Threat Analysis of Emanation-Based Attacks

Daniel Flower
University of Auckland, 2005

This paper gives an overview and threat analysis of eavesdropping attacks based on the emanations given off by computer hardware. Three distinct examples, as described in recently published research papers, are examined in order to give a better understanding of the types of possible emanation-based attacks, and the threat of attack. Specifically, the emanations of sound from keyboard clicks, light from LED indicators, and reflected light from monitors are looked at. It is suggested that although the threat is small for most organisations, these attacks are feasible and therefore must be considered by security teams.

1 Introduction

Much of computer security is concerned with making software secure, for example encrypting data to make it unreadable to the unauthorised, installing firewalls to keep out intruders, and of course the use of passwords for authentication. However, no matter how well a system has been secured, there will always be an inescapable weakness: data is transferred between the hardware and user in an unencrypted method (for example the user types on a keyboard, and the system displays output on the screen).

As (or if) computer systems become more secure from a technical standpoint, attacks that exploit the above fact are sure to become more popular. These attacks range from the very simple, for example watching somebody's fingers as they type in their password, to much more sophisticated attacks, such as reconstructing the contents of a computer screen by capturing the light reflected from it. It is this type of emanation-based attack which is the focus of this report.

1.1 Definition

The phrase "emanation-based attacks" as used in this report refers to those attacks whereby data is compromised by receiving the analogue outputs (e.g. light, sound etc) emanated from computer hardware. An emanation-based attack is a type of eavesdropping attack.

1.2 Characteristics

In [Lam04], Lampson identifies four key concerns of security: *secrecy*, *integrity*, *availability* and *accountability*. *Secrecy* is defined as "controlling who gets to read information", so emanation-based attacks compromise secrecy. *Integrity* deals with the use of data, and *availability* with access to data for authorised users, so an emanation-based attack does not compromise these areas. Finally, *accountability* is "knowing who has had access to information or resources". Because the data in an emanation-based attack is being accessed by an authorised user, from the system's point of view, there would normally be nothing unusual happening during the attack, so identifying and tracing the attackers is impossible through any auditing done by the system.

This is a dangerous aspect of emanation-based attacks, as it can bypass even the best-implemented security system. Figure 1 shows how an attacker bypasses an example of an implemented security system.

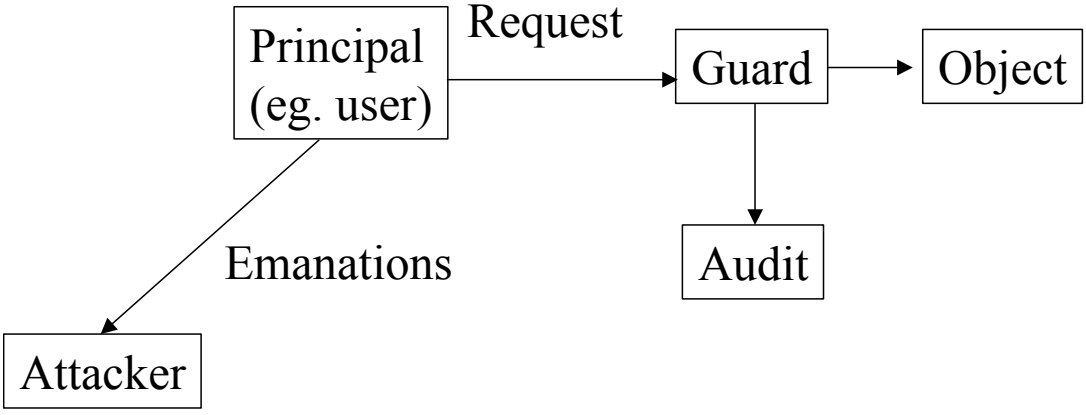

Figure 1: an attacker exploits the weakness of any security system: the emanations of a computer. In this case, a security system using an access control model, which involves a guard checking users before allowing access to an object and auditing all activity, is rendered irrelevant by an emanation-based attack.

Any attack that circumvents a security system appears to be very dangerous, yet little is heard about this type of attack compared to others, implying that it is not a common method. The obvious difficulty with this method is of course the fact that the attacker needs to physically be near an authorised user, whereas some other types of attacks (e.g. hacking into a computer system) can take place half a world away from the victim.

The aim of this report then is to better understand the threat posed by this type of attack. To do this, three different types of attack, as described in three separate research papers, will be evaluated: keyboard clicks to determine what is being typed, the blinking of LED indicators to determine what data is being processed by a device such as a modem, and reflected light to determine what is being displayed on screen.

The key in understanding the threats of each attack is to understand the cost of the attack (for example the cost of equipment required), who would be able to carry the attack out (for example, is the attack easy or could only a specialist implement it?), and whether easier attacks exist. All these points will be looked at, and it is hoped that an evaluation of these papers will give a better understanding of the threats posed by emanation-based attacks in general[1].

2 Case studies

This section studies three types of emanation-based attacks as described in three published papers. In each case an overview of the paper is given, along with the authors' findings and a critique of the paper, followed by a threat analysis.

2.1 Keyboard attacks

The paper being studied in this section is entitled “Keyboard Acoustic Emanations” by Dmitri Asonov and Rakesh Agrawal ([AA04]).

2.1.1 Overview

The main idea behind this paper is that although the difference in sound that computer keys make when pressed is inaudible to most humans, there *is* a difference and so a computer can be taught to recognise the differences and hence “listen” to what is being written on a computer by recording the sound of the key clicks.

[1] Another type of emanation-based attack, called a TEMPEST attack, whereby data is reconstituted by analysing the electromagnetic emanations of devices (particularly VDUs), is not being analysed in this report due to their expense and difficulty (see [NSA02]), and hence lower threat.

To achieve this, a neural network was trained with the input being certain extracted features of the sound wave created by the action of pressing a key, and the output being the key that was thought to be pressed.

Various tests were carried out, with attacks being taken from different distances (from around 0.5 meters up to 15 meters), different numbers of used keys (from two keys up to 30 keys), training on one person and testing on others, and training on one keyboard and testing on others of the same brand.

2.1.2 Authors' findings

Through experimentation, the authors found that different keys were distinguishable from one another because a key press hits a metal plate in the keyboard, and because different keys hit different parts of the plate, different sounds are produced. The authors found that the neural network could distinguish between 30 keys from a distance of 0.5 meters, where training and testing were done on the same keyboard, with an accuracy of around 79%. Around 88% of the time the correct key was one of the network's top three estimates.

The authors claimed that moving from 0.5 meters to 15 meters made no difference to the accuracy, which is an important result, as it would allow an attacker to maintain a fairly sizeable distance from the victim during an attack. They also found that one person could train the neural network, and it could then be used to mount an attack on another with a different typing style, which again is very important, as this method would be (nearly) useless if the trainer of the network could only attack themselves. A limitation they found is that training on one keyboard and mounting an attack on another of the same type gave very poor results (only around 52% of the time was the correct key in the network's top four guesses), and so it can be safely assumed that the attack would not work at all when trained on one keyboard and tested on another brand.

The method also worked nearly as well on notebook keyboards, telephones and ATMs.

2.1.3 Evaluation

A very important finding was that the accuracy for this attack did not change when the recording device moved from 0.5 meters away from the keyboard to 15 meters away. The authors made this claim based on a network that could distinguish between only two keys, however in the paper this distance was not tested when the full keyboard was being used. Distinguishing between two keys should be a lot easier for a neural network than distinguishing between 30, and so some doubt exists as to where the moving 15 meters away really does not hurt the accuracy.

When testing on different people, the authors had the people typing with a single finger versus multiple fingers, and with light force versus strong force, so any mention of typing speed was conspicuous in its absence. The authors stated that the noises that a single click makes last around 100 ms. A very fast type in can type up to 15 characters a second, and that this speed the sound waves created by the clicks would be overlapping each other (see figure 2).

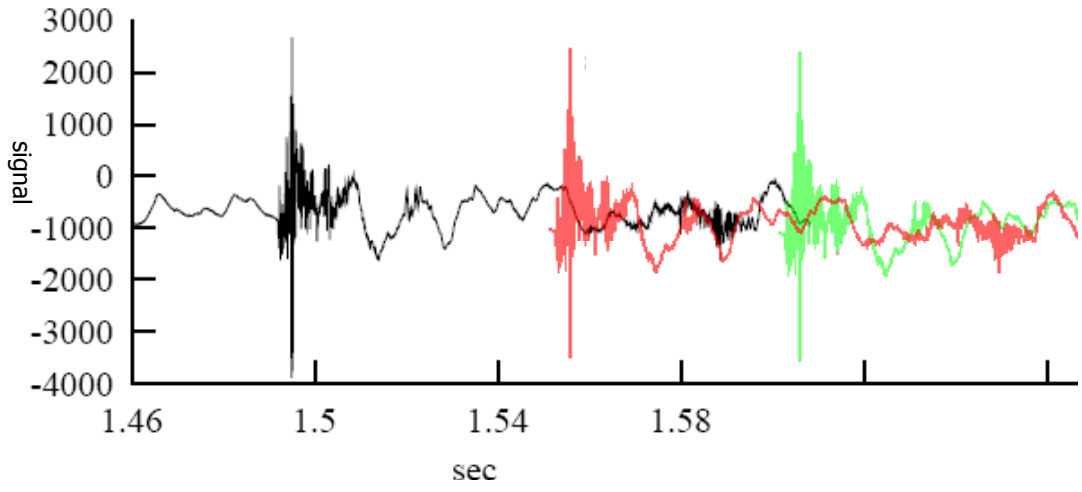

Figure 2: three keyboard-clicks overlapping each other. This is an alteration to a graph given by the authors, which showed the signal from a single key click. In this figure, two more clicks have been added with an assumed typing speed of 15 characters per second. You can see that the sound waves overlap each other, and some sound waves would cancel each other out while others would be amplified, and it is unclear whether the described attack could deal with this situation.

Although maintaining a typing speed of 15 characters per second is difficult, it is possible that this type of speed is attained when people type in passwords

and other frequently used pieces of data, which may be the type of data that this attack would be used for. No mention of this issue was raised in the paper, so it is unclear whether fast typing would thwart this type of attack.

2.1.4 Threat analysis

This attack is not expensive (a cheap, parabolic microphone and a computer is the only equipment required), and many people would be capable of carrying it out. Someone with experience in extracting features from sound waves and someone with a fairly basic knowledge of machine learning would be all that was required, so probably many postgraduate students could perform this attack. However, there are two crucial limitations to this attack: the requirement of physical proximity and the non-transferability of a trained neural network to another keyboard.

As stated above, the fact that the attacker - or at least the attacker's recording device - must be physically close to the victim greatly reduces the potency of this attack. The main reason for this, is that if an attacker is able to get near their victim, then an easier method would be to just discreetly watch them type on the keyboard, or perhaps use a small camera to record what is being written. This would restrict the usefulness of this kind of attack to a situation where the victim was typing in a public place (such as an Internet cafe) but was being careful enough to hide what they were typing from any onlookers.

The requirement that the attack take place on the same keyboard as it was trained obviously requires that the attacker have access to the keyboard before the attack takes place (again, this would be possible at a place such as an Internet cafe). While some information-gain can be had by carrying out the attack on only the same model of keyboard as was trained on, and hence increasing the possibility of attack to anybody whose keyboard model is known, the accuracy rate is quite low and more low-tech methods (such as peeking over someone's shoulder) may be more effective.

However, the fact that this method worked on ATM keypads does make this a dangerous attack. It is feasible that attackers could discreetly train their network on an ATM, and then sit around five meters away from it with their

microphone pointed at the keypad, waiting for people to come and reveal their PIN numbers.

Therefore, those organisations that have computer systems that hold sensitive information situated in public places should be the most concerned about this attack, and should consider investing in keypads that cannot be compromised in this way.

2.2 LED attacks

The paper being studied in this section is entitled “Information leakage from optical emanations” by Joe Loughry and David Umphress ([LU02]).

2.2.1 Overview

This paper examines the information leaked from LED indicators on devices such as modems. This is based on the premise that as data passes through such devices, LED indicators can flash on and off in a way that is closely correlated to the data passing through them, even to the point where plain text can be read from them.

The authors classified devices as: class one, where the LED lights indicate the status of the device, such as a "power on" light; class two, where the lights indicate activity; and class three, where the lights' status is closely related to the data being processed. Because LED lights have such a fast response time, class three devices are in danger of revealing information even when processing data at relatively high speed.

The authors divided their research into identifying the amount and kind of class three devices, and the feasibility and accuracy of an attack on such a device.

2.2.2 Authors' findings

In their survey of devices, 14 out of 39 devices (36%) were classified as class three. All but one of the dial-up modems tested were class three, but none of the network adapters were. The network adapters were safe because, the authors stated, the manufacturers stretched the time that an LED would flash on for to a minimum of a few milliseconds. This allows humans to see the lights more clearly, but in effect the speed of the lights compared to the device

becomes very slow, and so no data can be read. This idea would improve any class three device to be a class 2 device. In addition, on two routers, front panel LEDs were class two, but the LEDs on the back were class three. All surveyed storage devices and other network devices were safe.

The authors found that on class three devices, data could be read from a long distance, for example plain text from a modem could be read from up to 10 meters away. Error-free readings were taken at speeds of 56 kb per second, and the authors claim that it "ought to" work at speeds of up to 10 Mb per second, however this was not tested.

The authors claim anecdotally that many offices have computer systems facing glass windows with LED indicators visible from the street. This would potentially allow an attacker to eavesdrop on data in a business in a very discreet manner. The authors summed up the problem very well with the following sentences: “One of the advantages of LED displays is that they can be read from across a room. The disadvantage may be that they can be read from across the street.”

2.2.3 Evaluation

It would have been revealing if the authors had better illustrated their results, for example giving error rates of plain text from different distances for devices running at different speeds. Instead, we are given the correlation between the actual data and of the detected data at different distances for speeds up to 19,200 kb per second. The correlation for all speeds was 1.0 at five meters, 0.8 at 10 meters, and just 0.2 at 20 meters.

The claim that data could be read from "across the street" seems feasible, but unproven. Admittedly, the authors did state that this was a "proof of concept experiment", and they certainly seemed to prove their concept.

2.2.4 Threat analysis

This threat is both more expensive and more difficult to carry out than the previous attack due to the more difficult job of capturing and processing light. An engineer experienced in capturing and processing light waves would be required, although no machine learning would be required.

As with the previous attack, the requirement of physical proximity is a limitation of this attack. Based on their findings, it would seem that class three devices that can be seen from 10 meters away are at risk, however because this was a proof of concept experiment, it is quite possible that there are other techniques that could be used to extend this distance.

One fact that better techniques will not improve though, is that a line of sight is required between the attacker and the device. Therefore, this type of attack should only be the concern of those organisations that have computer systems visible from public areas. These organisations should consider classifying their devices as class one, two, or three devices, and should cover the LED indicators or move any class three devices to a more concealed location.

LED indicators exist so that humans can monitor device activity, however there is no point in having them flash as quickly as the data running through them as this is too fast for humans to capture. Therefore, the complete eradication of this threat should be borne by the manufacturers of devices, by ensuring that no devices are class three devices.

2.3 Image reconstruction attacks

The paper being studied in this section is entitled “Optical time-domain eavesdropping risks of CRT displays” by Markus Kuhn ([Kuh02]).

2.3.1 Overview

This paper describes a proof of concept experiment whereby an image on a cathode-ray tube is reconstituted by using a photosensor, even when a direct line of sight is not possible.

The gives a brief overview of line of sight attacks, for example it is shown that theoretically an amateur-level telescope could be used to read even a small font from 60 meters away, at an angle of 60° from the screen. The true purpose of the paper though is reconstructing images from light reflections, and the paper goes into some detail about relevant issues such as which sensors to use, types of display modes on CRT monitors, and the effect that different lighting has.

The results look at an actual experiment that the author performed, and the theoretical possibilities of this type of attack.

2.3.2 Author's findings

The actual experiment involved placing a CRT monitor one metre from a white wall, with a photosensor 0.5 meters behind the monitor (i.e. 1.5 meters from the wall, with no direct line of sight to the monitor), and when the background lighting was dark, the image was successfully reconstructed. Figure 3 shows the picture that was shown on the CRT monitor, and the reconstructed image.

While very impressive, the proximity of the sensor to the monitor was too close to simulate a real attack. However, the author had shown that such an attack is feasible.

The author then went on to use certain physical theories to estimate the limit of this attack, and stated that a monitor placed two meters from a wall may be susceptible to this attack with a photosensor placed up to 50 meters away, provided that it was very dark. It was shown that background light was the biggest obstacle in this kind of attack, and so in daylight the same scenario would not work with the photosensor even one metre from the wall.

2.3.3 Evaluation

The result achieved was very impressive, and it is certainly feasible that techniques could be developed to improve the attack further. However, presuming that this attack would work from 50 meters requires a great leap of faith, as from that kind of distance it seems that much of the light would bounce away into the surrounding environment.

In the experiment, white text on a black background was used, but the reason for this was not given. Due to the fact that generally text is black on a white background, there is some suspicion that a black background was used because a white background would give poor results (perhaps due to there being too much light, for example).

Further work needs to be done before this method becomes a credible threat.

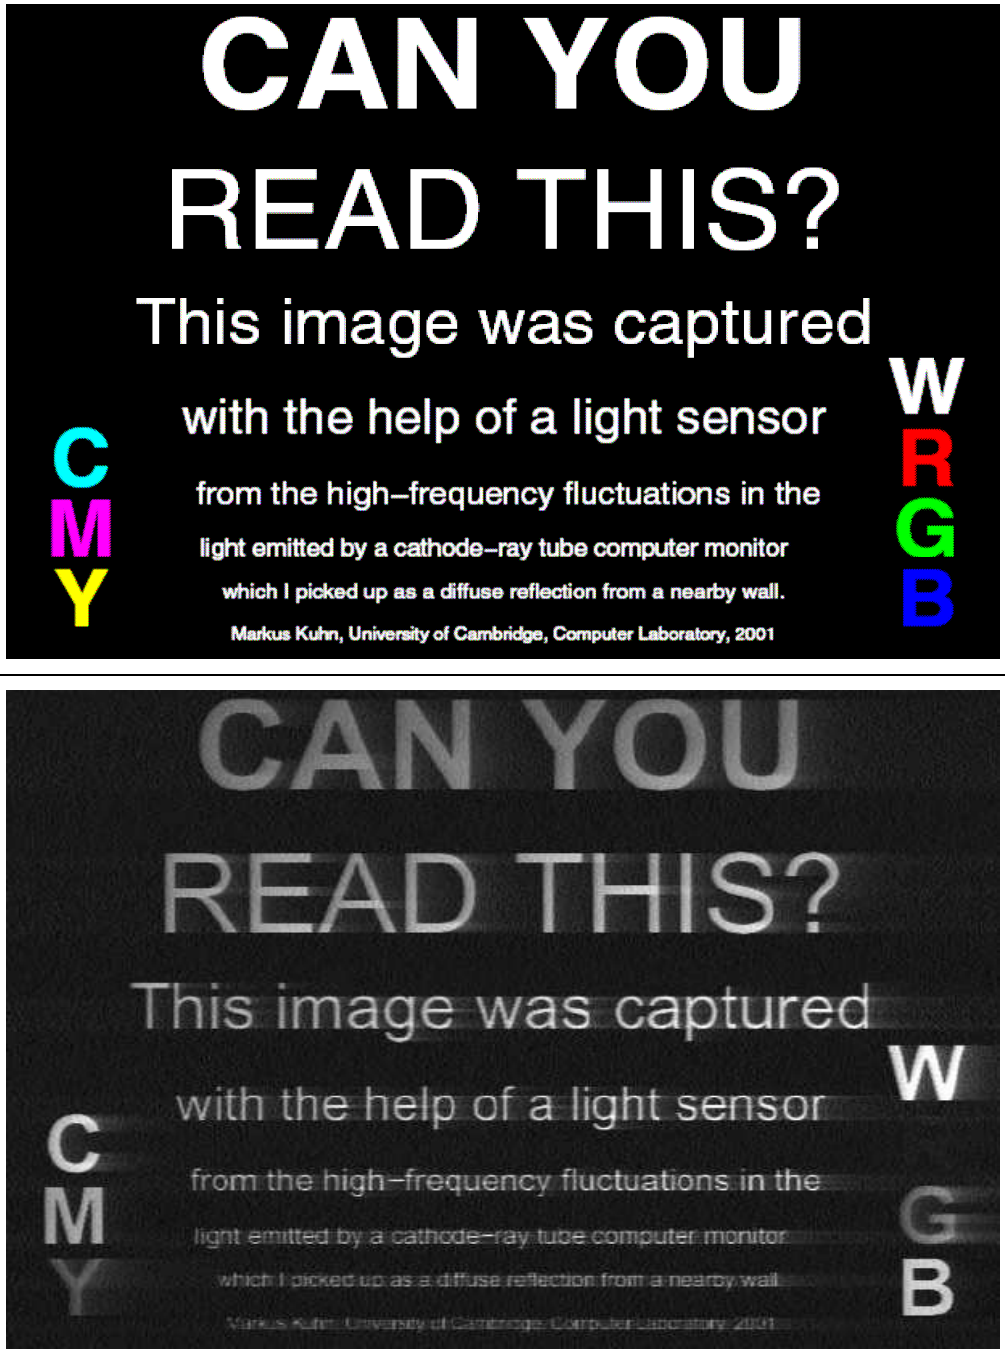

Figure 3: Above: a test image displayed on a CRT monitor, and below: the image recreated from the light reflections off a wall close by.

2.3.4 Threat analysis

This was by far the most complicated of the three attacks examined in this report. It required specialist equipment, and an expert understanding in the physics of light. In addition, the limited experiment conducted, while an

impressive start, is still a long way from being a useful attack method. These facts combined make the threat of such an attack very low at this time, and this method would probably be only of concern to only the most paranoid organisations.

Furthermore, this threat may already be becoming outdated, as CRT monitors are being replaced with LCD monitors, which as stated in the report are much less susceptible to this type of attack due to their physical nature.

3 Discussion

This section will discuss the strengths and weaknesses of emanation-based attacks in general, and concludes with an overall assessment of the threat of these attacks.

3.1 Strengths

The biggest strength of this type of attack was illustrated back in figure 1. This showed how these attacks completely bypass the implemented security system, and even with complete and detailed auditing the attacker cannot be traced.

Furthermore, because emanation-based attacks are relatively unknown compared to more conventional types of attacks, there is generally little security in place to prevent them.

3.2 Weaknesses

The requirement of physical proximity, and in most cases access to specialised equipment and/or detailed knowledge of some discipline, make these attacks difficult to carry out.

3.3 Evaluation

It is said that a chain is as strong as its weakest link, and so for an attacker, it makes sense for them to attack the weakest part of a system's security. It is likely then, that there are easier attacks to be made than emanation-based attacks. If an attacker is in the proximity of their victim, easier methods are probably possible, for example looking over someone's shoulder from a discreet position, rather than reconstituting the contents of their monitor. In

many other cases, social engineering (i.e. the manipulation of an authorised user to perform some action or reveal privileged information[2]) is probably easier than launching an emanation-based attack.

Having said that though, there probably are certain situations where the threat is higher, for example the ATM example given earlier. Therefore, it is important that security personnel are knowledgeable about such attacks, and that they consider them when designing and analysing systems.

4 Conclusion

A brief overview of emanation-based attacks has been given, along with three unique examples in order to give a general understanding of what can be achieved with these attacks, and how likely they are to occur.

While it was concluded that the threat is not high from these attacks, mainly due to the fact that easier methods do exist, their very existence should be a cause of concern for those implementing security in computer systems.

5 References

[AA04] Asonov, Dmitri, and Agrawal, Rakesh. 2004. Keyboard Acoustic Emanations, *Proc. of IEEE Symposium on Security and Privacy*: 3-11

[CNN05] A convicted hacker debunks some myths. 2005. *CNN.com*. http://www.liu.edu/cwis/cwp/library/workshop/citchi.htm (accessed October 14, 2005)

[Kuh02] Kuhn, Markus G., 2002. Optical time-domain eavesdropping risks of CRT displays, *Proc. of the 2002 IEEE Symposium on Security and Privacy*: 3-18

[Lam04], Lampson, Butler, 2004. Computer Security in the Real World, *Computer* 37 (6): 37-46

[LU02] Loughry, Joe, and Umphress, David A. 2002. Information leakage from optical emanations, *ACM Trans. Information Systems Security* 5 (3): 262-289

[NSA02] National Security Agency. 1992. NACSIM 5000 TEMPEST Fundamentals. *National Security Agency, Fort George G. Meade, Md*. http://cryptome.org/nacsim-5000.htm

[2] See [CNN05] for a short but interesting discussion.