MITSUBISHI ELECTRIC RESEARCH LABORATORIES
http://www.merl.com

Privacy-Enhanced Displays by Time-Masking Images

Yerazunis, William, and Carbone, Marco

TR2002-11   February 2002

Abstract

How to display an image with moderate security, in an unsecured area.

OzCHI 2001

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Copyright © Mitsubishi Electric Research Laboratories, Inc., 2002
201 Broadway, Cambridge, Massachusetts 02139

PRIVACY-ENHANCED DISPLAYS BY TIME-MASKING IMAGES

ABSTRACT

This paper describes a method for enhancing the privacy of computer displays in public and semipublic areas. By operating the display at a higher-than-usual frame rate and alternately displaying frames of an arbitrary private image and a computed mask image, unauthorized viewers perceive one image, while authorized viewers with appropriately keyed shutterglasses see an entirely different (and private) image. Although the technique can be defeated, it provides a measure of privacy against casual and opportunistic privacy penetrations.

KEYWORDS: data privacy, display technology, shutterglasses

0. INTRODUCTION

This paper describes a technique for enhancing the privacy of computer displays in public and semipublic areas. Our technique starts with a pair of desired images - a public image and a private image; our goal is to have unauthorized viewers of the system perceive the public image, while authorized users, wearing special eyewear, perceive the private image. We do this by computing a special mask image as some function of both the public and private images, and displaying single frames of this mask image interleaved with single frames of the private image. This technique requires the CRT display to operate at higher-than-usual rates, and that authorized users wear active eyewear.

1. PRIOR ART

Several prior systems have been published or marketed for privacy enhancement. One type of system is what might be termed a *static system* , where there are no moving parts or dynamic alterations to the video stream. An example of this is the security screen by Russel (1994), by the Designs and Controls Lab at UVA (1999), and sold by 3M, where a thin set of nearly microscopic louvers obscure the screen to any viewer not on the louver axis. As a side effect, the louvers enhance screen contrast by blocking side light. The system is moderately effective at preventing casual eavesdropping but fails against shoulder-surfing.

A more interesting static system is the polarization-based system produced by MMI (2001) as a modification to LCD screens used on laptops. In this system, the front polarizer of an LCD screen is made removable. In normal use, the polarizer is temporarily reinstalled. For “secure use”, the polarizer is removed from the LCD and the authorized user wears polarizing sunglasses. Since the information is polarization-encoded, users without polarizing sunglasses see only a very bright white screen. A common set of consumer-grade polarized sunglasses defeats the system and enables full resolution viewing of the display. Reflection of the display by a dielectric medium (eyeglasses, smooth plastic, etc.) provides weak polarization and partial display readability.

Ohtake and Aoki (1990) and McManus (1995) describe dynamic systems, using active eyewear on authorized viewers. Ohtake uses a system that interleaves three frames of intentionally misleading characters with one frame of private image, and uses shutterglasses (eyeglasses with electronically controlled high-speed shutters built in) to view the desired frame. Their system does not allow the separate control of the publicly viewed image, and does not conceal that a secret is on display.

McManus (1995) uses active deterrence of eavesdropping by interposing bright flashes of light between data frames. Authorized users wear shutterglasses which block the bright flashes;

unauthorized users see the flashes which are bright enough to obscure the data display. To provide good obscuration, the flashes must be much brighter than the data; approximately 20 dB. Note that a 20 dB is equivalent to a factor of 100 in brightness. In any case, at least 99% of the display power is now diverted to the obscuring flashes, which is inefficient and intractable for battery powered devices.

Related to the current work are works by Shoemaker and Inkpen (2001), and Needham and Koizumi (1998). These systems are collaborative rather than adversarial in that all users agree to wear glasses of various types (polarization, color-filter, shutterglasses, etc.). The users view a shared display on which multiple private images are displayed; each user sees the image tailored to the glasses they are wearing. So equipped, the users can work collaboratively, sharing data as needed. Cooperation is necessary in this system; if a user removes their glasses, they see all of the private images overlaid on the display.

John Carpenter's motion picture *They Live* (1988) exhibits a close but fictional representation of our system. Carpenter does not postulate how the “Hoffman lenses” operate in the film, nor can we substantiate his more extraordinary findings.

2. TECHNIQUE

Our technique for privacy enhancement is based on the principle of persistence of vision. The human eye integrates the incoming light for significant time, and cannot easily differentiate between light pulses at 60 Hz vs. 120 Hz. This allows us to mask an image by supplying a second image within the persistence of vision time; the eye combines the two images by summing. We call these “time-masked images”.

To show a time-masked image, we alternately display one frame of private data with one frame of mask. The mask frame is brighter where the private data is darker, and vice versa. Because of persistence of vision the casual onlooker sees the sum of the private data and the mask. We can manipulate the average value to provide an innocuous (or entertaining) public image, while the authorized user wearing appropriately synchronized shutterglasses sees only the private data.

2.1 Technique Details I: The Math

To introduce the basic concepts in time-masked images, consider the simplest version of time-masking where we are supplied with a grayscale private data image (the “supplied secret” or SS). We will display SS interleaved with some displayed mask (or DM) such that the average of SS and DM = gray. Assuming pixel values between 0.0 and 1.0:

$$(SS + DM) / 2 = 0.5$$
$$DM = 1.0 - SS$$

so the DM image is just the pixelwise brightness inverse of the SS image. An authorized viewer with shutterglasses will see SS, 0, SS, 0,... and so can perceive the private data image. Unauthorized viewers will see the display series SS, DM, SS, DM... which is perceived as a neutral gray.

This system is easily understood, theoretically clean, but fails when implemented. The problem is that most CRT displays are nonlinear in brightness; brightness = 1.0 is not perceived to be twice as bright as brightness = 0.5 . This nonlinearity is known as the characteristic “gamma” of the monitor.

Monitor gamma varies with manufacturer and technology. In order to get proper time-masking, the DM image calculation results must be corrected for the gamma of the display system in use. For pixel values scaled to the range 0.0 <= value <= 1.0 the relationship:

$$\text{out\_pixel\_value} = \text{in\_pixel\_value}^{(1/\gamma)} + 0.5$$

with γ of approximately 2.2 yields good results on the CRTs we have tested.

Now, we consider the case of both a supplied secret (SS) and the supplied public (SP) image. In this case, we must first rescale the images so that the combined SS and DM images are within the dynamic range of the monitor, as well as determine offsets to determine where in the dynamic range the images will be. Assume we are given α $(0 \le \alpha \le 1)$ for the dynamic range of the supplied secret (SS) image, and O_{ss}, the black-level offset of the SS image. Similarly, β $(0 \le \beta \le 1)$ is the dynamic range of the perceived public (PP) image, and O_{sp} is the black level offset of the PP image. Using a linear transformation to obtain the perceived secret (PS) and perceived public (PP) images:

$$PS = \alpha \cdot SS + O_{ss}$$
$$PP = \beta \cdot SS + O_{sp}$$

We can then compute the displayed secret image (DS) and the displayed mask (DM) images:

$$DS = \alpha \cdot SS + O_{ss}$$
$$DM = \beta \cdot SP - \alpha \cdot SS + O_{sp}$$

There are restrictions on the realizable values for α, β, O_{ss}, and O_{sp}. The monitor cannot show anything blacker than 0.0 nor brighter than 1.0 , so we can show:

$$\alpha + \beta \le 1$$
$$\alpha + O_{sp} \le 1$$

These inequalities show the tradeoff between the respective dynamic ranges of the perceived public (PP) and perceived secret (PS) images. A high dynamic range perceived public (PP) image forces a low dynamic range, dim perceived secret (SS) image, and vice versa.

With PP and PS allocated equal dynamic range ($\alpha = \beta = 0.5$, $O_{ss} = 0.0$, $O_{sp} = 0.5$) the public image appears perfectly acceptable, although slightly low in contrast and with an elevated black level. The PS image appears slightly dimmer than usual, but remains acceptable for viewing. In the case of 24-bit color displays, each of the red, green, and blue channels must be manipulated separately. Figures 1a, 1b, and 1c show example images for SS (supplied secret), SP (supplied public), and the resulting DM (displayed mask) images respectively.

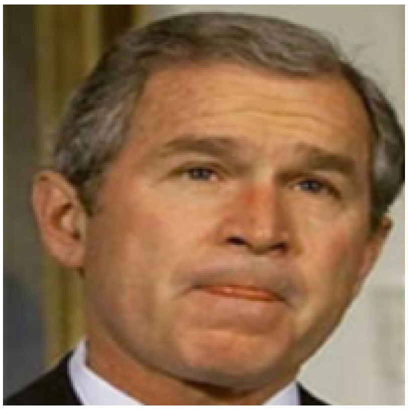

Fig. 1a: private image

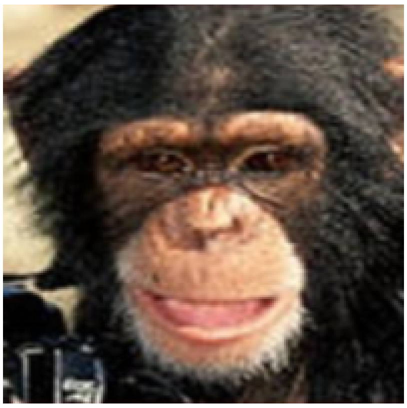

Fig. 1b: public image

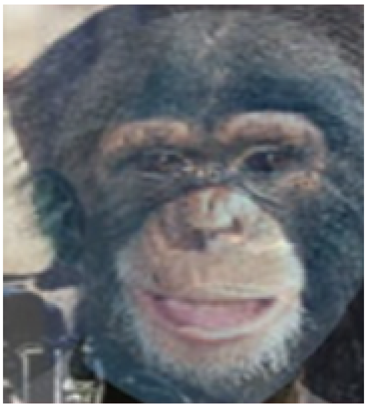

Fig. 1c: displayed mask

2.2 Technique Details II: The Software

The driving software for our time-masked image experimentation is written in ANSI C on top of SVGAlib and SDL. Each library allows the privileged-mode C code to catch the vertical retrace interrupt and flip the appropriate bitmapped image into the video display memory. The base system is Linux with X windows, version Red Hat 7.1, running on an Intel Pentium III at 700Mhz. We have achieved frame rates up to 150 Hz with SDL, and 105 Hz with SVGAlib. We have found it advantageous to precompute mask frames for some of the demonstrations, because we have not yet determined how to do the DM calculations in real-time on the cards. We hope to remedy this in further research.

2.3 Technique Details III: The Shutterglasses

Since many users report flickering with nematic liquid-crystal stereographic shutterglasses (typically limited to 60 Hz or less), we used ferroelectric liquid crystal (FLC) polarization rotator elements to construct much faster shutterglasses. The bistable FLC rotators can switch polarization rotation from +45 degrees to -45 degrees at approximately 100 KHz, when driven by a bipolar + -5 volt control line.

We have constructed both wired and infrared wireless shutterglasses. The wired glasses are driven by direct connection to the parallel port of a Linux PC, the wireless glasses carry a small battery and are keyed by a coded synchronization signal from an IR emitter also driven by the parallel port. Both sets of glasses operate on a commercially available CRT driven at 100 to 150 frames/sec.

3. RESULTS

We have constructed the system as described, and tested it in ad-hoc fashion against a significant number of laboratory members and visitors. The results are encouraging. Figure 2 shows the system in use; figures 3a and 3b show views through the operating shutterglasses.

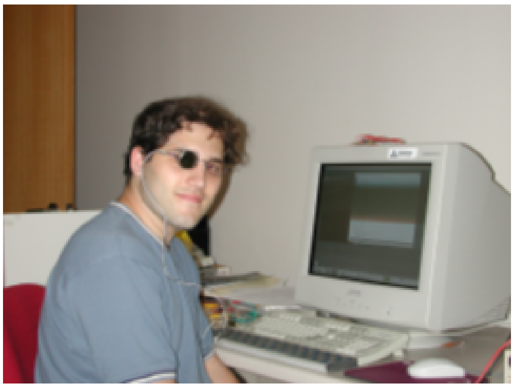

Fig. 2: system in use

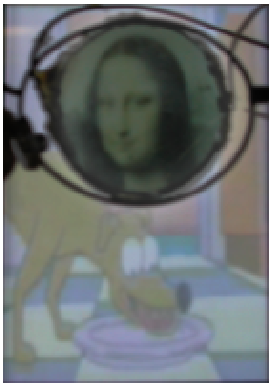

Fig. 3a

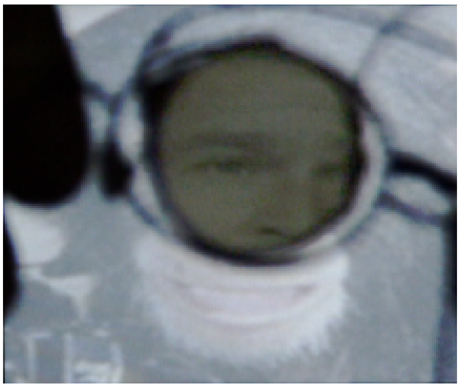

Fig. 3b

3.1 Static images

As described, the system provides nearly 100% effectiveness when both the private and public images are static images. Until they were informed otherwise, subjects did not note anything unusual in the displayed image beyond that it appeared slightly washed out and lower in contrast than usual.

Approximately forty subjects have been shown the display in our ad-hoc testing so far. No test subject was able to make an accurate statement as to the nature of the private image, even using the “comb attack” and “knock attack” described below, even with knowledge of how the system worked. The best that any test subject was able to state were on the level of “It's a picture of a person. I can't tell who”. This applied even when motivated by a reward of $20 to the first subject in a group to accurately identify the secret image. No subject had any trouble naming the famous person when allowed to use the authorized-user shutterglasses.

3.2 Dynamic images

Dynamically changing public images over a static (or near static) private images are as effective as the static public / static private case. Like the static / static case, no subject was able to note a difference between un-masked and masked sequences beyond the slightly lower contrast levels; no subject was able to make an accurate statement as to the nature of the private image in a dynamic public / static-private time-masked sequence, even with financial motivation, and none had any trouble correctly naming the famous person or item in our animation / static image tests when using the shutterglasses.

Dynamic private images are a more difficult problem. Beyond the issues of real-time computation of DM and DS images at 100+ Hz, strong action sequences in the private image often leave ghost images visible to unauthorized users. Although no subject was able to state what actions or objects were present in the action sequence, almost all subjects were able to note that "something was happening", "I can see something there", etc. We have tried several time-averaging methods but have not yet properly compensated for the short-term persistence curve of the human visual perception system.

3.3 Methods to Defeat Time-Masked Privacy-Enhanced Displays

We have found several methods to attack a time-masked sequence of images. The most obvious method is to use another pair of shutterglasses, with the shuttering speed externally controlled. This attack yields long periods of good viewing of the time-masked private image, but requires preparation and investment on the part of the snooper. This attack can be partially defeated if the time-masked sequence can be modulated either in speed, or some if a more random pattern of frame selection than strict alternation is used. A digital camera with a fast shutter can recover large fragments of the image which can be pieced back together after recording.

A simple attack is the "comb attack". By extending the fingers of one hand into the rough shape of a comb, and rapidly shaking this comb in front of an eye, the fingers act as unsynchronized shutters, causing inaccurate averaging of the DM and DS images. This simple attack yields some vague shape information about the display contents, but it is not enough to identify a person if the person is not yet known. This attack calls a lot of attention to the user, and so clearly is not suitable for covert snooping.

Another impromptu attack is the "knock attack". By striking the head with a fist, the eye is rapidly displaced across the scene, and the average of DS and DM images are no longer aligned on the retina. This attack works about as well as the "comb attack", but calls attention to the eavesdropper, and is physically uncomfortable. More stealthy, but almost as painful, is the "crushed ice attack", where the covert snooper chews on some very brittle material, such as crushed ice. Like the knock attack, the crushed ice attack causes rapid uncompensated displacements of the eye across the scene, yielding a low-quality, transient ghost image.

4. ARE FULLY SECURE DISPLAYS POSSIBLE?

With these limitations and attacks, we now consider if this technique for privacy-enhanced displays can be extended to secure a display to the extent that a well-funded, well-rehearsed attack fails to extract the private data. Unfortunately, the answer seems to be "not easily".

The controlling issue is that we must assume that the attacker is somehow able to capture (either on film or electronically) every pixel brightness in every frame. We can expect that the attackers will exploit any available mathematical or statistical relationship between pixels in an attempt to reassemble the private data (e.g. using a 2-D spin-glass algorithm to solve for most likely states, as in Mezard (1987)). To prevent these attacks, we must not display any pixel on the screen with any dependence on the private data.

We can accomplish this for one-bit images by operating the shutterglasses on a per-pixel rather than a per-frame basis. We display 100% white pixels, opening the shutterglasses to reveal white pixels and closing the glasses for black pixels. This method requires that the full video bandwidth data stream be transmitted to the shutterglasses over the data link for every frame. It would be simpler to issue head-mounted displays to each authorized user and regain the desk space taken up by the monitor.

To save bandwidth and avoid transmitting the interceptable video data stream, we can do the following for black-and-white images. We provide the display and the shutterglasses in advance with two identical copies of a difficult-to-predict sequence Δ. The shutterglasses open when the bit in Δ is a 1, and close when the bit is a 0. The display device uses its copy of Δ to know when to display correct data (1), and when to display inverted data (0) . Because the display is showing randomly-mixed correct and inverted data, an eavesdropper perceives the display as gray; authorized users see only the correct data. The well-funded attackers with a complete copy of the video stream still cannot reconstruct the correct image, because the bits in Δ are equally likely to be 0 or 1 and therefore each pixel displayed is equally likely to be correct or inverted.

Using FLC shutterglasses at 100 KHz, showing frames at 1/60th second, approximately 1600 black-versus-white pairs of pixels can be displayed: roughly equivalent to a 24x70 display, or four lines of 12 characters each, with complete security. By extension, 8-bit gray-scale requires 8 1-bit images scaled at 2^n brightness levels, and 24-bit (RGB color) images require 3 8-bit images resulting in only 70 pixels displayed. We conclude that fully secure shutterglass displays are limited in application.

5. FURTHER WORK

This work continues to be a work in progress, and we have identified several areas for further consideration. Display technology is an issue, especially in the area of hardware to generate the appropriate DM images in real time instead of precomputing. Determination of the causes of ghosting in dynamic private images is yet unsolved; further work in the proper dynamic filtering is needed.

The cooperative systems of Shoemaker and Inkpen (2001), and Needham and Kozumi (1998) could be extended to use time-masking with multiple secret images to give a standardized overview to all viewers, and controlling the shutterglass modulation to view detailed or personalized augmented data to users wearing shutterglasses; shutterglasses can have their modulation reprogrammed via the data link, to switch among a number of possible displays.

6. ACKNOWLEDGEMENTS

The authors would like to thank Drs. Darren Leigh and Paul Dietz for their great assistance in the design and construction of the privacy-enhanced display system.

7. REFERENCES

Shoemaker, B. D., Inkpen, K. M. (2001) Single Display Privacyware: Augmenting public displays with private information. *Proceedings of CHI, Conference on Human Factors in Computing Systems. Seattle, USA, April 2001*. (http://www.edgelab.sfu.ca/publications/chi2001_sdp.pdf)

MMI : (2001) http://www.man-machine.com/invisivw.htm , accessed 2001-July-24

Needham, B.H., Koizumi, D. H. (1998) *Method of Displaying Private Data to Colocated Users*, US Patent 5,963,371

Displays and Controls Lab (1998), Virginia Polytech. Inst. Blacksburg, VA. - *Proc. of the Human Factors and Ergonomics Society, 42nd Annual meeting,* vol 2, 1998, pp 1565-9 ISBN 0 94529 11 1

McManus, C. E. (1995), *System and Method for Data Security,* US Patent 5,629,984

Russel, A.R. (1994), *Privacy filter for a display device,* US Patent 5,528,319

Ohtake, S., Aoki, Y., (1990) Hokkaido Univ, Sapporo, Japan. A study of security for CRT display with liquid crystal shutter, *Transactions of the Institute of Electronics, Information, and Communication Engineers D-I*, Vol J73D-I No. 4, April 1990.

Carpenter, J. (1988), *They Live*, Universal Studios, ID423USDVD

Mezard D., Parisi M., Virasoro (1987) M. A., *Spin Glass Theory and Beyond*, World Scientific Publishing, ISBN 9971-50-115-5