

3 High-Resolution Interactive Displays

Mark Ashdown¹, Philip Tuddenham², and Peter Robinson²

¹Thales Research and Technology, Worton Drive, Reading RG2 0SB, UK.
mark@ashdown.me

²University of Cambridge Computer Laboratory, 15 JJ Thomson Avenue, Cambridge,
CB3 0FD, UK. {Philip.Tuddenham, Peter.Robinson}@cl.cam.ac.uk

Abstract

Tabletop displays are mostly used for casual applications that do not require intricate graphics or precise manipulation. Browsing photographs and maps are common applications. A higher resolution is required to support work involving detailed graphics and text. A display the size of a desk, with the resolution of a typical LCD monitor, will have around 14 megapixels. Tabletop displays are usually constructed from projectors, and the only way to achieve this size and resolution is to combine multiple projectors. We present techniques from multi-projector display walls and adapt them for tabletops. These high-resolution displays also require high-resolution input, and although touch is simple and natural, better accuracy can generally be achieved using a pen. We also review technologies for pen input on tabletops.

Introduction

Large displays have compelling benefits. For individuals they allow materials to be spread out and perused as on a physical desk, thus exploiting the kinaesthetic sense and styles of working used in traditional workspaces; they facilitate multitasking; and they have been shown to have cognition and productivity benefits [1]. For groups they allow participants to work together while adhering to social protocols of personal space without formal turn-taking, and support consequential communication: allowing users to remain aware of the actions of others through their peripheral vision.

Most tabletop displays have low resolution. In recent years there has been much research interest in tabletop displays, and hardware of various forms has been constructed. One of the key features of tabletop displays is that they have diagonal sizes of around 30 to 80 inches, much larger than a conventional monitor. Most designs use a single-projector front- or rear-projected display, which enables the large size, but typically have low resolution in terms of pixels per inch. This is

adequate to enable the standard tabletop demo application, photo browsing, but there are many applications that would benefit from a tabletop implementation but require high resolution to show detailed graphics.

Many applications require high-resolution to display the necessary information. Examples for which people currently use their desktop computers include document review, data analysis, and computer-aided design. Applications that have been implemented on a tabletop to demonstrate the use of high resolution include web browsing (the text must be legible), collaboration over documents (text and written annotations must be legible), programming (a typical integrated development environment display is very dense), unmanned aerial vehicle control (air-traffic control displays are dense), and command and control (maps with roads, buildings, symbols) [2]. Chapter [Isenberg et al] discusses information visualization on tabletops, which is another application that will benefit from high resolution.

This chapter describes techniques that have been used to achieve high-resolution output and input for tabletop displays. In particular, the tiling of projectors using techniques developed for multi-projector display walls, which has proved to be a powerful approach, and pen input, which complements touch input by providing more precision.

Background

Work on tabletop displays at the University of Cambridge Computer Lab started with the DigitalDesk [3], and continued with Origami [4], Escritoire [5], and T3 [2]. A common goal in these projects is the creation of a human-computer interface like a traditional desk, where items are arranged freely over a large surface and modified using a pen, as with traditional pen and paper. The goal has always been to provide high resolution displays to enable detailed work [6]. Many other tabletop systems have been constructed. Continuing advances in display technology, in particular the commoditisation of digital projectors, have brought such displays within the reach of researchers interested in the implications for human-computer interaction.

Resolution

As described in the introduction above, most tabletop displays have relatively low-resolution output. If they provided high resolution, new classes of applications could be implemented, and subjected to the research on novel multi-touch and pen-based interfaces that is being conducted on tabletops.

The Display Resolution section below (page 4) defines what we mean by low and high resolution. It defines a lower limit by considering the resolution necessary to display detailed content, and an upper limit by considering the properties of human vision and the more realistic resolutions of current LCD monitors. It then presents an analysis of how this range of resolutions can be provided by combining projectors, with particular attention to the cost of the hardware.

Display Wall Techniques

Since DigitalDesk [7] introduced the idea of projecting down onto a table to produce a 'tabletop display', many researchers have explored this type of computer interface, and the use of projectors as personal displays rather than devices for making formal presentations to an audience. Despite some issues such as lack of privacy, and generation of noise and heat, this new use for projectors has been shown to have benefits for individual work and collocated collaboration [8].

Improvements in calibration techniques and the falling cost of projectors have led to the construction of many multi-projector display walls for scientific visualization. They tend to be large expensive installations. The Princeton Scalable Display Wall [9] was one of the first large scale walls, which demonstrated a 4×2 projector array, followed by a 6×4 (24 projector) array on custom-made adjustable bases, driven by a cluster of PCs. Various other multi-projector display walls have followed [10].

When combining multiple projectors, the aim is to make a single seamless display. Achieving seamlessness involves compensating for two types of problem: geometric and photometric [10]. Geometric compensation fixes the position of every pixel in the display, while photometric compensation fixes its colour. Compensation techniques such as perspective warping, edge blending, automatic calibration using computer vision, and intensity and full-colour radiometric measurement and compensation have been studied extensively in the context of display walls.

Compared to display walls, tabletop displays have a different emphasis: the display of as many pixels as possible from a single computer, in a self-contained horizontal form. However, many of the techniques from display walls can also be applied to tabletops with either front or rear projection. The Multi-Projector Display Techniques section below (page 11) reviews these techniques, shows how some of them have been applied to tabletops, and lists others that could be applied.

Input

An important aspect of the allure of tabletop displays is the naturalness and simplicity of direct input, rather than indirect control with a mouse or touch pad. In particular, touch and multi-touch input are driving many of the new interfaces that are being created. Frustrated total internal reflection (FTIR) and diffuse illumination have been used in many tabletop displays because they are simple, robust, and combine well with rear projection. Other sensing technologies based on resistance, capacitance, and computer vision, have also been developed. These touch input technologies are described in Chapter [\[Interactive multi-touch surfaces\]](#).

A high-resolution display is best exploited with an accurate input method. Chapter [\[chapter by Wigdor et al.\]](#) discusses the issue of inaccuracy in touch input and proposed interaction techniques to alleviate that problem. [The High-Resolution Input section below \(page 22\)](#) below addresses another approach to the issue – giving the user a pen to allow accurate input and writing. It lists pen-input technologies, followed by descriptions of the issues intrinsic to pen input on tabletops.

Display Resolution

Before delving into techniques for creating high-resolution displays, we must define what resolution is, and how high we wish it to be. The resolution of a display is the number of pixels it has per unit distance. We assume a regular grid of square pixels, and state resolutions in pixels per inch (ppi), which is the convention for monitors and printers (1 inch = 25.4 mm). In this section we derive a resolution range for which to aim in new devices, and compare this to existing display devices and tabletop systems.

Minimum Resolution

A lower bound on the resolution that is required of a display can be defined by experimenting with various types of graphical content. Photographs degrade gracefully as resolution is reduced, but text imposes a hard limit, below which it is illegible.

Text on the Web has typically been displayed using a 10pt or 12pt font [11], and these are also typical font sizes for printed documents. A resolution of at least 2 pixels per millimetre (48 ppi) has been shown to be required to read 12pt text [2], which corresponds to approximately 60 ppi for 10pt text. Tullis et al. [12] imposed a minimum font size of 6pt on Microsoft Windows. Assuming the default

96 ppi, this is equivalent to a 10pt font at 57.6 ppi, thus roughly agreeing with the previous 60 ppi figure. We will therefore use 60 ppi as a lower limit on the resolution required on a tabletop display.

Maximum Resolution

To define an upper bound on the resolution of a display, one can consider human visual acuity. This results in a resolution above which it is pointless to go, because a human viewer will not be able to tell the difference.

‘Normal’ human visual acuity is widely taken to be 1.0 arc minute [13], that is, a person can discern a grating test pattern when the distance between the lines is 1 minute of arc (1/60 of a degree). Actually, the acuity of most young observers is somewhat higher, at 0.59 arc minutes [14]. This angle corresponds to one cycle of the grating, which requires two pixels. Assuming that a tabletop display is used at arm’s length, which is approximately two feet (61 cm), this leads to spatial resolutions of around 300 to 500 ppi.

The lower figure of 300 ppi is the standard resolution for professional printing. Note that printers often have higher dots-per-inch (dpi) ratings, because they use halftoning to display grey levels using patterns of binary dots. One might wish to create the ultimate display by matching the upper figure of 500 ppi, but this would probably be a waste of resources because most people would not be able to distinguish such fine detail. This prompts us to consider the desired resolution required for typical content.

Resolution of Monitors

Calculating the desired resolution of a display directly from the limit of visual acuity results in a somewhat excessive resolution. A more realistic one is obtained by matching commonly available computer monitors. Monitors display images via a regular grid of coloured pixels. The actual resolution of a monitor depends on the hardware, and is typically not known automatically by the software. The computer’s operating system therefore assumes a certain number of pixels per inch so that it can display items, such as a word processed document, at ‘life size’. Mac OS assumes the print standard of 72 ppi, the X Windows System defaults to 75 ppi or 100 ppi, and Microsoft Windows and Ubuntu Linux assume 96 ppi.

LCD monitors currently (2009) on the market have diagonal sizes of 15 to 30 inches, and display resolutions between 1024×768 and 2048×1152. Resolutions, also known as *dot pitch*, are between 80 and 110 ppi. Rather than go above 110 ppi, manufacturers choose to increase the physical size of the display. We therefore choose 110 ppi as the practical upper limit of resolution that will be

useful for a tabletop display, rather than the higher value in the subsection above, which is derived from the limit of human visual acuity.

Resolution of Projectors

The term resolution is used in two ways regarding projectors. One is *native resolution*, which states the number of pixels in the imaging element. The other is the density of pixels on the screen. In this section we will consider only the native resolution, because the pixel density is determined by the throw distance and is therefore dependent on a particular physical configuration.

The two main values that characterize the output of a projector are the native resolution (number of pixels), and the light output (brightness). For creating tabletop displays we are interested in good price-performance for both.

The native resolutions available in reasonably priced projectors are dictated by market forces. For a long time the only options were 800×600 and 1024×768. Table 3.1 below shows the most popular resolutions that are available now, with their names, width and height in pixels, and resulting number of megapixels. For example, an XGA projector has $1024 \times 768 = 786,432$ pixels = 0.79 megapixels.

Name	Width (pixels)	Height (pixels)	Megapixels
SVGA	800	600	0.48
XGA	1024	768	0.79
WXGA (720p)	1280	720	0.92
SXGA+	1400	1050	1.47
UWXGA (1080p)	1920	1080	2.07

Table 3.1 Common native resolutions of projectors

Fig. 3.1 illustrates the relationship between native resolution, brightness, and current (2009) market price, of a representative set of projectors from 20 manufacturers. Almost all the projectors had one of the five native resolutions listed in Table 3.1. The remaining few were near the centre of the distribution, so they were omitted to simplify the graph. Brightness of projectors is measured in ANSI lumens [10].

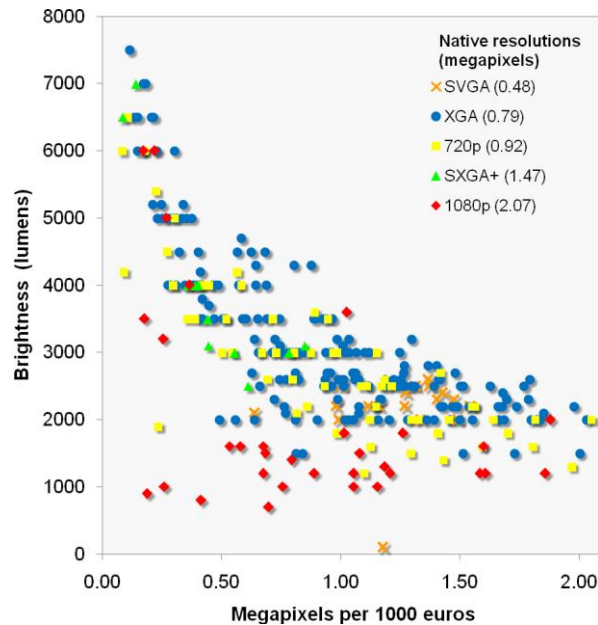


Fig. 3.1 Brightness of projectors against megapixels-to-price ratio. Native resolutions are listed in Table 3.1.

The horizontal axis of Fig. 3.1 indicates the pixel economy of the projectors, with higher numbers of megapixels per 1000 euros indicating better economy. The concave curve of the points shows that there is a trade-off between projector brightness and pixel economy. To make a good-value high-resolution display, one should use projectors on the right of the graph, which means there is a brightness limit of around 2500 lumens. Until recently XGA projectors offered the best pixel economy. These have now been joined by 720p and 1080p projectors, due to the demand for these resolutions driven by the uptake of high-definition television (HDTV). The graph shows that the most economical projectors, purely in terms of pixels, are the XGA, 720p and 1080p ones, with higher resolutions giving slightly lower brightness.

A native resolution of 1080p has more than twice as many pixels as 720p, and almost three times as many as XGA. Because there is a per-projector overhead in creating a multi-projector display, including the cost of graphics cards, cables, physical space, and calibration complexity, 1080p is a good choice. It provides good pixel economy while minimising the number of projectors. Given that 1080p projectors will be used, the number required to make a tabletop display is then determined by the size of the table and the desired resolution.

Resolution of Existing Tabletops

Table 3.2 below lists some tabletop systems that state their size and resolution attributes in their associated publications. The resolutions are calculated along the diagonal to even out any variation between horizontal and vertical.

Tabletop system	Front/ Rear	Native (pixels)			Physical (inches)				Res. (ppi)
		Horiz	Vert	Diag	Horiz	Vert	Diag	Mpxl	
DigitalDesk [7]	F	1120	780	1365	18	12	22	0.87	63
DigiTable [15]	F	1024	768	1280	34	26	42	0.79	30
DViT Table (med. res.) [16]	F	2048	1024	2290	60	48	77	2.10	46
DViT Table (high res.) [17]	R	2800	2100	3500	60	48	77	5.88	63
Escritoire (fovea) [18]	F	1024	768	1280	17	12	20	0.79	63
Escritoire (periphery) [18]	F	1024	768	1280	48	36	60	0.79	21
i-m-Top (fovea) [19]	R	1280	720	1469	12	8	14	0.92	106
i-m-Top (periphery) [19]	R	1280	768	1493	47	32	57	0.98	26
InteracTable [20]	R	1024	768	1280	33	26	42	0.79	30
Lumisight [21]	R	1024	768	1280	16	16	22	0.79	57
Origami [4]	F	1024	768	1280	16	12	20	0.79	64
Surface [22]	R	1024	768	1280	26	15	30	0.79	43
T3 [23]	F	3072	1536	3435	48	36	60	4.72	57

Table 3.2 Size and resolution of some tabletop displays. Values are rounded to the nearest integer, excluding megapixels. These displays have between one and six projectors each. The ‘front/rear’ column indicates front or rear projection.

The tabletop systems in Table 3.2 have resolutions between 21 and 106 ppi, with most of them below our 60ppi lower limit. This reflects the fact that most tabletop displays to date have been designed for informal graphical tasks such as photo browsing, or fun tasks, such as simple games, that do not display detailed information. The availability of higher resolutions will allow more serious applications.

Several systems have combined projectors to achieve higher resolutions. Escritoire [24] used overlapping projectors to create a display with a high resolution region near the user for detailed work, and a much larger, low-resolution periphery (Fig. 3.2a and Fig. 3.2b). Baudisch’s focus-plus-context screen [25] achieved a similar result in a vertical form. A system called i-m-Top [19, 26] enhanced this concept by using a steerable projector to create a moving fovea that can follow the user’s input (Fig. 3.2c).




Fig. 3.2 Focus-plus-context tabletop displays. (a) Escritoire full view, (b) Escritoire fovea, (c) i-m-Top fovea.

T3, the TableTop Toolkit [2], is an open-source software toolkit for creating multi-projector tabletop displays. Fig. 3.3a shows a 3×2 XGA projector configuration with a total of 4.72 pixels, as listed in Table 1. Fig. 3.3b shows an application that exploits the resolution. The DViT tabletop at the University of Calgary [16] tiled two projectors to create a 2.1 megapixel tabletop display, and a second-generation model [17] uses four projectors to create a 5.88 megapixel display (Fig. 3.3c).

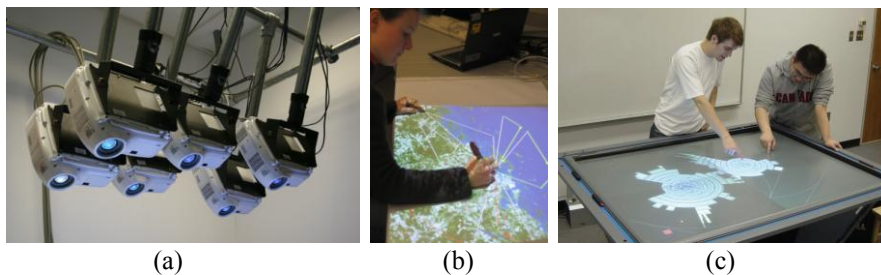


Fig. 3.3 Tiled tabletop displays. (a) Six-projector T3 display, (b) collaborative targeting task using T3, and (c) four-projector DViT display.

Tabletop Resolution Guide

A tabletop display should be a self-contained system where the projectors are driven from a single PC, using the best value projectors to achieve a good price-performance. The analysis of the resolution of projectors above (page 6) advocates using 1080p projectors, which will have around 2000 lumens. Modern PCs can have four PCI express slots, each of which can hold a graphics card with two digital outputs, allowing eight projectors to be driven from a single PC.

The number of projectors determines the total number of pixels on the display, and the resolution varies as the reciprocal of the size. Fig. 3.4 illustrates the relationship between resolution and diagonal size, assuming the display is composed of 1080p projectors. Our minimum and maximum resolutions of 60ppi

and 110ppi are shown, as are the typical size of various types of tables, such as a small coffee table, or an eight-person conference table. For a small coffee table (30-inch diagonal) one projector is sufficient, three are ample. For a six-person conference table, six to eight projectors are required. Note that the calculations that were used to generate Fig. 3.4 are only based on tabletop area. Given a particular projector aspect ratio, a regular array of projectors, and minimal projector overlap, only certain shapes of tabletop display are actually possible.

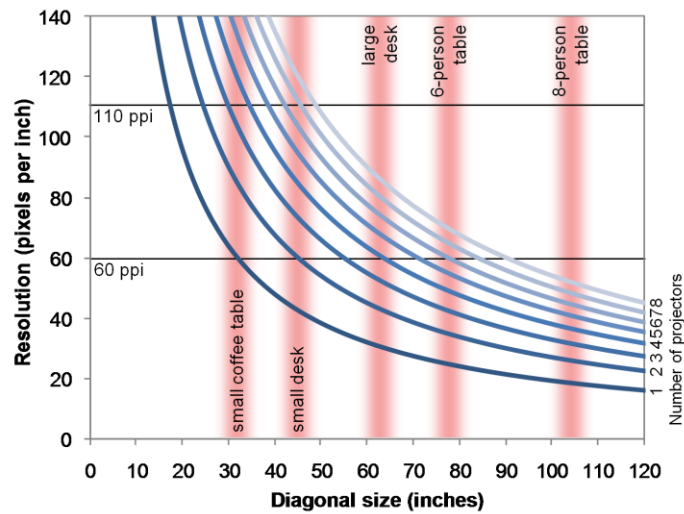


Fig. 3.4 Tabletop display resolution against size, assuming 1080p projectors. The horizontal lines show suggested minimum and maximum resolutions. The vertical regions show sizes of typical tables.

A tabletop display should have a resolution of at least 60 ppi to show detailed content, such as text. 110 ppi is a reasonable upper limit. Eight projectors can be driven from a single PC, which allows acceptable resolutions on tabletops up to a six-person conference table. As an example, Fig. 3.4 indicates that a large desk with a diagonal of 60 inches (1.52 metres) using eight 1080p projectors will have a resolution of 78 ppi, which is comparable to the lower end of LCD monitor resolutions. The cost of the projectors for this display, using the best-value 1080p projector from Fig. 3.1 with a cost of 1.88 megapixels per 1000 euros, will be 8823 euros. This places large high-resolution tabletop displays within the reach of research projects on the applications and interaction techniques suitable for such displays. Consumers will have to wait a bit longer for such devices to become affordable products, but we can work on the software now, in anticipation of the availability of the hardware.

Multi-Projector Display Techniques

Multiple projectors can be combined to create a large high-resolution display, but this requires more than simply placing the projectors next to each other. The projectors must be calibrated so their combined output appears as a single seamless display. Various techniques have been developed over the last decade to address this issue in vertical cinema-style multi-projector displays walls, and many of these can be directly applied to tabletops.

The techniques for calibrating multi-projector displays address two fundamental problems: geometric misalignment and colour variation [10]. When projectors are combined, and are not precisely mechanically aligned, perspective distortion occurs. This can be removed by warping the graphics before displaying them by passing them all through the 3D graphics card, in a process called *geometric compensation*. Colour variation due to variation in output within and between projectors is addressed by *photometric compensation*. These two processes are described below, and between them is a description of *edge blending*, which has both geometric and photometric aspects. This section is concluded with some system issues.

The methods described in this chapter are equally applicable to front and rear projected displays, both of which have been used in various tabletop systems (Table 3.2). Front projection means the space under the display is free, as with a normal table, and can be used with sensing technologies comprising an opaque tablet. Rear projection hides the projectors inside the device to create a more self-contained unit, and avoids occlusion of the projection by the hands of the users. Some of these methods can also be used in tabletops made from other display technologies, such as LCD panels.

Many of the concepts in this section are covered in detail in Practical Multi-Projector Display Design by Majumder and Brown [10]. The interested reader might also want to consult Spatial Augmented Reality by Bimber and Raskar [27], and review the proceedings of the IEEE International Workshop on Projector-Camera Systems (PROCAMS),¹ which started in 2003.

The text below assumes precise meanings for several terms: *characterization* is the measurement of how a device actually performs; *compensation* is the modification of inputs to the device so that the system as a whole behaves as an ideal device; and *calibration* is the combination of the previous two steps.

Geometric Compensation

If a projector is not perfectly aligned to its projection surface, the resulting display will not appear as a perfect rectangle, and will instead be distorted. For a single

¹ <http://www.procams.org/>

projector, small distortions are not a problem for people viewing the display, but when projectors are combined their outputs must align precisely with one another. For this reason, multi-projector displays originally required precise and laborious mechanical alignment [9], but now commodity graphics hardware can be used to compensate for the distortions.

Projective Transformation

A projector is the inverse of a camera. It contains an imaging element and a lens, but the light travels out rather than in. By assuming central projection, the model of a pin-hole camera, we can compensate for the distorting effects of oblique projection by taking the original image to be displayed, and warping it before it is projected [28].

If the projection surface is flat, two-dimensional (2D) points in the projector correspond to 2D points on the surface, and the relationship between these two sets of points is a projective transformation which can be represented by a planar homography [29]. Geometric compensation is achieved by applying a homography to the original image, using the equation

$$\mathbf{x}' = \mathbf{H}\mathbf{x}, \quad (1)$$

where \mathbf{H} is a 3×3 matrix and the vectors \mathbf{x}' and \mathbf{x} are two-dimensional points in homogeneous form [29]. Different forms of matrix \mathbf{H} produce different types of transformation (Fig. 3.5).

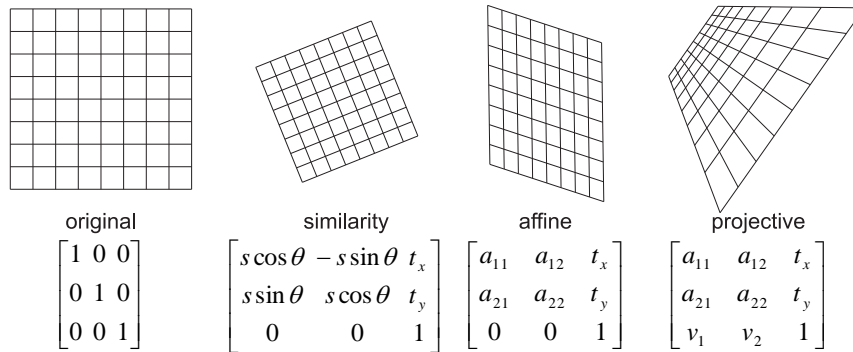


Fig. 3.5 Transformations and their homography matrices [29]. Value s is a scale, θ is a rotation angle, and (t_x, t_y) is a translation. The other variables are other degrees of freedom. The bottom right element is normalized to 1 because the matrix is defined only up to a scale.

Projection from a very acute angle will produce not only geometric distortions, but also differences in focus across the surface. A projective transformation can

remove the geometric distortions but the blurring caused by a projector's finite focal depth is more difficult to remove, so this effect should be avoided where possible.

Characterisation

In the geometric case, characterisation is also known as registration, and consists of measuring the transformations between various two-dimensional spaces. Ashdown [5] uses four such spaces: device-dependent pen input co-ordinates, device-independent tabletop co-ordinates, pixel locations in the projected image, and pixel locations in the texture on the graphics card. For ease of implementation, registration can be performed by selecting a series of projected points using the pen input device, and then the software can calculate a linear least-squares solution for the homography matrix in equation 3.1. If a calibrated camera is present, calibration of the projector can be performed automatically by projecting patterns such as dots or lines, and using image processing to locate them within the camera image [10].

Graphics Hardware

A very efficient way to perform a projective transformation on a typical computer is to put the original image into a texture on the graphics card, and draw it as a polygon. The transformations in Fig. 3.5 can be performed by manipulating the vertices of a polygon drawn with DirectX or OpenGL [5].

Modern windowing systems such as Windows Aero and Linux Compiz pass all graphical user interfaces through the 3D card, in a process known as *compositing*. When compositing is used, any of the transformations shown in Fig. 3.5 can be added to the transformation stack and performed on all content that is displayed with virtually no extra computational cost.

Even if projective transformations to compensate for projector alignment are not required, using compositing still offers a significant advantage. In particular, because tabletops are horizontal and people can stand around the outside of them, the ability to reorient content to the viewer is necessary. Orientation has been found to be important in collaborative tabletop use, affecting how individuals comprehend information, how collaborators coordinate their actions, and how they mediate communication [30]. The ability to rotate windows has not been required in conventional window managers, so it has not been implemented, and must be added as an extra feature in tabletop interfaces. Compositing allows the use of similarity transformations (Fig. 3.5) to arbitrarily rotate and translate any item, without incurring any extra processing overhead. Other transformations and effects can easily be added, such as scaling, alpha blending, and 3D effects.

Further Methods

Content-Dependent Transformation

Transforming rasterized content reduces the effective resolution of the result, but in some cases it is possible to alleviate the loss by tailoring the transformation to the content being transformed. When the projective transformation for geometric compensation is close to a simple translation, Hereld and Stevens apply a transformation to text that translates individual words separately [31], and Tuddenham and Robinson [32] have used this on a tabletop display.

Nonlinear and Nonparametric Methods

The projective transformation is a simple linear transformation on homogeneous points, and is typically sufficient for compensating for projection onto a planar surface. Other options are piecewise linear transformations [33], in which the screen is divided into a triangular mesh and each triangle is warped independently, and general nonlinear transformations, which can be implemented in the graphics card using fragment shaders [10].

Nonplanar Screens

Tabletop displays are generally planar (flat) surfaces. However, other options are available, such as parametric surfaces. General three-dimensional objects can be augmented with computer graphics via *two-pass rendering* [27] whereby the known 3D shape is first rendered from the viewpoint of the projector, then the result is projected onto the object, but this beyond the scope of this chapter.

Edge Blending

When geometric compensation is used, projectors can be positioned casually and graphics are aligned in software following calibration. In general, the projected images will not be exactly rectangular due to keystoneing, and will be partially overlapped to avoid gaps between them. Fig. 3.6 shows a typical configuration. The issue here, is that the resulting output will be much brighter in the overlap regions, because more than one projector is contributing.



Fig. 3.6 Arrangement of six projectors. The images are not rectangular, so two or more projectors overlap at some points.

Binary Masking

The most obvious solution to this problem is to use only a single projector at any point, with the other(s) displaying black. This is the approach used by Escritoire [5] which displays a small high-resolution *fovea* inside a large low-resolution *periphery*, and simply displays a black rectangle in the region of the periphery that is overlapped by the fovea.

Blending

In T3, which supports tiled arrays of projectors as shown in Fig. 3.6, Tuddenham uses an edge blending technique for multi-projector displays described by Raskar et al. [34] and applies it to a multi-projector tabletop display [2]. For each projector i and two-dimensional pixel location \mathbf{x} , an alpha mask $\alpha(i, \mathbf{x})$ is multiplied by the projector's pixel value to obtain its edge-blended value. The values of the masks are constrained by $0 \leq \alpha(i, \mathbf{x}) \leq 1$ and $\sum_i \alpha(i, \mathbf{x}) = 1$, and are defined as

$$\alpha(i, \mathbf{x}) = \frac{dist(i, \mathbf{x})}{\sum_m dist(m, \mathbf{x})}, \quad (2)$$

where $dist(i, \mathbf{x})$ is defined as the Euclidean distance on the projection surface of point \mathbf{x} from the nearest edge of projector i .

Using this alpha blending approach, the projector arrangement in Fig. 3.6 results in the blending masks in Fig. 3.7. The result is shown in Fig. 3.8.




Fig. 3.7 Edge blending masks corresponding to the six-projector configuration in Fig. 3.6, from T3 [23].




Fig. 3.8 (a) Geometric compensation ensures the squares line up, (b) blacking out the overlap region leads to visible seams, (c) using alpha blending the seams are no longer visible.

Further Methods

Smoother Blending Function

The simple linear edge blending function in equation 3.2 can be replaced with smoother one, such as one that has a continuous first-order derivative [35].

Projector Brightness Variation

Having the blending functions sum to unity assumes that all projectors have the same brightness, that is, the same white level. This condition can be met, at least approximately, by manually adjusting the brightness of the projectors so they look the same. A more precise result can be obtained by allowing brightness to vary spatially, limited by a perceptual threshold on the gradient. Many projectors leak light when they are displaying black [36], so the black level will be increased in the overlap regions. This can also be addressed using a perceptual uniformity method, or simply by raising the lower limit of all parts of the display. This is a much less severe effect than the increased white level, so only a small amount of contrast need be sacrificed.

Photometric Compensation

Colour output can vary considerably between projectors, which can lead to visible variation over a multi-projector display. Colour matching can be attempted by adjusting the image settings commonly found on monitors and projectors: brightness, contrast, colour balance (between ‘warm’ and ‘cool’), and individual controls for the red, green, and blue channels. However, these are crude controls, and matching them is a laborious manual process. Photometric compensation methods aim to alleviate this problem. They require an understanding of some colour spaces that are briefly introduced in the subsection below.

Additive Colour

Projectors, like monitors, use additive colour. The default colour of the display surface is black, light is added by the projector to display a pixel, and the light combines additively. Humans see in three-dimensional colour, so displays usually use three primary colours. The standard primaries are red, green and blue (RGB, see Fig. 3.9). A pixel in an image is stored by a computer as an RGB value: a tuple of three numbers between 0 and 1 that are sent to the projector to be displayed.




Fig. 3.9 Additive colour. Three primaries are used: red, green, and blue. All three together produce white.

RGB values are device-dependent. The exact colour of the three primaries varies between projectors, because of variations between manufacturers and models, and also because the output of a projector bulb will vary as it ages [37]. This means that compensation for this variation is necessary.

Characterization

To compare the output of multiple projectors we should use a device-independent colour space, and the standard one is CIE XYZ-space [38]. In an additive colour space the final colour is a linear combination of the primaries, so to characterize the range of possible outputs of the projector we need only measure the minimum and maximum values of each of the primaries, that is, the colour produced by RGB inputs $(0,0,0)$, $(1,0,0)$, $(0,1,0)$ and $(0,0,1)$. This results in a three-dimensional *gamut* in XYZ-space (Fig. 3.10a), which is often displayed in two-dimensional xy-space (Fig. 3.10b), by ignoring luminance, and only considering chrominance (also known as chromaticity coordinates [38]).

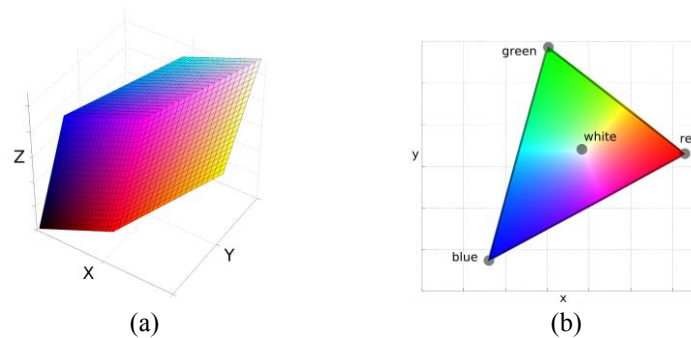


Fig. 3.10 Display device colour gamut in (a) XYZ-space and (b) xy-space.

The gamut is a bound on the colours that a device can produce, but it does not fully define the mapping from RGB input to colour output. The relationship

between each of the RGB channels and the intensity of its associated primary is called its transfer function. The original computer displays were cathode ray tubes (CRTs), which had gamma curves for transfer functions. That is, the intensity of a pixel I was related to the input value V by

$$I \propto V^\gamma. \quad (3)$$

Different values of γ produce differently shaped curves (Fig. 3.11a). LCD monitors and projectors now emulate the gamma curve, for the sake of backwards compatibility, so measuring the shape of this curve is also required during characterization. Some display devices have a mode designed to linearise their response and emulate $\gamma=1$, but this may not be sufficient, as shown in Fig. 3.11b, which shows the transfer function of a projector (Sanyo LP-XP45) measured with a colorimeter (Minolta CS-100). The curves for the three primaries are more linear than a typical gamma curve, but still noticeably curved at the top and bottom.




Fig. 3.11 (a) Gamma functions. The lines show $\gamma=1.0$ (linear), $\gamma=1.8$ (used by Mac OS), and $\gamma=2.2$ (used by Windows and the sRGB standard). (b) Response measured from a projector with gamma correction.

In addition to a per-projector gamut and transfer function for each RGB channel, Majumder and Stevens measure the maximum luminance achievable at each pixel on the display, so spatial variation can be removed [39]. The luminance is generally higher near the centre of a projected image than at the side, and much higher where projectors overlap.

The full photometric characterization can be measured using a calibrated camera. An uncalibrated camera can also be used, given some assumptions about the response of the camera.

Compensation

The response of each projector is linearised by inverting the transfer functions. For each RGB channel there are only 256 possible inputs, so the inverted function is stored in a lookup table.

Intra-projector variation in RGB primaries can be addressed using a colour mixing matrix: a 3×3 matrix multiplied by the original RGB values. This can be chosen to match all the projectors to a device-independent colour standard like sRGB, or just to match the projectors with each other.

The spatial variation of maximum achievable luminance can be addressed by applying a luminance attenuation map [39] to darken parts of the display that are too bright. This is implemented in the same way as the edge blending masks previously introduced (Fig. 3.7), and the two methods can be combined. To avoid wasting projector luminance, the limit of contrast sensitivity in human vision can be exploited to allow display brightness to vary spatially without the user's noticing [40].

Further Methods

Irregular Gamuts

Some DLP projectors use a colour wheel with a white component in addition to the red, green, blue components. This boosts the lumens rating of the projector at the expense of distorting the gamut from the parallelepiped in Fig. 3.10a to a more complex shape [36]. This could be addressed with more complex colour matching, or by sacrificing the extra brightness.

Intra-Projector Variation

The compensation method described above assumes that it is only a projector's luminance that varies spatially, which is a good approximation when using a dedicated projection surface [41]. If projectors are used to create a display on a patterned wall or table, compensation for the pattern can be implemented by using a spatially-varying gamut [42].

Content-Dependent Compensation

The luminance attenuation maps described in **the Compensation section above** hide the spatial variation in display brightness at the expense of some loss of contrast. The greatest possible contrast can be attained by allowing luminance clipping, bounded by a perceptual threshold, and performing content-dependent compensation [43]. However, a case must be made to justify this trade-off of computational complexity for display contrast.

System Integration

Tabletop displays have a different emphasis from typical multi-projector display walls in that they ought to be smaller, cheaper, more manageable systems. All of the geometric and photometric techniques described above can be implemented in a commodity 3D graphics card. PC motherboards are now available with four PCI-express slots, so with two digital outputs per graphics card, eight projectors can be driven from a single PC, which is enough even for a large multi-user tabletop display (Fig. 3.4).

Managing Multiple Graphics Cards

Current graphics cards now commonly have a gigabyte of memory, but even that large amount must be managed carefully if many high-resolution tiles (the tabletop display's analogue of windows) are to be stored as textures. Also, graphics card drivers may limit textures to sizes such as 1024×1024 or 2048×2048 pixels.

T3 splits a large tile into smaller sections, stores the sections in separate textures, and draws them as an array of polygons [2]. The textures are only stored on the graphics cards on which the corresponding sections are currently visible, thus reducing the memory requirements, and the time to update the textures when the tile content changes. This happens transparently to the user of the toolkit, freeing the application programmer from such considerations, and presenting a single large display device that supports tiles of any size.

Synchronisation of the updates of all projectors is also required. T3 addresses this simply by first preparing the double-buffered output for all projectors, then switching all of the buffers as quickly as possible. In practice synchronisation has not been a problem.

Rendering Performance

Modern graphics cards are designed to render thousands of texture mapped polygons per second, so they can easily handle projective transformations (page 12) with almost no latency, or load on the CPU. The performance of T3 was tested with 200 to 800 tiles, with approximately 10,000 pixels each, being continuously transformed and redrawn on a four-projector display. This resulted in between 32 and 22 frames per second (fps). The Buffer Framework [44] also addresses this issue, and achieves similar performance [45].

Graphics cards are optimised for drawing many textured polygons, but not as optimised for updating those textures. T3 achieves 20 fps while updating a 1 megapixel texture spanning four projectors without compression. As in any

graphical interface, it is good practice to design screen repainting code to only repaint what has changed.

Further Methods

Distributed Rendering

Chromium [46] intercepts OpenGL calls to present an existing application on a multi-projector display wall. However, scene complexity is generally not an issue for tabletop displays because the number of polygons is small, and it is desirable to perform all the rendering on a single PC, so distributed rendering should not be necessary.

Distributed Calibration

The iLamps of Raskar et al. [47] are a combination of a projector with a camera, wireless networking, a tilt sensor, and onboard computing, to form devices that calibrate with each other to create an ad-hoc multi-projector display. If such devices were mass produced, they could be used to create a large high-resolution display without the need for a high-end PC to drive it.

High-Resolution Input

Touch input is suitable for casual applications such as photo browsing, but basic input with the fingers is rather imprecise. If detailed work such as military command and control using maps is to be performed on a tabletop, more precise input is required. This can be provided in two ways: enhance touch input with new interaction techniques; or put a tool in the user's hand to provide more precision, with the most obvious tool being a pen. Another option for input is tangible interaction, that is, the interaction of multiple physical objects with the tabletop display. That is not discussed here because it is the focus of Chapter [chapter on Active Tangible Interactions]. Touch input technology will not be covered here because it is described in detail in Chapter [Interactive multi-touch surfaces]. Issues with touch input are mentioned briefly, followed by the technology and issues of pen input.

Touch Input Issues

Fingers are not very accurate for pointing on high-resolution displays. This is because of two issues. The first is occlusion. The finger is often wide compared to the items being manipulated, so information is occluded in what is known as the 'fat finger problem'. The second issue is precision. On a mouse pointer, the exact

point of interaction, at the tip of the arrow, is obvious, but the location of the ‘hot spot’ within the area of a finger is not so obvious. These two problems increase as display resolution increases. Various approaches have been developed, particularly for small mobile devices, and they can be adapted for tabletop displays. Chapter [Wigdor and Benko] describes the problems and potential remedies.

Pen Input Technology

People practise writing and drawing with pens from an early age. This type of natural free-form input has long been available on desktop and handheld computers. Below is a list of technologies that allow input with a pen, also known as a stylus, on a tabletop display.

Acoustic

Acoustic devices emit ultrasound from the stylus which is detected by microphones in a bar along the side of the display. Mimio is such a device. It is relatively inaccurate and is designed only to track a single stylus at a time [5], but it can easily be added to an existing display simply by attaching the bar with the microphones. Multiple bars and styli have been combined to form a very large pen-input surface [48].

Electromagnetic

Various sizes of electromagnetic tablets are available, from the graphics tablets by Wacom, to the large-format ‘digitisers’ by GTCO CalComp which have sizes up to 60×44 inches (74-inch diagonal). These devices provide accurate positioning and robust hardware, but they are large expensive devices and are only suitable for front-projection. They have been used in several tabletop displays [5, 2]. N-trig has recently released a display overlay that combines capacitive multi-touch input with an electrostatic pen, which has important implications for tablet computers, and possibly tabletops.

Optical pen

Anoto pens use a camera inside the stylus to detect its position on a surface printed with a special dot pattern. This provides high accuracy and precision. The coordinates are streamed over Bluetooth, and multiple styli can be used concurrently. Haller et al. [49, 50] demonstrated a front-projected tabletop

interface using these pens, and they were also used in T3 [23]. Recently these pens have been combined with rear projection and FTIR touch input [51]. This technology provides an economical way to make a surface of any size with any number of pens, although the hardware required in the pens makes them fairly bulky and requires them to be recharged regularly.

Lateral Cameras

SMARTBoard DViT overlays, designed for wall displays, use infra-red LEDs and cameras mounted at the table edge to provide sensing over a large area [52] and have been used in several tabletop systems [16, 17]. They can support one touch reliably, and two maximum. The overlay is added to an existing display, or incorporated into the bezel around the edge. It works equally well for pen or finger, but it is not as accurate as the graphics tablet or digital pens listed above. A limitation is that the user must not lean on the surface, or put any objects on it, because this obscures the cameras.

Rear camera

InfrActables [53] uses a camera behind a rear-projection screen to track multiple pens. The pens signal their id and state to the camera using a binary code from an LED in the tip. MightyTrace [54] replaces the single camera with an array of infra-red sensors, which allows the technique to be used in an LCD panel in a similar manner to how ThinSight [55] allows touch input on an LCD panel. There is more information about InfrActables and MightyTrace in Chapter [Kunz and Fjeld].

6DOF Tracking

Six-degree-of-freedom tracking provides continual updates for the pose of a tracked device: three distances specify a position, and three angles an orientation. Ashdown a Polhemus FastTrak to provide remote pointing for linked tabletop and wall displays [5]. Parker et al. [56] studied the efficiency of, and preference for, different pointing methods. Pointing from a distance allows users to interact with both close and distant items on a tabletop display without having to change modes or devices. To interact with a close object the user simply touches the stylus to the table. A disadvantage of the Polhemus technology is that it uses a magnetic field for tracking, so it is affected by metal objects and its accuracy degrades as the distance between the emitter and stylus increases.

Pen Input Issues

Although mouse, touch, and pen input are, at their most basic, simply ways to point to positions on the display, they allow different types of interaction, which affects the design of interfaces.

Using a pen instead of touch input largely solves the problem of precision, but some occlusion remains. The hand holding the pen can occlude information near the pen, depending on whether the user is left or right handed. A conventional scrollbar on the right may be problematic for a left-handed user, who needs to reach across to use it, thus occluding the information that is being scrolled. A conventional pop-up menu that appears to the bottom right of the cursor may be problematic for a right-handed user, because it appears underneath the user's hand. A pen-based interface should take account of occlusion by the hand, and could benefit from knowledge of which hand is being used, if this can be detected automatically. Brandl et al. [57] have implemented a circular menu that avoids occlusion by rotating automatically based on the positions of the hand and pen.

Buxton defines three input device states: *tracking*, *dragging* and *out of range* (OOR) [58]. A conventional mouse has the tracking state, plus one dragging state for each button. Touch input has OOR and tracking states (for each finger). A pen potentially has all three states. The Mimio, DViT, and Anoto mentioned above have the tracking state when they are touching the surface, and the OOR state. The GTCO CalComp digitizers have all three states: tracking is enabled when the pen is within about 2cm of the surface, and they have multiple buttons giving multiple dragging states. Other tablet technologies may also offer continuous pressure and tilt information in addition to position. The Polhemus 6DOF tracker cannot simply be taken out of range at will, so that state should not be required in any interaction technique.

These differences in device capabilities mean that an interface must be targeted at a specific input type, or adapted to work with a lowest common denominator. For instance, having different effects for 'hovering' over an item versus selecting it, as is common with mouse interaction, is not suitable for touch input or some pen input types, because there is no way to distinguish between hovering over an item and selecting it. A conventional graphical user interface ported to a system with touch or pen input may emulate the mouse, but what happens when the device goes out of bounds? Having the mouse cursor stay where it is may cause unwanted hover effects. Having it disappear may trigger unwanted mouse exit events. Similarly, ergonomic differences mean that performance will be different with different input types [59]: entering handwriting is easy with a pen, much more difficult with a finger or mouse.

We offer no ready-made solutions to these problems, but an awareness of the issues will allow them to be considered carefully at the design phase. Moving an application to a tabletop is not simply a matter of displaying an existing GUI on the display device.

Tabletops have typically provided either touch or pen input but not both simultaneously. Future systems should exploit the benefits of both. Guiard provides a model of bimanual action [60] that can inform the design of such input techniques. He gives handwriting as an example of bimanual action: the non-dominant hand successively repositions the paper, and the dominant hand then does all the writing within a relatively small portion of the table. More generally, the non-dominant hand acts first, it acts more coarsely in time and space, and it sets up a frame of reference in which the dominant hand works. Think of threading a needle, or hammering a nail. Brandl et al. [61] have embodied these ideas in a set of principles for two-handed input, and implemented some specific techniques that combine pen and touch input. We believe that, like multi-touch, pen-plus-touch calls for more research on interaction techniques, and convergence on interface conventions just as has occurred with graphical user interfaces controlled with a keyboard and mouse.

Future Trends

This section lists our predictions for the future of high-resolution interactive displays.

Smaller Cheaper Projectors

It seems that projector native resolutions will remain at 1080p (1920×1080 pixels) for now, because this is the standard for high-definition television. Projectors have steadily been getting cheaper and smaller for many years, and several manufacturers are now producing *pico projectors*, which are about the size of a mobile phone, and easily fit in a pocket. These devices use LEDs or lasers as the light source, and although native resolutions are not up to 1080p yet, they are increasing. Moreover, laser projectors do not use traditional objects and can be made with an infinite focal depth, thus removing focussing problems for short-throw and tilted surfaces. This continues the trend that is making multi-projector displays more convenient and economical.

Vector Graphics

Graphics displayed on tabletops are typically prepared in an off-screen raster image, and then subjected to one of the transformations shown in Fig. 3.5 to scale, reorient, and perform geometric compensation. This resampling causes aliasing,

and reduces the effective resolution of the display. Systems should transition to specifying graphics in a vector format like postscript, display postscript, Cairo,² or SVG. They will then be able to combine all necessary transformations into a single stack, and perform rasterisation in a single step.

Flat Panels

Projectors are currently useful for building tabletop displays in order to investigate the technical details, individual use, and collaborative use of tabletop displays. However, they make the systems unwieldy because of the throw distance required for projection, so before tabletop displays become mass-market products, the projectors should be replaced by flat panels.

Lambda Table [62] is an example of a tabletop display composed of an array of LCD panels. Fig. 3.12 shows that in this case the seams between the panels are very apparent, but the technology already exists to increase the visual area of such panels using a special lens so that they can be abutted with no visible seams.




Fig. 3.12 A variant of Lambda Table [62]. (a) Multiple panels form a tabletop display. (b) Objects are tracked to create a tangible user interface. Images courtesy of Jason Leigh.

In future, LCD panels will be replaced by organic light-emitting diode (OLED) displays, which have a better brightness, contrast, and colour gamut than LCDs, or by electrophoretic displays, which have much lower power consumption and can be viewed under bright light.

Some of the techniques described in the section above on multi-projector displays (pages 11–22) will still be applicable to tabletop displays composed of flat panels. Geometric transformations will be required to reorient content for multiple users, colour calibration may be required when multiple panels are used, and rendering will probably still be distributed over multiple graphics cards.

² <http://cairographics.org/>

Combined Sensors and Displays

Current systems typically combine separate input and output technologies to form an interactive tabletop display: either front projection with an off-the-shelf input device, or rear projection with FTIR. Future systems will combine input and output into a single self-contained device, as has been prototyped in ThinSight [55] and MightyTrace [54]. Ideally, each pixel of the display will provide both input and output, possibly by augmenting the conventional red, green, and blue components of each pixel with two more: an infra-red emitter and an infra-red sensor. The challenge will then be to create table-sized versions of these devices, or make ones with no bezels so they can be tiled to cover large areas.

Combined Touch and Pen

The ease of touch should be combined with the precision of pen input. This will allow interactions like the use of paper on a traditional desk, and will allow users to exploit the naturalness and efficiency of bimanual interaction.

Conclusion

Tabletop displays have typically provided low resolution output, which has limited the applications that have been possible. Resolutions between 60 and 110 ppi are desirable, but many of the systems listed in [Table 3.2 \(page 8\)](#) are below 60 ppi.

Most tabletop displays have been made using projectors, and the most economical native resolution for projectors has now improved from 1024×768 to the HDTV standard of 1920×1080, which provides slightly over 2 megapixels per projector. This means that one or two projectors are sufficient to create a high-resolution coffee-table-sized display, and four are sufficient for a large desk.

Various techniques that have been developed for multi-projector display walls are directly applicable to front or rear projected tabletop displays, making it possible to tile projectors seamlessly. The most useful and practical of these are: geometric compensation using projective transformations to align the images from multi projectors; edge blending using a simple alpha function to hide the overlaps between projectors; and photometric compensation for inter-projector variation. Over time, tiled projectors will be superseded by tiled flat panels that support both input and output, but some of multi-projector techniques described here will still be applicable. For example, applying transformations to reorient content for multiple users will be necessary.

The switch to high-resolution output also prompts the use of more accurate input, which could be achieved by augmenting touch with techniques to tackle occlusion and precision, or by adding pen input. We believe that touch and pen input should be combined for bimanual action because they are complementary.

References

- [1] Czerwinski M, Smith G, Regan T, Meyers B, Robertson G, Starkweather G (2003) Toward Characterizing the Productivity Benefits of Very Large Displays. In: Proc. IFIP Interact 2003, pp 9–16
- [2] Tuddenham P (2008) Tabletop interfaces for remote collaboration. PhD thesis, University of Cambridge, tech. rep. UCAM-CL-TR-734, December 2008
- [3] Wellner PD (1993) Interacting with Paper on the DigitalDesk. *Commun ACM* 36(7):87–97
- [4] Robinson P, Sheppard D, Watts R, Harding R, Lay S (1997) A framework for interacting with paper. In: Proc. Eurographics '97, vol 16, <http://www.cl.cam.ac.uk/research/origami/Origami1997c/>
- [5] Ashdown M (2004) Personal projected displays. PhD thesis, University of Cambridge, tech. rep. UCAM-CL-TR-585
- [6] Robinson P (1995) Virtual offices. <http://www.cl.cam.ac.uk/~pr10/publications/rsvr95.pdf>, proc. Royal Society discussion meeting on Virtual Reality, July 1995, British Telecom Publication number SRD/R5/1.
- [7] Wellner PD (1994) Interacting with Paper on the DigitalDesk. PhD thesis, University of Cambridge Computer Laboratory, tech. rep. UCAM-CL-TR-330
- [8] Bishop G, Welch G (2000) Working in the Office of “Real Soon Now”. *IEEE Comput Graph & Appl* 20(4):76–78
- [9] Li K, Chen H, Chen Y, Clark DW, Cook P, Damianakis S, Essl G, Finkelstein A, Funkhouser T, Housel T, Klein A, Liu Z, Praun E, Samanta R, Shedd B (2000) Building and Using a Scalable Display Wall System. *IEEE Comput Graph & Appl* 20(4):29–37
- [10] Majumder A, Brown MS (2007) Practical Multi-Projector Display Design. A. K. Peters
- [11] Bernard ML, Chaparro BS, Mills MM, Halcomb CG (2003) Comparing the effects of text size and format on the readability of computer-displayed Times New Roman and Arial text. *Int J Human-Computer Stud* 59(6):823–835
- [12] Tullis TS, Boynton JL, Hersh H (1995) Readability of Fonts in the Windows Environment. In: CHI'95 Conference Companion, pp 127–128
- [13] Kaiser PK, The Joy of Visual Perception: A Web Book. <http://www.yorku.ca/eye/acuity.htm>, Accessed 22 June 2009
- [14] Boff DR, Lincoln JE (1988) Engineering Data Compendium of Human Perception and Performance. <http://www.dtic.mil/dticasd/edc/EDCsec01/e01-0602.html>, Accessed 22 June 2009
- [15] Coldefy F, Louis-dit-Picard S (2007) DigiTable: an interactive multiuser table for collocated and remote collaboration enabling remote gesture visualization. In: Proc. PROCAMS 2007
- [16] Scott SD (2005) Territoriality in Collaborative Tabletop Workspaces. PhD thesis, University of Calgary
- [17] Isenberg P, Carpendale S (2007) Interactive Tree Comparison for Co-located Collaborative Information Visualization. *IEEE T Vis Comput Gr* 13(6):1232–1239, (Proc. Visualization / Information Visualization 2007)
- [18] Ashdown M, Robinson P (2005) Escritoire: A personal projected display. *IEEE Multimedia* 12(1):34–42, doi: 10.1109/MMUL.2005.18

- [19] Hu T, Chia Y, Chan L, Hung Y, Hsu J (2008) i-m-Top: An Interactive Multi-Resolution Tabletop System Accommodating to Multi-Resolution Human Vision. In: Proc. TABLETOP 2008, pp 177–180
- [20] Streitz NA, Geißler J, Holmer T, Konomi S, Müller-Tomfelde C, Reischl W, Rexroth P, Seitz P, Steinmetz R (1999) i-LAND: A Interactive Landscape for Creativity and Innovation. In: Proc. CHI 99, pp 120–127
- [21] Kakehi Y, Iida M, Naemura T, Shirai Y, Matsushita M, Ohguro T (2005) Lumisight Table: An Interactive View-Dependent Tabletop Display. *IEEE Comput Graph & Appl* 25(1):48–53
- [22] (2008) Microsoft Surface Datasheet. <http://download.microsoft.com/download/2/3/b/23b2282e-9562-40ee-910c-ad721b57217d/MicrosoftSurfaceDatasheet.pdf>, Accessed 22 June 2009
- [23] Tuddenham P, Robinson P (2007) T3: Rapid Prototyping of High-Resolution and Mixed-Presence Tabletop Applications. In: Proc. IEEE TABLETOP 2007
- [24] Ashdown M, Robinson P (2003) The Escritoire: A Personal Projected Display. In: 11th Int. Conf. in Central Europe on Computer Graphics, Visualization and Computer Vision (WSCG 2003), pp 33–40
- [25] Baudisch P, Good N, Stewart P (2001) Focus Plus Context Screens: Combining Display Technology with Visualization Techniques. In: Proc. UIST 2001, pp 31–40
- [26] Hsiao CH, Chan LW, Hu TT, Chen MC, Hsu J, Hung YP (2009) To move or not to move: A comparison between steerable and fixed regions of high-resolution projection in multi-resolution tabletop systems. In: Proc. ACM CHI 2009, pp 153–162
- [27] Bimler O, Raskar R (2005) Spatial Augmented Reality: Merging Real and Virtual Worlds. A. K. Peters, available as a free download from <http://www.uni-weimar.de/medien/ar/SpatialAR>
- [28] Sukthankar R, Stockton RG, Mullin MD (2001) Smarter Presentations: Exploiting Homography in Camera-Projector Systems. In: *IEEE Proc. ICCV 2001*, pp 247–253
- [29] Hartley R, Zisserman A (2003) *Multiple View Geometry in Computer Vision* 2nd edition. Cambridge University Press
- [30] Kruger R, Carpendale S, Scott SD, Greenberg S (2003) How people use orientation on tables: comprehension, coordination and communication. In: Proc. ACM GROUP '03, pp 369–378
- [31] Hereld M, Stevens R (2005) Pixel-aligned warping for multiprojector tiled displays. In: Proc. IEEE Int. Workshop on Projector-Camera Systems 2005
- [32] Tuddenham P, Robinson P (2007) Improved legibility of text for multiprojector tiled displays. In: Proc. PROCAMS 2007, doi: 10.1109/CVPR.2007.383464
- [33] Brown MS, Seales WB (2002) A practical and flexible tiled display system. In: PG '02: Proceedings of the 10th Pacific Conference on Computer Graphics and Applications, IEEE Computer Society, Washington, DC, USA, p 194
- [34] Raskar R, Brown M, Yang R, Chen WC, Welch G, Towles H, Seales B, Fuchs H (1999) Multi-Projector Displays Using Camera-Based Registration. In: Proc. IEEE Visualization '99, pp 161–168
- [35] Harville M, Culbertson B, Sobel I, Gelb D, Fitzhugh A, Tanguay D (2006) Practical Methods for Geometric and Photometric Correction of Tiled Projector Displays on Curved Surfaces. In: Proc. IEEE Int. Workshop on Projector-Camera Systems 2006
- [36] Stone MC (2001) Color and Brightness Appearance Issues in Tiled Displays. *IEEE Comput Graph & Appl* 21(5):58–66
- [37] Hereld M, Judson IR, Stevens RL (2000) Introduction to Building Projection-based Tiled Display Systems. *IEEE Comp Graph & App* 20(4):22–28
- [38] Wyszecki G, Stiles WS (1982) *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd ed. Wiley
- [39] Majumder A, Stevens R (2002) Lam: Luminance attenuation map for photometric uniformity in projection based displays. In: Proceedings of ACM Virtual Reality and Software Technology, pp 147–154

- [40] Majumder A, Stevens R (2005) Perceptual Photometric Seamlessness in Projection-Based Tiled Displays. *ACM Transactions on Graphics* 24(1):118–139
- [41] Majumder A, Stevens R (2004) Color nonuniformity in projection-based displays: analysis and solutions. *IEEE Trans Visualization and Computer Graphics* 10(2):177–188, doi: 10.1109/TVCG.2004.1260769
- [42] Grossberg MD, Peri H, Nayar SK, Belhumeur PN (2004) Making One Object Look Like Another: Controlling Appearance Using a Projector-Camera System. In: *Proceedings of CVPR 2004*, pp 452–459
- [43] Ashdown M, Okabe T, Sato I, Sato Y (2006) Robust content-dependent photometric projector compensation. In: *Proceedings of IEEE International Workshop on Projector Camera Systems (PROCAMS) 2006*, doi: 10.1109/CVPRW.2006.172
- [44] Isenberg T, Miede A, Carpendale S (2006) A buffer framework for supporting responsive interaction in information visualization interfaces. In: *Proc. Fourth Int. Conf. on Creating, Connecting and Collaborating through Computing (C5'06)*, IEEE Computer Society, pp 262–269
- [45] Miede A (2006) Realizing responsive interaction for tabletop interaction metaphors. Master's thesis, Otto-von-Guericke-Universität Magdeburg
- [46] Humphreys G, Houston M, Ng R, Frank R, Ahem S, Kirchner P, Klosowski J (2002) Chromium: A stream processing framework for interactive rendering on clusters. *ACM Transactions on Graphics* 21(3):693–702
- [47] Raskar R, van Baar J, Beardsley P, Willwacher T, Rao S, Forlines C (2003) iLamps: Geometrically Aware and Self-Configuring Projectors. In: *Proceedings of ACM SIGGRAPH 2003*, pp 809–818
- [48] Summet J, Somani R, Abowd G, Rehg J (2002) Interactive walls: Addressing the challenges of large-scale interactive surfaces. *Tech. Rep. git-gvu-02-35*, Computer Science Dept., Georgia Inst. of Technology
- [49] Haller M, Brandl P, Leithinger D, Leitner J, Seifried T, , Billinghurst M (2006) Shared design space: Sketching ideas using digital pens and a large augmented tabletop setup. *Advances in Artificial Reality and Tele-Existence* 4282/2006:185–196
- [50] Haller M, Brandl P, Leithinger D, Leitner J, Seifried T (2007) Large interactive surfaces based on digital pens. In: *Proc. 10th Int. Conf. on Humans and Computers (HC-2007)*, pp 172–177
- [51] Leitner J, Powell J, Brandl P, Seifried T, Haller M, Dorray B, To P (2009) Flux: a tilting multi-touch and pen based surface. In: *CHI'09 Extended Abstracts*, ACM, pp 3211–3216, doi: 10.1145/1520340.1520459
- [52] Morrison GD (2005) A Camera-Based Input Device for Large Interactive Displays. *IEEE Comput Graph & Appl* 25(4):52–57, doi: 10.1109/MCG.2005.72
- [53] Ganser C, Steinemann A, Kunz A (2006) Infractables: Multi-user tracking system for interactive surfaces. In: *Proc. Virtual Reality Conference*, pp 253–256, doi: 10.1109/VR.2006.86
- [54] Hofer R, Kaplan P, Kunz A (2008) Mighty trace: Multiuser technology on lcds. In: *Proc. CHI 2008*, pp 215–218
- [55] Hodges S, Izadi S, Butler A, Rrustemi A, Buxton B (2007) Thinsight: Versatile multi-touch sensing for thin form-factor displays. In: *Proc. UIST*, pp 259–268, doi: 10.1145/1294211.1294258
- [56] Parker JK, Mandryk RL, Inkpen KM (2006) Integrating Point and Touch for Interaction with Digital Tabletop Displays. *IEEE Comput Graph & Appl* 26(5):28–35
- [57] Brandl P, Leitner J, Seifried T, Haller M, Doray B, To P (2009) Occlusion-aware menu design for digital tabletops. In: *CHI 2009 Extended Abstracts*, ACM, pp 3223–3228, doi: 10.1145/1520340.1520461
- [58] Buxton W (1990) A three state model of graphical input. In: D. Diaper et al. (Eds), *Human-Computer Interaction - INTERACT '90*, Elsevier Science Publishers B.V. (North-Holland), pp 449–456

- [59] Forlines C, Wigdor D, Shen C, Balakrishnan R (2007) Direct-touch vs. mouse input for tabletop displays. In: Proc. CHI'07, ACM, pp 647–656, doi: 10.1145/1240624.1240726
- [60] Guiard Y (1987) Asymmetric Division of Labor in Human Skilled Bimanual Action: The Kinematic Chain as a Model. *J Motor Behaviour* 19(4):486–517
- [61] Brandl P, Forlines C, Wigdor D, Haller M, Shen C (2008) Combining and measuring the benefits of bimanual pen and direct-touch interaction on horizontal interfaces. In: Proc. AVI 08, pp 154–161
- [62] Krumbholz C, Leigh J, Johnson A, Renambot L, Kooima R (2005) Lambda table: High resolution tiled display table for interacting with large visu-alizations. In: Proc. Workshop on Advanced Collaborative Environments (ACE) 2005