

Handheld Projectors for Mixing Physical and Digital Textures

Paul Beardsley

Clifton Forlines

Ramesh Raskar

Jeroen VanBaar

Mitsubishi Electric Research Labs, Cambridge MA*

Abstract

Handheld projectors offer a new type of display modality, not tied to a physical screen or to a fixed projection area, yet providing a larger display than is available from a handheld device with fixed screen. This paper begins with a review of our prototype handheld projector, and describes our work on interaction using a cursor that can be tracked across the projection. The most immediate use for such a device is to support existing applications like web-browsing. We show examples of this type of application.

But there is a broader question too - does a handheld projector support new types of application that are not available with a physical screen or a fixed projection? We describe two applications that illustrate how the combination of handheld projection plus interaction supports ways to interact with the physical world that are much more natural than via conventional displays.

1 Introduction

Projectors are becoming smaller, while advances in self-calibration mean that an arbitrarily placed projector can automatically produce a projection that is keystone-corrected and upright [Sukthankar et al. 2000; Raskar and Beardsley 2001]. Portability and self-calibration will be key factors in the increasing deployment of projectors in opportunistic settings. But this is still a familiar type of scenario - a fixed projector making a fixed display. The use of projectors in handheld devices promises to be a much more significant innovation. A driving force is that the decreasing size of cellphones and PDAs is in direct conflict with the need for a display that is easily visible to the user. An attached projector can address this by being compact in size while still producing a good-sized display.

2 Prototype Handheld Projector

Current commercial data projectors are close to being small enough for handheld use. Figure 1 shows a prototype that we built based on the Plus V-1080 projector (original release date 2002).



Figure 1: Prototype handheld projector.

The device consists of

- a Plus V-1080 projector, 1024x768 pixels, 60Hz framerate, dimensions 7x5x1.5 inches, weight 1kg.
- a Basler A602F camera, 640x480 pixels, 100Hz framerate.
- four rigidly attached laser pens.
- hand-grip with click buttons under the trigger finger.
- umbilical to a computer.

The device weighs about 2.5 pounds, which would of course be excessive for a truly portable device, but it is nevertheless a genuinely handheld projector and has enabled a variety of experiments. We note that much smaller projection devices do exist - for example, the Canesta projected keyboard [www.canesta.com 2004], and the Siemens Mini-Beamer cellphone [www.siemens.com 2002].

The very first requirements that arise when considering handheld projection are keystone correction and image stabilization i.e. fixing the projection on the surface even under hand jitter. Keystone correction needs an estimate of the orientation of the projector relative to the surface. Image stabilization needs the full pose - orientation and location - of the projector relative to the surface.

Keystone correction is therefore the easier case, and one way to find the orientation is to use a calibrated projector-camera pair. The camera observes the current projection, and thereby infers orientation, and there is an iterative update of the projection at each time-step to remove keystoning. Image stabilization can also utilize a projector-camera pair, although there needs in addition to be some fixed texture on the display surface to provide a fixed external coordinate frame for the device. Of course, there are some well-known types of motion-sensor that could be attached to a projector that would avoid the need for physical texture, for example ultrasound [www.isense.com 2000], but we have concentrated on projector-camera systems as a likely route to cheap, self-contained devices.

3 Interactive Projection

Our initial work was built on the observation that purely passive handheld projection is of very limited use, and even a 'slide-show' mode that allowed a user to tab through a canned sequence of projections (using say next- and previous- buttons on the handheld projector) cannot support interesting applications. So how can we transpose familiar interaction methods from the desktop or handheld device to do interaction with a projection? The problem differs from laser pointer interaction with a display in a fixed installation [Olsen, Jr. and Nielsen 2001], because we are concerned with more opportunistic projection in arbitrary environments.

It might seem straightforward to provide mouse-interaction by putting a touch-pad on the handheld projector. But this adds bulk to the device, which we are trying to avoid. It also implies two-handed use. And there is a resolution issue, in doing fine control of a cursor for a relatively large projection using a small touch-pad.

Our technique avoids these problems, and allows a cursor to be tracked across the projection by a direct pointing motion of the handheld projector. Once the cursor is at the desired location, items are selected using buttons on the trigger finger of the device in Figure 1.

The approach is as follows. First assume there is some way to compute the 3D motion of the projector relative to the display sur-

*email:[pab,forlines,raskar,jeroen]@merl.com

face. We return to the details later. Knowing the projector’s 3D location at each instant, we can create any desired projection on the surface – in particular we can factor out the projector motion to create a projection that is static on the display surface. Figure 2 illustrates how the the projector image plane is modified under projector motion to maintain a static projection (green area) on the surface. Now consider what happens if part of the the projector image plane always has the same fixed graphic. Figure 2 shows a cursor graphic at the center of the projector image plane. The effect for the user is to see a static projection with a cursor moving across it in response to pointing motion of the projector.

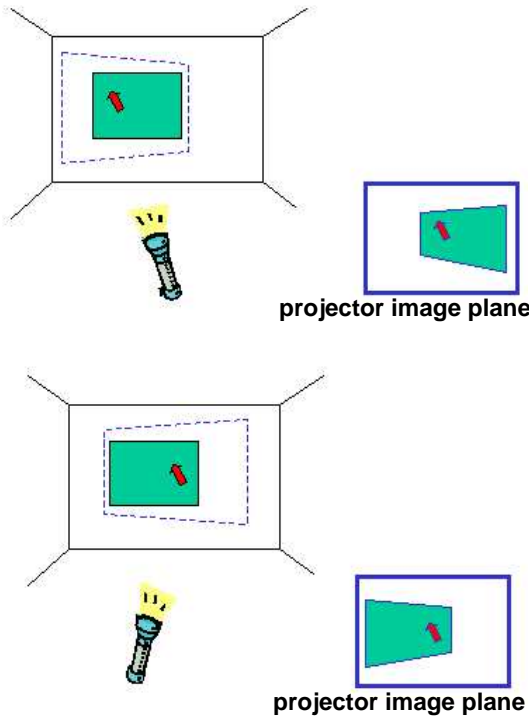


Figure 2: The handheld projector points toward the left- and then right-side of the display area. The main projection (green) stays static on the display surface, due to the adjustments made on the projector image plane. The cursor (red arrow) moves in direct correspondence with the motion of the projector, because the cursor graphic occupies a fixed position on the projector image plane. The user sees a static projection with the cursor tracking across it. The blue dashed line indicates the boundary of the full projection area.

A constraint of our technique is that the projection data can only occupy part of the projector image plane. But one-handed pointing motion feels natural as a way to direct the cursor. And it leaves the fingers free to hover over the mouse buttons, ready for a button click. This is in contrast to the touch-pad where a single finger must do a context-switch between touch-pad and buttons, or different fingers are brought into play, but which is anyway an inherently more complex interaction for the user.

4 Basic Techniques

Here we provide more detail about the algorithms that support the functionality in the previous section – Sections 4.1 and 4.2 describe image stabilization on the display surface under motion of the handheld projector. Section 4.3 describes interactive projection.

4.1 Image Stabilization

We have so far used fiducials (distinctive visual markers) on the display surface to define a coordinate frame for image stabilization. Figure 3 illustrates the basic task. The goal is to place the projection in the target area shown by a red dashed line, defined relative to four fiducials. This is achieved by using the camera to sense the fiducials, and hence to infer the target area in camera image coordinates. The target area is then transformed to projector image coordinates. Finally the projection data is mapped to these coordinates on the projector image plane, giving correct placement of the projection. The whole process is repeated for each new time-step. The detailed steps are

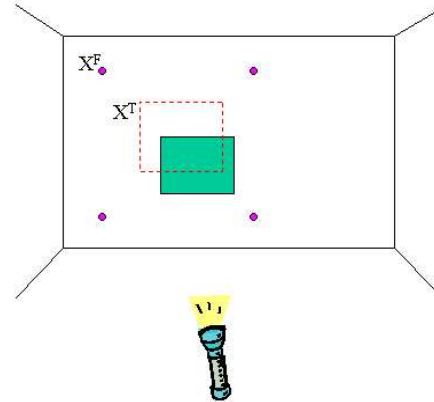


Figure 3: The purple points are fiducials on the display surface. The red dotted-line indicates the target area for a projection - this is defined relative to the fiducials (the target area is not physically marked on the surface). The green area indicates the current projection. Projection motion, e.g. due to hand-jitter, causes a discrepancy between the target area and the position of the projection, and a correction is applied at each new time-step.

- Detect the camera image coordinates c_i^F of the fiducials. Compute the homography H_{SC} between the display surface and the image, using the known surface coordinates X_i^F of the fiducials, and the points c_i^F .
- Apply H_{SC} to the known surface coordinates X_i^T of the target area, to obtain the camera image coordinates c_i^T of the target area.
- Detect the camera image coordinates c_i^P of the vertices of the current projection. Compute the homography H_{CP} between the camera image plane and the projector image plane, using the points c_i^P , and the known projector image coordinates of the vertices of the projection.
- Apply H_{CP} to c_i^T to generate p_i^T , the coordinates of the target area on the projector image plane.

The algorithm is projective. An alternative approach would be to work in euclidean space, explicitly computing the location and orientation of the projector relative to the display surface. This is sometimes useful (for example if we wish to make explicit the projector’s distance from the display surface, for some other use like determining content). But the projective algorithm is lighter computationally.

4.2 Using the Laser Pens

In practice, there is a modification to step (3) above and the computation of c_i^P . Detecting c_i^P in the camera image is difficult when the projection has arbitrary appearance. Instead we detect laser spots

from the four laser pens on the device, which is straightforward because the spots are bright. The laser rays are not concurrent and are not coincident with the focal points of the camera or projector. But we can still use the observed laser spots to compute H_{CP} provided there has been some pre-calibration.

First note that a laser ray projects to a line on the camera image plane and a line on the projector image plane. Also note that there is a line homography that describes how a point on a laser ray transforms between the camera and projector image planes. If we have the line homographies for each of the four laser rays, then we can transform the observed laser spots from the camera image plane to the projector image plane, and then use the four camera-projector correspondences to compute H_{CP} .

The pre-calibration method for computing the line homography for a specific laser ray is as follows. Point the projector-camera-laser system at a planar surface, with the projector showing a fixed pattern. Compute a homography H_{CP} between the camera and the projector using the pattern. Detect the laser spot c^L for the given laser ray on the camera image plane. Store the corresponding points c^L and $H_{CP}.c^L$. Repeat with the projector-camera at a further two or more distinct positions relative to the surface. Use c_1^L, c_2^L, \dots to compute the line that is the projection of the laser ray on the camera image plane. Use $H_{CP1}.c_1^L, H_{CP2}.c_2^L, \dots$ to compute the line that is the projection of the laser ray on the projector image plane. Use the correspondences to compute the line homography between the two lines.

Fuller details in [Beardsley et al. 2004].

4.3 Interactive Projection

At each time-step the projector image plane is updated to achieve a static projection on the display surface, as described in Section 4.1. In conjunction with this update, the center pixels of the projector image plane are overwritten with a graphic for the cursor.

The effect on the projector image plane is that the main projected content is continually being updated, to factor out projector motion, while the cursor graphic is fixed. The effect on the display surface is that the main projected content appears static, while the cursor moves across the surface in direct correspondence with projector motion.

5 Applications - (1) Conventional Personal Applications

The initial type of applications that we investigated are conventional desktop applications, transposed to a projected setting. Figure 4 shows an example of web-browsing. The update rate for the stabilization is about 70Hz. This is sufficient for good stabilization and a natural cursor interaction.

The attached video material shows an interaction with a dynamic game application. As an informal observation, the ability to 'pick up' parts of a projection gives them a very physical feel. This may partly arise from the fact that mouse interaction on a conventional, physical display is indirect - the mouse motion on the pad is transformed to a cursor motion on the screen. But with interactive projection, the pointing motion of the projector directly guides the cursor, so the feeling of a physical connection is much stronger.

6 Applications - (2) Projected Augmented Reality

The second class of applications that we have investigated are for augmented reality of a physical surface. Figure 5 shows projection of a room map next to a fuse box. Cursor selection of a specific




Figure 4: Interacting with a projection of a web browser. The main window remains static - note its position relative to the black fiducials on the wall. The cursor is tracked from its central position on the browser over to the left-side, by a pointing motion of the projector. For illustration purposes the full projector image plane is made visible by making the background red. Normally the background would be black and barely visible. The cursor is overly large, also for purposes of illustration.

room causes the corresponding fuse to be highlighted in the bank of fuses. Interactive projection has many possibilities for AR - for example it allows the augmentation overlay to be dragged to a desired position on the physical surface.




Figure 5: Augmentation of a fuse box. At left, a projected room map. The user tracks the cursor to a location on the map and clicks to invoke a projected highlighting of the corresponding fuse. The outer black rectangle defines the coordinate frame for the projection. The colored circular badges are used to determine appropriate projected content for this object.

7 Applications - (3) Mixing Physical and Digital Textures

We provide two examples of our current work on mixing physical and digital textures. The first is an application for attaching electronic data to selected positions in an environment, and the second is copy-paste applied to physical texture.

7.1 Attaching Electronic Data to Physical Texture

Figure 6 show how the user can do a mouse hold-and-drag operation to define a rectangular region-of-interest (ROI) on a physical surface. In this case there are no fiducials to provide a coordinate frame on the surface. But the texture of the object itself is sufficient to define a homography between successive camera images and hence place all successive cursor positions in a single camera coordinate frame. Knowing the cursor track in this camera coordinate frame, we can transform to projector image coordinates as before, and then present the user with the projected rectangle to indicate the selection.




Figure 6: Selecting a physical region of interest by mouse click-and-drag, analogous to the familiar windows operation.

Such a selection can be used to attach and retrieve electronic notes at arbitrary (textured) locations in an environment. The steps are

- The user defines a ROI around the required physical location.
- The camera image data for the ROI – a ‘texture key’ - is stored along with the electronic data which is to be attached.
- On a subsequent visit, the user selects a ROI at the same physical location.
- The image data for the new ROI is matched against all stored texture keys by template-matching, and the data associated with the best match is projected onto the surface.

The same functionality could be achieved by mousing on an image of the surface on a physical display. But the issues are the convenience of the interaction and the direct colocation of physical and digital data, versus indirect operation on a handheld display.

7.2 Copy-Paste of Physical Texture

Figure 5 shows a copy-paste operation. There is a control panel at left. The required texture is captured via a selected region-of-interest as in the previous section, and the texture is then projected onto a clean surface. This is one way to deal with the problem of how to project augmented reality onto a surface that does not support projection - we instead copy-paste the texture and show the augmented reality overlay on the pasted projection. Copy-paste can also allow different parts of an environment to be arranged side-by-side for comparison. As a practical matter, this kind of application would benefit from a high-resolution camera so that the copy-paste can show a high level of detail. Again this type of functionality can be done with a physical display. But projection brings some unique attributes, for example easy shared viewing between users, or the ability to attach a copy-paste to a fixed physical location for continued reference.




Figure 7: At left, a control panel; at right, selected texture is captured and redisplayed elsewhere.

8 Why Projection?

Projection is not suited for all applications. For example, visibility degrades under some lighting conditions. But handheld projection

also has some unique strengths, and the question is whether there is a class of applications for which it could prove the most appropriate technology. Our work is an investigation of mobile, opportunistic projection and we list some of its advantages. Firstly, handheld projection has the potential to make many everyday surfaces into displays. Device size is not linked to display size, making projectors a good technology for a portable handheld display. Furthermore, projection has some specific advantages when making augmented reality overlays – it avoids a context-switch between fixed display and physical scene, and avoids resolution problems when viewing or interacting with a scene via a small handheld display. Finally we showed how interactive projection can provide a direct interaction with physical features.

9 Future Work

- We plan to replace the current visible-light laser pointers on the device with infra-red lasers. The cost of the current laser pens was about \$30 each and the likely cost for infra-red lasers is less than double that.
 - We would like to avoid the current use of physical fiducials on the wall. One attractive option is a separate projector that is placed at a fixed position and projects a fixed pattern, maybe also IR. Thus we have a two device system rather than a completely self-contained device, but both the projector-camera and the fixed projector are potentially very small devices. As a real product one could envisage them snapping together when not in use.
 - There are several other areas awaiting investigation - using interactive projection to directly create augmented reality content on an object; using gesture to control projected content; integrated projection from multi-projectors etc.

10 Conclusion

Handheld projectors have the potential to extend the way we think about display, moving beyond fixed screens to allow opportunistic display on everyday surfaces around us. This paper has described our technique for interacting with projected data, and suggested how it can support new types of application that mix physical and digital texture. There are other ways to achieve the same functionality, but interactive projection provides a very natural and minimal interaction for selecting a physical area. It offers a novel reversal of the familiar desktop metaphor - while windows systems utilize icons like movable folders and trash-barrels, we are borrowing familiar mouse-interactions like selection and copy-paste and showing how they support interaction on physical surfaces.

References

- BEARDSLEY, P., VANBAAR, J., AND RASKAR, R. 2004. Augmenting a Projector-Camera Device with Laser Pointers. TR 2004/035, MERL.
- OLSEN, JR., D. R., AND NIELSEN, T. 2001. Laser pointer interaction. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, ACM Press, 17–22.
- RASKAR, R., AND BEARDSLEY, P. 2001. A Self-Correcting Projector. In *IEEE Comp. Soc. Conference on Computer Vision and Pattern Recognition (CVPR'01)*.
- SUKTHANKAR, R., STOCKTON, R., AND MULLIN, M. 2000. Automatic keystone correction for camera-assisted presentation interfaces. In *Proc. Intl. Conf. Multimedia Interfaces '00*.
- WWW.CANESTA.COM, 2004. Canesta projection keyboard.
- WWW.ISENSE.COM, 2000. InterSense: ultrasound-based tracker.
- WWW.SIEMENS.COM, 2002. Siemens Mini-Beamer.