

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
{mgk25,rja14}@cl.cam.ac.uk

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 transmits. This can be used for both attack and defence. To attack a system, malicious code can encode stolen information in the machine's RF emissions 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 these emissions. 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 which not only interferes with radio reception, but also leaks information about the data being processed. Known as *compromising emanations* or *Tempest* radiation, a code word for a U.S. government programme aimed at attacking the problem, the electromagnetic broadcast of data has been a significant concern in sensitive computer applications.

In his book 'Spycatcher' [1], 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 and failed. 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 [2]. 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

* Supported by a European Commission Marie Curie training grant

handles or transmits unclassified data. Equipment with both ‘red’ and ‘black’ connections, such as cipher machines and multilevel secure workstations, requires particularly thorough testing. The U.S. standard NACSIM 5100A that specifies the test requirements for Tempest protected equipment, and its NATO equivalent AMMSG 720B, are classified documents [3–5]. In Germany, even the names of the government standards on compromising radiation are kept secret.

So we lack full information about the measurement technology required for Tempest tests, but descriptions in published patents [6, 7] 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 can be suitable sensors, especially to test noisy 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 acoustic emanations.

When conducting attacks, similar periodic averaging and cross-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. Knowledge of the fonts used with video displays and printers allows maximum likelihood character recognition techniques to give a better signal/noise ratio for whole characters than is possible for individual pixels. Malicious software implanted by an attacker can also generate periodic or pseudorandom signals that are easy to detect.

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 smartcard 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 non-volatile PIN retry counter is updated. In this way, the PIN retry limit may be defeated.

Electromagnetic radiation as a computer security risk was mentioned in the open literature as early as 1967 [9]. One of the first more detailed public descriptions of the Tempest threat appears to have been a 1983 report in Swedish [10], but the problem was brought to general attention by a 1985 paper [11] 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 re-

sults were later confirmed by Möller, Bernstein and Kolberg, who also discuss various shielding techniques [12].

purely theoretical contributions are difficult due to the lack of published data about the information-carrying 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 peripherals are many times more expensive than 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.

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 [19], such equipment is not yet generally available. The graduate student's Tempest spying kit is more likely to be just a 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 designed a screen image that causes the VDU beam current to approximate a broadcast AM radio signal. If this latter has a carrier frequency f_c and an audio tone with a frequency f_t , then it can be represented as

$$\begin{aligned} s(t) &= A \cdot \cos(2\pi f_c t) \cdot [1 + m \cdot \cos(2\pi f_t t)] \\ &= A \cdot \left\{ \cos(2\pi f_c t) + \frac{m}{2} \cdot \cos[2\pi(f_c - f_t)t] + \frac{m}{2} \cdot \cos[2\pi(f_c + f_t)t] \right\}. \end{aligned}$$

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 centre of one pixel to the centre of its right neighbour. The pixel clock is an integer multiple of both the horizontal and vertical deflection frequencies, that is the rate $f_h = f_p/x_t$ with which lines are drawn and the rate $f_v = f_p/y_t$ 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 $x_t y_t - x_d y_d$ 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 $f_p = 95$ MHz, $x_d = 1152$, $y_d = 900$, $x_t = 1472$ and $y_t = 939$ are used on this system, which leads to deflection frequencies of $f_h = 64.5$ kHz and $f_v = 68.7$ Hz.

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

$$t = \frac{x}{f_p} + \frac{y}{f_h} + \frac{n}{f_v},$$

for all $0 \leq x < x_d$, $0 \leq y < y_d$ 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 $\lfloor \frac{255}{2} + s(t) + R \rfloor$ with amplitudes $A = \frac{255}{4}$ and $m = 1$, where $0 \leq R < 1$ is a uniformly distributed random number that spreads the quantization noise (dithering). See Fig. 1 for screen contents generated this way to broadcast an AM tone.




Fig. 1. Example screen contents that cause the authors' computer monitor to broadcast an $f_t = 300$ Hz (left) and 1200 Hz tone (right) on an $f_c = 2.0$ MHz carrier in amplitude modulation.

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 line or 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 involved, the power lines appear to propagate more RF energy than the parasitic antennas in the PC do. 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 metres.

The shortwave (HF) radio bands in the 3–30 MHz range seem to be the best for this attack. They are the highest bands that normal radios can pick up and that are well below the pixel frequency f_p . Although computer monitors and video cables are too small to be 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 FM. (Of course, as time passes, rising pixel frequencies 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_v . Fast switches between screen patterns and real-time amplitude modulation can also be accomplished using the colour lookup table. 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 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 as they are described in any digital communications textbook [20]. 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.

A side note for designers of unusual radio transmitters: a PC graphics adapter can be transformed into a digital short-wave transmitter by connecting a suitable antenna to the video output. With no monitor connected, we can set $x_d = x_t$ and $y_d = y_t$ to suppress the sync pulses and blanking intervals and get a clean spectral shape. With carefully optimized software, modern processors can fill video memory faster than the graphics hardware reads it out, opening the possibility of real-time voice and data transmission either using a standard laptop or with at most a very simple RF output stage.

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 [21]. It is basically a small 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 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.



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

The emissions related to the screen content are mostly found far above 30 MHz in the VHF and UHF band (unless we use pathological screen contents as in the audio broadcasting experiment described above). These 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 screens, provide any Tempest 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 video mode was well outside the range of displays for which the ESL 400 had been designed. We had to set the horizontal sync 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.

A Tempest monitor sufficient to repeat the following experiments can be built by interconnecting common household components: a PC multi-sync monitor provides the display, a PC graphics adapter card provides the sync pulses, and a VCR tuner connected to a VHF TV antenna and amplifier performs the demodulation.

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 [11] 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 because of 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 the derivative of the video signal. This is not usually a problem with text, because characters (in most languages) are identifiable from their vertical components; but it hinders eavesdropping 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 or halftoning is a technique that uses this to increase the number of colour shades available on displays with a small colour lookup table [22]. On modern high-resolution monitors, users cannot easily distinguish between a medium grey and a chequered halftoning 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.




Fig. 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 metres 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. The bars next to the seven triangles below the square were drawn in uniform colours (dark red, green, 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 that 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 control 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.

As Fig. 5 shows, this can be generalized. We can embed not only text but arbitrary greyscale images inside other cover images. Embedded images give an impression of the large bandwidth that is available to the attacker by dither modulation. Let $C_{x,y,c}$ be the value of a cover image at pixel coordinates (x, y) for primary colour $c \in \{\text{red, green, blue}\}$ and let $E_{x,y}$ be the pixel value of the image that shall be embedded covertly for reception by the eavesdropper. Then the colour component values that we have to display on the screen are

$$S_{x,y,c} = (C_{x,y,c}^{\tilde{\gamma}} + \min\{\alpha(1 - E_{x,y}), C_{x,y,c}^{\tilde{\gamma}}, 1 - C_{x,y,c}^{\tilde{\gamma}}\} \cdot d_{x,y})^{1/\tilde{\gamma}}$$




Fig. 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 broadcast 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.

where $d_{x,y} = 2[(x + y) \bmod 2] - 1 \in \{-1, 1\}$ is the dither function, $0 < \alpha \leq 0.5$ is a parameter that determines the maximum amplitude of the added dithering and $\tilde{\gamma}$ is as described below. Here, all pixel values are normalised values in the range 0 (black) to 1 (maximum luminosity), so with 8-bit displays the value written into the frame buffer is $\lfloor 255 \cdot S_{x,y,c} + R \rfloor$.

The colour component value chosen by the display software is usually mapped linearly to the video 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 $L = \text{const} \cdot V^\gamma$ for CRTs, where γ is usually in the range 1.5–3.0 and depends on the hardware design of the monitor. Software that compensates this non-linearity performs what is known as *gamma correction* [22, 23]. The overall luminosity of a two-colour dither pattern depends on the arithmetic mean of the luminosities L rather than the voltages V . To remain inconspicuous for the viewer, amplitude variations in the dither pattern must be performed such that the average luminosity is preserved.

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 chequered patterns, the restricted bandwidth of the beam current introduces many intermediate values. 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 signal [24]. We performed tests in which we determined the video voltage \bar{V} that gives rise to a colour of the same brightness as a dither mix of the voltages V_1 and V_2 . For a dither pattern of horizontal lines, the formula $\bar{V} = (\frac{1}{2}V_1^\gamma + \frac{1}{2}V_2^\gamma)^{1/\gamma}$ produced excellent predictions with $\gamma = 2.0$, the exponent for our CRT. For a




Fig. 5. The left photo shows what the user sees on her computer monitor. At the same time, the eavesdropping receiver shows in the right photo two greyscale images that we embedded. Bright and dark parts of the cover image reduce the amplitude that the embedding can utilize inconspicuously, so these areas become visible as bright shadows on the eavesdropping receiver ($\alpha = 0.4$).

chequered dither pattern, which looks much smoother, the same formula still worked, but the exponent changed to $\tilde{\gamma} = 1.28$. This is the value that we have to use to determine $S_{x,y,c}$. The gamma-correction parameters that computers can download from modern monitors are thus not sufficient to gamma-correct a high-amplitude chequered dither pattern.

The embedded image should be smoothed in order not to arouse the very sensitive edge detectors implemented in the human retina. Where the transmitted image must be very difficult to see, the correction parameter $\tilde{\gamma}$ 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 barcodes in uniformly coloured areas, but in structurally rich content like photos or the animations shown by screen savers.

The implication for systems with mandatory access control is that any software with pixel-level access to an unshielded display must either be part of the trusted computing base or be prevented from accessing protected data.

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 generators 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 pseudorandom bit sequence. A cross-correlator in the receiver gets one input from an antenna and sees at its other input the same pseudorandom bit sequence presented with the guessed pixel clock rate of the monitor. It 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 sequence depending on whether a 0 or 1 bit is to be broadcast. Readers familiar with direct sequence spread-spectrum modulation [20] will find the idea familiar, and many spread-spectrum engineering techniques are applicable.

The advantages of using spread-spectrum techniques are that higher data rates and reception ranges can be achieved, and that only the pixel clock frequency and (perhaps) the carrier frequency have to be selected. This enables fast lock-on and fully automatic operation.

A practical benefit is that it may only be necessary to use a small number of lines—perhaps in the toolbar, or even off the visible edge of the screen. If a spreading sequence 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 to avoid visible bright or dark spots.

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 signal that encodes the license serial number plus a random value [27]. 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 a software 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 signal helps distinguish echoes from messages received from different computers. Finally, if the signal were displayed by

the operating system, it could contain 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 6 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 centre. 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 $f_p/2$ for the points at the border of the zoneplate image.




Fig. 6. 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 centre). 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., $> 0.7 \cdot f_p/2$) cause a significant signal to be received from our monitor.

In the right part of Fig. 6, 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 give best reception). As one might expect, only the horizontal frequency of the signal determines what is received. Note that only the outer 30% of the zoneplate image area appears dark on the receiver. 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 $0.7 \cdot f_p/2 < f \leq f_p/2$ in the 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 7 shows in the upper left side a magnified pixel field that displays some text. On the upper right, 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.




Fig. 7. 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. The graphics in the upper row show the pixel luminosities, while below there are magnified screen photographs of a 21×5 mm text area. While the user can see practically no difference between the fonts, the filtered text disappears from our eavesdropping 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, as the magnified photos in the lower half of Fig. 7 show. 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, such filtering causes a text which could previously be received easily to vanish completely from the Tempest monitor, even when the antenna is right next to the VDU (Fig. 8). Filtered text display requires greyscale representation of glyphs, but this technology is already available in many display drivers in order to support anti-aliasing fonts. We are optimistic that if the low pass filtering is combined carefully with anti-aliasing techniques, readability can be better than with simple bi-level fonts. Simple low pass filtering could also be performed without software modifications by programming the filter available in the next generation of graphic adapters. Good anti-Tempest display routines will also apply the opposite of the techniques used in OCR fonts: there could be small random variations to the glyphs to make automatic character recognition by the eavesdropper more challenging.

Eavesdropping text from a monitor is only one of the Tempest risks associated with personal computers. Nevertheless, we still consider it the most significant one. The video display unit is usually the strongest source of radiation and due




Fig. 8. On the left we see what the eavesdropping monitor shows when text is displayed with normal fonts. The small screen size and the modulo four separation of image lines renders the text unreadable on our simple monitor, but the presence of the signal is clear. On the right, the screen content was low pass filtered as in Fig. 7 and the received Tempest signal has vanished except for the horizontal sync pulses.

to its periodic nature, a video signal can easily be separated from other signals and from noise 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 [28]. 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 and a random scan order 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.

DRAM refresh is another periodic process in every computer that deserves consideration. 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 convinced that our Soft Tempest techniques, and in particular Tempest fonts, can provide a significant increase in emanation 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 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 cost of physical shielding and the continuously increasing clock frequencies of modern computers ensure that the problem will not go away quickly. Things will be made worse by the arrival of cheap software radios—universal receivers in which all demodulation of the signal after the intermediate frequency conversion is done completely in software on high-speed DSPs [19]. This technology will allow low-budget attackers to implement sophisticated Tempest attacks which were previously only possible with very expensive dedicated equipment.

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.

The attack consists of implanting malicious software in a target computer to steal secret data and transmit it in a manner optimised for some combination of reception range, receiver cost and observability. This ‘Tempest virus’ can attack computers not connected to any communication lines and situated in rooms from which the removal of storage media is prohibited. It can also be used in commercial applications such as software copy protection.

On the defensive side, we have shown how fonts can be designed with spectral characteristics that significantly reduce the effective range of eavesdropping at a negligible cost in image quality.

References

1. Peter Wright: *Spycatcher – The Candid Autobiography of a Senior Intelligence Officer*. William Heinemann Australia, 1987, ISBN 0-85561-098-0
2. *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
3. Deborah Russell, G. T. Gangemi Sr.: *Computer Security Basics*. Chapter 10: TEMPEST, O’Reilly & Associates, 1991, ISBN 0-937175-71-4
4. A. J. Mauriello: Join a government program to unveil Tempest-spec mysteries. *EDN* vol 28 no 13, pp 191–195, June 23, 1983
5. Anton Kohling: TEMPEST – eine Einführung und Übersicht zu kompromittierenden Aussendungen, einem Teilaspekt der Informationssicherheit. In H.R. Schmeer (ed.): *Elektromagnetische Verträglichkeit/EMV’92*, Stuttgart, February 1992, pp 97–104, VDE-Verlag, Berlin, ISBN 3-8007-1808-1.

6. 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
7. 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
8. Ernst Bovenlander, invited talk on smartcard security, Eurocrypt '97, May 11–15, 1997, Konstanz, Germany
9. Harold Joseph Highland: Electromagnetic Radiation Revisited. *Computers & Security* vol 5, pp 85–93 and 181–184, 1986
10. Kristian Beckman: Läckande Datorer [Leaking Computers]. Cited in [9, 18]
11. Wim van Eck: Electromagnetic Radiation from Video Display Units: An Eavesdropping Risk? *Computers & Security* vol 4, pp 269–286, 1985
12. 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
13. Peter Smulders: The Threat of Information Theft by Reception of Electromagnetic Radiation from RS-232 Cables. *Computers & Security* vol 9, pp 53–58, 1990
14. Überkoppeln auf Leitungen [Cross-talk on cables], Faltblätter des BSI 4, German Information Security Agency, Bonn, 1997.
15. Schutzmaßnahmen gegen Lauschangriffe [Protection against eavesdropping attacks], Faltblätter des BSI 5, German Information Security Agency, Bonn, 1997.
16. Bloßstellende Abstrahlung [Compromising Emanation], Faltblätter des BSI 12, German Information Security Agency, Bonn, 1996.
17. Joel McNamara: The Complete, Unofficial TEMPEST Information Page. Internet Web page, URL <<http://www.eskimo.com/~joelm/tempest.html>>.
18. Harold Joseph Highland: The Tempest over Leaking Computers. *Abacus* vol 5 no 2, pp 10–18 and 53, 1998
19. Raymod J. Lackey, Donald W. Upmal: Speakeasy: The Military Software Radio. *IEEE Communications Magazine* vol 33 no 5, pp 56–61, May 1995
20. John G. Proakis: *Digital Communications*. 3rd ed., McGraw-Hill, New York, 1995, ISBN 0-07-051726-6
21. Operating Manual for DataSafe/ESL Model 400B/400B1 Emission Monitors. DataSafe Limited, 33 King Street, Cheltenham, Gloucestershire GL50 4AU, United Kingdom, June 1991
22. James D. Foley, Andries van Dam: *Fundamentals of Interactive Computer Graphics*, Addison-Wesley, 1982
23. Michael Bach, Thomas Meigen, Hans Strasburger: Raster-scan cathode-ray tubes for vision research—limits of resolution in space, time and intensity and some solutions. *Spatial Vision* vol 10 no 4, pp 403–414, 1997
24. Stanley A. Klein, Q. James Hu, Thom Carney: The Adjacent Pixel Nonlinearity: Problems and Solutions. *Vision Research* vol 36 no 19, pp 3167–3181, 1996
25. Lars Høivik: System for Protecting Digital Equipment Against Remote Access. United States Patent 5165098, November 17, 1992

26. John H. Dunlavy: System for Preventing Remote Detection of Computer Data from TEMPEST Signal Emissions. United States Patent 5297201, March 22, 1994
27. Markus G. Kuhn, Ross J. Anderson: Software Piracy Detector Sensing Electromagnetic Computer Emanations. UK patent application no 9722799.5, October 29, 1997
28. Markus G. Kuhn, Ross J. Anderson: Low Cost Countermeasures Against Compromising Electromagnetic Computer Emanations. UK patent application no 9801745.2, January 28, 1998