Soft Tempest: Hidden Data Transmission Using Electromagnetic Emanations

Markus G. Kuhn[⋆] and Ross J. Anderson

University of Cambridge, Computer Laboratory, New Museums Site,
Pembroke Street, Cambridge CB2 3QG, United Kingdom

Abstract. It is well known that eavesdroppers can reconstruct video screen content from radio frequency emanations. We discuss techniques that enable the software on a computer to control the electromagnetic radiation it emanates. This can be used for both attack and defence. To attack a system, malicious code can encode stolen information in the machine's Tempest emanations and optimise them for some combination of reception range, receiver cost and covertness. To defend a system, a trusted screen driver can display sensitive information using fonts which minimise the energy of RF emanations. There is also an interesting potential application to software copyright protection.

1 Introduction

It has been known to military organizations since at least the early 1960s that computers generate electromagnetic radiation that not only interferes with radio reception, but that also leaks information about the data being processed. Known as *compromising emanations* or *Tempest* radiation, a code word for a classified U.S. government research programme, the electromagnetic broadcast of data has been a significant concern in sensitive computer applications.

In his book 'Spycatcher' [7], former MI5 scientist Peter Wright recounts the origin of Tempest attacks on cipher machines. In 1960, Britain was negotiating to join the European Economic Community, and the Prime Minister was worried that French president De Gaulle would block Britain's entry. He therefore asked the intelligence community to determine the French negotiating position. They tried to break the French diplomatic cipher, but without success. However, Wright and his assistant Tony Sale noticed that the enciphered traffic carried a faint secondary signal, and constructed equipment to recover it. It turned out to be the plaintext, which somehow leaked through the cipher machine.

Sensitive government systems today employ expensive metallic shielding of individual devices, rooms and sometimes entire buildings [14]. Even inside shielded environments, the 'red/black' separation principle has to be followed: 'Red' equipment carrying confidential data (such as computer terminals) has to be isolated by filters and shields from 'black' equipment (such as radio modems) that handles or transmits unclassified data. Equipment with both 'red' and

[⋆] Supported by a European Commission Marie-Curie training grant

‘black’ connections, such as cipher machines and multilevel secure workstations, requires particularly thorough testing. The US standard NACSIM 5100A that specifies the test requirements for Tempest protected equipment, and its NATO equivalent AMSG 720B, are classified documents [6]. In Germany, even the names of the government standards on compromising radiation are kept secret.

So we can only speculate about the measurement technology required for Tempest tests, but descriptions in published patents [12, 13] suggest that the tools employed are orders of magnitude more sensitive than the spectrum analysers used in standard electromagnetic compatibility (EMC) and radio frequency interference (RFI) testing. Some tests involve long-term cross-correlation measurements between signals measured directly inside the target system and the noisy and distorted signals received from external sources including not just antennas but also power and ground lines, peripherals and network cables. Even microphones might be suitable sensors, especially to test equipment like line printers. By averaging correlation values over millions of samples, even very weak traces of the processed information can be identified in electric, electromagnetic, and even acoustic emanations.

When conducting attacks, similar averaging and correlation techniques can be used if the signal is periodic or if its structure is understood. Video display units output their frame-buffer content periodically to a monitor and are therefore a target, especially where the video signal is amplified to several hundred volts for cathode ray tubes. Special attack software that an attacker could implant in a system can also generate periodic or pseudo-random signals that are easy to detect. Knowledge of the fonts used with video displays and printers allows maximum-likelihood techniques to be used to get a better signal/noise ratio for whole characters than is possible for individual pixels.

Similar techniques can be applied when snooping on CPUs that execute known algorithms. Even if signals caused by single instructions are lost in the noise, correlation techniques can be used to spot the execution of a known pattern of instructions. Bovenlander reports identifying when a smartcard performs a DES encryption by monitoring its power consumption for a pattern repeated sixteen times [8]. Several attacks become possible if one can detect in the power consumption that the processor is about to write into EEPROM. For example, one can try a PIN, deduce that it was incorrect from the power consumption, and issue a reset before the PIN retry counter update and be completed. In this way, the PIN retry limit may be defeated.

The first public description of the Tempest threat appears to have been a 1983 report in Swedish [1], but the problem was brought to general attention by a 1985 paper [2] in which van Eck demonstrated that the screen content of a video display unit could be reconstructed at a distance using low cost home built equipment — a TV set whose sync pulse generators were replaced by manually controlled oscillators. His results were later confirmed by Möller, Bernstein and Kolberg, who also discuss various shielding techniques [5].

Smulders later showed that even shielded RS-232 cables can often be eavesdropped [4]. Connection cables form resonant circuits consisting of the induction of the cable and the capacitance between the device and ground; these are excited by the high-frequency components in the edges of the data signal, and the resulting short HF oscillations emit electromagnetic waves.

It has also been suggested that an eavesdropper equipped with fairly simple radio equipment could pick up both magnetic stripe and PIN data by standing near automatic teller machines, because card readers and keypads are typically connected to the CPU using serial links. A related risk is cross-talk between cables that run in parallel. For instance, the reconstruction of network data from telephone lines has been demonstrated where the phone cable only ran parallel to the network cable for two meters [15]. Yet another risk comes from 'active' attacks; an attacker who knows the resonant frequency of say a PC's keyboard cable can irradiate it with this frequency and then detect keypress codes in the retransmitted resonance signal thanks to the impedance changes they cause [16].

Considering the excitement that van Eck's findings created [3], and the enormous investment in shielding by the diplomatic and defence community, it is surprising that practically no further research about Tempest attack and defence has appeared in the research literature. However, an RF lab is expensive, while purely theoretical contributions are difficult due to the lack of published data about the emanations of modern hardware.

Commercial use of Tempest technology is also marginal. Attempts have been made by the UK and German governments to interest commercial firms in Tempest, in order to help maintain capabilities developed during the Cold War. This has been without success: Tempest-shielded PCs and workstations are many times more expensive then standard models, and sales are typically export controlled. So it is no surprise that shielded facilities and equipment are practically never used outside the diplomatic and defence communities.

This may change. In this paper, we describe a number of simple experiments that we have performed with a Tempest receiver and a cheap AM radio. This project started out of the curiosity of the authors and was not funded. We had no access to the expensive equipment that one would expect to find in a signals intelligence agency; even our elderly Tempest receiver is not much more sophisticated than a modified TV set. Our experiments thus show what kinds of attacks are practical in 1998 for a creative amateur eavesdropper. We have also developed some extremely low-cost protective measures.

2 Shortwave Audio Transmissions

If we want to write a computer virus to infiltrate a bank or certification authority, obtain key material and broadcast it to us over an improvised radio channel, then an important design criterion is the cost of the receiver. While intelligence services may already possess phased array antennas and software radios [18], such equipment is not yet generally available. The graduate student's Tempest spying kit is more likely to be just a worldband radio receiver connected to an audio cassette recorder, costing in total about US$100.

In order to get a computer VDU to produce audible tones on our radio, we have to design a screen that causes the VDU beam current to approximate a broadcast radio signal. If this latter has a carrier frequency f_c an audio tone with a frequency f_t it can be represented as

$$s(t) = A \cdot \sin(2\pi f_\mathrm{c} t) \cdot [1 - B \cdot \sin(2\pi f_\mathrm{t} t)]$$

The timing of a digital video display system is first of all characterised by the pixel clock frequency f_p, which is the reciprocal of the time in which the electron beam in the CRT travels from the center of one pixel to the center of its right neighbor. The pixel clock is an integer multiple of both the horizontal and vertical deflection frequencies, that is the rate $f_\mathrm{h} = f_\mathrm{p}/x_\mathrm{t}$ with which lines are drawn and the rate $f_\mathrm{v} = f_\mathrm{h}/y_\mathrm{t}$ with which complete frames are built on the screen. Here, x_t and y_t are the total width and height of the pixel field that we would get if the electron beam needed no time to jump back to the start of the line or frame. However the displayed image on the screen is only x_d pixels wide and y_d pixels high as the time allocated to the remaining $x_\mathrm{t}y_\mathrm{t} - x_\mathrm{d}y_\mathrm{d}$ virtual pixels is used to bring the electron beam back to the other side of the screen.

Attack software can read these parameters directly from the video controller chip, or find them in configuration files. For instance, on the authors' Linux workstation, a line of the form

```
ModeLine "1152x900"  95  1152 1152 1192 1472  900 900 931 939
```

in the X Window System server configuration file `/usr/lib/X11/XF86Config` indicates that the parameters $f_\mathrm{p} = 95$ MHz, $x_\mathrm{d} = 1152$, $y_\mathrm{d} = 900$, $x_\mathrm{t} = 1472$, and $y_\mathrm{t} = 939$ are used on this system, which leads to deflection frequencies of $f_\mathrm{h} = 64.5$ kHz and $f_\mathrm{v} = 68.7$ Hz.

If we define $t = 0$ to be the time when the beam is in the center of the upper left corner pixel ($x = 0$, $y = 0$), then the electron beam will be in the center of the pixel (x,y) at time

$$t = \frac{x}{f_\mathrm{p}} + \frac{y}{f_\mathrm{h}} + \frac{n}{f_\mathrm{v}},$$

for all $0 \leq x < x_\mathrm{d}$, $0 \leq y < y_\mathrm{d}$, and $n \in \mathbb{N}$. Using the above formula with the frame counter $n = 0$, we can now calculate a time t for every pixel (x,y) and set this pixel to an 8-bit greyscale value of $\lfloor 128 + s(t) \rfloor$ with amplitudes $A = 64$ and $B = 1$. See Fig. 1 for screen contents generated this way to broadcast an AM tone.

It is not necessary to fill the entire screen with the pattern, but the energy of the transmitted signal is proportional to the number of pixels that display it. Ideally, both f_c and f_t should be integer multiples of f_v to avoid phase discontinuities from one frame to the next.

We had no problems hearing a test melody broadcast by our PC, using a cheap handheld radio. This worked everywhere in our lab and in nearby rooms, while reception over longer distances was good so long as the receiver antenna was held close to power supply lines. As one might expect from the wavelengths

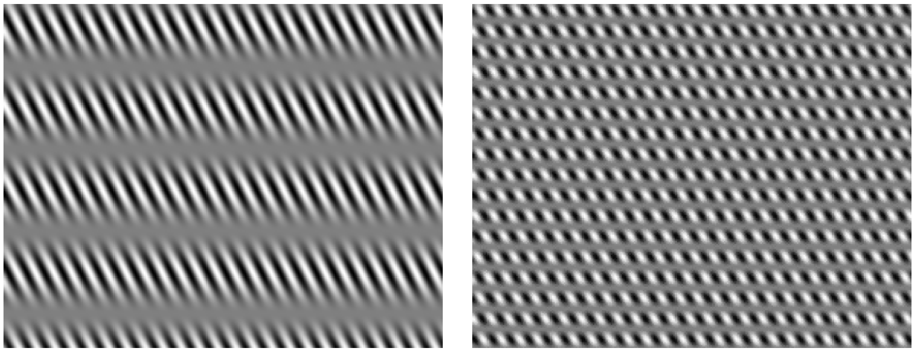

Figure 1. Example screen contents that cause the authors' computer monitor to broadcast an $f_{\mathrm{t}} = 300$ Hz (left) and 1200 Hz tone (right) on an $f_{\mathrm{c}} = 2.0$ MHz carrier in amplitude modulation.

involved, the power lines appear to propagate more RF energy than the parasitic antennas in the PC do. We suspect that if we connected the radio via a suitable HF bridge directly to the correct phase of the power network then we would receive the signal in neighbouring buildings. In addition, our handheld radio had only a simple untuned dipole antenna, so with a better antenna we would expect to get reasonable reception at several hundred meters.

The shortwave radio bands in the 1–30 MHz range seem to be the best for this attack. They are the highest bands that normal worldband radios can pick up and that are well below the pixel frequency f_{p}. Although computer monitors and video cables are too small to act as efficient antennas for these frequencies, the lower frequency bands would be even worse, while the VHF frequencies at which electronic components radiate well are too close to current pixel frequencies for software to modulate efficiently, especially using the frequency modulation supported by cheap radios in the VHF band. (Of course, within 5–10 years, rising pixel frequences might bring VHF FM radio within reach.)

The reception range depends largely on how noisy the radio spectrum is near the selected carrier frequency f_{c}, so this frequency should be selected to avoid nearby broadcast stations. Reception thus depends on the time of day, as the shortwave bands are crowded at night.

In a typical low-cost attack, the eavesdropper would place a radio and cassette recorder near the target and implant the attack software using standard virus or Trojan techniques. Since the broadcast patterns will be visible, the attack should take place after business hours while avoiding times when the chosen frequency is swamped by ionospheric propagation of interfering stations. Many PCs are not turned off at night, a habit encouraged by the power management features of modern systems. If monitors are also left powered up, then the attack software might monitor network traffic to detect the presence of people in the department. Where monitors are turned off but PCs are not, a serviceable signal can usually be picked up: as well as the power line, the VDU cable can be a quite adequate antenna. In these cases, the attack software can broadcast unobtrusively in the evening and early morning hours.

The attack software can use frequency shift keying, with 0 and 1 represented by tone patterns like those shown in Fig. 1. These would be loaded into two video buffers which would be switched at the frame rate f_c. The bit pattern would be encoded first to provide forward error correction before its bits are used to select the sequence of tones transmitted.

Our low-cost eavesdropper can then take the cassette with the recorded broadcast to her PC and digitise the signal with her sound card. The remaining steps involve symbol detection, synchronization and decoding and are described in any digital communications textbook [19]. Typical bit rates that can be obtained are of the order of 50 bit/s, so our attack software has to choose the data it transmits. Obvious targets include password files, key material and documents selected by text searching of the hard disk.

3 The Video Display Eavesdropping Receiver

We performed further experiments using an ESL model 400 Tempest monitoring receiver (Fig. 2) from DataSafe Ltd. of Cheltenham, UK. This device is not intended for signals intelligence missions; it was designed in the late 1980's as a test and demonstration tool to work with the video display technology of that period [9]. It is basically a normal black-and-white TV set with some modifications, of which the most important is that the sync signal recovery circuits have been replaced by two manually adjustable oscillators. The horizontal deflection frequency or line rate can be selected in the range 10–20 kHz with almost millihertz resolution, while the vertical deflection frequency or frame rate can be chosen in the range 40.0–99.9 Hz with 0.1 Hz resolution.

Like a normal TV set, this receiver performs an upper sideband linear demodulation with 8 MHz bandwidth and displays inverted video (a higher baseband voltage is shown darker on the 13 cm screen). Unlike a normal TV set, it can be freely tuned in four bands in the range 20–860 MHz and has a sensitivity ranging from 60 μV at 20 MHz to 5 μV at 860 MHz. A more expensive version of this receiver featured a larger screen, line frequencies up to 35 kHz, a demodulator that could be switched between linear AM, logarithmic AM and FM, a receiver bandwidth adjustable from 1.5–8 MHz, a notch filter and a manual override of the automatic gain control.

With a folded 4 m dipole antenna, we got the best image quality in the 100–200 MHz range. This antenna is by no means optimal; experiments with a borrowed spiral log conical antenna with a nominal 200–2000 MHz range gave much better reception results even at frequencies of 140–200 MHz. This more expensive antenna appears better suited to the elliptically polarised emanations from a typical video monitor.

The monitor used in our experiments is a common 43 cm Super-VGA PC monitor (model MT-9017E produced by *iiyama*, 160 MHz video bandwidth) that

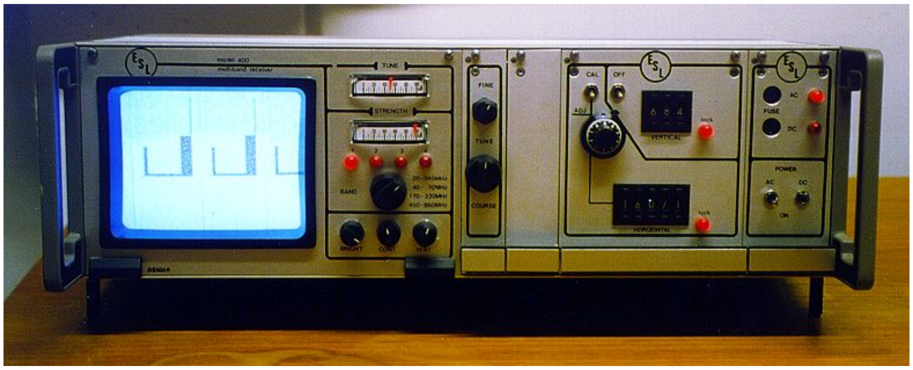

Figure 2. DataSafe/ESL Model 400 Tempest Emission Monitor used in our experiments.

fulfills the MPR II low-radiation requirements. The video mode is the same as that used in the audio broadcast experiment described in section 2.

The MPR and TCO low-radiation requirements specify only measurements in the bands up to 400 kHz. The fields emitted in these bands are mostly created by the deflection coils and do not carry much information about the screen content. The emissions related to the screen content are found mostly far above 30 MHz in the VHF and UHF band (unless we use pathological screen contents as in the audio broadcasting experiment described above). The MPR and TCO standards, which were introduced because of health concerns, do not require shielding in the VHF and UHF bands and are thus largely irrelevant from a Tempest point of view. Monitor buyers should not assume that so-called low-radiation monitors, or even LCD displays, provide any protection; we found that some modern TFT-LCD laptop displays give clearer reception than many cathode-ray tubes.

With a 64 kHz line frequency and 95 MHz pixel clock, our PC monitor was well outside the range of displays for which the ESL 400 had been designed. We had to set the horizontal synch generator to around 16.1 kHz, a quarter of the PC's actual frequency. This causes the screen content to be displayed in four columns on the receiver monitor; as successive pixel lines are now split up modulo four, normal text characters although visible are unreadable.

4 Hiding Information in Dither Patterns

We observed that our Tempest receiver mostly displays the high-frequency part of the video signal. The strongest useful spectral components are at frequencies close to the pixel frequency and its harmonics. However, monitor technology has changed critically over the past decade. The early 1980's terminals studied by van Eck in [2] switched the electron beam on and off for every single pixel. This improved image quality on the low video bandwidth CRTs of the time, as it made all the pixels in a line appear identical. Without this pixel pulsing, pixels in the middle of a horizontal line would appear brighter than those at the edge due to the slow voltage rise and fall times supported by early electronics. Thus short horizontal lines would have appeared as ovals.

Modern video display units have a much higher video bandwidth and so do not need pixel pulsing. As a result, all the eavesdropper can receive of a horizontal line on a modern monitor screen are two short impulses, emitted when the beam is switched on at the left end and switched off again at the right end. Indeed, the Tempest signal is roughly amplitude the of the video signal's derivative. This is not usually a problem with text, because characters (in most languages) are identifiable from their vertical components; but it hinders the reception of screen contents such as photographic images that cannot be reconstructed easily from their sharp vertical edges.

The human eye is less sensitive to high than to low spatial frequencies. Dithering is a technique that uses this to increase the number of colour shades available on displays with a small colour lookup table. On modern high-resolution monitors, users cannot easily distinguish between a medium grey and a checker board pattern of black and white pixels, especially as the distance between pixels is often smaller than the diameter of the electron beam focus. For the eavesdropper, on the other hand, the high-frequency black/white dither pattern creates the strongest possible signal while a constant colour results in the weakest.

We can use this difference in the spectral sensitivity of the user and the eavesdropper to present different information to them. Figure 3 shows on the left a test signal on the authors' workstation monitor, and on the right the image seen on our Tempest receiver.

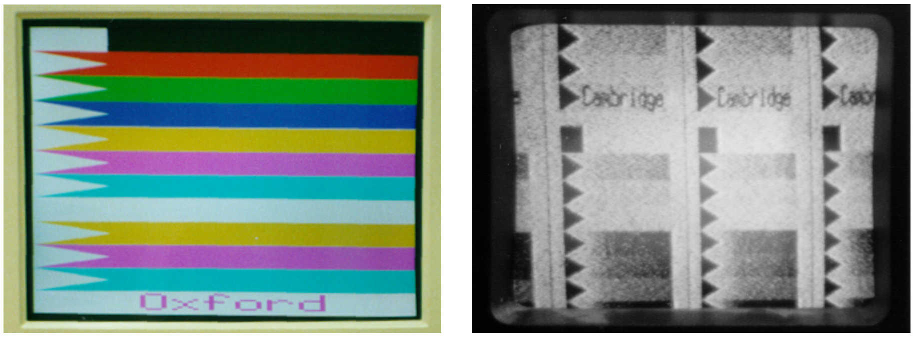

Figure 3. Test image as displayed on computer monitor (left) and the captured signal shown on the eavesdropping receiver. Our receiver supports only vertical deflection frequencies of 10–20 kHz, so we had to set it to 16.1 kHz, a quarter of the actual line frequency, and three copies of the image appear next to each other. (The fourth is lost during beam flyback.)

This test image contains on the left side one square and several triangular markers drawn with a dither pattern of vertical black and white lines. These markers help to locate other image features and even with our simple dipole antenna are very clearly visible on the receiver monitor, even in other rooms over 20 meters away. On the right side of every marker is a colour bar that looks uniform on the computer monitor but fades out on the left side of the Tempest image. These bars next to the seven triangles below the square were drawn in uniform colors (dark red, dark green, dark blue, yellow, magenta, cyan, and grey) on the left end, fading smoothly to dither patterns (red/black, green/black, blue/black, yellow/black, magenta/black, cyan/black, white/black) at the right. The next three bars below are again yellow, magenta, cyan on the left side, but this time the dither pattern shows a phase shift between the primary colours so the the dither pattern on the right end is red/green, red/blue, and blue/green. Between the left and right end of the bars, the amplitude of the dither pattern increases linearly. This test image enables us to see at a glance which of the three electron guns produces a usable Tempest signal and at which edge height.

One observation is that the signals generated with identical video input voltages for the three primary colours red, green, and blue show different Tempest amplitudes. One reason is that the white calibration of the monitor transfers equal input voltages into different acceleration voltages and beam currents. Another seems to be that the emissions for the three primary colours create different polarisations of the emitted waves; varying the antenna position changes the relative luminosity of the received test bars. Even the phase shift of one primary colour in the dither patterns of the second set of yellow, magenta, and cyan can be distinguished in some antenna positions. By evaluating polarisation modes with several antennas, it might even be possible for an eavesdropper to reconstruct some colours.

A fascinating application of the eavesdropper's sensitivity to dither amplitudes is given in the colour bar right of the eleventh triangle marker below the square. While the computer monitor clearly displays "Oxford" here in large letters, the eavesdropper sees instead the message "Cambridge". Figure 4 shows the magnified pixel field around the letters "Ox" that radiate as "Ca". While "Oxford" is drawn in magenta instead of grey by deactivating only the green component, "Cambridge" is embedded in the image by increasing the amplitude of the dithering.

The dither amplitude change must be smoothed in order not to arouse the very sensitive edge detectors implemented in the human retina. In order to make the change invisible, several physical effects of the monitor hardware have to be taken into account. The colour component value chosen by the display software is usually mapped linearly to the video input voltage supplied to the monitor. But the relation between the video voltage V and the luminosity L of the screen is non-linear and can be approximated as $L = \text{const} \cdot V^{\gamma}$, where γ is a hardware specific exponent that is usually in the range 1.5–3.0 depending on the design of the cathode ray tube. The programmer has to remember that the overall luminosity of a two colour dither pattern depends on the arithmetic mean of the luminosities rather than the voltages.

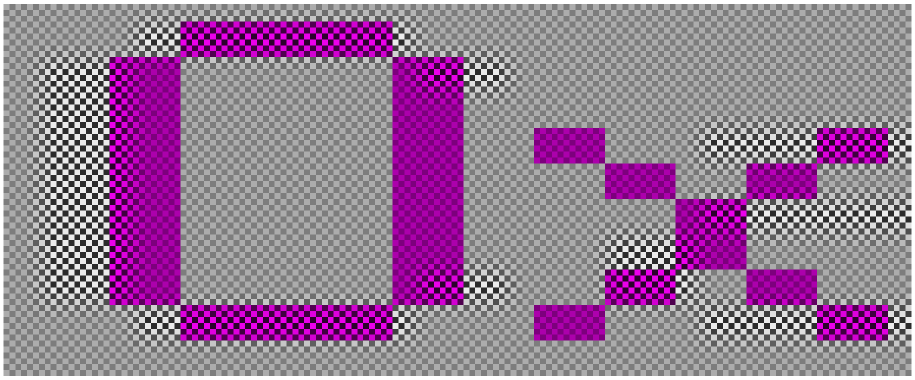

Figure 4. A magnification of the section that reads “Ox” on the computer monitor but “Ca” on the eavesdroppers screen (see Fig. 3) shows how the broadcasted message was hidden. The text made visible to the eavesdropper is present as gamma-corrected amplitude modulation in the background pattern, while the foreground message is just a low frequency signal.

The resulting calculations are known to TV and computer graphics specialists as *gamma correction*. We observed that the arithmetic average of the gamma corrected luminosities only predicts the luminosity accurately for a dither pattern consisting of horizontal lines. For dither patterns with vertical lines or checker board patterns, the finite bandwidth of the VDU beam introduces many intermediate voltages into the video signal supplied to the CRT. An accurate luminosity estimation for such dither patterns with high horizontal frequency components — the ones of interest for hiding information in emissions — would involve integration of the distorted gamma corrected video voltage. The gamma correction parameters that computers can download from modern monitors are thus not sufficient to correct a high-amplitude dither pattern.

Where the transmitted image must be very difficult to see, the dither parameters should be manually calibrated for a specific monitor. The calibration depends not only on the type of monitor, but also on its brightness, contrast, and colour temperature settings, which the user might modify. So a careful attacker will not attempt to hide readable text or bar code like information in uniformly coloured areas, but in structurally rich screen content like desktop background photos or the animations shown by screen saver programs. Such programs, as well as any software with access to the display, are therefore part of the trusted computing base, unless there is effective physical shielding.

5 Broadband Transmissions

Our dither amplitude modulation of large readable letters was designed to allow easy low-cost reception of hidden broadcast information with a modified TV set. A professional eavesdropper is more likely to select a method that affects only a small part of the screen layout and that is optimized for maximum range and robust reception with sophisticated equipment. In this section, we outline what such a system might look like.

Reception of monitor emanations with modified TV sets requires either exact knowledge of the horizontal and vertical deflection frequencies or a strong enough signal to adjust the sync pulse genereators manually. With larger distances and low signal levels, the emitted information can only be separated from the noise by averaging the periodic signal over a period of time, and manual adjustment of the synch is difficult.

In a professional attack, one might use spread-spectrum techniques to increase the jamming margin and thus the available range. The attack software would dither one or more colours in several lines of the screen layout using a pseudo-random bit sequence (PRBS). Trojan software, for example, might embed this dithering in its windows toolbar. A cross-correlator in the receiver gets one input from an antenna and sees at its other input the same pseudo-random bit sequence presented with the guessed pixel clock rate of the monitor. The cross-correlator will generate an output peak that provides the phase difference between the receiver and the target. A phase-locked loop can then control the oscillator in the receiver such that stable long-term averaging of the screen content is possible. Information can be transmitted by inverting the PRBS depending on whether a 0 or 1 bit is to be broadcasted. Readers familiar with direct sequence spread spectrum modulation will find the idea familiar, and many spread spectrum engineering techniques are applicable here.

If a simple PRBS coded as a series of black and white pixels is too different from the normal grey toolbar expected by the user, then phase modulation can be used instead. The amplitude of the dither pattern can be reduced smoothly for a few pixels at phase jumps in the dither pattern to avoid visible bright or dark spots in the toolbar. It is also possible to use only a small number of lines — perhaps only one unused line at the top of the toolbar (or even off the visible edge of the screen).

The advantages of using spread spectrum techniques are:

- only the pixel clock frequency and (perhaps) the carrier frequency have to be selected. This enables fast lock-on and fully automatic operation;
- higher reception ranges can be achieved as noise is suppressed by the cross-correlation and averaging;
- higher data rates can be achieved and automatic decoding of the received data is simplified.

An interesting commercial application of this could be in software license enforcement. Most software licenses allow the use of a product on only one computer at a time, but this condition is frequently violated. Commercial software vendors tackle piracy by forming trade associations which prosecute offenders, but the main enforcement problem is not so much identifying offending companies as getting the initial search warrant. This motivates the design of a system that will detect piracy from outside an offender's premises.

Our suggestion is that software packages include in their screen layout a few lines with a PRBS signal that encodes the license serial number plus a random value [20]. Just as "TV detector vans" circulate in countries with mandatory television license fees to discover unlicensed TV sets from their stray RF emissions, a "software detector van" can be used to patrol business districts and other areas where software piracy is suspected. If the van receives twenty signals from the same copy of Word from a company that has only licensed five copies, then probable cause for a search warrant has been established.

The random value encoded in the PRBS helps distinguish echos from messages received from different computers. Finally, if the PRBS were displayed by the operating system, it could broadcast the identities and license numbers of all currently active programs.

6 A New Protective Measure: Tempest Fonts

As we noted above, only the high-frequency components of the video signal can be picked up by the eavesdropper. Figure 5 shows on the left a test image that helps us to determine which part of the image spectrum actually produces a Tempest signal. This "zoneplate" signal is used by TV technicians, and is generated from the function $\cos(x^2 + y^2)$ where the coordinate system origin is in the image center. At every point of this test signal, the local spectrum has a single peak at a horizontal and vertical frequency that is proportional to the horizontal and vertical coordinates of this point. This frequency peak reaches the Nyquist frequency $f_\mathrm{p}/2$ for the points at the border of the zoneplate image.

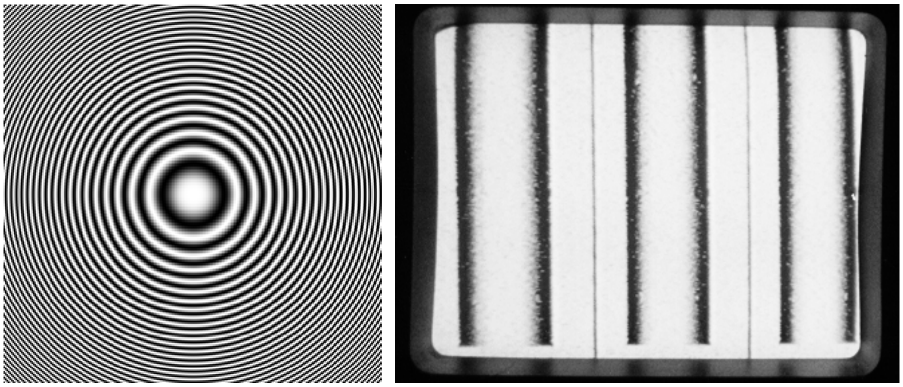

Figure 5. In the zoneplate test signal (left), every point has a local spectrum with a horizontal and a vertical frequency proportional to its coordinates (origin in the center). The received image (right) shows that only those parts of the zoneplate signal where the horizontal frequency is in the upper 30% of the spectrum (i.e., $> 0.7 \cdot f_\mathrm{p}/2$) cause the monitor to emanate a significant amount of energy.

In the right part of Fig. 5, we can see the Tempest signal received from a monitor showing the zoneplate image (for this and the other experiments described in this section, we brought our antenna as close as possible to the monitor to generate best case reception conditions). As one might expect, only the horizontal frequency of the signal determines what is being received. What is more interesting is that only the outer 30% of the zoneplate image area appears dark on the receiver monitor. This means that if we look at the Fourier transform of a horizontal sequence of pixels, only information present as frequencies f in the range $0.7 \cdot f_{\mathrm{p}}/2 < f \leq f_{\mathrm{p}}/2$ in the Fourier spectrum can be received in our setup. This value 0.7 clearly depends on the equipment used, but seems to be not untypical.

We wondered whether this leads us to a potentially very cheap software-based eavesdropping protection technique. Figure 6 shows on the left side a magnified pixel field that displays some text. On the right ride, the same pixel field is shown after we removed the top 30% of the Fourier transform of the signal by convolving it with a suitable $\sin(x)/x$ low-pass filter.

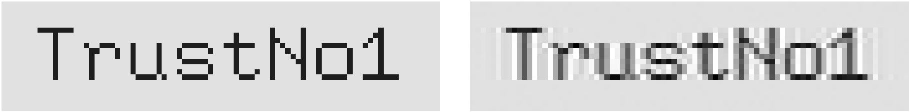

Figure 6. The text on the left is displayed with a conventional font, while the text on the right has been filtered to remove the top 30% of the horizontal frequency spectrum. While practically no difference between the fonts can be perceived by the user on a computer monitor, the filtered text disappears from the eavesdropper's monitor while the normal text can be received clearly.

The filtered text looks rather blurred and unpleasant in this magnified representation, but surprisingly, the loss in text quality is almost unnoticeable for the user at the computer screen. The limited focus of the electron beam, the limited resolution of the eye, as well as effects created by the mask and the monitor electronics filter the signal anyway.

While there is little visible change for the user, filtering the displayed text causes a previously easily receivable text to vanish completely from the Tempest monitor, even if the antenna is next to the VDU. Filtered text display requires greyscale representation of glyphs, but this technology is already available in many display drivers in order to support antialiasing fonts. We are optimistic that if the HF filtering is combined carefully with anti-aliasing techniques, readability can be better than with simple bi-level fonts.

The eavesdropping of text displayed on a monitor is only one type of Tempest risk associated with personal computers. Nevertheless, we still consider it the most significant risk. The video display unit is usually the strongest source of radiation and due to its periodic nature, a video signal can easily be separated from other signals by periodic averaging.

We have identified two more potential sources of periodic signals in every PC, both of which can be fixed at low cost by software or at worst firmware changes. Keyboard controllers execute an endless key-matrix scan loop, with the sequence of instructions executed depending on the currently pressed key. A short random wait routine inside this loop can prevent an eavesdropper doing periodic averaging. Secondly, many disk drives read the last accessed track continuously until another access is made. As an attacker might try to reconstruct this track by periodic averaging, we suggest that after accessing sensitive data, the disk head should be moved to a track with unclassified data unless further read requests are in the queue. Needless to say, if the keyboard and disk drive signals are strong enough to receive outside the controlled area, then their driver software (like the screen driver) is part of the trusted computing base.

The emanations from most other sources, such as the CPU and peripherals, are usually transient. To use them effectively, the eavesdropper would have to install software that drives them periodically, or at least have detailed knowledge of the system configuration and the executed software.

We are therefore convinced that our Soft Tempest techniques, and in particular Tempest fonts, allow a significant increase in security at a very low cost. There are many applications where they may be enough; in medium-sensitivity applications, many governments use a zone model in which computers with confidential data are not shielded but located in rooms far away from accessible areas. Here, the 10–20 dB of protection that a Tempest font affords is of major significance. There are also applications where Tempest fonts are the only option, such as when a nation suddenly has to buy large quantities of commercial off-the-shelf computers and field them in a sudden deployment such as Desert Storm. Finally, in applications such as diplomacy that require the highest levels of protection, users should install soft as well as hard Tempest protection; hardware shielding often fails due to dirty gaskets or to procedural security problems such as ambassadors refusing to keep doors closed on a hot day.

7 Conclusions

Compromising emanations continue to be a fascinating field of research, although they are mostly unexplored in the research literature. The high costs of physical shielding and the continuously increasing clock frequencies of modern computers ensure that the problem will not go away quickly, while the arrival of software radios on the hobby market will make things worse.

However, we have shown that Tempest is not just about RF engineering. Software techniques can make a huge difference: they can be used to mount new attacks, construct new defences and implement some quite novel applications. We believe that our Soft Tempest technology can significantly raise the whole Tempest game.

References

1. Kristian Beckman: Läkande Datorer [Leaking Data]. Cited in [3]

2. Wim van Eck: Electromagnetic Radiation from Video Display Units: An Eavesdropping Risk? Computers & Security **4** (1985) 269–286
3. Harold Joseph Highland: Electromagnetic Radiation Revisited. Computers & Security **5** (1986) 85–93 and 181–184.
4. Peter Smulders: The Threat of Information Theft by Reception of Electromagnetic Radiation from RS-232 Cables. Computers & Security **9** (1990) 53–58
5. Erhard Möller, Lutz Bernstein, Ferdinand Kolberg: Schutzmaßnahmen gegen kompromittierende elektromagnetische Emissionen von Bildschirmsichtgeräten [Protective Measures Against Compromising Electro Magnetic Radiation Emitted by Video Display Terminals]. Labor für Nachrichtentechnik, Fachhochschule Aachen, Aachen, Germany
6. Deborah Russell, G. T. Gangemi Sr.: Computer Security Basics. Chapter 10: TEMPEST, O'Reilly & Associates, 1991, ISBN 0-937175-71-4
7. Peter Wright: Spycatcher – The Candid Autobiography of a Senior Intelligence Officer. William Heinemann Australia, 1987, ISBN 0-85561-098-0
8. Ernst Bovenlander, invited talk on smartcard security, Eurocrypt 97
9. Operating Manual for DataSafe/ESL Model 400B/400B1 Emission Monitors. DataSafe Limited, 33 King Street, Cheltenham, Goucestershire GL50 4AU, United Kingdom, June 1991
10. Lars Høivik: System for Protecting Digital Equipment Against Remote Access. United States Patent 5165098, November 17, 1992
11. John H. Dunlavy: System for Preventing Remote Detection of Computer Data from TEMPEST Signal Emissions. United States Patent 5297201, March 22, 1994
12. Joachim Opfer, Reinhart Engelbart: Verfahren zum Nachweis von verzerrten und stark gestörten Digitalsignalen und Schaltungsanordnung zur Durchführung des Verfahrens [Method for the detection of distorted and strongly interfered digital signals and circuit arrangement for implementing this method]. German Patent DE 4301701 C1, Deutsches Patentamt, May 5, 1994
13. Wolfgang Bitzer, Joachim Opfer: Schaltungsanordnung zum Messen der Korrelationsfunktion zwischen zwei vorgegebenen Signalen [Circuit arrangement for measuring the correlation function between two provided signals]. German Patent DE 3911155 C2, Deutsches Patentamt, November 11, 1993
14. 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
15. Überkoppeln auf Leitungen [Cross-talk on cables], Faltblätter des BSI **4**, German Information Security Agency, Bonn, 1997.
16. Schutzmaßnahmen gegen Lauschangriffe [Protection against bugs], Faltblätter des BSI **5**, German Information Security Agency, Bonn, 1997.
17. Bloßstellende Abstrahlung [Compromising Emanations], Faltblätter des BSI **12**, German Information Security Agency, Bonn, 1996.
18. RJ Lackey, DW Upmal, Speakeasy: The Military Software Radio. IEEE Communications Magazine v 33 no 5 (May 95) pp 56–61
19. John G. Proakis: Digital Communications. 3rd ed., McGraw-Hill, New York, 1995, ISBN 0-07-051726-6
20. Ross J Anderson, Markus G Kuhn, Software Piracy Detector Sensing Electromagnetic Computer Emanations. UK Patent application no GB 9722799.5, 28th November 1997