

Multi-User Interaction using Handheld Projectors

Xiang Cao¹, Clifton Forlines^{2,1}, Ravin Balakrishnan¹,

¹Department of Computer Science
University of Toronto
www.dgp.toronto.edu
caox | ravin @dgp.toronto.edu

²Mitsubishi Electric Research Labs
Cambridge, MA, 02139, USA
www.merl.com
forlines@merl.com

ABSTRACT

Recent research on handheld projector interaction has expanded the display and interaction space of handheld devices by projecting information onto the physical environment around the user, but has mainly focused on single-user scenarios. We extend this prior single-user research to co-located multi-user interaction using multiple handheld projectors. We present a set of interaction techniques for supporting co-located collaboration with multiple handheld projectors, and discuss application scenarios enabled by them.

ACM Classification: H5.2 [User Interfaces]: Interaction styles. I.3.6 [Methodology and Techniques]: Interaction techniques.

General terms: Design, Human Factors

Keywords: Handheld projector, multi-user interaction.

INTRODUCTION

Recent advances in projection miniaturization will soon enable projectors to be carried in a pocket or even embedded in mobile devices such as cell phones and PDAs. With the ability to project information, handheld devices can overcome the inherent information-display limitations of their small embedded screens, and instead create larger displays on virtually any external surface. When coupled with appropriate tracking technologies, interaction with the displayed information can also move beyond the confines of the handheld device itself to encompass almost an entire physical environment.

Most handheld projector research to date [3, 8, 15, 16] has focused on supporting a single user. However, the larger displays generated by handheld projectors inherently afford multi-person viewing and thus have the potential to support co-located collaboration. In particular, when each user has a handheld projector, the interactivity between projectors can result in a rich design space for multi-user interaction.

Although many current handheld or portable devices have the ability to exchange data with other devices via wireless connections, the interaction required to facilitate such

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

UIST'07, October 7–10, 2007, Newport, Rhode Island, USA.
Copyright 2007 ACM 978-1-59593-679-2/07/0010...\$5.00.

exchange often requires cumbersome and explicit authentication procedures. While such procedures are generally unavoidable when devices (and their users) are not physically co-located, they may be unnecessary if we can design interaction that exploits the user and device co-locality to facilitate connectivity and collaboration.

Researchers have explored co-located collaboration between people using shared displays on tabletops [20, 24, 25] and walls [7, 9, 11, 22]. Since the workspace is shared between all users, information exchange and multi-user operations can be easily realized. However, these shared displays are not portable for ubiquitous use, and every user shares the same view of the workspace. Private and personalized information are not easily accommodated, and global conflicts [13] may occur, in which one user's action affects the entire shared display and disrupts other users.

In contrast, the use of multiple handheld projectors may open up a novel interaction paradigm for co-located users, in which they can share the same physical display and interaction space, while at the same time individually creating and controlling parts of the overall virtual display with their projectors. In this paper we explore the design space of multi-user interaction using multiple handheld projectors. Building on the single-user handheld projector interaction system we described in earlier work [8], we developed a set of interaction concepts and techniques to support multiple users working in a shared physical space, each using their own projector that is spatially tracked within the physical environment (Figure 1)




Figure 1. (a) System concept. (b) Handheld projectors. (c) System in use.

RELATED WORK

The Hotaru system [23] uses a static projector to simulate mobile projections. It also supports basic touch based annotation, rotation and file transfer. Beardsley et al. [3] and Raskar et al. [16, 17] introduced a computer vision approach to correct the projected image and create a stable distortion-free virtual desktop inscribed in it. Standard mouse-operations were achieved using a cursor displayed in the center of the projection. A flashlight metaphor has also been explored by several researchers [6, 8, 15], in which the projector reveals a portion of a large virtual workspace stationary on the projection surface. In particular, Cao and Balakrishnan [8] explored a set of generic interaction techniques using this metaphor and presented techniques for defining information spaces in a new environment on-the-fly.

Much of the previous research has focused on single-user single-projector interaction, with relatively little done on supporting interaction between users using multiple handheld projectors. Cao and Balakrishnan [8] discussed how the handheld projector may support both synchronous and asynchronous communication between people, but only focused on using one projector at a time. The Hotaru system [23] supports simple file transfer between users by overlapping the projections, but is preliminary otherwise.

On the other hand, co-located collaborative groupware has been widely investigated in other settings, especially with shared displays such as walls [7, 9, 11, 22] or tabletops [20, 24, 25]. Morris et al. [13] also discussed multi-user coordination policies for co-located groupware in general. Shoemaker and Inkpen [21] explored presenting private information on a shared display by letting users wear shutter glasses, trying to overcome the limitation of one shared view across users. However, the metaphors and techniques designed for these settings may not be applicable or sufficient for the distinct affordances of handheld projector interaction.

Researchers have also explored techniques to support direct information exchange between devices. Rekimoto's [18] Pick-and-Drop is a pen-based technique for transferring data by picking objects on one computer and dropping on another. Yatani et al. [26] present Toss-It, a technique for sending information from a PDA to other devices using a tossing gesture. Hinckley et al. [10] present stitching, pen-based gestures for binding multiple displays. Park et al. [14] present Touch-and-Play, enabling users to transfer data between devices by touching them. These techniques help connect information and display spaces on different devices. In contrast, multiple handheld projectors can seamlessly combine their display and information spaces by projecting onto a shared physical surface.

SYSTEM OVERVIEW

Building on our earlier single handheld projector system [8], our current prototype uses two Mitsubishi™ PK10 Pocket Projectors (Figure 1b), each weighing about 1 pound, with a resolution of 800×600 pixels. Each projector is augmented with two buttons for input (primary button for

selection and operations, and secondary button for triggering menus), and can be easily handled and moved using one hand. Two passive pens are also included for writing on surfaces. Both the projectors and the pens are tracked by a Vicon camera-based tracking system (www.vicon.com), which provides 6-dof (position + pose) information at millimeter precision. While we anticipate that ubiquitous high-quality spatial tracking will eventually become commonplace, this Vicon tracking system allows us to prototype future interaction scenarios using today's technology. Both projectors are connected to a 2.4GHz P4 PC, which produces the images and handles the interaction.

As in our earlier work [8], we use a flashlight metaphor as the basis for our interaction (Figure 1a). The image projected by each projector reveals a portion of a large workspace stationary on the projection surface. When the projector is moved, the projected image content changes accordingly, as if the projector is used as a flashlight to explore in darkness. This is implemented by changing and warping the projected image according to the projector's movement. Multiple workspaces can be associated to different physical surfaces in an environment, such as walls, tables, and bulletin boards. The workspaces are shared among the projectors. Different projectors may reveal different (or overlapping) regions of a workspace simultaneously. In addition, each projector's view of the workspace may also be personalized depending on which user is using it.

Tracking and calibration inaccuracies may result in slightly imperfect image alignment in overlapping projection areas. To avoid unpleasant double images, we provide the option to blank out one projector in the overlapping area, and let the other projector handle the display for both. We expect that advances in computer vision and calibration techniques will solve this problem in the long run.

Two users, each having their own projectors and pens, can interact with the system simultaneously. The system architecture is scalable to support three or more users. Each user is identified by a unique color, which is reflected by the cursor displayed by the projector, and the marks drawn by the pen. With two buttons, each user can independently manipulate virtual objects using the cursor, and trigger commands using crossing-based widgets. (details described in [8]). In this paper, we will focus on techniques to support interactions that involve multiple users.

INTERACTION CONCEPTS & TECHNIQUES

Ownership & Access Control

Each object in the workspaces may either have no ownership (accessible by all users), or be owned by a particular user. In the latter case, the owner of the object has full control over it. How other users can interact with it is determined by its access level, which is one of the following:

Public: The object is visible to all users (i.e. all projectors will display it), and all users can operate on it. Any object without an owner is implicitly public.

Semi-Public: The object is visible to all users, but only operable by its owner.

Private: The object is visible and operable only by its owner. It is not displayed in other users' projectors.

The ownership and access level of an object is indicated by the flags on the top right and top left corners of the object, respectively (Figure 2). The color of the ownership flag matches that of the owner, and the color of the access flag indicates the access level: green for public, yellow for semi-public, and red for private. The owner can cycle through the access flag levels by crossing it from outside the object to inside while holding the primary button down. Note that the term "visible" used within this access level context is not to be interpreted in the strictest sense, since other users can still peek at the object when it is being viewed by its owner. To completely hide the object's content, the owner can toggle the visibility flag by crossing it from outside the object to inside while holding the primary button down. The object will then be shown as a blank frame in its owner's view, and invisible to other projectors. Hiding an object's content using the visibility flag in this manner implies setting its access level to private.



Figure 2. Object with flags.

File Exchange

Exchanging files between users is a common task in multi-user interaction. Compared with current handheld devices, which rely on indirect procedures, in our system all file exchanges can be achieved by direct manipulation using several techniques each suited to different situations:

Passing Ownership

Users can pass the ownership of their objects to others, much like handing over a physical object in the real world. The receiving user can then operate on the object freely, or drag it into their personal folder (a container that stores all personal files) for later use. The ownership passing action is completed by the following steps:

Step 1: User A (owner) captures the object by clicking on it and holding the primary button down. (Figure 3a).

Step 2: User B clicks on the object and holds down the primary button.

Step 3: Both users dwell their cursors. (Figure 3b).

Step 4: After a brief period (≈ 1 sec), User B "captures" the object and becomes its owner (Figure 3c).




Figure 3. Passing object ownership. (a) before. (b) during. (c) after.

While dwelling is often undesirable in interactive systems, it serves an important purpose here in that the dwelling required of both users ensures quasi-explicit consent from both parties to perform the action. A handshake icon starts to fade in at Step 3, and reaches its full opacity by Step 4, giving users an indication of the upcoming passing action. During this stage, User A can either move away or release the object to prevent undesirable or unintentional passing.

Dropping into Personal Folder/Portal

For more efficient file exchange, objects can be directly dragged and dropped into one's personal folder. The dropping action is regulated by the access levels of both the object and the folder. Users can drop anything that they have operation access to into their own folder. In order to drop objects into other users' folders, either the folder needs to be set as public, or the folder's owner needs to hold down the primary button over the folder while the object is dropped into it, analogous to the real world action of holding a bag open to allow others to put things in it.

To provide access to folders far away, or to protect privacy of a folder's content, users can create a portal to their folder (Figure 4). Ownership is indicated by the portal's color. Anything dropped into the portal will be transported into the associated folder. The portal follows the same access control policies as the personal folder.




Figure 4. Dropping files into a portal.

Compositing Projections

When multiple handheld projectors are available, their displays can be composited to improve the viewing experience, beyond what is possible with a single projector.

Expanding the Display Area

As Raskar et al. [16] suggested, multiple projections can be aligned side by side to create a larger display area than a single projection without sacrificing image resolution. This is clearly helpful for viewing a large document or map.

In addition, our system can also intelligently adapt the view of the object to exploit the enlarged display area provided by multiple projectors. For example, when watching a movie, a cropped version of the movie is displayed when viewed with one projector, but seamlessly switches to a widescreen version when two projection displays are aligned horizontally to accommodate it. (Figure 5)




Figure 5. Viewing a movie.
(a) Cropped view. (b) Widescreen view.

Different projectors can also point at different regions of the workspaces, thus creating multiple viewing/operating areas, especially on different projection surfaces. Similar to Hinckley et al.'s [10] example using multiple tablet displays, one user can click on thumbnails of photos in an album projected on a surface convenient for operation (e.g. a table), while another user projects the full view of the selected photo on a larger surface (e.g. a wall). (Figure 6) Other applications include exploring related parts of a large graph, or transporting information between places.




Figure 6. Photo album browsing using two displays.

Accommodating and Combining Multiple Views

One unique characteristic of handheld projectors is that each user is creating their own display while sharing the same physical and information space. This can fluidly enable different views of the same object when displayed by different projectors, and allow users to see personalized information relevant to themselves. For example, a calendar shows appointments of the user who is projecting it as colored blocks (Figure 7a, b), and a photo frame shows a photo of the user who is projecting it (Figure 9a, b).

When multiple projections overlap on the same object, the different views (if applicable) are seamlessly blended by the optical overlaying of projection images. The result is a combined view that is relevant to all projecting users. For example, a calendar displayed by two overlaying projectors shows events for both users, and the empty spaces in it are timeslots available for both people to have a meeting (Figure 7c). This provides an intuitive and efficient way for scheduling meetings. In order to maintain privacy, text labels describing the events are hidden when the calendar is viewed by multiple users.

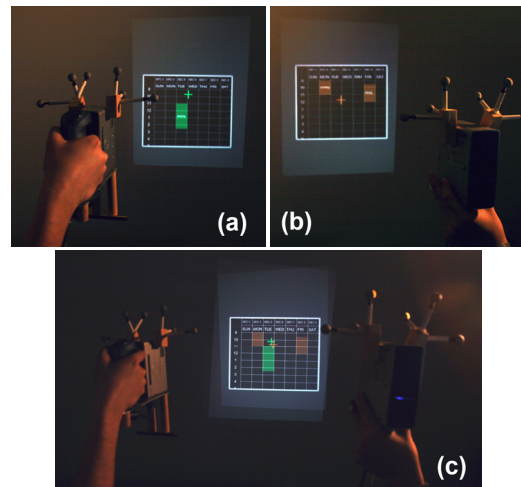


Figure 7. Calendar. (a) viewed by User A.
(b) viewed by User B. (c) viewed by both users.

As explored in the literature [8, 16], the handheld projector supports image resolution gradation and multiple information granularities depending on the distance between the projector and the projection surface. As the user comes closer to the surface, the projection area shrinks and a higher pixel density is achieved in the area, resulting in higher local resolution that can be used to display more detailed information. Utilizing this feature, multiple projectors can be combined to create a viewing experience similar to that of a focus plus context display [2]. One projector can be held afar to create the low-resolution coarse-granularity surrounding context in a larger area, and another projector is used close to the surface creating a focus region to explore high-resolution fine-granularity details within that context. Because the projection image also becomes brighter when the projector is nearer, the projection of the focus region automatically overlays the context information beneath it. Figure 8 shows this with a

multi-granularity city map. The context region shows main streets only, while the focus region shows all small streets. Compared with previous focus plus context screens [2], where the resolution and position of both focus and context displays are fixed, our solution is more flexible in that users can dynamically move, resize and change the resolution of both projections. We can also achieve nested focus regions by overlaying three or more projectors



Figure 8. Focus plus context display.

Direct blending of multiple views may not always be sufficient. The system can also render a semantic combination of different views when it detects multiple projectors overlaying on an object. For example, the photo frame shows a group photo of both users when two projectors overlap on it (Figure 9c). Alternatively, some critical information may only be revealed when multiple users look at them at the same time. For example, any single user can only see the cover page of a group assignment. Only when two or more users overlap their projections on it, can they read the content of the assignment. An extreme case of this is an object that is completely hidden in any single projector’s view, but become visible when projections overlap.




Figure 9. Photo frame. (a) viewed by User A. (b) viewed by User B. (c) viewed by both users.

Similar to our focus plus context usage, semantically combined views can also be useful when one projection is contained in another. A different view is rendered in the overlapping region only, as if using a magic lens [5]. Figure 10 shows a user exploring the inner structure of a car model using one projector like a magic lens.




Figure 10. Emulating a magic lens with semantically combined multiple projections.

Linkage between Objects

Two users can create linkages between their objects (one from each user) for information or operations that involve both (if applicable). Depending on their needs, two types of linkages can be created.

Snapping

Snapping provides a lightweight transient way to quickly view information that involves two objects. If two objects are compatible for snapping, then when they are moved close enough, they will snap to each other side by side. To unsnap them, users simply drag either or both of the objects in any direction past a small distance. When snapped together, either or both the objects will change their appearance to reflect information that relates to its partner. For example, when two users snap maps of their home addresses together, the maps change to show the direction between the two addresses, both in drawing and in text (Figure 11). When a clock is snapped to a city map, it changes its time zone to reflect the local time of the city.




Figure 11. Snapping. (a) Separate. (b) Snapped.

Docking

To perform operations that affect (possibly modify) multiple objects, two users can create a more explicit linkage between their objects by docking them together. Compared with snapping, which can be initiated by one user, docking is a more strict action that requires consent from both users to prevent operations not authorized by the object owner. In addition, only two objects of the same kind (two documents, two calendars, etc.) can be docked.

The docking action is completed by the following steps:

Step 1: Both users capture their own object by clicking on it and holding down the primary button (Figure 12a).

Step 2: Both users move their objects to roughly overlay them together, and dwell the cursors. (Figure 12b).

Step 3: After a brief period (≈ 1 sec), both objects become docked together (Figure 12c).




Figure 12. Docking action.
(a) before. (b) during. (c) after.

Similar to the ownership passing action, the dwelling ensures consent from both parties. A linkage icon starts to fade in at Step 2, and reaches its full opacity by Step 3. While the icon is fading, either user can move away or release the object to prevent unwanted docking.

Two docked objects become precisely superposed, and will be operated as a whole. Both ownership flags are shown side by side. The objects may also change their appearance to reflect information from their partners. A linkage flag on the bottom right corner indicates the linkage has been established. (Figure 12c) The pair of objects can be operated by both object owners, and any operation on the pair will affect both objects. For example, two half-completed documents (each worked on by a different user) can be docked to preview the combined document. Both users can then write annotations on it using their pens. The annotations are shown on both documents when undocked.

To undock a pair of objects, either user can toggle the linkage flag by crossing it. Then users can move their own objects away. Alternatively, two users can “tear apart” the objects by both capturing the docked pair at the same time, and dragging them in different directions. The “tear apart” action can also be performed immediately after Step 3 before either user releases the button. This enables “transient” docking to get a glimpse of the docked view.

Other examples include docking two personal calendars. Either user can then click in an empty timeslot to create a meeting for both people. For each user, the meeting will be labeled as “Meet X”, with X being the name of the other user (the label will not be shown until the calendars are undocked). Two users may also dock their portals, resulting in a two-way portal between their personal folders. Objects dropped into the docked portal by either user will be transported to the other user’s folder. This provides an efficient way for them to quickly exchange files.

Snapshot

While working in a shared workspace, sometimes a user may want to record information for later reference, especially when the information is from other users or created collaboratively. Triggered by a menu command, a user can take a snapshot of his/her region of interest. A translucent square (“viewfinder”) inscribed within the projection area indicates the region to shoot at, which can be moved and resized by moving and rotating the projector respectively (Figure 13). Pressing the primary button takes the shot as an image copy of what the user’s projector displays inside the viewfinder, which can be then manipulated and saved as an ordinary object owned by the shooter. The snapshot can also be used to take small parts from a large object (e.g. a document or a map) to reflect the point of interest. Note that a user cannot take peep shots of private information displayed by others, since that information is not shown in the shooter’s projector.




Figure 13. Snapshot.

Spatial Relationship between Users

The spatial proximity between people plays an important social role in terms of privacy. People get close to each other to have private conversations. Conversely, people feel uncomfortable if somebody else comes nearby when they are viewing private information. Our system can estimate user proximity and face orientation from the spatial locality of the projectors, and use this information to facilitate subtle interpersonal interaction that exploits real-world social protocols.

When a user is viewing a private object, if another user comes nearby, the private object becomes blurred so as to prevent the second user from peeking at its content (Figure 14a). Similarly, private objects also get blurred when other users cast their projection onto it, which also suggests that they are looking in that direction (Figure 14b). On the other hand, only when two users are close to each other can they

perform private communications such as passing ownership of a private object. For another example, when a private letter is addressed to both users, its content will only be revealed when the two people stand close to each other, and overlap their projection on the letter (Figure 15).

Another feature that exploits the spatial relationship between users is to avoid shining the projector into people’s eyes. When a projector is pointing at any other users’ estimated position, it temporarily turns off to avoid hurting their eyes.




Figure 14. Blurring private information. (a) when another user is nearby. (b) when within another user’s projection.




Figure 15. Private letter for two users. (a) Content hidden when they are afar. (b) Content revealed when they come close.

Independent Work

Users sharing the same physical space may not always be collaborating with each other. Compared with systems using a single shared display, the use of multiple handheld projectors allows users to see information relevant only to them. This largely reduces the likelihood of interpersonal conflict when people are working independently in a shared space. To further facilitate independent work, we provide the ability to create a “fence” around one’s work.

Triggered by a menu command, a user can sketch a line in the workspace using the cursor, which turns into a fence between his/her territory and another user’s. Once the drawing is finished, all objects in the workspace that belong to the user will be pulled back to his/her side of the fence, making room for the other user to work (Figure 16). The other user’s objects will stay where they are. The fence will stay visible as an informal demarcation between two users’ territories. However, it does not actually prevent users from moving objects beyond it. It is up to the users to maintain the notion of the boundary. The reason for this design is that we provide this feature only to facilitate but not override or enforce social protocols. Users will still have the flexibility to dissolve the boundary between them, or completely ignore it and start collaboration. This is also why the fence only pulls back the owner’s objects but does not push back other users’ objects. By doing so, people accommodate and do not compete with each other. The use of the fence should only be the result of a well-negotiated common understanding between people. Both users can draw a fence multiple times, only the most recent one in the workspace will stay visible.

A user can also set the access level of all their objects at once to prevent other users from operating or viewing any of their belongings. Both this and the fence feature can be particularly helpful when a user starts by working alone, and is subsequently joined by another user who then shares the workspace.




Figure 16. Drawing a fence. (a) before. (b) after.

USAGE SCENARIOS

The interaction concepts and techniques discussed above can support a variety of potential usage scenarios, including but not restricted to the following:

Casual Communication

Given the portability of handheld projectors, it is natural to use them to facilitate casual communication between people when they encounter one another. For example, exchanging contact information can be as easy as dragging avatars of people between personal folders (Figure 17a). Scheduling a meeting can be done by docking two personal calendars, without the hassle of separately comparing each other's schedule (Figure 17b). Sharing information such as music and photos can be achieved in several ways, including passing ownership and dragging into a personal folder or portal (Figure 17c). A user can also use the pen to write a note and drop it into another person's folder as a reminder (Figure 17d). Any of the above activities may take several steps using current mobile devices, sometimes even involving manual copying (e.g. phone numbers).




Figure 17. Casual communication. (a) Exchanging contacts. (b) Scheduling a meeting. (c) Music sharing. (d) Writing a reminder for the other person.

Group Meeting

Although mostly designed for mobile usage, multiple handheld projectors may also support more organized group meetings such as presentations or brainstorming. One projector can be dedicated to display the presentation slides or project a virtual whiteboard to write on. Attendees can use their personal projectors to access additional information on the table (Figure 18a). They can also use the pen to write comments on the table first, and then drag it to the wall to post it when desired (Figure 18b). In addition, they may take snapshots of the presentation slides using the projector, and then write notes on them (Figure 18c). The former has become a common practice using digital cameras, but the latter is currently inconvenient to do without the slides available beforehand.

Games

Even without dedicated design for gaming, the current system features can already support a few simple interesting games. For example, utilizing the access levels and multiple views of objects, we can create a collaborative

treasure hunting game. Players need to complement and combine their projectors' viewing powers to discover and collect treasures scattered around the physical space. Some treasures are only discoverable by a particular user, others only when projections overlap (Figure 19a). Utilizing the snapshot function, people can also play an ad hoc jigsaw puzzle game. Players take parts from a large picture using snapshots, and try to reassemble them later (Figure 19b).




Figure 18. Group meeting. (a) Looking up additional information. (b) Posting a comment. (c) Taking a snapshot of the presentation.




Figure 19. Games. (a) Treasure hunting. (b) Jigsaw puzzle.

Mobile social games [1, 4, 12] in which players physically walk around a city to complete tasks such as collecting treasures according to directions given by mobile devices, have attracted attention recently. These games encourage players who are not familiar with each other to collaborate face to face, therefore promoting social interaction between people. With the assistance of handheld projectors, the experience of mobile social games may be altered and arguably improved by projecting game information into the physical environment, thus further blurring the boundary between the game and reality.

INFORMAL USER FEEDBACK

For preliminary user feedback, we asked four individuals, working in pairs, to try the prototype system. Three of the four participants are regular cell phone users, and one of them owns a PDA. We demonstrated all the system features to each pair of participants, and then asked them to freely try out the techniques, especially those which involve interaction between them. Each session lasted about an hour. We observed participants' behaviors, and conducted a post-study interview.

All participants grasped the system concepts quickly, and did not show any difficulty learning the interaction techniques. As we expected, the feature that they found most appealing is the ability to easily exchange information in a shared workspace. The multi-view calendar also especially resonated with users, as it largely simplifies one of their most frequent tasks – scheduling meetings. Other highly welcomed features include the movie player that adapts to multiple projections, the snapshot function, the focus plus context map, and the fence to support independent work.

The participants' experience seemed to be affected by the imperfect alignment between projectors, as well as the somewhat jittery projection caused by unideal image update rate (≈ 25 Hz). These could be reduced with technical advances. They also had some reservations about projecting private information in public space, although they all agreed that the system designs surrounding privacy protection alleviated their concerns to some extent. Exploiting the embedded small screen on handheld devices for highly private information and operations may be one way to address this concern.

Some participants asked for more advanced support for collaboratively authoring and annotating text. Another participant suggested having a selection box that can be moved and resized similar to the snapshot viewfinder to quickly select and operate on multiple objects. For independent work support, participants suggested expanding the fence to other shapes such as a circle.

DISCUSSION

Handheld projectors provide interesting design challenges compared to other co-located collaborative settings such as a shared tabletop display. For example, users can create their individual displays with their projectors, allowing for easy support of personalized views, which is seldom the case in other settings. This also enabled the three-level access control we proposed as opposed to simply public v.s. private in most other systems. Users can also easily point the projector to virtually anywhere in the workspace with little physical constraints, whereas in the tabletop setting the reachability is constrained by the user's sitting position and arm length. This makes some social protocols that work in the tabletop scenario less applicable in our setting, therefore we considered alternative ways to coordinate users, especially for supporting independent work. Further, when working with handheld devices, the somewhat stable demarcation of personal and group territories used in

tabletop interaction [19] is less applicable because of the constant change of users' positions. However, changes in spatial relationship between users can be exploited to facilitate subtle interpersonal interaction.

In our prototype, both projectors are connected to the same computer, thus all data transmission is done locally. In the real world, we can expect the data exchange between handheld projectors to be backed either by peer-to-peer connections such as Bluetooth or infrared, or by centralized services such as WiFi or cell phone networks. The shared workspace created by the projectors can make the background connection mechanism transparent to the users. Identity verification is achieved simply by looking at the person's face, thus eliminating the need for passwords or other complex authentication schemes.

An interesting issue is that although our system provides various ways to support privacy, in some social contexts the very fact that the user is projecting data may be perceived as an indication that the information is public, and viewed as an invitation for other people to participate. This disparity indicates that more delicate design may be needed to convey subtle privacy cues to others, without changing current social protocols. One possible solution may be complementing the projection display with the small screen embedded on the handheld device to accommodate different scenarios. However, the implicitly public nature of projected imagery suggests that handheld projectors may be an ideal platform for mobile social games, which encourages ad hoc participation and initiates social interaction between strangers.

To explore interactions of the future with current technology, we used a commercial motion tracking system to track the projectors with high precision and low latency. However, we anticipate upcoming wireless location tracking systems such as indoor GPS (possibly combined with on-projector sensors such as tilt sensors) will soon enable such tracking more cheaply and ubiquitously, so as to allow our designs to be widely deployed in near future.

CONCLUSIONS AND FUTURE WORK

This work explored concepts and techniques to support interaction between multiple co-located users using handheld projectors. Interpersonal communication and collaboration may be supported more intuitively and efficiently compared to current handheld devices. Informal user feedback indicated that our designs were promising. Our work is the first systematic exploration of the design space of multiple handheld projectors, and may provide a basis for further investigations in this area.

In addition to the general features we developed, in the future we would like to identify and experiment with higher-level collaborative applications supported by handheld projectors. We are also interested in empirically investigating how social protocols between people may evolve with the usage of handheld projectors. Finally, we plan to extensively explore the rich design space of mobile social games using handheld projectors, which may change the way people currently think of games.

ACKNOWLEDGMENTS

We thank John Hancock, Noah Lockwood, Daniel Vogel, Khai Truong, Tomer Moscovich, Ryan Schmidt, Michael Jurka, Gerry Chu, and colleagues at the Dynamic Graphics Project lab (www.dgp.toronto.edu) at the University of Toronto and Mitsubishi Electric Research Laboratories.

REFERENCES

1. Barkhuus, L., Chalmers, M., Tennent, P., Hall, M., Bell, M., Sherwood, S., and Brown, B. (2005). Picking Pockets on the Lawn: The Development of Tactics and Strategies in a Mobile Game. *UbiComp*. p. 358-374.
2. Baudisch, P., Good, N., and Stewart, P. (2001). Focus plus context screens: combining display technology with visualization techniques. *ACM UIST Symposium on User Interface Software and Technology*. p. 431-440
3. Beardsley, P., Baar, J.V., Raskar, R., and Forlines, C. (2005). Interaction using a handheld projector. *IEEE Computer Graphics and Applications*, 25(1). p. 39-43.
4. Bell, M., Chalmers, M., Barkhuus, L., Hall, M., Sherwood, S., Tennent, P., Brown, B., Rowland, D., and Benford, S. (2006). Interweaving mobile games with everyday life. *ACM CHI Conference on Human Factors in computing systems*. p. 417-426.
5. Bier, E., Stone, M., Pier, K., Buxton, W., and DeRose, T. (1993). Toolglass and Magic Lenses: The see-through interface. *ACM SIGGRAPH Conference on Computer Graphics and Interactive Techniques*. p. 73-80.
6. Blaskó, G., Coriand, F., and Feiner, S. (2005). Exploring Interaction with a Simulated Wrist-Worn Projection Display. *ISWC'05 IEEE International Symposium on Wearable Computers*. p. 2-9.
7. Brignull, H., Izadi, S., Fitzpatrick, G., Rogers, Y., and Rodden, T. (2004). The introduction of a shared interactive surface into a communal space. *ACM CSCW Conference on Computer Supported Cooperative Work*. p. 49-58.
8. Cao, X. and Balakrishnan, R. (2006). Interacting with dynamically defined information spaces using a handheld projector and a pen. *ACM UIST Symposium on User Interface Software and Technology*. p.225-234.
9. Hilliges, O. and Terrenghi, L. (2006). Overcoming mode-changes on multi-user large displays with bi-manual interaction. *MU3I Workshop on Multi-User and Ubiquitous User Interfaces*.
10. Hinckley, K., Ramos, G., Guimbretiere, F., Baudisch, P., and Smith, M. (2004). Stitching: Pen gestures that span multiple displays. *AVI Conference on Advanced Visual Interfaces*. p. 23-31.
11. Izadi, S., Brignull, H., Rodden, T., Rogers, Y., and Underwood, M. (2003). Dynamo: a public interactive surface supporting the cooperative sharing and exchange of media. *ACM UIST Symposium on User Interface Software and Technology*. p. 159-168.
12. Joffe, B. (2005). Mogi: Location and Presence in a Pervasive Community Game. *Ubicomp Workshop on Ubiquitous Gaming and Entertainment*.
13. Morris, M.R., Ryall, K., Shen, C., Forlines, C., and Vernier, F. (2004). Beyond "social protocols": multi-user coordination policies for co-located groupware. *ACM CSCW Conference on Computer Supported Cooperative Work*. p. 262-265.
14. Park, D.G., Kim, J.K., Sung, J.B., Hwang, J.H., Hyung, C.H., and Kang, S.W. (2006). TAP: touch-and-play. *ACM CHI Conference on Human Factors in computing systems*. p. 677-680.
15. Rapp, S., Michelitsch, G., Osen, M., Williams, J., Barbisch, M., Bohan, R., Valsan, Z., and Emele, M. (2004). Spotlight Navigation: Interaction with a handheld projection device. *International Conference on Pervasive Computing, Video paper*.
16. Raskar, R., Baar, J.v., Beardsley, P., Willwacher, T., Rao, S., and Forlines, C. (2003). iLamps: geometrically aware and self-configuring projectors. *ACM Transactions on Graphics*, 22(3). p. 809-818.
17. Raskar, R., Beardsley, P., Baar, J.v., Wang, Y., Dietz, P., Lee, J., Leigh, D., and Willwacher, T. (2004). RFIG lamps: interacting with a self-describing world via photosensing wireless tags and projectors. *ACM Transactions on Graphics*, 23(3). p. 406-415.
18. Rekimoto, J. (1997). Pick and drop: A direct manipulation technique for multiple computer environments. *ACM UIST Symposium on User Interface Software and Technology*. p. 31-39.
19. Scott, S.D., Sheelagh, M., Carpendale, T., and Inkpen, K.M. (2004). Territoriality in collaborative tabletop workspaces. *ACM CSCW Conference on Computer Supported Cooperative Work*. p. 294-303.
20. Shen, C., Vernier, F., Forlines, C., and Ringel, M. (2004). DiamondSpin: An extensible toolkit for around the table interaction. *ACM CHI Conference on Human Factors in Computing Systems*. p. 167-174.
21. Shoemaker, G. and Inkpen, K. (2001). Single display privacyware: Augmenting public displays with private information. *ACM CHI Conference on Human Factors in Computing Systems*. p. 522-529.
22. Simon, A. (2005). First-person experience and usability of co-located interaction in a projection-based virtual environment. *ACM Symposium on Virtual Reality Software and Technology*. p. 23-30.
23. Sugimoto, M., Miyahara, K., Inoue, H., and Tsunesada, Y. (2005). Hotaru: Intuitive Manipulation Techniques for Projected Displays of Mobile Devices. *INTERACT 2005*. p. 57-68.
24. Wigdor, D., Leigh, D., Forlines, C., Shipman, S., Barnwell, J., Balakrishnan, R., and Shen, C. (2006). Under the table interaction. *ACM UIST Symposium on User Interface Software and Technology*. p. 259-268.
25. Wu, M. and Balakrishnan, R. (2003). Multi-finger and whole hand gestural interaction techniques for multi-user tabletop displays. *ACM UIST Symposium on User Interface Software and Technology*. p. 193-202.
26. Yatani, K., Tamura, K., Hiroki, K., Sugimoto, M., and Hasizume, H. (2005). Toss-it: intuitive information transfer techniques for mobile devices. *Extended Abstracts of the ACM CHI Conference on Human Factors in Computing Systems*. p. 1881-1884.