Eavesdropping attacks on computer displays

Markus G. Kuhn

http://www.cl.cam.ac.uk/~mgk25/

Computer Laboratory, University of Cambridge
15 JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom

Abstract

Electromagnetic information leakage from computer displays was first demonstrated to the general public by van Eck in 1985. Nearby eavesdroppers can pick up compromising emanations from computer hardware with directional antennas and wideband receivers. The basic phenomenon is easily demonstrated with modified TV sets. However, to separate practically readable text shown on modern high-resolution displays from interfering background noise, special digital wideband signal-processing systems are needed. Thanks to Moore's law, access to such hardware is no longer restricted to well-funded spies, but has come within the reach of amateurs. The problem is not restricted to cathode-ray tubes; some contemporary flat-panel systems are at least as vulnerable.

Keywords: Compromising emanations, emission security, Tempest, eavesdropping, video displays

1 Introduction

Computers emit energy in various forms, mostly as unintended side effects of normal operation. Where these emissions take the form of radio waves, they may become noticeable when they cause interference in nearby radio receivers. Some of these emissions carry information about processed data. Under good conditions, a sophisticated and well-equipped eavesdropper can intercept and analyze such signals to steal information at a distance.

The problem has been known since the early days of electronic computing. Some military organizations, concerned about these *compromising emanations*, started research on *emission security* around 1960. They established a set of test standards and management procedures for especially shielded equipment, known under the codename *Tempest* [1].

The concept of compromising emanations was brought to the attention of the broader public by a 1985 paper [2] and a 5-minute TV demonstration on the BBC programme "Tomorrow's World", in which van Eck demonstrated that the screen content of a video display unit could be reconstructed at a distance, using low-cost equipment, namely a TV set whose sync-pulse generators were replaced by manually controlled oscillators.

There are few published cases so far where compromising emanations were exploited with practical benefits for the attacker. One commonly-quoted example is the World War I practice of eavesdropping on single-wire enemy field telephones that use the ground for the return current. Opponents connected earth rods to portable valve amplifiers, in order to make the voltage drops generated across the terrain audible as intercepted speech signals [3]. Another example is former MI5 scientist Peter Wright's recollection of an eavesdropping attack on a diplomatic cipher machine, which leaked plaintext telex signals as weak high-frequency pulses on cables coming out of the French embassy in London [4, pp. 109–112].

The field of compromising emanations received renewed interest in the late 1990s through work on extracting cryptographic secrets from smartcard processors by analyzing their supply-current fluctuations [5]. This inspired a range of academic eavesdropping demonstrations, including one in which secret modular-exponentiation parameters were extracted from a cryptographic SSL accelerator module installed inside a server, using a radio antenna at a distance of 5 m [6].

Ongoing technological changes affect both the nature of accessible emissions, as well as the technology available to exploit them. In this brief survey, we will focus on video signals, which remain – thanks to their highly repetitive and redundant nature – one of the easiest to demonstrate emission-security risks. Pixel frequencies and video bandwidths have increased by an order of magnitude since van Eck's demonstrations with modified TV sets, and analog signal transmission is in the process of being replaced by Gbit/s digital video interfaces. Various flat-panel display (FPD) technologies are well on their way to replacing the cathode-ray tube (CRT). All these developments make it necessary to reevaluate the emission-security risks identified decades ago. Likewise, Moore's law has substantially increased the capabilities of low-budget attackers. Specialized wideband signal-processing equipment that, a decade or two ago, filled a large instrument rack and cost a fortune, can now be replaced with an FPGA DSP development board for a few hundred euros.

2 Eavesdropping on cathode-ray tube displays

Even though the days of CRT displays seem to be numbered, the analog video cables originally designed for them remain widely used. For this reason, and because of their conceptually simpler nature, the type of compromising emanations generated from CRT systems continues to be of interest.

2.1 Radio-frequency emissions

To display text or graphics, a microprocessor writes the required pixel brightness values into a frame-buffer memory. A graphics-controller chip periodically reads the entire content of this memory periodically, between 60 and 85 times each second. It converts one pixel colour after another into a voltage triplet (0–0.7 V for each of red, green, and blue) and passes these via a video cable to the monitor electronics. There, the video signal is amplified by a factor of about 100 and applied to the control grids inside the cathode-ray tube, which modulate the three electron beams that illuminate the screen phosphors.

Many parts of such a system can act as unintentional transmission antennas: data lines from the frame buffer to the video controller, the video cable to the monitor, the video amplifiers inside the monitor, and finally the control grids. Each time an electric current is switched on or off in one of these conductors, a brief electromagnetic impulse is emitted. If we mathematically split an infinitely short impulse into a sum of sine waves of different frequencies (Fourier analysis), each frequency shares an equal part of the pulse's energy. Real impulses from computer cables have a finite duration. As a result, their energy is spread across the radio spectrum mostly below the frequency that is the inverse of the impulse duration. For example, an impulse that lasts one billionth of a second (1 ns) will contain most of its energy at radio frequencies between zero and one billion hertz (0–1 GHz).

Any sufficiently sensitive radio receiver tuned to an arbitrary frequency within this range will be able to detect such an impulse. The question is, whether the impulse will stand out sufficiently against other noise signals. What background noise the eavesdropper has to cope with depends, among other factors, on the bandwidth of the receiver: the width of the slice of radio spectrum that it cuts out of the antenna signal and presents to the user. The wider the bandwidth of the radio receiver, that is the larger the difference between the lowest and highest frequency to which it is sensitive at a given tuning frequency, the shorter will be each output impulse.

Figure 1 illustrates the need for high-bandwidth receivers to capture high-bitrate digital signals. It shows the output of a radio receiver, at five different bandwidth settings, in response to a single nanosecond-long impulse at its antenna input, much like the impulse emitted by a voltage transition on a wire. The left-hand figure shows the output of the receiver's intermediate-frequency (IF) filter, before any demodulation takes place. The right-hand figure shows the output of an AM demodulator, which has rectified and smoothed the signal further.

The output pulse width is approximately the inverse of the receiver's IF filter bandwidth. Therefore, the bandwidth of a radio receiver used to eavesdrop on emissions from digital sources should be about as large as the clock rates involved, if individual transition impulses are to be distinguished. For example, in order to distinguish two impulses that are only one-tenth of a microsecond apart (10 MHz clock frequency), we need to use a receiver with a bandwidth of at least 10 MHz. Computer video signals have transmission rates of typically 20 to 150 million pixels per second, therefore a receiver bandwidth of at least 20 or 50 MHz is desirable. Otherwise the impulses from individual pixels could not be resolved in the output. Such bandwidths are much larger than the 6 MHz of a television receiver. On the other hand, the bandwidth cannot be chosen arbitrarily large either. The wider it is, the more difficult it becomes to find a quiet frequency range where no active radio transmitter overlaps the reception band.

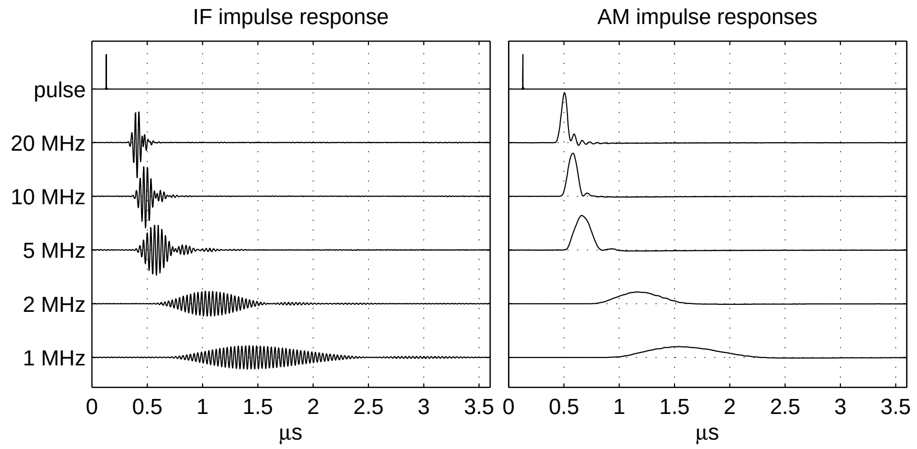

Figure 1: Receiver output resulting from a single nanosecond-short impulse at the antenna input.

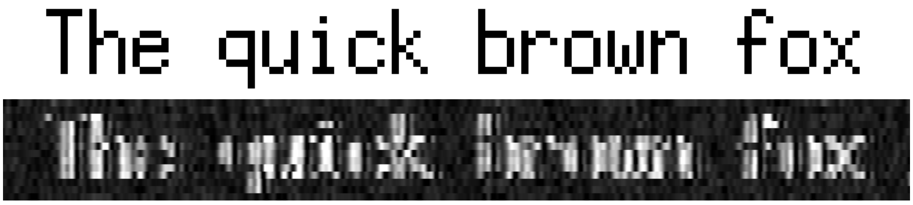

Figure 2: Text displayed on a cathode-ray tube (top) and signal seen by eavesdropper (bottom).

The upper part of Figure 2 shows some text as it is displayed on a cathode-ray tube monitor with analog VGA-cable input. Below, we see the same text reconstructed from the output of an AM receiver tuned to 480 MHz (50 MHz bandwidth). The signal from the output of the AM radio was digitized, using a computer-controlled storage oscilloscope, and converted into a raster image by converting the signal voltage into brightness values, pixel by pixel and line by line, in the same fashion in which the eavesdropped display does [7, p. 49]. The text eavesdropped this way remains readable, but it is distorted compared to the original. In particular, we no longer see the difference between a foreground and a background colour. Instead, we see a bright impulse wherever there is a change between foreground and background in a horizontal direction, that is wherever the electron beam in the eavesdropped monitor is being switched on or off. Text remains easy to recognize in this form for two reasons. Firstly, the shape of most letters carries a lot of redundancy; secondly, most of the information in text glyphs is carried by vertical high-contrast edges (even when anti-aliased). Continuous-tone photographs, with their much smoother edges, are far more difficult to recognize.

Why does an eavesdropper need to use a radio receiver here at all? The radio is not used for its originally intended purpose, namely to isolate the narrow frequency spectrum of a carrier wave from a transmitter and to undo the modulation by which the transmitter packs information onto this signal, such as frequency or amplitude modulation (FM, AM). With compromising video emanations, no such modulation takes place. Instead, a video-signal eavesdropper uses a radio receiver merely as a tuneable filter that suppresses all signals except those within an easily adjustable frequency band. Some frequency ranges of the electromagnetic spectrum are teeming with powerful transmitters that are orders of magnitude stronger than the compromising emanations. This is especially the case in frequencies reserved for mobile phones, pagers and local radio broadcast stations. Other areas of the spectrum are much quieter most of the time, including the frequency ranges reserved for civilian and military aviation signals and distant television stations. Since impulse signals have their energy spread over a wide frequency range, an eavesdropper can tune a radio receiver to any quiet part of the spectrum where such impulses stand out particularly well from the background noise. Most compromising video signals can be received in the 200–800 MHz range.

As Figure 1 shows, impulses appear as compact oscillations at the output of the intermediate-frequency filter of a radio receiver. In order to convert them into a visible video signal, they merely have to be rectified (converted into a positive voltage), which is essentially what an AM demodulator in a radio does. If such a signal is to be shown in real-time on a video monitor, it has to be brought to the correct voltage and augmented with synchronization pulses, which – unlike with TV signals – are not present in a usable form in the receiver output. An easy way to obtain sync pulses is to regenerate them independently, at exactly the same frequency at which they are generated in the eavesdropped system. For a stable image, these frequencies need to match to at least seven to eight decimal digits. If the horizontal deflection frequency were off by only one millionth, then after the electron beam in the original and the eavesdropper's monitor have both drawn the about one million pixels of a full video frame, they would already be an entire pixel column apart. At 60 frames per second, the eavesdropper's image would therefore roll 60 pixels per second to the left or right due to this frequency mismatch.

2.2 Optical emissions

Apart from radio waves, computer displays also emit light. This is, after all, what they were designed for. Optical telescopes can obviously be used to read screen surfaces from a distance. Surprisingly though, optical spying is not only a risk if the screen surface is directly visible to the eavesdropper. More technical effort makes it also possible to read the text shown on a cathode-ray tube monitor from merely a diffuse reflection of the light that it generates. An eavesdropper could point a telescope at the wall behind a computer screen, or even at the face of a user sitting in front of it. If the illumination conditions are such that enough light received this way comes from the monitor, a very fast light sensor (photomultiplier) can be connected to the telescope in order to convert the received high-frequency light flicker into a current, which will essentially be the video signal. It will be blurred by the afterglow of the screen phosphors and further distorted by added noise from other light sources, but there are processing techniques available to undo most of these distortions [7, 8].

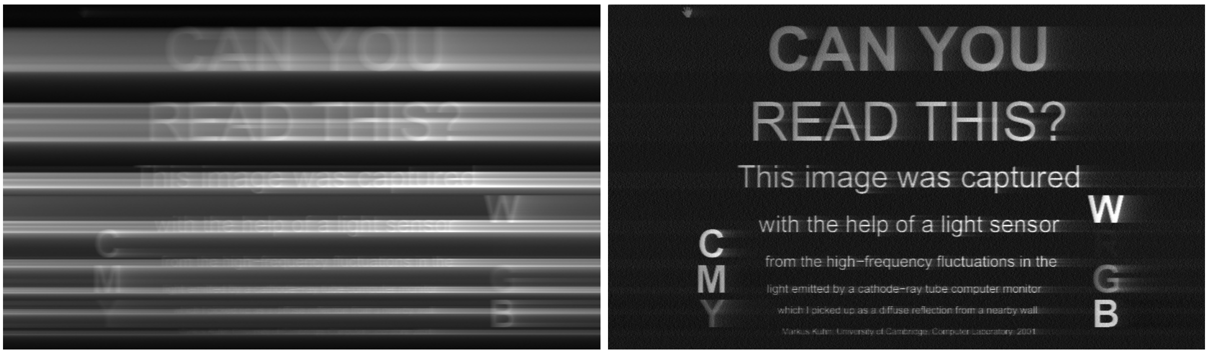

CAN YOU
READ THIS?
This image was captured
with the help of a light sensor
from the high−frequency fluctuations in the
light emitted by a cathode−ray tube computer monitor
which I picked up as a diffuse reflection from a nearby wall
Markus Kuhn, University of Cambridge, Computer Laboratory, 2001

Figure 3: Raw (left) and processed (right) photomultiplier signal from diffusely reflected CRT light.

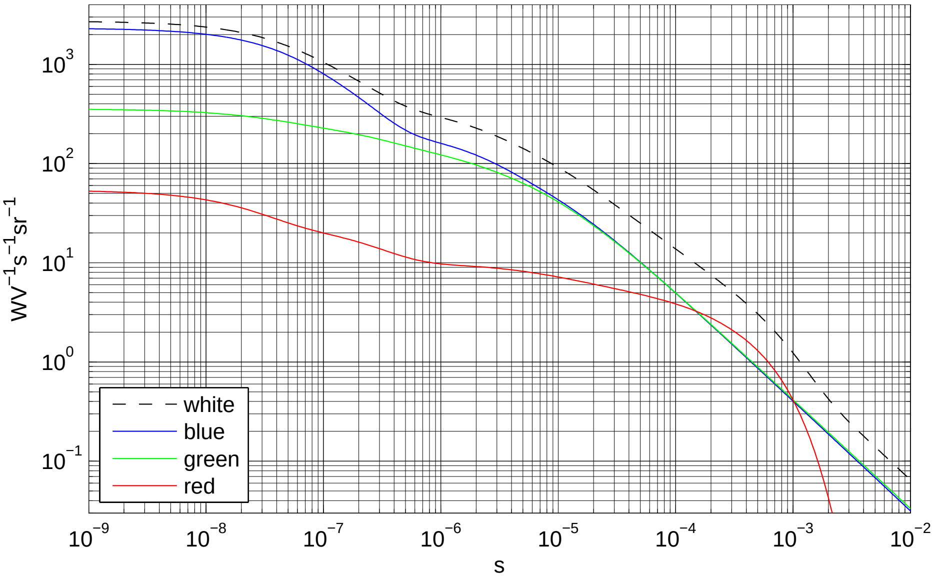

Figure 4: Phosphor afterglow (impulse response) plotted on logarithmic time and intensity scales [8].

Why does this work? Each time an electron beam is switched on to draw a bright area of the displayed image, the light coming from the CRT intensifies. If it is switched off again to draw a dark area, the light output is reduced correspondingly. Figure 3 shows how the signal from such a photosensor looks if it is mixed with sync signals and fed into a video monitor. It also shows how the blurring can be reversed using a specially designed filter. This amplifies the higher frequencies in the video signal in the opposite way to that in which the afterglow of the screen phosphors attenuated them.

The process by which the afterglow of the screen phosphors blurs the signal can mathematically be described as a convolution operation. By Fourier transforming the received signal into the frequency domain, the convolution operation becomes equivalent to a complex-number multiplication for each frequency component, which can be easily undone by division. The only obstacle to be overcome is that a noise-free mathematical model of the afterglow of the phosphors is needed, to reduce the risk of dividing some frequencies by a number near zero, thereby amplifying them far too much [7, 8]. Since the reconstruction amplifies the higher frequencies in the signal, which can substantially increase the noise, it is beneficial to start with a very low-noise recording of the blurred video signal. In practice, this can be obtained by averaging many recorded video frames first.

After such reconstruction, it is possible to read text from an optically eavesdropped video signal, even at small font sizes. In the monitor used in this demonstration, this worked particularly well with the signal in the blue channel. This is because, as Figure 4 shows, the blue phosphor gives off a very bright and short (100 ns) flash of light right after the electron beam strikes it, whereas the red phosphor stores the energy it has received and gives it off much more slowly through red photons.

Colour images can be reconstructed by repeating the entire process with a red, green and blue colour filter in front of the photo sensor.

3 Eavesdropping on flat-panel displays

The optical eavesdropping technique described in Section 2.2 does not work with flat-panel displays, where all pixels in a row are addressed simultaneously by the electronics. It is therefore not possible there to get a sequential video signal, one pixel after another, optically.

However, this does not mean that flat-panel displays suffer no eavesdropping risk. On the contrary, some that we examined were considerably easier to eavesdrop on by radio than many cathode-ray tubes, and also gave a much clearer signal. These emissions come not from the display module itself, but from the digital Gbit/s serial cables that are increasingly used to connect a flat-panel display to the video controller.

Unlike CRTs, which ultimately need an analog voltage to control the electron beam, flat-panel displays (TFT, plasma, etc.) are inherently digital devices. They need to store at least an entire row of pixels in memory to generate the necessary row and column signals that drive individual pixels. Therefore, any digital-to-analog conversion of the video signal would have to be undone in the flat-panel display. An entirely digital signal path saves both these steps plus the necessary pixel-clock recovery, thereby avoiding the associated loss of image quality. In some embedded devices, where space permits, the video controller and display module are directly connected via an 18- or 24-bit parallel data bus (for 6- or 8-bit red/green/blue pixel values). In notebooks, however, it is mechanically inconvenient to feed that many wires through the hinges that attach the display to the body. Therefore, the pixel values are fed through parallel-to-serial converter chips and transported through a thin video cable, on only three or four high-speed twisted-pair serial links.

Figure 5 shows a compromising video signal from a notebook (Toshiba 440CDX) received at a distance of 10 m through two intermediate offices (three intermediate plasterboard walls), even without the use of a directional antenna, in an office building with well over hundred other computers [9]. The calibration bar shows the equivalent voltage at the antenna input of the receiver. The reception technique is the same as was described for CRTs in Section 2.1.

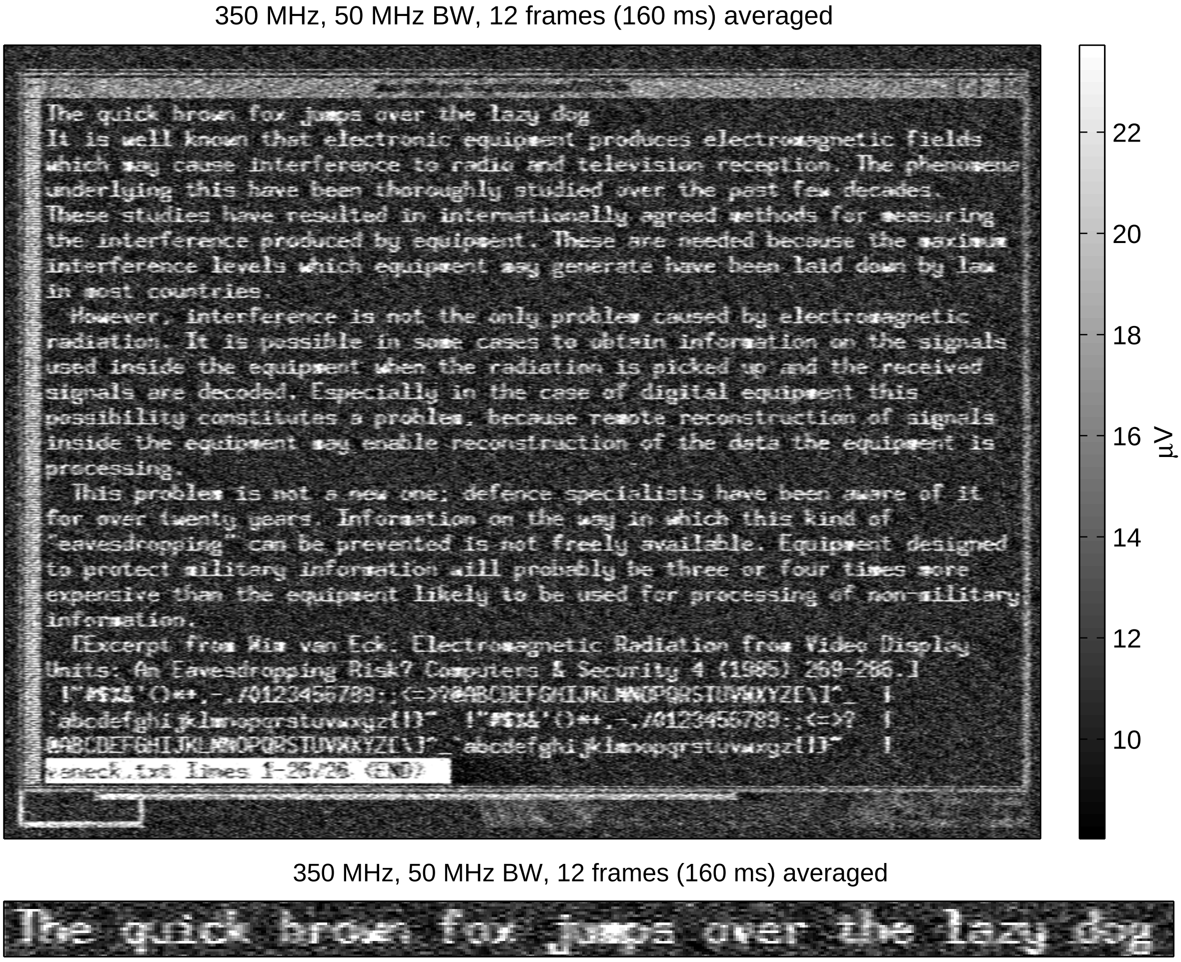
350 MHz, 50 MHz BW, 12 frames (160 ms) averaged
22
20
18
16
14
12
10
µV
350 MHz, 50 MHz BW, 12 frames (160 ms) averaged

```
The quick brown fox jumps over the lazy dog
It is well known that electronic equipment produces electromagnetic fields
which may cause interference to radio and television reception. The phenomena
underlying this have been thoroughly studied over the past few decades.
These studies have resulted in internationally agreed methods for measuring
the interference produced by equipment. These are needed because the maximum
interference levels which equipment may generate have been laid down by law
in most countries.
  However, interference is not the only problem caused by electromagnetic
radiation. It is possible in some cases to obtain information on the signals
used inside the equipment when the radiation is picked up and the received
signals are decoded. Especially in the case of digital equipment this
possibility constitutes a problem, because remote reconstruction of signals
inside the equipment may enable reconstruction of the data the equipment is
processing.
  This problem is not a new one; defence specialists have been aware of it
for over twenty years. Information on the way in which this kind of
"eavesdropping" can be prevented is not freely available. Equipment designed
to protect military information will probably be three or four times more
expensive than the equipment likely to be used for processing of non-military
information.
  [Excerpt from Wim van Eck: Electromagnetic Radiation from Video Display
Units: An Eavesdropping Risk? Computers & Security 4 (1985) 269-286.]
 !"#$%&'()*+,-./0123456789:;<=>?@ABCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_  !
`abcdefghijklmnopqrstuvwxyz{|}~  !"#$%&'()*+,-./0123456789:;<=>?  !
@ABCDEFGHIJKLMNOPQRSTUVWXYZ[\]^_`abcdefghijklmnopqrstuvwxyz{|}~   !
vaneck.txt lines 1-26/26 (END)
```

```
The quick brown fox jumps over the lazy dog
```

Figure 5: Text signal received from a notebook through two intermediate offices (3 plasterboard walls).

The bottom of Figure 5 shows a magnification of the first sentence. In comparison with Figure 2, what we see here is no longer a switching impulse that occurs only between horizontally neighbouring pixels of opposite colour. Instead, each screen colour is represented by a characteristic luminosity. A crude explanation would be that the brightness of each pixel seen by the eavesdropper is proportional to the number of 0-to-1 and 1-to-0 transitions in the data word that represents the pixel's colour on the serial line. This would be the case, if the receiver's bandwidth were at least as large as the bit rate (i.e., > 1 GHz). In practice, due to the presence of strong other transmitters, eavesdroppers will only be able to afford bandwidths in the region of the pixel rate, i.e. a few tens of megahertz. So we have to think of each screen colour being represented by a digital waveform, and consider the Fourier analysis (spectral composition) of the bit patterns of both the foreground and background colour. The spectra of these two waveforms will have frequency regions where their signal amplitude is particularly different. These are the frequencies where a radio receiver will see a particularly good contrast between a foreground and a background colour. As a result, the best tuning frequency for the eavesdropping receiver may dependent on the foreground and background colour combination used.

When eavesdropping on a flat-panel display connected to a PC via a digital (DVI) cable, two very different combinations of horizontal and vertical deflection frequencies may lead to a stable image. This is because most contemporary flat-panel displays contain not only a display module, but also a circuit board that implements a scan-rate conversion circuit (plus an on-screen menu for adjusting parameters). The display modules are only designed to be driven with a single combination of frequencies (usually a frame rate of 60 Hz for TFTs). However, historically, computers have used many different video frequencies. In the interest of usability and compatibility, manufacturers now add a frame buffer to monitors, to ensure that the display can be fed with most video modes and resolutions. The eavesdropper can then get a first signal, in the video mode set on the PC, from the DVI cable, and a second signal, in the video mode supported by the display module, from the internal interconnect between the scan-rate converter and the actual display module.

4 Tools of the eavesdropper

How can an attacker target one particular display, in the midst of emanations from dozens of other nearby computers and other sources of noise and interference?

Video signals can be separated very well from background noise, because they are highly redundant. The video controller repeats all data 60–85 times per second. Displayed information usually remains unchanged for many seconds or minutes. An eavesdropper can use periodic averaging of the received signal, at exactly its vertical repetition frequency, to average out noise and other unrelated signals. The frequency accuracy needed to obtain a stable image (more than eight decimal digits) means that even if there are two identical computers located right next to each other, operating even in the same video mode, their signals can still be separated. The frequency differences of the crystals in the pixel clock oscillators are large enough to cause the image of one computer to roll-over rapidly while the deflection frequencies remain adjusted to the other computer. With periodic averaging, a rolling image will quickly disappear as a uniform background blur and can then be subtracted.

Independent of which of the three previously-described types of compromising video emanations an eavesdropper wants to exploit (analog RF, CRT optical, digital RF), the problems to be overcome are mostly the same:

- the desired signal has to be separated sufficiently from background noise;
- a large parameter space (crude tuning frequency, high-accuracy horizontal and vertical deflection frequency) has to be scanned to find the parameters of one particular emitter;
- the phase relationship of the deflection frequencies must be maintained once these parameters have been identified.

Thanks to advances in semiconductor technology, the hardware needed to build a real-time signal-processing module that can perform the necessary steps is now coming within reach of even hobby eavesdroppers with a modest budget. Particularly well-suited as a platform for such designs are high-speed DSP FPGA development boards, as they are available today from several manufacturers. In order to gain a better understanding of these new risks, at our lab we are currently building the COVISP-1 (COmpromising VIdeo Signal Processor) system, based on an Altera Stratix II DSP development kit. Using field-programmable gate array (FPGA) technology, such boards can be turned into highly-specialized signal-processing systems that can apply several complex filtering and averaging steps in real-time to a video signal, at sampling rates of more than 100 MHz. This makes them suitable for processing computer video signals, where comparable pixel-clock frequencies are used.

For RF attacks, we need an antenna for the lower end of the UHF spectrum (about 200–800 MHz). For initial assessments and short-range demonstrations, the sort of logarithmic-periodic dipole antennas commonly used in electromagnetic-compatibility laboratories is well suited, because they are compact, and a single one covers the entire frequency range of interest. But such measurement antennas with broad frequency range have little directional gain. Once an attacker knows from an initial survey the best eavesdropping frequency for a particular target, the wideband antenna can be replaced with a directional Yagi antenna (similar to UHF TV roof antennas) for exactly that frequency. For even higher gain, an entire array of such antennas could be deployed [10]. Particularly useful for this purpose are adjustable Yagi antennas that can be reconfigured for different UHF frequencies.

Also needed is a sensitive tuner that covers not only the frequency range of interest, but that also offers a large enough bandwidth (at least 20–50 MHz), preferably adjustable. Commercially available dedicated Tempest laboratory measurement receivers (e.g., Rohde & Schwarz FSET, Dynamic Sciences) are certainly suitable, but very heavy and expensive, and provide far more functionality than what is needed merely for video eavesdropping. Far more affordable UHF wideband tuners are available commercially for radio-spectrum surveillance applications, and are increasingly becoming available for software-defined radio applications. Such receivers typically have an intermediate-frequency output in the 30–70 MHz range, which can be directly connected to the 120 MHz analog-to-digital converters found on signal-processing boards designed for digital-radio applications.

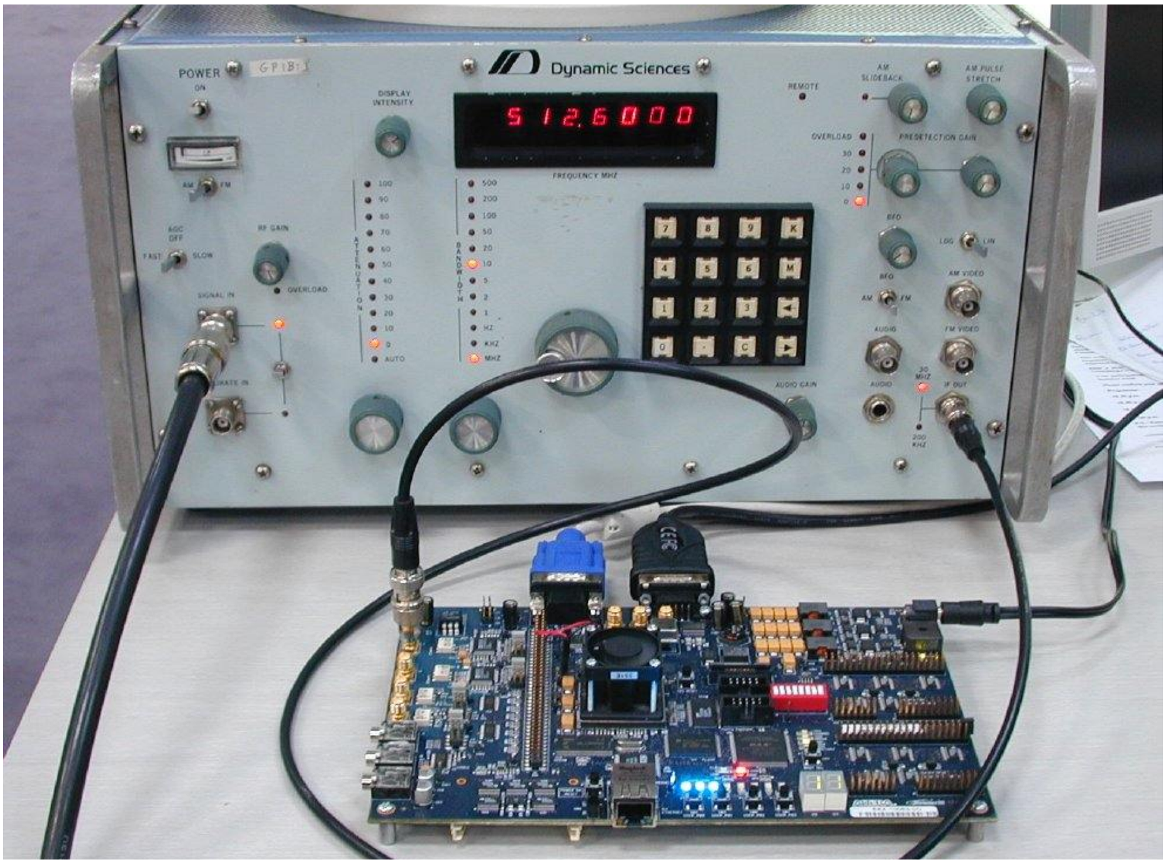

Figure 6: COVISP-1 compromising video signal processor (front), connected to wideband tuner (back).

The remaining functionality can all be implemented in software on either the FPGA chip or on a connected notebook computer. Our COVISP-1 prototype is, at present, able to generate sync pulses for all standard VESA video modes, as well as for arbitrary user-defined ones, and outputs these over the board's VGA connector. These frequencies can be conveniently adjusted, with a resolution of nine decimal digits. A hardware lookup table maps the digitized 12-bit IF input values to 8-bit VGA output voltages, and can be configured to perform rectification as well as offset and gain adjustment. The board also implements an input-level meter and a blanking function that switches the VGA video signal off near the sync pulses, as not doing so can interfere with the correct operation of the connected display electronic.

Our COVISP-1 prototype uses so far only ≈ 6% of the available FPGA gates, leaving room for many improvements. Future versions may add a frame buffer for 16-bit periodic averaging, a special correlation circuit for the rapid automatic recognition of received video modes, a digital input filter that can be configured to suppress narrow interference sources within the input band, a digital output filter for real-time deblurring of optical signals or for echo-cancellation in RF signals, an on-screen user interface, and perhaps even a web server for convenient remote operation and easy archiving of eavesdropped snapshots. Ultimately, such a signal-processing board can be turned, with suitable software, into an eavesdropping device that is nearly as easy to use as a wireless network: the user just switches it on and gets, within seconds, a ranking list of all available video signals in the vicinity, with thumbnail screenshots and the ability to click on any for real-time monitoring and recording. Combined with a suitable tuner daughter board and antenna kit, the result could become a video eavesdropping system whose bill of materials is not more than the price of a good notebook computer, but with functions that were previously only within the reach of well-funded intelligence agencies.

5 Defences

It may only be a question of time before powerful and increasingly easy to use eavesdropping tools for compromising emanations, as outlined in the previous section, can be easily downloaded from the Internet. What protective countermeasures are available for computers that regularly display really sensitive data?

5.1 Jamming

Jamming devices that deliberately increase the environmental background noise would be one possible countermeasure. Although such devices can be built to comply with radio regulations, jamming is rarely practised today for two reasons. Firstly, the jamming signal would have to be very carefully chosen and synchronized with the signal to be covered, otherwise, it can be filtered out by periodic averaging or other signal-processing techniques. Secondly, jamming devices may draw the attention of eavesdroppers to the location of equipment that its owner considers to be a particularly worthwhile target.

5.2 Shielding

Sensitive government systems today employ expensive metallic shielding of individual devices, cables, rooms, and sometimes entire buildings [11] to protect against compromising emanations. The relevant test standards and their rationales are, unfortunately, still secret and conforming products often remain export controlled. Existing civilian electromagnetic-emission standards were not designed to control compromising emanations and are not at all suited for this purpose [10].

5.3 Zoning

It is not easy to shield an individual device such that any chance of a successful close-range eavesdropping attack (from an adjacent room) is eliminated. This is likely to require substantial modifications to a product and an in-depth analysis of all emanations. Many governments have, therefore, adopted a more economic approach known as "zoning". Both devices and their locations are classified to indicate at what distance an eavesdropper may have access. A typical example of such a classification scheme would be:

- Zone 0: eavesdropper could be within 1–20 m;
- Zone 1: eavesdropper could be within 20–100 m;
- Zone 2: eavesdropper could be within 100 m to 1 km;
- Zone 3: eavesdropper cannot get closer than a kilometre.

These distance measures are assuming that there is only free space between the eavesdropper and the target. If there are, for example, from all sides two walls between the eavesdropper and the target, and these attenuate the frequencies of interest by 20 dB, then the distances to the eavesdropper become correspondingly shorter (e.g., 10–100 m for Zone 2, instead of 100–1000 m). For a given site, a protected space can often be identified that is not accessible to an eavesdropper (e.g., the area inside a guarded perimeter fence). If this is the case, a zoning measurement can be performed to classify each room of the site according to the RF attention between this room and the best receiver spot outside that perimeter.

Likewise, a device can be certified to be suitable for secure use in a Zone 2 room, if no useful compromising emanations can be found 100 m (line of sight) from the device, at the lowest plausible background noise level. In practice, almost any device that fulfils the requirements of civilian radio-frequency-interference standards will be secure in Zone 3 and can with, modest modifications, be made secure for Zone 2 usage. Some products will require no modifications for Zone 2, but a customer will not be able to identify these without suitable measurements having been performed in a controlled environment (shielded chamber). Many such products can be modified economically even for secure use in Zone 1. Such modifications may involve completing and improving the contacts of an already existing metal enclosure, the improvement of grounding arrangements, the replacement of wired connections with fibre-optical ones, and the installation of additional low-pass filters behind interface and power-supply connectors. Devices that can even be safely used in Zone 0, very close to the eavesdropper, may have to undergo quite substantial modifications and tests, the cost of which can easily exceed the price of the original product several times. The zones described above are not necessarily exactly the zones used today in any particular country, as the exact definitions of these zones remain secret. They are, however, indicative of the zoning scheme used in many NATO countries.

5.4 Soft Tempest

In systems where compromising emanations from digital video cables are the primary concern, these could of course be scrambled, for instance using similar technology as is already being introduced for copy protection and digital restrictions management. Finally, we have also proposed a number of more experimental software countermeasures [7]. For analog displays, for example, filtered fonts can be used where the contrast of vertical glyph edges is softened deliberately to reduce the amplitude of emitted impulses. Another example is the randomization of less significant bits in the frame buffer of purely digital display systems. The latter effectively generates a jamming signal that is repeated and updated at exactly the same rate as the covered signal, and is therefore not easily removed by averaging.

6 Conclusion

Compromising video signals and low-cost proof-of-concept demonstrations of them have been known for several decades. However, the signal-processing equipment needed for practical and effective exploitation was not easily available. This is about to change with the availability of powerful FPGA-based digital signal processing systems. These can be configured by software into hardware implementations of complex algorithms for the real-time processing of signals with at least 20–50 MHz bandwidth.

In addition, thanks to emerging markets for software-defined radios and ultra-wideband (UWB) communication systems, components designed for processing wideband signals and weak radio impulses are now finding their way into low-cost consumer products. As a result, it is becoming increasingly easy to build, with hobbyist resources, eavesdropping equipment for compromising emanations whose capabilities were a decade ago only available to well-funded military and intelligence organizations.

At the same time, the trend towards using twisted-pair Gbit/s serial links for periodic uncompressed video signals ensures that modern devices are even more likely to emit intelligible data. The old topic of compromising video emanations may well be facing a renaissance.

References

[1] D. Russel, G.T. Gangemi: Computer security basics. Chapter 10: TEMPEST. O'Reilly, 1991, ISBN 0-937175-71-4.

[2] Wim van Eck: Electromagnetic radiation from video display units: An eavesdropping risk? Computers & Security, Vol. 4, pp. 269–286, 1985.

[3] Arthur O. Bauer: Some aspects of military line communications. The History of Military Communications, Proceedings of the Fifth Annual Colloquium, Centre for the History of Defence Electronics, Bournemouth University, 24 September 1999.

[4] Peter Wright: Spycatcher – The candid autobiography of a senior intelligence officer. William Heinemann Australia, 1987, ISBN 0-85561-098-0.

[5] P. Kocher, J. Jaffe, B. Jun: Differential power analysis. In Michael Wiener (Ed.), Advances in Cryptology – CRYPTO'99, LNCS 1666, Springer, pp. 388–397, 1999.

[6] Suresh Chari, Josyula R. Rao, Pankaj Rohatgi: Template attacks. 4th International Workshop on Cryptographic Hardware and Embedded Systems, LNCS 2523, Springer, 2002, pp. 13–28.

[7] Markus G. Kuhn: Compromising emanations: eavesdropping risks of computer displays. Technical Report UCAM-CL-TR-577, University of Cambridge, Computer Laboratory, December 2003.

[8] Markus G. Kuhn: Optical time-domain eavesdropping risks of CRT displays. Proceedings 2002 IEEE Symposium on Security and Privacy, Berkeley, California, 12–15 May 2002, IEEE Computer Society, pp. 3–18, ISBN 0-7695-1543-6.

[9] Markus G. Kuhn: Electromagnetic eavesdropping risks of flat-panel displays. 4th Workshop on Privacy Enhancing Technologies, 23–25 May 2004, Toronto, LNCS 3424, Springer, pp. 88–105.

[10] Markus G. Kuhn: Security limits for compromising emanations. J.R. Rao and B. Sunar (Eds.): Workshop on Cryptographic Hardware and Embedded Systems (CHES 2005), Springer, LNCS 3659, pp. 265–279.

[11] Electromagnetic pulse (EMP) and Tempest protection for facilities. Engineer Pamphlet EP 1110-3-2, 469 pages, U.S. Army Corps of Engineers, Publications Depot, Hyattsville, December 31, 1990.