Electromagnetic Eavesdropping Risks of Flat-Panel Displays

Markus G. Kuhn

University of Cambridge, Computer Laboratory,
15 JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom
`http://www.cl.cam.ac.uk/~mgk25/`

Abstract. Electromagnetic eavesdropping of computer displays – first demonstrated to the general public by van Eck in 1985 – is not restricted to cathode-ray tubes. Modern flat-panel displays can be at least as vulnerable. They are equally driven by repetitive video signals in frequency ranges where even shielded cables leak detectable radio waves into the environment. Nearby eavesdroppers can pick up such compromising emanations with directional antennas and wideband receivers. Periodic averaging can lift a clearly readable image out of the background noise. The serial Gbit/s transmission formats used by modern digital video interfaces in effect modulate the signal, thereby making it even better suited for remote reception than emanations from analog systems. Understanding the exact transmission format used leads to new attacks and defenses. We can tune screen colors for optimal remote readability by eavesdroppers. We can likewise modify text-display routines to render the radio emanations unreadable.

1 Introduction

Electronic equipment can emit unintentional signals that allow eavesdroppers to reconstruct processed data at a distance. This has been a concern for the design of military hardware for over half a century. Some governments handle highly confidential information only with equipment that is especially shielded against such compromising electromagnetic emanations. The exact “TEMPEST” emission limits and test procedures applied in the procurement of these systems are still secret. Anecdotal evidence suggests that they are several orders of magnitude stricter than, for example, civilian radio-interference regulations.

Electromagnetic radiation as a potential computer security risk was mentioned in the open literature as early as 1967 [1]. The concept was brought to the attention of the broader public in 1985 by van Eck [2], who showed that the screen content of a cathode-ray tube (CRT) display can be reconstructed at a distance using a TV set whose sync pulse generators are replaced with manually controlled oscillators. Several more studies of the compromising video emanations of late 1980s CRT displays appeared [3,4,5,6,7], with advice on electromagnetic shielding as a countermeasure. Steganographic embedding of information into CRT emissions and the use of low-pass filtered fonts as a simple software countermeasure have been demonstrated as well [8].

Presented at 4th Workshop on Privacy Enhancing Technologies, 26–28 May 2004, Toronto, Canada

Display technologies have evolved rapidly since then. Additional shielding has become standard, not only to meet stricter international electromagnetic compatibility requirements [9], but also to address health worries associated with non-ionizing radiation [10]. Pixel frequencies and video bandwidths have increased by an order of magnitude since [2,3,4,5,6,7] 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 of replacing the cathode-ray tube (CRT) monitor. All these developments make it necessary to reevaluate the emission-security risks identified in the 1980s.

A new form of compromising emanations from video displays was discovered more recently. The high-frequency variations of light emitted by a CRT can carry enough information about the video signal to permit the reconstruction of readable text [11]. Under low background illumination, this is practical even after diffuse reflection from nearby surfaces. LCDs are not vulnerable to this particular risk, not only because their pixels react much slower than CRT phosphors, but also because these technologies update all pixels in a row simultaneously. This makes it impractical to separate the contribution of individual pixels in a row to the overall light emitted.

Discussions following the publication of [11] suggest that flat-panel displays are widely believed to pose no electromagnetic eavesdropping risk either. Two facts may contribute to such an assumption. Firstly, FPDs lack deflection coils, which makes them – compared to CRTs – "low radiation" devices in the frequencies below 400 kHz, where field strengths are limited by a Swedish ergonomic standard [10]. Secondly, LCDs operate with low voltages and – unlike CRTs – do not amplify the video signal by a factor of about 100 to drive a control grid that modulates an electron beam.

The experiments reported here demonstrate that some types of flat-panel display do pose a realistic eavesdropping risk. In particular, with some modern video interfaces, it is quite easy to configure the display of text in a way that maximizes the leaking signal strength. This makes emanations from these displays even easier to receive than those of modern CRTs. We begin with a brief description of video, eavesdropping and measurement technology in Sect. 2 and 3. The two case studies presented in Sect. 4 and 5 analyze the compromising radio emanations first from a laptop LCD and then from a desktop LCD that is connected to its PC graphics card with a *Digital Visual Interface (DVI)* cable. In both cases, the video cable used to connect the display panel with the graphics controller turned out to be the primary source of the leaking signal. An understanding of the digital transmission format used helped to optimize the choice of screen colors to raise or reduce the feasibility of an eavesdropping attack significantly.

2 Video Display Interfaces

Early video terminals contained the frame buffer and CRT in a single unit, avoiding the need for a user-visible video interface. With the modular PC architecture introduced by the IBM PC, displays and graphics cards turned into exchangeable components, available from multiple vendors with standardized connectors. The signalling techniques used on these interfaces were initially parallel digital interfaces. With 1, 4, and 6 TTL-level lines, respectively, the IBM PC's MDA, CGA, and EGA video controllers signalled the color of each pixel to the monitor. With the 15-pin VGA connector introduced in 1987, the dominant personal computer display interface turned to using three analog voltages (0–0.7 V), one to control each primary color.

More recently, the industry moved back to digital video signalling for two reasons. The first is related to signal quality limits. The geometry of the old 15-pin VGA connector was not designed for very-high-frequency signals. The 640×480@60Hz video mode used by the original VGA card had a pixel clock frequency of merely 25 MHz, whereas more recent high-end displays use pixel rates of 300 MHz or more. As signal wavelengths drop below typical cable lengths, the lack of a properly impedance-matched coaxial feedthrough in the VGA connector causes increased inter-pixel interference.

The second reason is the advent of flat-panel technologies, such as liquid-crystal, plasma, or organic electroluminescence displays. These devices have to sample the video signal, in order to assign to each discrete pixel on the display surface its current color via row and column access lines. They maximize contrast by buffering an entire line of the video signal, to drive all pixels in a row concurrently.

As flat-panel displays have to store video lines in digital memory, they require video information not only as binary encoded color shades, but also as a sequence of discrete pixel values. All recent digital interface standards therefore include a pixel clock line, avoiding the reconstruction of the pixel clock signal that has to be performed in FPDs with VGA input.

Current flat-panel displays buffer digitally only a few pixel rows. The entire image is still stored only in the frame buffer of the video controller. Modern flat-panel video interfaces therefore still have to continuously refresh the entire image content between 60 and 85 times per second, just as with CRTs. This continuous refresh ensures that the signals on the video interface are periodic, at least between changes in the displayed information. A periodic signal has a frequency spectrum that consists of narrow lines spaced by the repetition frequency. A receiver can attenuate all other spectral content by periodic averaging with the exact same repetition frequency.

3 Eavesdropping Instrumentation

Any signal carried by a conductor can, at least in principle, be eavesdropped electromagnetically, by simply connecting a nearby antenna to an amplifier and recording device, for example a digital storage oscilloscope. While this approach can be useful in attempts to record a waveform in the largest possible bandwidth, it is in practice not feasible, unless the signal is strong, or the experiment is performed with very low background noise. Outside special shielded chambers, waveforms picked up by antennas will be dominated by the many radio broadcast services that populate the spectrum from below 10 kHz to above 10 GHz, not to mention numerous other sources of radio noise.

An eavesdropper of compromising emanations, therefore, must selectively amplify only those parts of the radio spectrum that provide the best signal-to-noise ratio. Unlike radio transmissions, most compromising RF emanations are baseband signals, that is, they are not modulated with a carrier frequency to shift them into a narrow and reserved frequency slot of the radio spectrum. However, digital signals consist of discrete symbols (bits, pixels, etc.) transmitted at some rate f. From the sampling theorem we know that the frequency spectrum up to $f/2$ contains already all information carried by the signal. If the individual symbols have spectral energy beyond that frequency, for example because they contain sharp edges with a raise time much shorter than the bit or pixel duration, then the information in the signal will be repeated in several $f/2$ wide bands at higher harmonic frequencies. It is therefore sufficient for an eavesdropper to find any frequency range with good signal-to-noise ratio that is merely at least half as wide as the bit or pixel rate.

The frequency range with the best signal-to-noise ratio depends equally on the targeted device and on the background noise, both of which can vary significantly with the device, video mode and location. Building good analog bandpass RF filters that can be adjusted over a wide range of frequencies is not easy. A more practical approach than direct filtering is the use of a superheterodyne AM receiver that multiplies the input signal with a sine wave of adjustable frequency to shift the frequency band of interest to a fixed intermediate frequency where it can then be filtered easily to the required bandwidth. The subsequent rectification and low-pass filtering in the AM demodulator will destroy some phase information and with it valuable information, such as the difference between positive and negative edges in the eavesdropped signal. But it will also lead to a much lower frequency signal that can be digitized comfortably with a sampling rate of not much more than twice the bandwidth.

The particular receiver used to acquire the example images shown in this paper was a *Dynamic Sciences R1250*, an instrument that was specifically designed to meet the (confidential) requirements of the “TEMPEST” measurement standard NACSIM 5100A. Its center frequency can be tuned from 100 Hz to 1 GHz and it offers intermediate-frequency (IF) filters with bandwidths ranging from 50 Hz to 200 MHz. The length of the shortest impulse that can be recognized at a receiver output is the inverse of the IF filter bandwidth, which therefore has to be comparable to the pixel clock frequency of modern displays. Most other commercially available AM radio receivers (including TV tuners) are not designed for bandwidths larger than about 8 MHz. Another important feature of the R1250 is that its automatic gain control can be disabled. This makes it possible to compare the amplitude of any input signal with that of a reference sine-wave generator. This way, it was possible to provide an antenna input voltage scale for all the received video images shown here. The output of the AM receiver was for adjustment purposes displayed in real-time on a normal computer monitor, whose sync lines were driven by a programmable arbitrary-waveform generator, to reproduce the line and frame rate of the targeted display. Special care was necessary to set up the sync-pulse generators such that the refresh rate they generated was adjustable to match that of the targeted display with less than 10^{-7} relative error, which is smaller than the stability and sometimes even resolution of many standard function generators.

The images shown in this paper were recorded with a digital storage oscilloscope (8-bit resolution, 16 MB acquisition memory, up to 1 GHz sampling frequency) directly from the output of the AM demodulator and converted with specially written software into raster images. The antenna used was a log-periodical broadband antenna designed for a frequency range of 200–1000 MHz, as it is commonly used for electromagnetic compatibility measurements. All recordings were performed without any shielding in a normal modern office building in a semi-urban environment with over a hundred other computers operating in the same building. Further details about the instrumentation are given in [18].

4 Case Study: Laptop Display

Figure 1 shows an amplitude-demodulated and rastered signal as it was received from the first example target, a Toshiba Satellite Pro 440CDX laptop that shows a Linux boot screen in an 800×600@75Hz video mode. The antenna was located at 3 m distance in the same room as the target device. A quick scan through different frequencies in the 50–1000 MHz range showed that setting the AM receiver to a center frequency of 350 MHz and an intermediate-frequency bandwidth of 50 MHz gave one of the clearest signals. The image shown is the average of 16 recorded frames, in order to reduce noise. For comparison, the lower right corner shows one of these frames without any averaging. Even there, readable text stands out clearly from the background noise. The frames were recorded with a sampling frequency of 250 MHz.

A number of observations distinguish the signal seen Fig. 1 from those typical for CRTs:

- The low-frequency components of the video signal are not attenuated. Horizontal bright lines appear in the reconstructed signal as horizontal lines and not just as a pair of switching pulses at the end points, as would be the case with CRTs.
- Font glyphs appear to have lost half of their horizontal resolution, but are still readable.
- In the 800×600@75Hz video mode used, the clearest signal can be obtained at a center frequency of about 350 MHz with 50 MHz bandwidth, but weaker signals are also present at higher and lower frequencies, in particular after every step of 25 MHz.
- The mapping between displayed colors and the amplitude of the signal received for a pixel turned out to be highly non-monotonic. A simply gray-bar image resulted in a complex barcode like display, as if the generated signal

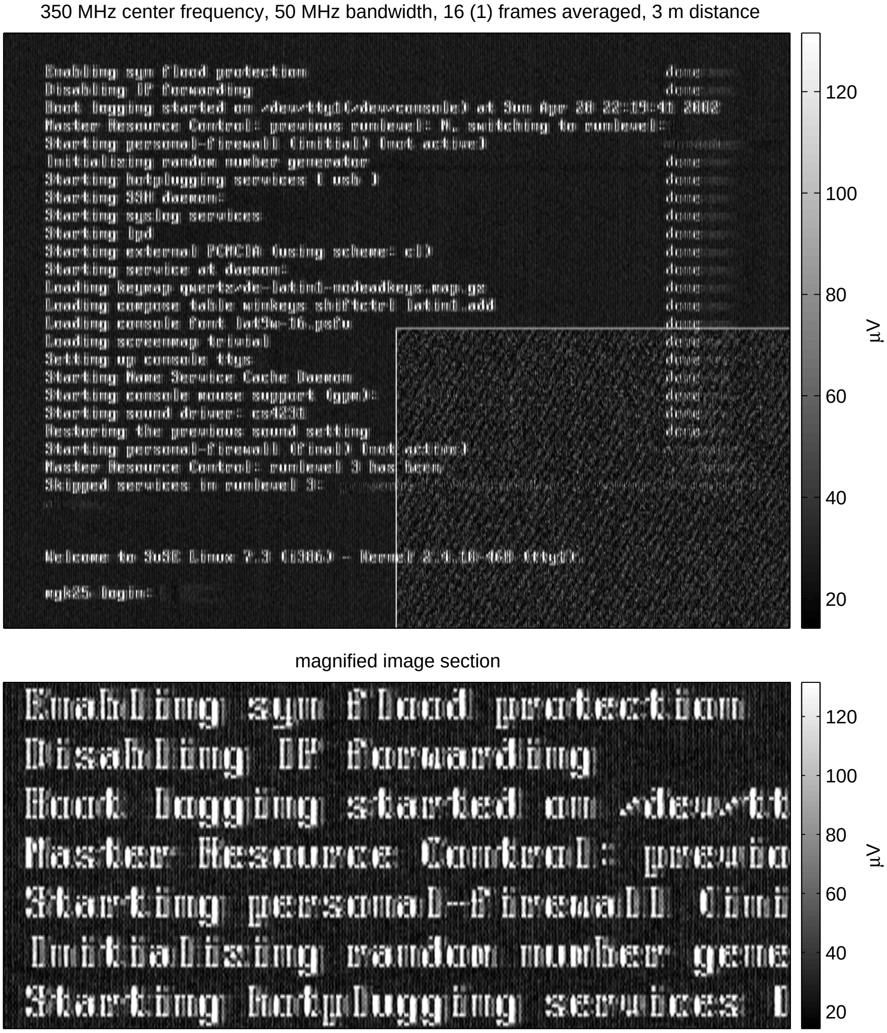
350 MHz center frequency, 50 MHz bandwidth, 16 (1) frames averaged, 3 m distance
120
100
80
60
40
20
µV
magnified image section
120
100
80
60
40
20
µV

```
Enabling syn flood protection                                         done
Disabling IP forwarding                                               done
Boot logging started on /dev/tty10(/dev/console) at Sun Apr 28 22:19:41 2002
Master Resource Control: previous runlevel: N, switching to runlevel:
Starting personal-firewall (initial) [not active]                     unused
Initializing random number generator                                  done
Starting hotplugging services [ usb ]                                 done
Starting SSH daemon:                                                  done
Starting syslog services                                              done
Starting lpd                                                          done
Starting external PCMCIA (using scheme: c1)                           done
Starting service at daemon:                                           done
Loading keymap qwertz/de-latin1-nodeadkeys.map.gz                     done
Loading compose table winkeys shiftctrl latin1.add                    done
Loading console font lat9w-16.psfu                                    done
Loading screenmap trivial                                             done
Setting up console ttys                                               done
Starting Name Service Cache Daemon                                    done
Starting console mouse support (gpm):                                 done
Starting sound driver: cs4232                                         done
Restoring the previous sound setting                                  done
Starting personal-firewall (final) [not active]
Master Resource Control: runlevel 3 has been
Skipped services in runlevel 3:


Welcome to SuSE Linux 7.3 (i386) - Kernel 2.4.10-4GB (tty2).

mgk25 login:
```

```
Enabling syn flood protection
Disabling IP forwarding
Boot logging started on /dev/tt
Master Resource Control: previo
Starting personal-firewall (ini
Initializing random number gene
Starting hotplugging services [
```

Fig. 1. Eavesdropped Linux boot screen visible on the LCD of a Toshiba 440CDX laptop (log-periodic antenna, vertical polarization).

amplitude were somehow related to the binary representation of the pixel value.

- Using a simple improvised near-field probe (a coaxial cable whose ends are shaped into a 50 mm dipole) instead of an antenna, to scan the immediate vicinity of the laptop, it became clear that no significant emissions came from the display module itself, but that the source appeared to be the interconnect cable between the LCD module and the mainboard.

A closer examination of the laptop reveals a digital video link as the origin of these emanations. The display module (Sharp LM12S029 FSTN) used in this laptop is connected to the video controller via eight twisted pairs, each about 30 cm long. They originate on the mainboard in two integrated parallel-to-serial converters and LVDS transmitter chips designed for linking to flat-panel displays (NEC DS90CF581 [12]). The 18-bit color data that the video controller provides for each pixel on its parallel output port has to be serialized into fewer lines, to fit through the hinges, which is exactly the task that these two "FPD-Link" chips perform. They multiply the clock signal supplied from the video controller by seven, and each transmits per clock cycle on three twisted-pair channels $3 \times 7 = 21$ data bits, which consist here of 18 data bits for the pixel color and three bits for horizontal sync, vertical sync and a control signal. The fourth pair carries the clock.

The video controller outputs 50 million pixels per second. However, since it transmits the data for two consecutive pixels simultaneously over two independently operating FPD-Link chips, each of these receives a clock frequency of only 25 MHz, which it multiplies to a data rate of 175 MHz, resulting in an overall data rate of 1.05 Gbit/s transmitted on all six channels through the hinges.

LVDS (low voltage differential signaling [13]) is a generic interface standard for high-speed data transmission (up to 655 Mbit/s). It uses symmetric twisted transmission lines and was designed to minimize RF interference.

However, as Fig. 1 shows, such precautions are not sufficient for emission security. The approximately 100 μV amplitude that the log-periodic antenna receives for the BIOS default colors used in this screen at 3 m distance corresponds to a field strength of 57 dBμV/m (50 MHz bandwidth) and an equivalent isotropic radiating power would be about 150 nW.

A signal of this amplitude is strong enough to permit a simple and realistic eavesdropping demonstration across several rooms. In the next experiment, the same laptop and antenna are located about 10 m apart in different office rooms, separated by two other offices and three 105 mm thick plaster-board walls.

In this setup 12 consecutive frames were acquired with a sampling rate of 50 MHz in one single recording of 160 ms (eight million samples). The exact frame rate necessary for correctly aligned averaging was determined with the necessary precision of at least seven digits from the exact distance of the first and last of the recorded frames. It was determined with an algorithm that calculated starting from a crude estimate of the frame rate the cross-correlation of these two frames, and then corrected the estimate based on the position of the largest peak found there (Fig. 2). (The process is not fully automatic, as due to other video signals in the vicinity, echos, and multiple peaks, it can sometimes be necessary to manually chose an alternative peak.)

Figure 3 shows the result, an easily readable view of an `xterm` window that shows some test text. The received signal amplitude of about 12 μV corresponds with this antenna to a field strength of 39 dBμV/m. This drop by 18 dB compared to the 57 dBμV/m in the previous 3 m line-of-sight measurement can in part be attributed to the 10 dB free-space loss to be expected when tripling the

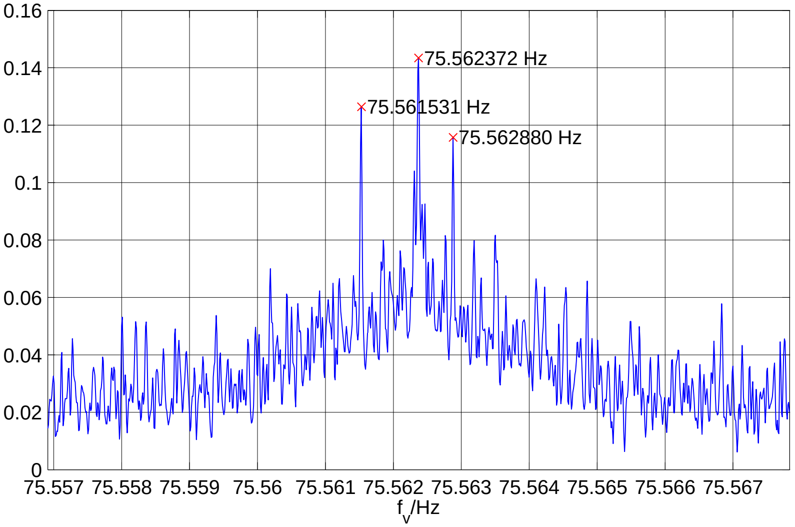

Fig. 2. Determination of the frame rate f_{v} for the multi-frame signal recorded in Fig. 3 through crosscorrelation between the first and last frame in the recorded series.

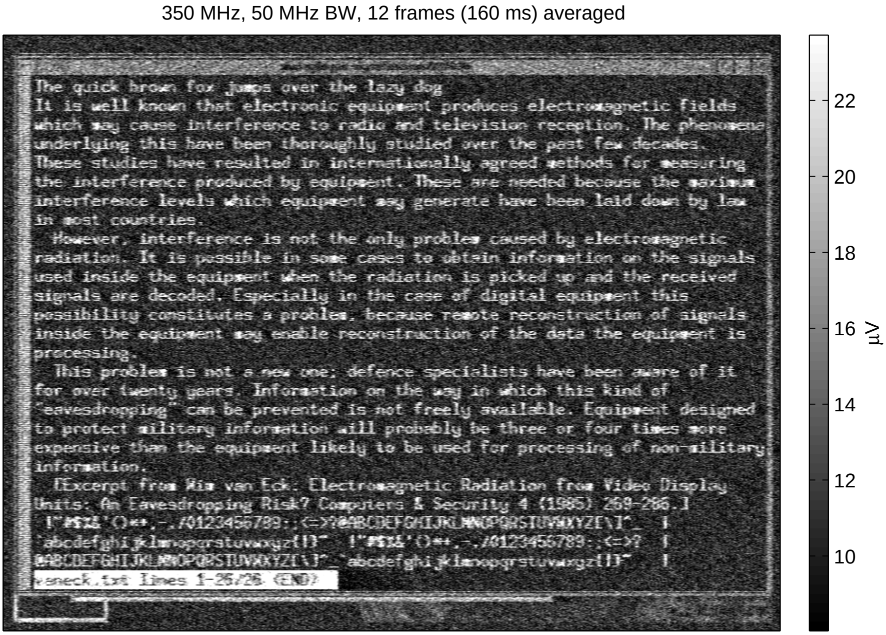
350 MHz, 50 MHz BW, 12 frames (160 ms) averaged

```
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)
```

Fig. 3. Text signal received from a 440CDX laptop at 10 m distance through two intermediate offices (3 plasterboard walls).

distance between emitter and antenna. The remaining drop suggests that each of the plasterboard walls contributes 2–3 dB additional attenuation, which appears to be a typical value, judging from the UHF building-material attenuation values described in the literature [14].

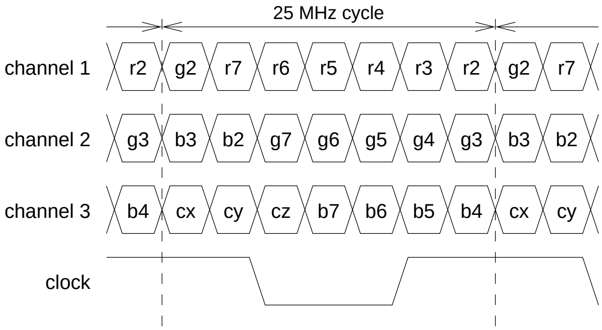

Fig. 4. Bit assignment in the FPD-Link transmission cycle.

In order to better understand the relationship between the signal displayed on the target device and that seen on the rastered output of an AM receiver, it is worth having a closer look at the exact transmission format. The details are very specific to the particular product targeted here, but the principles explained can easily be transferred to similar designs. Application software typically provides the display driver with 24-bit color descriptions of the form $(r_7 \ldots r_0, g_7 \ldots g_0, b_7 \ldots b_0)$. Figure 4 shows, how these bits are packed in a 440CDX laptop into the pixel cycle of three FPD-Link channels[1]. One of the FPD-Link chips transmits all pixels in odd-numbered columns, the other one the pixels in even-numbered columns.

Armed with an understanding of what choice of colors elicits which waveform from the channel drivers, we can now experiment with various combinations, in particular those that promise to maximize or minimize the contrast between the foreground and background of text in the emitted signal.

Figure 5 shows a test text in various color combinations, together with the corresponding RGB values specified by the application program and the resulting bit patterns on the three transmission channels. Line 1 is simply the black-on-white combination commonly used in word processing software. Line 2 is an attempt to find the signal with the largest number of bit transitions in the foreground and the smallest number in the background, in order to maximize

[1] Being an 18-bit per pixel interface, the two least significant bits of each byte are not represented. A further restriction is that the video memory of this laptop supports the 800×600@75Hz video mode only with a 16 bits per pixel encoding (5 red, 6 green, 5 blue), in which the video controller hardware fills in the values $r_2 = r_7 \wedge \ldots \wedge r_3$ and $b_2 = b_7 \wedge \ldots \wedge b_3$ automatically.

line	description	foreground		background	
		RGB	signal	RGB	signal
1	black on white	00 00 00	000000x 0x00000 xxx0000	ff ff ff	111111X 1X11111 xxx1111
2	maximum contrast	a8 50 a0	010101x 0x01010 xxx1010	00 00 00	000000x 0x00000 xxx0000
3	maximum contrast (gray)	a8 a8 a8	010101x 1x10101 xxx1010	00 00 00	000000x 0x00000 xxx0000
4	minimum contrast	78 00 00	001111x 0x00000 xxx0000	00 f0 00	000000x 0x11110 xxx0000
5	minimum contrast	78 60 00	001111x 0x01100 xxx0000	30 f0 00	000110x 0x11110 xxx0000
6	minimum contrast (phase shift)	70 70 00	001110x 0x01110 xxx0000	38 e0 00	000111x 0x11100 xxx0000
7	text in most significant bit, rest random	—	r1rrrrx rx1rrrr xxx1rrr	—	r0rrrrx rx0rrrr xxx0rrr
8	text in green two msb, rest random	—	rrrrrrx rx11rrr xxxrrrr	—	rrrrrrx rx00rrr xxxrrrr
9	text in green msb, rest random	—	rrrrrrx rx1rrrr xxxrrrr	—	rrrrrrx rx0rrrr xxxrrrr

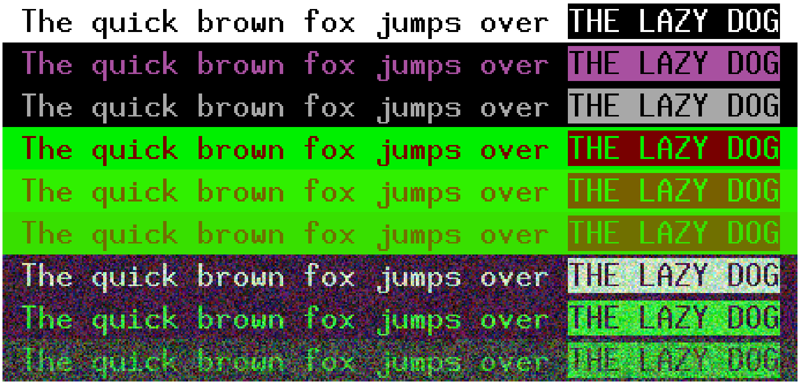

The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG

Fig. 5. Test text to compare the emission characteristics of selected foreground and background color combinations.

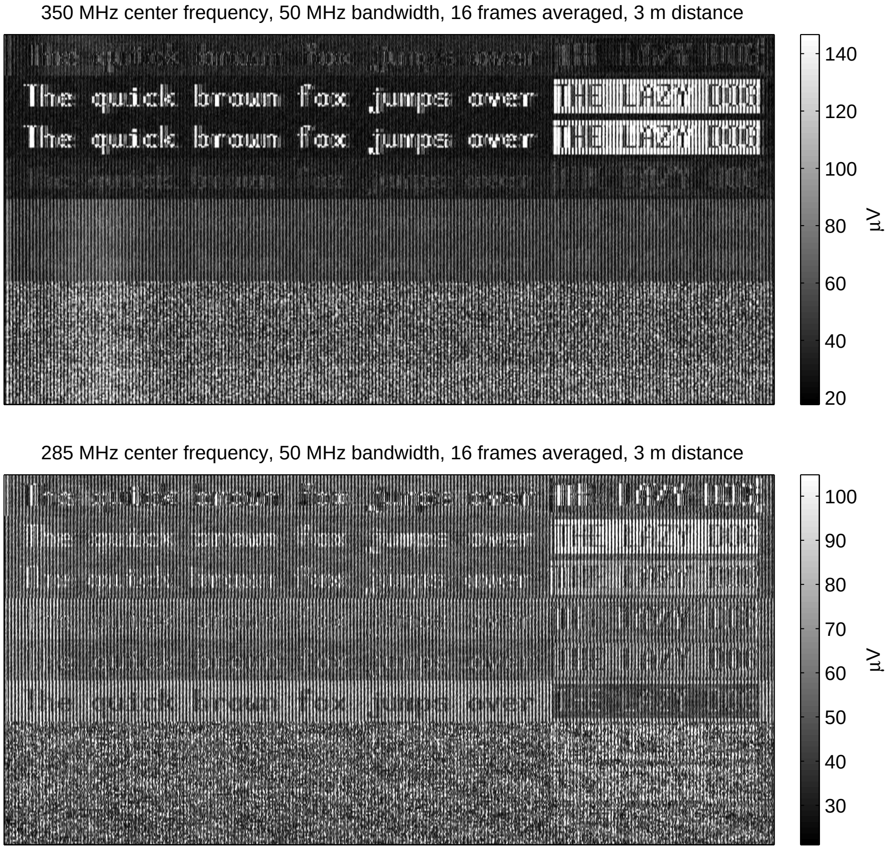
350 MHz center frequency, 50 MHz bandwidth, 16 frames averaged, 3 m distance

285 MHz center frequency, 50 MHz bandwidth, 16 frames averaged, 3 m distance

Fig. 6. Signals received from the test display in Fig. 5.

contrast and readability for the eavesdropper. Line 3 attempts the same, but maximizes the visible contrast in favor of having identical signal polarity on all three lines for the foreground pixels. (In a symmetric transmission channel, signal polarity should in principle not make a difference for an eavesdropper.)

Line 4 is a first attempt to find a combination of two colors whose radio signature is difficult to distinguish under the assumption that the eavesdropper can evaluate only the total number of bit transitions that happen on all channels together. The idea is to let bit transitions always happen at the same time during the cycle, but in different channels.

Line 5 is a variant that keeps even the total number of transitions in each line constant and line 6 keeps in addition the length of positive pulses constant and encodes the difference between foreground and background color only as a one-bit phase shift in two of the channels.

The last three lines finally demonstrate what happens if most of the bits are filled randomly, in order to jam the eavesdropper's periodic averaging process with a meaningless signal of exactly the same period. This jamming should be particularly effective if the neighbor bits of each data carrying bit are selected randomly, as this will randomize whether the data bit will contribute a transition pulse to the compromising emanations or not.

Figure 6 shows the signal received with 50 MHz bandwidth at two frequencies. The first at 350 MHz is the one where the maximum-contrast colors in lines 2 and 3 offer the strongest signal. They result in a $175/2 = 87.5$ MHz square wave, but this particular frequency and its first harmonic collide in Cambridge with signals from local broadcasting stations. 350 MHz is one of the first harmonics in a quieter band, and a $\lambda/4$ monopole for that frequency is with 40 cm also quite close to the length of the twisted pair, leading to more efficient far-field emissions. In this band, the maximum bit-transition patterns in lines 2 and 3 generate field levels of 59 dBμV/m at 3 m (240 nW EIRP). The black-on-white text in line 1 causes a significantly weaker signal, because only a single bit transition is generated in each channel by a transition between full black and white levels (except for the blue channel which also contains control bits).

The first attempt at finding a protective color combination in line four is not fully effective, which suggests that edges in different transmission lines cause noticeably different electromagnetic pulses and can therefore be distinguished. This could be caused either by tolerances in LVDS driver parameters or by impedance differences between conductor pairs. Lines 5 and 6, which use a constant number of bit transitions in each channel and vary only their relative phases, provide the eavesdropper at this frequency band practically no usable contrast, as do all the test lines in which random-bit jamming is applied.

Even though a 50 MHz wide band captures enough information to resolve horizontally pixel pairs accurately, it does not quite cover the entire $175/2 = 87.5$ MHz wide spectrum that contains (according to the sampling theorem) the full information present in the 175 Mbit/s bitstream. Tuning to a different center frequency provides a different extract of the entire signal to the demodulator, effectively applying a different filter to the video signal. The bottom half of Fig. 6 shows one center frequency (285 MHz), where the low-contrast color combinations suddenly become readable.

We can conclude that the only effective software protection technique against compromising emanations of FPD-Links, as used in numerous laptops, appears to be the addition of random bits to the color combinations used for text display. When implementing such a technique, it is critical to understand that these random bits must be randomly selected each time a new character is placed on the screen.

If the random bits were selected, for example, in a glyph rendering routine that is connected to a glyph cache, to ensure that an already generated bitmap is reused whenever the same character is used multiple times on the screen, then this merely assists the eavesdropper. If the addition of random bits were done identically at each location where a glyph is used, then the random bits merely increased the values in a glyph-signal distance matrix, which would only reduce the error probability during automatic radio character recognition.

5 Case Study: Digital Visual Interface

The NEC FPD-Link interface technology appears to be mainly used in embedded display systems, such as laptops. For connecting flat-panel displays to desktop computers, three other interface standards that define connector plugs have been defined: VESA Plug & Display (P&D) [15], VESA Digital Flat Panel (DFP) [16], and Digital Visual Interface (DVI) [17].

These three standard connectors differ only in auxiliary interfaces (USB, IEEE 1394, VGA) that are carried on the same cable, but that are not relevant here. All three interfaces use, in mutually compatible ways, a digital video transmission technology called *Transition Minimized Differential Signaling (TMDS)*, also known as *PanelLink*, developed by Silicon Image Inc.

A TMDS link consists of three channels, similar to the FPD-Link system described in the previous section. Each is formed by a symmetric twisted-line pair and carries 8-bit values for one of the three primary colors. A fourth twisted-pair channel provides a byte clock for synchronization.

What distinguishes TMDS most from FPD-Link is the encoding used. Each 8-bit value transmitted over a channel is first expanded into a 10-bit word. The encoding process consists of two steps, each of which has one of two options to change the eight data bits, and each signals its choice to the receiver by appending another bit.

In the first step, the number of "one" bits in the 8-bit data value $d_7d_6\ldots d_0$ is counted. A new 9-bit value q is generated by setting $q_0 = d_0$ and

$$\begin{aligned} q_i &= q_{i-1} \oplus d_i \qquad \text{for } 1 \le i \le 7 \\ q_8 &= 1 \end{aligned}$$

if there are more zeros in d ($\oplus$ is *exclusive or*), and

$$\begin{aligned} q_i &= \neg\, q_{i-1} \oplus d_i \qquad \text{for } 1 \le i \le 7 \\ q_8 &= 0 \end{aligned}$$

if there are more ones in d. In case of four zeros and ones each, only d_0 is counted.

In the second step, either the bits $q_7q_6\ldots q_0$ are all inverted and $q_9 = 1$ is added, or all bits remain as they are and $q_9 = 0$ is added instead. The decision is made by taking into account how many "zero" and "one" bits have been transmitted so far and the choice is made that leads to a more equal count.

The first step aims at reducing the maximum number of bit transitions that can occur per value on the channel, as the following examples illustrate (d_0 and q_0 are at the right end, respectively):

$$\begin{aligned} 10101010 &\longrightarrow 0\,11001100, \qquad & 01010101 &\longrightarrow 1\,00110011, \\ 00000000 &\longrightarrow 1\,00000000, \qquad & 11111111 &\longrightarrow 0\,11111111. \end{aligned}$$

While an 8-bit word can contain up to eight bit transitions, after this recoding, only a maximum of five transitions is possible in any of the resulting 9-bit words (including one transition between consecutive words). This is, because the less frequent bit can only appear up to four times in a byte, and each presence of it is signalled by a transition in the generated 9-bit word.

The purpose of the second step is to limit the difference between the total number of “zero” and “one” bits. This keeps the signaling scheme DC balanced, which simplifies the use of transformers for galvanic separation of transmitter and receiver. For an exact description of the encoding algorithm see [17, p. 29].

The following examples show how in the full encoding the DC-balancing mechanism adds longer repetition cycles to sequences of identical bytes. The binary words are this time shown in Littleendian order (q_0 and d_0 at the left end), in order to match transmission order, which is least significant bit first. For example, encoding a sequence of zero bytes leads to a cycle of nine 10-bit words, whereas for the byte 255, the cycle length is only seven:

$$\begin{array}{l}
00000000, 00000000, 00000000, 00000000, 00000000, \ldots \longrightarrow \\
\quad 0000000010, 1111111111, 0000000010, 1111111111, 0000000010 \\
\quad 1111111111, 0000000010, 1111111111, 0000000010, \\
\quad 0000000010, 1111111111, 0000000010, 1111111111, 0000000010 \\
\quad 1111111111, 0000000010, 1111111111, 0000000010, \\
\quad \ldots \\
11111111, 11111111, 11111111, 11111111, 11111111, \ldots \longrightarrow \\
\quad 0000000001, 1111111100, 1111111100, 0000000001, 1111111100 \\
\quad 0000000001, 1111111100 \\
\quad 0000000001, 1111111100, 1111111100, 0000000001, 1111111100 \\
\quad 0000000001, 1111111100 \\
\quad \ldots
\end{array}$$

To find a color combination that provides the best possible eavesdropping reception of TMDS encoded video signals, we can try to look for one with as many bit transitions as possible in one color and as few as possible in the other. A second consideration is that the extended cycles added by the DC-balancing algorithm might reduce readability and that it is therefore desirable to find maximum contrast bytes with a cycle length of one. This can only be achieved if the resulting 10-bit words do not affect the difference in the bit balance counter maintained by the DC-balancing algorithm. In other words, the 10-bit words selected should contain exactly five “one” bits, and there exist 52 byte values that will be encoded in such a DC balanced TMDS word.

For example, the bytes hexadecimal `10` and `55` fulfil these criteria:

$$\begin{array}{l}
00001000, 00001000, \ldots \longrightarrow 0000111110, 0000111110, \ldots \\
10101010, 10101010, \ldots \longrightarrow 1100110010, 1100110010, \ldots
\end{array}$$

These TMDS bit patterns will be used irrespective of the previous bit balance, because the full encoding algorithm specified in [17, p. 29] contains a special case.

It sets $q_9 = \neg q_8$ whenever the rest of q contains exactly four "zero" and four "one" bits, which is the case here. The encoding of any pixels encoded with one of the 52 balanced words will therefore remain unaffected by any other screen content.

		foreground	background
line	description	RGB	RGB
1	black on white	`00 00 00`	`ff ff ff`
2	maximum bit transition contrast	`00 00 00`	`aa aa aa`
3	half bit transition contrast	`00 00 00`	`cc cc cc`
4	balanced word, max contrast	`10 10 10`	`55 55 55`
5	minimum signal contrast	`ff 00 00`	`00 ff 00`
6	low nybble random	`0r 0r 0r`	`fr fr fr`
7	text in msb, rest random	—	—
8	text in green two msb, rest random	—	—
9	text in green msb, rest random	—	—

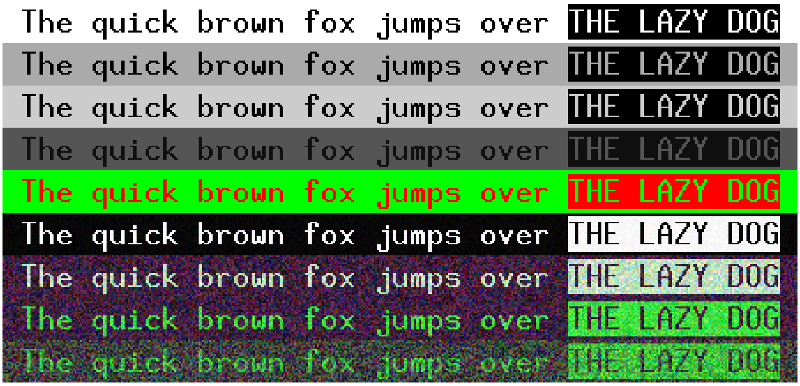
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG
The quick brown fox jumps over THE LAZY DOG

Fig. 7. Test image for text contrast in compromising emanations from DVI cables.

Figure 7 shows a number of different foreground/background color combinations, including the black-on-white text in line 1 and two naïve approaches to obtain maximum reception contrast in lines 2 and 3. The color combination for high-contrast reception just suggested is used in line 4, and the rest represents a number of attempts to find minimum contrast signals and to add random bits for jamming.

Figure 8 shows the signals received from a DVI display system that shows the test display of Fig. 7. The graphics card in this setup was an "ATI Rage Fury Pro" and the display a "Samsung SyncMaster 170T". The 1280×1024@60Hz video mode used in this setup has a pixel clock frequency of 108 MHz.

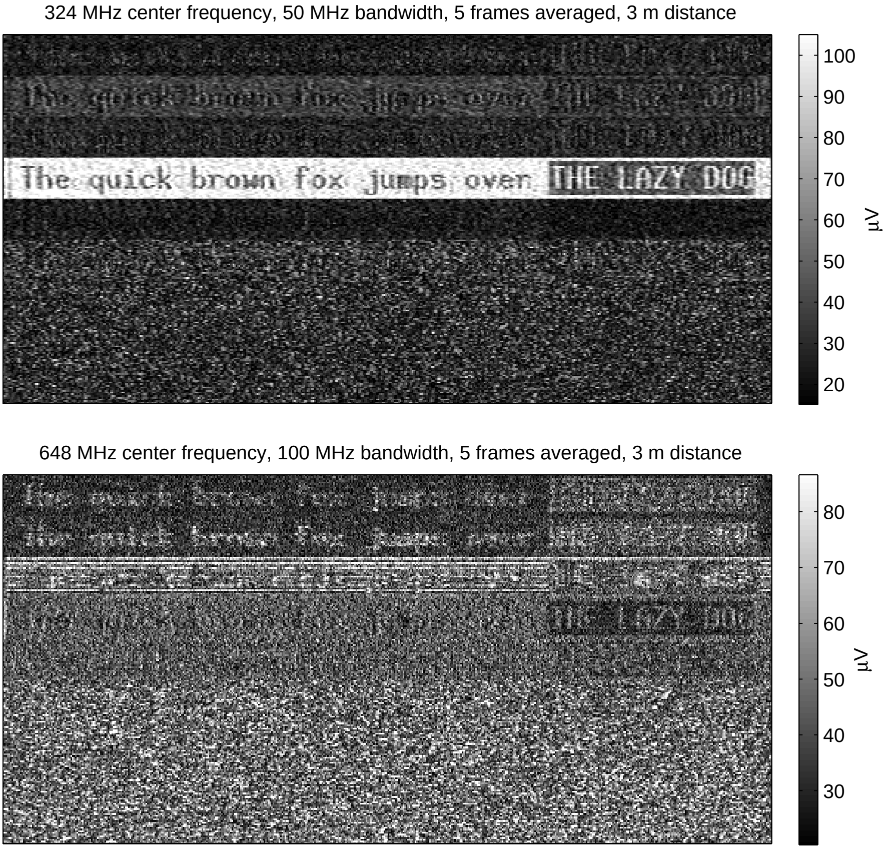

Fig. 8. Received emanation in two frequency bands from a DVI cable transmitting the text image of Fig. 7.

While an excellent signal can be obtained with the `55/10` color combination, other color combinations, including black/white are either considerably weaker or provide a noticeable reception contrast only on a different frequency. The transitions and DC-balancing cycles added by the TMDS encoding are not sufficient to make emanations from DVI cables entirely unreadable, but the signal quality is noticeably degraded compared to simpler transmission formats. In particular, thanks to the TMDS encoding, a much smaller number of least-significant random bits added for jamming already is sufficient to eliminate even weakest traces of the displayed text in the received signal.

An additional property of the TMDS encoding that might be of use for a radio-frequency eavesdropper is that during blanking intervals, four special 10-bit words 0010101011, 1101010100, 0010101010 and 1101010101 represent the four possible combinations of the horizontal and vertical sync signals. These words contain eight bit transitions each and can this way be distinguished from any normal color.

It might be worth noting that the DVI standard is prepared for two optional extensions that, even though not intended for this purpose, might also be of use for reducing emanation security concerns. The first is *selective refresh*, a mode of operation in which the display has its own frame buffer and refreshes the display with the desired frequency, without overloading the transmission capacity of the DVI link. The DVI link can then operate at a lower speed and might even become active only when data in the display's frame buffer needs to be updated. The absence of a continuous periodic signal would be likely to make radio-frequency eavesdropping on the interface cable impractical.

The second option under development is *High-bandwidth Digital Content Protection (DVI/HDCP)*, an encryption and key negotiation layer designed to be used over the DVI interface between digital video players and television sets. Intended to prevent unauthorized copying of uncompressed video signals by placing the decryption step into the display device, it would also render interface cable emanations unreadable.

Even a cryptographically weak key exchange protocol is likely to provide sufficient protection against a passive compromising-emanations eavesdropper, who can see the communication only in a noisy and restricted form. In the presence of significant noise, a computationally secure key negotiation scheme can be built using simple anti-redundancy techniques. One party sends out a several thousand bits long random string R. Both sides then use a hash $h(R)$ as the session key to encrypt the remaining communication. Even a moderate amount of bit errors in an eavesdropped copy of R will make it computationally infeasible to find from that the key $h(R)$.

6 Conclusions

The eavesdropping risk of flat-panel displays connected via Gbit/s digital interfaces to their video controller is at least comparable to that of CRTs. Their serial transmission formats effectively modulate the video signal in ways which provide eavesdroppers with even better reception quality. A detailed understanding of the encoding algorithms and bit arrangement used in digital video links allows programmers fine-grained control over the emitted signal. In a simple serial transmission system, like NEC's FPD-Link, the strongest signal can be obtained by choosing colors that result in alternating bits on the transmission line. In interfaces involving TMDS encoding, only a careful analysis of the encoding algorithm leads to a maximum contrast color combination. Using colors that result in bit-balanced code words prevents a state change in the encoder. This avoids distortions to the transmitted signal and can be used to improve the quality of intentional emissions. Combinations of foreground and background colors can be selected to reduce the readability of text in the compromising emanations. Much better protection can be achieved by randomizing the less-significant bits of the transmitted RGB values. This emits a jamming signal that cannot be eliminated via periodic averaging, because it has exactly the same period as the text signal.

References

1. Harold Joseph Highland: Electromagnetic Radiation Revisited. Computers & Security, Vol. 5, pp. 85–93 and 181–184, 1986.
2. Wim van Eck: Electromagnetic Radiation from Video Display Units: An Eavesdropping Risk? Computers & Security, Vol. 4, pp. 269–286, 1985.
3. Anton Kohling: TEMPEST – eine Einführung und Übersicht zu kompromittierenden Aussendungen, einem Teilaspekt der Informationssicherheit [TEMPEST – an introduction and overview on compromising emanations, one aspect of information security]. In H.R. Schmeer (ed.): Elektromagnetische Verträglichkeit/EMV'92, Stuttgart, February 1992, pp. 97–104, VDE-Verlag, Berlin, ISBN 3-8007-1808-1.
4. Erhard Möller, Lutz Bernstein, Ferdinand Kolberg: Schutzmaßnahmen gegen kompromittierende elektromagnetische Emissionen von Bildschirmsichtgeräten [Protective measures against compromising electromagnetic emissions of video displays]. 1. Internationale Fachmesse und Kongreß für Datensicherheit (Datasafe '90), Karlsruhe, Germany, November 1990.
5. Gerd Schmidt, Michael Festerling: Entstehung, Nachweis und Vermeidung kompromittierender Strahlung [Origin, detection and avoidance of compromising radiation]. MessComp '92, 6. Kongreßmesse für die industrielle Meßtechnik, Wiesbaden, 7–9 September 1992.
6. Sicurezza Elettromagnetica nella Protezione dell'Informazione, ATTI SEPI'88, Rome, Italy, 24–25 November 1988, Fondazione Ugo Bordoni.
7. Symposium on Electromagnetic Security for Information Protection, SEPI'91, Proceedings, Rome, Italy, 21–22 November 1991, Fondazione Ugo Bordoni.
8. Markus G. Kuhn, Ross J. Anderson: Soft Tempest: Hidden Data Transmission Using Electromagnetic Emanations. Information Hiding, IH'98, Portland, Oregon, 15–17 April 1998, Proceedings, LNCS 1525, Springer-Verlag, pp. 124–142.
9. Information technology equipment – Radio disturbance characteristics – Limits and methods of measurement. CISPR 22, International Electrotechnical Commission (IEC), Geneva, 1997.
10. TCO'99 – Mandatory and recommended requirements for CRT-type Visual Display Units (VDUs). Swedish Confederation of Professional Employees (TCO), 1999. `http://www.tcodevelopment.com/`
11. 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.
12. LVDS Transmitter 24-Bit Color Flat Panel Display (FPD) Link, National Semiconductor Cooperation, 1998. `http://www.national.com/pf/DS/DS90CF581.html`
13. Electrical characteristics of low voltage differential signaling (LVDS) interface circuits, ANSI/TIA/EIA-644, Electronic Industries Alliance, 1996.
14. Homayoun Hashemi: The Indoor Radio Propagation Channel. Proceedings of the IEEE, Vol. 81, No. 7, July 1993, pp. 943–968.
15. VESA Plug and Display Standard. Version 1, Video Electronics Standards Association, 11 June 1997.
16. VESA Digital Flat Panel (DFP). Version 1, Video Electronics Standards Association, 14 February 1999.

17. Digital Visual Interface – DVI. Revision 1.0, Digital Display Working Group, April 1999. `http://www.ddwg.org/`
18. Markus G. Kuhn: Compromising emanations: eavesdropping risks of computer displays. Technical Report UCAM-CL-TR-577, University of Cambridge, Computer Laboratory, December 2003.

A Spectral Analysis of TMDS Signals

Fourier theory and the convolution theorem can be used to explain the spectral composition of the signal on a TMDS channel in the example from Sect. 5. Let the function t_{55} denote the waveform that we obtain if we repeat the 10-bit word representing the byte value hexadecimal `55` with 108 MHz. The Fourier transform $\mathcal{F}\{t_{\mathsf{55}}\}$ is a line spectrum with lines at 108 MHz, 216 MHz, 324 MHz, . . . , 972 MHz. Let v be a binary video signal with a pixel frequency of 108 MHz, which equals 1 during bright pixels and 0 while a dark pixel is transmitted. So if we transmit bright pixels as the value `55` and dark pixels as a value `10`, the resulting waveform is

$$w = v \cdot t_{\mathsf{55}} + (1 - v) \cdot t_{\mathsf{10}} = v \cdot (t_{\mathsf{55}} - t_{\mathsf{10}}) + t_{\mathsf{10}} \ . \tag{1}$$

Multiplication in the time domain corresponds to convolution in the frequency domain, hence we end up for the waveform transmitted on the TMDS channel with the spectrum

$$W = V * \mathcal{F}\{t_{\mathsf{55}} - t_{\mathsf{10}}\} + \mathcal{F}\{t_{\mathsf{10}}\} \ . \tag{2}$$

In other words, the spectrum of the pixel-value waveform V will be copied in W centered around each of the spectral lines of the Fourier transform of the difference between the two data words. The signal intensity of the various frequency-shifted incarnations of V depends on the amplitude of the respective spectral lines of $\mathcal{F}\{t_{\mathsf{55}} - t_{\mathsf{10}}\}$. Figure 9 illustrates the relative intensity of the spectral lines of $|\mathcal{F}\{t_{\mathsf{10}}\}|$, $|\mathcal{F}\{t_{\mathsf{55}}\}|$, and $|\mathcal{F}\{t_{\mathsf{55}} - t_{\mathsf{10}}\}|$. It also shows the line spectrum $|\mathcal{F}\{t_{\mathsf{55}}\}| - |\mathcal{F}\{t_{\mathsf{10}}\}|$, which better approximates the contrast that an AM demodulating receiver can see, as it discards phase information received. Since w is a discretely sampled waveform, its spectrum will be copied at all multiples of the sampling frequency (1.08 GHz here), attenuated by the spectrum of a single bit pulse.

The center frequency of 324 MHz used in Figure 8 is not the strongest line in the spectrum of $|\mathcal{F}\{t_{\mathsf{55}}\}| - |\mathcal{F}\{t_{\mathsf{10}}\}|$, but it was the strongest located in a quieter part of the background-noise spectrum during this measurement. It still results in a signal strength in the order of 100 µV at the receiver input, comparable to what was measured earlier in Sect. 4 for the laptop.

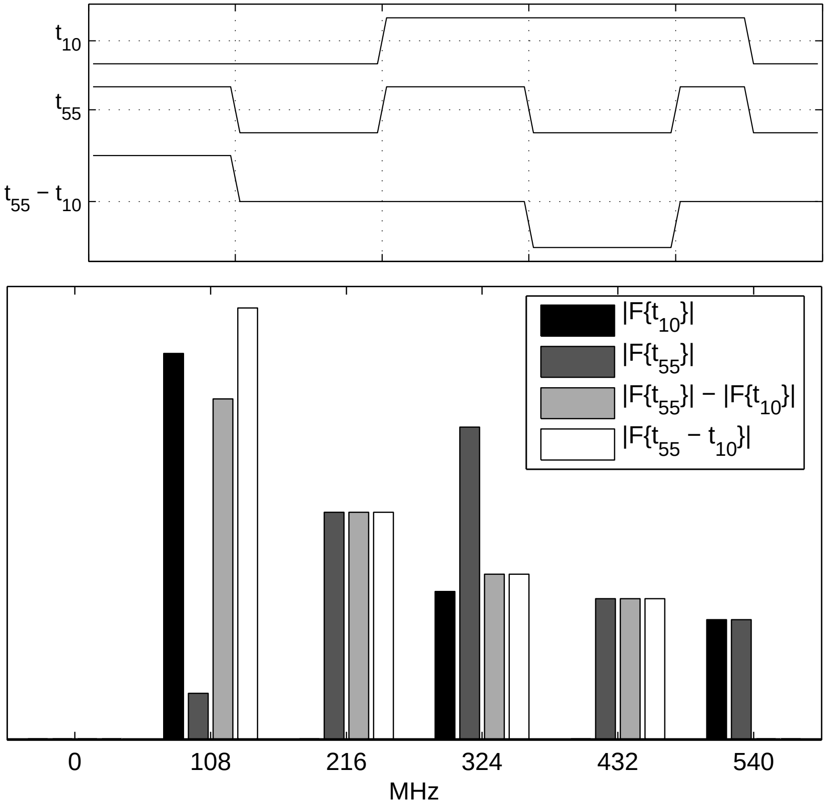

Fig. 9. Time and frequency domain representation of the TMDS-encoded maximum contrast byte combination hexadecimal `10` and `55` as well as their difference signal.