

Going Beyond the Display: A Surface Technology with an Electronically Switchable Diffuser

Shahram Izadi¹, Steve Hodges¹, Stuart Taylor¹, Dan Rosenfeld²,
Nicolas Villar¹, Alex Butler¹ and Jonathan Westhues²

¹Microsoft Research Cambridge
7 JJ Thomson Avenue
Cambridge CB3 0FB, UK

²Microsoft Corporation
1 Microsoft Way
Redmond WA, USA

{shahrami, shodges, stuart, danr, nvillar, dab, jonawest}@microsoft.com




Figure 1: We present a new rear projection-vision surface technology that augments the typical interactions afforded by multi-touch and tangible tabletops with the ability to project and sense both through and beyond the display. In this example, an image is projected so it appears on the main surface (far left). A second image is projected through the display onto a sheet of projection film placed on the surface (middle left). This image is maintained on the film as it is lifted off the main surface (middle right). Finally, our technology allows both projections to appear simultaneously, one displayed on the surface and the other on the film above, with neither image contaminating the other (far right).

ABSTRACT

We introduce a new type of interactive surface technology based on a switchable projection screen which can be made diffuse or clear under electronic control. The screen can be continuously toggled between these two states so quickly that the switching is imperceptible to the human eye. It is then possible to rear-project what is perceived as a stable image onto the display surface, when the screen is in fact transparent for half the time. The clear periods may be used to project a second, different image through the display onto objects held above the surface. At the same time, a camera mounted behind the screen can see out into the environment. We explore some of the possibilities this type of screen technology affords, allowing surface computing interactions to extend ‘beyond the display’. We present a single self-contained system that combines these off-screen interactions with more typical multi-touch and tangible surface interactions. We describe the technical challenges in realizing our system, with the aim of allowing others to experiment with these new forms of interactive surfaces.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. Graphical user interfaces.

General terms: Design, Human Factors, Algorithms.

Keywords: Surface technologies, projection-vision, dual projection, switchable diffusers, optics, hardware

INTRODUCTION

Interactive surfaces allow us to manipulate digital content in new ways, beyond what is possible with the desktop computer. There are many compelling aspects to such systems – for example the interactions they afford have analogies to real-world interactions, where we manipulate objects directly with our fingers and hands. Some systems play on these analogies further by associating tangible objects with the digital interface, again leveraging our skills from the real world to interact in the digital.

Many different interactive surface technologies have been developed over the past few decades. One major category is the *rear projection-vision* system used in many research prototypes (e.g. [7, 16, 22, 31, 32]) and now even commercial products (e.g. [23, 24]). Such systems support capabilities beyond that of regular touch screens, permitting multiple fingers, and even other tangible objects near the surface to be sensed. Because the optical path is behind the display, occlusions are greatly mitigated compared with front projection and vision. The *diffuser* is a key part of these systems, displaying the projected image and ensuring that the camera can only detect objects close to the surface. However, the use of a diffuser also means that the displayed image, the sensing, and the hence user interaction will inherently be bound to the surface.

In this paper we present a new surface technology called SecondLight which carries all the benefits of rear projection-vision systems, but also allows us to extend the interaction space beyond the surface. Like existing systems we can display, sense, and therefore interact on the surface, but

we can also simultaneously project and image through the projection screen. This key difference stems from the use of a special type of projection screen material, described in detail in this paper, which can be rapidly switched between two states under electronic control. When it is *diffuse*, projection and imaging *on* the surface is enabled; when *clear* projection and imaging *through* is possible. This opens up the three-dimensional space above (or in front of) the surface for interaction. Because projection is no longer limited to the surface, it can be used to augment objects resting on or held above the primary display. Furthermore, both objects and user gestures can be sensed and tracked as they move around in 3D space. It is worth reiterating that these novel features of SecondLight coexist simultaneously with conventional surface-based interactions such as touch and tangible input.

SecondLight represents a new approach to support these types of extended surface computing interactions, bringing together ideas from diverse research areas, and integrating them into a single self-contained system. We feel that the use of switchable diffusers is particularly relevant for the interactive surface and tabletop communities providing systems with the ‘best of both worlds’ – the ability to leverage the benefits of a diffuser and rear projection-vision for on surface interactions, with the potential to instantaneously switch to projecting and seeing through the surface. In this paper we describe the construction and operation of our current SecondLight system in full to allow practitioners to experiment with this new form of interactive surface.

In the next section we begin by introducing the reader to some of the possibilities that this technology enables by way of a number of proof-of-concept examples. This structure for the paper is slightly unconventional, but we feel that the new functionality introduced by SecondLight should be presented up-front to motivate our work. The related work section that follows reviews some common concepts from the literature that we have adapted and integrated into our work. We also highlight previous work which uses similar switchable diffusers and we explain the benefits switchable technology has when compared to other diffuse materials. This leads on to a more detailed description of our technology which delves into the inner workings of the hardware, optics and sensing, uncovering some of the key challenges we faced. We end by discussing some of the current limitations of the system, and our plans for future work.

MOTIVATING SECONDLIGHT

A Trick of the Light

The key technical component in SecondLight is an electronically controllable liquid crystal material similar to that used in “privacy glass”, an architectural glazing material which can be switched between transparent and diffuse states as shown in Figure 2. When transparent the material is clear like glass, and most of the light passes straight through. When diffuse the material has a frosted appearance, and light passing through in either direction will scatter. In SecondLight this material is used as a rear-projection screen, resulting in a system which can display digital con-

tent *on* the surface whilst it is diffuse, or project *through* the surface when switched to its clear state.

Using the custom electronics described later, we can continuously switch the screen between these two states so quickly that it does not flicker – it looks as if it is continually diffuse. It is then possible to rear-project what is perceived as a stable image onto the display surface, when the screen is in fact transparent for half the time. During the clear periods a second image can be projected onto any suitably diffuse objects held on or above the display. As shown in Figure 1, by careful control in this way two different projections can be displayed on and through the surface, seemingly at the same time. Further, the images projected through the surface do not contaminate (or bleed onto) the ones on the main screen, and vice versa. This provides two independent projection channels.



Figure 2: A switchable screen, in clear state (left) and diffuse state (right). Note: the hand is very close to the underside of the surface.

A camera placed behind the switching screen can capture images when the surface is diffuse. The light scattering property of the screen in this state makes touch detection much easier. Additionally, when the screen is clear, the camera can see right through the surface into the space above (or in front of) the display. This allows the accurate up-close touch input to be augmented with richer data sensed at greater depths. A variety of diverse vision techniques can then be used, for example recognizing hand gestures from a distance, tracking diffuse objects with markers in order to project onto them, or detecting faces to ‘see’ the number and position of people around the surface.

The ability to “simultaneously” project both onto a surface and through it enables a number of interesting scenarios which we explore further in the remainder of this section.

‘On the Surface’ Interactions

With the projection screen in its diffuse state, SecondLight exhibits the established properties of multi-touch and tangible surfaces. Two examples are shown in Figure 3.

Aside from allowing an image to be projected on the surface, the diffuser plays a key role in detecting when fingers and other tangible objects are touching or very close to the display. This is an essential feature for direct input surfaces because the interaction relies on robustly detecting touch. Since the diffuser causes light to scatter, only objects very close to the surface will be clearly imaged and this simplifies touch detection, mitigating the need for computationally expensive methods such as stereo vision. We can detect a wide range of objects such as fingers, hands, brushes, game pieces, mobile devices and so forth, and also support

the unique identification of objects using retro-reflective markers such as those proposed in [16, 32].




Figure 3: Some of the typical on-surface interactions afforded by SecondLight include multi-touch input (left) and detection of physical objects such as a real paint brush & tangible interface widget (right).

Projection ‘Beyond the Surface’

SecondLight allows the combination of traditional surface interactions with more advanced features that extend interaction *beyond* the surface. This is made possible because we can project through the display surface, allowing physical objects that have suitable surfaces to be augmented with projections emanating from the SecondLight unit.

As shown earlier in Figure 1, one example is a thin sheet of diffuse film which is augmented by projection from below, as it rests on the surface and even when lifted up. This projected image is maintained on the film whilst projecting an entirely different image on the primary surface, without cross contamination of the two images. The dual projection capabilities of SecondLight can be used to create interesting layering and magic lens [2] effects. Instead of projecting two entirely unrelated images, the image projected through is visually connected to the one being projected on the surface. For example, as shown in Figure 4 top left, the image projected on the surface could be a car, with an associated image that reveals its inner workings being projected through (Figure 4 middle). In this scenario, if a user passes a piece of translucent material over the display, otherwise hidden information is revealed, creating a two-layer effect. Different translucent objects of varying forms and shapes can be used to exploit this capability, each effectively act-

ing as a physical magic lens. SecondLight inherently supports this notion of tangible UI layering in a simple and yet powerful way.

Two things are worth noting about this scenario. Firstly it does not require any tracking or computer vision algorithms or even any software to be written – it is all happening optically by virtue of the switchable diffuser and the two projected images. (Although as demonstrated later, this does not preclude the use of such algorithms for more sophisticated interactions.) Secondly, the object does not have to be resting on or even in contact with the surface. It can be lifted off the surface and the second projected image will still be maintained. This is different from the physical magic lenses previously used for tangible interaction [3] and is a unique capability when compared to other rear-projected surfaces. This allows us to interact using a magic lens which may be manipulated in six degrees of freedom. Clearly, some issues in terms of depth of focus and distortion exist with current projector technologies. The former we discuss towards the end of this paper, but the concept of distortion is one we explore in the next section.

Tracking mobile surfaces

If the position of the magic lens described in the previous section is tracked it becomes possible to support more sophisticated interactions. For example, a zoom effect can be applied as the lens is moved towards or away from the surface – making its behavior more analogous to a real magnifying glass – or new layers of information could be revealed as it rotates. SecondLight supports such tracking if the magic lens is augmented with either passive (retro-reflective) or active (powered infrared LED) tags. Indeed, by tracking the position and orientation of a mobile display surface such as a magic lens, it is possible to alter the projected image in real time so that it appears centered and without any foreshortening or other distortion, even as this mobile surface is manipulated in three dimensional space.

This allows us to create extremely cheap and lightweight peripheral surfaces that can be used in conjunction with the



Figure 4: Creating tangible layering effects is extremely simple with SecondLight. A translucent sheet of diffuse film is placed above an image of a car to reveal its inner workings (middle) or above the night sky to reveal the constellations (right). Plain images are shown on left. Circular acrylic discs with diffuse topsides can be used to create a magic lens effect. Note how the magic lens image is maintained even if the disc is lifted well away from the surface.

SecondLight display. Users can shift their attention between the primary, shared display and one or more smaller, mobile displays, viewing and interacting with content on both as they please. The mobile surfaces can be tilted towards the user or even held in the hand and the rendered content will track them accordingly, as illustrated in Figure 5. Further, as also shown in this figure and described in detail later, a novel active tagging scheme is applied to the surface which not only supports tracking of the mobile surface, but allows this object to support multi-touch input.



Figure 5: Mobile projection surfaces with passive and active markers being held at different orientations above the surface. The markers define the location of the object, which is tracked using the camera imaging through the surface. This allows correction of the through projection enabling it to be distortion-free and appear centered on the mobile surface. Top: a flexible sheet with passive retro-reflective marker strips. Middle: an actively tracked surface with its own battery-powered IR light source which also allows multi-touch input on the mobile surface to be sensed through the SecondLight surface (middle right). Bottom: the projected image can be corrected for distortion as it is moved and tilted, thereby supporting quick and natural reorientation into more conformable viewing positions. For example, tilted vertically towards the user or oriented towards another user for viewing.

Other Interesting Tangible Possibilities

Tangible objects can enrich the way we interact with tabletops and other surfaces [11]. Systems such as Reactable [16] and Microsoft Surface [23] track the position of tangible objects placed on the surface in order to project visual content immediately below or around the object. They also support user input through direct manipulation of these tangible objects. We can support this type of interaction with SecondLight but the additional ability to project light through the surface allows us to explore new designs for tangible objects. In the simplest scheme, a transparent ob-

ject with a diffuse top surface allows a projected image to be displayed on top of the object. In a more complex embodiment, shown in Figure 6, circular prisms built into the object allow the projected image to be totally internally reflected onto the sides of the object. In this example, we see a rolling tickertape that provides feedback to the user using the vertical sides of the object. Again this is only a proof-of-concept, but it demonstrates how a cheap (e.g. injection molded) tangible object can be illuminated and render graphics on its surfaces. Imagine game pieces, such as chess pieces, designed in this way allowing animated graphics to be rendered onto their faces and bodies.



Figure 6: The ability to shine light through the display, gives rise to other novel tangible object designs. In this example, we demonstrate an object that uses internal prisms to project the incoming light onto its sides. The prism inside the object is shown right. The effect of the prism on the projection is shown left. This middle image shows the screen in a non-switching diffuse state – illustrating the behavior of a typical projection screen.

Gesturing and Input from a Distance

With the projection screen clear, a completely unattenuated image can be captured by the camera. With sufficient illumination it is possible to track the users' hands from a distance and identify hand gestures and poses using computer vision techniques. One simple illustrative example is shown in Figure 7.




Figure 7: SecondLight allows gesture-based interactions with the primary surface from greater distances than many back projected systems (left).

A stereo camera, 3D camera technology, or structured light could be used to support depth based interactions (although we have yet to explore these possibilities). Alternatively, as we have shown, tags and markers can also provide approximations for depth. Although interacting from a distance breaks the direct manipulation metaphor, it does open up yet another input modality. Feedback can also be provided during these depth-based interactions by projecting onto the underside of interacting objects. It may even be possible to provide coarse feedback to the user when gesturing from a distance by illuminating the underside of hands, using for example changes in color to indicate if a gesture has been recognized by the system (see Figure 7).

This concludes the broad overview of some of the possibilities that SecondLight enables. It is important to stress that SecondLight has not been developed to address a specific problem or application, but rather is an exploration of a new technology which we believe has the potential to deliver a range of interesting and compelling user experiences. Realizing the system and even the simple proof-of-concept demonstrators presented above has been challenging. We have needed to address issues in optics, mechanical design, electronics, computer vision and graphics, before considering the UI or application layer. Our aim in the rest of this paper is to share our insights, so that others can develop and explore the possibilities afforded by this type of surface computing experience. First however, we review some of the relevant research that has inspired our work.

RELATED WORK

SecondLight combines concepts from diverse research areas into a single self-contained system. This section describes key related work starting with interactive surfaces.

Thin Form-factor Interactive Surfaces

A key technical challenge for interactive surfaces has been the low-level sensing techniques employed to detect the movements of fingertips and objects on the display surface. A variety of resistive, capacitive and inductive schemes have been employed in the past. Both resistive and capacitive touch screens have been scaled to support multi-touch [1, 4, 12]. However, one major issue is that these systems cannot recognize a wide range of untagged objects in addition to fingertips. Such capabilities are an essential part of interactive surfaces, particularly ones that need to support tangible input.

A variety of optical sensing techniques that support touch and objects in various ways have also been developed. The Philips Entertaible [20] detects multiple touch points plus certain object shapes using IR emitters and detectors mounted in a bezel around the edge of the screen. Thin-Sight [8] uses an array of IR proximity sensors embedded behind the backlight of an LCD panel to detect when fingertips, hands and other object outlines are in close proximity to the display. Although optical sensing using discrete IR components in this way can be flexible, in particular allowing the optics to be collapsed down into a thinner form factor, currently these systems have drawbacks in terms of scalability and resolution.

Projection-vision Surfaces

Arguably, if form factor is not an issue, a more powerful approach for optical sensing is to use cameras instead of discrete sensing components. Camera-based systems provide a high resolution means for capturing richer information about arbitrary objects near to the display. One main distinction between different camera-based systems is whether the optics is mounted at the *rear* or in *front* of the projection screen (for tabletops this would imply *below* or *above* the surface respectively). In most cases, the optics we refer to comprise not only the camera for sensing but a projector for display. Researchers have used the term *projection-vision* to describe such systems. In these systems a

diffuse or opaque screen is used to display the projected image.

Front-based approaches tend to capture richer data regarding the interactions occurring *on* and *above* the projection screen as the image captured by the camera is not scattered by a diffuser. Any object in the field of view of the camera can be imaged and in theory sensed. However in practice, people interacting with the surface will inevitably lead to parts of the camera or projector image becoming occluded. The latter does however allow objects to be directly augmented by this projection making front-based approaches popular for certain tangible interfaces [11, 30]. In such systems it is also far more difficult to accurately detect when an object is close to or touching the surface as opposed to further away, and solutions such as stereo vision [21] tend to be computationally expensive. One notable exception is PlayAnywhere [32], which uses a short-throw projector and a single off-axis camera. Here touch is determined by detecting both fingertips and their shadows.

Rear projection-vision systems (e.g. [16, 22, 23]) address both this occlusion problem and touch discrimination. Illumination is required in these systems to detect objects in front of the diffuser. Typically IR light is used as it is invisible to the human eye and will not adversely affect the displayed image. A diffuse light source shines IR light out through the display surface, and a rear mounted camera detects any reflections from fingers and other IR reflective objects. The diffuser acts to hide reflections from any objects beyond a few millimeters from the surface, improving the accuracy of touch detection by reducing chances of false positives caused by accidentally sensing objects far from the surface.

Han [6] describes a different technique for multi-touch sensing using frustrated total internal reflection (FTIR). Here a clear sheet of acrylic is placed on top of the rear projection screen and edge-lit with IR light such that the light is repeatedly totally internally reflected. When a finger is pressed up against this sheet, it causes IR light to be emitted away from the finger through the projection screen whereupon it can be detected with a rear mounted IR camera. This gives a more ‘binary’ indication of touch – until the fingertip is actually in contact with the surface there will be no signal. However, this approach has two main downsides. First, it cannot sense a wide variety of objects beyond fingers. This issue can be resolved in part by adding a layer of silicone as described by [28], but only objects exerting pressure above a certain threshold can be sensed in this way. Second, because the sensor is ‘binary’ no hover or proximity information can be detected.

Through-surface Imaging and Projection

The previous section described systems that support “on the surface” interactions. In this section we describe surface technologies that can go beyond the display surface for input or even output. TouchLight [31] employs a holographic rear projection screen that can be projected onto and at the same time allows imaging at far greater depths than typical diffusers. This is demonstrated by using a digi-

tal stills camera to support high resolution color imaging through the display. The lack of a diffusers however means stereo vision (and hence two cameras) are required to support touch detection.

Perhaps the most similar rear projection-vision systems to our work are those built around Lumisty film [15]. This material is also used for privacy glass, but its appearance changes from transparent to translucent and vice versa depending on the angle it is viewed from. This property of the screen has been combined with a Fresnel lens by Kakehi et al. to support a number of novel uses, e.g. creating an interactive tabletop which can show different images depending on which direction the user approaches [13], and more recently to project onto vertical surfaces resting on the display [14].

There are however practical advantages to using a switchable diffuser instead of Lumisty. First, because we are switching the screen from diffuse to clear we have finer grained control of the optical configuration. Therefore we can ensure that both on and through projections are completely independent of one another, with no cross contamination of these signals occurring. Second, although the view dependent nature of Lumisty provides some interesting multi-view scenarios, it also means that people will not observe a single consistent image on the tabletop surface as they move around it – each side would require a separate projector and the images will need to be fused and synchronized. We feel that both switchable and directional diffusers provide distinct features (and limitations) and ultimately their utility depends on the particular context of use.

Projection onto Tracked Surfaces

One of the key features of privacy film diffusers is their ability to support projection onto mobile peripheral surfaces through the display. Lee et al. [17, 18] presents an evolution of techniques for tracking and projector calibration that really demonstrates the power and wonderment that comes from using projections onto arbitrary passive surfaces. Lee’s novelty comes in combining the sensing and projection spaces thus greatly improving the ease of calibration.

Other compelling systems include PaperWindows [9], Urp [30], and work by Raskar et al. [26] which further demonstrate the power of supporting real-time projection onto passive mobile surfaces such as sheets of paper. In all these systems, as with SecondLight, distortion-free projections are rendered onto these mobile surfaces in real-time by tracking the bounds of the mobile surface using a camera or other light sensor. However none of these support projection from behind the display surface. For example, in PaperWindows top-down projection and a Vicon tracker embedded in the environment are used. SecondLight provides a single self-contained unit that also minimises occlusion because it uses rear projection and vision.

Interactive Systems with Switchable Diffusers

Finally, there has also been work on the use of switchable diffusers for interactive applications. Early pioneering work by Shiwa et al. [27] presents the use of a switchable dif-

fuser with a camera embedded behind to support gaze corrected video conferencing. Other compelling interactive systems built using switchable diffusers include blue-c [6] which uses the screen for capturing images of participants inside an immersive CAVE. All these systems capture images through the surface rather than considering projection through the display onto other diffuse objects. They all use a switchable diffuser technology with slow switching speeds, which results in low transmittance when driven at the higher rates needed to avoid flicker perception. This is sufficient for image acquisition, but will result in issues when projecting through the surface, most significantly inadequate brightness of the through projection and cross contamination between images projected on and through the surface. One other notable system is the DepthCube 3D Volumetric Display [29], a solid state, rear projection, volumetric display that consists of a high-speed video projector, and a multi-layer stack of switchable diffusers. The high-speed video projector projects a sequence of slices of the 3D image into the multi-layer optical element where each slice is rendered at the corresponding depth.

SecondLight explores the use of switchable diffusers in the context of interactive surfaces. It brings together many of the concepts in the broad range of research covered in this section, and integrates them into a single self-contained system which benefits from the use of rear projection-vision but also carries some of the affordances of front projection (or top-down) systems. These broad topics have made developing SecondLight into a considerable endeavour, involving a number of specific technologies which have not previously been used for surface computing. In the remainder of the paper we describe the hardware, optics, and sensing aspects of our system in order for others to explore such possibilities.

THE HARDWARE IN MORE DETAIL

The switchable diffuser

We use a polymer stabilized cholesteric textured liquid crystal (PSCT-LC) optical switch from LC-Tec [19] as our switchable projection screen. PSCT-LC is similar to polymer dispersed liquid crystal (PD-LC), a material that is commonly used as privacy glass in offices and store fronts (these are windows that can change their optical properties from frosted to clear at the flick of a switch). Both PD-LC and PSCT-LC are made from a special material containing liquid crystal molecules which are normally randomly-oriented and which therefore scatter light in all directions. However, they become untwisted and therefore aligned in response to a suitable electric field which may be generated by applying a voltage across two parallel, transparent substrates on either side of the screen.

PSCT-LC combines the liquid crystal molecules with lower concentration of polymer that splits the liquid crystal material into separate domains that can individually respond quicker than the larger regions present in PD-LC, thus reducing switching time and improving viewing angle. This is critical in supporting switching above the threshold of flicker perception whilst at the same time supporting the acceptable levels of transmittance to support projection on

and through. The result is a material that can be switched between clear and diffuse states in less than 0.5ms. With a suitable voltage applied, the PSCT-LC screen we use has 82% transmittance; this compares to less than 3% with no excitation.

Driving the diffuser

We continually switch our PSCT-LC screen between diffuse and clear states at 60Hz, which we found was a sufficient frequency to avoid flicker perception when looking directly at the surface. Each cycle consists of around 8.3ms when 150V is applied to the screen to make it clear followed by 8.3ms with no applied voltage, at which point it returns to its natural diffuse state. The exact proportion of time in each state (i.e. the duty cycle) can be varied according to specific needs of the system design. We opted for a 50% duty cycle because we are interested in both on and off surface interactions. Increasing the diffuse interval at the expense of the clear interval, for example, will increase display brightness on the surface at the cost of reducing brightness of the through projection. It will also decrease the available light to the camera for imaging through the surface.




Figure 8: The waveform that drives the PSCT-LC screen. A 150V potential is applied for 8.3ms at a rate of 60Hz; to prevent any DC offset the polarity of this is reversed on alternate cycles.

In order to prevent premature aging of the PSCT-LC through electrolysis, the manufacturer specifies that it is driven with a DC-balanced signal. For this reason, we reverse the polarity of the 150V on alternate cycles. The display is in diffuse state during the zero-voltage period between polarity-reversals, and is transparent when 150V is applied (with either polarity). An example waveform is shown in Figure 8.

Our driver circuit is based on an H-bridge architecture. Each side of the panel is connected to one half-bridge, capable of switching between 0 and +150 V. A potential of 0V, +150V or -150V may therefore be applied across the PSCT-LC depending on whether neither, the left or the right half-bridges are enabled (respectively). Each half-bridge is implemented as a complementary emitter follower, made from NPN and PNP power audio transistors. These transistors are capable of delivering the high transient current (~4A) required to rapidly switch the PSCT-LC panel, which is effectively a nonlinear capacitor of around 6 μ F. Electrical isolation between the high-voltage circuitry and the rest of the system is achieved through the use of an opto-coupled level shifter. Figure 9 shows the driver circuit.




Figure 9: The PSCT-LC H-bridge driver circuit.

Projector setup

In order to project different content on the surface versus through it, we need to alternate between two different images in sync with the switching diffuser. The ‘on’ image is displayed for 8.3ms followed by the ‘through’ image for 8.3ms and so on. The frame rate for each of these two images is 60Hz, so when interleaved in this way we have an effective frame rate of 120Hz. Whilst 120Hz projectors are available, for the initial SecondLight prototype we chose what we believe is a more flexible solution – we use two off-the-shelf Hitachi CPX1 60Hz projectors in combination with fast optical shutters to create the two interleaved 60Hz images. Like the switchable diffuser, the shutters are liquid-crystal based, but in this case they switch to a black, light-blocking state when they are not clear.

Blocking the light from the first projector whilst the projection surface is clear causes the image from the second projector to pass through the PSCT-LC; on the next part of the cycle we reverse the shutters so that the ‘through’ projector is blocked and the light from the first projector is displayed on the surface.

The ferroelectric liquid crystal (FLC) shutters used are LV-4500P-OEM units from DisplayTech [5] and are driven with ± 5 V. Again, these are driven differentially to remove the need for a negative supply rail – the output stages of a microcontroller effectively form the half bridges. We found that the performance of the FLCs tends to slowly deteriorate during the operation of the system – after a couple of hours of continuous operation they are not so effective at blocking light. We believe that this is due to a combination of heat and possibly exposure to ultraviolet light leaking from the projector, and as a result we fit an IR reflective hot mirror and a UV blocking filter between each FLC and its respective projector, and we also cool them with forced air from a fan.

Camera configuration

In addition to projecting both onto and through the surface, we also image what is on the surface and beyond using two ImagingSource DMK 21BF04 [10] cameras mounted behind the diffuser. Whilst SecondLight allows the capture of full colour images, to date we have fitted IR pass filters to limit imaging and sensing to the infrared spectrum.

We use both diffuse and FTIR IR light sources in conjunction with the first camera to sense multiple fingers and other objects. The FTIR light source consists of 264 Osram SFH 4255 high power IR LEDs which are distributed on a 6mm pitch along the edges of a 490 x 390mm sheet of

4mm thick clear acrylic. The LEDs are wide angle ($\pm 60^\circ$), 850nm devices which are surface mounted to a custom-made PCB at right angles. They are driven continuously at around 80mA in chains of 6 devices from a 12V PSU. We use the same LEDs as the source of diffuse illumination.

The second camera is triggered so that it grabs images when the PSCT-LC is clear, and can therefore also see far beyond the surface. Any IR sources in the field of view, such as IR LED markers, will be clearly visible.

Putting it all together – a small matter of timing

All the components of SecondLight are held together using a lightweight frame made from a modular extruded aluminium system from Bosch Rexroth. Note that we have adopted a tabletop configuration, but a vertical setup is equally feasible. The PSCT-LC, clear acrylic overlay and IR LEDs are held in place using a black acrylic bezel which is secured to the top of the frame. The various power supplies for the system rest on a shelf at the bottom of the frame.




Figure 10: The layout of the main SecondLight components; on the left is a cross-section through the side of the unit and on the right is a photo from the front.

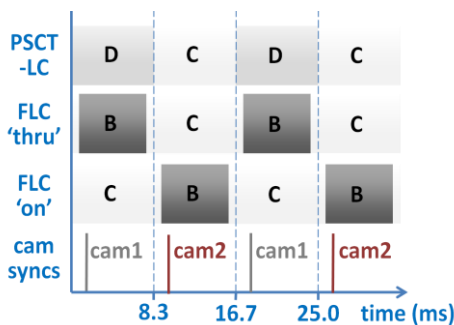


Figure 11: The timing used to drive the PSCT-LC diffuser, the FLC shutters and the camera 'sync' trigger inputs. Key: D=diffuse, C=clear, B=blocking.

A cross section through the frame, which depicts the location of the cameras and projectors relative to the display surface, is shown in Figure 10. Note that the projectors, which are designed to sit on a table or hang from a ceiling, do not emit light symmetrically and as a result are mounted off-axis which leaves space for the cameras in between. This has the added bonus of preventing light from the

'through' projector from shining into users' eyes – we configure the system so that this projector is the one at the front of the unit as users approach, and the light projected from it will not be seen unless they explicitly lean over the display and look directly into the projector. A photo of the assembled frame is also shown in Figure 10.

A custom PCB with an Atmel AT-Mega8 8 bit microcontroller operating at 4 MHz acts as the master signal generator and timing controller for the SecondLight system. The board provides a control signal for each half of the H-bridge to determine PSCT-LC state, camera sync signals to trigger image capture, and drive signals to open and close the FLC shutters. Figure 11 shows the timing control for the various elements of the system.

SENSING

The images captured during FTIR illumination can be used to accurately detect if one or more fingertips are touching the surface. An example image is shown in Figure 12 left. We have been pleasantly surprised by the quality of the FTIR sensing in our setup. Fingertips appear as extremely bright pixels in the image and are easy to extract using computer vision. We only need to correct the lens distortion, threshold the image, run connected component analysis and determine the centre of mass for each contact to determine touch points – not even background subtraction is required if the diffuse illumination is modulated and no other IR sources are nearby. This allows us to achieve close to full camera frame rates (currently 60Hz), which seems to give an acceptable user experience. We currently use component tracking in order to continuously identify individual points. Major and minor axes are also derived in order to determine orientation of each tip. The size and intensity of touch contact can also be used to estimate (albeit coarsely) the pressure applied once touch is detected.



Figure 12: two hands imaged using FTIR (left) and diffuse illumination with the PSCT in diffuse (middle) and clear states (right)

The image captured with diffuse illumination during clear states can be used for detection of a wider range of IR reflective objects, such as brushes, game pieces, mobile devices or even visual markers [32] – see Figure 13. We also found that using diffuse illumination with the screen also in a diffuse state reduces the background signal significantly. The light scattering causes objects to blur as they move away from the camera as shown in Figure 12, until they disappear from the scene. In practice, there is a sharp change in diffusion as an object moves millimetres off the surface. This effect can be used to our advantage to determine when objects are actually touching the surface.

In typical back projection-vision systems each touch point only appears as a small ellipsoid-like shape in the captured

image. For example, if we look at the FTIR image shown in Figure 12, here it is unclear which component belongs to which user. Therefore during the processing, a probabilistic approach must be used to determine which touch points belong to a specific hand. With SecondLight, the image captured when the screen is diffuse can be combined with the ‘through’ image to enrich the data. For example, the users’ hands and possibly arms are likely to be clearly visible when they are touching the surface – see the example image in Figure 12. By mapping the touch points sensed in the diffused FTIR touch image to each user’s hand we get a sense of the number of people who are touching the surface. Further, we have noticed occasions that the camera looking through the surface in clear states can image a user’s face, as highlighted in Figure 13. This data could be combined with face detection algorithms to determine the number of users around the display and potentially if they are looking at the surface.

Diffuse illumination can also be used to track objects off the surface. Here the vision problem does become harder; in particular using centre of mass of objects is at odds with the particularly rich shape information we obtain. We have been exploring the use of contours extracted from the image by computing the spatial gradient using a Sobel filter and optical flow techniques [32] to more suitably model our depth based interactions.



Figure 13 Left: an example image captured in the PSCT clear state. Here we obtain a far greater depth of field than other rear projected surfaces. A mobile phone, magic lens and paint brush are touching the surface, but further away we can image the user’s entire arm and even face through the surface. Right: a simple mobile surface that uses FTIR to generate a signal for tracking to enable perspective correction. The LED strip which is mounted on the side is shown resting on the center of the device. This device also generates a multi-touch signal, which, along with the edge markers, can be imaged by the SecondLight IR camera. The output captured by the camera is shown inset, here three fingers are touching the surface and the two strips highlight the edges of the display surface.

One final aspect of our system is the need to track mobile surfaces from a distance. This is required in order to correct the projection so that it is perfectly aligned and undistorted when it appears on the mobile surface, even as that surface moves and tilts. To achieve this we must define a projective transform [25] which pre-distorts the image so that it ends up rendered with the correct perspective on the mobile sur-

face plane. To solve such a projective transform we essentially need to detect and track four points on the mobile surface.

One approach is to use diffuse IR illumination and passive markers which reflect the IR and can be detected when the main surface is clear. Although this works it is noisy. We have therefore opted to track mobile surfaces without diffuse illumination and instead use active markers on the object. Figure 13 shows a mobile surface embedded with a row of IR LEDs (the same components used in the main SecondLight surface). They are powered by 2 AAA batteries, and placed such that they shine horizontally into the mobile surface causing TIR to occur. Two lines laser-etched on opposite ends of the surface causes FTIR to occur continually in that region, generating a signal that provides a good approximation to the bounds of the object (see Figure 13 right inset) and allows the projective transform to be solved. Furthermore, the mobile surface supports multi-touch input at no extra cost because any touch points will cause additional FTIR which is detected by the SecondLight camera.

LIMITATIONS AND FUTURE WORK

We already have a number of enhancements in mind for our system which we plan to address in future work; some of these are presented here. Currently we only image in the IR spectrum, but imaging visible light could enrich the interactions that the system supports. For example, a high resolution digital stills camera located behind the screen could be triggered to capture images when the surface is clear, in a similar way to TouchLight. This would allow us to image or ‘scan’ objects such as documents in color (both on the glass using a projector for illumination, and potentially off the glass given enough ambient light). We have also verified the feasibility of imaging both ‘on’ and ‘through’ at 120Hz using two cameras, each of which is triggered twice during the relevant 8.3ms timeslot, once at the start and once towards the end, although again, we have not yet incorporated this into our prototype.

In practice we found the projector configuration described in the previous section, with the ‘through’ projection shining away from the front of the tabletop, limits the likelihood that light is projected directly at the user. For other scenarios it may be possible to use the through camera to detect and track users faces so that we can ‘project black’ in the region of any eyes detected. The use of two projectors which run continuously but whose output is optically shuttered is very flexible – for example it allows us fine-grained control of the duty cycle of the ‘on’ and ‘through’ projectors. But a solution which uses a single, high frame rate projector would be a lot less wasteful of light energy, and hence potentially brighter as well as being more compact and probably cheaper.

We use projectors with a high depth of focus, and adjust the ‘through’ projection to give a crisp image anywhere within around 150mm of the surface – beyond this it starts to gradually blur. A motorized focus system, combined with depth sensing of the surface being projected onto, would be

useful. However, this would still be limited to focusing on a single surface, whereas the image from a laser projector is always in focus, so this is an alternative technology we would like to explore.

CONCLUSIONS

In this paper we have introduced SecondLight, a novel surface technology which allows interactions beyond the display alongside more traditional on-surface interactions. The technology brings together diverse ideas and concepts from many different research areas and integrates these into a single, self-contained solution. The specific contributions of our work are as follows:

- The use of switchable diffusers for interactive surfaces.
- Simultaneous projection on and through the surface without cross contamination of these projections.
- Tracking of and projection onto objects in real-time through a rear projected tabletop whilst maintaining entirely different content on the primary surface.
- Projecting images through a tabletop onto perpendicular sides of tangibles including non-planar surfaces using prisms.
- FTIR multi-touch on the surface and on secondary displays above the surface with the sensing and processing integrated in the tabletop unit.
- Combined imaging on the surface, for robust touch and hover detection, with imaging through the display to enrich the sensing. For example, to track objects and hand gestures from a greater distance or 'see' the number of people and their orientation around the surface.

Switchable diffusers present an exciting technology for interactive surfaces and tabletops – allowing us to combine the benefits of a diffuse display surface with the ability to project and see through the surface, thereby extending interaction into the space in front of the display. We have described the hardware and software we used to construct our prototype in some detail, allowing others to explore these new forms of interactive surfaces.

ACKNOWLEDGMENTS

We thank our UIST reviewers and also colleagues at Microsoft for their input and feedback.

REFERENCES

1. Apple iPhone Multi-touch. <http://www.apple.com/iphone/>
2. Bier, E. A., et al. Toolglass and magic lenses: the see-through interface. In Proceedings SIGGRAPH '93.
3. Brown, L. D. and Hua, H. Magic Lenses for Augmented Virtual Environments. *IEEE Comput. Graph. Appl.* 26, 4, 2006.
4. Dietz, P. and Leigh, D. DiamondTouch: a multi-user touch technology. In *ACM UIST 2001*.
5. DisplayTech. FLC shutter datasheet. http://www.displaytech.com/pdf/photonics_shutters_ds.pdf.
6. Gross, M. et al. blue-c: a spatially immersive display and 3D video portal for telepresence. In ACM SIGGRAPH 2003 (San Diego, California, July 27 - 31, 2003). 819-827.
7. Han, J.Y. 2005. Low-Cost Multi-Touch Sensing through Frustrated Total Internal Reflection. In *ACM UIST 2005*.
8. Hodges, S., Izadi, S., Butler, A., Rrustemi, A., and Buxton, B.. ThinSight: Versatile Multi-touch Sensing for Thin Form-factor Displays. In *ACM UIST 2007*.
9. Holman, D., et al. Paper windows: interaction techniques for digital paper. In Proceedings of CHI '05. Oregon, USA, April 02 - 07, 2005, 591-599.
10. ImagingSource 2007. http://www.theimagingsource.com/en/products/cameras/firewire_mono/dmk21bf04/overview/.
11. Ishii, H. and Ullmer, B. 1997. Tangible bits: towards seamless interfaces between people, bits and atoms. In Proceedings of CHI'97 Georgia, United States, March 22 - 27, 1997.
12. JazzMutant Lemur. http://www.jazzmutant.com/lemur_overview.php
13. Kakehi, Y., et al. Lumisight Table: An Interactive View-Dependent Tabletop Display, *IEEE CGA*, 25(1) 48-53, 2005.
14. Kakehi, Y., Iida, M., and Naemura, T. Tablescape plus: up-standing tiny displays on tabletop display. In TABLETOP'07.
15. Lumisty, <http://www.glassfilmenterprises.com/lumisty.htm>
16. Jordà, S. Et al. 2005. The ReacTable. *International Computer Music Conference (ICMC2005)*, 2005.
17. Lee, J., Dietz, P., Aminzade, D., and Hudson, S. Automatic Projector Calibration using Embedded Light Sensors, *ACM UIST*, October 2004.
18. Lee, J. and Hudson S., Foldable Interactive Displays. In review. http://www.youtube.com/watch?v=nhSR_6-Y5Kg
19. LC-Tech, 2007. FOS-PSCT OPTICAL-SHUTTER datasheet. <http://www.lctecdisplays.com/files/datasheets/FOS-PSCT.pdf>
20. van Loenen, E. et al. Entertaible: A Solution for Social Gaming Experiences. In *Tangible Play workshop, IUI Conference*, 2007.
21. Malik, S. and Laszlo, J., 2004. Visual Touchpad: A Two-handed Gestural Input Device. In *ICMI 2004*.
22. Matsushita, N. and Rekimoto, J., HoloWall: designing a finger, hand, body, and object sensitive wall. In *ACM UIST 1997*.
23. Microsoft Surface, <http://www.surface.com>
24. Perceptive Pixel, <http://www.perceptivepixel.com>
25. Projective Transformation, http://en.wikipedia.org/wiki/Projective_transformation
26. Raskar, R., Beardsley, P.A., A Self-Correcting Projector, *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2001.
27. Shiwa, S., Ishibashi, M.: A Large-Screen Visual Telecommunication Device Enabling Eye Contact. In Digest of technical papers of Society for Information Display International Symposium: SID1991. pp. 327-328, 1991.
28. Smith, D., Graham, N., Holman, D., Borchers, J. Low-Cost Malleable Surfaces with Multi-Touch Pressure Sensitivity, *tabletop*, pp. 205-208, TABLETOP'07, 2007.
29. Sullivan, A. A Solid-State Multi-Planar Volumetric Display, *SID Symposium Digest Tech Papers* 34, 1531-1533 (2003).
30. Underkoffler, J. and Ishii, H. 1999. Urp: a luminous-tangible workbench for urban planning and design. In Proceedings of CHI'99, May 15 - 20, 1999, 386-393
31. Wilson, A.D. TouchLight: An Imaging Touch Screen and Display for Gesture-Based Interaction. In *ICMI 2004*.
32. Wilson, A.D. 2005. PlayAnywhere: A Compact Interactive Tabletop Projection-Vision System. In *ACM UIST 2005*.