ThinSight: Integrated Optical Multi-touch Sensing through Thin Form-factor Displays

Shahram Izadi, Steve Hodges, Alex Butler, Alban Rrustemi and Bill Buxton
Microsoft Research Cambridge
7 JJ Thomson Avenue
Cambridge, CB3 0FB
{ shahrami, shodges, dab, v-albanr, bibuxton }@microsoft.com

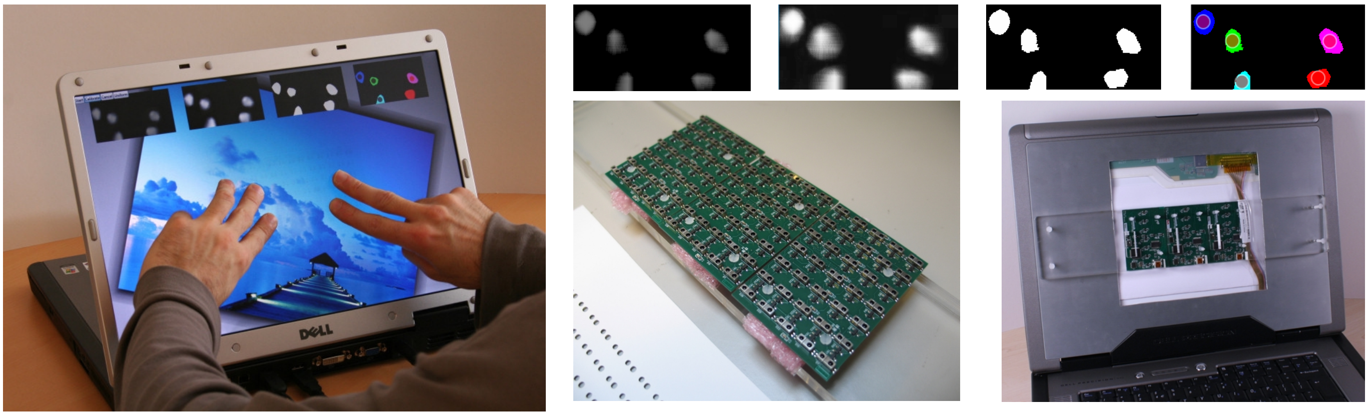

Figure 1: ThinSight enables multi-touch sensing using novel hardware embedded behind an LCD. Left, multi-touch input on a regular laptop LCD. Top row, the sensor data after interpolation, normalization, binarization, and component analysis. Bottom middle, front of three sensor boards tiled. Bottom right, sensors in position behind the display.

ABSTRACT
ThinSight is a novel optical sensing system, fully integrated into a thin form factor display, capable of detecting multiple objects such as fingertips placed on or near the display surface. We describe this new hardware, and demonstrate how it can be embedded behind a regular LCD, allowing sensing without compromising display quality. Our aim is to capture rich sensor data through the display, which can be processed using computer vision techniques to enable interaction via multi-touch and physical objects. A major advantage of ThinSight over existing camera and projector based optical systems is its compact, low profile form factor making such interaction techniques more practical and deployable in real-world settings.

Keywords: Novel hardware, infrared sensing, multi-touch, physical objects, thin form-factor displays.

INTRODUCTION
Multi-touch provides a more intuitive and direct way of interacting with rendered digital content. Widely disseminated work from Wilson (e.g. [20]) and Han [5], and products such as Apple's iPhone [1] and now Microsoft's Surface [11] have showcased this input technique, captivating online audiences and helping to encourage a fundamental rethink of how we interact with computers. Multi-touch has an even longer history however; the first systems appeared well over two decades ago (see [2] for an overview of the major landmarks).

Single touch displays, such as those based on resistive or capacitive overlays, cannot scale to supporting multi-touch robustly (see [17] for a full discussion), which has led to various alternative technologies. Some of these emerging systems [1,3,9,13,17] are based on purpose built sensing electronics, typically using an arrangement of capacitive electrodes on the display surface (although other sensors have also been used [4,7,10,15]). Other systems [8,12,18,19,20] employ cameras either in front of or behind the display to capture and process images of hands and other objects on the surface. These optical sensing approaches have proven to be particularly powerful in the richness of data they capture, and the flexibility they can provide in processing and detecting arbitrary objects including multiple fingertips. However, such camera and projector systems typically require a large optical path in front of or behind the display, which limits their portability and restricts the places where they can be deployed.

ThinSight is a novel optical sensing system, fully integrated into a thin form factor display, capable of detecting multiple objects such as fingertips placed on or near the display surface. The system is composed of an array of infrared (IR) emitters and detectors, which can be readily embedded behind a regular LCD, allowing IR based sensing without loss of display capabilities. See Figure 1.

In this paper we describe some of our experiences and experiments with ThinSight to date. Note that a fuller discussion of the underlying hardware is described in [6].

Our current prototype has fairly low-resolution sensing and covers only a portion of the display. However, even at this low-resolution, fingertips and hands are clearly identifiable through the ThinSight display, allowing multi-touch applications to be rapidly prototyped by applying simple image processing techniques to the data. Another compelling aspect of our optical approach is the ability to sense more than just hands and fingertips. Outlines of objects can already be observed through the display. With viable improvements in resolution, we envisage supporting detection of such physical objects near the surface for tangible input.

The major advantage of ThinSight over existing, projected and camera based optical systems is its compact, low profile form-factor making multi-touch and tangible interaction techniques more practical and deployable in a variety of real-world settings.

IMAGING THROUGH AN LCD USING IR LIGHT

A key element in the construction of ThinSight is a device known as a retro-reflective optosensor. This is a sensing element which contains two components: a light emitter and an optically isolated light detector. It is therefore capable of both emitting light and, at the same time, detecting the intensity of incident light. If a reflective object is placed in front of the optosensor, some of the emitted light will be reflected back and will therefore be detected.

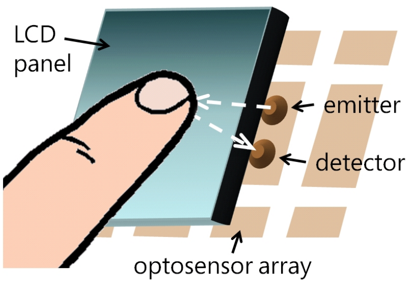

Figure 2: The basic construction of ThinSight. An array of retro-reflective optosensors is placed behind the LCD panel. Each of these contains two elements: an emitter which shines IR light through the panel; and a detector which picks up any light reflected by objects such as fingertips in front of the screen.

ThinSight is based around a 2D grid of retro-reflective optosensors which are placed behind an LCD panel. Each optosensor emits light that passes right through the entire panel. Any reflective object in front of the display (such as a fingertip) will reflect a fraction of the light back, and this can be detected. Figure 2 depicts this arrangement. By using a suitable grid of retro-reflective optosensors distributed uniformly behind the display it is therefore possible to detect any number of fingertips on the display surface. The raw data generated is essentially a low resolution monochrome "image" of what can be seen through the display in the infrared spectrum. By applying computer vision techniques to this image, it is possible to generate information about the number and position of multiple touch points.

A critical aspect of ThinSight is the use of retro-reflective sensors that operate in the infrared part of the spectrum, for a number of reasons:

- Although IR light is attenuated by the layers within commercially available LCD panels, some signal is detectable through the display, which is unaffected by the image displayed on screen.
- A human fingertip typically reflects around 20% of incident IR light and is therefore a quite passable 'reflective object'.
- IR light is not visible to the user, and therefore doesn't detract from the image being displayed on the panel.

ThinSight is not just limited to detecting fingertips in contact with the display; any suitably reflective object will cause IR light to reflect back and will therefore generate a 'silhouette'. Not only can this be used to determine the location of the object on the display, but potentially also its orientation and shape (within the limits of the sensing resolution). Furthermore, the underside of an object may be augmented with a visual marking (a barcode of sorts) to aid identification.

In addition to the detection of passive objects via their shape or some kind of barcode, it is also possible to embed a very small infrared transmitter into an object. In this way, the object can transmit a code representing its identity, its state, or some other information, and this data transmission can be picked up by the IR detectors built into ThinSight.

THE SENSING ELECTRONICS

The prototype ThinSight board depicted in Figure 3 uses Avago HSDL-9100 retro-reflective infrared sensors. These devices are designed to be used for proximity sensing – an IR LED emits infrared light and an IR photodiode generates a photocurrent which varies with the amount of incident light. Both emitter and detector have a centre wavelength of 940nm. With nothing but free space between the HSDL-9100 and an object with around 20% reflectivity, the object can be detected even when it is in excess of 100mm from the device.

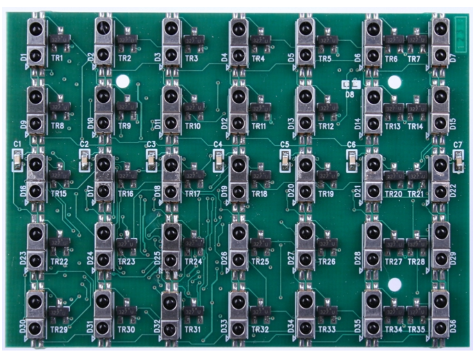

Figure 3: The front side of the sensor PCB showing the 7x5 array of IR optosensors. The transistors that enable each detector are visible to the right of each optosensor. Three such PCBs are used in our ThinSight prototype.

When mounted behind the LCD panel, the IR light is attenuated as it passes through the combination of layers from the emitter through to the front of the panel and then back again to the detector. As a result, with the current prototype a fingertip has to be at most around 10mm from the LCD surface in order to be detected.

The prototype is built using three identical custom-made 70x50mm 4-layer PCBs each of which has a 7x5 grid of HSDL-9100 devices on a regular 10mm pitch. This gives a total of 105 devices covering a 150x70mm area in the centre of the LCD panel. The pitch was chosen because it seemed likely that at least one sensor would detect a fingertip even if the fingertip was placed in between four adjacent sensors. In practice this seems to work well.

A PIC microcontroller collects data from one row of detectors at a time to construct a 'frame' of data which is then transmitted to a PC over USB via a virtual COM port.

Integration with an LCD panel

We constructed the ThinSight prototype presented here from a Dell Precision M90 laptop. This machine has a 17" diameter LCD with a display resolution of 1440 by 900 pixels. In order to mount the three sensing PCBs directly behind the LCD backlight, we cut a large rectangular hole in the lid of the laptop. The PCBs are attached to an acrylic plate which screws to the back of the laptop lid on either side of the rectangular cut-out as depicted in Figure 1 right.

To ensure that the cold-cathode does not cause any stray IR light to emanate from the backlight, we placed a narrow piece of IR-blocking film between it and the backlight. Finally, we cut small holes in the white reflector behind the LCD backlight to coincide with the location of every IR emitting and detecting element. Accurate alignment of these holes and the sensors is of course critical. The modifications involved removing the LCD panel from the lid of the laptop and then carefully disassembling the panel into its constituent parts before putting the modified version back together. During re-assembly of the laptop lid, the hinges were reversed, resulting in a 'Tablet PC' style of construction more appropriate for testing touch based interaction.

THINSIGHT IN OPERATION

The raw sensor data undergoes several simple processing and filtering steps in order to generate an image that can be used to detect objects near the surface. These include bicubic interpolation to scale up the raw sensor data by a factor of 10 and further smoothing using a Gaussian filter. This results in a 150x70 greyscale image. Once this image is generated, established image processing techniques can be applied in order to determine coordinates of fingers, recognise hand gestures and identify object shapes.

The images we obtain from the prototype are very promising, particularly given we are currently using only a 15x7 sensor array. With ThinSight, fingers and hands within proximity of the screen are clearly identifiable. Some examples of images captured through the display are shown in Figure 4.

Fingertips appear as small blobs in the image as they approach the surface, increasing in intensity as they get closer. This gives rise to the possibility of being able to sense both touch and hover. To date we have only implemented touch/no-touch differentiation, using thresholding. However, we have been able to reliably and consistently detect touch to within around 1mm for a variety of skin tones, so we believe that disambiguating hover from touch will be possible.

In addition to fingers and hands, optical sensing allows us to observe other IR reflective objects through the display. In practice many objects reflect IR. Figure 4 illustrates how the display can distinguish the shape of reflective objects in front of the surface. Currently we do not have the resolution to provide sophisticated imaging or object recognition, although large details within an object such as the dial widget are identifiable. Nonetheless this is an exciting property of our optical approach which we are investigating further.

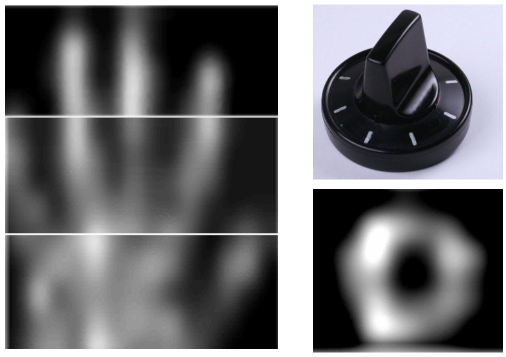

Figure 4: Left, an image of a user's hand captured through a ThinSight display (based on concatenation of three images captured from the display). A physical dial interface widget (top right) sensed on the surface (bottom right). Note these are the rescaled interpolated images.

A logical next step is to attempt to uniquely identify objects by placement of visual codes on their bottom surface. Such codes have been used effectively in tabletop systems to support tangible interaction [20]. We have also started preliminary experiments with the use of such codes on ThinSight, which look promising, see Figure 5.

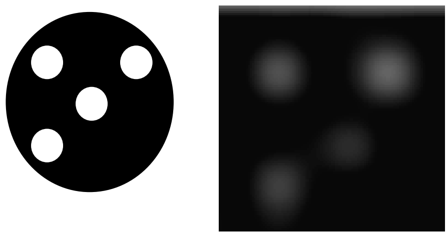

Figure 5: An example 2" diameter visual marker and the resulting ThinSight image as sensed through the display. Note this is the rescaled interpolated image.

Active electronic identification schemes are also feasible. For example, cheap and small dedicated electronics containing an IR emitter can be placed onto or embedded in-side an object that needs to be identified. These emitters will produce a signal directed to a small subset of the display sensors. By emitting modulated IR it is possible to transmit a unique identifier to the display.

Interacting with ThinSight

As shown in Figure 1 top, multiple fingertips are straightforward to sense through the ThinSight display. In order to locate multiple fingers we threshold the sensor image to produce a binary image. The connected components within this are isolated, and the centre of mass of each component is calculated to generate representative X, Y coordinates of fingers. A very simple homography can then be applied to map these fingertip positions (that are relative to the sensor image) to onscreen coordinates.

Using these established techniques, fingertips are sensed to within a few mm, currently at 8-10 frames per second. Both hover and touch can be detected, and disambiguated by defining appropriate thresholds. A user can therefore apply zero force to touch and interact with the display, with no apparent issues of occlusion or ambiguity. It is also possible to sense fingertip pressure by calculating the increase in the area and intensity of the fingertip 'blob' once touch has been detected.

CONCLUSIONS AND FUTURE WORK

We believe that the prototype presented in this paper is an effective proof-of-concept of a new approach to multi-touch sensing for thin displays. We have shown how this technique can be integrated with off-the-shelf LCD technology. The optical sensing allows potential for rich "camera-like" data to be captured by the display and processed using computer vision techniques. This allows new type of human computer interfaces that exploit multi-touch and tangible interaction on thin form-factor displays, making such interaction techniques more practical and deployable in real-world settings.

We have many ideas for potential refinements to the ThinSight hardware, firmware and PC software, including re-architecting each of these to improve the resolution of sensing, framerate, and power consumption. We have begun to expand the sensing area to cover the entire display, which has been relatively straightforward given the scalable nature of the hardware. In addition to such incremental improvements, we are also exploring new applications and interaction techniques that can truly benefit from multi-touch, multi-user and tangible input. These include tabletop configurations of ThinSight for the home to support media management and gaming using multi-touch interaction and physical objects; making such activities more intuitive, engaging, social and fun.

We also believe that ThinSight provides a glimpse of a future where emerging displays [14, 16], will cheaply incorporate optical sensing pixels alongside RGB pixels in a similar manner, resulting in thin form factor multi-touch sensitive displays becoming more ubiquitous.

REFERENCES

1. Apple iPhone Multi-touch. http://www.apple.com/iphone/technology/
2. Bill Buxton, 2007. Multi-Touch Systems that I Have Known and Loved. http://www.billbuxton.com/multitouchOverview.html
3. Paul Dietz and Darren Leigh. DiamondTouch: a multi-user touch technology. In ACM UIST 2001.
4. J. Y. Han, 2005. Low-Cost Multi-Touch Sensing through Frustrated Total Internal Reflection. In ACM UIST 2005.
5. J. Y. Han, http://www.perceptivepixel.com/
6. Hodges, S., Izadi, S., Butler, A., Rsustemi, A., Buxton, B. ThinSight: Versatile Multi-touch Sensing for Thin Form-factor Displays. ACM UIST 2007 (to appear)
7. JazzMutant Lemur. http://www.jazzmutant.com/lemur_overview.php
8. Myron W. Krueger, Thomas Gionfriddo, and Katrin Hinrichsen. Videoplace: an artificial reality. In ACM CHI 1985, 35--40.
9. S.K. Lee, W. Buxton, K.C. Smith, A Multi-Touch 3D Touch-Sensitive Tablet. In ACM CHI 1985.
10. Evert van Loenen et al. Entertaible: A Solution for Social Gaming Experiences. In Tangible Play workshop, IUI Conference, 2007.
11. Microsoft Surface, http://www.surface.com
12. Nobuyuki Matsushita and Jun Rekimoto, 1997. HoloWall: designing a finger, hand, body, and object sensitive wall. In ACM UIST 1997.
13. Rekimoto, J., SmartSkin: an infrastructure for freehand manipulation on interactive surfaces. In ACM CHI 2002, 113-120.
14. SensOLED technology. http://www.sensoled.ch/.
15. Tactex Controls Inc. Array Sensors. http://www.tactex.com/products_array.php
16. Toshiba Matsushita, LCD with Finger Shadow Sensing, http://www3.toshiba.co.jp/tm_dsp/press/2005/05-09-29.htm.
17. Wayne Westerman, 1999. Hand Tracking, Finger Identification and Chordic Manipulation on a Multi-Touch Surface. PhD thesis, University of Delaware.
18. Pierre Wellner, 1993. Interacting with Paper on the Digital Desk. CACM 36(7): 86-96.
19. Andrew D. Wilson, 2004. TouchLight: An Imaging Touch Screen and Display for Gesture-Based Interaction. In Conference on Multimodal Interfaces 2004.
20. Andrew D. Wilson, 2005. PlayAnywhere: A Compact Interactive Tabletop Projection-Vision System. In ACM UIST 2005.