How People Partition Workspaces in Single Display Groupware

Edward Tse[1], Jonathan Histon[2], Stacey Scott[1] and Saul Greenberg[1]
Department of Computer Science[1], Psychology[2]
University of Calgary
Calgary, Alberta, Canada, T2N 1N4
[tsee, sdscott, saul]@cpsc.ucalgary.ca[1], jhiston@mit.edu[2]

ABSTRACT
Single Display Groupware (SDG) lets multiple people, each with their own input device, interact simultaneously on a single display. With two or more people potentially working in the same or nearby areas of the display, the actions of one could interfere with others, e.g., by raising menus and bringing tool palettes into areas where others are working. Interaction techniques could be used to mitigate the interference; however, other approaches might be more suitable if collaborators were to naturally partition their workspace into distinct areas when working on a particular task. To determine the realistic potential for interference, we investigated people performing a set of collaborative drawing exercises in a co-located setting, paying particular attention to the locations of their interactions in the shared workspace. We saw that spatial division occurred consistently and naturally across all participants, rarely requiring any verbal negotiation. Particular divisions of the space varied, influenced by seating position and image semantics. These results have several implications for the design of SDG workspaces, including the consideration of peoples' seating positions at the display, the use of moveable Local Tools and in-context menus, and the use of dynamic transparency to mitigate interference.

Author Keywords
Single Display Groupware (SDG), CSCW, Co-located collaboration, design implications

ACM Classification Keywords
H5.3. Group and Organization Interfaces: Computer Supported Cooperative Work.

INTRODUCTION
The majority of real time CSCW research has focused on how computers can support distributed collaboration. There are many potential benefits associated with these systems because distributed work interactions are difficult – if not impossible – without technological support. Yet the vast majority of work meetings, whether scheduled or casual, are between co-located people. Co-located interactions can also benefit from technological support since a computerized work surface can significantly enhance the capabilities of people working together.

In a physical workspace, people's use of space is constrained by the location of their arms and bodies. For example, people standing in front of a whiteboard turn-take or choose a side to work on simply because their bodies get in the way [10]. Yet these limitations disappear in an electronic workspace because people are controlling a virtual hand (the cursor) through a mouse. This environment enables people to work atop each other, and to easily reach any part of a large workspace.

While this ability to simultaneously work in the same area of the workspace increases interaction capabilities, it also introduces the possibility of *interference*, where one person's 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, obscuring another person's working area underneath the raised menu or dialog [13]. As a result some researchers have developed tools that mitigate interference effects. However, to our understanding no one has carefully examined whether people work concurrently within the same location in a co-located, shared workspace, or if they naturally partition their activities in a way that minimizes the potential for interference.

This paper investigates the natural behaviour of co-located collaborators sharing a virtual workspace. If workspace partitioning extends to the virtual workspace, it would have implications for the design of SDG applications, including the position and design of shared tools and objects, and the development of interaction techniques that recognize people's natural use of space.

RELATED WORK

Designing Shared Virtual Workspaces
The shared nature of an SDG workspace presents the potential for interference between collaborative actions. Several solutions have been proposed to help reduce the potential for interference. Bederson et al. [2] proposed the

Submitted to ACM CHI 2004

use of *Local Tools* where all application functionality is accessed through tools represented as separate icons positioned in the shared workspace. Zanella and Greenberg [13] investigated the effects of using *transparent interface components* in SDG workspaces. They found that transparency helped reduce the interference caused by the invocation of menus in the workspace.

Another approach to reducing interference has explored enabling users to share a physical display, while having access to private information via personal shutter glasses [8]. Although this approach could eliminate interference it requires the use of special glasses that can be ergonomically and socially awkward, as it introduces a visual barrier between the co-located collaborators.

Each of these approaches have the underlying assumption that people using SDG may often work atop each other and, thus, it is important to mitigate possible interference effects. But is this assumption correct? We believe that understanding the realistic likelihood of interference in a shared virtual workspace will provide valuable insights into the design of SDG interfaces.

Partitioning of Spaces

Partitioning of office workplaces is a very familiar concept. For example, workplaces are typically divided into interior offices, meeting rooms, and cubicles. 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, 11].

Territorial behaviour extends to physically shared workspaces as well. Several observational studies of tabletop collaboration have shown that people using traditional media (pen and paper) over a shared tabletop surface partition their workspace into several areas, or territories. Tang [9] found that people use the area immediately in front of them to define a personal space, while Scott [7] further differentiates people's partitioning into personal, group, and storage territories. These partitions seem to help group members organize their collaborative activities. Because collaborators in physical spaces typically spend much of their time interacting in different areas, the actual potential for interference is fairly minimal.

Collaborating in a virtual workspace with indirect input devices, such as mice, removes physical constraints from the shared workspace yet the social and behavioural norms defining personal spaces may still be in effect. In order to determine whether partitioning emerges in a shared virtual environment, we investigated pairs completing a shared drawing activity in an SDG desktop environment.

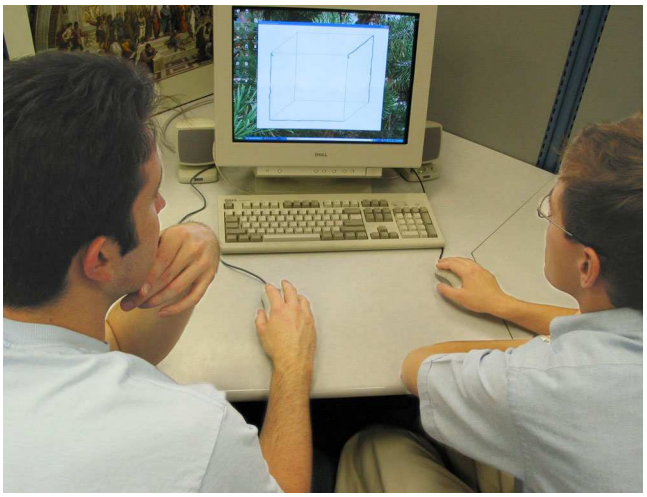

Figure 1. Our experimental set up.

EXPERIMENTAL METHODOLOGY

In this experiment, 48 university students (41 males and 7 females) worked in pairs to complete a series of collaborative tracing and drawing exercises. 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 three working hypotheses:

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

Equipment

Two mice and a single keyboard were attached to a standard computer with a single upright 19 inch CRT display. This was arranged on a desk so that each person had their own mouse and a clear view of the display, as pictured in Figure 1. A custom single display groupware application, created atop the SDGToolkit [12], presented either a pre-defined image or an empty canvas to participants. The software let participants simultaneously draw atop these images with their mice (Figure 2), and let them advance to the next image once they had completed a trial.

Data collection. The application software logged all participants' drawing actions, recording the time, particular movements and button presses of each mouse, and a corresponding user ID. A pre-test questionnaire was

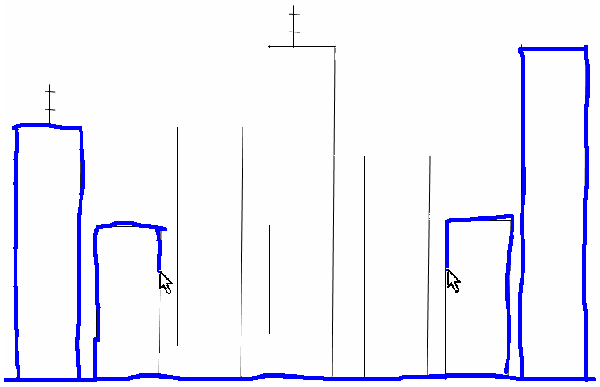

Figure 2. A typical tracing exercise.

administered to collect demographic information and previous experiences with collaborative computer software. A post-test questionnaire inquired about participants' awareness of any partitioning that occurred, or of any particular partitioning strategies used. All participants' verbal comments were captured through audio recordings. Finally, experimenters took field notes of any salient events.

Procedure

After completing consent forms and pre-test questionnaires, 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 (the first three of which were practice trials),
- their goal was to completely trace these images as quickly and as accurately as possible with their mouse,
- they could each draw on the image at the same time,
- only one trace was required, i.e., if one person traced over a line, the other participant would not have to retrace it,
- when they judeged the tracing task to be complete, they could advance to the next exercise,
- 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 3. 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 4 (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.

Images for Testing Hypothesis 3

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

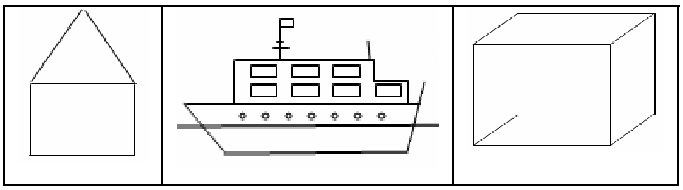

Figure 3. The Trial Images.

types of workspace 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 was often drawn in pilot studies using an *upper/lower* division, likely because it had many long continuous horizontal lines which 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 (using tracing drawings from the pilot study) to see if peoples' partitioning behaviour would be affected (Hypothesis 3).

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.

Figure 4 shows the images we created, with two images per category (top two rows) and how they are classified (bottom row). These images were then given to subjects after the trials.

Of course, many tasks are creative, where groups begin with a blank sheet of paper and a vague outcome. One well-studied example is the collaborative design activities of small groups [4]. Consequently, we included a free-form drawing task for the last two trials. Since there are no *a*

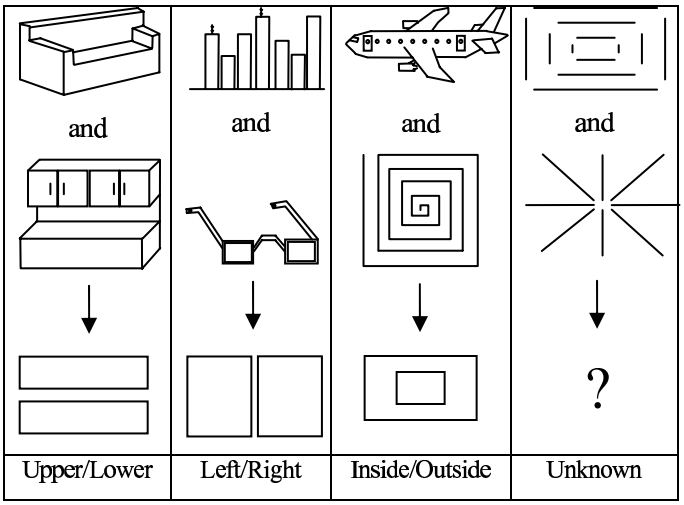

Figure 4. Image Generalization.

	Tracing	Drawing	Overall
Simultaneous drawing	41% (8.9%)	14% (6.9%)	36% (9.0%)

Table 1. Average percentage of time in each trial where drawing activity was concurrent, grouped by task type (standard deviations are shown in parentheses).

priori division semantics in this task, we wanted to see whether people would still divide the workspace, and if so, would those divisions reflect the semantics of the drawing as it is being created and would participants verbally negotiate the partitioning.

RESULTS AND DISCUSSION

Do participants draw simultaneously?

Data analysis. In order to assess how much concurrent drawing activity occurred, mouse movements and button states were logged at 0.166 second intervals. Each time sample that contained a mouse-drawing event for both participants was identified as an example of concurrent activity. The total concurrent drawing activity was determined by multiplying the number of identified concurrent samples by 0.166 seconds. In order to make consistent comparisons across different pictures and different groups, we expressed the results in terms of the percentage of time in each trial where concurrent drawing activity was observed.

Results. As reported in Table 1, over all trials, collaborators were drawing concurrently an average of 36% of the time. This suggests that participants were not always engaging in turn-taking behavior but were comfortable engaging in concurrent activity in the workspace.

We observed almost three times as much concurrent activity in the tracing task compared to the free form drawing task, as shown in Table 1. The data for a particular pair illustrates this tendency: they had a mean concurrency time of 12 seconds (37%) in the tracing trials, but had only 4.32 seconds (20%) of concurrent interaction in one of the free form drawing tasks and didn't draw concurrently at all in the other.

We believe that the reduced task structure in the free form drawing task contributed to the observed difference; this

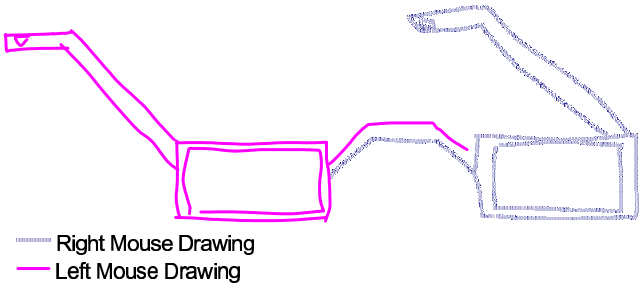

Figure 5. A Typical Drawing Visualization.

	Tracing									Drawing		
	Cupboard	Sofa	Glasses	Cityline	Spiral	Airplane	Star	Broken Boxes	**Total**	Car	House	**Total**
Partitioned	21	20	23	23	23	21	23	19	**173**	13	16	**29**
Unpartitioned	1	0	1	1	1	1	1	2	**8**	2	1	**3**
Discrepant	2	4	0	0	0	2	0	3	**11**	5	4	**9**

Table 2. Number of groups showing partitioning behavior for each image.

phenomenon will be discussed in more detail later in the paper. In general, it appears collaborators took advantage of the multiple input capabilities provided in the SDG environment, substantiating the potential for interference during collaboration in an SDG workspace.

Do participants partition the workspace?

Data analysis. In order to determine whether participants partitioned the workspace, data from the log files was used to create visualizations of each participant's drawing activity. These visualizations separated the mouse actions of each participant by reproducing each in a different colour and shade. Figure 5 depicts the result for a typical drawing. Three people 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, and the predominant dimension the partitioning occurred along.

We also calculated the average cursor separation distance (in pixels) between participants while they were both concurrently drawing in each trial. In order to understand which strategies, if any, participants used to coordinate their workspace activities, post-test questionnaire responses were synthesized and then categorized by comparable strategies.

Results. The results in Table 2 reveal that across all of the trials, regardless of the pair or diagram, the coding of the visualizations revealed that participants overwhelmingly partitioned the workspace (86.7%), with only a small number of images identified as unpartitioned (4.2%). Inter-coder reliability was high as all three coders agreed on the presence or absence of partitioning (and in the former case, on the type of partitioning used) in 213[1] out of the 233[2] coded visualizations (91.4%).

[1] There was 100% agreement of a discernable partitioning in 20 discrepant images, but raters disagreed on the exact type of partitioning used. These cases often involved more complex partitioning strategies than simply 'one area versus another area,' such as upper versus lower or left versus right. 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.'

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

Table 3. Mean separation distance by task type (standard deviations shown in parentheses) and the proportion of the total workspace size represented by this distance.

Analysis of the average cursor separation during concurrent drawing corroborates these coding results. Table 3 summarizes the results of this analysis, showing the mean pixel separation, by task type, as well as the proportional size of the 1100x900 pixel workspace this distance represents. These findings show that participants' were typically interacting at quite a distance from each other; on average, close to half the workspace away from each other (Mean=413, SD=124 pixels). The average separation of their interactions was higher in the tracing trials (Mean=446, SD=65 pixels) compared to the free form drawing task (Mean=380, SD=119 pixels). The latter result is likely due to the increased need for coordinated drawing interactions in the free form task because the exact structure of the diagrams were not explicitly presented.

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 majority of the responses could be grouped into three main categories, shown in Table 4: (1) sides of the screen, (2) opposite areas of the drawing, and (3) partner avoidance. The first category refers to one person typically working on one particular side of the screen and his or her partner working on the other side. The second category focuses on how each person typically works on opposite sides of the drawing. The third category refers to people's comments of generally trying to work somewhere other than where their partner was currently working.

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 (i.e., "you do this, I'll do that") attempts to do so."

What factors influence workspace partitioning?

Data Analysis. The coding results of the visualizations were used to investigate the relationship between the semantics, or underlying structure of the images and how collaborators partition their interactions in the workspace. Field notes and responses to the post-test questionnaire were also analyzed to determine whether people consciously used this underlying structure to help coordinate their actions.

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 side 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."

Table 4. Summary of coordination strategies given by participants.

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 coding for each visualization to determine the correlations between the partitioning that occurred 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*. The coding results identified examples of partitioning corresponding to each of these categories. In addition, a *diagonal* category emerged during the coding process. The *diagonal* category describes a division of activity along a diagonal axis of the image, such as the upper-left/lower-right split shown in Figure 6. The visualization coding results are shown by image type in Table 5.

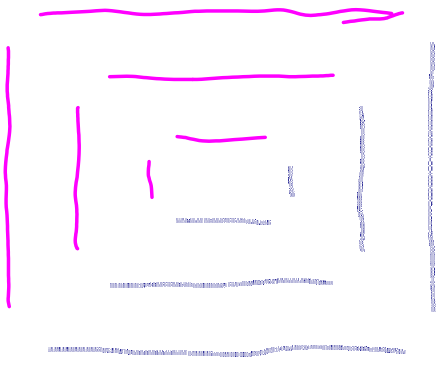

Figure 6. Diagonal Partitioning

[2] Due to technical problems, data was missing for 7 of the free form drawings.

Task		Tracing								Drawing	
Image Category		U/L		L/R		I/O		N/A		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,N/A = Unknown Partitioning, FF = Free Form

Table 5. Number of groups that performed each type of partitioning for each trial[3].

As expected, in the tracing trials the image category appeared to influence how partners partitioned the workspace. The coding results show that any natural underlying structure suggested by the diagram, such as the vertical separation of the top and bottom shelves in the *cupboard* image, appear to strongly affect how the tracing task was divided. The ambiguous underlying structure of the *star* and *broken boxes* diagrams is also reflected in the coding results. Partitioning of these images varied from *left/right* to *upper/lower* to *inside/outside*.

Typically, the clearer the underlying structure of the task, the more likely participants were to leverage this structure. As shown in Table 5, the clear *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 workspace actions. In contrast, the *upper/lower* structure of the *sofa* image that we had identified on the basis of the long continuous horizontal lines, and from the partitioning tendencies observed in the pilot trials, was not always taken advantage of by the participants. We believe that the numerous connections between the upper and lower areas enabled participants to easily apply a *left/right* partitioning scheme.[4]

When the diagram was not initially visible, such as in the free form task, the conceptual structure of the task often helped collaborators coordinate their interactions. For instance, partners frequently divided the *house* drawing into the roof and main body components, which is indicated by the predominance of *upper/lower* partitioning found across the *house* trials (34.8% of the trials).

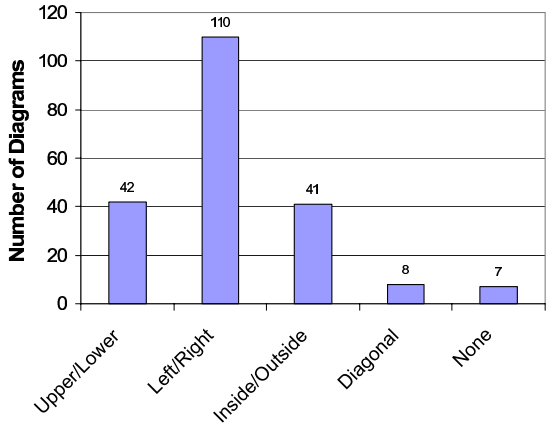

Figure 7. Number of diagrams that were partitioned by semantic type

Overall though, the visible structure of the tracing task simplified coordination of activities in the shared workspace. People tended to use more complex, object-based partitioning strategies in the free form task, such as drawing a "window", "wheel", or "chimney." This strategy was frequently used in the *car* trials most likely because the different components of a car needed to be tightly integrated. Post-test questionnaire responses confirmed that more communication was necessary in the free form trials, as illustrated by the comment: "Doing the scenario we didn't talk much but when we did the free form we were forced to talk." Often, distinct *leader* and *follower* roles would develop during these trials to help partners coordinate their activities. These roles contributed to the reduced concurrency that was seen in the free form task, as reported earlier in the paper.

While the underlying structure of the diagrams seemed to influence workspace partitioning, Table 5 shows that participants did not always use this structure to divide their workspace activities. Figure 7 presents the coding results organized by partitioning type. This figure shows the predominance of the *left/right* partitioning strategy (47.2% of the total images), which was used by participants over 2.5 times more often than any other strategy.

The overwhelming prevalence of the *left/right* partitioning in the shared workspace, along with the many participant responses mentioning the use of "sides" of the screen, led us to further investigate the influence of the seating position of the people at the shared display. For the 110 tracings and drawings that were coded as a *left/right* partitioning, 96 (87.3%) were found to have a direct correlation between the side of the display a person was sitting and the side of his or her on-screen drawing activity.

The *left/right* partitioning strategy appeared to be the default strategy that most pairs relied on for smoothly

[3] Instances with no partitioning and discrepant cases are not included in this table.

[4] Similarly with the airplane image. Line continuity analysis and pilot testing suggested that the most efficient approach to tracing the airplane would be to use an *inside/outside* partitioning.

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 comment describes the typical *inside/outside* partitioning strategy used to trace the *spiral* diagram, the strategy that was used in 91.3% 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. Furthermore, when components of the diagram needed to be tightly integrated people tended to utilize turn-taking more and worked closer together in the workspace.

IMPLICATIONS

The results of this study suggest that collaborators in an SDG workspace can be expected to naturally organize their interactions to take advantage of any underlying structure in the task that minimizes any overlap in the shared workspace. Designers of SDG workspaces can leverage this natural tendency by analyzing shared workspace tasks and designing them to facilitate natural partitioning.

Consider the physical location of users. Partitioning is strongly influenced by the physical location of the collaborators sitting at the display. Designers can leverage this aspect of partitioning tendencies by appropriate positioning of multiple interface components or task activities. Further investigation using more than two participants needs to be performed in order to determine how this implication scales with the number of collaborators or the orientation of the display.

Consider using small, in-context menus. Participants

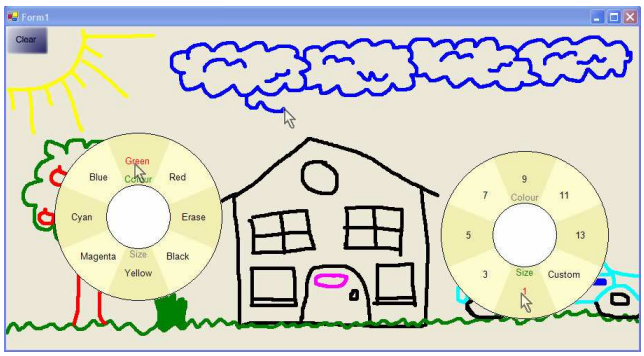

Figure 8. SDG Designers could leverage the natural division of space with the use of in-context flow menus.

partitioned the workspace into non-overlapping, distinct, work areas. Therefore, there is minimal potential for interference if a user invokes a menu in the area in which they are currently working. SDG designers can expect that, in general, small, in-context menus, such as, standard pop-up menus, Tool Glasses [3], or Flow Menus [5] invoked by one collaborator should rarely interfere with any other users. Also, Local Tools [2], which can be placed on the workspace in or near a user's current working area, would facilitate easy access to system functionality, without getting in the way of others' workspace activities.

Consider dynamic transparency. 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 interface component. *Dynamic transparency* [6] could be used to compensate for these situations. In the context of SDG applications, dynamic transparency could be used to lessen the impact of interface components that appear over someone else's activities. For example, by default an interface component would appear opaque, except when it would obscure someone else's current working area, or if someone else moved into that area. While someone is working beneath the component it would be semi-transparent; then, when the region became clear, the component would return, gradually, to its opaque state. Therefore, readability of items on the interface component is compromised only when an interference situation has occurred.

Designing personal spaces in a shared workspace. People often make use of both *personal spaces* and *group spaces* within a shared workspace [7]. 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 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.

CONCLUSIONS AND FUTURE WORK

We have explored collaborators' interaction behaviours in an SDG workspace to determine the realistic potential for interference in such a shared virtual environment. Our contributions include empirical evidence that people often partition their collaborative activities into separate regions of the workspace. As well we have shown that several

factors influence this workspace partitioning, including the seating position of users and the underlying structure of the collaborative task.

We have also shown that when the underlying structure of the collaborative task is ambiguous, participants need more explicit efforts to coordinate their actions. For example in the free-form drawing task, explicit communication increased between participants and they often developed *leader* and *follower* roles.

We have also presented several implications based on these findings that can be applied to the design of shared virtual workspaces for co-located collaboration. These implications can be used by SDG application designers to take advantage of the realisticly low potential for interference.

The next stages of this research will be to incorporate these insights into our own designs of workspaces in SDG applications. Also we plan to explore collaborative interactions involving more than two participants as well as SDG environments beyond the desktop, such as tabletop computers and wall-displays. Before we can provide people with tools that complement their co-located collaborative activities we first need to further our understanding of the strategies that people rely on to coordinate their interactions when working together.

ACKNOWLEDGMENTS

We would like to thank the members of the Interactions Laboratory at the University of Calgary for their thoughtful support. We would also like to thank the Natural Sciences and Engineering Research Council (NSERC), Alberta Ingenuity, the Informatics Circle of Research Excellence (iCore), and the Federal Aviation Administration through the National Center of Excellence for Aviation Operations Research (NEXTOR), for their financial support.

REFERENCES

1. Altman, I. (1975). *The Environment and Social Behavior*. Monterey, California: Brooks/Cole Publishing Company.
2. Bederson, B. B., Hollan, J. D., Druin, A., Stewart, J., Rogers, D., & Proft, D. (1996). Local Tools: An Alternative. *Proc. ACM UIST '96*, 169-170.
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.
4. Bly, S. (1988) A use of drawing surfaces in different collaborative settings. *Proc. ACM CSCW'99*, 250-256.
5. Guibretière, F., Winograd, T. (2000) Flow Menu: Combining Command, Text, and Data Entry. *Proc. ACM UIST '00*, 213-216.
6. Gutwin, C., Dyck, J., and Fedak, C. (2003) The Effects of Dynamic Transparency on Targeting Performance. *Proc. GI '03*, 105-112
7. Scott, S.D. (2003). Territory-Based Interaction Techniques for Tabletop Collaboration. In *Conference Companion of UIST'03*, in press.
8. Shoemaker, G., Inkpen, K. (2001) Single Display Privacyware: Augmenting Public Displays with Private Information. *Proc. ACM CHI '01*, 286-293
9. Tang, J.C. (1991). Findings from observational studies of collaborative work. *Int. J. Man-Machine Studies*, 34, 143-160.
10. Tang, J.C. (1989) Listing Drawing and Gesturing in Design: A Study of the Use of Shared Workspaces by Design Teams. PhD Thesis, Stanford University, 1989.
11. Taylor, R.B. (1988). *Human Territorial Functioning*. New York: Cambridge University Press.
12. Tse, E., Greenberg, S. (2003) *Rapidly Prototyping Single Display Groupware through the SDGToolkit.* Report 2003-721-24, Department of Computer Science, University of Calgary, Calgary, Alberta Canada. March.
13. Zanella, A. and Greenberg, S. (2001) Reducing Interference in Single Display Groupware through Transparency. *Proc. ECSCW '01*, 339-358