Single Display Privacyware: Augmenting Public Displays with Private Information

Garth B. D. Shoemaker and Kori M. Inkpen
EDGE Lab
School of Computing Science
Simon Fraser University
Burnaby, BC V5A 1S6, Canada
garths@cs.sfu.ca, inkpen@cs.sfu.ca
http://www.edgelab.sfu.ca

ABSTRACT

The research area of Single Display Groupware (SDG) confronts the standard model of computing interaction, one user working on one computer, by investigating how to best support groups of users interacting with a shared display. One problem that has arisen in SDG research concerns access to private information. Previously, private information could not be displayed on a shared display, it could only be accessed on external devices, such as private monitors or Personal Digital Assistants (PDAs). This paper discusses Single Display Privacyware (SDP), an interaction technique that allows private information to be shown within the context of a shared display. A description of the hardware and software components of our prototype SDP system is given, as are the results of a user study performed to investigate users interacting in the environment. Conclusions concerning future research in the area of SDP are discussed.

Keywords

Single Display Privacyware (SDP), Single Display Groupware (SDG), privacy, awareness, collaboration, CSCW, CSCL

INTRODUCTION

Research in the area of Single Display Groupware (SDG) investigates how to best support groups of users collaborating around a shared display. SDG research is based on the belief that the physical proximity of users, combined with the use of a shared display, allows users more natural and effective communication than either collaboration with separate displays or remote collaboration. The use of a shared display and the physical proximity of group members closely models our group interactions with shared artifacts in the physical world.

While SDG systems are good at supporting collaboration around public artifacts, they do not naturally support access to private information. If users of a shared display wish to refer to private information, they are normally required to use a different display. Any information shown on the public display is visible to all users, and therefore cannot be private.

Several recent SDG research projects have addressed the issue of privacy through the use of Personal Digital Assistants (PDAs) networked with a shared display. Such systems provide each user with a PDA which is used to view and manipulate private information, or to interact with the shared display. While the ability to display private information in such a manner is valuable for many reasons, there are drawbacks and limitations. One drawback of the technique is that users are required to split their attention between multiple displays. The requirement that users monitor and interact with multiple displays can result in unnecessary cognitive overhead. One limitation of the technique is that private information cannot be shown in the context of related public information. Mixing the display of public and private information could be used to highlight relationships between the two.

This paper discusses a novel interaction technique, known as Single Display Privacyware (SDP), that takes a different approach towards allowing access to private information in conjunction with shared displays. The interaction technique dictates that private information be shown on the shared display, instead of on a separate display, as with other techniques. In order to keep the information private, either the display output must be filtered or the display surface augmented in a manner such that each user can only see private information appropriate for him or herself. All users share the public information, but each user also has private access to private information.

The following section summarizes important research in the area of SDG, and then focusses on research dealing with privacy and awareness issues. SDP and the problems it addresses are then discussed in detail. Following that, a description of the hardware and software components of our prototype SDP system is given. The results of a user study that was performed to explore users' interactions with our prototype system are then discussed. Finally, conclusions are drawn concerning SDP's potential, and issues that need to be addressed in future research are discussed.

RELATED LITERATURE

Single Display Groupware

Research in the area of SDG investigates how to best support multiple users working on a shared computer display. The research area of SDG was first introduced by Stewart [16], who later gave the subject a more formal treatment [15]. Stewart's introduction, however, did not pre-date all work in the area of SDG. Many researchers have investigated groupware systems that support co-located users.

One aspect of SDG research concerns the technical problems associated with supporting the interactions of multiple users working on a shared display. Because of the dominant usage model of "one person per computer," most computer platforms are not designed to accept independent input from multiple mice or multiple keyboards. Making it possible for multiple users to interact simultaneously with a computer system is a significant problem. Research that confronts this problem includes that of Bier and Freeman [2], who's MMM toolkit supported input from multiple mice. Also in this category is the research of Bricker [3], and Hourcade and Bederson [8]. Bricker's research introduces the concept of Cooperatively Controlled Objects, while Hourcade and Bederson's work presents a low-level Java framework for the development of SDG systems using arbitrary input devices.

Another important aspect of SDG research deals with the effects on behaviour when users are working in an SDG environment, as opposed to a distributed groupware environment, or a single-user environment. While it is hypothesized that working in an SDG environment can improve inter-personal communication and awareness, there are almost surely other potential effects of working in such an environment that have not been identified. Studies by Inkpen [10, 9] and Scott [13] have investigated the behaviour of school children working under different collaborative conditions. Several beneficial effects to working in SDG environments were identified.

Privacy and Awareness

Privacy and awareness is an important topic of research for all groupware environments. Awareness, individual group members' understanding of what other group members are doing, has a large impact on the efficiency and coordination of the group when working on a shared task. Privacy, limiting the availability of information to a single user, serves to reduce the level of group awareness. This can often be beneficial. In his research on distributed groupware systems, Gutwin observes that an excess of awareness information can result in awareness overload [7]. Often, when large amounts of information is presented, users have trouble discerning between useful information and unimportant information. Gutwin also identifies a tradeoff that must be made when designing groupware systems. He argues that an increase in group awareness is accompanied by a decrease in the power of each individual user [6]. For example, if users are allowed to freely navigate a workspace, independent of other users, the individual user has significant power. However, group awareness suffers, as each user no longer necessarily knows the location of other users. The power of individual users increases at the cost of group awareness, and group awareness increases at the cost of individual user power.

Several groups have begun to look at SDG systems that allow for privacy. These systems generally involve the use of Personal Digital Assistants (PDAs) as private displays for each user. The PDAs are networked with the shared display, and information can be passed back and forth. Notable among these systems are those discussed by Rekimoto [12], Greenberg [5], and Myers [11]. Rekimoto's novel interaction technique, pick-and-drop, allows users to use a special stylus with which they can "pick-up" and "drop" information between the shared display and a private PDA. Greenberg discusses a system for use by members of a group who occasionally meet and must coordinate information that has been gathered outside of the meeting. Myers has developed a suite of tools under the "Pebbles" moniker, that allow users to manipulate information on a shared display using a PDA. The common thread between these three research projects is that the systems investigated allow users to access and alter private information on a private display, while public information is available on a shared display.

A preliminary discussion of Single Display Privacyware was previously presented by Shoemaker [14]. The issues involved with displaying private information on a shared display were briefly outlined, and an early ancestor of the prototype system presented in this paper was discussed.

Privacy in Virtual Reality

Researchers have also investigated systems that provide privacy support in virtual reality scenarios. Agrawala [1] introduced the Two-User Responsive Workbench, a collaborative system providing independent stereo views for two users. Other works by Butz [4] and Szalavari [17] also investigate VR systems that support independent views. Szalavari's work explores the importance of privacy in gaming environments, and Butz's work discusses different approaches towards privacy management. These papers discuss privacy among co-located users, but do not deal with issues specific to users in an SDG environment. As well, while they present intricate implementations of example systems, no user studies are discussed.

CONCEPT AND MOTIVATION

This section will first discuss how SDP is defined, and then will examine problems exhibited by SDG systems that can potentially be alleviated or solved by the use of SDP systems.

Single Display Privacyware

The concept of Single Display Privacyware (SDP) is a direct extension of the Single Display Groupware (SDG) concept. Stewart's definition of an SDG system requires that it support a shared user interface, that there be shared user feedback, and that navigation be coupled between users (Figure 1). SDP systems loosen the last two requirements. Through the introduction of privacy, essentially a private output channel, users need not necessarily share feedback, and navigation is not necessarily coupled between users (Figure 2). In addition to these qualities, an SDP system must have the ability to arbitrarily intermix private and public information anywhere on the shared display surface. Public information is any information that is visible to all members in the group using the display. Private information is any information that is visible to a subset of the members of the same group. The means by which the system displays the private information is undefined. The users may wear glasses that filter display output by polarization or frame interleaving, the display may emit light directionally towards individual users, or users may wear see-through head-mounted displays that overlay private information on the shared display. Any of these systems would qualify as SDP systems.

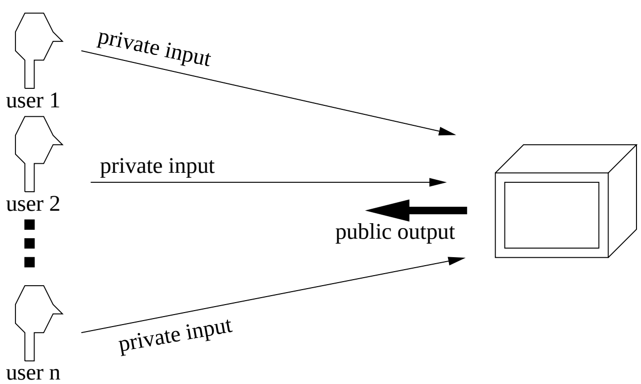

Figure 1: I/O channels in Single Display Groupware.

The definition of SDP leads to the definition of different types of private information that can be displayed on a shared display. The first kind of private information is "contextually private" information. This is information that is private to one user, but is significant in its physical proximity to some public information. Examples of contextual private information would be cursors and contextual menus. Cursors only hold meaning if they have a position relative to other onscreen objects. Having a cursor on a private PDA would be useless if you wanted to click on an object residing on a shared display. Similarly, contextual menus are only meaningful if associated with a particular onscreen object. It is important that they exist in their context, otherwise they are not of much use. The second kind of private information is "private area" information. This encompasses situations in which a particular area of the display is dedicated to private information (i.e. one area of the display is private for all users). An example of this would be a collaborative system that provides an area in which a private notebook can be displayed. A special case of the "private area" scenario is if the private area covers the entire display. In this case each user sees entirely different content on the display. While this may not be ideal for collaborative tasks, it may be suitable in other scenarios, such as competitive gaming.

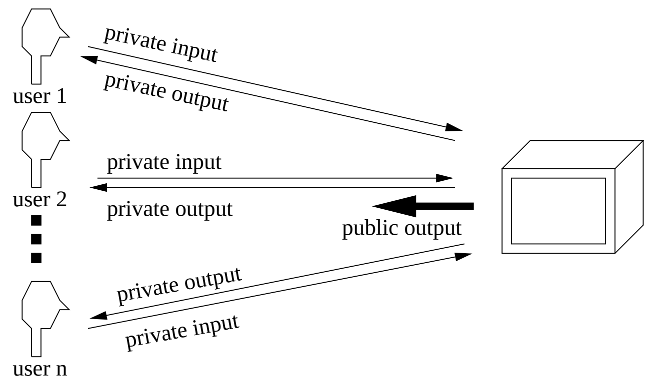

Figure 2: I/O channels in Single Display Privacyware.

The Screen Real-Estate Problem

The issue of screen real-estate is a key problem for SDG systems. When multiple users share a display, it is necessary that information widgets for every user be displayed. This may include tool palettes, cursors, menus, and home areas. Many of these widgets need to be duplicated for each user. As a result of this duplication, it is easy for the display to become overly cluttered. There is usually a hard limit on the available display area, and this can easily be met if several users are sharing a display.

The screen real-estate problem can be alleviated by using privacy techniques. Many of the widgets that exist for each user are user specific. They need not be visible to other users, and can therefore be made private. Once private, these widgets only occupy space for one user, making that space available for other uses by other users. The type of privacy support used to combat screen real-estate issues can vary. Cursors and menus might employ contextual privacy, for example, while dedicated private areas might be defined for tool palettes and home areas.

The Awareness Overload Problem

An issue closely associated with the screen real-estate problem is that of awareness overload. Group awareness is important when working in a collaborative environment, but too much awareness information can confuse a user. Distributed groupware systems usually suffer from a lack of awareness information, but SDG systems are more likely to suffer from a surplus of awareness information. Limiting the amount of awareness information presented to a user can improve the user experience. In addition, as observed by Gutwin [6], limiting group awareness in SDG systems may also increase individual user control.

SDP systems allow the level of awareness to to be controlled by either hiding or showing awareness widgets to specific users. Awareness information vital to effective collaboration can be broadcast to the whole group of users, while more personal awareness information will be shown to one user. This allows for a balance to be reached, where users have access to the information they need, but are not overwhelemed by unneeded awareness information.

The Privacy in Context Problem

In many situations it is important that users of a shared display see user-specific information within the context of the public information. An example of this is given by Tani [18], who investigated workers in industrial control room environments. Workers in control rooms monitor system status on a large shared display that holds general contextual information. When an event of interest occurs, individual workers must refer to specific information regarding system state. This information is generally different for each worker. Even though this information has strong contextual relevance, workers must access the information on small private displays. Displaying all the information for every worker on the shared display would be too confusing.

SDP provides a possible solution for the privacy in context problem. Again considering the control room scenario as an example, some of the private information for each worker could be displayed in the context provided by the large shared display. Each worker would only see the private information directly relevant to his or her specific job.

IMPLEMENTATION

A prototype system was implemented for use in an investigation into the feasability and usefulness of displaying private information on a shared display. The prototype system uses a combination of technologies that is only one of several possible for achieving privacy on a shared display.

Hardware

There are two main components to the hardware of the prototype system. The input component of the system allows two users to input simultaneously into the computer, while the output component of the system allows private information to be mixed with public information and displayed to each of the users.

The input component of the system consists of two Universal Serial Bus (USB) mice connected to the USB bus of the computer. Accessing the USB bus through the DirectInput API allows signals from the two mice to be accessed separately. Using this information to track the individual mice, two cursors can be simulated in software. The result is that each user has an independent cursor controlled by an independent mouse.

The output component of the system allows private information to be shown on the shared display to each of the users. The hardware consists of a StereoGraphics CrystalEyes system that has been adapted to produce private output instead of stereo output. The standard CrystalEyes system consists of two main components. First, the “synch-doubling emitter” fits between the monitor and video card on the computer. The emitter alters the signal from the video card so that the display refreshes at twice the normal rate. A byproduct of this is that the video image is “stretched” to be twice as tall as normal. Every odd-numbered display refresh frame shows a stretched version of the top half of the display image, while every even-numbered display refresh frame shows a stretched version of the bottom half of the display image. The second component of the CrystalEyes system is the CrystalEyes glasses. These glasses have liquid crystal element lenses that can alternate from opaque to transparent states very quickly. The glasses plug into the synch-doubling emitter and synchronize with the display refresh. The lenses of the glasses alternate from opaque to transparent states at the same rate as the display refresh.

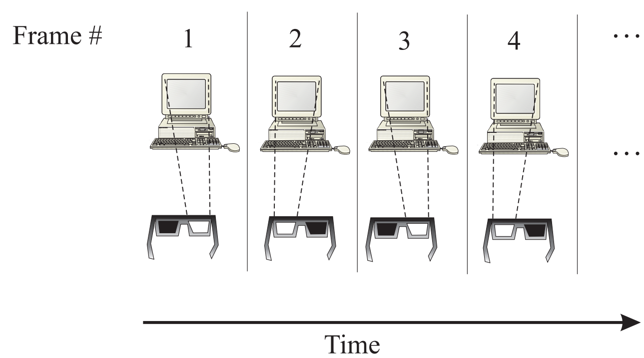

Figure 3: Shuttering sequence for producing stereo images.

The normal operation of CrystalEyes glasses allows one user to view a stereographic image on a normal computer display. The left and right lenses of the glasses shutter between opaque and transparent states, synchronized with the display refresh. At any moment in time, one lens is opaque and the other is clear, as shown in Figure 3. The result is that one eye sees all the odd-numbered refresh display frames, while the other eye sees all the even-numbered refresh display frames. If display content is drawn so that the even and odd display frames show a slightly different perspective of the same scene, appropriate for each eye, the user perceives a three dimensional stereo view of that scene.

For our prototype, we altered the normal operation of the CrystalEyes glasses in order to provide privacy for two users, instead of a stereo view for a single user. First, two sets of CrystalEyes glasses were attached to the synch-doubling emitter. Each of these glasses were then altered. One pair of glasses was altered so that both lenses opened during odd-numbered refresh display frames, and closed during the other display frames. The other pair of glasses was altered so that both lenses opened during even-numbered refresh display frames, and closed during the other display frames. Thus, when two users are wearing these glasses, each user sees the display output during different refresh display frames, as shown in Figure 4. One user sees odd-numbered frames, and the other user sees even-numbered frames. Public information is drawn on both even-numbered and odd-numbered frames, while private information for each of the users is only drawn on the appropriate frames. This technique allows for the realization of an SDP system. Two users can work on a shared display, with private information shown in the context of public information.

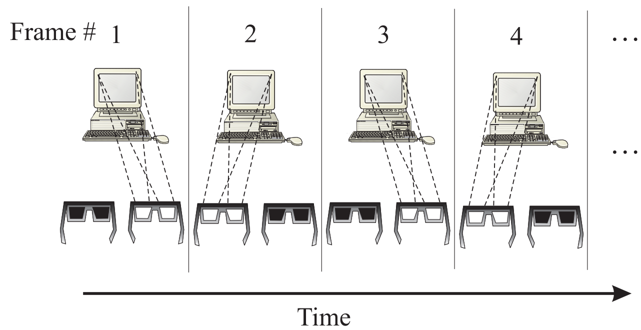

Figure 4: Shuttering sequence providing privacy for two users.

Software

The software component of the prototype system that we developed is a collaborative construction environment using LEGO-like blocks, shown in Figure 5. Two users are given the task of constructing three different predetermined block structures in a shared work area. The construction area shows a perspective view of a surface on which blocks can be placed by either user. Users can be in one of two modes: place mode, or edit mode. Place mode allows users to drop blocks in the construction area or change the properties (colour, orientation) of the block they are about to place. Edit mode allows users to delete or change properties of blocks that have already been placed. Changes to the properties of blocks are done through contextual menus that pop-up over the block being edited. The instructions area of the workspace shows step-by-step instructions for building each structure. Each structure has six construction steps. The instruction area can be toggled to show the instructions for any one of three different structures. Buttons along the bottom of the workspace are used to change between the two modes, or to rotate the view of the construction area and instructions.

For study purposes, two versions of the software were developed. The "public version" of the software operates as a normal SDG system. Both users see exactly the same output

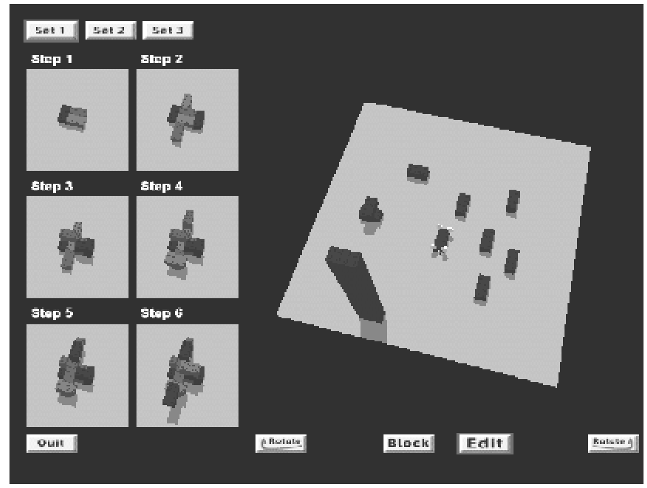

Figure 5: Screenshot of the prototype software application.

from the display. The "private version" of the software has three differences with respect to the "public version." First, user cursors are private. While users can still see operations being performed by the other user, their partner's cursor is invisible. Second, instructions are private. Each user's instruction area is independent of the other user's. For example, User A can look at instruction sequence #1, while User B is looking at instruction sequence #3. If they choose to, of course, the two users can look at the same instruction sequence. In the "public version," both users must look at the same instructions. Third, contextual menus are private. Each user can only see his or her own menus. In the "public version" users can see, but not interact with, the other user's menus.

Along with the two versions of the software, two different three-structure instruction sets were created for use with the study.

THE USER STUDY

We conducted an exploratory study to investigate the usefulness of displaying private information on a shared display for a collaborative activity. Given that this is a novel interaction paradigm, we sought to gain preliminary insights into users' interactions in this type of environment, and to investigate interface issues for the design of SDP.

Method

Participants and Setting This study was conducted on-site at Electronic Arts Canada, in Vancouver, Canada (where this research was being undertaken). Sixteen employees (8 male and 8 female) took part in the study. All participants were volunteers, recruited by an email sent to all employees at Electronic Arts Canada. The participants included artists, programmers, and producers.

Experimental Design The participants were assigned a partner of the same gender, and each pair was observed working with both the "private version" and the "public version" of our prototype system. To reduce learning effects, two different instruction sets, of similar difficulty, were created. Each instruction set included three instruction sequences. Each instruction sequence described how to build one LEGO-like structure. To reduce order effects, the order of prototype version and instruction set were both counterbalanced, resulting in four different combinations. For each combination, we observed two pairs (one male and one female), for a total of eight different pairs.

Qualitative and quantitative data were collected using video, computer logs, and pre and post session questionnaires. The video camera was positioned in front of the participants to best capture their interactions during the session. The computer logs recorded the participants' interactions with the system and logged the amount of time each user spent looking at each instruction set. The pre-session questionnaire was used to gain knowledge related to the participants' experience with computers. The post-session questionnaire was used to gather participants' reactions to both versions of the prototype system, including opinions and preferences towards the two versions and individual features.

Procedure Before beginning a session, participants were asked to complete a colour-blindness test and fill out the pre-session questionnaire. The ability to discern colour was important for the experimental task, as different blocks were different colours. We then introduced the task to the participants and explained the two versions of the prototype system. They were also given an instruction-sheet, detailing different actions possible in the software. Following this, each participant was asked to spend five minutes completing an example task with the software, in order to become familiar with the interface. The participants were then asked to work together with their partner for ten minutes, building three LEGO-like structures using either the "private" or "public" version of the prototype. Following this, they completed a second ten-minute trial, building another set of three LEGO-like structures, using the alternate version of the prototype. In both trials, they were asked to complete the task as quickly and as accurately as possible. Upon completion of the two trials, all participants filled out the post-session questionnaire.

Results

Effectiveness of the Private Condition Observations of the participants' interactions for each version of the game (private and public) revealed interesting findings. An important first point to address is that all participants, when using the private version, were able to effectively interact and progress towards the goal. In fact, pairs of participants performed a marginally significant higher number of actions (placing blocks, rotating blocks, and changing the colour of blocks) in the private condition as compared to the public condition, $F(1,7) = 3.668, p = 0.097, r^2 = 0.344, power = 0.380$. On average, participants performed 73 actions in the private condition and 62 actions in the public condition. We further examined the total number of blocks placed by each pair (not including the blocks that were subsequently removed because of errors), for each of the two conditions. This gave an indication of how many blocks were correctly placed. Pairs placed more blocks on average in the private condition (28 on average), than in the public condition (23 on average), however this difference was not statistically significant, $F(1,7) = 2.344, p = ns, r^2 = 0.251, power = 0.264$.

Results from the post-session questionnaires revealed that participants were relatively comfortable using both conditions. On a five-point scale (with one representing very comfortable and five representing very frustrating), 13 of 16 participants ranked the private condition as a one or two and 9 of 16 participants ranked the public condition as a one or two (see Figure 6). This result was not statistically significant ($Wilcoxan Z = -0.962, ns$).

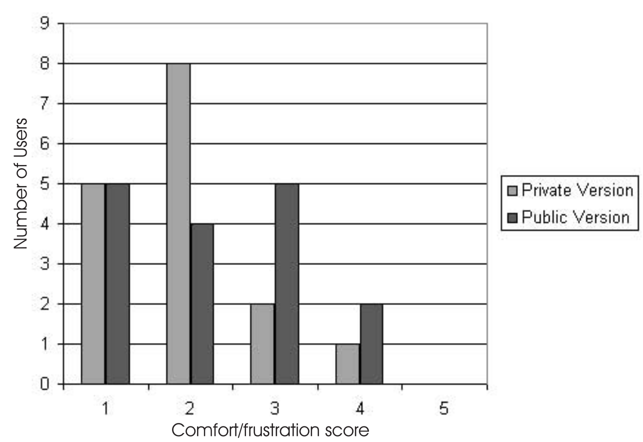

Figure 6: User experience with private and public version (1=Very comfortable, 5=Very frustrated).

Collaborative Strategies Another important result regards the collaborative strategies used during the trials. The private version of the prototype allowed the two users to view either the same instruction sequence, or different instruction sequences. This made it possible for participants to either work collaboratively, or independently. This is in contrast to the public version of the prototype, where users were limited to looking at the same instruction sequence. The computer logs recorded the amount of time the different instructions were viewed by each participant. These logs indicate that some participants took advantage of the flexibility of the private version of the prototype. Figure 7 shows the percentage of time the two users in each pair spent looking at the same instruction sequence when using the private version. It can be seen that four pairs worked almost exclusively together, spending between 98% and 100% of their time looking at the same instruction set. Two pairs of users worked mostly independently, spending 6% and 7% of their time looking at the same instruction set. The remaining two pairs developed a hybrid approach, sometimes looking at the same instruction sequence, sometimes looking at different instruction sequences. Comments from these two pairs of users shed light on some reasons for adopting the hybrid approach:

> "Even though we were working on different parts we could still proof the other guy's work easily and quickly."

> "Private condition rocked simply because it was easier to make better use of our time because we could both work on different sets and then when we started working on the same set we could still talk about what we were going to do."

> "For the private sets we worked in parallel on different data sets and then worked together on the third."

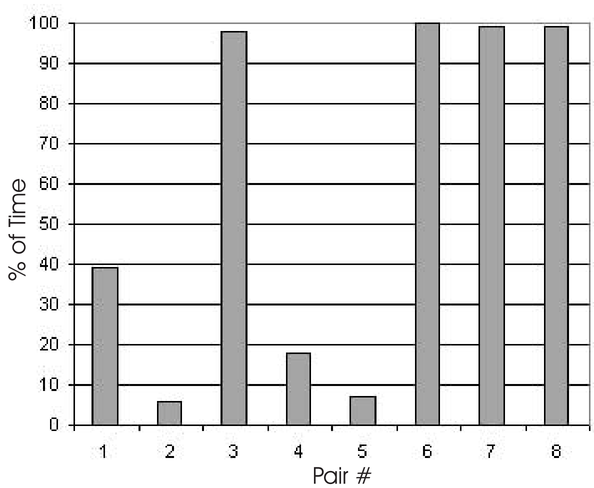

Figure 7: Percentage of time each pair spent looking at the same instruction sequence.

Privacy and Awareness In terms of awareness and privacy, the participants in our study expressed a preference for having menus and instruction sets private, while for cursors, their preferences were mixed (see Figure 8). This issue was also raised frequently in the participants' feedback on the post-questionnaire. Ten people made explicit statements that they disliked certain components being public. Statements were made against public menus (7 statements), public cursors (2 statements), and public instructions (3 statements). Comments included: "overlapping obstructs view" and "think the other active mouse is yours." One person made an explicit statement that they liked public cursors, "could use them to point at what I want to talk about." This evidence, that different users have different preferences as to what should be private, suggests that user customization of SDP systems might be valuable.

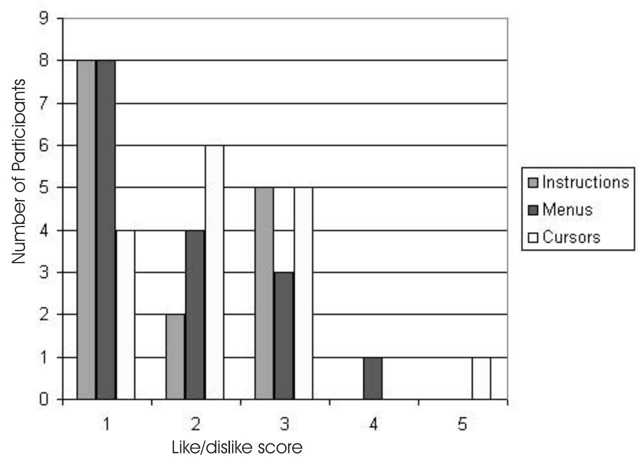

Figure 8: User response to instructions, menus, and cursors being private (1=I liked it a lot, 5= I disliked it a lot).

Observations revealed that displaying instructions, cursors, and menus in the private output channel occasionally resulted in confusion on the part of the participants. For example, one participant would gesture to an object with their private cursor and comment "look at this." This is natural behaviour given that two people looking at a shared physical artifact (e.g. the computer display) expect to see similar things. This phenomenon is potentially cognitively dissonant and may need to be addressed with additional awareness tools, such as tags indicating privacy, to facilitate users' interactions in this unique environment.

CONCLUSIONS

Single Display Privacyware, as introduced in this paper, is a new interaction technique that confronts the problem of how to display private information when a group of users is working on a shared display. SDP allows private information to be shown to specific users on a display that is shared by a group. The display, in essence, is augmented with private information suitable for each individual user.

The results of a user study using a prototype SDP system indicated that SDP is a valid interface technique that is well-accepted by users. Only one of the sixteen subjects indicated discomfort working with the SDP system. This is impressive, considering that this was a first attempt at choosing a domain and task suitable for an SDP system, as well as a first attempt at implementing such a system.

The analysis of the study resulted in several new questions regarding SDP. One important question regards how privacy can be used to foster either collaboration or independent investigation. Privacy made it possible for users in the experiment to work more independently than if they were working without privacy support. Privacy also allowed users to switch back and forth between collaborative and independent approaches to the task. Whether or not it is important to support independent work in SDP systems is a question that must be investigated. Another question that was raised regards how users are to identify what information is private and what is public. Showing private information within the context of public information can be advantageous, but it also blurs the distinction between the two for the user. Methods of keeping the user informed as to what information is private must be developed. Finally, it was observed that different users have different preferences concerning what information should be kept private, and what should be made public. An investigation into what level of privacy customization should be supported would be valuable.

ACKNOWLEDGEMENTS

We would like to thank EA Canada, BC ASI, and NSERC for funding this research, as well as Keith Shu for his help with the code, Colin Swindells for his EE help, Tim Crane from StereoGraphics for help customizing the hardware, Tom Shermer for supervising, Daryn Mitchell for art and proofreading, Kellogg Booth and Christine MacKenzie for the loan of equipment, and members of the EDGE Lab and GRUVI Lab for support.

REFERENCES

1. M. Agrawala, A. C. Beers, B. Fröhlich, P. Hanrahan, I. McDowall, and M. Bolas. The two-user responsive workbench: Support for collaboration through independent views of a shared space. In *SIGGRAPH 97 Conference Proceedings*, pages 327–332. Addison Wesley, 1997.

2. E. A. Bier and S. Freeman. MMM: A user interface architecture for shared editors on a single screen. In *Proceedings of UIST '91*, pages 79–86. ACM Press, 1991.

3. L. J. Bricker, M. J. Bennett, E. Fujioka, and S. L. Tanimoto. Colt: A system for developing software that supports synchronous collaborative activities. In *EdMedia*, 1999.

4. A. Butz, T. Höllerer, S. Feiner, B. MacIntyre, and C. Beshers. Enveloping users and computers in a collaborative 3D augmented reality. In *Proceedings of IWAR 99 (International Workshop on Augmented Reality)*, pages 35–44, 1999.

5. S. Greenberg, M. Boyle, and J. LaBerge. PDAs and shared public displays: Making personal information public, and public information personal. *Personal Technologies*, 3(1), 1999.

6. C. Gutwin and S. Greenberg. Design for individuals, design for groups: Tradeoffs between power and workspace awareness. In *Proceedings of CSCW '98*, pages 207–216. ACM Press, 1998.

7. C. Gutwin and S. Greenberg. Effects of awareness support on groupware usability. In *Proceedings of CHI '98*, pages 511–518. ACM Press, 1998.

8. J.-P. Hourcade and B. B. Bederson. Architecture and implementation of a Java package for multiple input devices (MID). Technical Report CS-TR-4018, University of Maryland, May 1999.

9. K. M. Inkpen, W. Ho-Ching, O. Kuederle, S. D. Scott, and G. B. D. Shoemaker. "This is fun. we're all best friends and we're all playing!": Supporting children's synchronous collaboration. In *Proceedings of CSCL '99*, 1999.

10. K. M. Inkpen, J. McGrenere, K. S. Booth, and M. Klawe. The effect of turn-taking protocols on children's learning in mouse-driven collaborative environments. In *Proceedings of Graphics Interface '97*, pages 138–145, 1997.

11. B. A. Myers, H. Stiehl, and R. Gargiulo. Collaboration using multiple PDAs connected to a PC. In *Proceedings of CSCW '98*, pages 285–294. ACM Press, 1998.

12. J. Rekimoto. A multiple device approach for supporting whiteboard-based interactions. In *Proceedings of CHI '98*, pages 18–23. ACM Press, 1998.

13. S. D. Scott, G. B. D. Shoemaker, and K. M. Inkpen. Towards seamless support of natural collaborative interactions. In *Proceedings of Graphics Interface 2000*, pages 103–110, 2000.

14. G. B. D. Shoemaker. Supporting private information on public displays. In *CHI 2000 Extended Abstracts*, pages 349–350. ACM Press, 2000.

15. J. Stewart, B. B. Bederson, and A. Druin. Single Display Groupware: A model for co-present collaboration. In *Proceedings of CHI '99*, pages 286–293. ACM Press, 1999.

16. J. Stewart, E. M. Raybourn, B. Bederson, and A. Druin. When two hands are better than one: Enhancing collaboration using Single Display Groupware. In *Proceedings of CHI '98*, pages 287–288. ACM Press, 1998.

17. Z. Szalavari, E. Eckstein, and M. Gervautz. Collaborative gaming in augmented reality. In *Proceedings of VRST 98*, pages 195–204, 1998.

18. M. Tani, M. Horita, K. Yamaashi, K. Tanikoshi, and M. Futakawa. Courtyard: Integrating shared overview on a large screen and per-user detail on individual screens. In *Proceedings of CHI '94*, pages 44–50. ACM Press, 1994.