**		**I/O**		**?**		**FF**	
Trial:		Cupboard	Sofa	Glasses	Cityline	Spiral	Airplane	Star	Broken Boxes	Car	House
Partitioning	**U/L**	**15**	4	0	1	0	5	8	0	1	**8**
	L/R	5	**13**	**23**	**22**	1	**13**	**15**	**8**	**7**	3
	I/O	1	1	0	0	**21**	3	0	5	5	5
	D	0	2	0	0	0	0	0	6	0	0

Legend: U/L = Upper/Lower, L/R = Left/Right, I/O = Inside/Outside, D = Diagonal, ? = Unknown Partitioning, FF = Free Form

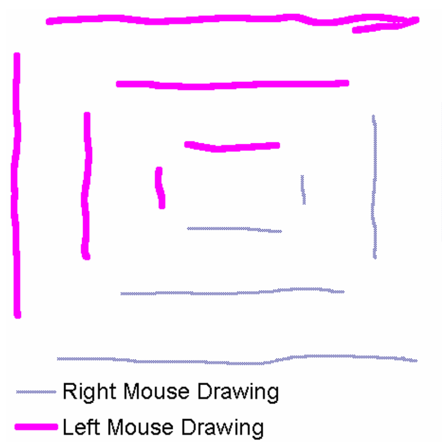

Figure 8. Diagonal partitioning

workspace partitioning. However, ambiguity in other image structures led to alternate partitioning strategies. For example, while we pre-categorized the *sofa* image as upper/lower on the basis of the long continuous horizontal lines, most of the pairs interpreted it as left/right. Similarly, most pairs interpreted the airplane by left/right rather than our inside/outside category.

When the image was not initially visible, such as in the free form task, the conceptual structure of the emerging drawing often helped collaborators coordinate their interactions. Although not captured in the coding, we noticed that people tended to use more complex, object-based partitioning strategies in the free form task, such as drawing a "roof" or "window" or "chimney" for the house, or a "wheel" or "body" for the car. Pairs could then choose to draw these objects following the partitions suggested by our standard categories in Table 5: this is how the free-form drawing in Figure 9 evolved from people drawing objects to using a (roughly) upper-lower division. However, as the numbers for the drawing task show, there is diversity in the way this is done. Post-test questionnaire responses confirmed that more communication was necessary in the free form trials, as illustrated by the comment: "Doing the [tracing] scenario we didn't talk much but when we did the free form we were forced to talk." Often, distinct

[2] Discrepancies between coders usually occurred when coders identified more complex partitioning strategies than simply 'one area versus another area'. For example, in the free form drawing trials the partitioning was often object-based, such as 'wheels of the car' or 'roof of the house.'

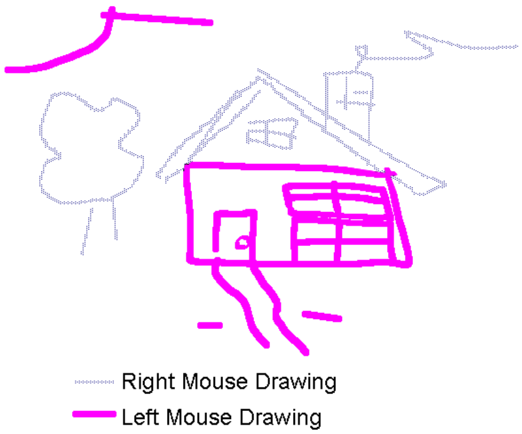

Figure 9. Partitioning in a free-form drawing

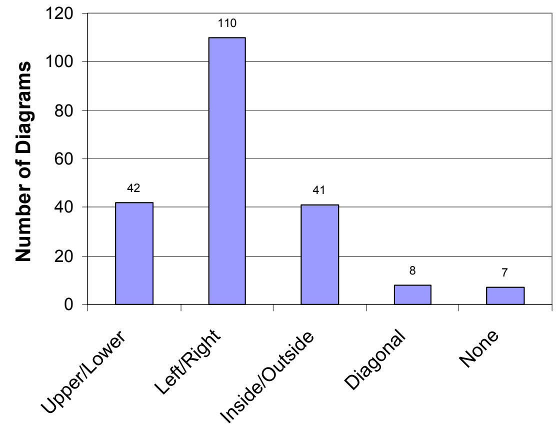

Figure 10. Number of diagrams partitioned by semantic type

leader and *follower* roles would develop during these trials to help partners coordinate their activities. We suspect these roles contributed to the reduced rate of concurrent drawing actions but increased gesturing actions that was seen in the free form drawing task (Table 1).

How did seating position influence what people did? As already mentioned, Table 5 shows that participants did not always use the expected semantic structure of an image to divide their tracing activities. Figure 10 presents the same data, but now organized by partitioning type. We see the predominance of the *left/right* partitioning strategy (47% of the total images) used by participants. Left/right was chosen 2.5 times more frequently than the 2nd most popular partitioning strategy. We then looked at the 110 tracings and drawings that were coded as a *left/right* partitioning. Of these, in 96 (87.3%) cases a person drew on the side of the display that directly corresponded with their seating position. Participant responses in the questionnaire confirmed this, as many mentionied partitioning by the use of "sides" (Table 4, 1st row).

This *left/right* match to seating position appeared to be the default strategy that most pairs relied on for smoothly coordinating their drawing activities. Collaborators typically seemed to abandon this default strategy only when the underlying structure of a diagram suggested a much more efficient division of labour. One participant's post-test questionnaire response, exemplifies this behaviour: "I took the left side of the screen. He took the right. With the *spiral*, we ran into each other so I went to the end and worked backward" (this *inside/outside* partitioning strategy was used in 21/22 of the *spiral* trials).

In summary, participants predominantly partitioned the workspace based on where they were sitting at the display, each claiming their respective "side". The underlying structure (visible or conceptual) of the drawing task also influenced their division of labour in the workspace. Obvious partitioning occurred more often in drawings that had unambiguous underling structure.

5. IMPLICATIONS TO PRACTITIONERS

We now know that interference is rare. We saw that collaborators naturally organize their interactions in the shared SDG workspace to minimize spatial overlap. They implicitly divide the space into partitions that more or less define a person's working region. Partition location is influenced by any underlying structure in the task and is also influenced by a person's seating position: one tends to favour the workspace area immediately to the front of them. From [13] we also know that people implicitly and explicitly negotiate coordination of shared access to the space and the things within it through the mechanics of collaboration. We also know that people will be still or will turn-take (i.e., work sequentially rather than concurrently) for some of their time (as seen in Table 1): interference cannot happen during these moments.

This has major implications to the design of interaction techniques for SDG, as described below.

Reveal spatial division in task semantics. Many tasks have semantics containing inherent spatial divisions. By judiciously organizing the workspace, we can exploit these spatial divisions so that people will use them to naturally partition their own activities. For example, many tasks are 'construction tasks' where people are assembling or modifying an artifact from various elements (e.g., drawings, blueprints, puzzle games, arranging nodes into a graph structure). To promote partitioning behavior and thus lessen interference, elements could be separated into logical piles located on different parts of the workspace, perhaps reflecting seating position. Similarly, if the actual artifact being constructed has different parts to it, then these can be located on the workspace in a manner that spatially separates one part from the other.

Exploit the physical location of users. We know that partitioning is strongly influenced by the physical location of the collaborators sitting at the display. Designers can leverage this tendency by appropriately positioning interface artifacts. If people are sitting side by side, artifacts should be separated to the left and right. If people are seated across a table, artifacts should be arranged on opposite sides. Frequently chosen artifacts could be copied so they are available on both sides. For example, the inset shows an SDG concept map application with two copies of a tool palette located for side by side work. Because these tool palettes will likely become 'personal' resources instead of group resources, they can even be designed to respond to only the person on that side.

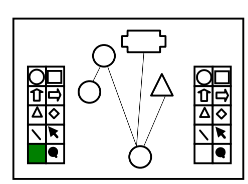

Reposition transient popup windows to exploit separation. Because people try to stay away from each other, the likelihood of interference occurring when a user invokes a transient window (such as a dialog box or pop-up menu) is small. Even so, there are certain design decisions that can be taken to minimize interference possibilities. First, instead of appearing at the center of the screen, transient dialog boxes resulting from an individual's action should appear near that individual's cursor. Second, pop-up menus and dialogue boxes can be further positioned so that they are as far away as possible from the other person while still being near the local cursor. For example, the inset shows two people working on the left and right side of the SDG workspace, and how their popups appear to the left (or right) of cursor owned by that person to increase separation.

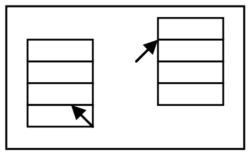

Let people move artifacts (such as tools) to their preferred locations. Instead of anchoring tools and artifacts to particular locations in a workspace, people should have the ability to move them. They will likely put it in a place in or near their current working area. This idea is already implicit in Local Tools [2], where individual tools are scattered around the workspace, and can be repositioned at the user's discretion. Repositioning facilitates easy access to system functionality, without getting in the way of others' workspace activities.

Another approach uses floating palettes, where people can move palettes of tools and objects around the workspace. For example, Tse and Stavness implemented an SDG version of Xerox PARC's Tool Glass [3] atop the SDGToolkit [22]. As illustrated in Figure 11, each person has their own tool glass, in this case a see-through palette of colors. Each person also has two pointing devices, one for each hand. The user positions the tool glass over an object with their non-dominant hand, and applies an action to the underlying object by 'clicking through' a palette icon with their other hand. Because these tool glasses are easily movable with the non-dominant hand, we expect people would keep them away from the other's space.

Divide the workspace into personal and group spaces. People often make use of both *personal spaces* and *group spaces* within a shared workspace (see Figure 1) [10][15][20]. Within a personal space, people typically perform individual work, which may or may not be later integrated into the group space. Within the group space collaborators usually work on or communicate about the group product. While the tasks that participants performed in this study were simple and only required the use of the group space, their partitioning behaviour within the group space suggests several implications for the design of personal spaces in an SDG workspace. First, the default location of a personal space should be based on the associated person's physical location at the display. Second, a personal space should be mobile. If the structure of the collaborative task suggests a partitioning different from the physical position of collaborators, a person may want to move the personal space near his/her claimed part of the group space. This idea is now being developed by Scott [15].

Mitigate interference effects, even if they are rare. While partitioning reduces the likelihood of interference in an SDG workspace, there remains some potential for obscuring someone else's view of the workspace when invoking a menu or other

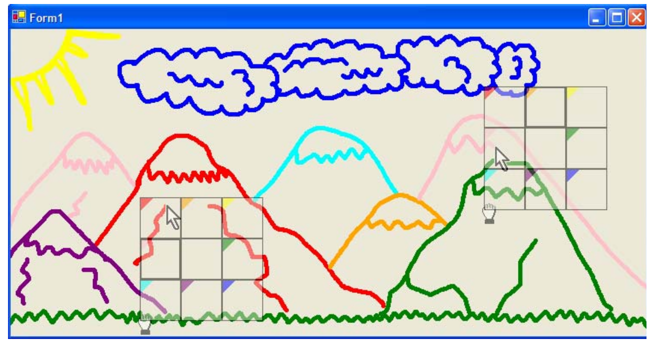

Figure 11. SDG Tool Glass, showing two users and their movable click-through tool palettes.

Figure 12. SDG Flow Menus, where rapid gestures eventually replace the visible menus

interface component. Thus judicial uses of techniques that mitigate interference effects are still valuable.

For our first example, we can modify Zanella and Greenberg's transparent menus to exploit *dynamic transparency* [8]. By default the pop-up menu would appear opaque. However, the menu would become semi-transparent if it were raised near someone else's working area, or if someone moved close to it. When that person moved away, the menu would return, gradually, to its opaque state. Therefore, readability of items on the menu is compromised only when a (rare) interference situation is imminent.

Our second example exploits gesture-based input techniques, such as those used in marking menus and flow menus [7]. These menus encourage users to learn how to select items through rapid gestures – if done quickly enough, the menu does not even appear and thus interference possibilities would be reduced even further. Figure 12, for example, shows our SDG implementation of Flow Menus, also implemented in the SDGToolkit [22]. We see two people, each with a hierarchical flow menu for selecting colors and line thicknesses. After people learn the gestural equivalent of selecting particular colors and thicknesses (shown as the stylized lines near the cursor), these menus would no longer appear.

6. CONCLUSIONS AND FUTURE WORK

We began this paper by questioning the actual risk of *interference* within Single Display Groupware. If the risk is high, then we would need to develop special purpose interaction components to mitigate its effects. If the risk is low, then SDG versions of conventional interface components could prove more suitable. To determine this risk, we studied how co-located people partitioned their collaborative drawing activities within a shared SDG workspace.

Our analysis of the study results contributes clear empirical evidence that people often partition their collaborative activities into separate regions of the workspace. We saw that partitioning is influenced by factors such as the semantics of the underlying task as well as seating position. These results suggest that people naturally avoid interfering with one another by spatially separating their actions in the workspace.

These findings can be used by SDG application designers to take advantage of the realistically low risk of interference. In particular, we contributed several ideas (some which we have implemented) that reconsider the design of shared virtual workspaces and of interaction techniques for co-located collaboration.

Of course, much is left to do. While we believe that our tracing and drawing tasks have properties that generalize to many types of SDG situations, there are undoubtedly tasks that do not fit this mold – these need to be studied further. We have also initially concentrated on only two collaborators, but it is easy to envision situations where several people are working around a shared display. In those situations, it is not clear what effect the increased numbers will have on partitioning behaviour and the likelihood of interference. Finally, we need to evaluate the design techniques we have suggested to see how well our ideas bear out in practice.

7. ACKNOWLEDGMENTS

Thanks to Interactions Laboratory members at the University of Calgary for their support. Funding was partially supplied by Natural Sciences and Engineering Research Council, Alberta Ingenuity, the Informatics Circle of Research Excellence, and the Federal Aviation Administration through the National Center of Excellence for Aviation Operations Research. We particularly thank reviewers of an earlier version of this paper, whose valuable advice we followed.

8. REFERENCES

[1] Altman, I. *The Environment and Social Behavior.* Monterey, California: Brooks/Cole Publishing Company, 1975.

[2] Bederson, B. B., Hollan, J. D., Druin, A., Stewart, J., Rogers, D., & Proft, D. Local Tools: An Alternative. *Proc. ACM UIST '96*, 169-170, 1996

[3] Bier, E., Stone, M., Pier, K., Buxton, W. and DeRose, T. Toolglass and Magic Lenses: The See-Through Interface. *Proc. ACM SIGGRAPH '93*, 73-80, 1993.

[4] Bly, S. A use of drawing surfaces in different collaborative settings. *Proc. ACM CSCW'99*, 250-256, 1988.

[5] Dietz, P. and Leigh, D. DiamondTouch: A multi-user touch technology. *Proc ACM UIST'01*, 219-266, 2001,

[6] Greenberg, S., Boyle, M. and LaBerge, J. PDAs and Shared Public Displays: Making Personal Information Public, and Public Information Personal. *Personal Technologies*, 3(1), 54-64, Elsevier, March 1999.

[7] Guibretière, F., Winograd, T. Flow Menu: Combining Command, Text, and Data Entry. *Proc. ACM UIST '00*, 213-216, 2000.

[8] Gutwin, C., Dyck, J., and Fedak, C. The Effects of Dynamic Transparency on Targeting Performance. *Proc. GI '03*, 105-112, 2003.

[9] Gutwin, C. and Greenberg, S. Design for Individuals, Design for Groups: Tradeoffs between power and workspace awareness. *Proc. CSCW'98*, 207-216, 1998.

[10] Kruger, R., Carpendale, M.S.T., Scott, S.D., & Greenberg, S. How People Use Orientation on Tables: Comprehension, Coordination and Communication. In *Proc. GROUP'03*, pp. 369-378, 2003.

[11] Myers, B., Stiel, H. and Gargiulo, R. Collaboration Using Multiple PDAs Connected to a PC. *Proc. CSCW'98*, 285-294, 1998.

[12] Nardi, B. *A Small Matter of Programming: Perspectives on End User Computing*. Cambridge: MIT Press, 1993.

[13] Pinelle, D., Gutwin, C. and Greenberg, S. Task Analysis for Groupware Usability Evaluation: Modeling Shared-Workspace Tasks with the Mechanics of Collaboration. *ACM TOCHI* 10(4), 281-311, 2003.

[14] Rekimoto, J. A multiple device approach for supporting whiteboard-based interactions. *Proc. CHI '98*, 18–23, 1998.

[15] Scott, S.D. Territory-Based Interaction Techniques for Tabletop Collaboration. *Conference Companion of UIST*'03, 2003.

[16] Shoemaker, G., Inkpen, K. Single Display Privacyware: Augmenting Public Displays with Private Information. *Proc. ACM CHI '01*, 286-293, 2001.

[17] Smart Technologies, Inc. *DViT Digital Vision Touch Technology: White Paper*, www.smarttech.com/dvit/DViT_white_paper.pdf, (2003)

[18] Stewart, J., Bederson, B. and Druin, A. Single display groupware: A model for co-present collaboration, *Proc ACM CHI'99*, 286-293, 1999.

[19] Tang, J.C. Findings from observational studies of collaborative work. *Int. J. Man-Machine Studies*, 34, 143-160, 1991.

[20] Tang, J.C. Listing Drawing and Gesturing in Design: A Study of the Use of Shared Workspaces by Design Teams. PhD Thesis, Stanford University, 1989.

[21] Taylor, R.B. *Human Territorial Functioning*. New York: Cambridge University Press, 1988.

[22] Tse, E., Greenberg, S. Rapidly Prototyping Single Display Groupware through the SDGToolkit. *Proc 5th AUIC Australasian User Interface Conference*, CRPIT Conferences in Research and Practice in Information Technology Series (28), Australian Computer Society Inc., p101-110, 2003.

[23] Zanella, A. and Greenberg, S. Reducing Interference in Single Display Groupware through Transparency. Proc. ECSCW '01, 339-358, 2001.