Enabling Concurrent Dual Views on Common LCD Screens

Seokhwan Kim[1,2], Xiang Cao[1], Haimo Zhang[1,3], Desney S. Tan[1]

[1]Microsoft Research
Beijing, China; Redmond, WA, USA
{xiangc, desney}@microsoft.com

[2]University of Tsukuba
Tsukuba, Ibaraki, Japan
skim@iplab.cs.tsukuba.ac.jp

[3]National University of Singapore
Singapore
z-haimo@comp.nus.edu.sg

ABSTRACT

Researchers have explored a variety of technologies that enable a single display to simultaneously present different content when viewed from different angles or by different people. These displays provide new functionalities such as personalized views for multiple users, privacy protection, and stereoscopic 3D displays. However, current multi-view displays rely on special hardware, thus significantly limiting their availability to consumers and adoption in everyday scenarios. In this paper, we present a pure software solution (i.e. with no hardware modification) that allows us to present two independent views concurrently on the most widely used and affordable type of LCD screen, namely Twisted Nematic (TN). We achieve this by exploiting a technical limitation of the technology which causes these LCDs to show varying brightness and color depending on the viewing angle. We describe our technical solution as well as demonstrate example applications in everyday scenarios.

Author Keywords

Dual-view display, LCD, twisted nematic.

ACM Classification Keywords

H.5.m. [Information interfaces and presentation (e.g., HCI)]: Miscellaneous.

General Terms

Design, Human Factors.

INTRODUCTION

Display devices that are capable of presenting two or more different views concurrently for different viewing angles and/or different viewers, or multi-view displays, have attracted increasing attention in recent years. Such displays provide interesting affordances beyond conventional single-view displays. For example, they may support multiple people viewing personalized information, protect private information from bystanders, or enable natural stereo 3D viewing experiences. To support these applications, a variety of multi-view display technologies have been developed, some that have viewers wear special glasses as selective filters, and others that focus on special optical designs to manipulate light routes so as to present varying information in different directions.

Despite the appeal of these technologies, their requirement for specialized (and often expensive as well as cumbersome) display hardware has limited their adoption by general consumers for use in daily scenarios. To address this challenge, we present a simple and pure software solution that enables two independent views for different viewing angles concurrently on the most common and affordable type of LCD (Liquid Crystal Display) screen, namely Twisted Nematic (TN), without any hardware modification or augmentation (Figure 1). Known for its cost effectiveness and power efficiency, as well as good response speed, TN type LCDs are the current default for low-to-mid-end computer monitors, especially laptop screens. For example, in the first quarter of 2010, 92% of shipped laptops were using TN type LCD screens [17]. Our solution can be easily employed on such LCDs that are already ubiquitous with no additional cost, and potentially make multi-view display applications truly available every day and everywhere.

Figure 1. A common laptop screen based on TN LCD showing two images concurrently for different viewing angles. (a) Bottom view. (b) Top view.

Our solution is made possible by deliberately exploiting a limitation of the TN LCD technology, namely that the observed brightness and color of these LCDs vary when viewed from different angles. This well-known effect results in their so-called “narrow view” and is generally deemed as a drawback of the TN technology. However, by carefully examining the characteristics of such changes, we can intentionally manipulate the pixel colors of an image so that its observed contrast is maximized or minimized, effectively showing or hiding it, at different viewing angles. By spatially or temporally multiplexing two such images optimized for alternate angles, we are able to display two independent views concurrently, each for a different viewing angle. We have tested our solution to work robustly on a variety of TN LCDs. In this paper, we describe our technical solution for displaying dual views on off-the-shelf TN LCDs, as well as present demonstrations for potential daily applications.

RELATED WORK

Multi-View Display Technologies

There exist a variety of technologies that enable two or more concurrent views on a single display device. Here, we discuss representative categories.

Technologies Requiring Special Glasses

Polarizer glasses are perhaps the most familiar technology, widely used for stereoscopic 3D displays. By polarizing the light emitted from the screen in one of two perpendicular directions, one of the two displayed images becomes visible and the other blocked to the viewer. Along the same principle of filtering out the irrelevant views, shutter glasses work by having the screen switching between displaying multiple images in alternate frames. Unlike polarizer glasses which are restricted to two views, with high refresh rate displays it is possible to provide more than two views using shutter glasses [9].

These glasses-based technologies are not dependent on the viewing angle thus extremely flexible. However, they require the viewer to have and wear special glasses, which are not only intrusive but also impractical sometimes, e.g. in privacy protection scenarios. Compared to these, other technologies attempt to display different views according to the viewing angle, thus not requiring augmentation on the viewer side.

Technologies Built into Displays

In a parallax barrier display, fine vertical gratings are placed in front of the screen [4]. Well aligned gratings can block or pass the light emitted from specific spatial elements in the display depending on the viewing angle. By spatially interlacing columns from the two images on the screen behind, the viewer sees only one image from a certain angle. Following a similar principle to parallax barrier, a lenticular lens display places a thin sheet of cylindrical lens array in front of the screen [4]. Again columns from two images are spatially interlaced on the screen. Unlike parallax barrier which blocks half of the light, lenticular lens refracts the light in different directions for different viewing angles, thus offers higher brightness [1]. Besides the above more well-known technologies, there are also other more specialized solutions. For example, time-sequential display shows each view in sequence and use an optical lens to direct each of them to a different viewing angle [10,15]. Two-view two-projector display [4] uses a half-silver mirror to produce two viewing regions in space.

As can be seen, all the multi-view display technologies described above require specialized hardware, hence accessibility to them is very limited for the general public. In contrast, our solution works on widely available TN LCD screens without any hardware modification or augmentation.

Multi-View Display Applications

Currently the most popular and mature application of multi-view displays is perhaps to enable stereo 3D viewing experience. Beyond the more traditional glasses-based 3D displays [2], there has been increasing interest in auto-stereoscopic displays (or so called "naked eye 3D display"), which present the stereo image pair by leveraging the viewing angle difference between the two eyes [3].

Matsusik et al. suggested interesting applications using multi-view displays for a single user [9]. These include discrete views, where the user can watch different domains of information according to the view angle; and layered display, where the user can move their head to see more or less details in the information. This prevents interference between multiple overlapping pieces of information and overload to the user.

In addition to single user scenarios, researchers have explored presenting multiple personalized views on a shared display, especially to enable private vs. public interaction. For example, Shoemaker and Inkpen [12] used a shutter glasses system to support both sharing a public area and seeing private information in their research on single-display groupware. Smith and Piekearski [13] proposed a tabletop system using a lenticular-style lens, which displays personalized content to be viewed from different sides of the table, and clear demarcation of public and private content on the tabletop.

We share the perspective of these works that multi-view displays can be useful for various application domains. In this paper, our focus is not on inventing new applications, but a technology that aims to make such applications widely accessible to all users. On the other hand, because of this almost universal availability, some daily applications that were not considered in the past due to limited availability and benefic/cost ratio now becomes worthwhile, of which we will give examples in the paper.

Exploiting TN LCD Property for Special Display Effects

Most recently and based on a similar principle to ours, Harrison and Hudson [7] presented how to exploit the optical properties of TN LCDs to display one piece of single-colored supplementary information in addition to an existing image in oblique angles, which is invisible from straight-on angles. Their technique is suitable for unidirectional privacy scenarios. In comparison, we develop a generic dual-view solution that enables displaying two independent and arbitrary images at customizable angles, thus are able to support general multi-view applications.

TN-BASED LIQUID CRYSTAL DISPLAY (LCD)

Here we first introduce basic principles and characteristics of LCDs to help understand how our solution is possible.

An LCD comprises of a matrix of LC (liquid crystal) molecules between two polarizers, and a uniform backlight beneath them (Figure 2). The two polarizers are polarized in perpendicular directions so that by default the backlight cannot pass through. However, when the polarized light coming from the first polarizer passes through the LC matrix, its polarization direction rotates according to the direction of the LC molecules, making it no longer perpendicular to that of the second polarizer. Thus the resulting light is able to pass through the second polarizer. The exact amount of light passing through depends on the angle between the LC molecules and the two polarizers. Varying the voltage applied to the LC molecules controls their direction, and in turn the light intensity eventually emitted from the screen. Extending this principle, each screen pixel consists of three color filters (red, green, and blue, or R, G, and B) and three independently controlled groups of LC molecules so that it can produce various colors.

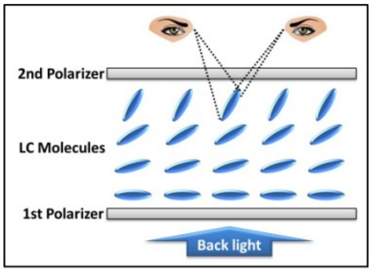

Figure 2. Principle of (TN) LCD.

Depending on the specific type of LCD technology, the LC molecules are rotated in different fashions. In particular, in TN LCDs, the LC molecules are rotated within a plane perpendicular to the screen as Figure 2 illustrates. Because of this, when the viewer looks at the screen from different angles, the line of sight (hence the line of light transmission) is also at different angles with regard to the direction of the LC molecules. This results in the light polarization directions being rotated differently by the LC molecules, leading to different light intensities emitted from the same pixel to different angles. In addition, because R, G, and B lights respond to the LC molecules slightly differently, this may also result in color shift. This is the reason for the well-known phenomenon of these LCDs' varying brightness and color depending on the viewing angle, often described as a "narrow view."

Figure 3 plots the image brightness (in terms of pixel values in a photo of the LCD taken by a digital camera, as detailed in the "Measuring brightness curves" section) we measured from a range of viewing angles, both vertically and horizontally, on an HP TouchSmart tm2 tablet PC, which uses a common TN LCD. The LCD was measured while placed statically in a landscape orientation. For vertical viewing angles, negative angles mean viewing from the bottom and upward the screen (denoted "bottom views") and positive angles mean viewing from the top and downward the screen ("top views"). As a precaution for potential confusion, note that in situations of a laptop with tiltable screen and the viewer sitting statically in front of it, bottom views are observed when the screen is tilted facing upwards, and top views are observed when the screen is tilted facing downwards (as seen in Figure 1).

Similarly, for horizontal viewing angles, negative angles mean viewing from the left to the screen and positive angles mean viewing from the right. We can see that the R, G, and B channel curves all follow the same trends, with slight differences in the exact numbers. Although we measured numbers in these graphs from a single LCD, theoretical analysis on LCD performance [8,14,16] suggests that the general trends of these curves generalize to all TN LCDs, even though the exact numbers may vary between devices.

As we can see, the vertical viewing angles show much more dramatic changes in light intensity than the horizontal viewing angles do. This is because when the line of sight is within the same plane as the LC molecule rotation (as shown in Figure 2), the angle between these two also changes dramatically along with the viewing angle, while if the line of sight is perpendicular to the rotation plane, this correlation is much less drastic. LCD manufactures usually set the LC molecule rotation plane to optimize for a "wider view" horizontally as this is the direction in which viewers are more likely to be moving or distributed.

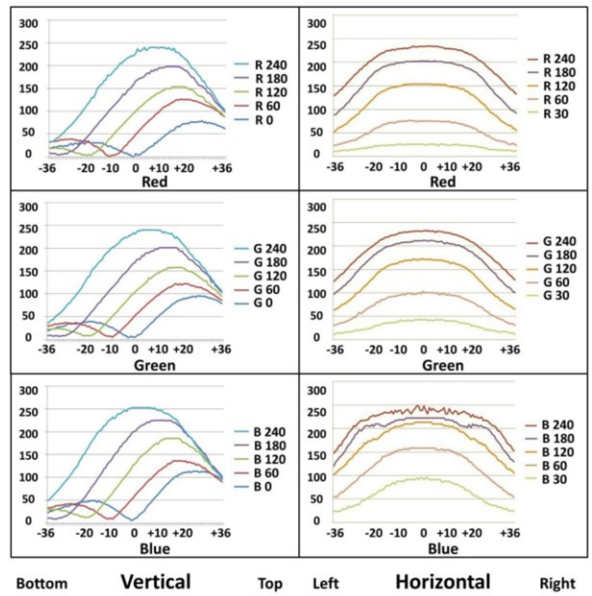

Figure 3. Image brightness measured from different viewing angles of a typical TN LCD, separated by color channel (R, G, B). X-axis represents the viewing angle, and Y-axis represents observed image brightness. Each brightness curve represents a different pixel value being displayed (e.g. R 240 means a pixel value of RGB (240, 0, 0), each color channel displays in the range of 0-255, here only showing representative curves). Y-coordinate of the curve at 0° represents the "true" brightness seen from the front.

Besides TN, there exist a number of other LCD technologies, such as VA (Vertical Alignment) and IPS (In Plane Switching). These technologies are less prone to varying brightness/color across viewing angles (i.e. provides a "wider view"), but are generally more expensive. As such, they are mostly used for higher-end displays. TN with its cost and power efficiency, as well as faster response, remains the mainstream for low-to-mid-end computer monitors, especially laptop screens.

SHOWING AND HIDING A SINGLE IMAGE

We now explain how we can show and hide a single image at different viewing angles by leveraging the characteristics of TN LCDs (Figure 3). In the next section we will describe how to extend it to concurrently showing two different images at different angles.

General Principle

In general, in order to show the same image at viewing $angle_{show}$, and hide it at viewing $angle_{hide}$, the image should consist of pixel colors that maximize their observed contrast at $angle_{show}$, and at the same time have an observed contrast at $angle_{hide}$ below the threshold t of perceivable contrast. Hence the question is reduced to finding such a combination of pixel colors on a given LCD for a given pair of $angle_{show}$ and $angle_{hide}$. Note here contrast can be conveniently represented as the difference between observed brightness values, which may also be equivalently converted to the contrast ratio in terms of luminance via a logarithmic relationship [6].

Single Color Channel

We take a divide-and-conquer approach and first focus on enabling showing and hiding an image consist of a single color channel (R, G, or B). By examining the curves in Figure 3, we can see that over the range of the negative vertical viewing angles, many curves intersect with one another. Simply put, wherever two curves intersect, it indicates these two pixel color values appear exactly the same from this viewing angle, thus can be used to hide the image perfectly. On the other hand, each pair of curves also diverges quickly beyond the intersection point, meaning they are indeed capable of showing the image at other angles. Similarly, many curves converge quickly when the vertical viewing angle moves towards larger positive angles, which are also promising candidates for hiding information at these angles. On the contrary, in the horizontal viewing angles, all the curves are roughly parallel and do not intersect, meaning it might be infeasible to hide an image by changing the horizontal viewing angle only (after all, the LCDs are optimized for maintaining more visibility in horizontal viewing angles).

Thus in this paper we focus on providing dual views at different vertical viewing angles. Nonetheless, with rotatable desktop monitors and tablet PCs becoming commonplace, one can easily rotate the LCD to a portrait setting in scenarios where dual views at horizontal viewing angles are more appropriate, as will be demonstrated in some of our applications later. Based on these observations, we now explain how to find the optimal combination of pixel values for a single color channel.

Binary Image

We start from the simplest case, where we show and hide a binary image that consists of only two pixel values and in a single color channel (R, G, or B). To find such a pair of values, we developed a simple automatic algorithm, which takes $angle_{show}$, $angle_{hide}$, t, as well as the brightness curves for the color channel of the particular LCD as input. The algorithm first searches for all possible pairs of pixel values that have an observed contrast (i.e. difference in observed brightness) $< t$ at $angle_{hide}$. Then, among these pairs, it searches for the pair that has the largest observed contrast at $angle_{show}$. This pair is then chosen to render the binary image.

Figure 4 illustrates a binary image being shown and hidden at alternate viewing angles (+/−25°) using the G channel. For the purpose of illustrating the technique, we give examples using the G channel, but the same effects can be achieved using R or B channels as well. The pixel value pairs are found by the algorithm using the curves in Figure 3, and t is empirically determined to be 10. Table 1 lists these values and their respective observed intensities at the two angles, where Pair *a* is used to render the image to be shown the image at +25° and hidden it at −25°, and Pair *b* is used to do the opposite. Similarly, Pair a' is for the image to be shown at +10° and hidden it at −10°, and Pair b' for the opposite.

Figure 4. Binary image. (a) Image shown at +25° and hidden at −25°. (b) Image hidden at +25° and shown at −25°.

		Pair a			Pair b		
Pixel Value (G)		1	190	Observed contrast	202	255	Observed contrast
Observed Intensity	+25°	96	166	70	163	173	10
Observed Intensity	−25°	35	38	3	23	138	115
		Pair a'			**Pair b'**		
Pixel Value (G)		1	105	Observed contrast	241	255	Observed contrast
Observed Intensity	+10	59	133	74	233	244	11
Observed Intensity	-10	27	32	5	198	234	36

Table 1. Optimal pixel value pairs (G channel) for the TN LCD mentioned in Figure 3.

As has become clear in Table 1, one inevitable result of images being rendered this way is the reduction of observed contrast at $angle_{show}$ compared to the full contrast supported by the LCD. This is caused by the smaller range of pixel values that can be used, and in some cases (e.g. Pair a at +25°) also by the oblique viewing angle. Take Pair b at -25° for example, the observed contrast 115 is approximately half (0.46) of the maximally possible contrast (248) of this LCD (which when measured from 0° has an observed brightness of 253 when displaying pixel value 255, and 5 when displaying pixel value 0).

We should also note that the image is not only hidden at exactly $angle_{hide}$, but also the neighborhood around it in which the observed contrast remains unperceivable. The range of this neighborhood varies by device and $angle_{hide}$ itself, and is generally between 5-10°. At viewing angles between $angle_{show}$ and $angle_{hide}$, the observed contrast is usually between those at $angle_{show}$ and at $angle_{hide}$ which means the image is generally still visible, albeit with a further reduced contrast.

Continuous Brightness Range

When examining Figure 3, we notice that although each pair of curves may intersect at a different point, in general neighboring curves intersect at neighboring positions both in terms of viewing angle and in terms of observed intensity. This suggests that if instead of only using the optimal pixel value pair found for $angle_{show}$ and $angle_{hide}$, we use the continuous range of values between the pair, the observed contrast may still be low enough to hide the image. To do so, we take an existing grayscale image, and perform a linear transform of its pixel values to envelop them between the pair, so that the original maximal pixel value maps to the higher value in the pair, and vice versa, i.e.:

$$Pixel_{render} = Pair_{min} + (Pixel_{orignal} - Orignal_{min})\,(Pair_{max} - Pair_{min})\,/\,(Original_{max} - Original_{min})$$

where $Pair_{min}$ and $Pair_{max}$ are the lower and higher value in the optimal pair, $Original_{min}$ and $Original_{max}$ are the minimal and maximal pixel values in the original image, and $Pixel_{original}$ and $Pixel_{render}$ are the original and rendered value for each pixel.

Figure 5. Image rendered with continuous brightness range (G channel only). (a) Image shown at +25° and hidden at −25°. (b) Image hidden at +25° and shown at −25°.

Figure 5 illustrates an image being shown and hidden at alternate viewing angles (+/−25°) with continuous G levels, deriving from the optimal pixel value pairs in Table 1. As can be seen the image now displays more subtle details, while at the same time still perfectly hidden at $angle_{hide}$.

Combining Three Color Channels

As R, G, B color channels are perceived independently by human (as well as by cameras), combining them to enable showing and hiding of colored images is relatively straightforward. Taking an arbitrary colored image as input, for each one of its 3 color channels, we can separately and independently determine the rendered pixel values (binary or continuous range) according to the methods in the previous section, and simply remix the 3 rendered channels into the resulting colored image.

8-Color Image

Combining only the optimal pairs for each color channel (Table 2), we can display a collection of 8 colors (2×2×2) in total (red, green, blue, yellow, cyan, magenta, black, white) at $angle_{show}$, which is sufficient for many applications. To further increase the color expressiveness, we adopt the common technique of image dithering (e.g. used in inkjet printers), which simulates continuous colors by using spatial dot patterns from a small set of colors. To do this, we use the Floyd–Steinberg dithering algorithm [5].

	Pair a (show at +25°, hide at +25°			Pair b (hide at +25°, show at +25°		
	R	G	B	R	G	B
Pixel Value	1, 202	1, 190	1, 198	241, 255	202, 255	161, 255
Observed intensity at +25°	77, 175	96, 166	113, 167	161, 168	163, 173	167, 168
Observed intensity at-25°	29, 30	35, 38	41, 41	3, 122	23, 138	11, 157

Table 2. Optimal pixel value pairs (R, G, B) for showing/hiding at +/−25° on the TN LCD mentioned in Figure 3.

Note that these colors are not shown at their full saturation or brightness, as the observed brightness for each color channel in each pixel can be neither 255 nor 0. Slight hue shift from the intended color may also happen, as the observed brightness may not be uniform across the 3 channels. Nonetheless, the resulting view at $angle_{show}$ is sufficiently clear and vivid, and the image is still well hidden at $angle_{hide}$. Figure 6 illustrates this.

Figure 6. 8-color image with dithering. (a) Image shown at +25°, hidden at −25°. (b) Image hidden at +25°, shown at −25°.

"Full-Color" Image

Further extending the displayed color space, we can combine the continuous brightness range from each color channel to render the image with the maximal possible continuous color range at $angle_{show}$. Figure 7 illustrates this. Compared to 8-color dithering, which preserves maximal possible contrast/saturation by using only saturated colors, such "full-color" images present more subtle color resolution, at the cost of lesser contrast/saturation. The exact image contrast and saturation is dependent on the available observed contrast from the optimal pixel value pairs in each channel. On the other hand, as each color channel has slightly different fluctuations in terms of the observed brightness within the rendered pixel value range, now the image can no longer be perfectly hidden at $angle_{hide}$. Nonetheless, such barely observable contrast at $angle_{hide}$ will be easily overwhelmed by the other image to be shown at this angle, as we describe below.

Figure 7. "Full-Color" image. (a) Image shown at +25° and hidden at −25°. (b) Image hidden at +25° and shown at −25°.

SHOWING DUAL IMAGES BY MULTIPLEXING

Having described the full capability of showing and hiding a single image, we now explain how to simultaneously present two different images already rendered for showing/hiding at alternate viewing angles. This would achieve our goal of simultaneously displaying two images to different viewing angles.

Obviously, simple alpha blending of these two rendered images would not work as the interpolated pixel values no longer comply to either of the two rendering color sets. We therefore need a way to concurrently display these two images, while maintaining the pixel values for each of them. To do this, we employ multiplexing, a common technique for dividing a medium into several mutually separated (often interlaced) segments to transmit signals. Two types of multiplexing were used: spatial and temporal.

Spatial Multiplexing

Spatial multiplexing interlaces two images in the spatial domain by using alternating pixels. For example, having Image a and b rendered using Pair a and b respectively, the final rendered image contains half of the pixels from Image a, and the other half from Image b (Figure 8a). Thus when viewed from either of the two angles, one image becomes visible while the other image becomes a uniform color (nearly black or white). As the two are interlaced on a fine spatial granularity (pixel-level), the viewer simply sees one continuous image (Figure 8b).

Temporal Multiplexing

Temporal multiplexing interlaces two images in the temporal domain by displaying one image at every odd-numbered frame and the other at every even-numbered frame (at 60Hz in our implementation). Similarly to spatial multiplexing, at either viewing angle, the odd (or even) frames show one image while the other frames are blank (black or white). Human visual persistence creates the perception of a single continuous image (Figure 9).

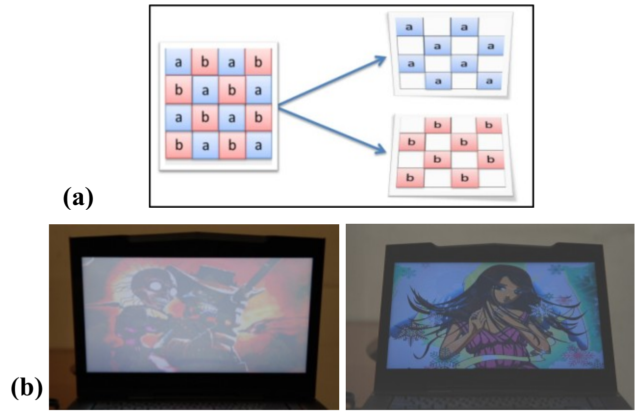

Figure 8. Spatial multiplexing. (a) Interlacing two images in alternate pixels. (b) Two images seen from different angles.

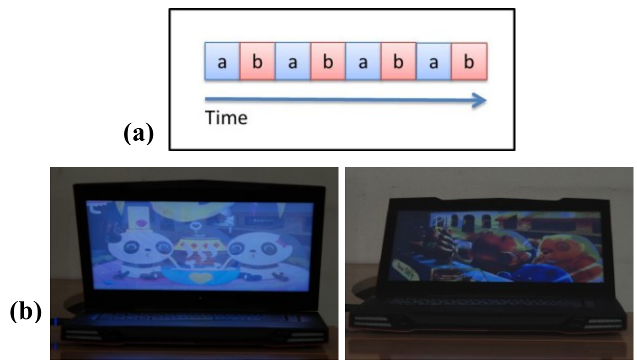

Figure 9. Temporal multiplexing. (a) Interlacing two images in alternate frames. (b) Two images seen from different angles.

Both multiplexing methods sacrifice resolution in one domain in exchange of maintaining the resolution in the other. Comparatively, spatial multiplexing is more advantageous as it does not introduce intrusive visual flickering, and the full procedure is embedded into a single static image that can be shown without special programs. All later examples in the paper were created using spatial multiplexing unless otherwise indicated.

One issue for both multiplexing methods is the reduction of image saturation, brightness, and/or contrast, as the image being shown is effectively blended with a black or white background resulting from the hidden image. To address this issue, the rendering algorithm may intelligently determine if the available contrast becomes too low according to the brightness curves, and where applicable switch from rendering in "full-color" to rendering in 8-color dithering to compensate for the loss of contrast/saturation.

MEASURING BRIGHTNESS CURVES

We have presented our solution based on the assumption that we know the exact brightness curves of the particular LCD (such as shown in Figure 3). We now explain how we gain such specific knowledge. While previous research has shown it is in theory possible to calculate such curves through simulation [14], this requires tens of intrinsic parameters of the LCD that are generally inaccessible without elaborate calibration and calculation [8,16]. We developed practical methods to empirically measure these curves.

Measuring using a Camera

A digital camera is essentially a multi-channel light sensor array, which can be used to measure the brightness of the LCD as viewed from different angles. We started with a brute-force method by setting up the camera at a fixed distance from the LCD in a dark room, and with all automatic settings turned off. We rotated the LCD screen in front of the camera vertically and horizontally between −60° and +60° and at 10° intervals, with the rotation angle measured by a protractor. At each rotation angle, the LCD displays a sequence of pure R, G, and B colors and covering the full range of pixel values (0-255) at 30 intervals for each of the three channels. The camera takes a photo of each of these colors, and samples the captured color in the center of each photo as the observed brightness. Aggregating all these samples results in curves similar to those in Figure 3, and in-between curves may be interpolated. Needless to say, this method is very time-consuming, and also requires precise control and measurement of the rotation angle. To increase efficiency and precision of the camera-based measurement, we developed the following method.

(a) (b)

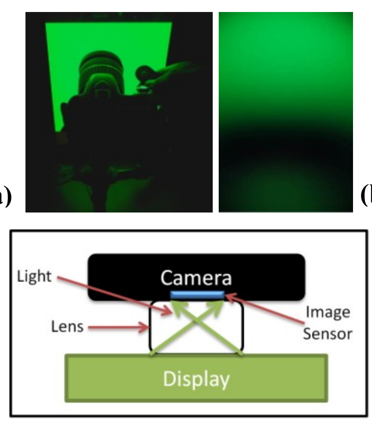

(c)

Figure 10. Measuring brightness curves using a digital camera. (a) Setup. (b) Photo taken. (c) Principle.

Inspired by methodology used by several LCD reviewers, we take a digital camera and push its lens directly against the LCD screen (Figure 10a). What this enables is that each pixel on the image sensor is essentially observing the LCD at a different angle, resulting in a wide and continuous range of both vertical and horizontal viewing angles (Figure 10c). Therefore a single photo incorporates sufficient brightness information to generate two complete brightness curves, one vertical and one horizontal, for the color being displayed by the LCD (Figure 10 b, one can easily see how this photo matches the trends in Figure 3). This not only significantly increases the efficiency of the measurement, but also results in a very high resolution in terms of the angles being measured for.

Interactive Calibration with Naked Eyes

The above camera-based measurement method allows comprehensive recovery of the brightness curves, based on which we can automatically extract optimal pixel color combinations for any viewing angles. However, in many cases end users only need to quickly find rendering parameters that work for one particular scenario or setup (i.e. two particular viewing angles) with their LCD at hand. To provide an even lighter-weight way of calibrating our system, we developed an interactive program that helps the user easily find two approximate pixel color pairs for displaying dual views in two particular angles, using their naked eyes for judgment.

Examining the values in Table 1 and 2 reveals that for all 3 color channels, pixel value = 1 is always among the optimal pixel value pair for showing at top views and hiding in bottom views (Pair a). Similarly, pixel value = 255 is always among the optimal pair for the opposite case (Pair b). Based on this empirical finding, we assume that we only need to find the opposite R, G, B values in the two pairs. To do so, the user first looks at the LCD from the bottom viewing angle s/he would want to use. The program displays a black block on a black background (both with RGB=1,1,1) which is surely indistinguishable. The user uses a slider to increase the R value of the block until it just becomes distinguishable from the background. The program records this R value for Pair a. The same is repeated for collecting G and B values for Pair a. Similarly, for Pair b the user looks from the top viewing angle s/he wants to use, and the program starts from a white block on a white background (both with RGB = 255, 255, 255), and user decreases the R, G, or B values of the block until it becomes distinguishable, recording the values for Pair b. The entire calibration process can usually be done easily within a few minutes. Figure 11 illustrates.

(a)

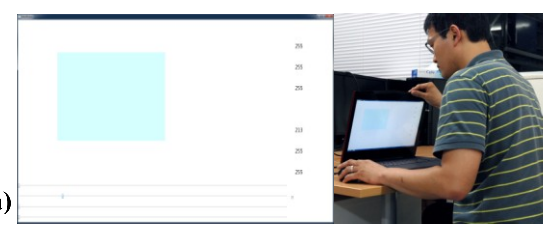

(b)

Figure 11. Interactive calibration. (a) Program interface. (b) A user collecting pixel values for a top viewing angle.

Generalizability across Different LCDs

Using the camera-based method, we measured three LCDs in addition to that in Figure 3, which included (1) Dell E2310Hc desktop LCD (2) Dell M15x laptop and 3) Dospara BL212 laptop. All four LCDs showed similar trends in the brightness curves, while the exact brightness and viewing angle vary. We were able to extract optimal color pairs for all of them and display clean dual views.

Especially, we were interested in examining how well the pixel value pairs optimized for one LCD may work on another. Figure 12 illustrates the same image pair shown on the Dell M15x laptop with three different setups: (a) using its own optimal color pairs for +/−25°, viewed from +/−25° (b) using HP tablet PC's optimal color pairs for +/−25°, viewed from +/−25° (c) using HP tablet PC's optimal color pairs for+/−25°, but viewed from +17°/−43°. (a) unsurprisingly looks best, while (b) shows slight ghosting from the other image. However, by slightly adjusting the viewing angle, (c) shows clean images again. We found that for the majority of time, a dual-view rendering for one laptop can still work well without modification on other laptops with readjusted viewing angles. This generalizability indicates that for many common use cases, measurement of the particular LCD is not necessary, and further assures the wide applicably of our solution.

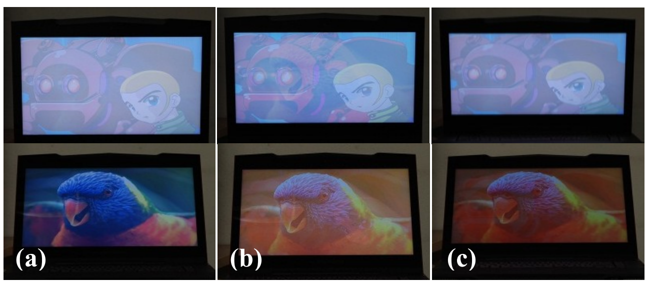

Figure 12. Various setups on a Dell M15x laptop.

EXAMPLE APPLICATIONS

Unlike previous multi-view display applications which are mostly targeted at higher-end applications, given the wide availability of TN LCDs and that they are already possessed by the majority of users, our solution may be easily incorporated into many daily application scenarios.

Movie Player

By allowing a movie player to play two different movies simultaneously, multiple people can enjoy different programs on the same screen. The default setup of presenting dual views for different vertical angles may interestingly make for family scenarios where adults and children may see different movies suited for their interest depending on their height (Figure 13). Hence the viewing angles are not just abstractly mapped to the content but convey sematic meanings.

Figure 13. Two people, one standing and one sitting, watching different programs on the same screen.

Personalized Game Perspective

Current video game players usually rely on split screen views when they play multi-player first-person perspective games with co-located friends. This is not only an inefficient usage of screen real estate, but also does not allow showing game information private to one user. Figure 14 illustrates using a tablet PC in portrait setup to play the classic pong game. Our solution allows two players standing side by side in front of it and each sees an exclusive first-person perspective of the game, taking up the entire screen but without interfering with the other player's view.

Figure 14. Two players both see personalized first-person perspective full-screen game views.

Figure 15. Two players play a card game. A spectator watches from the side. All three people have different views.

Figure 15 shows another interesting setup, where two players facing each other play a card game on a touchscreen tablet PC laid flat between them, similar to an interactive tabletop setup. Each player only can see the information on their own cards in the area near themselves, whereas they can only see the back of the cards in their opponent's area. The region between the two is public and is visible to both. This private/public interaction demarcation somehow resembles that in [13], however only requires a regular touch tablet and no cumbersome hardware. Further, a spectator sitting between the two players is able to see cards from both players, as both players' views are visible (albeit with a lower contrast) from such an intermediate viewing angle. Without deliberate design, our simple technology effectively supports three different views that inherently suit the three roles in the game.

Protecting Privacy

Although the above example touches upon private information in a game, for more critical privacy application such as banking it requires the private information to be visible only to the user but nobody else. Considering that with our current solution we only *hide* each view in a relatively small angular range, we need to devise a way to do the opposite, i.e. only *show* the information in a small angular range. To serve this need, we adopt a solution using a random dot pattern. This is inspired by random dots patterns employed in some non-digital copyright protection solutions [11]. The principle is that the critical information is surrounded by a random dots pattern, making it extremely difficult for humans to segment, effectively hiding it. Only when the random dots pattern is removed does the critical information becomes visible. Applying this principle, if we want to present private information at one specific angle only, we not only render the information with this angle as $angle_{show}$, but also surround it with a random dot pattern (here black and white) rendered with this angle as $angle_{hide}$. The effect is that only from this specific viewing angle, the random dot pattern disappears and the information becomes perfectly visible, while from all other viewing angles the random dot pattern makes the information undetectable to the observer. By doing so, we essentially swapped the showing and hiding range. As shown in Figure 16, the private text is illegible from all angles (including horizontal) except a narrow viewing angle from the bottom.

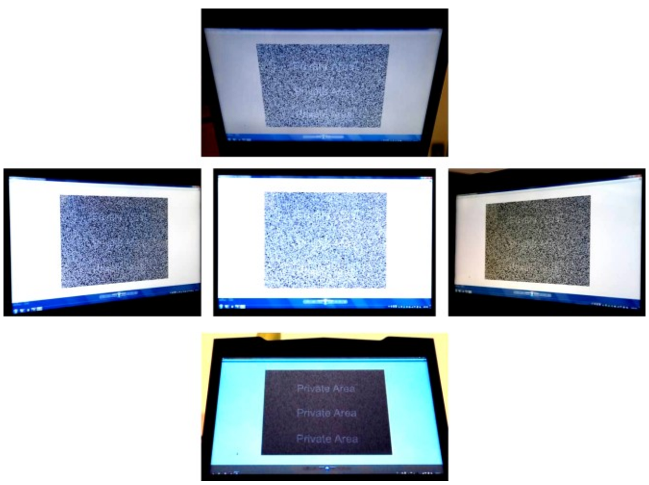

Figure 16. Privacy information surrounded by a random dot pattern, only visible from one specific viewing angle.

Autostereoscopy

Similar to other multi-view display technologies, our solution may also be potentially used for presenting autostereoscopic images to enable 3D viewing experience with naked eyes. This essentially turns any TN LCD into a 3D display when they are set to portrait orientation. Figure 17 demonstrates this on a tablet PC by presenting a pair of stereo images to the two eyes respectively. Like other autostereoscopic displays, the 3D sensation is dependent on the viewers' distance and position [3,9]. To assist the user find such optimal distance and position easily, we displayed "L" and "R" characters in the left-eye and right-eye views respectively, so that when users see these letter with different eyes they know they have reached the optimal point. One limiting factor here is because the two eyes are looking from angles that differ by a small amount only, in order to render the two independent views, the image contrast needs to be further reduced compared to other situations.

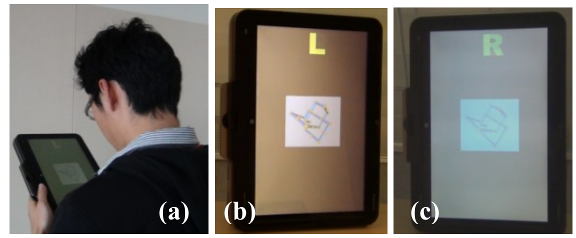

Figure 17. (a) A regular tablet PC serves as an autostereoscopic display. (b) Left eye view (-15 °) (c) Right eye view (+15 °)

Mirror Effect

An interesting phenomenon happens when placing the LCD sideways near a mirror. This results in the LCD itself and its virtual image in the mirror being at different angles from the user's eyes, hence enables the user to see different content inside and outside the mirror. Besides the apparent magical effect, this could also have practical applications such as creating a virtual second monitor, an extremely cheap solution for extending the screen real estate. Figure 18 illustrates.

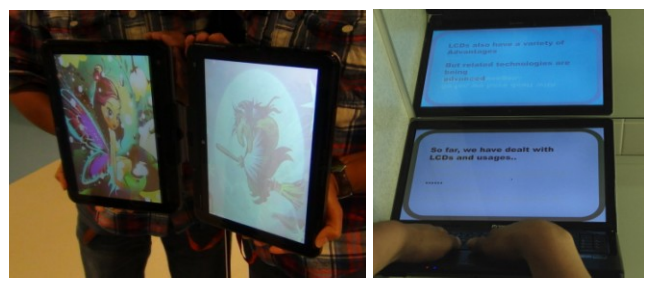

Figure 18. A mirror reveals the alternate view.

INFORMAL USER FEEDBACK

Although we did not perform a formal user study (mostly because this is a purely optical phenomenon and we expect variance between users to be quite minimal), we have demonstrated our work to more than 50 users in various occasions and under various lighting conditions from typical daylight to pure darkness. All users were very positive and pleased about the effects they observed. Most did not notice the reduction of contrast until prompted, and in general found the images natural to watch. Given that the viewing range for each independent view is relatively tolerant (5-10°), most users did not have a problem positioning themselves or the screen to see the different views.

DISCUSSION AND FUTURE WORK

Several interesting and challenging issues are worth further investigations in the future.

We have described our solution that allows us to present two concurrent views on TN LCDs. Whether it is possible to present three or more views remains an open question. The card game example seems to shed light on this, although in that case the three views are not mutually independent. Based on our understanding of the brightness curves, a more sophisticated algorithm might be possible to simultaneously optimize for three views. A related question is whether it is possible to use a similar global optimization method to achieve more than one view without resorting to multiplexing, i.e. truly reusing the same pixels for both views. This is something that requires deeper investigation.

For measuring the brightness curves and acquiring optimal pixel color combinations, we have developed both a comprehensive objective method using a camera, and a lightweight subjective method through user interaction. We would like to explore methods that get the best of both worlds. On the one hand, we plan to design a composite calibration pattern that incorporates multiple colors in a single camera shot to further increase the measurement efficiency. On the other hand, we would like to find a simplified parameterized model for the brightness curves, and using an improved interactive calibration procedure to recover these parameters and in turn the full curves.

We have shown that a single camera photo can capture screen brightness observed from multiple angles. However, this may also be true for human eyes. Especially when the screen is very large or very close to the viewer, s/he is indeed looking at different locations on the screen at slightly different angles. This may result in imperfect hiding/showing effects as the image is optimized for a single viewing angle. However, if we anticipate such effect, it may be pre-compensated by optimizing different parts of the image for slightly different viewing angles. We would like to investigate this in detail in the future.

In our current experiments we have used an empirically determined threshold t for perceivable contrast. In the future we would like to ground this with deeper understanding of human visual cognition. In addition, we are also interested in investigating the interaction between the LCD brightness and environmental lighting, and how it may affect human perception of the visibility of the images.

Finally, while the TN type LCD shows most dramatic brightness/color change across viewing angles and thus most suitable for our purpose, other type of LCDs such as VA and IPS may also have their own variances, although to a much lesser extent. As future work, we would like to investigate the optical properties of other LCD technologies to determine whether similar solutions may exist.

CONCLUSION

We presented a pure software solution for providing concurrent dual views on common TN-based LCD screens without any hardware modification or augmentation. We described our complete solution in detail, including the rendering algorithm, measurement and calibration methods, as well as demonstrated potential applications. Our solution is shown to work robustly across a variety of LCDs, and can be easily reproduced or integrated into various applications. By providing a solution that is widely applicable to existing devices at no additional cost, we hope it will contribute to more mainstream adoptions of multi-view displays in the real world.

REFERENCES

1. Chen, C., Chang, M., Ke, M., Lin, C., Chen, Y. (2008) A novel high brightness parallax barrier stereoscopy technology using a reflective crown grating. *Microwave and Optical Technology Letters*, 50(6), 1610-1616.
2. Cruz-Neira, C., Sandin, D. J., DeFanti, T. A. (1993) Surround-screen projection-based virtual reality: the design and implementation of the CAVE. *SIGGRAPH*, 135-142.
3. Dodgson, N. A. (2002) Analysis of the viewing zone of multiview autostereoscopic displays. *Proc. SPIE*, (4660), 254.
4. Dodgson, N. A., Autostereoscopic 3D Displays, *Computer*, 38(8), 31-36.
5. Floyd, R. W., Steinberg, L. (1976) An adaptive algorithm for spatial grey scale. *Proc. Society of Information Display*, (17), 75–77.
6. Halsted, C. P. (1993) Brightness, Luminance, and Confusion, Information Displays. http://www.crompton.com/wa3dsp/light/lumin.html
7. Harrison, C., Hudson S. E. (2011) A new angle on cheap LCDs: making positive use of optical distortion. *UIST*, 537-539.
8. Lien, A. (1990) Extended Jones matrix representation for the twisted nematic liquid-crystal display at oblique incidence, *Applied Physics Letters*, 57(26), 2767-2769.
9. Matusik, W., Forlines, C., Pfister, H. (2008) Multiview user interfaces with an automultiscopic display. *AVI*, 363-366.
10. Moore, J. R., Dodgson, N. A., Travis, A. R. L., Lang, S. R. (1996) Time-multiplexed color autostereoscopic display. *Proc. SPIE*, (2653), 10-19.
11. Phillips, G. K. (2004) New digital anti-copy/scan and verification technologies. *Optical Security and Counterfeit Deterrence Techniques V*, 133-141.
12. Shoemaker, G. B. D., Inkpen, K. M. (2001) Single display privacyware: augmenting public displays with private information. *CHI*, 522-529.
13. Smith, R. T., Piekarski, W. (2008) Public and private workspaces on tabletop displays. *Proc. AUIC*, (76), 51-54.
14. Toyooka, T., Yoda, E., Yamanashi, T., Kobori, Y. (1999) Viewing angle performance of TN-LCD with hybrid aligned nematic film, *Displays*, 20(5), 221-229.
15. Travis, A. R. L. (1990) Autostereoscopic 3-D Display, *Applied Optics*, 29(29), 4341-4342.
16. Yu, F. H., Kwok, H. S. (1999) Comparison of extended Jones matrices for twisted nematic liquid-crystal displays at oblique angles of incidence, *Optics InfoBase*, 16(11), 2772-2780.
17. Q2`10 Quarterly Wide Viewing Angle Technology Shipment and Forecast Database, (2010) *DisplaySearch*.