

Ballot printer – protection against eavesdropping attacks – guidance for system designers

1 Background

In 2007, the Dutch government reviewed the security and use of electronic voting machines in elections, after the campaigning group “Wij vertrouwen stemcomputers niet” had published a report [1] demonstrating a number of attacks against such devices.

One of these demonstrations exploited electro-magnetic radiation that was unintentionally generated by the electronics inside a voting machine that was withdrawn in 2007, which accidentally revealed information about what political party the user is choosing. This information could be exploited by anyone with a suitable radio receiver tuned to the right frequency, as had been demonstrated at 25 m distance. The particular problem exploited there was that the device’s software continuously updated a liquid-crystal display with the currently displayed text, and due to the particular encoding required for that text, the repetition time of that software loop depended on the number of (non-ASCII) characters with diacritical marks displayed. One of the political parties (Christen-Democratisch Appèl, CDA) listed on the electronic ballot sheet has an “è” character in its name, and the display of that character reduced the rate at which the software loop executed. That rate was audible as a buzzing sound with an AM radio, on many tuning frequencies throughout the HF and VHF radio spectrum, and that buzz lowered its frequency whenever the CDA party name was displayed. A demonstration video [2] released by the campaigning group featured a “CDA detector”, a user-friendly improvised detector device that indicated visually whether a nearby voting machine was currently displaying a diacritical character or not.

Subsequent testing by a government laboratory [3] revealed that a second type of voting machine that was withdrawn in 2006 also was vulnerable to radio-frequency eavesdropping, revealing much more useful information, at even larger distances. This voting machine featured an integrated 1280×1024-pixel flat-panel video display that was refreshed at 60 Hz by an internal analogue VGA-cable video interface. With a suitable wide-band receiver tuned to multiples of the pixel frequency (108 MHz), e.g. at 432 MHz, and the help of raster-display software, a remarkably clear copy of the video image could be reconstructed tens of metres away. Software-defined radio equipment that can be programmed to visualise such compromising video signals is now widely available to hobbyists for less than 800 euros. A closer investigation by the author at the time revealed that the graphics controller used on the embedded PC main board found in the second type voting machine was the source of these emissions. Its three (red, green, blue) digital-to-analogue converter circuits together emitted the sum of the red, green and blue voltage generated for the VGA output as an amplitude-

modulated signal, via the power-supply lines of the main board, which then acted as transmitting antennas.

As a result of these events, the requirement emerged that future electronic voting equipment would have to be designed to minimise such “compromising emanations”, and would have to be tested for emission security.

2 TEMPEST standards

The phenomenon that computers can accidentally emit secret data, as unintended electromagnetic radiation, is not new, and has been a concern to military computer users since the 1960s. The particular emission-security risk posed by raster video displays, as was the case with the voting machine that was withdrawn in 2006, for example, is known as van-Eck effect, named after the Dutch telecoms engineer who first described it publicly in 1985.

NATO countries use a family of so-called TEMPEST standards to limit such compromising emanations in equipment destined to handle government secrets. The current incarnations of these standards are known as SDIP-27, 28, and 29. SDIP-27/1¹ specifies three security levels against which equipment can be tested. These levels make different assumptions about how close an eavesdropper can get to the equipment. Roughly speaking, the strictest Level A appears to be aimed at protecting against eavesdroppers that are at least about 8 m away, Level B against those at least 20 m away, and Level C against those at least 100 m away. Where equipment is installed at fixed locations in buildings, walls can provide additional shielding (depending on the building materials used). SDIP-28 defines three zones of electromagnetic shielding that rooms in buildings might offer against an eavesdropper (zones 0, 1 and 2), and these are roughly equivalent to the above equipment levels. SDIP-29 finally specifies detailed rules for which SDIP-27 level of equipment is allowed to be installed in what SDIP-28 building zone, and what additional measures have to be taken for various combinations of equipment level, room zone and information classification level (such as minimum separation to other cables or transmitters).

Some countries have adapted these standards to reduce cost and match their particular security concerns. For example, the German government uses a “zoning model”, which tests common office equipment according to simplified variants of the SDIP-27 Level B or C tests. These do not require the test engineers to understand the inner workings of the tested device, and provide simple pass/fail criteria that can be established in an automated test.

3 TEMPEST market, suppliers, and testing

The SDIP standards are classified documents (SDIP-27 is NATO CONFIDENTIAL, SDIP 28 and 29 NATO RESTRICTED) and therefore they play hardly any role so far outside the government ICT market. A small number of TEMPEST manufacturers specialize in producing equipment that complies with the SDIP-

¹ For brevity, we may omit the /1 etc. revision suffix of SDIP standards here.

27 standard. Operating in this market requires authorization by a national TEMPEST authority. This involves security-clearance vetting for all employees involved in assembling and testing the products, secure facilities for handling classified documents, and physical security for the laboratory and production facilities. Many national TEMPEST authorities publish lists of TEMPEST products that they have approved.

Most TEMPEST products on the market are regular personal computers and associated peripherals (laser printers, scanners, etc.) for office use, ruggedized laptops for military use, as well as network infrastructure components (routers, switches, firewalls, encryption units, etc.). However, the SDIP standards can equally be applied to other information-technology devices.

Products designed to comply with the SDIP-27 Level A standard typically use a metallic enclosure to provide a high level of electromagnetic shielding. Likewise, the power supply is chosen to include low-pass filters that minimize conducted signal leakage along the power-supply cable. TEMPEST-specific features may include especially shielded cable assemblies, or removable media and USB ports located behind a shielded door, which additional interface electronics automatically deactivates as soon as a switch contact reports that the flap door is opened. TEMPEST keyboards and mice have shielded enclosures, as well as shielded and filtered cables. Non-standard connectors can be used not only provide good shielding but also to discourage users from later replacing such components with non-TEMPEST versions.



The SDIP-27 standard specifies upper limits for permitted field strengths of compromising emissions. These are verified in a shielded room (semi-anechoic chamber) with measurement antennas connected to a spectrum analyser or measurement receiver. The measurement equipment used is similar to that required by civilian radio-frequency interference prevention standards, but the detailed parameters, requirements and limits differ significantly. National TEMPEST authorities often define additional quality assurance requirements, or make provisions for simplified tests in some situations. For example, the German TEMPEST authority distinguishes between two different types of tests applied during the design and production of TEMPEST products:

- The **certification test** is a comprehensive investigation of the entire relevant radio spectrum, with many different antennas in different orientations, as well as power-line taps, using the most sensitive measurement radio receivers and antenna amplifiers available. Preparing a certification test involves writing a detailed test plan that identifies the

signals of particular concern inside the device (known as “RED signals”), as well as their characteristics (bit rates, frequencies, etc.). The required measurements can sometimes take several days, depending on the individual device and in particular the types of RED signals identified inside the device. The actual measurement times are established as part of the detailed test plan. These measurements may have to be repeated frequently during the design and prototyping stages, as well as each time a modification is made later on to a product design.

- The **short measurement procedure** is instead a quick scan of the emissions of each individual device at the end of the production line, or after maintenance. It can usually be performed with a simpler electromagnetic-interference measurement receiver and fewer antennas and may focus on signal types and antenna positions that have been identified as particularly critical during the certification test. It aims to spot outliers caused by individual manufacturing defects in an already well-understood product, for example defects in the assembly of the shielded enclosure, in the wiring arrangements, or in installed filter components. The design of a suitable short measurement procedure for each product is part of its type approval process. A typical short measurement procedure can take about 25 minutes per device (including setup time). The design of a short-measurement procedure is also NATO CONFIDENTIAL, and is specific to the setup and calibration available in one particular laboratory.

Some manufacturers of TEMPEST equipment are only equipped to perform short measurement procedures, and leave the certification measurement to a national reference laboratory. Others are authorized and equipped to also perform the certification test, and the TEMPEST authority then only reviews their test plans and reports.

4 Applicability of SDIP-27

In 2007, the Dutch government first considered the possibility of requiring compliance with NATO’s SDIP-27 specification for future electronic voting machines. The initial proposal was to require compliance with the SDIP-27 Level B requirements, but with permitted signal levels reduced by 12 decibels (dB). This is the improvement in signal level experienced by eavesdroppers who reduce their distance to the target by a factor four, and was motivated by a minister at the time asking for adequate protection at a distance of 5 metres (4 times closer than the 20 m assumed by Level B). A later proposal was to require compliance with the Level A limits instead.

Several concerns were raised regarding the applicability of SDIP-27 to electronic voting equipment:

- A) As SDIP-27 is classified NATO CONFIDENTIAL, it is not available to the general public, and therefore independent sceptics of proposed future voting technology will not be able to convince themselves of the adequacy of its test requirements. They would have to trust that any trade-offs made by the military designers of this standard are also adequate for civilian applications, such as voting machines.
- B) SDIP-27 only specifies limits for permitted signal levels for various frequencies, bandwidths and target signal characteristics, but it does not provide any explanation or information about how these limits were decided. There is no rationale for the permitted emission levels available, nor any list of assumptions that went into their choice, other than the eavesdropping distances mentioned above. Therefore, even reviewers with the security clearance needed to read SDIP-27 cannot easily determine the suitability of this standard for applications, such as voting machines, that may have rather different requirements than what is needed to protect regular office equipment or military communications kit (see Section 5).
- C) The simple 12 dB reduction to the SDIP-27 Level B limits originally proposed would only work for frequencies for which the measurement antenna used in the test is already in the “far field” of the emitted radiation, where electric and magnetic components of the field are related to each other by a fixed factor (377Ω) and where signal voltages drop linearly with distance. At the SDIP-27 measurement distance of 1 metre, this is only the case for frequencies above about 50 MHz, and the 12-dB adjustment would not have correctly converted limits for lower frequencies, including some at which eavesdropping had been demonstrated successfully against the voting machine that was withdrawn in 2007. At lower frequencies, electric and magnetic fields may have to be measured independently to predict signal levels at other distances, and the required conversion factors will depend on the frequency. Therefore, the SDIP-27 Level A limits seemed more appropriate.
- D) One of the assumptions that may have gone into the design of the SDIP-27 limits is that the eavesdropper is trying to reconstruct and then read the content of an unknown text message, e.g. an email being typed on a keyboard or being displayed. This is a much more challenging task than that faced by the eavesdropper of an electronic voting device, because on the latter only a very small number of possible messages can be processed, e.g. the choice of one of about a dozen political parties, or the choice from a few dozen known candidates from the chosen party.

5 Specific concerns for electronic voting equipment

Concern D) in Section 4 is perhaps best illustrated by the fact that the demonstration that so successfully raised initial concerns about radio emissions from the voting machine that was withdrawn in 2007 [1][2] was actually only able to recover *one single bit* of information, namely whether a party name containing a non-ASCII character was being displayed or not, which at the time identified one single party. This is not only a reminder that, to discredit a technology in the public eye, sometimes a far less than complete attack demonstration may be quite sufficient. It is also a reminder that the target signal of a voting-machine eavesdropper is potentially of an extremely low bitrate.

The choice of one out of 16 known party names, for example, represents just 4 bits of information ($2^4 = 16$), and if that party name is looked at by the user (and eavesdropper) for a duration of, say, 4 seconds, eavesdroppers could achieve their goal by merely recovering a 4 bit / 4 s = 1 bit/s data signal.

It is impractical to devise a simple, SDIP-27 style, field-strength limit curve, to be measured with a spectrum analyser, that reliably excludes any possibility of such a low bitrate information leakage being possible from a tested device. To explain why, let us take a brief detour into communications theory. The Shannon–Hartley channel capacity theorem

$$C = B \log_2(1 + S/N)$$

states the best-case data rate C that could be transmitted over a communication channel of a given bandwidth B and signal-to-noise power ratio S/N . If we wanted to ensure, for example, that nobody can recover a $C = 1$ bit/s signal from a $B = 50$ MHz wide radio channel (for which software-defined radio receiver equipment is available today for a few hundred euros), then we would end up with a maximum permitted signal-to-noise ratio of $S/N = 2^{C/B} - 1 = 1.4 \times 10^{-8} = -79$ dB. In other words, under ideal circumstances, an attacker who knows exactly what to look for might still be able to extract 1 bit/s worth of data if the received signal power is a hundred million times (!) weaker than the background noise. However, if we don't know what the emitted waveform looks like, as is the case when we use a spectrum analyser that only measures power levels like those defined in SDIP-27, then such weak signals are impossible to distinguish from antenna noise. Instead, the eavesdropping receiver has to know how the information was encoded, in order to be able to use statistical techniques (e.g., calculating cross-correlation coefficients) to detect it against the noise.

One can argue, of course, that the Shannon–Hartley limit describes merely what an optimal, intentional transmitter (e.g., a digital modem) can get across a communications channel, and in practice, accidental emissions like those of concern here are usually orders of magnitude less easy to decode than ones that were specifically designed to be easy to decode.

However, there are some types of unintentional signals that, by accident, can be quite close to being optimal encodings. Of particular concern are signals that are periodically repeated, and those where the target signal is modulated by replacing each of its values with a known, very long, apparently random, sequence of bits, transmitted at a rate of many Mbit/s. In such circumstances,

accidental emissions start to resemble those of a direct-sequence spread-spectrum modulator, a type of transmitter optimized to send low bitrates over very noisy channels. (Satellite navigation is a prominent application of this technique.)

Imagine a computer that displays one of a small set of photographs that fill a large part of the display, such as candidate portraits or scanned party logos, and the eavesdropper merely wants to distinguish which known photo is displayed. Photos can contain a lot of random-like bits (due to camera sensor noise, dithering algorithms, decompression artefacts), and therefore the digital representation of these images, as it appears on a high-speed digital serial video interface would likely differ in nearly half of all bits and therefore form orthogonal signals, just like the symbols in a spread-spectrum modulator.

6 Proposed protection measures for a ballot printer

The Dutch government considers deploying a new generation of electronic voting machines where the functions of recording the voter's choice and the function of counting these choices at the end of the election day is split across two types of devices. A ballot printer presents the voter on a display the ballot sheet(s), guides the voter through the balloting process and records the choice made by the voter on a piece of paper, which might, for example, have a format similar to an airline boarding pass. These ballot papers are collected in a traditional ballot box and form an easily human-verifiable paper trail of the election. At the end of the day, the ballot box is emptied and the pieces of paper are fed into a vote counter that helps to count the result. This count can again be independently verified by manual recounting. The ballot printer is responsible for preserving the secrecy of the vote, and its designers are therefore expected to implement TEMPEST countermeasures. The vote counter deals only with ballots that are no longer linked to individual voters. It therefore does not have to protect the confidentiality of any processed data and therefore does not require TEMPEST countermeasures.

6.1 Use of SDIP-27 Level A

Complying with the SDIP-27/1 Level A limits for electric radiated, magnetic radiated and conducted power-line emission limits is a sound foundation for the TEMPEST protection concept for a ballot printer. These limits are known to be technically and economically achievable using conductive shielding and following good electromagnetic-compatibility design practice when planning the architecture and internal layout of such a device. There is a pool of suppliers available experienced in implementing this standard who have the necessary calibrated equipment, and there is an infrastructure of national TEMPEST authorities that provides quality assurance for these laboratories. There exist at the moment no other readily applicable emission-security standards that fulfil these requirements.

SDIP-27 requires the test engineer to make a catalogue of "RED" to be protected signals in the system, and characterize their nature (bit rate, frequency spectrum, parallel/serial, repetitive or non-repetitive, etc.). Most of the

interfaces found in modern PCs and similar embedded ICT devices operate at data rates ranging from tens of MHz to a few GHz, and for these, SDIP-27 measurements have to be made at correspondingly high bandwidths.

However, in light of the concerns raised above regarding low-bitrate risks for ballot printers, it would seem prudent to also apply, at least during the certification measurement, those SDIP-27 tests that are targeted at the lowest bit rates. The provisions of the standard do not go all the way down to 1 bit/s signals, as postulated above, but its tests for the nearest supported bitrate could still be applied usefully here. They would help to ensure that a narrow-band audible signal, such as the display loop buzz in the voting machine that was withdrawn in 2007, if it occurs at all, is not again as easily exploitable.

In addition, ballot-printer designers should not only aim to remain under the relevant SDIP-27 limit curves, but should also try to avoid the internal occurrence of RED signals that might be attractive eavesdropping targets. The following subsections give some guidance to product designers as to how that can be achieved.

Regarding environmental factors, devices tested according to the SDIP-27 standard should be installed according to the requirements of the SDIP-29 standard, which are quite easy to apply in case of a Level A shielded ballot printer, as they merely exclude the presence of other cables and transmitters within 50 cm of the device.

6.2 Avoiding compromising internal signals

6.2.1 Video system considerations

The periodic nature of video refresh signals make them a particularly attractive and easy to exploit target for radio-frequency eavesdroppers. The designers of a ballot printer should therefore consider internal or external video interfaces as a RED signal that warrants particular attention. Choosing the right type of interface (e.g. DisplayPort) can help to significantly reduce emission-security problems at the source.

We assume here that the video-display system of the ballot printer is likely to use similar OEM components as those found in a modern personal computer or mobile computing device (tablet computer). That has two main components: the graphics adapter (or GPU) and the display panel.

The graphics adapter (historically located in desktop PCs on a separate “graphics card”) is today often a part of the peripheral chipset located on the main board near the CPU and DRAM. It’s job is to periodically read the frame buffer content from the video RAM (which is in modern chipset graphics controllers a part of the main DRAM) and supply it via a video interface link, typically 60 times per second, to the display panel.

The display panel contains a timing controller (TCON) chip that receives the video signal from the graphics controller, buffers at least one line of it, and then forwards it in parallel to the many row or column decoder chips that actually drive the individual displayed pixels.

The intra-panel links between the TCON and row/column decoders are usually less of interest to a video eavesdropper, as the data they process is transmitted

in parallel, parts of a row at a time, although they might be an exploitable source in voting equipment where only a few bit about the displayed information may have to be inferred. These links are usually also quite well shielded, as they are located close to the ground plane within a printed circuit board inside the display panel.

The DRAM access pattern of the graphics controller is also usually also a less attractive target, for similar reasons. Its timing is less predictable and data is transferred in bursts, as the RAM bus is shared with the main CPU, and the links are usually also shielded within the printed circuit board.

The most attractive eavesdropping target in the video system is usually the link between the graphics controller and the TCON chip in the display panel. There are at least three reasons for this: it uses a serial interface that transmits one pixel after another, the conductors involved are much longer, and these video cables are prone to forming ground-return loops.

Many interface standards exist for this video link:

- **FPD-Link (LVDS)** – This has so far been the dominant input interface of liquid-crystal display panels in desktop computers. The digital video signal is transferred via three or more twisted-pair low-voltage differential signalling (LVDS, 1.1/1.4 volt) wires into the TCON chip. A separate pair supplies a pixel-clock signal, usually at a seventh of the bit-rate. This system was introduced as “Flat Panel Display Link” by National Semiconductor in the mid 1990s, but is colloquially known today more commonly as LVDS, an acronym for a standard that only describes the binary voltage levels used on the wires. There is no line encoding on these links: the red/green/blue bits plus some sync-pulse bits are transmitted over the wire directly as binary values, as they appear in the frame buffer [5]. The display panels accept exactly one video mode (display resolution and refresh timing). If a display panel is tightly integrated with a main board, as in a laptop or smartphone, the FPD-Link goes all the way from the graphics controller to the TCON chip. Where display panels and the main board are located in separate devices, as in a desktop computer, the FPD-Link is only used within the display to link with a **display controller** on a separate PCB. The job of the latter is to convert the other video-interface standards described below into FPD-Link, and also to resample the incoming video signal, in case it does not have the exact resolution and timing required by the panel TCON, to guarantee backwards compatibility.
- **Analogue RGB (VGA connection)** – On these 15-pin connectors, red, green and blue intensity is encoded for each pixel as an analogue voltage (0–0.7 V), along with horizontal and vertical synchronization pulses on separate lines. The VGA interface is a 1980s anachronism originally designed to drive cathode-ray tube displays. It is badly suited for driving flat-panel displays, which ultimately have to regenerate the pixel-clock signal, sample the analogue voltages and digitize them, all steps that can introduce errors, emissions, and add cost to the circuitry. Nevertheless,

the VGA interface is still very widely used today, in particular in industrial embedded PC components, which have a long support period and therefore have to remain backwards compatible with decades-old technology.

As the experience with the voting machine that was withdrawn in 2006 has shown, the digital-to-analogue converters required in a graphics controller to drive the VGA interface can be a very attractive source of modulated compromising video signals for an eavesdropper.

Flat-panel display modules with VGA input contain an extra display-controller circuit board that converts the VGA signal into FPD-Link, as mentioned above. As a result, the attacker has now *two* video links on offer, the one from the graphics adapter to the display controller, and the one from the display controller to the panel. This was the case with the voting machine that was withdrawn in 2006.

- **Digital Visual Interface (DVI)** – This is an older (mid 1990s) high-speed digital serial interface for flat-panel displays. It is in many ways similar to FPD-Link, but is better standardized for compatibility, and uses a very simple patented 8-bit to 10-bit line encoding scheme called Transition Minimized Differential Signalling (TMDS), which recodes transmitted bytes to remove DC currents from the binary signal [5]. Display panels usually do not directly receive DVI input: the DVI signal is usually first converted by a display controller into LVDS. For that reason, DVI is an external connector standard only and is usually not used inside integrated devices such as laptops or smartphones.
- **High-Definition Multimedia Interface (HDMI)** – This is technically the same TMDS interface as used by DVI, but also supports audio and uses a different connector, specified by the television industry, along with an optional content-protection protocol for movie copyright protection. DVI-HDMI adapters just change the plug shape and require no active electronics.
- **MIPI Alliance Display Pixel Interface (DPI)** – This is an older standard developed by the mobile device industry for use in phones and tablet computers. It uses a digital parallel interface and is still found on many smaller display panels. Its rigid timing and simple interface make it attractive to eavesdroppers, and its many parallel data lines can even cause electromagnetic-interference problems.
- **MIPI Alliance Display Bus Interface (DBI)** – This is an older standard developed by the mobile device industry for display panels that include a display controller and frame buffer. The host computer does not have to provide a refresh signal and can use the interface to update the displayed image only when needed. Little information is publically available about this standard, but the absence of a periodically repeating video image on this interface could mean that it poses a lower eavesdropping risk.

- **MIPI Alliance Display Serial Interface (DSI)** – This interface is commonly used in smartphones and tablet computers to drive the display panel. Although it is a considerably more complicated protocol than LVDS, with packet headers, checksums and bidirectional communication, and provides a number of options, many commonly used DSI implementations use the non-burst mode that offers the same highly predictable timing as LVDS, TMDS or DPI-based video interfaces. From an eavesdropper’s perspective, DSI is therefore quite similar to LVDS or TMDS.
- **DisplayPort (DP)** – This is the most recent digital video interface developed by the computer industry [4]. The underlying technology is completely different from either FPD-Link or DVI/HDMI. Several of its characteristics make it particularly favourable for TEMPEST applications and make eavesdropping on DisplayPort interfaces a far more challenging task:
 - DisplayPort uses a scrambler as part of its line encoding in order to flatten the Fourier spectrum of its emissions and suppress spectral peaks caused by particular image contents. This reduces the chances of any particular image content causing a problem spectral peak during SDIP-27 and EMI spectrum measurements. According to the standard, the scrambler reduces spectral peaks by about 7 dB.
 - As a side effect, the scrambler also makes it far more difficult, probably even impractical, for an attacker to reconstruct any information about the displayed image from the DisplayPort emissions. The DisplayPort scrambler uses a 16-bit linear-feedback shift register (LFSR) in order to generate a pseudo-random bit sequence that is then XORed into the image data by the transmitter. After that scrambling step, an 8-bit-to-10-bit line encoding is applied that further complicates the relationship between the displayed information and the transmitted bit sequence. The LFSR output repeats itself every $2^{16} - 1 = 65535$ bits, and is, in addition, reset every 512 lines, using a special scrambler-reset (SR) symbol sent by the transmitter. Unless the number 512 and the total number (including blanking interval) of lines per frame share any prime factors (which is not the case if the total number of lines in the video mode is an odd number), or the number 65535 and the number of bits per line do the same, the interval at which a scrambled video signal repeats itself should be 512 frames (≈ 8.5 s), too long, for example, to enable practically useful periodic averaging of the signal by the attacker in order to reduce noise.

- DisplayPort uses a small number of fixed bit rates, independent of the video mode used. Unlike with most other digital interfaces, video data is transmitted in data packets with header and padding bytes, and not continuously with a television-like timing. As a result, DisplayPort cables are not a common source of van-Eck-style video emanations and this again will make it very hard for an eavesdropper to synchronize to the transmitted data.

Recommendations:

- **Avoid using the VGA interface.** It was meant for CRTs and needs to be converted first into a digital video signal before any contemporary display panel can use it. The multiple signal conversions involved today when the VGA interface is still used can pose significant eavesdropping risks.
- **Preferably avoid using FPD-Link/LVDS, DSI, or DPI.** Their simple and predictable line encoding can cause easy to reconstruct compromising video signals. If the better alternative (embedded DisplayPort, see below) is not available or practical, they may be acceptable if the countermeasures against ground-return loops described later in this report are implemented carefully, including: separation from other conductors, star grounding topology, and use of ferrite-ring chokes.
- **Avoid DVI or HDMI.** Everything said for FPD-Link also applies to DVI and HDMI. In addition, like with VGA, these standards are not accepted directly by display panels, but are first converted into FPD-Link/LVDS or DPI via a display controller, resulting again in two copies of the video signal becoming available to the attacker. In other words, from a TEMPEST perspective, the DVI and HDMI interfaces combine the disadvantages of LVDS and VGA.
- **Avoid using a multi-standard display controller chip for converting between different video interfaces and video modes.** Instead, make sure that the graphics controller on the main board supplies directly the signal type required by the TCON chip in the display panel. Two different video links can double the chance that at least one of them provides an exploitable signal.
- **Use the embedded DisplayPort (eDP) interface standard to directly link the graphics controller with the display panel (or Direct Drive Monitor).**

While the full implications of the DisplayPort line encoding for eavesdroppers are still the subject of on going research, our current understanding suggests that it makes a successful eavesdropping attack on DisplayPort cables highly unlikely. DisplayPort uses signalling rates of either 1.62 or 2.7 Gbit/s in each of the 1–4 twisted-pair lanes, therefore any eavesdropper using a receiver bandwidth of less than in the order of 1 GHz will suffer too much inter-symbol interference to be able to recover the line code with a low-enough error rate to be able to invert the line encoding. Such ultra-wide receiver bandwidths are at the moment not practical for many reasons, including the required processing power, lack

of affordable off-the-shelf equipment, and interference from radio transmitters.

When configuring or evaluating a DisplayPort device, if possible, check that the LFSR scrambler is not generating output sequences that repeat at the frame rate, or some small integer multiple thereof. With a fixed image content, the bit pattern on the DisplayPort interface should ideally only repeat every 512 frames. Such a test could be performed by tapping into the eDP links with a fast oscilloscope with enough acquisition memory to record hundreds of frames, and then looking for peaks in the auto-correlation function of the recorded signal.

- **Activate frequency modulation of clock signals.** Some graphics adapters can apply “spread spectrum” frequency modulation to the generated pixel-clock, for example with a 30 kHz triangular wave [6]. This is a trick aimed at making compliance with the “quasi-peak detector” measurements required in international radio-interference standards (CISPR 22, CE mark) easier. Where such an option is available in the graphics driver, it should be activated. Such an option is unlikely to increase the chances of passing an SDIP-27 test, which does not use a quasi-peak detector, but it adds non-determinism to the signal that may thwart an attacker’s attempt to detect the signal against the background noise using cross-correlation techniques. This frequency modulation can also affect the timing of intra-panel signals, and would then make them more difficult to eavesdrop as well.

Market availability: The term eDP means “embedded DisplayPort”, that is DisplayPort used inside an integrated device, such as a laptop or tablet computer, without an externally visible DisplayPort connector. Panels with eDP input are already used in recent mobile devices, such as, for example, the LG LP097QX1-SPA1 panel (2048x1536 pixels) in the iPad 3. Industry forecasters expect that the eDP interface will gradually replace FPD-Link/LVDS in the coming years in the display panel market, especially for higher resolutions, and it can, therefore, be hoped that a range of display panels with eDP input and mainboards with corresponding eDP outputs can be sourced for the ballot-printer display system. Nevertheless, this is relatively new technology that may not yet be as mature and easily available in lower volumes as traditional panels with LVDS or DSI input, or monitor modules with integrated multi-standard display controllers. Likewise, operating-system drivers for the graphics adapter will have to support eDP. This has in the past held back in desktop computers the use of “Direct Drive Monitors”, displays with DisplayPort connector and without a separate display controller, where the DisplayPort cable connects directly to the panel input.

6.2.2 Printer considerations

Most TEMPEST printers on the market are A4-sheet office laser printers. The modifications required for them to pass tests are therefore well understood by TEMPEST manufacturers. For example, particular focus has to be given to shielding the laser-diode drive current, which can be eavesdropped in ways similar to an analogue video signal. While in a video signal, the periodic refresh of the image leads to redundancy that makes such a signal easy to eavesdrop,

with laser printers, the signal redundancy that makes eavesdropping of unknown text possible is the high vertical resolution. Many successive pixel rows in high-resolution printer output are nearly identical, which means that an eavesdropper can average them, to reduce noise, into an image of lower vertical resolution that is still readable.

The main other low-cost printer technologies available are ink-jet and thermal printing.

Thermal printers may be particularly attractive for a ballot printer, due to their robustness, simplicity and ease of maintenance (no ink or toner needed). We have not yet performed TEMPEST measurements on any thermal printer. Therefore, the following considerations are somewhat speculative:

Thermal printers use an array of resistors to heat up thermal paper in order to release ink that is embedded in this type of paper. The printing process involves a controller rapidly switching on and off of per-column heating currents, and these switching events may cause electromagnetic and power-line emissions.

On the positive side: thermal printers should be highly parallel devices, which print entire rows of pixels at the same time. Parallel processes tend to provide much less information than serial processes to eavesdroppers, because of the higher overlap of activity in time. However, this needs to be confirmed with measurements on a particular product under consideration. Some thermal printer designs might well have arranged heating elements in some kind of scan matrix, to reduce the number of current drivers required, heating up one at a time in rapid succession.

In contrast to laser printers, thermal printers that drive their heating elements in parallel would use low-frequency current signals. Typical print speeds are 1000–2000 rows of pixels per second. This implies low-impedance near-field conditions for the measurement, and therefore radiated magnetic fields should be tested in the certification measurements, something that is not commonly done for other office equipment that operates at much higher frequencies.

In case there are problems with getting thermal printer emissions below the magnetic Level A limit lines, improved cable layout and similar geometric measures should be tried first. Metallic enclosures are much less effective at shielding magnetic rather than electric fields. Special high-permeability alloys (Mu-metal) are available to improve shielding of low-frequency magnetic fields, if needed.

Smaller paper-ballot sheet sizes are preferable, as the paper input and output apertures could leak RF signals in the gigahertz range. Placing the high-frequency video electronics and the low-frequency printer electronics into separately shielded compartments, or even separate enclosures connected by a shielded combined power and data cable, would make this less of a concern. (Splitting the device into two enclosures may also be advantageous for the design of the enclosure, as a printer, a paper reservoir, and a large-format video display have very different form factors, and any single enclosure that tries to integrate all of these is likely to be quite large.)

The paper compartment of the printer should ideally be kept outside the main electromagnetic enclosure of the printer, to reduce the risk of it being damaged

during the paper reloading process (e.g., dirt interfering with high-frequency gaskets).

6.2.3 Touch-screen considerations

TEMPEST displays are usually fitted with a conductive transparent panel across the display surface, with careful continuous conductivity to the metallic enclosure around the edge. This adds a transparent window to the enclosure without interrupting conductivity across the surface, and there are conductive transparent sheets available that do not affect the viewing comfort negatively.

The most popular touch-panel technologies in mobile devices are capacitive sensors that sense a conductive finger. These cannot work if a conductive transparent layer is inserted between the finger and the panel. Alternatives include resistive (pressure sensitive) touch panels, infrared beam systems, imaging systems, or fitting the touch panel outside the shielded enclosure.

6.2.4 Software and user-interface considerations

Recommendation: Avoid busy-wait loops – The application software on the ballot printers should avoid endless loops that deal with confidential data, as these could accidentally modulate this data in radio emissions (as was the case with the voting machine that was withdrawn in 2007). When the software awaits a new user interaction, it should instead place the CPU into an idle mode, as is common practice on all modern operating systems. This way, the eavesdropper could only observe brief bursts of CPU activity, with much less opportunity to modulate data (i.e., produce easy to receive periodic signals).

An attacker might be able to see whether the CPU is currently in idle mode or not, for example by recording the power consumption of the device. The user interface should, therefore, be designed such that the number of times that the CPU awakes from its idle state in response to a user interaction does not reveal the vote. To give a simple example: if the user had to use a cursor-down key to move a cursor to the location of a candidate's name on a screen, always starting with the cursor pointing to the first candidate, then the number of cursor-down presses, and equivalently the number of CPU awakenings from the idle state, can identify the selected candidate. We expect this to be less of a problem with a touch-panel, where the user directly enters display coordinates rather than navigating with keys. (Otherwise, a simple countermeasure would be to place the cursor initially at a location that has been picked uniformly at random, and to allow the cursor to wrap around between the first and the last entry.)

Avoiding that the number of key presses or touch-screen taps reveals useful information also reduces the risks of acoustic eavesdropping, e.g. determining the vote by listening to the number of taps onto a screen or key.

Randomizing less-significant pixel bits – One software technique for reducing the risk of compromising emanations that can be picked up from digital video displays is to replace in the frame-buffer data, each time the display needs to be updated by the application, the less-significant bits of the red/green/blue values with freshly generated random bits [5]. To make this practical, the application software would have to render its display layouts into a separate frame buffer, which is then randomized, for example by replacing the least 4 significant bits in each 8-bit RGB value with random bits, across the entire frame buffer area,

before transferring this updated image into the frame buffer used by the graphics adapter. A practical way of implementing this is to use texture maps of a 3D GPU library (e.g., OpenGL). The display layout and the random noise can be kept in separate texture maps, and the random-noise map can be made almost transparent and placed on top of the display layout. This way, the CPU never has to touch the memory locations that store the rendered party list while the voter makes their choice.

This randomization should only occur when the application software updates the display, to ensure that the rate at which the random bits are replaced with new ones is comparable to the rate at which the intended video image is updated. If the random bits are updated very frequently (e.g., at 60 Hz), an attacker might be able to remove this added random noise by periodic averaging. If the bits are updated too rarely, an attacker might be able to see image changes by subtracting a previously recorded video signal from a current one. This technique, similar to what we have remarked above regarding the CPU idle state, can leak to the eavesdropper when and how often user interactions and associated display updates occur. Therefore, the user interface should avoid elements where the number of interactions leaks useful information (e.g., no use of up/down keys to select entries).

Using a 3D API and GPU to display screen masks (such as party lists) that could leak the voter's choice also enables easy implementation of another countermeasure: by slightly randomly realigning the texture map of the party list in the GPU's 3D space, the list can be very slightly rescaled, translated, rotated, and sheared each time it appears on the screen in response to a voter's choice. Such variation should be small enough to not be perceived by the voter (each corner just moves a few pixels). This way, the eavesdropper never knows the exact alignment of the bitmap displayed, which can further help to reduce the accuracy they can achieve with classification algorithms used to distinguish between the signals emitted by different party lists. At the same time, GPU-based rescaling can improve rendering quality through the anti-aliasing filters applied.

We would advise considering the use such a pixel randomization technique in the application software if the video system used involves an unscrambled digital link. With the availability of a scrambled video link in the form of embedded DisplayPort (eDP), this software technique is less crucial, and therefore mentioned here only as a recommendation, rather than as a major requirement.

6.3 Electromagnetic-compatibility design practice

The design of TEMPEST equipment generally benefits from following the same design guidance that is also taught and recommended as best practice to avoid electromagnetic-compatibility (EMC) problems. As there are many good textbooks available on EMC design, we summarize here only a few basic principles.

6.3.1 Shielding

The ballot printer should be housed in a metallic enclosure that was specifically designed to attenuate electric and electromagnetic fields ("Faraday cage"). The basic requirement for an effective shielded enclosure is that electric currents can

flow unhindered across it, at low impedance, along the shortest possible path between any two points on its surface. This requires continuous electrical connectivity along any joints between two metal parts, avoiding long non-conductive gaps in between. This can be achieved, for example, by using many screws or spring contacts at regular short distances, never more than a small number of centimetres (e.g., 40 mm) apart. If there is a gap between metal parts, then any electric current that wants to flow across it has to take a detour around the gap and the shielding then fails for wavelengths shorter than the length of that detour.

6.3.2 Avoiding ground-return loops

Any electrical current should return to its source along almost the same path over which it came, otherwise the resulting circuit acts like a loop antenna and generates a magnetic field. Some approaches for reducing ground-return loops include:

- Twist supply and return conductors, to avoid the formation of loops.
- Attach them very closely to any metal structures that might offer alternative low-impedance ground-return routes.
- Add a (coaxial) metallic shield around twisted pairs, with a low-impedance ground connection on both sides, to provide an alternative low-impedance ground return that is not an effective loop antenna.
- Preferring symmetric/balanced signal interfaces over asymmetric ones, (avoid cables with TTL or RS-232 signals in favour of USB, LVDS, etc.).
- Install ferrite ring chokes across balanced communication lines, in order to suppress common-mode currents.
- Implementing a star-shaped grounding concept, where there is only one single low-impedance ground connection for any component. (This may require the use of non-conductive spacers to separate components from the metal chassis. Also beware of capacitive coupling.)

6.3.3 Reduce or filter cables that penetrate the shielded enclosure

Any cable that crosses a metallic enclosure can act as a receiving antenna on the inside and as a transmitting antenna on the outside, over which compromising signals can escape. Therefore, such cables should best be avoided. Where they are unavoidable, a low-pass filter can be installed at the boundary in order to attenuate all parts of the radio spectrum that are not required on this cable. This is usually only practical for power-supply lines and some low-speed interfaces. For high-speed interfaces, if they are necessary at all, fibre-optic interfaces are much preferable, as their cables lack conductive parts.

USB port – The specification of the ballot printer may include a USB port. Due to its small aperture, the USB slot itself is usually not a concern for TEMPEST emissions, as long as nothing is plugged into it. But cables connected to such

slots can form accidental transmission antennas that bring internal signals outside the shielded enclosure where they can radiate. Therefore, users must be discouraged from plugging anything into a USB port while the device processes confidential data, unless that configuration was tested during the TEMPEST certification. In some TEMPEST desktop PCs, USB ports are located behind an electromagnetically shielded flap. While the flap is open, a switch disconnects the USB ports. A USB memory stick can still be operated inside the shielded enclosure when the flap is closed, but the flap prevents that anyone connects and successfully uses a USB cable. USB ports cannot be filtered very effectively, due to the high differential-mode bandwidth that this interface requires to work normally. Therefore, the flap eliminates the risk that a connected USB cable leaks radiation from inside the shielded enclosure. This makes sense in a general-purpose PC, where there are many reasons for why a user might want to connect a USB cable without authorization. With a ballot printer, we would hope that the use of the USB port is restricted by the software to its only function, namely to upload a configuration file from a memory stick, and then without any incentive to connect a USB cable, such a flap mechanism may not be required. In this case, the test plan would have to explain why it does not require plugging a USB cable into the port during the measurements, as is usually required.

6.4 Radio transmitters and RFID readers

Including a radio transmitter into a device can significantly increase the complexity of TEMPEST testing, as the test plan now has to confirm that the transmitter is not accidentally picking up or modulating any RED signal.

Therefore, if the ballot printer includes a transmitter, for example an ISO 14443 proximity-card reader (which emits a 13.5 MHz carrier wave), then it should power down the card reader (no carrier emitted) before the voter can enter their choice. It should power up again only after the ballot printer has deleted from its memory any trace of the voter's choice. This way, the emitted carrier wave will not pose an additional eavesdropping risk, as the printer has no operational transmitter while it handles confidential data. Then TEMPEST tests only have to verify that the transmission antenna is not passively leaking internal signals through the shielded enclosure, like any other cable penetrating it.

6.5 Other eavesdropping protection considerations

Radio-frequency eavesdropping, for example in the 100 Hz to 10 GHz spectrum, is only one prominent channel that may allow an eavesdropper to gain information about the vote. This section looks briefly at some other channels, for which currently no protection standards exist.

6.5.1 Optical leaks

The light emitted by a computer display can leak to an observer, even if there is no direct line of sight, either through diffuse reflection, such as from the voter's face, hands or clothes, or through specular reflection, such as via eyes or glasses. This is a particular risk if the room in which the ballot printer is operated is not brightly lit. Therefore, it is prudent to ensure that the layout and average display colour and brightness do not leak the voter's choice. A particularly bad idea, for example, would be to display a party list on a background colour that clearly identifies the political party. The face of a voter looking for a candidate of the

“Green Party” might then light up in green, or a green square could appear in the reflection of the screen in their eyes or glasses. In comparison, making a pencil cross on a paper ballot does not noticeably affect the average brightness or colour of the ballot sheet, and therefore does not pose such a risk.

As a rule of thumb, the average brightness in each of the three colour channels (red, green, blue), averaged across all pixels of the display of one party list, should vary by less than 5% from the average brightness of all party-list displays. For example, candidate portraits should not occupy a large fraction of the screen surface and should preferably share a neutral (e.g., grey) background colour. The length of the party list displayed should not be easily recognizable from a distance. This could be achieved by padding the end of any party list with empty entries of similar colour.

6.5.2 Power consumption

Apart from radio-frequency emissions on the power line, which can be controlled with low-pass filters in the power supply, there may also be lower-frequency variations in the power supply current (well below 1 kHz) that could leak the voter’s choice.

For example:

- A thermal printer activates a small heating element to darken a pixel on the paper, therefore the overall power consumption of a thermal printer can leak the number of black pixels on the ballot printout, which in turn may indicate the voter’s choice. If tests suggest that this is a realistic threat, the printer could be supplied via a constant-current circuit, which can be built by connecting a current regulator in series with a voltage regulator.
- Modern CPUs consume significantly different currents depending on whether they are busy or idle. CPU activity also can also reveal a voter’s choice. For example, if party lists are rendered each time from scratch, using computing-intensive graphics API calls (anti-aliased fonts, rescaling portrait photos, etc.), then the amount of CPU time spent on rendering the party list can leak the length of the party list being displayed, and thus the voter’s choice. Therefore, any code that is executed while the voter makes their choice should be written with constant execution time in mind. For example, party lists could be rendered in advance into an in-memory bitmap (say an OpenGL texture map), and when it comes to actually displaying a specific party list, the CPU merely executes the constant-time operation of instructing the GPU to switch the display to that particular texture map. This way, hardly any information related to the voter’s choice appears in the CPU’s activity level or on the CPU’s memory interface.

6.5.3 Headphones

The specification envisages a separate interface for visually impaired users, using voice menus presented via headphones. Headphones can create additional

eavesdropping opportunities in two ways. Their cables and voice coils can act as unintentional transmission antennas, in particular for high-frequency signals picked up from inside the ballot printer. In addition, the sound they produce can leak into the environment, which could be picked up by nearby human ears, as well as by directional microphones and signal-classification algorithms trained on known dialogue fragments. The voice interface has to confirm the choice of the voter and therefore must remain confidential.

The following considerations aim to mitigate eavesdropping risks created by the headphone plug-in interface:

- Headphones are commonly available with either analogue interface (2.5 mm, 3.5 mm or 6.35 mm plug) or digital interface (USB connector, USB audio device class).
- If the ballot printer provides an analogue headphone interface, then a band-pass filter should be installed inside its metallic enclosure close to the socket, to limit both differential and common mode currents across the headphone wires to the voice frequency range (approximately 100 Hz to 7 kHz). The designer of such an analogue filter should keep in mind several functions: (a) to prevent the headphone cable emitting RF signals from inside the printer (e.g., up to 10 GHz), (b) to limit the audio-frequency range for which the headphones connected need to provide good acoustic leakage suppression, and perhaps even (c) to reduce the risk of covert signal emission by malicious ballot-printer software (e.g. at ultrasonic frequencies).
- Unintended emissions are easier to suppress on analogue audio interfaces, whereas USB interfaces depends on differential-mode signals of well over 10 MHz to get through.

Many different types of earphones are available (terminology: one or two earphones along with a headband form a headphone) [7]:

- insert earphones – designed to be inserted into the ear canal
- intra-concha earphones – designed to fit the concha cavity (outside the ear canal) with an acoustic exit close to the entrance of the ear canal
- supra-concha earphones intended to rest on the ridges of the concha cavity
- circumaural earphones – having a cavity large enough to cover the region of the head including the ear

(For completeness: there exist also “ear shells” – earphones hanging on the ear, and “stethoscopic headphones” – insert headphones coupled to the ears by a pair of rigid tubes, like a stethoscope. The latter would require no conductor outside the ballot printer, making them safest from a purely electromagnetic-emission point of view.)

Earphones can also be classified into these types of construction:

- acoustically closed/open – intended to prevent/provide acoustic coupling between the external environment and the ear canal
- closed/open back – does not/does emit significant sound radiation from the back of the transducer to the external environment

In the interest of leakage prevention, the headphones chosen for the ballot printer should obviously be both acoustically closed and also feature a closed back. Insert earphones can be favourable for privacy reasons because they need to move least air to achieve a certain sound-pressure level inside the ear canal. However, hygiene considerations make circumaural earphones more practical for an electronic voting device with many users, unless the operator is prepared to hand out single-use insert phones to each voter who requires them (some of which can be sourced for less than 1 euro per pair).

Headphones under consideration for use with ballot printers should be tested for their ability to suppress unwanted sound radiation to the environment. International standard IEC 60268-7 [7] specifies such a test, among others. It involves mounting the headphones on a head/ear simulator, adjusting the loudness measured inside the artificial ear to some reference value (e.g., 94 dB sound pressure level), sweeping the input frequency across the frequency range of interest (e.g. 100 Hz to 7 kHz) with a sine-wave signal generator connected to the headphones under test, and measuring the loudness of the leaking sound with a microphone located at 0.1 m distance, facing the back of the earphone. The result is a chart that shows the sound-pressure-level difference between the microphone in the simulated ear canal and the one outside, for each tested frequency.

Many manufacturers offer headphones specifically designed to shield the ear against environmental noise, for example targeted at users in recording studios or aviation. Their data sheets often document the attenuation of environmental sound as perceived by the ear. This is *not* the same thing as the attenuation of unwanted sound radiation into the environment (although these two measures may be correlated). For example, headphones with active cancellation of environmental noise do not actively cancel unwanted sound leakage into the environment. Only the attenuation of unwanted sound radiation into the environment is of interest to privacy and should therefore be measured separately.

Other considerations:

- The audio dialogue should explain to the voter at the start of an audio session how to control the volume throughout the voting process, and how to repeat a message they did not understand. It should then encourage them to use the lowest useable audio volume setting, to minimize the risk of audible signal leakage.
- Ballot printers that are used with headphones should be positioned as far away as practical from other people.
- A ballot printer could also implement measures to increase environmental noise when used with earphones, in order to mask

audible signal leakage and thereby reduce the signal-to-noise ratio for an eavesdropper. (For example, if it includes a variable-speed cooling fan, the software could reconfigure that fan to maximum speed during an audio session, to raise the background noise level. If a loudspeaker is used to generate background noise, it could output noise generated by an unpredictable random source, like those used to generate cryptographic keys, which can then be band-pass filtered to match the peak spectral composition of typical voice dialogues. Alternatively, it could also compose background noise by mixing and superimposing many fragments of dialog waveforms stored in the system, each time newly chosen and realigned at random, to be as difficult to separate as possible from the actual voice dialogue. However, adding a loudspeaker also opens another potential covert channel for unauthorized software modifications to deliberately leak data via steganographic techniques.)

Other options, such as increasing the environmental background noise by playing music in the room, installing soundproof cabins, or dedicated ballot printers for audio users in separate rooms, were not considered practical.

7 Summary and conclusions

This report highlighted the applicability and limits of the SDIP-27/1 Level A standard to ballot printers, in particular if we treat the abstract signal that encodes the choices made by the voter as a RED signal of very low bit rate processed inside the device. While applying SDIP-27 Level A is a reasonable and practical element of a TEMPEST strategy for such a device, in light of the low bitrate to be protected, and considering that the “inspectable space” required by SDIP-29 around a Level A device may not be enforceable at many voting stations, it is also worth following some additional design guidelines that we have presented in this report.

The designer of a TEMPEST ballot printer should in particular consider the advantages of the embedded DisplayPort (eDP) video interface, and how its scrambler and correct use can help to significantly reduce the risk of eavesdropping on the video link.

We have also highlighted some lessons learned from the failures of the previous generation of electronic voting machines, in particular regarding the design of the software and the video system.

References

- [1] Rop Gonggrijp and Willem-Jan Hengeveld: Studying the Nedap/Groenendaal ES3B voting computer – a computer security perspective. USENIX/ACCURATE 2007 Electronic Voting Technology Workshop (EVT '07).
https://www.usenix.org/legacy/events/evt07/tech/full_papers/gonggrijp/gonggrijp.pdf
- [2] Voting computer tempest attack. Youtube video, uploaded October 2006. <https://www.youtube.com/watch?v=B05wPomCjEY>
- [3] Aanvallen op het stemgeheim via elektromagnetische effecten. Algemene Inlichtingen- en Veiligheidsdienst, October 2006. https://www.aivd.nl/binaries/aivd_nl/documenten/publicaties/2006/10/31/aanvallen-op-het-stemgeheim-via-elektromagnetische-effecten/aanvallenophetstemgeheim.pdf
- [4] VESA DisplayPort standard, Version 1, Revision 1a. Video Electronics Standards Association, 11 January 2008.
- [5] M.G. Kuhn: Electromagnetic eavesdropping risks of flat-panel displays. In: Privacy Enhancing Technologies, Lecture Notes in Computer Science, vol. 3424, pp. 88–105, Springer-Verlag, 2004.
- [6] M.G. Kuhn: Compromising emanations of LCD TV sets. IEEE Transactions on Electromagnetic Compatibility, Vol. 55, No. 3, pp. 564–570, June 2013.
- [7] International Standard IEC 60268-7: Sound system equipment – Part 7: Headphones and earphones. Edition 3.0, International Electrotechnical Commission, 2010-01.