

Number 577



UNIVERSITY OF
CAMBRIDGE

Computer Laboratory

Compromising emanations: eavesdropping risks of computer displays

Markus G. Kuhn

December 2003

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500
<http://www.cl.cam.ac.uk/>

© 2003 Markus G. Kuhn

This technical report is based on a dissertation submitted June 2002 by the author for the degree of Doctor of Philosophy to the University of Cambridge, Wolfson College.

Technical reports published by the University of Cambridge Computer Laboratory are freely available via the Internet:

<http://www.cl.cam.ac.uk/TechReports/>

ISSN 1476-2986

Summary

Electronic equipment can emit unintentional signals from which eavesdroppers may reconstruct processed data at some distance. This has been a concern for military hardware for over half a century. The civilian computer-security community became aware of the risk through the work of van Eck in 1985. Military “Tempest” shielding test standards remain secret and no civilian equivalents are available at present. The topic is still largely neglected in security textbooks due to a lack of published experimental data.

This report documents eavesdropping experiments on contemporary computer displays. It discusses the nature and properties of compromising emanations for both cathode-ray tube and liquid-crystal monitors. The detection equipment used matches the capabilities to be expected from well-funded professional eavesdroppers. All experiments were carried out in a normal unshielded office environment. They therefore focus on emanations from display refresh signals, where periodic averaging can be used to obtain reproducible results in spite of varying environmental noise.

Additional experiments described in this report demonstrate how to make information emitted via the video signal more easily receivable, how to recover plaintext from emanations via radio-character recognition, how to estimate remotely precise video-timing parameters, and how to protect displayed text from radio-frequency eavesdroppers by using specialized screen drivers with a carefully selected video card. Furthermore, a proposal for a civilian radio-frequency emission-security standard is outlined, based on path-loss estimates and published data about radio noise levels.

Finally, a new optical eavesdropping technique is demonstrated that reads CRT displays at a distance. It observes high-frequency variations of the light emitted, even after diffuse reflection. Experiments with a typical monitor show that enough video signal remains in the light to permit the reconstruction of readable text from signals detected with a fast photosensor. Shot-noise calculations provide an upper bound for this risk.

Acknowledgments

I would like to thank my former supervisor Ross Anderson for making this entire project possible and for encouraging my initial experiments with his ESL 400 emission monitor. I am also much in debt to Tony Kruszelnicki of TK Electronics in Lincoln for an extended loan of a Dynamic Sciences R-1250 receiver and various accessories. The technical staff of the Computer Laboratory and in particular Piete Brooks earned my thanks for their patience and competent assistance. Robert Watson contributed useful GPIB/TCP gateway software for instrument control during his brief stay as a visiting student. Simon Moore provided a storage oscilloscope and Richard Clayton helped reverse-engineer an annoying problem with its firmware. Richard Clayton, Ross Anderson, and Gareth Evans provided valuable comments on the draft text and were along with Sergei Skorobogatov and David Wheeler available for useful discussion. The purchase of some of the equipment used was made possible through the TAMPER hardware security laboratory support provided by NDS and Hitachi. I was supported by a European Commission Marie Curie training grant.

Contents

1	Introduction	9
1.1	Historic background and previous work	10
1.1.1	Military activities	10
1.1.2	Open literature	12
1.2	Motivation and scope	14
2	Foundations and test equipment	19
2.1	Antenna types	19
2.2	Receivers	23
2.3	Receiver calibration	26
2.3.1	Impulse bandwidth	28
2.3.2	Impulse strength	31
2.4	Signal correlation	33
3	Analog video displays	37
3.1	Video-signal timing	37
3.2	Analog video-signal spectra	40
3.3	Eavesdropping demonstration	45
3.3.1	Realtime monitoring	45
3.3.2	Experimental setup	47
3.3.3	Results	47
3.4	Radio character recognition	54
3.5	Hidden transmission via dither patterns	57
3.6	Filtered fonts as a software protection	60
4	Digital video displays	67
4.1	Case study: Laptop display	68
4.2	Case study: Digital Visual Interface	77

5	Emission limits	85
5.1	Existing public standards	86
5.1.1	Ergonomic standards	87
5.1.2	Radio-frequency interference standards	87
5.2	Considerations for emission security limits	89
5.2.1	Radio noise	91
5.2.2	Radio signal attenuation	92
5.2.3	Power-line noise and attenuation	94
5.2.4	Antenna gain	95
5.2.5	Processing gain	96
5.3	Suggested emission limits	97
6	Optical eavesdropping of displays	105
6.1	Projective observation with telescopes	105
6.2	Time-domain observation of diffuse CRT light	106
6.3	Characterization of phosphor decay times	107
6.3.1	Instrumentation	108
6.3.2	Measurement method	110
6.3.3	Results	111
6.4	Optical eavesdropping demonstration	116
6.5	Threat analysis	120
6.5.1	Direct observation	120
6.5.2	Indirect observation	122
6.5.3	Observation of LEDs	123
6.6	Receiver design considerations	124
6.7	Countermeasures	126
7	Review, outlook and conclusions	129
A	Electromagnetic fields	143
A.1	Maxwell's equations	143
A.2	Quantities and units	146
A.3	Electromagnetic emanations	146
A.4	Transmission lines and antennas	148
A.5	Time-domain characterization of antennas	151

B	Notes on experimental setups	157
B.1	Impedance-matched attenuators	157
B.2	Video sync signal generation	158
C	Glossary	161

Chapter 1

Introduction

“It has long been a dream of cryptographers to construct a ‘perfect’ machine [...] The development in the last twenty years of electronic machines that accumulate data, or ‘remember’ sequences of numbers or letters, may mean that this dream has already been fulfilled. If so, it will be the nightmare to end all nightmares for the world’s cryptanalysts. In fact, the people who live in the vicinity of the National Security Agency think that there already are too many cipher and decoding machines in existence. The electronic equipment plays havoc with their television reception.”

— D.T. Moore, M. Waller: Cloak & Cipher, 1965 [1, p. 153]

Computer and communication equipment receives and emits energy in various forms, such as electrical currents, heat, light, conducted and radiated electromagnetic waves, sound and vibrations. Most energy consumed will be released as heat or is used to form intended symbols on communication channels. Some of the rest is correlated in various ways to processed data and can form unintended information leaks. This opens unconventional opportunities for technically sophisticated outsiders to get unauthorized access to processed confidential information.

Compromising electric, electromagnetic, optic, acoustic, ultrasonic, mechanic, etc. emanations can be a potential computer security threat if information is emitted in a form that can be practically separated from background noise and decoded at sufficient distance using compact and available equipment. It can then be used to bypass commonly employed physical, cryptographic, and software access-control mechanisms at the operating-system, network, and application level.

Such exploitable emanations can occur as a result of:

- the normal operation of a system
- deliberate or accidental exposure of a device to an unusual environment
- the execution of software that was designed to modulate data into emitted energy

Carefully chosen software measures can sometimes be applied to control the emitted signals. They can emit data in a form particularly suited for easy remote reception or

they can render the otherwise feasible remote reconstruction of data far less practical. In the latter case, such measures can be a welcome low-cost alternative for, or an additional protection layer to, hardware shielding.

Digital computers and telecommunication equipment have penetrated our civilization over the past half century, and at the same time, many aspects of human society have become significantly dependent on the availability, integrity, or confidentiality of automatically processed information. The risk posed by malicious abuse of this information infrastructure is widely recognized today, as is demonstrated by the significant current media interest in computer security incidents and the growth of an information security industry that offers a wide range of countermeasures.

Historically, most practical outsider attacks on information systems have involved some form of access to communication links, in order to eavesdrop exchanged data or impersonate communication partners. Highly effective protection techniques against both these threats, namely cryptographic protocols and authentication tokens, are on the way to becoming ubiquitous product features. Other attacks require either knowledge of vulnerabilities in the specific implementation or configuration of the target system, or physical access to it. Both tend to be available only opportunistically and can then be used to either get hold of data from storage media or install a backdoor function that enables remote access. The restriction of physical access to critical systems is a reasonably well-understood problem. Vulnerabilities in software can usually only be abused temporarily, until the problem becomes more widely known and fixed. Containment mechanisms in computer architectures and operating systems designed to prevent the violation of security policies by vulnerable or maliciously designed application software are another important form of protection, although the use of such techniques in mass-market standard software is still in its infancy.

Simpler routes of unauthorized access are becoming less feasible, while at the same time the value of electronically processed financial, intellectual, military, administrative, and personal information continues to grow. As a result, the focus of determined attackers can be expected to shift towards more technically advanced subversive technologies, such as the exploitation of compromising emanations. Techniques that have so far been economically feasible only for well-funded intelligence agencies may well become attractive for criminals.

1.1 Historic background and previous work

1.1.1 Military activities

It has been known to military organizations since at least the early 1960s that computers generate electromagnetic radiation that not only interferes with radio reception, but also leaks information about the data being processed. Known as *compromising emanation* or *Tempest* radiation, the unintentional electromagnetic broadcast of data has been a significant concern in sensitive military and diplomatic computer applications since then. The US military code word *Tempest*¹ referred originally to a classified US government

¹Names for classified US military programs such as “Tempest” are usually random dictionary words, although numerous attempts to resolve the word as an acronym can be found.

program aimed at studying such *emission security (EMSEC)* problems and at developing protection standards. It has since then become a synonym for compromising emanations.

An early example of the military exploitation of compromising emanations from communications equipment occurred during World War I, when the German army successfully eavesdropped on enemy voice communication from the earth loop current of allied battlefield phone lines [2, 3, 4]. Only a single insulated wire was used to interconnect field phones at the time, in order to reduce the weight of cable drums that the signal troops had to carry. The return current went through the ground and eavesdroppers managed to pick up the resulting voltage drop with valve amplifiers connected to well-spaced ground spikes. In 1915, the French and British forces noticed the problem and started to experiment with this effect. They eventually implemented countermeasures such as placing earth connections only hundreds (and later even thousands) of meters behind the front trenches, using twisted-pair cables, reducing the line currents and limiting the sensitivity of information communicated via field phones. Their attempts to use the same eavesdropping technique later revealed that the German army already had implemented such precautions as well.

In his book “Spycatcher” [5, pp. 109–112], former MI5 scientist Peter Wright recounts the origin of “Tempest” attacks on cipher machines. In 1960, Britain was negotiating to join the European Economic Community, and the Foreign Secretary was worried that France’s president De Gaulle would block Britain’s entry. He therefore asked the intelligence community to determine the French negotiating position. They tried to break the French diplomatic cipher and failed. MI5 and GCHQ scientists then installed a broad-band radio-frequency tap on the telex cable that came out of the French embassy in London and routed the signal to a room in the nearby Hyde Park Hotel. There they noticed that the enciphered traffic carried a faint secondary signal and constructed equipment to recover it. It turned out to be the plaintext, which somehow leaked through the cipher machine.

Sensitive government systems today employ expensive metallic shielding of individual devices, cables, rooms and sometimes entire buildings [6]. Even inside shielded environments, the *red/black separation principle* has to be followed: *Red* equipment carrying confidential data (such as a computer terminal) has to be isolated by filters and shields from *black* equipment (such as a radio or modem) that handles or transmits unclassified data. Equipment with both *red* and *black* connections, such as cipher machines and multilevel secure workstations, requires particularly thorough testing.

The US government defined its first national compromising-emanations test standards “NAG1A” and “FS222” in the 1950s and 1960s [7]. A revision titled “National Communications Security Information Memorandum 5100: Compromising Emanations Laboratory Test Standard, Electromagnetics” followed in 1970 and a later revision “NACSIM 5100A” was defined in 1981. The names of the standards keep changing. “NSTIS-SAM TEMPEST/1-92” appears to be the current incarnation, of which extracts were declassified in 1999 after a Freedom-of-Information-Act request made by John Young [8]. The released parts, however, reveal mostly only material that can also be found in the open computing, security, and electromagnetic-compatibility literature, while the actual emanation limits, test procedures, and even definitions of some terms remain classified as military secrets.

The recently declassified tutorial “NACSIM 5000: Tempest Fundamentals” [9] lists further US standards such as the rationale for the NACSIM 5100A limits (NACSIM 5002), the

“Compromising Emanations Laboratory Test Standard, Acoustics” (NACSEM 5103) and its rationale (NACSEM 5104) [10, 11, 12, 13, 14]. All these are, along with the NATO equivalent “AMSG 720B”, still classified documents and were therefore not accessible to the author [15, 16].

In Germany, even the names of government standards on emission security, which are administered by the *German Information Security Agency (BSI)*, are kept secret and only a few brief leaflets and articles with overview information are available [17, 18, 19, 20, 21, 22, 23]. Descriptions in some published German patents [24, 25] suggest that the measurement techniques employed in some of the tests might include not only receivers, spectrum analyzers and oscilloscopes, but also devices to perform real-time cross-correlation measurements between *red* signals measured directly inside the target system and the noisy and distorted signals received from external sensors such as antennas, power-line taps, or microphones.

1.1.2 Open literature

Electromagnetic radiation as a computer security risk was mentioned in the open literature as early as 1966 [4, 26], but without technical details on specific risks and eavesdropping techniques. Probably the first more detailed public description of compromising emanation risks appeared in Sweden around 1982 and in 1984 a Swedish government committee issued an 18-page booklet [27] in Swedish that informed the wider business community about the threats of acoustic, radiated, conducted and, in particular, video emissions [28].

The concept was brought to the attention of the broader public by a 1985 paper [29] and a 5-minute TV demonstration on the BBC program “Tomorrow’s World”, in which van Eck demonstrated that the screen content of a video display unit could be reconstructed at a distance using low-cost home built equipment, namely a TV set whose sync-pulse generators were replaced by manually controlled oscillators. A couple of conference presentations followed that discussed the finding and some shielding techniques [16, 30, 31].

Smulders showed that shielded RS-232 cables can be eavesdropped at a distance [32]. Connection cables form resonant circuits consisting of the induction of the cable and the capacitance between the device and ground; these are excited by the high-frequency components in the edges of the data signal, and the resulting short HF oscillations emit electromagnetic waves.

The designer of a commercial anti-Tempest jamming transmitter [33] claimed that an eavesdropper standing near an automatic teller machine with fairly simple radio equipment could pick up both magnetic stripe and PIN data [34, 35, 36]. Card readers and keypads are typically connected to the processor unit using serial links, which might emanate sufficiently strong pulses during bit transitions, but no detailed description of an attack has been published so far.

The two *SEPI* conferences organized by the Italian telecommunications research organization *Fondazione Ugo Bordoni* in 1988 and 1991 [37, 38] focused on electromagnetic information security and left a collection of proceedings papers, most of which are in Italian. Briol [39] shows brief examples of power-line conducted and acoustic compromising emanations from a 5×7 dot matrix printer, Hodara [40] discusses techniques for tapping into optic cables by bending fibers, and Demoulin et al. [41] provide another discussion of the electromagnetic radiation emitted by shielded data cables. Köksaldı et al. [43] discuss

some preprocessing steps towards automatic character recognition from compromising video emanations.

Interest in compromising emanations and other aspects of hardware security surfaced again in the 1990s with the mass-market introduction of tamper-resistant cryptographic modules. These devices are intended for applications that need to protect stored cryptographic keys, and the execution of algorithms that operate on them, against unauthorized access.

The most popular form of portable cryptographic module is the smartcard [44, 45], a credit-card shaped plastic card with embedded microcontroller. The typical interfaces are either five electrical surface contacts (for ground, power supply, reset, clock, and a bi-directional serial port) or an induction loop. Typical smartcard processors are 8-bit microcontrollers with a few hundred bytes of RAM and 5–30 kilobytes of ROM and NVRAM. Other popular form factors for tamper-resistant modules include battery-like small steel cans (e.g., the Dallas Semiconductor *iButton*), CardBus/PCMCIA modules, and various PCI plug-in cards (e.g., the IBM 4758 Cryptographic Coprocessor).

Cryptographic modules are used both in applications with tamper-resistance and tamper-evidence requirements [46]. Tamper resistance means that stored information must remain protected, even when the attacker can work on several samples of the module undisturbed for weeks in a well-equipped laboratory. Tamper evidence is a weaker requirement in which the regular holder of the module must merely be protected against unnoticed access to information stored in the module.

There are several popular applications for tamper-resistant smartcards. Operators of some pay-TV conditional-access systems hand out millions of cards to customers, each of which contains the key necessary to descramble some subscription TV service [47]. Pirates who manage to extract the key from one single issued card can use it to produce and sell illicit clone cards. Most proposed forms of digital rights management (DRM) mechanisms are based on some form of tamper-resistant element in the user's system.

The designer of a tamper-resistant module must assume that the attacker has full control over the physical environment in which the card operates. Larger form factors like PCI plug-in cards offer enough space to implement comprehensive protection measures. This can include a tamper-sensing membrane that blocks all mechanical access paths to the integrated circuits and auxiliary electronics [52], metal shielding against electromagnetic emanations, as well as low-pass filters on connections to block conducted compromising emanations and provide immunity against interference attempts. Secrets are typically stored in battery-backed static RAM, such that they are only preserved as long as power to also operate a tamper-sensing alarm mechanism is available. Additional sensors for temperature and ionizing radiation can be added to ensure that the SRAM is only operated under conditions in which shorting the power-supply pins will lead to a rapid destruction of stored information [54, 55]. Continuous timer-triggered rotation of SRAM content helps further to reduce long-term data remanence risks.

Smartcards, on the other hand, are particularly vulnerable. Their 0.8 mm thick package requires tamper-sensing membranes to be implemented directly on the chip surface, extremely close to the protected circuitry. The absence of a continuous power supply means that stored secrets must be kept in non-volatile memory, and, as a result, the preservation of stored secrets is not intrinsically tied to the availability of power to operate a tamper alarm with a zeroization mechanism. Smartcards receive not only their unfiltered power

supply via the outside interface, but also their clock and reset signals. In particular, the power supply of CMOS circuits is a rich source of information, which provides for every clock edge an estimate for the number of gates that change their output state. This not only allows the type of instruction currently executed to be determined – a useful capability for triggering glitch attacks and tracing execution paths through conditional branches – but also provides information about accessed memory values.

Research interest in compromising emanations from smartcards increased significantly when Kocher, Jaffe, and Jun [56] demonstrated the power of combining the analysis of high-frequency current fluctuations with cryptanalytic techniques on block ciphers. In their *Differential Power Analysis* attack, they demonstrated the reconstruction of DES subkey bits merely from access to a number of known plain or cipher texts, the corresponding power-line current curves and knowledge of the cipher algorithm being used. They showed that it is feasible to evaluate power-line information without prior reverse engineering of the low-level design of the executed software and that it is instead sufficient to look for correlations with single bits in intermediate results of the executed algorithm. The correlation process takes care of locating the specific machine instructions that leak the compromising energy. A number of improvements of the attack, attacks on other algorithms and discussions of countermeasures have been published since then [57, 58], including variants that measure magnetic-field fluctuations above the chip surface [59, 60, 61, 62], as well as an attack on an SSL accelerator module inside a closed server from 5 m distance [63].

1.2 Motivation and scope

My interest in compromising emanations was sparked by a TV report on the “Inslaw/PROMIS affair” by the investigative journalist Egmont Koch [64, 65].² It contained an interview with Michael J. Riconosciuto, who claimed to have developed for the CIA a revision of the PROMIS database software that the US Department of Justice was to sell to law-enforcement and intelligence agencies as well as banks in about 80 countries. According to Riconosciuto, his job was to introduce backdoor functions into PROMIS that allowed the US intelligence agencies to access the information stored in the foreign databases that PROMIS installations all over the world accessed. He explained that PROMIS installations typically had no network connection, therefore his modifications continuously broadcast the content of database files as electromagnetic radiation in a way that allowed reception at a distance. However, no implementation details were provided. In particular, it was not even clear whether the claimed broadcast backdoor was implemented in software only or involved modifications to the host hardware (e.g., DEC

²The details of this case have been widely discussed in the US media and various publications [66, 67, 68]. In the 1970s, the US Department of Justice funded the software company Inslaw to develop the public-domain software package PROMIS (Prosecutor Management Information System). This half-million lines of code COBOL program was designed to integrate data from a heterogeneous collection of databases, in particular in order to track information about individuals. After the initial publicly funded development phase had ended, Inslaw continued with the development of proprietary extensions and then discovered that the US Department of Justice had not only used this proprietary version without licensing it, but also fielded it in other government agencies and even sold it abroad. The case went to court and became the subject of a congressional hearing.

VAX).³

The information provided by Riconosciuto on the actual use of Trojan Horse software versions to enable the electromagnetic theft of confidential information by US intelligence agencies remains dubious and obscure. Nevertheless, the concept of software-controlled modulation of electromagnetic emissions from information processing equipment should clearly be of great interest to the computer security research community, and I described the first demonstration experiments of software-controlled emanations in 1998 [72]. The fact that the potential of such techniques had been recognized by the US military research community is at least reflected by the definition of a code word for this field of study in an NSA document that was partially declassified in 1999 [70]:

“TEAPOT: A short name referring to the investigation, study, and control of intentional compromising emanations (i.e., those that are hostilely induced or provoked) from telecommunications and automated information systems equipment.”

The implementation of a broadcast backdoor in system or application software could allow attackers to by-pass conveniently a large number of otherwise very solid computer security measures, such as physical access control (locked rooms with intruder alarms), “air gaps” (physical separation of hardware and network components), virtual networks and firewalls (network-infrastructure enforced limitations on reachability), mandatory access control (operating-system enforced restrictions on information flow between programs, users, and interfaces) and the encryption and authentication of data on communication channels. Suitably modulating broadcast information and augmenting it with forward-error-correction information would simplify reception significantly compared to the evaluation of information carried in merely accidental emanations. Like any other malicious backdoor or *Trojan Horse* functionality, compromising broadcast algorithms could be

- introduced by the original software developer,
- added somewhere in the distribution or maintenance chain,
- installed during temporary access to a system,
- introduced on less carefully protected systems by executing software of untrustworthy origin with potentially self-replicating malicious components (viruses, worms).

There might also be room for legal and even commercially viable applications of “Teapot” technology. For instance, Ross Anderson suggested an application for copyright enforcement [72]. Commercial software could deliberately modulate the software license number

³I contacted Koch soon after the broadcast and obtained several handwritten pages of diagrams and technical explanations that Riconosciuto had provided. Unfortunately, these diagrams did not turn out to be descriptions of actual software or hardware modifications related to PROMIS or reception and decoding equipment to pick up such information. They were just hand drawn copies of figures from an early 1970s textbook on information transmission using non-sinusoidal carrier signals [69, pp. 92,170,240,287–289], an early precursor technology of what is today better known as *direct sequence spread-spectrum modulation*, which forms the basis of wireless LANs, satellite navigation systems and military low-probability-of-intercept radios. Such techniques might in general be relevant for modulating information into RF computer emanations (see p. 75), but this did not provide further specific insights into the nature of the claimed modifications.

under which it is used into the PC's RF emissions, and detector vans operated by anti-software-piracy associations patrolling in business districts could look for buildings where the emanated signals suggested that several copies of a single-user license were being used concurrently. Similar RF detection systems already are used in some countries to find unlicensed TV viewers.

Considering the excitement that van Eck's findings created [26, 28, 71], and the significant investment in shielding by the diplomatic and defense community, it is surprising that so little further research on "Tempest" and "Teapot" attacks and defense has appeared in the open literature. There are several possible reasons. Well equipped radio-frequency laboratories are expensive and commercially available narrow-band receivers for radio amateurs or electromagnetic compatibility (EMC) testing are not well suited for work on compromising emanations. Specialized broadband receivers designed for the military market are somewhat expensive, export controlled, and rarely found in academic research environments. Purely theoretical contributions on the other hand are difficult due to the lack of published data about the information-carrying emanations of modern hardware.

Commercial use of "Tempest" technology is also marginal. Shielded PCs and peripherals are many times more expensive than standard models, and sales are typically export controlled. So it is no surprise that shielded facilities and equipment are practically never used outside diplomatic, defense, and perhaps some public-key infrastructure applications. There currently exist no civilian protection and test standards concerning compromising emanations and the occasional public demonstration of eavesdropping technology is mostly limited to repetitions of van Eck's experiments with modified TV sets, which do not usually provide a realistic demonstration of the eavesdropping risk of contemporary hardware. This makes it difficult even for computer security experts who are familiar with the available academic literature to find realistic data on the threats of compromising emanations. In particular display technology has undergone significant changes and data rates have increased by more than an order of magnitude since van Eck's paper.

The present study of compromising emanations has several goals:

- Compromising emanations are only briefly mentioned in today's computer-security textbooks, therefore it seems desirable to provide a brief introduction to related phenomena, measurement techniques, and equipment for readers with a background in computer science but not in radio communication and high-frequency technology.
- Experiments published with insufficient details in the 1980s need to be repeated with more recent display hardware, reception and signal processing equipment. This should provide the computer-security community with updated example data, to stimulate discussion on realistic threat assumptions and the need for further research and protection standards. I was in the lucky situation of obtaining from various sources measurement equipment that should at least be able to approximate what researchers in secret signal-intelligence laboratories might have had available less than a decade ago.
- Inspired by the claims surrounding the PROMIS affair, I wanted to perform and document a number of experiments on how software can affect electromagnetic emanations, looking at both attack and defense.
- Compromising emanations in the visible bands seem to have been neglected so far and I therefore became interested in studying whether, for instance, diffuse light

emitted by cathode-ray tubes carries enough high-frequency content of the video signal for the reconstruction of readable text, and what theoretical upper bounds for this eavesdropping risk could be established.

The eavesdropping demonstrations that I will describe are mostly laboratory simulations and theoretical discussions. Even though the equipment that I used allowed me to obtain the data necessary to discuss the feasibility of practical attacks, it lacked the portability, real-time signal processing features and directional gain necessary for evaluating emanations in a realistic covert setting.

All presented experiments were conducted on some of the PC hardware that was easily available in our department. They do not represent an effort to provide a comprehensive survey of the properties of a wide range of products and should be seen only as examples for test techniques that could equally be applied to a wider range of device types.

Without access to a shielded room, I had to rely on periodic averaging techniques and short antenna distances to reduce external interference in measurements. As a result, I restrict myself to discussing primarily the periodic refresh signals appearing in video display systems, even though initial experiments with printers, serial ports and system buses suggest that interesting demonstration experiments can be performed with these as well.

Chapter 2 and Appendix A provide a brief review of the foundations of electromagnetic theory, its units of measurement, transmission lines, broadband antennas, as well as a description of the receiver used for the experiments in later chapters. It also covers some concepts useful for the discussion of broadband impulse signals. Chapter 3 discusses the emanations of cathode-ray tube monitors driven with analog signals, including a more detailed treatment of some of the ideas that I presented originally in [72]. Chapter 4 discusses the emanations of the digital video interfaces that have been used more recently to connect flat-panel displays. Chapter 5 suggests and discusses RF emission-security laboratory test limits. Finally, Chapter 6 demonstrates a new optical eavesdropping technique, a slightly extended version of what I presented in [73], followed by a concluding Chapter 7.

Chapter 2

Foundations and test equipment

Two theoretical foundations – Maxwell’s equations and Fourier analysis – combined with practical engineering knowledge about antennas, transmission lines, radio receivers, and digital signal processing form the basis for understanding compromising radio-frequency emanations and their exploitation. Readers who do not feel particularly familiar with the basic notions of electromagnetic fields, high-frequency electronics and some of the associated units of measurement ($\text{dB}\mu\text{V}/\text{m}$, etc.) may first want to consult, for a brief revision, sections A.1, A.2, and A.3 in the appendix. This section introduces some of the radio-frequency measurement equipment used in later experiments. It also discusses some notions and practices of particular use for the quantitative characterization of unmodulated broadband impulse signals.

2.1 Antenna types

Simple dipole antennas are well suited for the reception of narrow-band signals, where a good estimate for the frequency of interest is known and the signal has a bandwidth of not more than about a tenth of the center frequency. The directional gain of dipoles can be increased by up to about 12 dB with the addition of director and reflector dipoles around the receiving dipole. This construction is known as a *Yagi-Uda antenna*, and forms – usually with a folded reception dipole – the most common VHF/UHF household TV antenna.

Dipoles and Yagi antennas tuned to an exact frequency are not necessarily the best choice for compromising emanations experiments though. Most forms of compromising emanations are broadband signals, which means that the lower and upper frequency limit of information carrying emanations can differ by a factor of two or more. Several broadband antenna types have been developed, which offer a reasonably constant impedance over a wide frequency range.

Yagi antennas are nevertheless an excellent choice for actual attacks in circumstances where the VHF/UHF center frequency with the highest information content is known precisely, the bandwidth is not much higher than 1/10 of the center frequency, and the attacker has the time to construct an antenna for a particular eavesdropping target. Where compactness of the antenna is not a major concern, entire arrays of Yagi antennas can be connected together to increase the effective aperture. This might be feasible, for




Figure 2.1: Examples of different broadband antenna types: log-periodic antenna (200–1000 MHz), discone (200–1300 MHz), active monopole (100 Hz–30 MHz), bi-conical antenna (30–300 MHz), and active ferrite loop (H-field, 100 Hz–30 MHz with four different ferrite rods).

example, if the eavesdropper can operate behind a number of large windows from a nearby building.

Broadband antennas, on the other hand, have a good impedance match over a wide frequency range, but lack good directional gain. They are best suited for laboratory work, where the emanations of a device have to be characterized for a threat assessment or for the preparation of an eavesdropping operation. Some compact examples can be seen in Fig. 2.1.

The first approach for constructing a broadband antenna is a dipole with very thick wires, for example, shaped as cylinders or cones. A dipole is an LC oscillator, whose frequency selectivity (quality factor) depends on the L/C ratio. Thick dipoles have larger capacitance and lower inductance; the resulting low L/C ratio makes them less frequency selective. An implementation of this idea is the bi-conical antenna (30–300 MHz) commonly used for VHF EMC field-strength measurements.

The second approach is the careful combination of a number of dipoles of varying length to cover the entire desired band, leading to the logarithmic-periodic antenna, which is commonly used for EMC field-strength measurements in the lower UHF range (200–1000 MHz). It consists of two metal rods, to which in alternating polarity a series of dipoles is connected. The entire antenna has a triangular shape as the length of a dipole

is proportional to its distance from the triangle's tip, which points in the direction of maximum gain. The gain factor of the antenna depends on the *periodicity* τ , the length ratio of neighboring dipoles, and it is typically about 3–6 dB (for $0.60 < \tau < 0.85$) better than that of a single dipole.

The third approach are horn antennas. They smooth the impedance mismatch between a transmission line and free space by gradually extending the geometry of the transmission line until its impedance approaches that of free space. This avoids the reflection of energy caused by any abrupt change of impedance. The discone antenna is an example that combines the principles of a horn antenna with that of a broadband dipole. It is widely used as an omnidirectional reception antenna for receivers that work over large parts of the UHF spectrum. The ground shield of a coaxial cable is connected to a metallic cone with an angle of 60° and a length $l = 0.35c/f_l$, where f_l is the lowest supported frequency. The inner conductor is connected to a metal disc above the cone, whose diameter is 70 % of the maximum diameter of the cone. According to [74], such a discone antenna provides a SWR < 2 to a 50Ω cable over a frequency range from f_l to $10 \cdot f_l$.¹ An ultra-broadband (2 kHz to 0.8 GHz) horn antenna designed for low-distortion analysis of time-domain waveforms is described in [42].

The fourth approach is the active antenna. Instead of relying on the inductance and capacitance of a dipole matching the impedance of the transmission line, a monopole much shorter than the wavelength is used to probe the electric field strength. An amplifier circuit with high input impedance (e.g., $10 \text{ M}\Omega$) is connected directly to the base of the monopole. It converts the high-impedance voltage appearing on the antenna rod into a 50Ω output impedance signal suitable for transmission via a coaxial cable.

Active antennas are the only compact antennas available for sensitive field-strength measurements in the 100 Hz–30 MHz range. Depending on whether a metallic rod or a ferrite core loop is connected to the amplifier, an active antenna can measure either E- or H-field strength, a distinction that is of particular interest in the frequencies below the VHF range, where an eavesdropper might well be located in the near-field of an emitting device. The active E- and H-field antenna that is part of the Dynamic Sciences R-1150-10A Portable Antenna Kit that I used is a two-stage zero-gain amplifier using two n-channel JFET transistors (2N5397 and MPF820) [80]. It features two separate amplifiers, one with $10 \text{ M}\Omega$ input impedance for connection of the E field monopole (1 m), the other with $200 \text{ k}\Omega$ input impedance for connection with one of four interchangeable ferrite loops (100 Hz–50 kHz, 50 kHz–1.5 MHz, 1.5–10 MHz, 10–30 MHz).

The choice of antenna not only affects the resulting directional gain of narrow-band signals in different frequency bands; it also determines the waveform of a pulse provided to the receiver. Figure 2.2 shows examples of the response pulses and their relative amplitude spectra for a number of different antennas, when they received a subnanosecond pulse. The amplitude spectra illustrate the frequency-domain filtering performed by these antennas. The bi-conical transfers hardly any energy outside the 20–400 MHz range, the log-periodic antenna receives only above 180 MHz and the simple dipole provides in general lower voltages and peaks near 300 MHz. All these spectra are of course also affected by the transmission characteristics of the transmitting discone antenna at 2 m distance, as well as by reflections from any nearby metal surfaces. The time-domain signals show that with a larger antenna bandwidth, the overall envelope of the impulse becomes shorter. Most

¹although that performance was apparently not quite reached by the low-cost antenna tested with time-domain reflectometry in Fig. A.4 (c) shown on page 155

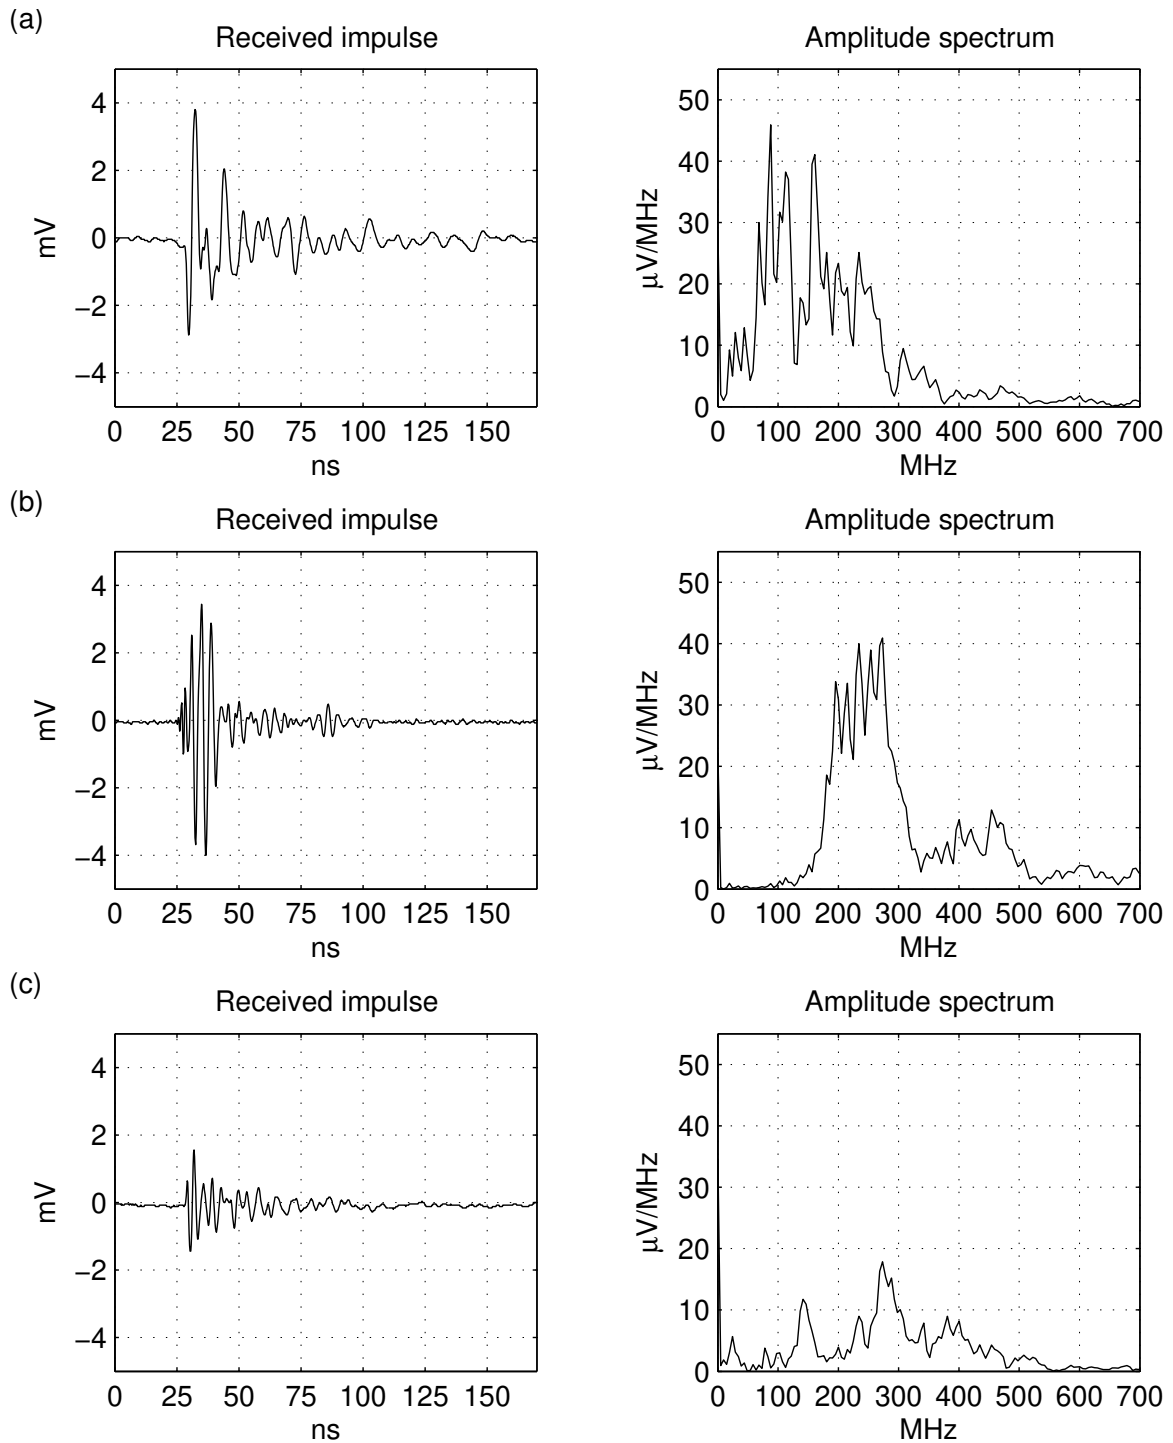


Figure 2.2: A subnanosecond impulse with $60 \text{ dB}\mu\text{V}/\text{MHz}$ is transmitted via a Watson WBD-40 discone antenna. These diagrams show the output signals provided at 2 m distance by three different receiving antennas connected to an oscilloscope with 500 MHz bandwidth: (a) bi-conical 20–200 MHz and (b) log periodic 200–1000 MHz antenna in the Dynamic Sciences R-1150-10A Portable Antenna Kit, as well as (c) a simple 393 mm ($\lambda/2$ for 357 MHz) dipole.

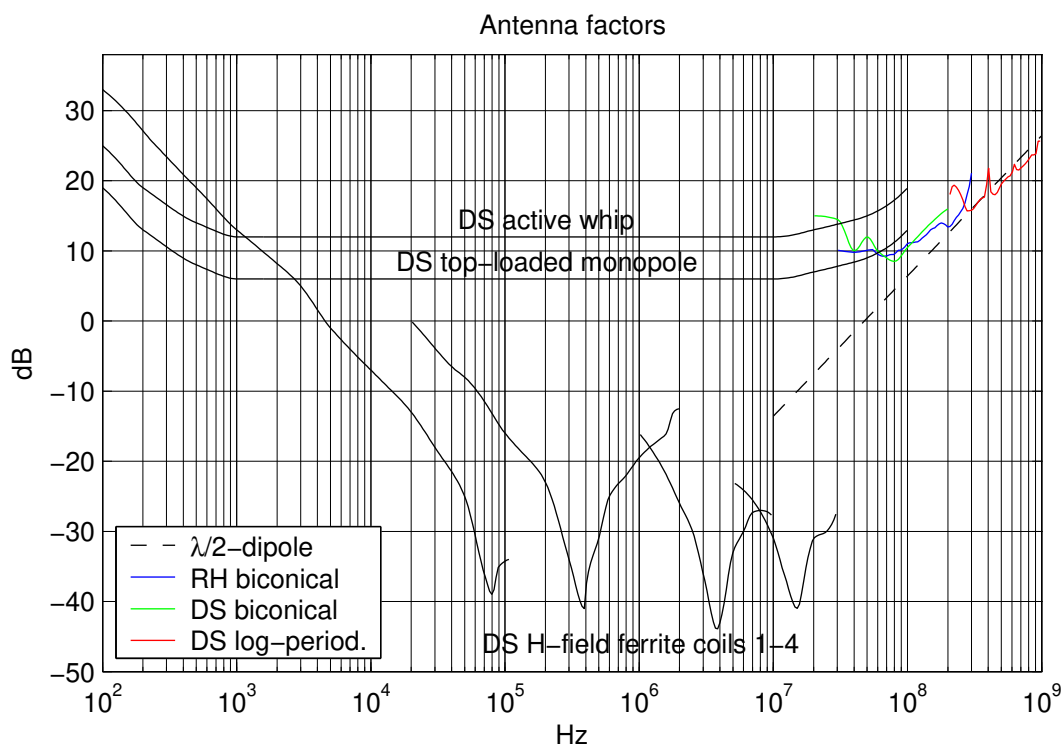


Figure 2.3: Antenna factors for the Rolf Heine biconical antenna and the components of the Dynamic Sciences R-1150-10A Portable Antenna Kit according to the product documentation. The figure has to be added to the measured $\text{dB}\mu\text{V}$ receiver input voltage to obtain the measured $\text{dB}\mu\text{V}/\text{m}$ in the case of an electric field antenna or $\text{dB}\mu\text{A}/\text{m}$ for a magnetic field coil.

of the pulse energy is located within the first 15 ns for the log-periodic antenna, while the equivalent decay takes more than twice as long for the simple dipole.

Commercially available antennas for EMC field strength measurements like those shown in Fig. 2.1 come with calibration charts that provide an *antenna factor* for each frequency (Fig. 2.3). The “unit” of this antenna factor is dB/m and this number has to be added to the $\text{dB}\mu\text{V}$ measured with the help of a receiver to obtain the measured $\text{dB}\mu\text{V}/\text{m}$ field strength. The antenna factor found in the product literature is nothing but the effective length $l_e = \frac{U}{E}$ of the antenna written as $(-20 \cdot \log_{10} l_e/2)$ dB/m .

2.2 Receivers

It is in principle possible to connect an antenna directly via an amplifier to an analog-to-digital converter (ADC) and perform any further processing digitally. I used this approach to record the impulses shown in Fig. 2.2. The signal actually received was distorted by a lot of background noise, but due to the close proximity of the transmitter, the intended signal was clearly recognizable against the background. By averaging a sequence of 1000 of these impulses, I managed to eliminate the background noise from Fig. 2.2 almost completely, without using a shielded chamber for the experiment.

For weaker signals of interest, the signal-to-noise ratio at the antenna output can be significantly lower than the dynamic range of the ADC, and analog preconditioning prior to digitization becomes essential before information buried in noise can be retrieved.

If $S(f)$ is the Fourier spectrum of the signal to be received and $N(f)$ is the spectrum of the unwanted background noise, then a filter with the characteristic

$$H(f) = \frac{|S(f)|^2}{|S(f)|^2 + |N(f)|^2}, \quad (2.1)$$

also known as a *Wiener filter*, will optimize the signal-to-noise ratio if applied to a received signal $c(t) = s(t) + n(t)$, where s and n are the signal and noise, respectively. The Wiener filter is least-square optimal in the sense that the filtered signal $\tilde{S}(f) = C(f) \cdot H(f)$ minimizes the error $\int |\tilde{s}(t) - s(t)|^2 dt$. [76, 77]

In practice, it is too tedious to build a good analog approximation of a Wiener filter for each encountered signal and noise spectrum. Usable results can be obtained with a crude approximation. The simplest approach is a band-pass filter that selects out of the entire spectrum the range with the highest signal-to-noise ratio (SNR). A compromise has to be made, because while smaller bandwidths might maximize the average SNR in the pass band, larger bandwidths lead to shorter pulse widths, which determine the maximum digital signal bitrate that can be reconstructed after the filter.

Apart from applying an approximation of a Wiener filter, rectification and low-pass filtering are further useful steps that can be performed conveniently before digitization. The received impulse waveform in Fig. 2.2 (b) shows oscillations with a half-period of 2 ns. This means that in order to average several repetitions of the waveform, their recordings have to be synchronized much better than 2 ns, otherwise phase correlation is lost and the instances of the recorded impulse will cancel each other out during the averaging process. Rectifying and low-pass filtering the result not only significantly reduces the requirement for exact synchronization, it also allows the use of much lower sampling frequencies during data acquisition, without risking the emergence of aliasing frequencies.

The preprocessing steps suggested in the previous paragraphs, bandpass filtering, rectifying and low-pass filtering, are exactly the steps performed in an AM radio. AM receivers are usually used to receive amplitude-modulated narrow-band signals of the form

$$s(t) = A \cdot \cos(2\pi f_c t) \cdot [1 + m \cdot v(t)] \quad (2.2)$$

where f_c is the carrier frequency, and $v(t)$ is the broadcast signal. Common bandwidths are in the order of 10 kHz for voice communication and 10 MHz for video transmission.

Both tunable and non-tunable forms of AM detectors are possible. In a non-tunable receiver, a series of (potentially switchable) fixed-frequency filters are used to band limit the antenna signal before it is amplified and provided to a detector circuit that rectifies and low-pass filters it. Since it is difficult to construct high-quality tunable filters, tunable receivers usually apply the superheterodyne principle. The antenna signal is filtered only by a very wide preselector and then multiplied with a sine wave generated in a local oscillator (LO). The product of two harmonic waves

$$\cos(2\pi f_1 t) \cdot \cos(2\pi f_2 t) = \frac{1}{2} \sin[2\pi(f_1 + f_2)t] + \frac{1}{2} \sin[2\pi(f_1 - f_2)t] \quad (2.3)$$

is the sum of another two harmonic waves, with frequencies corresponding to the sum and difference of the frequencies of the multiplied waves. So after multiplication of the received signal with the output of the local oscillator, the resulting waveform contains the received frequencies shifted to two intermediate frequency ranges that depend on the




Figure 2.4: Bottom: Dynamic Sciences R-1250 wide-range receiver. Middle: R-1250-30 HF preselector for 15 passbands in the 0.25–20 MHz range (left) and R-1250-20A wide-band AM detector (right). Top: R-1160C calibrated impulse generator.

frequency of the local oscillator. An intermediate frequency filter selects one of these two terms and reduces the bandwidth of the signal to the desired size. By tuning the local oscillator, the fixed intermediate frequency (IF) filter can select a frequency window out of the spectrum that can be tuned conveniently over a wide frequency range with the accuracy of a PLL-synthesized local oscillator.

The receiver I used is a model R-1250 produced by *Dynamic Sciences* in California (Fig. 2.4). According to the product literature [78], it fulfills the (confidential) requirements of the NSA Tempest measurement standard NACSIM 5100A. It differs from more commonly available radio receivers in a number of parameters:

- This receiver can be tuned continuously from 100 Hz to 1 GHz. The frequency resolution is 0.1 Hz up to 250 kHz, 1 Hz up to 20 MHz and 100 Hz above that.
- It offers 21 intermediate-frequency bandwidth selections, increasing in 1-2-5 steps from 50 Hz to 200 MHz. Other commercially available receivers that are not purpose-built for the analysis of compromising emanations usually do not support

IF bandwidths above 8 MHz (TV), which is more than an order of magnitude below the pixel clock of current video display systems.

- It provides input attenuators in 10 dB steps from 0 to 100 dB and has an especially robustly designed 50 Ω RF input. This makes the receiver suited for monitoring signals on power lines (via a simple high-pass filter), where high-voltage spikes could destroy more fragile input stages in other receivers.
- The automatic gain control circuit can be deactivated. The RF gain can be manually adjusted over 50 dB and the pre-detection gain over 30 dB. Together with the attenuator settings, this allows the overall voltage gain of the receiver to be adjusted over 180 dB. In other words, the highest and lowest amplification factors differ by 10^9 .

The R-1250 is a superheterodyne receiver and the intermediate frequency filters are centered at either 200 kHz or 30 MHz, depending on the selected tuning frequency and bandwidth. It switches automatically between different preselection filters and mixers. These cover the seven bands 100 Hz–250 kHz–20 MHz–200 MHz–350 MHz–550 MHz–750 MHz–1 GHz. In addition, there is a 20 MHz–1 GHz band which is upwards mixed to an intermediate frequency of 1.47 GHz for the extra-wide bandwidths of 50, 100, and 200 MHz, which are useful for the analysis of emanations from video displays and other high-speed digital systems.

The 200 kHz or 30 MHz intermediate frequency signal is directly provided as an output signal, but can also be fed into a linear AM, logarithmic AM, BFO, or FM demodulator. The demodulated signal is made available both at the full video bandwidth (about half the IF bandwidth), as well as a filtered audio signal for headphone monitoring of narrow-band signals. The beat-frequency oscillator (BFO) mode makes continuous wave signals (constant sine tones) in the spectrum audible by adding a pseudo-carrier signal before the AM demodulation, such that the signal appears as an audible tone and not as a DC voltage on the audio output. Most functions of the receiver can be remote-controlled via an IEEE-488 bus connector.

The receiver also has two built-in analog post-processing functions referred to as *slideback* and *pulse stretch* that can be applied to the demodulated signal. One of these was meant to assist the operator of an analog oscilloscope connected to the AM video output in comparing voltage levels, the other helps to make short spikes more visible. These functions are probably obsolete today, since any modern digital oscilloscope provides with on-screen cursors and peak-detection acquisition modes far more accurate and convenient alternatives.

Comparable specialized wide-band receivers for compromising emanations analysis are also manufactured by *Rohde & Schwarz* in Germany (e.g., FSET7, FSET22) [79].

2.3 Receiver calibration

There can be many different center-frequency and bandwidth dependent signal paths through a measurement receiver. Each of these varies in gain, adds non-linearities, and involves a vast number of electronic components that affect the signal characteristics. The

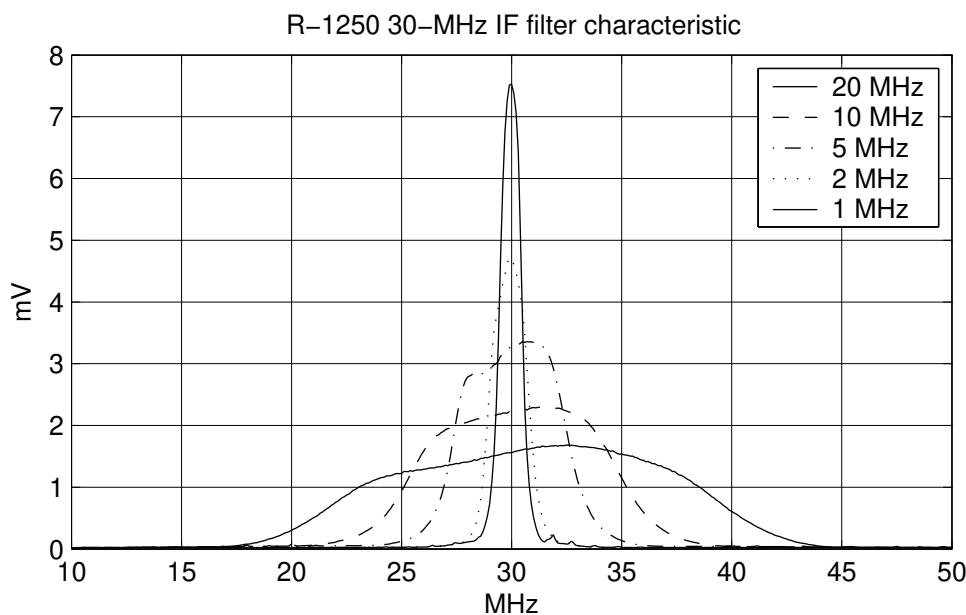


Figure 2.5: Transfer function amplitudes of five IF filters in the R-1250 receiver.

parameters of these components can vary considerably with temperature, stress damage, and age. As a result, the actual signal amplification provided by a receiver cannot be estimated reliably from how its gain is currently set. Signal strength measurements have to be made in relation to a calibration source. To do so, both the signal transducer (antenna, power-line clamp, etc.) and the calibration source are connected alternately to the measurement receiver, which should have two switchable input ports for that purpose. The signal level is adjusted at the calibration source until it results in the same output voltage on either the intermediate-frequency or demodulated output as the measured signal. The reading on the calibration source then provides the measurement.

The function $H(f) = \frac{U_o}{U_i}$ is the characteristic transfer function of a filter if, when applying an input voltage $u_i(t) = \text{Re}(U_i \cdot e^{2\pi j f t})$, we observe an output voltage $u_o(t) = \text{Re}(U_o \cdot e^{2\pi j f t})$. It is related to the impulse response $h(t)$ of the filter, that is the function with which the filter convolves its input, via the Fourier transform:

$$H(f) = \int_{-\infty}^{\infty} h(t) e^{2\pi j f t} dt = \mathcal{F}\{h\}(f) \quad (2.4)$$

$$h(t) = \int_{-\infty}^{\infty} H(f) e^{-2\pi j f t} df = \mathcal{F}^{-1}\{H\}(t) \quad (2.5)$$

If we consider instead of a simple band-pass filter an entire superheterodyne receiver, then the frequencies of the antenna input and IF output signals will be shifted by the receiver's local oscillator frequency and we might find $|H(f)| \gg 1$ due to the RF and IF amplifiers found in a receiver. Apart from that, a receiver's IF output can still be approximated in the Fourier domain as just a linear filter transfer function that is multiplied with a frequency-shifted version of the antenna input signal spectrum.

Figure 2.5 shows the transfer function of an R-1250 receiver for five different 30-MHz IF filter bandwidth settings. For this measurement, I connected the receiver to an impulse generator (40 dB μ V/MHz pulse strength, 1 kHz pulse rate), which provided a test signal

with uniform power distribution around the tuning frequency (50 MHz). An Anritsu MS2601B spectrum analyzer recorded the amplitude spectrum at the receiver's IF output port, which is what the diagram shows. The spectrum analyzer was set to a resolution bandwidth (100 kHz) larger than the pulse rate, such that the line spectrum of the periodic pulse input appears to be continuous. The video bandwidth was not limited, such that we see the actual peak voltage of the spectrum analyzer's IF signal. The RF gain setting on the receiver is identical for all five shown IF bandwidth settings.

Alternatively, the transfer function could be measured without a spectrum analyzer by supplying a sine-wave signal to the input and observing the IF output voltage at different tuning frequencies, where either the signal source or the receiver can be tuned over the frequency range of interest.

2.3.1 Impulse bandwidth

Ideally, the transfer function of a band-pass filter would be symmetric around its center frequency \hat{f} , but as Fig. 2.5 shows, in practice this is only approximated due to component tolerances. The 3 dB bandwidth of a filter is the frequency range over which $|H(f)|$ does not fall below $\frac{1}{\sqrt{2}} \cdot |H(\hat{f})|$.

To measure the overall field strength of a narrow-band signal, such as a continuous wave (CW) tone or an AM/FM broadcast signal, we adjust the 3 dB IF-filter bandwidth of the receiver such that it entirely covers the bandwidth of the examined signal, but still suppresses any out-of-band noise. As Fig. 2.5 shows, the frequency characteristic of the IF filter in a receiver, which usually consists of two critically coupled LC circuits, has a more or less flat plateau. Therefore, as long as the measured signal has a smaller bandwidth than the IF filter, all its spectral components will be amplified by approximately the same factor, characterized by the height $|H(\hat{f})|$ of the IF filter plateau. The exact roll-off shape of the filter characteristic outside the 3 dB bandwidth will not significantly affect the output signal strength, as there is no input signal in that part of the spectrum. So we just have to connect a sine wave generator adjusted to the tuning frequency of the receiver to its antenna input and then adjust its amplitude (usually expressed in μV_{rms} or dBm) until the peak-to-peak voltages seen with an oscilloscope on the intermediate frequency (IF) output of the receiver are identical to those seen with the antenna. The resulting amplitude setting on the sine-wave generator becomes the measured signal level at the antenna input. It can be converted into a field strength by applying the antenna factor for the respective frequency range, as explained in Sect. 2.1.

A sine-wave generator alone, however, is not a suitable calibration source for broadband pulses, which are of particular interest in the investigation of compromising emanations. Since the Fourier transformed representation of a sine wave is just a (theoretically infinitely narrow) spike in the frequency spectrum, it samples the transmission characteristic of a filter only at one single frequency. A pulse signal on the other hand (see Fig. A.2) has in its Fourier transformed representation roughly uniform voltages up to a frequency of at least $1/(4T)$, where T is the pulse width. As a result, the entire area of the receiver's IF filter characteristic (see Fig. 2.5) including the roll-off outside the 3-dB bandwidth determines the output voltage generated for a pulse, and not just the input/output voltage ratio at the center frequency \hat{f} .

In order to quantify the strength of a pulse signal with the help of a sine-wave generator as a calibration source, we first have to characterize the *impulse bandwidth* [81, 90, 8] B_i of




Figure 2.6: This figure shows the IF (30 MHz) and AM demodulator output that the R-1250 receiver provides after receiving a subnanosecond impulse for five different IF bandwidths.

the receiver and its IF filter. This is the bandwidth of an idealized filter with rectangular characteristic that has, at its center frequency, the same output/input voltage ratio as the filter to be characterized. It can be calculated by dividing the area under the plot of the filter transfer function by the peak height:

$$B_i = \frac{1}{|H(\hat{f})|} \int_{-\infty}^{\infty} |H(f)| df. \quad (2.6)$$

For example, numerically integrating the transfer function amplitudes measured with a spectrum analyzer in Fig. 2.5 leads to impulse bandwidths of 1.3, 2.2, 5.8, 10.1, and 16.1 MHz.

The NSA “Tempest” measurement standard [8] describes in its declassified section K.1 as part of the calibration procedure for impulse generators an alternative time-domain technique for determining the impulse bandwidth of filters and receivers. It is based on the fact that the length of a filter’s impulse response is inversely proportional to its bandwidth. Figure 2.6, for example, shows the IF and AM demodulator output impulse responses of the R-1250 receiver at five different bandwidth settings, when a less than one nanosecond long voltage pulse is applied to the antenna input. A 1 MHz filter results in an about $(1 \text{ MHz})^{-1} = 1 \text{ μs}$ long impulse, etc.

Figure 2.7 demonstrates this time-domain technique for determining the impulse bandwidth. We use an oscilloscope to record the waveform at the IF output of the receiver or the output of a band-pass filter. We integrate the area covered by the waveform pattern and divide it by the maximum height of the pattern (the positive half each time). The result is the width of a rectangular pulse that occupies the same area as the envelope of the positive waveform half. If the waveform shows several lobes, we add the odd-numbered ones and subtract the even-numbered ones. The envelope shown in Fig. 2.7 was calculated numerically via single-side-band demodulation of the recorded IF waveform. The dashed rectangular pulse has the same peak voltage and area as the envelope; for the 5 MHz filter it is $T = 177 \text{ ns}$ long. The reciprocal value $(177 \text{ ns})^{-1} = 5.7 \text{ MHz}$ is the resulting impulse




Figure 2.7: An impulse passed through a superhet receiver with 5 MHz IF bandwidth results in these signals on the IF and AM demodulator outputs. Like in Fig. 2.8, we determine the width T of the rectangular pulse with equal area. From that, we can determine the *impulse bandwidth* of the receiver: $2F = 1/T = (177 \text{ ns})^{-1} = 5.65 \text{ MHz}$.




Figure 2.8: An impulse passed through a rectangular low-pass filter with cut-off frequency F and therefore bandwidth $2F$ results in a waveform $\frac{\sin 2\pi Ft}{2\pi Ft}$. The area of the filtered pulse is identical to that of a rectangular pulse of the same amplitude with width $T = \frac{1}{2F}$.

bandwidth, which is almost identical to the 5.8 MHz determined in the frequency domain in Fig. 2.5. The equivalent rectangular pulse could, in principle, also be determined from the AM demodulator output, but the additional distortions caused by the baseband filter just add another error source, therefore this is not recommended.

In order to understand why the time and frequency domain techniques are roughly equiv-

alent, we consider an ideal rectangular low-pass filter

$$H(f) = \begin{cases} 1 & \text{if } |f| < F \\ 0 & \text{otherwise} \end{cases} \quad (2.7)$$

which has an impulse response

$$h(t) = \int_{-\infty}^{\infty} H(f) e^{-2\pi ft} df = \int_{-F}^F e^{-2\pi ft} df = \frac{\sin 2\pi Ft}{2\pi Ft} \quad (2.8)$$

as shown in Fig. 2.8. The width of a rectangular pulse of equal height and area is

$$T = \frac{1}{\max_t h(t)} \int_{-\infty}^{\infty} h(t) dt = \int_{-\infty}^{\infty} \frac{\sin 2\pi Ft}{2\pi Ft} dt = \frac{1}{2F} \quad (2.9)$$

The low-pass filter H with the cut-off frequency F has a bandwidth $2F$ (from $-F$ to F). The half-bandwidth with negative frequencies is for a low-pass filter just a mathematical convenience, but it becomes explicitly visible as a range of positive frequencies when a rectangular band-pass filter \tilde{H} with center frequency \hat{f} and bandwidth $2F$ is written as $\tilde{h}(t) = h(t) \cdot \sin 2\pi \hat{f} t$. So we have shown that the bandwidth of a rectangular low-pass or band-pass filter is always exactly the reciprocal value of the width of the rectangle in the time domain with equal area and height as the envelope of the filter's impulse response. If the transfer function of a real filter $F(f)$ has a less sharp roll-off than that of an idealized rectangular filter $H(f)$ then it can be approximated by the convolution $(G * H)(f)$ of a rectangular filter characteristic $H(f)$ with that of some narrower-bandwidth transfer function $G(f)$. In the time domain, this corresponds to the multiplication of the impulse response $h(t)$ with $g(t)$, and since $G(f)$ is assumed to be a much narrower spectrum than $H(f)$, $g(t)$ has to be much broader than $h(t)$ and therefore does not affect the shape of $h(t) \cdot g(t)$ significantly. Therefore, the described time-domain measurement of the impulse bandwidth can, within limits, also be applied to not-quite-rectangular filter characteristics such as those shown in Fig. 2.5.

With this technique, we can now calibrate the impulse bandwidth of a filter or receiver and use that in turn to quantify the spectral density of any impulse source.

2.3.2 Impulse strength

An impulse is a short transient rise and fall of voltage (not necessarily in that order). This waveform is of particular interest in the examination of compromising emanations, as every switching activity results in an electromagnetic impulse of some form.

The mathematically idealized representation of an impulse is Dirac's delta function

$$\delta(t) = \int_{-\infty}^{\infty} e^{\pm 2\pi j t f} df = \begin{cases} 0 & \text{if } t \neq 0 \\ \infty & \text{if } t = 0 \end{cases} \quad (2.10)$$

which is a single infinitely narrow and high pulse with finite area

$$\int_{-\infty}^{\infty} \delta(t) dt = 1. \quad (2.11)$$

The Fourier representation of an impulse is a continuous spectrum of cosine functions that all go through phase zero at the same location, where they all add up to form the

peak voltage of the impulse. At any other location, their respective phases are uniformly distributed and their voltage contributions cancel each other out. Since an ideal impulse contains energy at every part of the frequency spectrum, the output of a filter that results if an ideal impulse were applied at its input will characterize the filter completely. However, an idealized pulse would contain an infinite amount of energy and therefore any real-world impulse is necessarily band-limited and has a width larger than zero.

For very broadband impulses, the bandwidth of the measurement equipment can be smaller than that of the impulse, and as a result, the measured peak voltage of the impulse will depend equally on both the spectral density of the pulse as well as the filter characteristic of the measurement instrument.

The strength of a single impulse signal $v(t)$ can be quantified by integrating its voltage over time (unit: 1 Vs):

$$\sigma = \int_{-\infty}^{\infty} v(t) dt. \quad (2.12)$$

Figure A.2 shows a rectangular pulse of width T with an impulse strength $\sigma = UT$. The absolute value of the Fourier transform of that pulse is (within 10 % or 1 dB) constant at UT for frequencies $f < 1/(4T)$.

Commercially available impulse generators are calibrated in terms of a spectral density 2σ with unit 1 V/Hz [81]. Doubling the result of the Fourier transform takes into account that half of the spectrum amplitude of a real-valued signal is located in the negative frequency range. The spectral density of a pulse obtained in this way can be multiplied with the impulse bandwidth of a filter in order to predict the peak voltage (corrected for any gain that the filter might have at its center frequency) of the waveform that we obtain when we pass that pulse through this filter.

In order to characterize the strength I of an impulse, we can therefore use a receiver or filter for which we have determined the impulse bandwidth B as described in the preceding section. In addition, we need a calibrated sine-wave generator adjusted to the center frequency of the receiver or filter. We connect alternately the impulse source and the sine-wave generator to the receiver or filter and observe with an oscilloscope the output of the filter (preferably the output of the intermediate-frequency filter in the case of a receiver). We then adjust the amplitude A of the sine-wave generator until the peak-to-peak voltage of the filter output is equal for both the impulse and sine-wave source. The root-mean-square amplitude setting A of the sine-wave generator divided by the impulse bandwidth of the filter or receiver B is the spectral density of the pulse. Common units are $\mu\text{V}/\text{MHz}$ for voltage signals or $\mu\text{V}/(\text{m} \cdot \text{MHz})$ for electric field impulses, as well as the corresponding decibel scales.

In practice, the use of a sine-wave generator as a reference signal can be cumbersome, because measurement receivers contain a large number of filters and they switch between them automatically, such that the user cannot easily be certain of the current IF filter's impulse bandwidth. A more convenient alternative to a sine generator as a comparison source is the use of an impulse generator that has been calibrated in $\text{dB}\mu\text{V}/\text{MHz}$. It provides a direct reference for impulse strength, independent of the exact impulse bandwidth of the currently used IF filter. I used in my experiments a Dynamic Sciences R-1160C impulse generator for that purpose. The subnanosecond impulses it produces do not have an exactly uniform spectral density up to 1 GHz, therefore a calibration chart has to be used to convert the impulse strength setting into the actual spectral density for given frequency ranges.

An important characteristic of an impulse spectrum is that the various frequency components are phase correlated to each other. They are all approximately in phase near the location of the pulse and as a result, the peak voltage of an impulse will grow linearly with the bandwidth of the filter. Additional frequency components simply add their contribution to the peak.

In contrast, the various frequency components of real noise are not phase correlated. Examples of real noise are thermal noise or the superposition signal from a large number of independent sources, but not, for example, just a single unwanted impulse source. As a result, it is not the voltage of noise but only its power or squared voltage that increases linearly with the filter bandwidth. In other words, the voltage of noise will grow with the square root of the filter bandwidth. Consequently, in order to detect broadband impulse energy against a uniform noise background, the filter bandwidth should be selected to be as large as possible.

2.4 Signal correlation

An eavesdropper records some accessible (“black”) signal $b(t)$ in order to gain information about a supposedly inaccessible (“red”) signal $r(t)$. The black signal will usually be an analog physical quantity transformed into a voltage by a signal transducer (antenna, field probe, cable tap, microphone, hydrophone, photosensor, seismometer, laser interferometer, etc.). The red signal can either be analog or digital and both can be vector signals (multiple channels). The shielding of a system can be described as *unconditionally secure* against eavesdropping with a specified sensor type and position if it is impossible to find any statistical dependence between $r(t)$ and $b(t)$. In a system that is only *practically secure* against eavesdropping, such statistical dependencies might well be possible to find, but they are not sufficiently strong to enable the decoding of intelligible information from the red channel.

Statistical independence of two random variables X and Y is in practice difficult to verify with laboratory measurements. Tests for a hypothesis of statistical independence such as the χ^2 test require measurements of not only the relative frequencies $P(X)$ and $P(Y)$ of the two random variables, but also the relative frequencies $P(X, Y)$ of all values of their product space, in order to check whether $P(X) \cdot P(Y)$ is a good estimate for $P(X, Y)$. Because signal dependency is not just instantaneous but can also involve time delays and many other physically plausible transforms (e.g., frequency modulation), the random variables X and Y are in practice rather large vectors. They can be, for example, sliding sample windows over $r(t)$ and $b(t)$, as well as relevant transforms of these. In order to verify a hypothesis of statistical independence, frequency histograms over potentially huge state spaces would have to be recorded, which would quickly reach impractical memory requirements.

A far simpler signal security criterion than statistical independence is the lack of correlation. Two random variable vectors X and Y are uncorrelated, if their covariance matrix $E[(X - \bar{X})(Y - \bar{Y})^T]$ equals zero, where $\bar{X} = E(X)$ denotes the expected value or mean of X . Given a sample $(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)$ of observations of X and Y , the correlation of these random variables can be quantified by the correlation-coefficient




Figure 2.9: These plots demonstrate the use of cross-correlation to find even weak traces of a *red* signal in noisy sensor output.

matrix

$$R = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})^\top}{\sqrt{\sum_{i=1}^n (x_i - \bar{x}) \cdot \sum_{i=1}^n (y_i - \bar{y})^\top}}. \quad (2.13)$$

Even the covariance matrix of sliding sample windows over $r(t)$ or $b(t)$ can still be a rather large data structure, but if we assume that the statistical dependence of $b(t)$ from $r(t)$ is in the form of a noisy causal linear time-invariant channel

$$b(t) = (r * h)(t) + n(t) = \int_0^\infty r(t - t') h(t') dt' + n(t) \quad (2.14)$$

where $n(t)$ is uncorrelated noise, then the correlation test can be further simplified.

Figure 2.9 illustrates this technique. It shows in curve (a) a short part of a 1000-bit long random test signal $r(t)$ as it might occur inside a computer system, for example, on a bus or data transmission line. Curve (b) is a bandpass filtered version $(r * h)(t)$ of that signal, which is what an eavesdropper would receive in the absence of any other noise. The shape of the individual pulses in (b) depends, for example, on the frequency and phase characteristics of any transmitting and receiving antennas involved [see Fig. 2.2 for examples for the influence of the receiving antenna on $h(t)$]. Finally, curve (c) is the full received signal $b(t)$ with noise added, in which the individual impulses have visually disappeared.

In order to recover a noise reduced estimate of $h(t)$, we first create a preconditioned version $\tilde{r}(t)$ of the original signal $r(t)$, which we assume to be accessible in this test. The preconditioned signal $\tilde{r}(t)$ shown in curve (d) marks a positive edge with a positive pulse and a negative edge with a negative one. This signal convolved with the shape of an individual edge response from curve (b) will lead to the entire curve (b). This signal

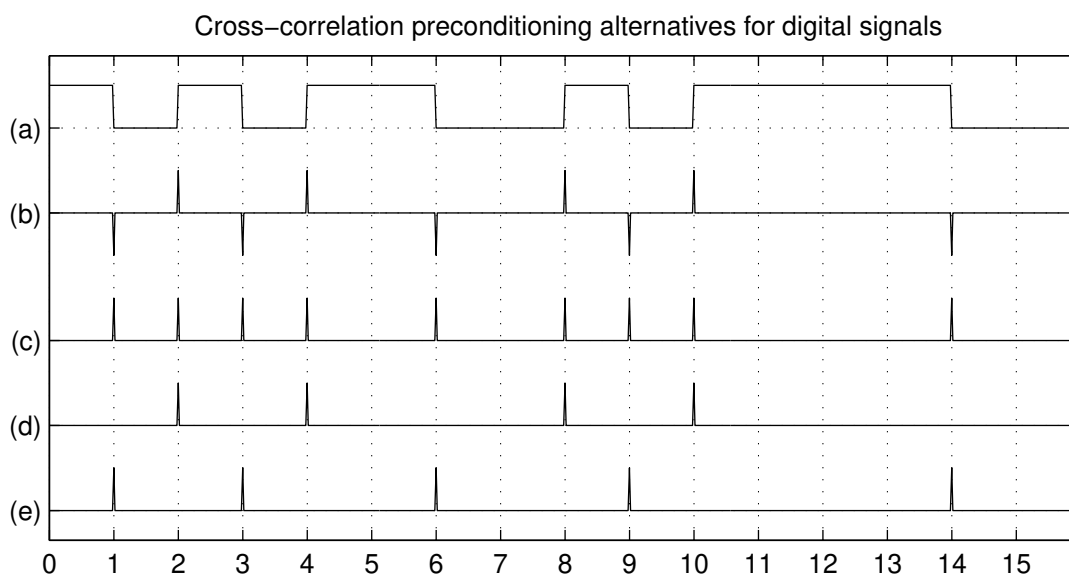


Figure 2.10: Some possible preconditioning techniques that can be applied to a digital *red* signal before performing a cross-correlation with an analog *black* signal.

cross-correlated with the received signal (c) leads to curve (e), which shows as the result of the cross-correlation

$$\phi_{\tilde{r}b}(l) = \int_{-\infty}^{\infty} \tilde{r}(t+l) b(t) dt \quad (2.15)$$

an estimate of a single edge response.²

Figure 2.9 illustrates how cross-correlation can detect a known signal within noise and in this way provides a sensitive measure for the signal strength available in $b(t)$. The preconditioning step to derive $\tilde{r}(t)$ from $r(t)$ made it possible for $\phi_{\tilde{r}b}(l)$ to approximate the step response of the observed time-invariant linear system.

Figure 2.10 shows several other useful preconditioning methods, which can be applied if the transmission channel is one of several commonly occurring types of time-invariant non-linear systems. Curve (a) is the original signal $r(t)$, which is the appropriate cross-correlation partner for $b(t)$ if the transmission channel is a low-pass filter, and curve (b) is the previously used form that marks rising and falling edges with positive and negative pulses. If some form of rectification occurs in the transmission channel, for example because an AM detector is used as part of the receiver circuit, curve (c) becomes the appropriate preprocessing step, as the responses to both falling and rising edge will now be found with equal polarity in $b(t)$. Often, rising and falling edges in $r(t)$ lead to different edge responses. A data line connected to open collector drivers and a pull-up resistor will show rapid fall times but considerably slower rise times, an effect that can, for example, be observed in the data line of a PC keyboard. The preprocessing steps in curves (d) and (e) measure these two edge responses separately in a cross-correlation.

Cross-correlation provides an easy quantification of energy emitted by each individual red

²The $\phi_{\tilde{r}b}(l)$ curve in Fig. 2.9 (e) shows, in addition to the main edge response (at zero lag), also two weaker inverted pulses at lag ± 1 bit duration. These emerge, because the autocorrelation function of $\tilde{r}(t)$ shows negative spikes at the same positions. Even though the bits of $r(t)$ are independent and equally distributed in this example, the edges in $\tilde{r}(t)$ are not. A rising edge, for example, cannot immediately be followed by another one.

signal from inside a system, even if noise from other sources cannot be suppressed. Cross-correlation experiments can be performed, for example, with the help of a two-channel digital storage oscilloscope that is used to record simultaneously a sufficiently long segment of both a red and a black signal. It is preferable for the device under test to be operated in a mode, where the red signal shows random activity. Care has to be taken that the signal probe connected to the red line does not act as an additional antenna itself. The cross-correlation can then be performed after the data has been transferred onto a PC. Faster visualization of the cross-correlation can be achieved with hardware correlators. An example of a product designed for such Tempest measurements is the *KOR 4 Correlator* developed by *Bosch/ANT* and the *German Information Security Agency*, the design of which is described in [24].

Chapter 3

Analog video displays

Wim van Eck's 1985 paper [29] on eavesdropping video terminals with modified TV sets made compromising emanations of cathode-ray tube (CRT) based computer displays widely known in the information security community. Several reports [16, 30, 31, 37, 38] about the video emanations of late 1980s display technology have been published, providing advice for the use of electromagnetic shielding techniques as countermeasures.

However, the data provided in these publications consists only of spectrum-analyzer curves that give no indication of what distortions of readability are to be expected for remotely reconstructed screen content. No examples of eavesdropped video signals that allow the reconstruction of readable text and the use of automatic character recognition have been published so far.

So there are several good reasons to revisit the subject. Display technology has evolved rapidly, pixel frequencies and video bandwidths have increased by an order of magnitude, analog signal transmission is in the process of being replaced by Gbit/s digital video interfaces and various flat-panel display (FPD) types have become serious competitors for the cathode-ray tube. In spite of these developments, the continued practical applicability of van Eck's attack and possible extensions have not been discussed in the open literature.

A certain amount of electromagnetic shielding of CRT monitors has become industry standard [85, 86] and new legal electromagnetic compatibility requirements [88, 89, 91] have been introduced since then, but the application of shielding and other design options specifically tested for information-security purposes is still very much restricted to applications where such protection is legally required, namely for facilities that have government licenses to operate digital-signature key certification authorities [93] or to handle classified "national security" information.

3.1 Video-signal timing

Practically all modern computer video displays are raster-scan devices. As in a television receiver, the image is transmitted and updated as a sequence of scan lines that cover the entire display area with constant velocity. The pixel luminosities in this sequence are a function of the video-signal voltage. (Vector displays are an alternative technique, in which not only the intensity but also the path of the electron beam in a cathode-ray tube is controlled by the displayed data. They were used during the last century in some early

computers, measurement equipment, and avionics displays, but have disappeared almost completely from the market, except perhaps in the form of analog oscilloscopes.)

The time and frequency characteristic of the video signal in a digital display system is first of all characterized by the pixel clock frequency f_p , which is – in the case of a CRT – the reciprocal of the time in which the electron beam travels from the center of one pixel to the center of its right neighbor. The pixel clock is an integer multiple of both the horizontal and vertical deflection frequency, that is the rate $f_h = f_p/x_t$ with which lines are drawn and the rate $f_v = f_h/y_t$ with which complete frames are built on the screen. Here, x_t and y_t are the total width and height of the pixel field that we would get if the electron beam needed no time to jump back to the start of the line or frame. However, the image actually displayed on the screen is only x_d pixels wide and y_d pixels high, because the time allocated to the remaining $x_t y_t - x_d y_d$ invisible pixels is needed to transmit synchronization pulses to the monitor and bring the electron beam back to the other side of the screen.

PC software can read these timing parameters either directly from the video controller chip, or it can determine them via setup tools (e.g., `xvidtune`) or from configuration files. For instance, on a typical Linux installation, a line of the form

```
ModeLine "1280x1024@85" 157.5 1280 1344 1504 1728 1024 1025 1028 1072
```

in the X Window System server configuration file `/etc/X11/XF86Config` indicates that the parameters $f_p \approx 157.5$ MHz, $x_d = 1280$, $y_d = 1024$, $x_t = 1728$ and $y_t = 1072$ are used, which leads to deflection frequencies of $f_h = 91.146$ kHz and $f_v = 85.024$ Hz.

In order to facilitate the correct factory adjustment of the image geometry in displays over the wide range of video timings used in personal computers, the *Video Electronics Standards Association (VESA)* has standardized a collection of exact timing parameters [82], which most PCs use today. Table 3.1 lists the VESA timing parameters most relevant for video signal eavesdroppers.

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

$$t = \frac{x}{f_p} + \frac{y}{f_h} + \frac{n}{f_v}, \quad (3.1)$$

for all $0 \leq x < x_d$, $0 \leq y < y_d$ and $n \in \mathbb{N}$.

As a very simple demonstration of software controlled radio transmission of information via a CRT, we can generate an on-screen image whose video signal resembles that of a broadcast AM modulated audio tone [72]. To do this, the video signal has to approximate the waveform

$$\begin{aligned} s(t) &= A \cdot \cos(2\pi f_c t) \cdot [1 + m \cdot \cos(2\pi f_t t)] \\ &= A \cdot \left\{ \cos(2\pi f_c t) + \frac{m}{2} \cdot \cos[2\pi(f_c - f_t)t] + \frac{m}{2} \cdot \cos[2\pi(f_c + f_t)t] \right\} \end{aligned} \quad (3.2)$$

if the chosen carrier and audio tone frequencies are f_c and f_t , respectively. Using the above two formulae with the frame counter $n = 0$, we can calculate a time t for every pixel (x, y) and set this pixel to an 8-bit grayscale value of $\lfloor \frac{255}{2} + s(t) + R \rfloor$ with amplitudes $A = \frac{255}{4}$ and $m = 1$, where $0 \leq R < 1$ is a uniformly distributed random number to spread

x_d	y_d	f_v/Hz	f_h/kHz	f_p/MHz	x_t	y_t
640	350	85.080	37.861	31.500	832	445
640	400	85.080	37.861	31.500	832	445
720	400	85.039	37.927	35.500	936	446
640	480	59.940	31.469	25.175	800	525
640	480	72.809	37.861	31.500	832	520
640	480	75.000	37.500	31.500	840	500
640	480	85.008	43.269	36.000	832	509
800	600	56.250	35.156	36.000	1024	625
800	600	60.317	37.879	40.000	1056	628
800	600	72.188	48.077	50.000	1040	666
800	600	75.000	46.875	49.500	1056	625
800	600	85.061	53.674	56.250	1048	631
1024	768	43.479	35.522	44.900	1264	817
1024	768	60.004	48.363	65.000	1344	806
1024	768	70.069	56.476	75.000	1328	806
1024	768	75.029	60.023	78.750	1312	800
1024	768	84.997	68.677	94.500	1376	808
1152	864	75.000	67.500	108.000	1600	900
1280	960	60.000	60.000	108.000	1800	1000
1280	960	85.002	85.938	148.500	1728	1011
1280	1024	60.020	63.981	108.000	1688	1066
1280	1024	75.025	79.976	135.000	1688	1066
1280	1024	85.024	91.146	157.500	1728	1072
1600	1200	60.000	75.000	162.000	2160	1250
1600	1200	65.000	81.250	175.500	2160	1250
1600	1200	70.000	87.500	189.000	2160	1250
1600	1200	75.000	93.750	202.500	2160	1250
1600	1200	85.000	106.250	229.500	2160	1250
1792	1344	60.000	83.640	204.750	2448	1394
1792	1344	74.997	106.270	261.000	2456	1417
1856	1392	59.995	86.333	218.250	2528	1439
1856	1392	75.000	112.500	288.000	2560	1500
1920	1440	60.000	90.000	234.000	2600	1500
1920	1440	75.000	112.500	297.000	2640	1500

Table 3.1: The 43 VESA standard modes listed here cover the vast majority of video timing parameters used in current personal computer displays [82]. The frequency tolerance specified for f_p is $\pm 0.5\%$.




Figure 3.1: Example screen contents that cause a monitor in the $1280 \times 1024 @ 85\text{Hz}$ video mode to broadcast an $f_t = 300$ Hz (left) and 1200 Hz (right) tone on an $f_c = 1.0$ MHz carrier in amplitude modulation.

the quantization noise (dithering). See Fig. 3.1 for example screen contents generated this way to broadcast an AM tone.

If the carrier frequency is suitably chosen in the shortwave bands (roughly 2–20 MHz), this technique will allow a cheap world-band radio several meters from the monitor to be used to pick up a melody played. With $f_c = \frac{1}{2}f_p = 78.75$ MHz and $f_t = 1$ kHz, I measured 1 m above a “Dell D1025HE” 43-cm color monitor a field strength of about $32 \text{ dB}\mu\text{V}/\text{m}$ ($= 40 \mu\text{V}_{\text{rms}}/\text{m}$) with the help of an R-1250 measurement receiver set to 2 kHz bandwidth, an RF signal generator used as a reference source, and a calibrated biconical antenna. In this experiment, the highest field strength occurred when the biconical antenna was oriented perpendicular to the monitor’s display surface. For comparison, the mean noise level in a residential area in this band is with typically about $-8 \text{ dB}\mu\text{V}/\text{m}$ ($= 0.4 \mu\text{V}_{\text{rms}}/\text{m}$) 40 dB lower (see section 5.2.1), leading to a clearly receivable signal. Assuming the signal power drops proportional to the square of the distance to the receiver, which holds if we assume that the reference measurement has not been influenced significantly by near-field effects, the signal-to-noise ratio drops to 0 dB within a distance of about 10^2 m. This range estimate can be confirmed in practice, as an AM radio can pick up, under good conditions, an audible melody a few tens of meters from the monitor, for instance in neighboring rooms. The same melody can also be received at frequencies $i \cdot f_p \pm f_c$ for small natural numbers i , which is explained in the following section.

3.2 Analog video-signal spectra

According to Nyquist’s sampling theorem, the full information content of a digitally generated video signal with pixel-clock frequency f_p and pixel values v_i ($i \in \mathbb{Z}$) can be represented as a band-limited continuous waveform $v(t)$ whose Fourier transform $V(f)$ is zero for all frequencies $|f| > \frac{1}{2}f_p$. Given a sequence of sample values $v_i = v(i/f_p)$, this band-limited equivalent waveform can be obtained by interpolation as

$$v(t) = \sum_i v_i \cdot \frac{\sin \pi(f_p t - i)}{\pi(f_p t - i)}. \quad (3.3)$$

Conversely, the discrete time sampling of a continuous band-limited video signal $v(t)$ can mathematically be represented by multiplication with a series of equidistant Dirac impulses

$$\hat{v}(t) = v(t) \cdot \sum_{i=-\infty}^{\infty} \delta\left(t - \frac{i}{f_p}\right). \quad (3.4)$$

The Fourier transform of an equidistant impulse series results in another one, with reciprocal pulse distance:

$$\mathcal{F}\left\{\sum_{i=-\infty}^{\infty} \delta(t - ik)\right\}(f) = \int_{-\infty}^{\infty} \sum_{i=-\infty}^{\infty} \delta(t - ik) e^{2\pi i f t} dt = \frac{1}{k} \sum_{i=-\infty}^{\infty} \delta\left(f - \frac{i}{k}\right). \quad (3.5)$$

In order to understand the spectral content of a video signal, we will also refer to the convolution theorem, according to which multiplication of two functions g and h in the time domain corresponds to convolution of their respective Fourier transforms in the frequency domain

$$\mathcal{F}\{g \cdot h\} = \mathcal{F}\{g\} * \mathcal{F}\{h\} \quad (3.6)$$

and vice versa

$$\mathcal{F}\{g * h\} = \mathcal{F}\{g\} \cdot \mathcal{F}\{h\}, \quad (3.7)$$

where “*” denotes convolution, defined as

$$(g * h)(x) = \int_{-\infty}^{\infty} g(y) h(x - y) dy. \quad (3.8)$$

We obtained in (3.4) the sampled video signal $\hat{v}(t)$ by multiplying the band-limited analog version $v(t)$ with a series of impulses that are spaced $1/f_p$ apart, and this corresponds in the frequency domain to a convolution of the band-limited signal $V(f) = \mathcal{F}\{v\}(f)$ with a sequence of impulses that are f_p apart. The resulting Fourier transform of the sampled video signal therefore has copies of the original band-limited spectrum in the frequency range $-f_p/2$ to $f_p/2$ repeated throughout the spectrum with a repetition distance f_p :

$$\hat{V}(f) = \mathcal{F}\{\hat{v}\}(f) = \sum_{i=-\infty}^{\infty} V(f - i f_p). \quad (3.9)$$

The actual video signal produced by a graphics card does not consist of an infinitely sharp pulse for every pixel as in $\hat{v}(t)$. Instead, every pixel is represented, approximately, by a rectangular pulse of width $t_p = 1/f_p$. Let

$$\Pi(t) = \begin{cases} 1 & \text{if } -\frac{1}{2} < t \leq \frac{1}{2} \\ 0 & \text{otherwise} \end{cases} \quad (3.10)$$

denote a unit pulse with width and amplitude 1. The Fourier transform of a pulse $A\Pi(t/t_p)$ with amplitude A and width t_p is

$$\mathcal{F}\left\{A \cdot \Pi\left(\frac{t}{t_p}\right)\right\}(f) = A t_p \cdot \frac{\sin \pi t_p f}{\pi t_p f}. \quad (3.11)$$

We can describe the actual video signal $\tilde{v}(t)$ as a convolution of the impulse signal $\hat{v}(t)$ and the pulse shape of a single pixel, idealized as the rectangular pulse $p(t) = \Pi(t \cdot f_p)$. With the convolution theorem, we get the Fourier transform of the result:

$$\tilde{V}(f) = \mathcal{F}\{p * \hat{v}\}(f) = P(f) \cdot \hat{V}(f) = t_p \cdot \frac{\sin \pi t_p f}{\pi t_p f} \cdot \sum_{i=-\infty}^{\infty} V(f - i f_p). \quad (3.12)$$

A rectangular pulse with exact width t_p as well as any other pulse shape $p(t)$ where $\sum_{i=-\infty}^{\infty} p(t - i t_p)$ is constant over t will have a Fourier transform with $P(i f_p) = 0$ for all integers $i \neq 0$. This is the pulse shape one might expect from modern digital-to-analog converters. However, van Eck [29, Fig. 8c] noted that in some of the early 1980s video terminals that he tested, the pulse width is *shorter* than the pixel time, in other words, these video signals were a return-to-zero (RZ) encoding of the pixel values because the video voltage alternates between on and off, even during a continuous sequence of bright pixels. As a result, what van Eck described are systems in which the video-signal generator effectively amplitude modulates the band-limited representation of the pixel sequence $v(t)$ by multiplying it with a square wave of frequency f_p . This replicates the entire spectrum of $v(t)$ around all of the harmonic frequencies contained in this blanking signal, making it possible to reconstruct the original video signal throughout the radio spectrum with an amplitude modulation receiver with a bandwidth of at least f_p (resulting in an upper frequency limit in the output signal of $f_p/2$). TV receivers with bandwidths of about 6–8 MHz are therefore only suitable for accurately reproducing the information from emanated video signals with pixel frequencies in the same range. Nevertheless, the demonstration in the next section will show that, in practice, the redundancy of glyph shapes will allow an eavesdropper to visually read text at pixel frequencies up to five times the receiver bandwidth, given a good signal-to-noise ratio, and for larger font sizes perhaps more than that.

Unless this blanking of the video signal between pixels is performed in a graphics card, which was not the case in any of the products that I examined, the low-frequency end of $v(t)$ will not be repeated anywhere else in the generated spectrum $\tilde{V}(f)$ other than near 0 Hz. This means that the low-frequency components of $v(t)$ will be the most difficult to receive for a radio eavesdropper, because where an emitting conductor is significantly shorter than the wavelength, the emitted field strength grows with increasing frequency. For example, equations (A.23) and (A.24) show that the field strength grows by 40 dB/decade for differential-mode cable currents and 20 dB/decade for common-mode currents. The dimensions of the video cable and monitor circuits are small compared to the wavelength for frequencies below about 30 MHz and therefore prevent a significant amount of energy from leaving their immediate near-field environment, which is the volume around the circuit within about a sixth of the wavelength ($\lambda/2\pi$). In the far-field, compromising video emanations can therefore in practice only be received at frequencies significantly above 30 MHz, which means that eavesdroppers can be expected to pick up one or more of the replicated baseband video spectra at the harmonics of the pixel frequency, and not the baseband signal itself. Low-frequency signals are primarily a concern where the eavesdropper has access to a conductor (power line, water pipe, network cable, building material, etc.) that passes through the near field of the targeted circuit and picks up the signal there.

The analog video signal generated by a real-word video card does not represent pixels as exact rectangular pulses. Line drivers (and in some products also passive filters added




Figure 3.2: The pixel pulse shapes and signal spectra of analog VGA signals produced by graphics adapters can differ significantly between products. Top: Voltage curve of a single full-intensity pixel for two graphic cards and a laptop. Bottom: Frequency spectrum measured for a random-bit test image with a spectrum analyzer (pixel frequency f_p in parenthesis).

for interference-suppression purposes) limit the rise and fall time and thereby the high-frequency content of the pulse. In addition, some digital-to-analog converters switch to a new output voltage in several smaller steps, which can cause further deviations from the idealized model of an impulse sequence convolved with a rectangular pulse.

For example, the data sheet of the Texas Instruments TVP3026-220 RAMDAC (a digital-to-analog converter chip with built-in color lookup table RAM) used on the Matrox Millennium II graphics card specifies a typical video-output rise and fall time of 2 ns, which is reduced further to 4 ns in this example product by a low-pass filter located between the RAMDAC output and the VGA connector. The transition time of an edge is commonly specified as the time the signal voltage spends in the range 10 to 90 % of the full voltage swing. The output of a low-pass filter with a 3 dB cut-off frequency f has a time constant $\tau = \frac{1}{2\pi f}$ and a rise/fall time of

$$\Delta t = -\ln \frac{0.1}{0.9} \cdot \tau = -\ln \frac{0.1}{0.9} \cdot \frac{1}{2\pi f} \approx \frac{1}{\pi f}. \quad (3.13)$$

The output of this video card therefore seems to be limited to frequencies below 80–100 MHz.

Figure 3.2 shows how significantly the pulse shapes and frequency spectra can vary between different products, in this case two graphic cards and one laptop. The pulse shapes in the top diagram were recorded with an active oscilloscope probe (< 1 pF input capacitance, 500 MHz bandwidth) across a 75Ω termination resistor that connected the green and ground pins of the 15-pin VGA connector. The bottom diagram shows spectra of a test image, in which pixels are randomly set with equal probability to black or white, recorded with a spectrum analyzer that was connected via an impedance-matched attenuator to the VGA connector. (Due to limitations of the laptop hardware, its video mode $800 \times 600 @ 75 \text{ Hz}$ has a $3 \times$ lower pixel frequency than that used for the cards, which ran in the widely used $1280 \times 1024 @ 85 \text{ Hz}$ mode.)

The repetitive nature of the spectrum of a video signal is clearly visible in the spectrum analyzer curves for both the ATI card and the Toshiba laptop. The second (157–236 MHz) and third (236–315 MHz) harmonic of the ATI card's signal show about a 20 dB drop and the sixth harmonic (472–551 MHz) vanishes almost in the noise. The output of the Matrox card, on the other hand, seems to have gone through a much steeper low-pass filter, as the second harmonic is already attenuated by more than 40 dB. Finally, the laptop is an example of a device whose video signal shows little sign of successful low-pass filtering. Signal levels increase after an initial drop again above 300 MHz and even the 30th harmonic (742–767 MHz) is still clearly recognizable against the noise.

Another feature visible in the spectrum-analyzer curves of Fig. 3.2 are significant spikes at the integer multiples of the pixel-clock frequency. This component is most likely to be caused by cross-talk between the DAC output and the pixel clock signal, for instance via the converter's power supply lines, because these spikes are also present in the spectrum of a completely black test image. These AM-carrier like tones play no role in the demonstration experiments described below, but they are nevertheless worth mentioning, because a very sophisticated eavesdropper might isolate one of them with a notch filter and use it with a phase-locked loop to reconstruct the original pixel clock and obtain exact synchronization information.

3.3 Eavesdropping demonstration

3.3.1 Realtime monitoring

Compromising video signals can be visualized conveniently in real-time on a normal VGA multisync monitor, provided that their signal-to-noise ratio is good enough. We can simply connect its green 0.7 V/75 Ω input via an impedance-matched attenuator (see Section B.1) to the 3 V/50 Ω AM demodulator output of the R-1250 wide-band receiver shown in Fig 2.4. All the received signals shown in the following figures are not screenshots from such a monitor, but were captured with a digital storage oscilloscope from the receiver output and converted into raster graphics files on a PC. Nevertheless, the real-time observation of a received signal on a monitor is invaluable for manually adjusting the various reception parameters including the antenna position, center frequency, and bandwidth for best results.

The eavesdropper's monitor needs to be supplied with very close approximations \tilde{f}_h and \tilde{f}_v of the horizontal and vertical synchronization frequencies f_h and f_v that drive the targeted display. If, for example, these frequencies differ by even only one part per million (1 ppm, i.e. $\tilde{f}_h = f_h \times 1.000001$), and the number of pixel times per frame refresh $x_t y_t$ is about 10^6 , then after a single refresh, the positions of the electron beams on the two monitors will already have moved relative to each other by one pixel. The effect for the eavesdropper is that the image seen on the monitor will roll horizontally with a speed of

$$v_h = \frac{\tilde{f}_h - f_h}{f_h} \cdot x_t y_t \cdot f_v \cdot r \quad (3.14)$$

where r is the width of a pixel. With a refresh rate of $f_v = 85$ Hz and \tilde{f}_h being just 1 ppm too high, as in the above example, the image on the eavesdropping monitor would roll to the right with a speed of 85 pixels per second.

In order to achieve an image with acceptable stability, the oscillator used for generating the horizontal sync pulses for the eavesdropping monitor needs to be adjustable in frequency with a resolution of at least seven to eight digits. This way, the horizontal image drift can be adjusted with a resolution better than one pixel per second. The PLL pixel-clock generators found in typical graphics cards provide only an orders of magnitude cruder frequency resolution, which makes them unsuitable as sync signal generators for this application. (Recall that the VESA standard permits a frequency tolerance of 0.5 ‰ = 5000 ppm.)

I attempted several ways to program an available arbitrary-waveform generator to produce both the horizontal and vertical sync signals with the required accuracy. I eventually succeeded by using a second function generator that can produce a 10 MHz square wave signal with a frequency resolution of 1 Hz, which I fed into the 10 MHz reference-clock input of the arbitrary-waveform generator, in order to fine-tune its output frequency (see Section B.2 for details).

Even with a sufficiently high resolution for adjusting \tilde{f}_h , the image on the eavesdropping monitor will not remain in place for longer than a few seconds or minutes. The crystal oscillators in both the video adapter of the targeted system as well as the eavesdropper's time base are temperature dependent and their relative frequencies can drift in practice up to several ppm within just a few minutes, making it necessary for the eavesdropper

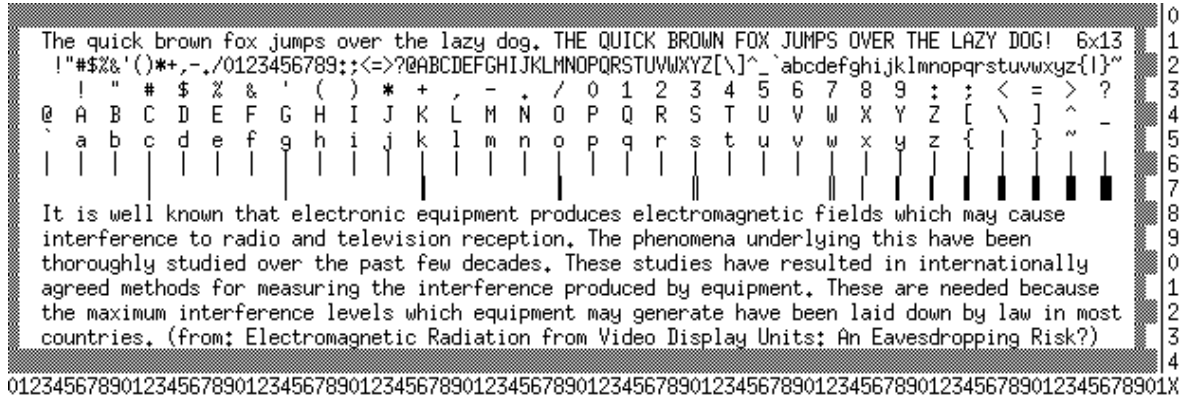


Figure 3.3: Test text on the target monitor when Figs. 3.4–3.6 were recorded.

to continuously adjust the \tilde{f}_h setting. Using a somewhat more stable time base (e.g., a temperature-controlled crystal oscillator, radio time-signal receiver, or atomic clock) can help the eavesdropper to ensure that this image drift is primarily caused by frequency fluctuations in the targeted system, but will not reduce the problem significantly.

The sensitivity of the reconstructed image position to tiny frequency errors in the sync signals might seem like a nuisance to the eavesdropper, but it is actually an advantage that helps single-out signals from an individual target device. Each PC video card uses its own independent crystal oscillator for generating the pixel clock signal. Even if there are several video displays within reception distance that run in exactly the same video mode, their respective clock frequencies are very likely to differ by at least several parts per million. The eavesdropper can therefore adjust \tilde{f}_h such that the signal of only one target system at a time appears stable on the monitor, while the signals from other systems will drift by at significant speed.

The human eye can quite easily separate on the real-time eavesdropping monitor the targeted system's stable signal from the moving signals of other monitors. The periodic averaging of multiple frames that needs to be applied to reduce noise when the signal is recorded digitally equally ensures that unwanted signals from other nearby monitors are blurred significantly when there is even a tiny frequency mismatch.

Where the exact video mode – and in particular parameter y_t – are known, the refresh frequency $\tilde{f}_v = \tilde{f}_h/y_t$ can be derived directly from the line frequency, and only one single frequency needs to be controlled. If the video mode of the targeted system is not one of the common standard modes, separately adjustable \tilde{f}_v and \tilde{f}_h controls significantly simplify a manual search for the right frequencies. A helpful observation for such manual adjustments is that for $\tilde{f}_h \ll f_h$ the image will be sheared such that vertical lines lean by an angle

$$\theta = \arctan \left(\frac{f_h - \tilde{f}_h}{f_h} \cdot x_t \right) \quad (3.15)$$

to the right, because the electron beam in the eavesdropping monitor falls behind (to the left) relative to the target monitor beam, as both progress down the screen. Equivalently, for $\tilde{f}_h \gg f_h$ vertical lines will lean to the left. The image will appear stable but distorted whenever $\tilde{f}_h = f_h \cdot m/n$ for small integers $m \neq n$.

3.3.2 Experimental setup

The target monitor used for the following tests was an “iiyama Vision Master Pro450”, which was designed for pixel-clock frequencies up to 300 MHz and all the video modes listed in Fig. 3.1. It displayed the test text shown in Fig. 3.3, provided from the VGA output connector of a Toshiba 440CDX laptop in the same 800×600@75Hz VESA video mode as was used for this signal source in Fig 3.2 ($f_p = 49.5$ MHz). The font used in this test text is the 6×13-pixel large **fixed** font from the X11 Window System distribution, a relatively small and widely used default font for applications such as `xterm`.

The eavesdropping transducer in this test was a vertically polarized log-periodical antenna (200–1000 MHz frequency range) placed about 3 m from the side of the monitor. At this distance, the field strength is already being measured in the far field. It can therefore be converted easily into the equivalent value at larger antenna distances ($\frac{1}{10}$ voltage at 30 m, etc.).

I recorded these signals from the AM demodulator output of the R-1250 receiver using an 8-bit digital storage oscilloscope (Tektronix TDS 7054) operating at a sampling rate of 500 MHz. The recording time per frame was 4 ms, leading to two million samples per image. The oscilloscope used is able to average the data acquired from a specifiable number of trigger events with 16 bits resolution, but unfortunately, to perform the averaging operation, it has to transfer the data over a slow interface after each acquisition and can therefore trigger only every 2.3 seconds. As a result, averaging 256 frames, with 2×10^6 samples each, takes about 10 minutes. A more suitable real-time averaging instrument could acquire the same data within $256/f_v = 3$ s. This instrument limitation made the recording significantly more sensitive to any frequency error in the reconstructed sync signal, which can easily be adjusted to drift for less than a pixel for a few seconds, but not for several minutes. It was therefore necessary to trigger on the sync pulse from the laptop VGA output, and not (as an eavesdropper would have to) on the sync pulse from the potentially slightly drifting function generator. A better-equipped eavesdropper would use a digital acquisition system that can average recorded frames in real time without missing a frame, can therefore accumulate the same number of frames in a much shorter time interval and the result will depend far less on the \tilde{f}_v . All the averaged signals were recorded by triggering on a vertical sync signal from the target device, except for the one in Fig. 4.3, which demonstrates how an actual eavesdropping attack could be performed in spite of this equipment limitation.

The software routine that turns the acquired and averaged signal into a raster graphic also needs to use a quite exact estimate of f_h , preferably with a relative error of less than 10^{-5} , as otherwise vertical lines in the source image will deviate from the reconstructed vertical by more than $\approx 1^\circ$. Even with the quite generous oversampling used in this experiment (500 MHz oscilloscope sampling frequency for a 49.5 MHz pixel clock), it is still necessary to align consecutive lines with subpixel resolution. This can be done by calculating for each pixel in the generated raster graphic the time according to equation (3.1) and interpolating the pixel value from the acquired data accordingly.

3.3.3 Results

Figures 3.4–3.6 illustrate the quality and readability that an eavesdropper can achieve from this signal source at receiver bandwidths ranging from 10 to 200 MHz. The images

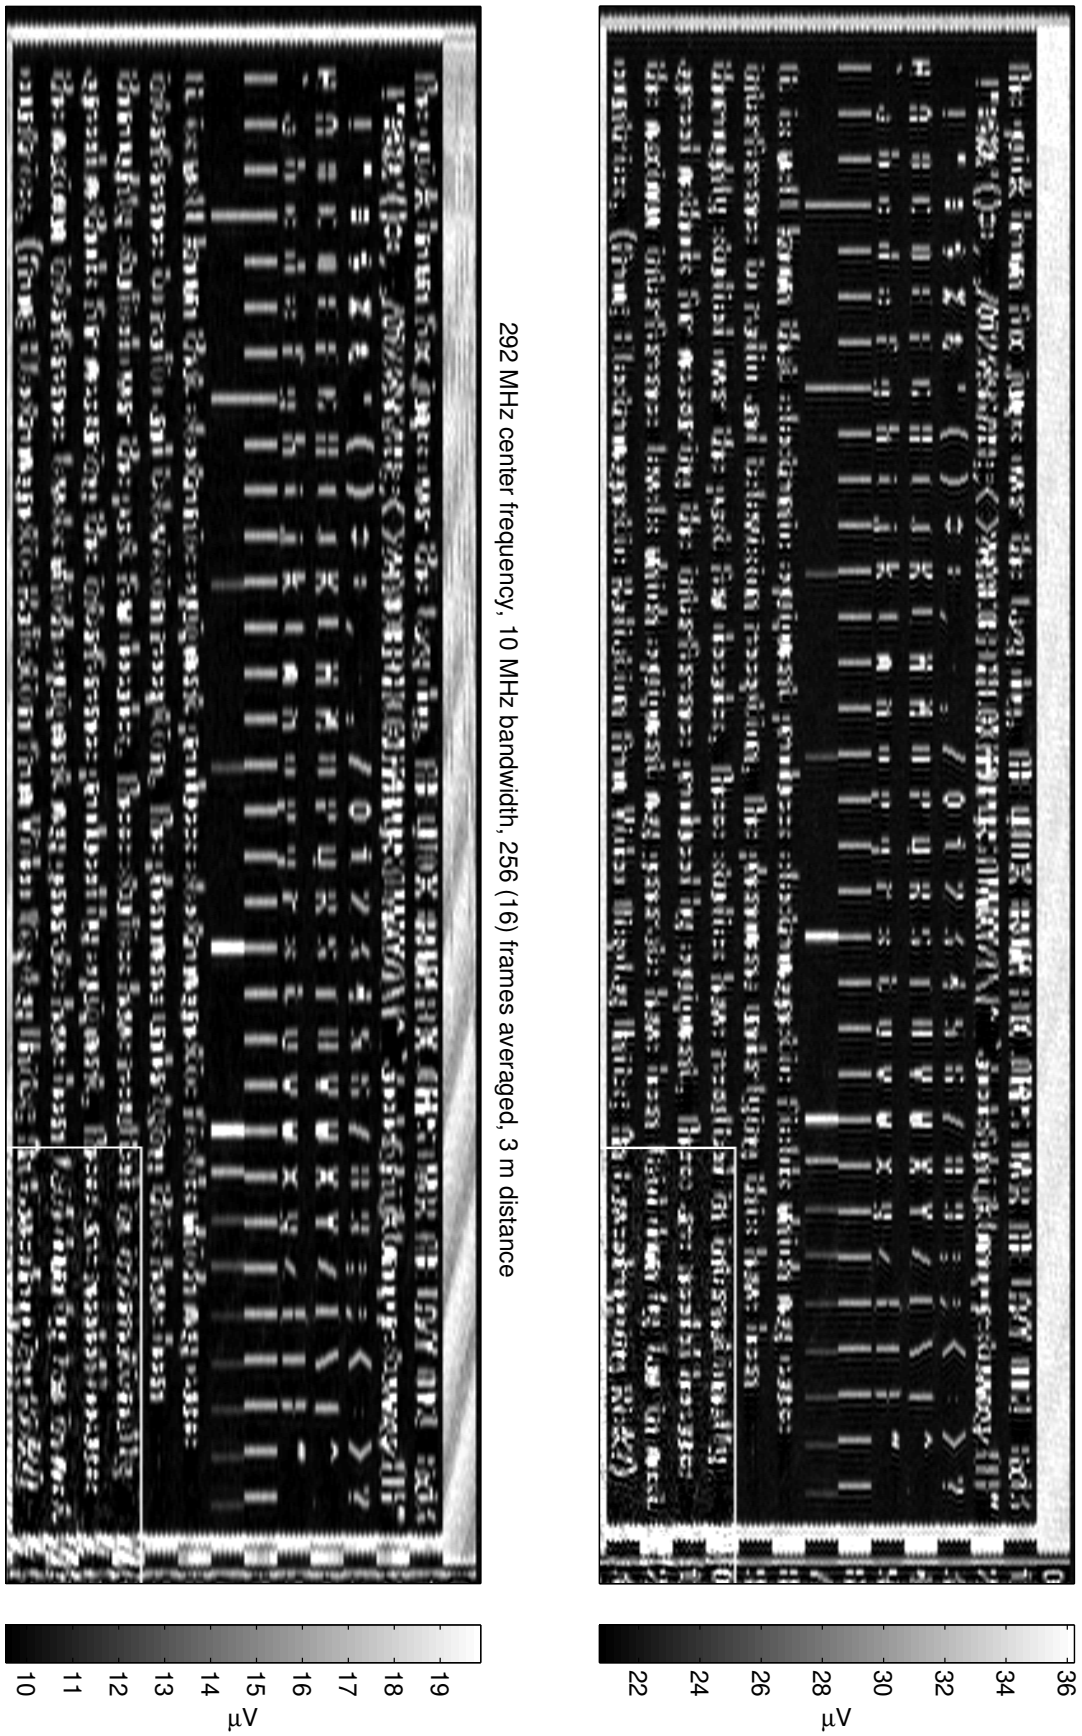


Figure 3.4: Examples of linear AM demodulated signals from a CRT (low bandwidth), with 256 frames averaged (only 16 in the bottom right corner for comparison).

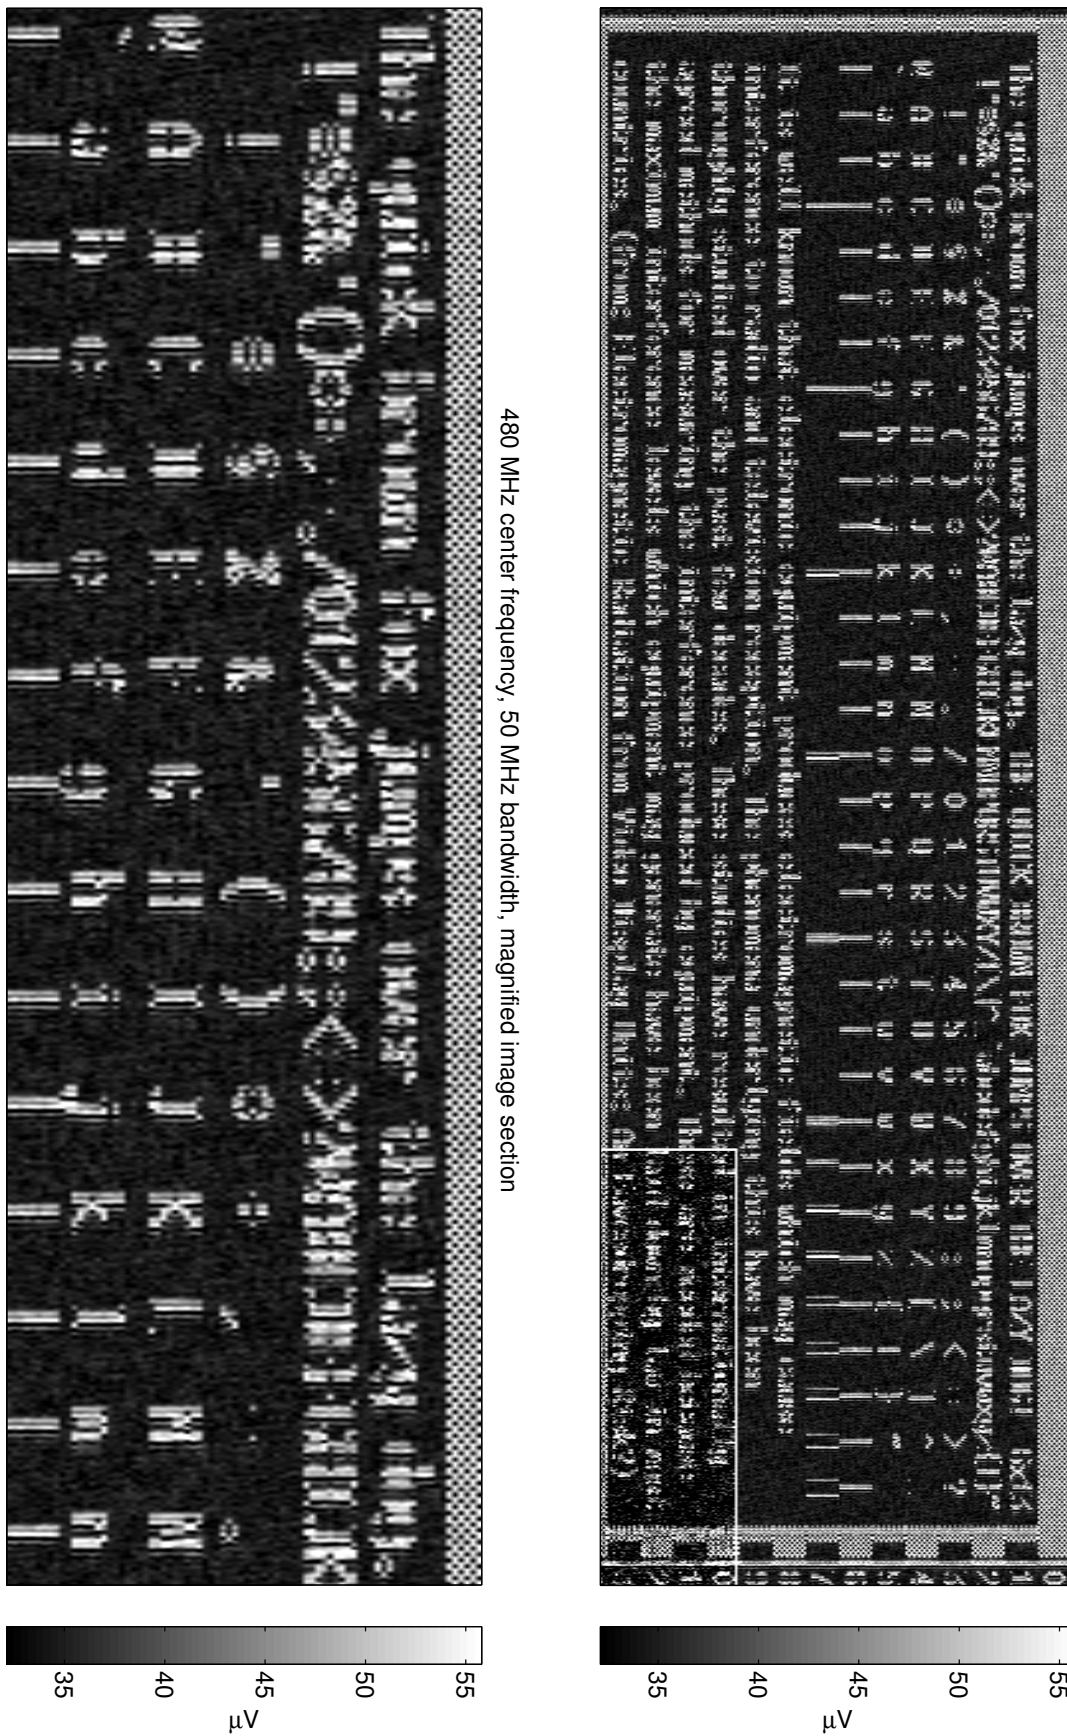


Figure 3.5: Examples of linear AM demodulated signals from a CRT

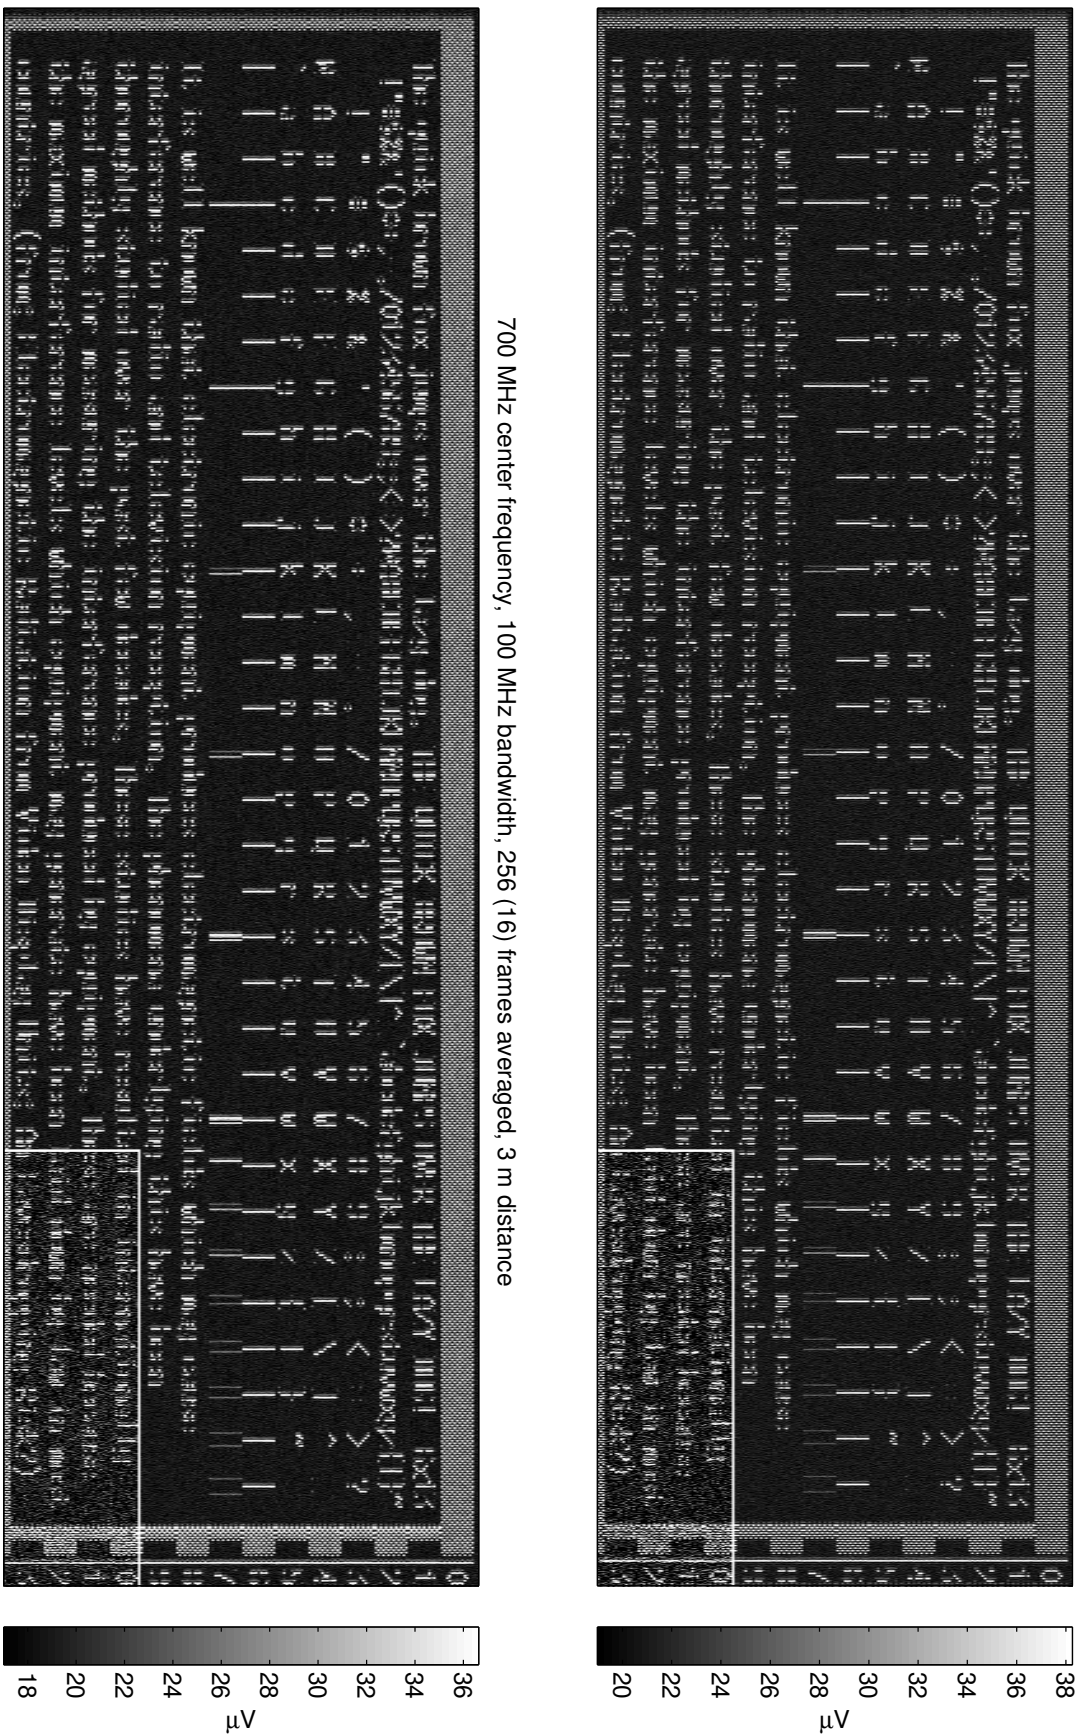


Figure 3.6: Examples of linear AM demodulated signals from a CRT (high bandwidth)

freq. [MHz]	BW [MHz]	IBW [MHz]	voltage [μV]	level [dB μV]	AF [dB]	e-field [dB $\mu\text{V}/\text{m}$]	pulse strength [(dB $\mu\text{V}/\text{MHz})/\text{m}$]	EIRP [nW]
292	10	8	9	19	16	35	17	1
292	20	14	18	25	16	41	18	4
480	50	45	23	27	19	46	13	12
700	100	89	18	25	22	47	8	15
740	200	141	18	25	22	47	4	15

Table 3.2: Comparison of signal strength measures for Figs. 3.4–3.6.

show the rastered signals for both 256 and 16 (lower right corner) averaged frames. The center frequency used for each bandwidth, respectively, was the one that, judging from the signal on the real-time monitor, promised the best signal-to-noise ratio in the 30 MHz to 1 GHz range. The measurements took place in an unshielded ground-floor office¹.

For each measurement, I used an RF sine-wave generator tuned to the center frequency of the receiver, to measure the receiver gain (i.e., which RF input voltage results in which AM demodulator output voltage). This way, I was able to add a calibration bar to each rastered image that shows the root-mean-square antenna-input voltage that corresponds to a shade of gray. The image brightness and contrast are scaled such that the 0.025 % lowest and highest sample values get mapped to black and white respectively.

The images show that the textual information appears with amplitudes of about 9–23 μV (19–27 dB μV) at the antenna input. Taking the log-periodic antenna factors from Fig. 2.3 into account, this translates into electric field strengths in the order of 35–47 dB $\mu\text{V}/\text{m}$ in the measured bands, and taking the receiver’s impulse bandwidth into account, the equivalent spectral densities of the impulses observed go up to 18 (dB $\mu\text{V}/\text{MHz})/\text{m}$. Using (A.32), the measured field strength can be converted into a flux density and by multiplying with the area of a sphere whose radius is the antenna distance, this can be converted into the equivalent signal power of an isotropic radiator (EIRP).

The results for each of the measured bands are listed in Table 3.2. They show that extending the bandwidth to 200 MHz in the last measurement did not capture any additional impulse energy and that the spectral density of the signal falls with increasing frequency from 300 to 700 MHz. The peak signal strength of this monitor in the lower UHF spectrum can be compared with a transmitter that emits with a power of about a dozen nanowatts. This makes the claim in [33] plausible that a 5 mW transmitter that emits a suitably shaped jamming signal can be used as an effective eavesdropping protection for video displays.

Figure 3.7 shows the spectrum analyzer curve (100 kHz resolution bandwidth, averaged output of peak detector) for the video signal fed into the target monitor. It also shows a spectrum of the signal from the eavesdropping antenna, which contains ambient noise from numerous other radio sources. The signal emitted by the target monitor is, at this narrow bandwidth, below the noise floor of the spectrum analyzer and therefore not visible. The crosses at the bottom represent most of the receiver center frequencies that led, during a manual scan of the spectrum, to readable text at various bandwidths (top row: 10 MHz, bottom row: 200 MHz). They show that narrower bandwidths can be used

¹William Gates Building in Cambridge, a modern office block with plasterboard walls separating rooms, filled with well over a hundred active computers, located in a semi-rural environment




Figure 3.7: This figure shows the spectral composition of the video signal from Figs. 3.3–3.6, compared with a random dot signal shown in the same video mode ($800 \times 600@75\text{Hz}$, $f_p \approx 49.5 \text{ MHz}$), as well as the signal at the log-periodic antenna output, which is dominated by other signal sources such as radio transmitters.

to obtain a signal even in parts of the spectrum between two noise sources that would not be suppressed sufficiently with a wider input bandwidth.

Particularly good center frequencies in this example setup turned out to be quieter areas of the local radio spectrum near 300, 430, and 700 MHz. The selection of an eavesdropping center frequency is a tradeoff between several factors:

- The radiation efficiency of connections and circuitry generally increases with frequency according to (A.23) and (A.24), as long as the conductors involved are small compared to the wavelength.
- The video signal in the monitor is passed through an amplifier to increase its amplitude to about 60–80 V before applying it between the cathode and the control grid to regulate the electron-beam current in a CRT. The limited bandwidth of the amplifier reduces the high-frequency content that can emanate from the control grid.
- As a mostly linear amplifier, the monitor will only emit spectral components that are present in the signal generated by the graphics card. These drop off in general with increasing frequency, depending on the design of the output filters.
- The reception band should naturally be free of strong continuously transmitting radio stations and other unwanted noise sources.




Figure 3.8: AM demodulator output from part of a single line in Figs. 3.4 and 3.5.

Figure 3.4 shows that with a bandwidth of a fifth of the pixel-clock frequency, it already is possible to distinguish individual characters at this small font size, although their shapes are still severely distorted and reading the text requires significant guesswork. Vertical edges are represented as single pulses and pairs of vertical edges cannot be resolved, which is still the case at 20 MHz bandwidth ($BW \approx f_p/2$), see also Fig. 3.8.

Starting with 50 MHz bandwidth ($BW \approx f_p$), as seen in Fig. 3.5, individual pixels can be resolved. The magnification shows, that the received signal contains a pulse for each image location where the electron beam is switched on or off. This is the result of the lack of low-frequency components in the replicated spectra of the video signal ($P(if_p) = 0$), as predicted by equation (3.12). The dither pattern in the border around the test text is visible at this bandwidth, as are neighboring vertical bars. Characters are now clearly recognizable, even though the reader has to get used to the changed glyph shapes. Vertical lines are doubled and horizontal lines are only represented by pulses marking their end points. Certain character pairs are more difficult to distinguish than normally, for example ‘E’ and ‘F’, which differ only in horizontal strokes.

Increasing the bandwidth to multiples of f_p does not add any further information to the reconstructed image, as the additionally incorporated spectrum only contains repeated copies of the same band-limited video signal $v(t)$. In a quiet RF environment, a larger bandwidth can improve the signal-to-noise ratio and make the beam-on/off pulses narrower. At higher bandwidths, the on and off pulses in this experiment could be distinguished by their different amplitudes.




Figure 3.9: These single-glyph signals extracted from Figs. 3.4 and 3.5 demonstrate the benefits of a larger receiver bandwidth for the reduction of inter-character interference. The two images on the left were received with 200 MHz bandwidth, the others with 50 MHz. Within these groups, the left “W” is from position (61,1) and the right one from position (58,2) in Fig. 3.3, respectively.

3.4 Radio character recognition

Manually evaluating and transcribing captured video signals from compromising emanations is feasible, as the examples in the previous section show, but slow and labor intensive. A well-equipped eavesdropper can, therefore, be expected to use pattern recognition software, in order to automatically transcribe a signal into plaintext. The example presented in this section demonstrates that this is feasible, with a very simple algorithm, at least for monospaced fonts.

If all characters have the same width, as is still the case with terminal emulators and video terminals that imitate the behavior of a teletype machine, then the received and rastered image can simply be cut into character cells of equal size, as shown in Fig. 3.9. The rigid timing of the video signal causes all characters to end up aligned identically in these cells. If the eavesdropper can correctly guess the exact font used, then the only two parameters that need exact adjustment for radio character recognition are the line rate \tilde{f}_h and the time when the first pixel of a text field is transmitted. Most other unknown variables that make optical character recognition a complex problem (rotation, scaling, glyph shapes, pixel alignment, glyph separation, etc.) are not an issue with video characters.

The choice of the best bandwidth for radio character recognition is a tradeoff. In lower bandwidths, the impulse response of the AM demodulator is longer than a single pixel, and some content of each character cell is influenced by the left neighbor character; see Fig. 3.9 for two examples of the character ‘W’ with different left neighbors at 200 and 50 MHz bandwidth. Inter-character interference is reduced at higher bandwidths. On the other hand, the sharper glyph shapes that come with increased bandwidth also make the comparison algorithm more sensitive to misalignment caused by errors in \tilde{f}_h .

In my experiments, choosing the bandwidth close to the pixel frequency turned out to provide the best results. Merely cutting the version of Fig. 3.5 with 256 averaged frames

into character cells, using those in lines 3–5 of the test text as a reference, and using the sum of all absolute values of pixel-value differences in a character cell as a decision metric, leads to the following recognition result for the remaining text:

```
The quick brown fox jumps over the lazy dog. THE QUICK BROWN FOX JUMPS OVER THE LAZY DOG! 6x13
!#$%&'()*+,-./0123456789:;<=>?@ABCDEFGHIJKLMNQRSTUvwxyz[\]^_`abcdefghijklmnopqrstuvwxyz{|}~
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. (from: Electromagnetic Radiation from Video Display Units: An Eavesdropping Risk?)
```

Only one single character (“electromagnetic”) is wrong in this example, which corresponds to a character error rate of 0.13 %. This result depends, of course, significantly on a good signal-to-noise ratio. When we apply the same matching algorithm on the signal generated from only 16 averaged frames, the recognized text reads as:

```
lhc quick bcown fox_jumps-vec-toe lazg dsg=TOE_QHICK-DROWM-EHX JUHPS Q?ER iUE LOZY DH6! -6zi3=
!#$%&'()*+,-=ZO!?'3'56709:;< >?@ADcDEFCHIJKLHncPQRHThVQ%YZ[\]^='abedcBg6Ijkmndpqcstuvw:yz{|}"
it Ic weII=kocwn=tHat-clctronic=cguipmct e_dduces-electrpmugmctc_fidldc_who-may euuse _-
= icce-feceae te-radix-and teIcvisicn cecpticc=-l6e phncmcna uedcrllyigg tcic=have=bcec= --
_-tncceughIy ctuHid=dvcc the eust few=decudes, ihsc stvdics'have =ecuItcd io_inteceutiocu_iy -
_ugrceH=mct6edc=foc meacuciny t6c icterfcsesce pouduccd_bg eeuipmct. These are-nccded bccouse
toc=meximum intcrferencnc ievcls which-eguipmct may gesc-atc-6ave oecn la7d=dewc=by law in mcsc
ceuntricc=-(fcm: Flectromegntic-Radiatibn f_om Video Dispiey_Hsitc:=Hn Eavcsdccc=pimg-Risk?)-
```

With a character error rate of 34 %, it is very severely distorted and not directly usable, for example, for full-text searching. But because most of the mis-recognized characters graphically resemble the correct character, it is still possible to guess most of the text.

Even though only 66 % of the characters were recognized correctly, half of the remaining ones (16 %) came second, and another 6 % third, on a list that shows, for each received pattern, all candidate symbols sorted according to how closely they match. In addition to the character error rate, we can introduce further performance quantities that take also into account, whether the correct character came as a close second or third.

If \mathbf{s} is a random vector representing the received signal for a character or symbol to be recognized, and $\{\mathbf{r}_1, \dots, \mathbf{r}_n\}$ are the reference signals available for the n characters to be distinguished, where \mathbf{r}_c is the correct one, then we can produce for each output of \mathbf{s} a sorted list (d_1, \dots, d_n) such that $\|\mathbf{s} - \mathbf{r}_{d_i}\| < \|\mathbf{s} - \mathbf{r}_{d_j}\| \Rightarrow i < j$, where $\|\cdot\|$ is the distance metric used (e.g., the vector norm $|\cdot|_1$ in the above experiment). Let $p_i = P(\mathbf{r}_c = \mathbf{r}_{d_i})$ be the probability that the correct character appears at the i -th position in that list. The *average depth of the correct character*² is then $\sum_{i=1}^n p_i \cdot i$ and the *entropy of the depth of the correct character* is $-\sum_{i=1}^n p_i \cdot \log_2 p_i$. The entropy measure, in particular, illustrates, how many bits of uncertainty the eavesdropper still has for each character. It can help to estimate, for example, the complexity remaining for a brute-force search to find a correct password, given that a distorted copy of it has already been recognized, and candidate passwords can therefore be tested in order of falling probability.

The average depth of the correct character in the above two experiments was 1.0013 and 2.0995, respectively. The entropy of the depth of the correct character was 0.0141 bit and 1.7778 bit, respectively.

²The term *average depth of correct symbol (ADCS)* is also mentioned in the NSA Tempest standard [8]. Its definition there remains classified, and we can only speculate whether it refers to the same order statistic used here.




Figure 3.10: This matrix indicates the pairwise difference of the radio signals of the ASCII characters in Fig. 3.5. Brighter positions indicate larger differences and dark spots signal character pairs that automatic recognition systems are more likely to confuse.

Figure 3.10 shows, for the font used in this example, the difficulty of distinguishing between all possible character pairs. A list of character pairs sorted by increasing distinguishability starts with

```
E!- cI /B0.c-8Te=-1 !OE""6 c_-."H8ICg!h=oaI69F,:7,I6F9/1''8Q,,eT.8
-_F|_ =e1'ZDU:o=0|o_'.|.TQL^'8',s'. _~UUiGqln'su|SSL.;Z;TGP07|^_HU-_s1=S
```

and ends in

```
Nn-1H8 y1R))QV% b;MNO(?R1!) '9Y1. !K'.Q1MNUhn)W(h@HMq;M!NHHUUQNh,MUM),)MN
|~NNj}M^d{UKTpyN{M1}{Nn}MUHN{pUMMjMN{Q_}j|_hjM{j|'}N|Q'}-{|}|{-}Mj{-}NNM{j
```

A font designer might use such difference matrices, in order to generate security fonts that minimize the success of automatic radio character recognition. But this is likely to also affect safety properties; glyphs that are too similar might be mixed up more easily by the regular viewer.

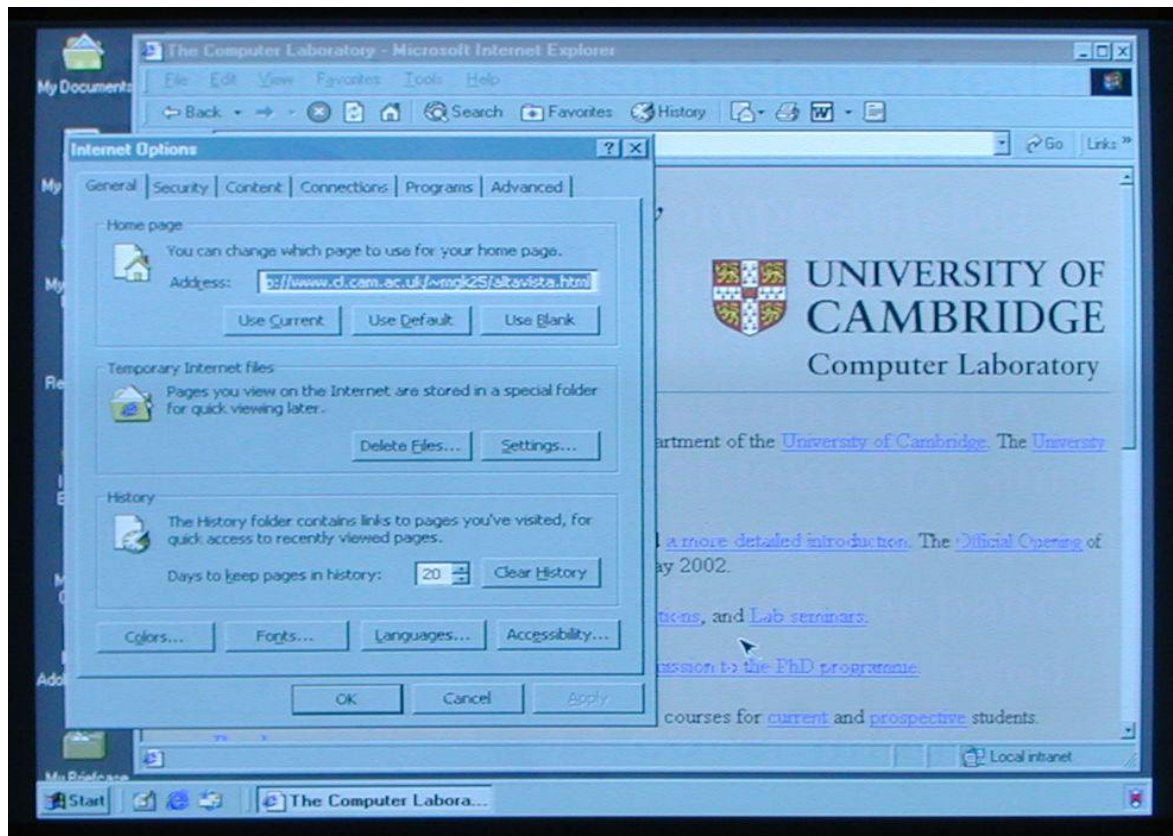


Figure 3.11: Photograph of the CRT screen display seen by the user while the hidden message from Fig. 3.14 is being transmitted.

3.5 Hidden transmission via dither patterns

The human eye becomes less sensitive for contrast as spatial frequency increases. Dithering or half-toning is a graphical technique that uses this effect to increase the number of color shades available on displays with a small color repertoire. On modern high-resolution CRT monitors, users cannot easily distinguish between a medium gray and a checkered halftoning pattern of black and white pixels, especially as the distance between pixels is often smaller than the diameter of the electron beam focus. Compromising video emissions on the other hand mainly carry the high-frequency part of the video signal. For the RF eavesdropper, a high-frequency black/white dither pattern on a CRT screen creates the strongest possible signal while a constant color minimizes emissions.

We can use this difference in the spectral sensitivity of the user and the eavesdropper to present different information to either. Figure 3.11 shows a test screen content on the same monitor as used in Section 3.3.2. Figure 3.12 shows the image captured at the same time with a receiver at 445 MHz center frequency and 10 MHz bandwidth.

The trick becomes clear in Fig. 3.13, which is a magnified representation of the pixel field seen on the target monitor. This example shows how software can covertly embed a grayscale image inside displayed screen content. It amplitude modulates the embedded signal using a maximum-frequency checkerboard pattern as the carrier, which it then adds to the cover image. The hidden example image shown in Fig. 3.14 demonstrates how sufficiently large embedded text can be made readable, even by merely using a 10 MHz bandwidth AM receiver, comparable to the one found in TV sets. The embedded photo

445 MHz center frequency, 10 MHz bandwidth, 1024 frames averaged, 3 m distance

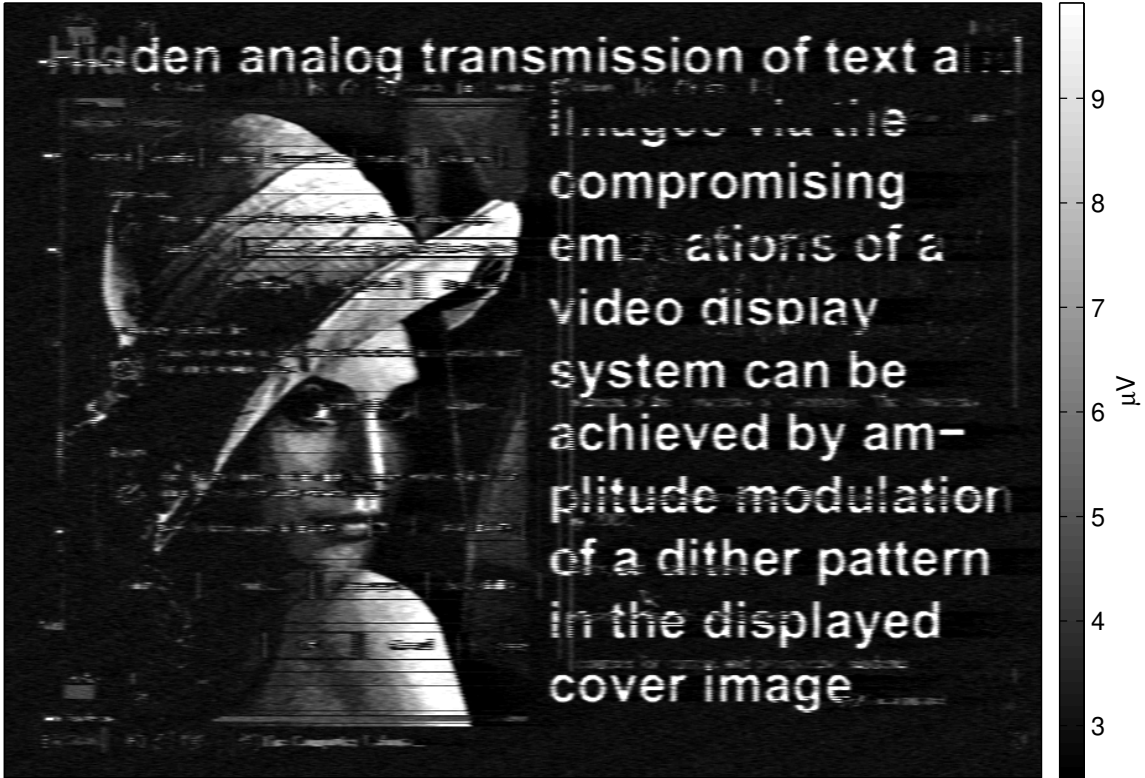


Figure 3.12: The hidden message as seen by the eavesdropper from a monitor that looks to the user in front of it as in Fig. 3.11.

gives an impression of the amount of information that can be transmitted in just a single static image.

The embedding software encounters a constraint in the form of the maximally possible range of pixel values, which might be fully used by the cover image. Adding the modulated embedded image could cause an overflow. If we simply truncated values outside the allowed range, parts of the embedded image would become visible. The following embedding algorithm handles this problem by reducing the amplitude of the embedded signal where necessary to avoid an overflow. This distorts parts of the embedded image in favor of visually affecting the cover image in very bright or dark areas. It also handles the luminosity non-linearity of CRTs.

Let $C_{x,y,c}$ be the value of a cover image at pixel coordinates (x, y) for primary color $c \in \{\text{red, green, blue}\}$ and let $E_{x,y}$ be the pixel value of the image that shall be embedded covertly for reception by the eavesdropper. Then the color component values that we have to display on the screen are

$$S_{x,y,c} = (C_{x,y,c}^{\tilde{\gamma}} + \min\{\alpha \cdot E_{x,y}, C_{x,y,c}^{\tilde{\gamma}}, 1 - C_{x,y,c}^{\tilde{\gamma}}\} \cdot d_{x,y})^{1/\tilde{\gamma}} \quad (3.16)$$

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

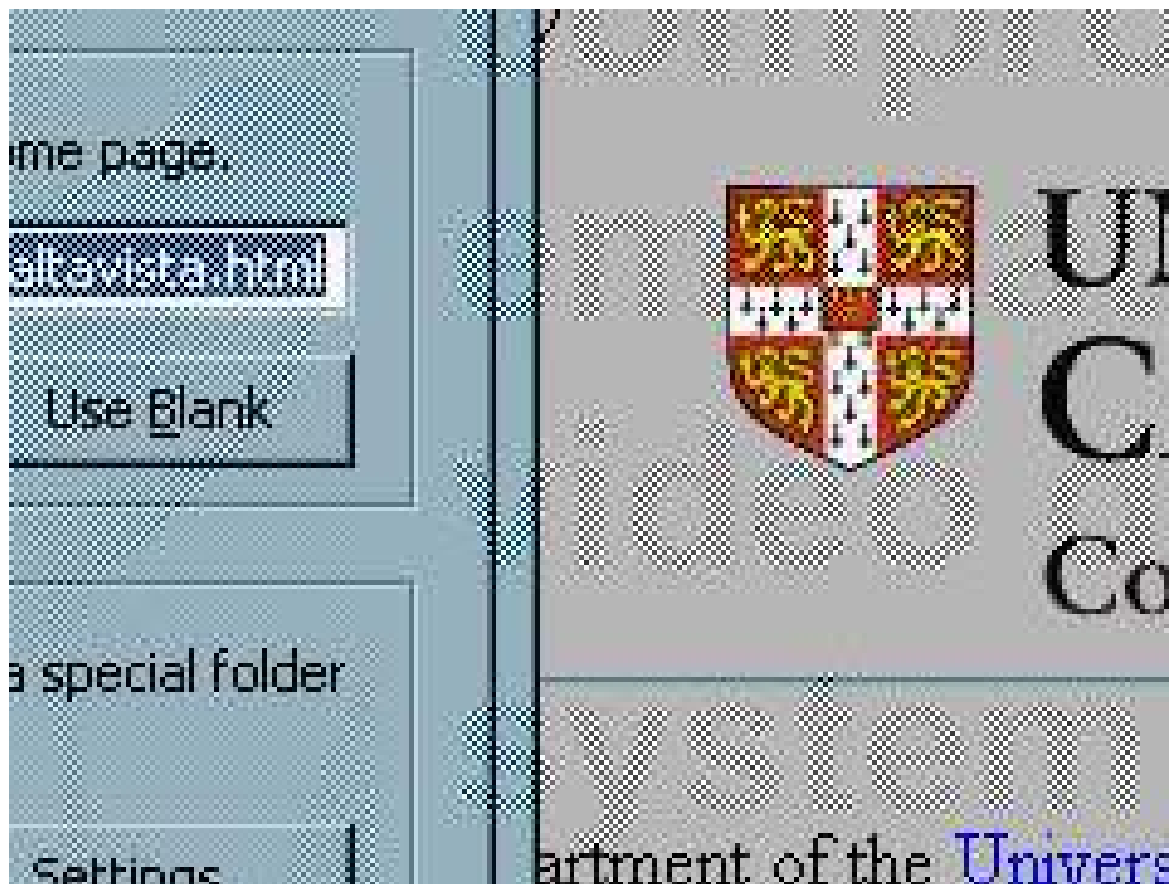


Figure 3.13: A magnification of the pixel field from Fig. 3.11 makes the luminosity-neutral amplitude modulation visible ($\alpha = 0.4$, $\tilde{\gamma} = 1.7$).

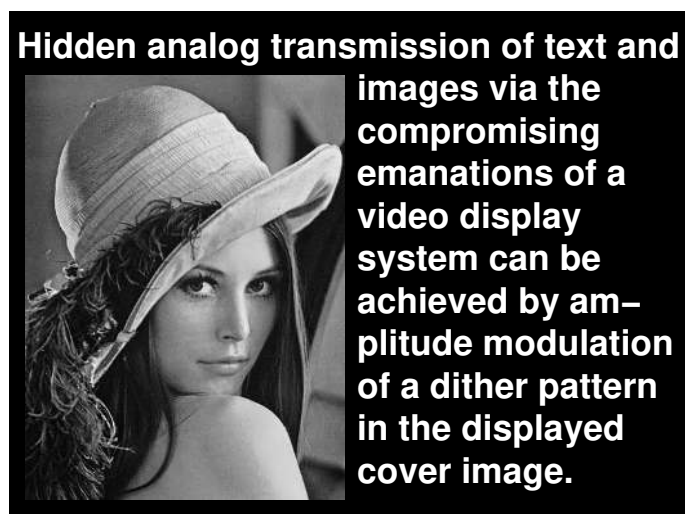


Figure 3.14: The embedded message that is covertly transmitted in this experiment.

The color component value chosen by the display software is usually mapped linearly to the video voltage supplied to the monitor. But the relation between the video voltage V and the luminosity L of the screen is non-linear and can be approximated as $L = \text{const} \cdot V^\gamma$, where γ is usually in the range 1.5–3.0 and depends on the design of the CRT. Software that compensates this non-linearity performs what is known as *gamma correction* [83]. The overall luminosity of a two-color dither pattern depends on the arithmetic mean of the luminosities L rather than the voltages V . To remain inconspicuous for the viewer, amplitude variations in the dither pattern must be performed such that the average luminosity is preserved.

I observed that the arithmetic average of the gamma-corrected luminosities only predicts the luminosity accurately for a dither pattern consisting of horizontal lines. For dither patterns with vertical lines or checkered patterns, the restricted bandwidth of the control-grid voltage introduces many intermediate values. An accurate luminosity estimation for such dither patterns with high horizontal frequency components – the ones of interest for hiding information in emissions – would involve integration of the distorted gamma-corrected video signal [84]. I performed tests in which I determined the video voltage \bar{V} that gives rise to a color of the same brightness as a dither mix of the voltages V_1 and V_2 . For a dither pattern of horizontal lines, the formula $\bar{V} = (\frac{1}{2}V_1^\gamma + \frac{1}{2}V_2^\gamma)^{1/\gamma}$ produced good predictions with $\gamma = 2.2$, the exponent for this CRT. For a checkered dither pattern, which looks much smoother, the same formula still worked, but the exponent changed to $\tilde{\gamma} = 1.7$. This is the value that we have to use to determine $S_{x,y,c}$. Modern monitors store gamma-correction parameters in their firmware, which computers can download. However, these are not sufficient to gamma-correct a high-amplitude checkered dither pattern.

The embedded image should be smoothed in order not to arouse the very sensitive edge detectors implemented in the human retina. Where the transmitted image must be very difficult to see, the correction parameter $\tilde{\gamma}$ should be manually calibrated for a specific monitor. The calibration depends not only on the type of monitor, but also on its brightness, contrast and color-temperature settings, which the user might modify. Readable text or barcodes can be better hidden in structurally rich content, such as photos or the animations shown by “screen saver” programs, rather than in uniformly colored areas.

3.6 Filtered fonts as a software protection

We can also use the difference in sensitivity to spatial frequencies of the CRT user and the RF eavesdropper for a simple software protection technique. It can be applied in many existing systems without any hardware changes. As the RF eavesdropper can only receive the upper portion of the baseband signal spectrum $V(f)$ (for $|f| < f_p/2$), attenuating these frequencies in the displayed pixel field can noticeably reduce the emitted radio signal without affecting visual readability severely.

Figure 3.15 shows, as an example, a widely used scalable word processing font (Microsoft’s “Arial”, scaled to 15 pixels-per-em) rendered in a number of different ways. The glyphs of scalable outline fonts, such as this one, are encoded as a set of control points for Bézier splines, which are rastered into a pixel field on request for a specified resolution by a rendering function. In the first line of Fig. 3.15, the font renderer was instructed to use only black and white pixels and to apply the TrueType “hinting” algorithm. Associated

- (1) The quick brown fox jumps over the lazy dog
- (2) The quick brown fox jumps over the lazy dog
- (3) The quick brown fox jumps over the lazy dog
- (4) The quick brown fox jumps over the lazy dog
- (5) The quick brown fox jumps over the lazy dog
- (6) The quick brown fox jumps over the lazy dog
- (7) The quick brown fox jumps over the lazy dog
- (8)

Figure 3.15: Arial rendered with different options and smoothing filters.




Figure 3.16: Impulse response curves of low-pass filters used in Fig. 3.15.

with the control points in a scalable font are special “hinting” instructions, in which the designer can specify, how to align the control points and the width of a glyph to the pixel grid. This is to improve the appearance of the rastered fonts, in particular at low resolutions. In the second line, the font was rendered with an anti-aliasing option enabled, in which partially covered pixels are rendered with a shade of gray, to reduce staircase effects. In the third line, the anti-aliasing option is still active, but hinting has been disabled, which can lead to more blurred edges, and also prevents changes to the size of glyphs.

Rendered with 1 bit per pixel, the glyphs show the largest number of maximum-contrast edges, which will be particularly well seen by an RF eavesdropper. In the second line, the hinting prevents that the anti-aliased font has significantly smoother edges. Vertical lines are still aligned for maximum edge contrast. With anti-aliasing activated and hinting disabled in line three, the contrast of the edges is varying randomly, as the pixel alignment of edges is no longer adjusted.

In order to obtain smooth edges on all occasions, it is necessary to low-pass filter a higher-resolution version of the rendered glyphs, before sampling them down to the de-

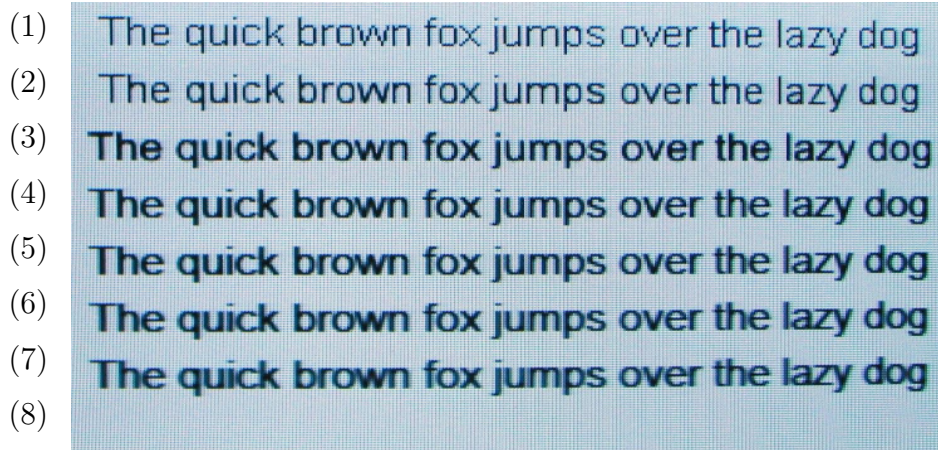


Figure 3.17: Photograph of the filtered-font test image (Fig. 3.15) as it appears on a CRT screen (120×60 mm).

sired resolution. This filtering only has to be applied along the horizontal axis (scanning direction of the electron beam), because only vertical edges are visible to the eavesdropper. The glyphs in lines (4) to (7) were produced this way. They were rendered with hinting enabled, such that the appearance of horizontal edges can benefit from the adjusted alignment. They were also rendered with a horizontal resolution four times larger than the vertical one, and they were subsequently convolved with one of the four finite-impulse-response filter-tap sets shown in Fig. 3.16. So if $0 \leq C_{x,y} \leq 1$ is the pixel value at coordinates x, y of the normally rendered text image, where x can be accessed in increments of $1/4$, then the filtered text image is

$$\tilde{C}_{x,y} = \alpha + (1 - 2\alpha) \cdot \sum_{i=-n}^n C_{x+i/4,y} \cdot r_i, \quad (3.17)$$

where $\alpha = 0.053$ is chosen to prevent a range overflow.

The coefficients r_i define a low-pass filter of order $2n = 20$ with a cut-off frequency of $\omega \in \{0.8/4, 0.7/4, 0.6/4, 0.5/4\}$ times the Nyquist frequency $f_p/2$, which leads, after horizontal subsampling by a factor of four, to signals that have the top 20 %, 30 %, 40 %, and 50 % of their spectrum attenuated, respectively. The filter coefficients were generated using the classical method of designing a digital linear-phase FIR filter, namely limiting the length of the $\sin(x)/x$ impulse response of an ideal low-pass filter by multiplying it with a Hamming window:

$$r'_i = \frac{\sin(\pi i \omega)}{\pi i \omega} \cdot [0.54 + 0.46 \cdot \cos(\pi i/n)], \quad r'_0 = 1 \quad (3.18)$$

$$r_i = r'_i / \sum_{j=-n}^n r'_j \quad i \in \{-n, \dots, n\} \quad (3.19)$$

The filtered text looks slightly blurred in the pixel-field representation of Fig. 3.15, but, as the screen photograph in Fig. 3.17 shows, the loss in edge contrast does not translate into a significant loss of text quality at the CRT screen. The limited focus of the electron

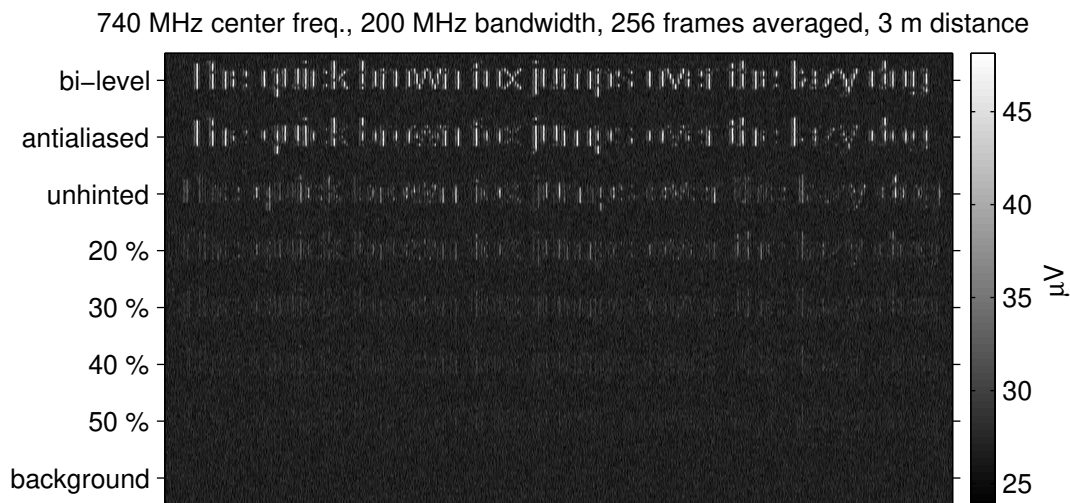


Figure 3.18: Received signal from the filtered-font test image.




Figure 3.19: Comparison of the received peak signal levels for the different lines of text in Fig. 3.18.

beam, the limited resolution of the eye, as well as effects created by the mask and the monitor electronics all filter and distort the pixel image anyway.

While there is little visible change for the user, the signal seen by an RF eavesdropper is deteriorated considerably, as Fig. 3.18 shows. Both the black/white and the hinted anti-aliased signals in lines (1) and (2) are clearly readable. The unhinted anti-aliased version is significantly more difficult to recognize, but still contains numerous clearly distinguishable edges that give away most character identities. Only the systematic horizontal low-pass filtering applied in lines (4) to (7) gradually attenuates the signal emitted by all edges equally.

Fig. 3.19 shows the peak voltage levels in each of the eight lines of the test image. Compared to removing only the top 20 % of the spectrum, the 30 % filter provides a significant reduction in peak signal strength. Filtering away even more, however, does not result in




Figure 3.20: Spectrum analyzer measurement of the video signal from the individual lines in figure Fig. 3.15 ($f_p \approx 49.5$ MHz).

very significant further reduction of signal peaks in this experiment. In practice, removing the top 30 % of the spectrum seems to be a good compromise between protection and rendering quality. With 50 % removed, the additional blurring effect starts to become quite noticeable on the screen (Fig. 3.17). Fig. 3.20 shows where in the spectrum on the video cable the various rendering techniques compared in this section remove energy.

In terms of peak signal power after AM demodulation, the 30 % filtered text was attenuated in this example by 11.7 dB compared to the bi-level representation. On the other hand, the attenuation offered by anti-aliasing and unhinted anti-aliasing alone is, at 0.4 and 1.1 dB, negligible in comparison. An attenuation by 10 dB means that, because of free-space attenuation, the eavesdropper has to come a factor three closer to the target to obtain a comparable signal level. In a mostly two-dimensional environment, this reduces the area in which potential hiding places can be located by a factor of 10; and if the eavesdropper can move in all three dimensions, the corresponding volume is even reduced 30-fold.

Filtered text display requires grayscale representation of glyphs, but this technology already is available in many display drivers, in order to support anti-aliased fonts; and a carefully filtered font can – especially at higher display resolutions – even be more visually appealing than a bi-level font.

Anti-eavesdropping display routines can also try to vary the shapes of individual glyphs slightly each time a character is displayed. Adding this to the use of proportional fonts and sub-pixel glyph widths will lead to further difficulty in the application of automatic radio-character recognition. Finally, to address the emanation risks of the digital components of a video system, in particular the data connection between the frame buffer and the ADC on the video card, the least-significant bits of a filtered text image should be randomized

individually for each displayed glyph occurrence. This will add a noise signal that cannot be removed by periodic averaging, a jamming approach that will be explored in more detail in the next section.

Chapter 4

Digital video displays

Most video terminals, the predominant human/computer interface from the mid 1960s to the mid 1980s, integrated both the frame buffer and the CRT with all associated circuitry inside a single unit. No user-visible interface between the two had to be defined by the manufacturer. This changed with the advent of modular personal-computer architectures, in particular the IBM PC and its successors. Displays and graphics cards turned into exchangeable components that are available from multiple vendors with standardized connectors. The signaling technique used on these interfaces evolved from digital to analog, and is now heading back to digital, with significant effect upon the nature of compromising video emanations.

For example, the original IBM PC contained the *Monochrome Display Adapter (MDA)*. It provided, along with the horizontal and vertical synchronization signals, a simple on/off voltage to signal whether each pixel is bright or dark. Soon afterwards the *Color Graphics Adapter (CGA)* represented the video information on four TTL-level digital lines, three for activating each of the primary colors red/green/blue separately, and a fourth to switch them all together between high and low brightness. This allowed the monitor to represent 16 color shades. During 1984–87, an extended version of the same 9-pin connector was used by IBM's *Extended Graphics Adapter (EGA)*. It provided six TTL lines to distinguish 64 color shades, using four intensity levels for each primary color.

Dropping memory prices then made it feasible to dedicate at first six bits and later an entire byte of memory to encode the intensity of each primary color. In April 1987, IBM introduced the *Video Graphics Adapter (VGA)*. To avoid the need for bulky 18 or 24-pin video connectors and cables, its designers decided to move the digital-to-analog converter from the monitor into the graphics card. This gave birth to the now ubiquitous 15-pin analog video connector, which represents the brightness of each primary color as a voltage in the range 0–0.7 V with 75 Ω impedance.

More recently, the industry has started to move back to digital video signaling between video controller and display 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. Whereas the 640×480@60Hz video mode used by the original VGA card had a pixel-clock frequency of merely 25 MHz, more recent high-end displays are driven with pixel rates of 300 MHz or more. As signal wavelengths drop below typical cable lengths, the lack of a properly shielded and impedance-matched coaxial feedthrough in the VGA connector starts to cause inter-pixel interference via cross-talk and reflections, limiting significant further increases in display resolution.

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. In order to maximize contrast, they buffer an entire line of the video signal and drive all pixels in a row concurrently for time intervals of length f_h^{-1} .

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 the more recent digital interface standards therefore include a pixel clock line. Flat panel displays that are provided with only an analog VGA signal first have to reconstruct the pixel clock signal using a phase-locked loop oscillator from the synchronization pulses, after which they can digitize the video voltages back into pixel values.

4.1 Case study: Laptop display

Compared to CRTs, which amplify the video signal by about a factor of 100 before applying it to the control grid to modulate the beam current, the components of an LCD seem to be far less likely candidates for forming significant sources of compromising radiation. Even though the transparent indium-tungsten oxide electrodes that form the row and column access matrix are several decimeters long, and could therefore form suitable dipoles for radiating at typical pixel-clock frequencies, they are only driven by a small voltage, and, more importantly, all column lines are switched in parallel. Therefore, pixels within the same row cannot be distinguished by the timing of any associated electromagnetic pulses.

In spite of this, I started eavesdropping experiments on laptops and encountered very strong and easily readable video emissions, in some cases with field strengths significantly larger than those encountered with CRTs. In order to understand the origin and nature of these emissions better, I decided to reverse engineer the display system of my own laptop, the Toshiba Satellite Pro 440CDX that already had served as a VGA signal generator in the previous chapter.

Figure 4.1 shows an example of a rastered amplitude-demodulated signal received from this laptop. Even without any frame averaging applied (bottom right corner) the text is clearly readable.

A number of observations distinguish these emanations from those of 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.
- Font glyphs appear to have lost half of their horizontal resolution, but are still readable.
- In the $800 \times 600 @ 75\text{Hz}$ default mode¹, 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 in steps of 25 MHz.

¹which is with $f_h = 47.45317$ kHz and $y_t = 628$ not a VESA standard mode

350 MHz center frequency, 50 MHz bandwidth, 16 (1) frames averaged, 3 m distance



magnified image section

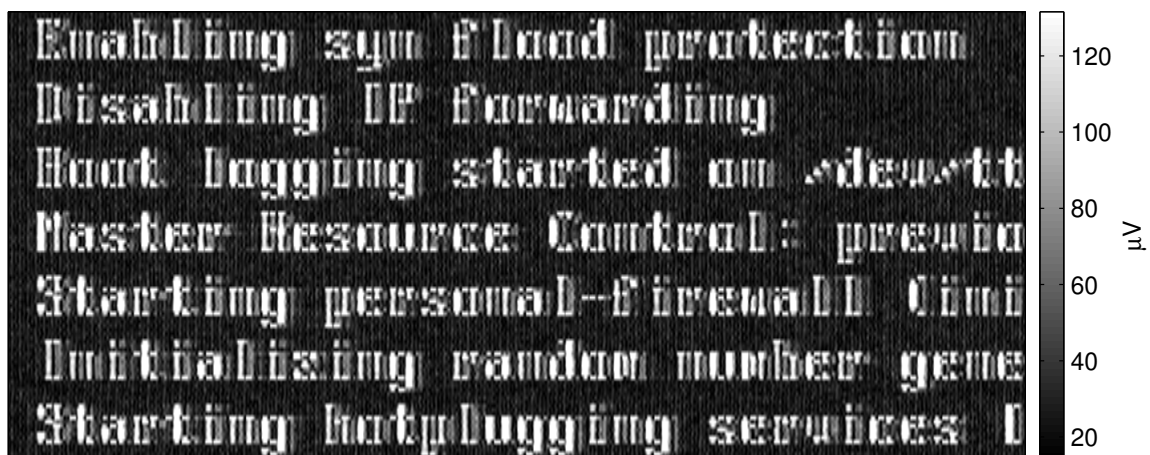


Figure 4.1: Eavesdropped Linux boot screen visible on the LCD of a Toshiba 440CDX laptop (log-periodic antenna, vertical polarization).

- 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 amplitude were somehow related to the binary representation of the pixel value.
- Using a simple near-field probe (a coaxial cable whose ends are shaped into a 50 mm wide dipole) instead of an antenna, I found no significant emissions from the display module itself, but located the source to be the interconnect cable between the LCD 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 [94]). The 18-bit color data provided by the video controller per pixel on its parallel output port has to be serialized into fewer lines, to fit through the hinges that connect the display panel to the laptop's main body. These two "FPD-Link" chips perform exactly this task. They multiply the clock signal supplied from the video controller by a factor of 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 data enable control signals. The fourth twisted pair carries the clock signal.

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 [95]) is a generic interface standard defined by the *Electronic Industries Alliance (EIA)* for high-speed data transmission (up to 655 Mbit/s). An LVDS channel is a twisted wire pair that is terminated with a 100 Ω resistor and driven with a 3 mA current source to generate an about 300 mV differential signal swing. The receiver consists of a differential amplifier across the termination resistor. One of the design goals of LVDS was to reduce electromagnetic interference by transmitting data such that the sum of the currents in the two conductors of a channel (the common-mode current) is approximately zero at any time. This technique, also known as symmetric or balanced transmission line, aims to prevent the lines from acting like a dipole antenna. Twisting the lines also reduces the area enclosed by the two conductors, to avoid the connection forming an effective loop antenna for the differential-mode current on the lines.

However, as Fig. 4.1 shows, these design precautions alone are not sufficient for providing emission security. The approximately 100 μV amplitude that I measured with the log-periodic antenna for the BIOS default colors used in this screen at 3 m distance corresponds to a field strength of 57 dB $\mu\text{V}/\text{m}$ (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. I placed the target laptop and the log-periodical antenna in two different offices² with two other offices in between. The target device and the antenna were separated in this experiment by a distance of 10 m as well as three 105 mm thick plaster-board walls.

Like a real eavesdropper, in this setup I did not use any access to the sync pulses of the laptop to overcome the performance restrictions of the averaging mode in the acquisition equipment (see Section 3.3.2). By operating the storage oscilloscope with a sampling frequency of 50 MHz, I started from a single manual trigger and recorded within 160 ms eight million samples, which covered 12 consecutive frame refreshes. I implemented an algorithm that calculates the cross-correlation of the first and the last frame in the recorded

²William Gates Building, Cambridge, ground floor




Figure 4.2: Determining f_v for the multi-frame signal recorded in Fig. 4.3 through crosscorrelation between the first and last frame in the recorded series.

series. It then searches for a large peak in the result based on a user-provided crude estimate of the expected refresh frequency f_v (see Fig. 4.2). This peak position is converted into a far more precise estimate of f_v , based on which it is then possible to align the individual recorded frames and sum them up for averaging. Due to other video signals in the vicinity, echos, and multiple peaks that can appear in the auto-correlation of a single video signal, it can sometimes be necessary to manually choose an alternative peak.

Figure 4.3 shows the result of this demonstration, an easily readable view of an `xterm` window that shows some test text. The received signal amplitude of about $12\ \mu\text{V}$ corresponds with this antenna to a field strength of $39\ \text{dB}\mu\text{V}/\text{m}$. This drop by 18 dB compared to the $57\ \text{dB}\mu\text{V}/\text{m}$ 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 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, which is discussed later in Section 5.2.2.

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 data transmission format used in this digital video interface. 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 $(r_7 \dots r_0, g_7 \dots g_0, b_7 \dots b_0)$. Figure 4.4 shows, how these bits are packed in a 440CDX laptop into the pixel cycle of three FPD-Link channels. Lacking a circuit diagram of the laptop, I determined this mapping by observing with an oscilloscope on the LVDS channels changes in the encoding when I varied the color of individual pixels. 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

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

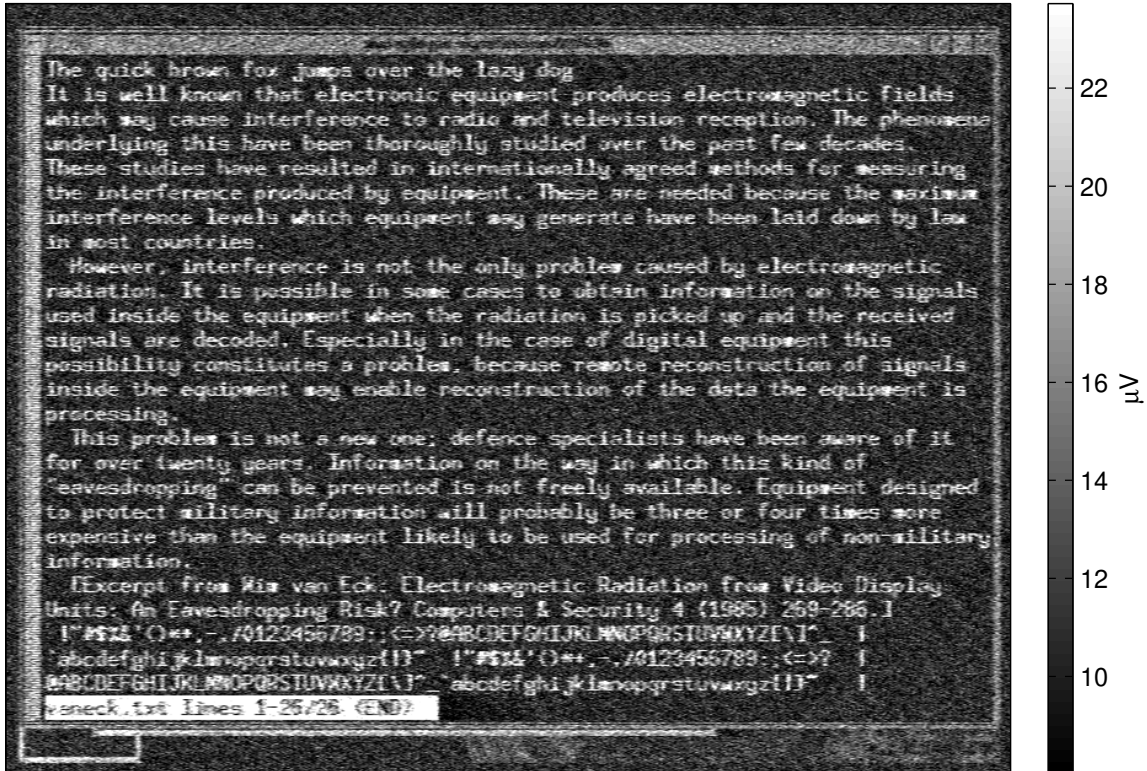


Figure 4.3: Text signal received from a 440CDX laptop at 10 m distance through two intermediate offices (3 plasterboard walls).

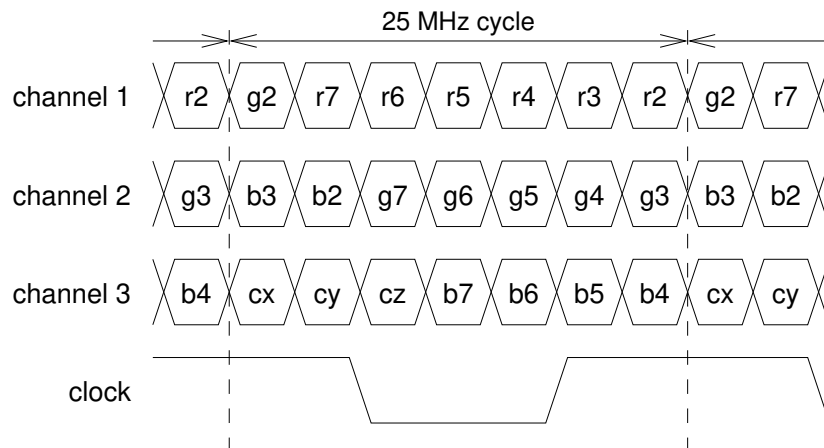


Figure 4.4: Bit assignment in the FPD-Link transmission cycle.

memory of this laptop supports the $800 \times 600 @ 75\text{Hz}$ 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 $r_2 = r_7 \wedge \dots \wedge r_3$ and $b_2 = b_7 \wedge \dots \wedge b_3$ automatically. One of the FPD-Link chips transmits all pixels in odd-numbered columns, the other one all 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 back-

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



Figure 4.5: Test text to compare the emission characteristics of selected foreground and background color combinations.

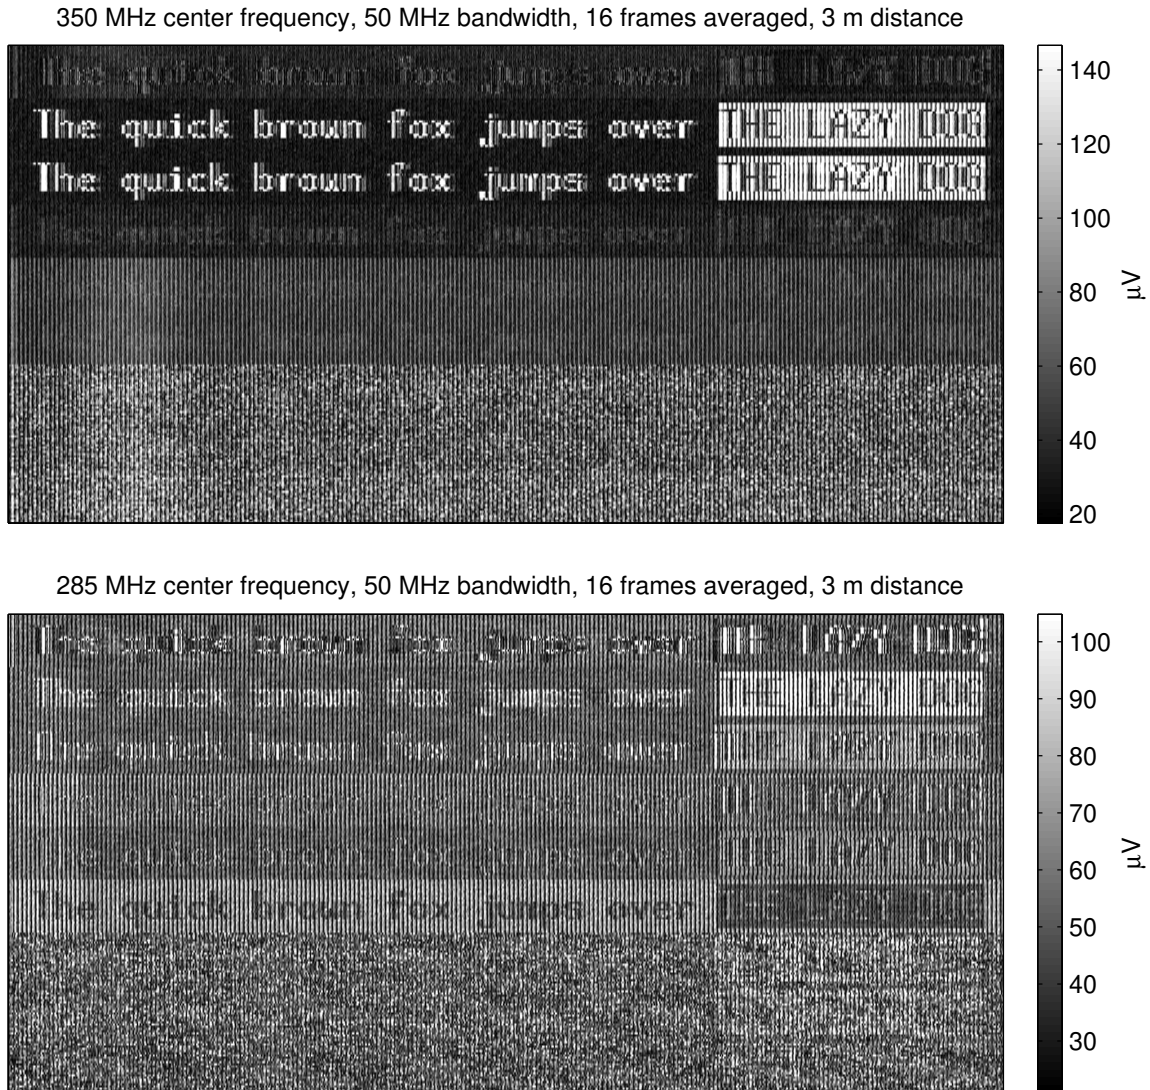


Figure 4.6: Signals received from the test display in Fig. 4.5.

ground of text in the emitted signal.

Figure 4.5 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 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 4.6 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 is occupied by strong FM radio transmitters, and the first harmonics collide in Cambridge with what appear to be signals from local VHF paging systems and so they cannot be used either. 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 $59 \text{ dB}\mu\text{V}/\text{m}$ 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. 4.6 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 (see Fig. 3.10), which would only reduce the error probability during automatic radio character recognition.

An application for maximum-contrast color combinations is the intentional broadcast of data. This can take place in the form of readable text, but the encoding symbols are far more reliably recognized automatically if standard digital modulation techniques are



Figure 4.7: The text “Tempest in a Teapot” encoded in ASCII (padded with start- and end-of-text control characters) DSSS-modulated.




Figure 4.8: Result of convolving the signal received from the display in Fig. 4.7 with the same random-bit sequence used there. Negative peaks mark the start of each byte and positive peaks are the 1 bits (MSB first).

utilized. As a simple example of how this can look like in practice, Fig. 4.7 shows a short ASCII string (“Tempest in a Teapot”) encoded using an approximation of a direct-sequence spread-spectrum modulation technique. I selected this modulation technique here mainly because it simplifies synchronization for the decoder.

A single bit is encoded as a 256-bit long random sequence, which is displayed in maximum emission contrast colors. A single line of Fig. 4.7 encodes a single bit. The presence of the random sequence signals a “one” bit, a line filled entirely in the background color signals a “zero” bit. The inverted random sequence appears as a start bit before every byte of the encoded text in order to assist in character synchronization.

Figure 4.8 shows the cross-correlation of the received signal (same parameters as in Fig. 4.1) and the random sequence. The result are positive peaks for “one” bits, and negative peaks marking the start of every byte. A simple decoder routine searches for the

start-bit peaks and then samples at the eight intermediate positions for the presence of the positive peak. It converts this curve back into the original text without any decoding errors: *Tempest in a Teapot*.

4.2 Case study: Digital Visual Interface

The NEC FPD-Link interface technology is currently 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 proposed in the past few years:

- VESA Plug & Display (P&D) [97]
- VESA Digital Flat Panel (DFP) [98]
- Digital Visual Interface (DVI) [99]

The P&D standard integrates in a single 34-pin connector a collection of interface standards: digital video (TMDS), analog video (VGA), Display Data Channel (DDC), IEEE 1394 (a high-speed serial bus for digital video cameras, etc.), and Universal Serial Bus (USB). The DFP standard defines a much simpler subset of the P&D standard, which carries in a 20-pin connector only digital video (TMDS) and DDC signals. The DVI standard finally includes in its 28-pin connector only digital video (TMDS), analog video (optional), and DDC. So far, the DVI connector appears to have won the best manufacturer support.

The differences between these three standards are not particularly relevant for the discussion in this section. We will only look at the properties of the digital video aspect of these interfaces, and all three interfaces use, in mutually compatible ways, a technology called *Transition Minimized Differential Signaling (TMDS)*, which is also known as *PanelLink* and was 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 twisted-line pair and carries 8-bit values for one of the three primary colors. A fourth twisted-pair channel provides a byte clock signal for synchronization. TMDS receivers connect each line in a pair via a $50\ \Omega$ terminator to a 3.3 V supply source and amplify the voltage difference between the two lines. TMDS transmitters contain a 10 mA current source that sinks current from one of the two lines – depending on the bit value transmitted – into ground, such that the voltage on that line drops 0.5 V below the supply source at the other end.

As far as the spectral content of the generated signal is concerned, the DVI standard requires that the rise/fall time (80/20 %) of any edge is at least 75 ps. This means in practice that the signals must be filtered to attenuate frequencies higher than 3 GHz. On the other hand, the rise/fall time has to be at least 40 % of a bit transmission time, which means that the cut-off frequency of a smoothing filter can be set as low as 55 % of the bit rate according to equation (3.13).

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 $d_7d_6 \dots d_0$ is counted. A new 9-bit value q is generated by setting $q_0 = d_0$ and

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

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

$$\begin{aligned} q_i &= \neg q_{i-1} \oplus d_i && \text{for } (1 \leq i \leq 7) \\ q_8 &= 0 \end{aligned}$$

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 $q_7q_6 \dots q_0$ 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):

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

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), because the minority bit can only appear up to four times in a byte, and each presence is signaled 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 [99, 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:

$$\begin{aligned} &00000000, 00000000, 00000000, 00000000, 00000000, \dots \longrightarrow \\ &0000000010, 1111111111, 0000000010, 1111111111, 0000000010 \\ &1111111111, 0000000010, 1111111111, 0000000010, \end{aligned}$$

```

0000000010, 1111111111, 0000000010, 1111111111, 0000000010
1111111111, 0000000010, 1111111111, 0000000010,
...
11111111, 11111111, 11111111, 11111111, 11111111, ... →
0000000001, 1111111100, 1111111100, 0000000001, 1111111100
0000000001, 1111111100
0000000001, 1111111100, 1111111100, 0000000001, 1111111100
0000000001, 1111111100
...
```

The name of the TMDS encoding suggests that its design was motivated by the desire to reduce the number of transitions, which should also reduce radio interference. However, it fails to achieve this goal. When encoding a sequence of 10^4 uniformly distributed random bytes with TMDS, I observed that the number of transitions actually increased by 3 %. Adding two bits increases the transition count, especially for the many byte values whose regular encoding has only few transitions. As the above examples also illustrate, in the very common white and black pixel sequences, two transitions are added per pixel.

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 fulfill these criteria:

```

00001000, 00001000, ... → 0000111110, 0000111110, ...
10101010, 10101010, ... → 1100110010, 1100110010, ...
```

These TMDS bit patterns will be used irrespective of the previous status of the bit balance, because the full encoding algorithm specified in [99, p. 29] contains a special case that 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.

Figure 4.9 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 4.10 shows the signals received from a DVI display system that shows the test display of Fig. 4.9. The graphics card in this setup was an “ATI Rage Fury Pro” and the display a “Samsung SyncMaster 170T”.

The $1280 \times 1024 @ 60\text{Hz}$ video mode used in this setup has a pixel clock frequency of 108 MHz. This pixel frequency is below the 165 MHz limit above which the DVI standard

line	description	foreground RGB	background 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	—	—



Figure 4.9: Example test image for text contrast in compromising emanations from DVI cables.

requires the second TMDS link to be activated to carry half of all pixels, therefore only one TMDS link is used, and the bitrate in each of its three channels is 1.08 GHz.

Fourier theory and the convolution theorem can again be used to explain the spectral composition of the signal on a TMDS channel in this example. 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 $\mathcal{F}\{t_{55}\}$ is a line spectrum with lines at 108 MHz, 216 MHz, 324 MHz, \dots , 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

$$w = v \cdot t_{55} + (1 - v) \cdot t_{10} = v \cdot (t_{55} - t_{10}) + t_{10}. \quad (4.1)$$

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

$$W = V * \mathcal{F}\{t_{55} - t_{10}\} + \mathcal{F}\{t_{10}\}. \quad (4.2)$$

In other words, the spectrum of the pixel-value waveform V will be copied into 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 $\mathcal{F}\{t_{55} - t_{10}\}$. Figure 4.11

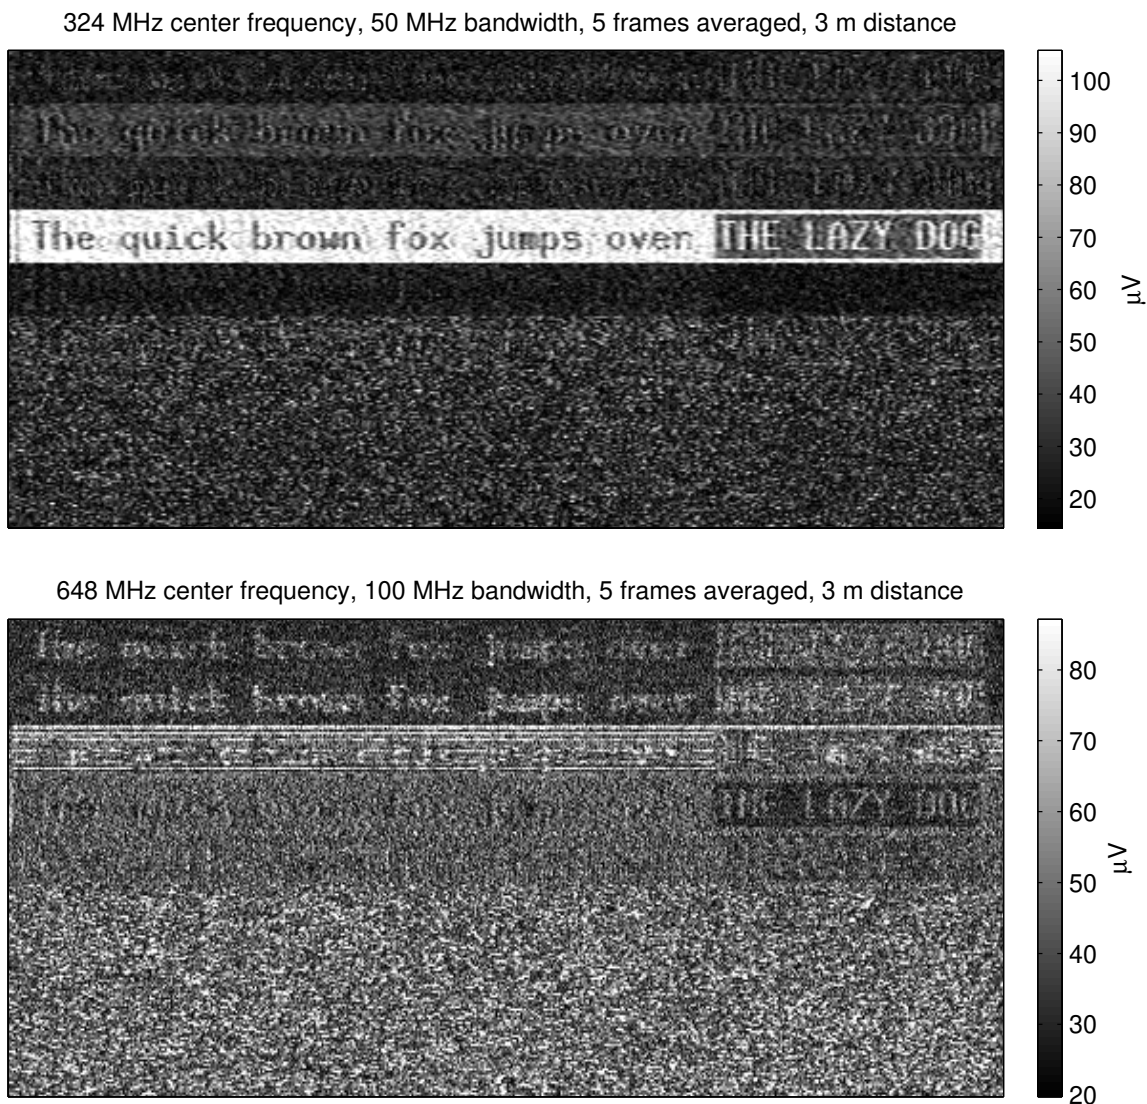


Figure 4.10: Received emanation in two frequency bands from a DVI cable transmitting the text image of Fig. 4.9.

illustrates the relative intensity of the spectral lines of $|\mathcal{F}\{t_{10}\}|$, $|\mathcal{F}\{t_{55}\}|$, and $|\mathcal{F}\{t_{55} - t_{10}\}|$. It also shows the line spectrum $|\mathcal{F}\{t_{55}\}| - |\mathcal{F}\{t_{10}\}|$, 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 (see Section 3.2).

The center frequency of 324 MHz used in Figure 4.10 is not the strongest line in the spectrum of $|\mathcal{F}\{t_{55}\}| - |\mathcal{F}\{t_{10}\}|$, 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 of about $100 \mu\text{V}$ at the receiver input, corresponding to a field strength at 3 m of $56 \text{ dB}\mu\text{V}/\text{m}$ (50 MHz), which is again equivalent to an isotropic radiated power of about 120 nW, comparable to what I measured earlier in Section 4.1 for the laptop.

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




Figure 4.11: Time and frequency domain representation of the TMDS-encoded maximum contrast byte combination hexadecimal 10 and 55 as well as their difference signal.

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 are used to represent the four possible combinations of the horizontal and vertical sync signals. These words contain eight bit transitions each, therefore they can never be generated by a normal color value and they unambiguously identify a blanking signal as a result. This makes them well suited for an eavesdropper to lock onto using cross-correlation, in order to reconstruct synchronization information, which is further simplified by the fact that the DC-balancing logic is deactivated during the blanking period.

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 op-

erate 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. Accessing the frame buffer in the display would still be a periodic process that might cause exploitable emanations, but as it is located closer to the device itself, shielding becomes simpler, and for flat-panel displays, a parallel readout of all pixels in a row can be implemented to entirely avoid the risky serialization of pixels.

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. The decryption in the device and subsequent processing would still be a periodic process that might leak some information unless the circuitry is specifically designed with emanation security in mind.

Even a cryptographically weak key exchange protocol, such as the one published in a first HDCP draft [100], 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 $h(R)$ using a brute-force search over all possible bit errors.

Chapter 5

Emission limits

Since about 1960, the military organizations of NATO countries have paid considerable attention to the problem of limiting compromising electromagnetic and acoustic emanations by developing test standards and conforming protected products. For example, according to estimates quoted in [7] in 1990, about half of all US government purchases of information technology equipment are made by military or diplomatic organizations, and about 20 % of these are “Tempest” certified products. This has been estimated to amount to several billion US dollars annually for the US alone.

In spite of these significant investments, the relevant US and NATO test standards, as well as their rationales, are still classified documents. This is undesirable in a number of ways:

- Users are unlikely to demand conformance of purchased products to secret standards unless they are legally required to do so. They have no idea what exact level of protection is tested and how the unknown cost-versus-protection tradeoffs made in these specifications fit into their overall security concept and budget. Demanding conformance to secret military specifications also severely restricts the choice of suppliers to a small number of defense contractors with the necessary security clearances and complicates independent verification.
- The traditional high level of secrecy surrounding the “Tempest” research, which seems to have been conducted primarily by employees of signal intelligence agencies and defense contractors, has led to somewhat unusual gaps in laws and civilian government specifications. For example, government regulations for civilian information-processing facilities with very high protection requirements, such as top-level certification authorities for digital signature infrastructures, either lack emission-security requirements completely or mention them only in generic terms without providing actual conformance requirements (e.g., [93]). Due to the lack of suitable public emission security standards, this can in some cases be equivalent to a de-facto requirement for using equipment conforming to secret specifications from government-licensed defense contractors.
- Secret “Tempest” specifications will not enjoy the continued quality assurance offered by public scrutiny and open academic research. Such peer review and feedback has led in the past repeatedly to significant improvements of technical standards,

for example, in cryptology, even where open research initially lags a decade or two behind the classified state of the art.

- The underlying experimental data and measurement techniques that led to the design of the “Tempest” test standards are currently not accessible for open academic and industrial research. This hinders further development of the field, for example, towards innovative low-cost solutions suitable for economic mass production by regular office-equipment and consumer-electronics vendors.
- As a consequence of the lack of related public research literature, civilian product designers are not educated in emission security. Opportunities for simple lowest-cost countermeasures that merely require awareness of the nature and mechanisms of emission security risks early in a design process are therefore missed.

Recent requests to have the US “Tempest” standards declassified via the Freedom of Information Act have only resulted in the publication of some excerpts. These cover merely standard terminology as well as test and calibration methods that are already widely documented in the electromagnetic compatibility literature. The actual conformance limits and exact test procedures remain unavailable [8, 9, 10, 11, 12].

There are no indications that the community of “Tempest” researchers linked to NATO signals intelligence agencies will release its standards in the foreseeable future. The reasons for this have not been published and we can only speculate. Currently exploited sources of intelligence may involve techniques that are part of the rationale for the emission limits. It is to be expected that signal intelligence agencies are primarily interested in preserving their capabilities, which would be helped by the lowest possible awareness of emission-security risks among users and product designers. If there is such a conflict of interest, it would justify the more independent public emission-security research, as the threat is then not merely a theoretical possibility, but a routinely exploitable vulnerability today.

The number of civilian information-processing facilities that have high protection requirements, and whose infrastructure is based on commercial off-the-shelf technology procured according to open standards, increases continuously, as does our society’s reliance on their security. This calls, in my opinion, for a new generation of emission-security test standards that is based entirely on published data and experimental techniques. Their development should follow the established procedures of international standardization organizations. Any underlying data should be open to scrutiny by academic peer review, in order to prevent that any tradeoffs that have to be made might be influenced by conflicting concerns of the signal-intelligence community. A model for such an effort could be the work that led to the international standardization of emission limits for electromagnetic compatibility [90].

5.1 Existing public standards

No public emission-security standards exist today. Two types of electromagnetic-emission limits for information technology have been widely accepted by the market, but neither was designed to reduce the risk of information-carrying emanations, or is even remotely suited to do so.

5.1.1 Ergonomic standards

Shortly after 1990, many manufacturers of CRT computer monitors introduced new “low radiation” models with improved electromagnetic shielding. These products conform to ergonomic/hygienic standards, aimed at reducing the exposure of humans to electromagnetic fields and their potential biological effects [85, 86, 87]. The TCO’92 specification developed by the Swedish Confederation of Professional Employees (TCO) imposes the following emission limits:

- The electrostatic potential of the screen surface must not exceed 0.5 kV.
- Alternating electric fields must not exceed 10 V/m in a 5–2000 Hz passband and also must not exceed 1.0 V/m in a 2–400 kHz passband, both 30 cm in front of the screen and 50 cm around the unit.
- Alternating magnetic fields must not exceed 200 nT in a 5–2000 Hz passband 30 cm in front of the screen and 50 cm around the unit and also must not exceed 25 nT in a 2–400 kHz passband 50 cm around it.

This standard limits only low-frequency fields below 400 kHz, which are generated by CRT deflection coils. Compromising emanations are typically significantly weaker and occur at much higher frequencies in the HF/VHF/UHF bands (3 MHz–3 GHz). Therefore, a TCO’92 conformance test will not provide any information about the emission-security properties of a device.

5.1.2 Radio-frequency interference standards

The second class of publicly available electromagnetic emanation standards is aimed at minimizing interference with radio communication services. Electromagnetic fields are also generated by devices that were not designed to transmit information. These can cause distortions in radio receivers and this became an increasing concern during the first half of the last century, when a number of countries started to introduce legally binding test standards and limits (e.g., the German VDE standards in 1949 or the US FCC regulations in 1954).

Thanks to the work of the *Comité International Spécial des Perturbations Radioélectriques (CISPR)*, the various national electromagnetic compatibility standards have been harmonized during the past decade. Manufacturers of information technology equipment now only have to ensure that their products conform to the CISPR 22 specification or the equivalent European standard EN 55022 [89], in order to fulfill the legal requirements on electromagnetic emissions in most countries. This standard imposes the following radiated emission limits:

- Electric fields must not exceed 30 dB μ V/m at 10 m distance in any 120 kHz passband in the frequency range 30–230 MHz.
- Electric fields must not exceed 37 dB μ V/m at 10 m distance in any 120 kHz passband in the frequency range 230–1000 MHz.

The field strength is determined with a special AM measurement receiver with a *quasi-peak (QP)* detector specified in CISPR 16-1. The output of this detector rises with a time constant of 1 ms, falls with a time constant of 550 ms and is displayed with a critically dampened mechanical indicator with an oscillation period of 100 ms or an equivalent implementation.

In addition, during the measurement, the mains power connector is plugged into a line impedance stabilization network (LISN) which provides a well-defined impedance of both live and neutral against ground, namely $50\ \Omega$ in parallel with $50\ \mu\text{H}$. The voltage measured across this impedance must

- not exceed 66–56 dB $\mu\text{V}/\text{m}$ with a quasi-peak detector and 56–46 dB $\mu\text{V}/\text{m}$ on average in any 9 kHz passband in the frequency range 150–500 kHz, where the limit decreases linearly with the logarithm of the frequency,
- must not exceed 56 dB $\mu\text{V}/\text{m}$ with a quasi-peak detector and 46 dB $\mu\text{V}/\text{m}$ on average in any 9 kHz passband in the frequency range 0.5–5 MHz,
- must not exceed 60 dB $\mu\text{V}/\text{m}$ with a quasi-peak detector and 50 dB $\mu\text{V}/\text{m}$ on average in any 9 kHz passband in the frequency range 5–30 MHz.

On communications lines, the common-mode current measured with a current probe and specified decoupling measures must

- not exceed 40–30 dB $\mu\text{A}/\text{m}$ with a quasi-peak detector and 30–20 dB $\mu\text{A}/\text{m}$ on average in any 9 kHz passband in the frequency range 150–500 kHz, where the limit decreases linearly with the logarithm of the frequency,
- must not exceed 30 dB $\mu\text{A}/\text{m}$ with a quasi-peak detector and 20 dB $\mu\text{V}/\text{m}$ on average in any 9 kHz passband in the frequency range 0.5–30 MHz.¹

These are the limits for “Class B” devices, which are sold for general domestic use. The standard also specifies less strict limits for “Class A” devices that are for use in industrial environments, where for instance the radiated limits are 10 dB higher, or equivalently measured a factor $\sqrt{10}$ further away (30 m).

With the European Union’s electromagnetic compatibility directive [88], manufacturers became responsible for guaranteeing conformance to the CISPR 22 interference limits. This led to a significant increase in training and general awareness of radio interference and immunity problems among product designers and it also resulted in the purchase of EMC test and measurement equipment by vendors. As a result, many information technology products sold in Europe became more carefully designed with regard to radio-frequency emission in general, in order to eliminate early in the design phase the risk of failure in EMC compliance tests later. As a side effect, some European post-1990 products show slightly improved emission security as a result of the increased EMC awareness, even though the legally enforced test standards were not designed for that purpose.

A brief look at the motivation and design of the EMC test standards helps to understand why they are not suited for emission security purposes. Radio broadcasters aim at ensuring

¹The corresponding voltage limits are 44 dB higher, which is equivalent in the case of a $150\ \Omega$ line impedance.

a minimum field strength of about 50–60 dB μ V/m throughout their primary reception area [126]. The CISPR limits were selected about 20 dB below that level to ensure that, at 10 m distance, the interference from a device will not limit the received signal-to-noise ratio to less than 20 dB.

Radio receivers are also quite sensitive to high-frequency interference signals that reach them via the power-supply cable. Both mains network and communication lines can act as large antennas for the emission of interference signals. Conducted emissions are only limited below 30 MHz, because for higher frequencies, radiated emissions are considered to be the dominant effect and signals on mains and communication cables will show up in measurements as radiated emissions as well.

For frequencies below 30 MHz, the local mains network will become part of the emitting antenna. In the interest of better reproducibility of test results, the signal emitted into the cable is measured directly across a well-defined impedance at the end of the power supply cable provided with the product. The conducted emission limits for mains power lines are lower than those for telecommunication ports, as radio receivers are exposed to power-line noise directly, whereas communication lines just contribute to radiated emissions.

The quasi-peak detector is used as a psychophysical estimation tool. It provides a measure of the approximate annoyance level that impulses of various strengths and repetition frequencies cause for human users of analog audio and television receivers. Strong disturbance impulses are tolerated if they occur sufficiently rarely, and even weak disturbances can be annoying at high repetition rates. The quasi-peak limits are higher than the average limits, as tests have shown that occasional broadband impulses are less annoying to AM/FM radio and TV users than a continuous narrow-band interference signal.

5.2 Considerations for emission security limits

Eavesdroppers can work with significantly lower signal levels than what might cause interference with radio and TV reception. They are concerned about how the emitted compromising signal compares in strength to the background noise, not to a broadcasting station. They can be expected to

- use high-gain antennas directed towards the emitting target device,
- look for broadband impulse signals in a quiet part of the spectrum without interference from broadcasting stations,
- use notch filters to suppress strong narrow-band broadcasting stations that interfere with the eavesdropped signal,
- use signal processing techniques such as periodic averaging, cross-correlation, digital demodulation, and maximum-likelihood symbol detection, in order to separate the wanted information-carrying signal from unwanted background noise,
- connect (with suitable high-pass filters and protection circuits) the input of sensitive receivers directly to mains and communications cables, or any other metal structures that connect to points near the target device, in order to pick up signals orders of magnitude weaker than those that could cause conducted interference in a normal radio receiver.

The emission limits therefore have to be based on an understanding of reasonable best-case assumptions of

- the minimal background noise that the eavesdropper faces even under good receiving conditions,
- the gain from antenna types that can be used covertly,
- the gain from the use of suitable detection and signal processing methods for the signal of interest,
- the closest distance between antenna and target device for which protection is needed.

One approach for designing an emission-security test standard is to provide an upper bound for the signal-to-noise ratio S/N that a radio-frequency eavesdropper could achieve. The major contributing factors to consider are

$$S/N = \frac{\hat{E}_B \cdot G_a \cdot G_p}{a_d \cdot a_w \cdot E_{n,B} \cdot f_r}, \quad (5.1)$$

where

- \hat{E}_B is the maximum field strength that the test standard permits,
- B is the impulse bandwidth used in the test,
- a_d is the free-space path loss caused by placing the eavesdropper's antenna at distance d from the target device, instead of the antenna distance \hat{d} used during the test,
- a_w is the additional attenuation caused by building walls,
- G_a is the gain of the best directional antenna that is feasible for use by the eavesdropper,
- G_p is the processing gain that can be achieved with techniques such as periodic averaging,
- $E_{n,B}$ is the field strength of natural and man-made radio noise at the location of the eavesdropping antenna within a bandwidth B ,
- f_r is the noise factor of the eavesdropper's receiver.

The signal-to-noise ratio S/N available to an eavesdropper on a conductor can be estimated in a similar way as

$$S/N = \frac{\hat{U}_B \cdot G_p}{a_c \cdot U_{n,B} \cdot f_r}, \quad (5.2)$$

where

- \hat{U}_B is the voltage of the signal of interest within a bandwidth B as it leaves the device on a short cable into a typical impedance,²

² \hat{U}_B can also be derived from a current determined with a current probe or absorber clamp.

a_c is the attenuation that the signal suffers along the network of conductors on its way to the receiver (due to impedance mismatches, radiation loss, resistance, filters, etc.),

$U_{n,B}$ is the root-mean-squared background noise voltage at the receiver input,

and f_r and G_p are as before.

The expected noise levels and attenuation values in the above equations are random variables, which, in the absence of better data, have to be modeled as being normally distributed with some mean and variance determined from the statistical evaluation of a large number of measurements in various environments. For the other parameters, reasonable estimates based on practical demonstrations have to be made, so that emission limits \hat{E}_B and \hat{U}_B can be selected that will keep the eavesdropper's signal-to-noise ratio below an acceptable level with sufficient probability. Different types of target signal are located in different frequency bands and permit different processing gains. Therefore, the above parameters will have to be estimated separately for each signal type of interest. General emission limits will have to consider for each frequency band the lowest acceptable source signal strength \hat{E}_B and \hat{U}_B for all types of compromising signals.

5.2.1 Radio noise

A standard survey-data reference for the noise levels to be expected in various environments throughout the radio spectrum exists in the form of ITU-R Recommendation P.372 [101], which summarizes the results from numerous noise intensity measurements and categorizes their origin.

The mean noise levels are provided in ITU-R P.372 in form of the *external noise factor*

$$f_a = \frac{P_{n,B}}{kT_0B} \quad (5.3)$$

and the *external noise figure*

$$F_a = 10 \text{ dB} \cdot \log_{10} f_a \quad (5.4)$$

where $P_{n,B}$ is the available noise power from a (hypothetical) lossless antenna, $k = 1.38 \times 10^{-23}$ J/K is Boltzmann's constant, $T_0 = 290$ K is a reference temperature, and B is the bandwidth of the receiver. In other words, the noise level is described relative to the thermal noise level of the antenna. Any resistor R at temperature T produces within a bandwidth B a root-mean-squared thermal-noise voltage $U_{n,B} = \sqrt{4kTRB}$. By connecting R to a matched load resistor R , the maximum available power $U_{n,B}^2/4R = kTB$ can be extracted, which makes kTB the *available noise power* of this source.

In order to convert the power available from an isotropic antenna into an equivalent field strength, we use equations (A.32) to (A.34) and get

$$E_{n,B} = \sqrt{\frac{4\pi P_{n,B} Z_0}{\lambda^2 G}} = f \cdot \sqrt{f_a \cdot B \cdot \frac{4\pi k T_0 Z_0}{c^2 G}} \quad (5.5)$$

where $Z_0 = \sqrt{\mu_0/\epsilon_0}$ is again the impedance of free space, $G = 1$ is the gain of an isotropic antenna, and $f = c/\lambda$ is the frequency. Expressed in the logarithmic units that are used

in [101], this is with (5.4) equivalent to

$$\begin{aligned}
 20 \text{ dB} \cdot \log_{10} \frac{E_{n,B}}{1 \text{ V/m}} = & \\
 F_a + 20 \text{ dB} \cdot \log_{10} \frac{f}{1 \text{ Hz}} + 10 \text{ dB} \cdot \log_{10} \frac{B}{1 \text{ Hz}} + & \\
 20 \text{ dB} \cdot \log_{10} \sqrt{\frac{4\pi k T_0 Z_0}{c^2 G} \cdot \frac{1 \text{ Hz}^3 \times 1 \text{ m}^2}{1 \text{ V}^2}} & \quad (5.6)
 \end{aligned}$$

and with constant expressions evaluated and scaled to more convenient SI prefixes, we obtain

$$\begin{aligned}
 20 \text{ dB} \cdot \log_{10} \frac{E_{n,B}}{1 \text{ } \mu\text{V/m}} = & \\
 F_a + 20 \text{ dB} \cdot \log_{10} \frac{f}{1 \text{ MHz}} + 10 \text{ dB} \cdot \log_{10} \frac{B}{1 \text{ MHz}} - 36.8 \text{ dB}. & \quad (5.7)
 \end{aligned}$$

One of the external noise figure curves provided in [101] is for atmospheric noise caused by remote lightning, which occurs up to about 30 MHz. The minimum expected F_a level that is exceeded 99.5 % of time starts at 150 dB at 10 kHz and drops proportional to the logarithm of the frequency to 0 dB at 1 MHz. Starting from about 80 kHz upwards, man-made noise is the dominating factor, even at a quiet rural receiver site, with a mean noise figure of

$$F_a = 53.6 \text{ dB} + 28.6 \text{ dB} \cdot \log_{10} \frac{f}{1 \text{ MHz}}. \quad (5.8)$$

In more populated areas, [101] also provides mean noise figures

$$F_a = a + 27.7 \text{ dB} \cdot \log_{10} \frac{f}{1 \text{ MHz}}. \quad (5.9)$$

with $a = 67.2$ dB for rural, 72.5 dB for residential, and 76.8 dB for business areas in the range 0.3–250 MHz, but it can be up to 8 dB lower for the lowest decile of measurements. Above 4 MHz, galactic noise becomes the most significant contributor at quiet receiver sites, but not in more populated areas.

Figure 5.1 uses these values, and others from [101], to estimate electric field levels at both quiet rural sites and business districts for 1 MHz bandwidth. In measurements, the receiver bandwidth will have to be smaller than the center frequency, therefore the shown curves are directly applicable only for frequencies of about 2 MHz and higher. For lower frequencies, lower bandwidths will have to be used.

It is worth noting that these are out-door levels and that this background noise might be attenuated if the eavesdropper and the target device are both located in the same building.

5.2.2 Radio signal attenuation

In free space (vacuum, dry air, etc.) the power flux density (power per area) of a radio signal drops with the square of the distance from a point source, because preservation of energy requires that the power flux density remains constant when it is integrated over




Figure 5.1: Expected electric field noise levels $E_{n,1 \text{ MHz}}$ excluding transmitter stations based on data from ITU-R P.372 [101]. The curves are for a receiver bandwidth of 1 MHz. Add 3/7/10/13/17/20 dB for bandwidths of 2/5/10/20/50/100 MHz, respectively. Subtract 37/20/9 dB for the 0.22/9/120 kHz bandwidth used in CISPR 16 measurements.

the closed surface of a volume that contains the transmitter. The power flux density is, according to (A.32), proportional to the square of the electric field strength, therefore, the electric field strength will at distance d be reduced by a factor

$$a_d = d/\hat{d} \quad (5.10)$$

compared to the value at a reference distance \hat{d} . In other words, increasing the distance to a transmitter in free space by a factor of 10 will reduce the signal strength by 20 dB.

Free-space loss is the dominating form of attenuation when there are no obstacles within a zone known as the *Fresnel ellipsoid*. This volume has the two antennas as its focal points and the path from one antenna to the other via a point on this ellipsoid is a half-wavelength longer than the distance between the antennas. With objects intersecting the Fresnel ellipsoid, diffraction effects have to be considered.

In terms of signal levels at antenna ports, the free-space loss between two isotropic antennas (a hypothetical construction with uniform gain $G = 1$ in all directions) is

$$\begin{aligned} L_{\text{fs}}(d) &= 20 \text{ dB} \cdot \log_{10} \frac{4\pi d}{\lambda} = \\ &= 20 \text{ dB} \cdot \log_{10} \frac{f}{1 \text{ MHz}} + 20 \text{ dB} \cdot \log_{10} \frac{d}{1 \text{ m}} - 28 \text{ dB}. \end{aligned} \quad (5.11)$$

Compared with the available references on out-door radio noise, I found far less clear data in the literature on the in-door radio signal attenuation by building materials. Two

survey publications [102, 103] provide data for the frequency range of 900 MHz to 100 GHz, which is of particular interest to designers of mobile personal communication systems and wireless networking applications. However, this data shows only a few trends and mostly documents a significant variability between different buildings.

The site general model in [102] suggests for 900 MHz a total path loss between two isotropic antennas of

$$L_t(d) = 20 \text{ dB} \cdot \log_{10} \frac{f}{1 \text{ MHz}} + N \cdot \log_{10} \frac{d}{1 \text{ m}} + L_f - 28 \text{ dB} \quad (5.12)$$

where $N = 33 \text{ dB}$ is a loss coefficient that takes into account typical wall spacings and materials for an office environment ($a_d \cdot a_w = (d/\hat{d})^{1.65}$). L_f is an additional floor penetration loss factor with values of 0/9/19/24 dB if the two antennas are 0–3 floors apart in a building. In large open rooms, attenuation is dominated by free-space loss ($N = 20 \text{ dB}$). In corridors it is slightly lower ($N = 18 \text{ dB}$, $a_d \cdot a_w = (d/\hat{d})^{0.9}$), because walls can reflect the signal or act like a wave guide.

Apart from this simple distance power law, the literature survey in [103] lists a number of alternative models that have been used to describe attenuation in buildings. One is an exponent that changes with distance (e.g., $N = 20 \text{ dB}$ for $1 < d < 10 \text{ m}$, $N = 30 \text{ dB}$ for $10 < d < 20 \text{ m}$, $N = 60 \text{ dB}$ for $20 < d < 40 \text{ m}$ and $N = 120 \text{ dB}$ for $d > 40 \text{ m}$). Another approach is to assume exponential attenuation within each wall or floor. Example values from published measurements mentioned in [103] include 1.4 dB for a cloth-covered office partition, 3 dB for wood and brick sliding, 3.8 dB for a double plasterboard wall, 7 dB for a 200 mm concrete block wall, 13 dB for another concrete wall, and 12 and 16 dB for floors in different buildings. Also listed are measurements between various points inside and outside a building, where at 900 MHz, in addition to free-space loss, attenuations of 10–25 dB have been reported. The attenuation has often been observed to decrease slightly with increasing frequency and higher floor number.

For the VHF frequency range, which is of particular interest for video signal eavesdropping, I found only a single study [104], which looked at 35 and 150 MHz signals from a far away station and found that signal levels inside buildings are in the order of 20–25 dB lower than outside in the street and that the building attenuation was in the range 5–45 dB in about 90 % of all measurements made, with a slightly lower attenuation for 150 MHz. The author notes that field strength can vary inside buildings by as much as 20 dB within a meter, which agrees with my own experience from eavesdropping demonstrations with handheld antennas inside office buildings.

5.2.3 Power-line noise and attenuation

Unfortunately, no comprehensive statistics are available so far in the literature for mains power noise and attenuation, especially not for the frequencies above 0.5 MHz that are relevant for broadband eavesdropping of video signals. A power-line model with a small number of parameters is proposed in [105].

An example for a small wide-band noise and attenuation study involving three New York City suburban homes and one New York City apartment is [106]. It found in a 1–60 MHz band the attenuation a_c between two outlets in a house to be on average about 10 dB if the outlets are on the same circuit, and 40 dB if they are on different circuits. Across

the frequency range, the attenuation can vary by as much as 60 dB because the lack of termination in a mains network means that some narrow frequency bands will be affected severely by cancellations from reflections. The root-mean-squared spread of the impulse response of the communication channel between two power outlets has a mean of 0.5 μ s, according to the same study.

The background noise levels shown in the same study indicate that appliances generate mostly noise below 20 MHz, whereas much of the energy in higher frequencies on a power line comes from radio broadcast services. Unfortunately, the noise levels documented are those larger than 95 % of all measurements, as the motivation of the study was to provide data for designers of power-line communications systems, whereas for security emission purposes, a minimal expected noise level would be more interesting, for instance the level below which the noise drops in less than 5 % of all measurements. Example spectrum analyzer plots suggest that with no appliances activated nearby, the power-line noise can be expected to be about 40 dB above thermal noise in the 1–20 MHz range and 25–30 dB in the 20–60 MHz range.

5.2.4 Antenna gain

The compact broadband antennas that are commonly used for EMC measurements, such as biconical, log-periodic, log-spiral, or double-ridged horn designs, have only little directional gain G_a , typically about 2–6 dBi (“dBi” refers to a decibel gain compared to an isotropic antenna).

One of the most practical families of high-gain antennas for the UHF and higher VHF frequency range is the Yagi-Uda type, some forms of which are well known through their widespread use for domestic terrestrial UHF TV reception. A folded dipole is preceded by a series of slightly shorter parallel director elements in the direction of maximum gain, and followed by one or more slightly longer reflector elements. Such an antenna is $\lambda/2$ wide and can be designed such that its gain is

$$G_a = 7.8 \text{ dB} \cdot \log_{10} \frac{l}{\lambda} + 11.3 \text{ dBi}, \quad (5.13)$$

where l is the length of the antenna [74, p. 458]. Increased gain and length of an antenna comes with reduced bandwidth and for the frequencies (200–400 MHz) and bandwidths (50 MHz) best suited for video signal eavesdropping, Yagi antennas with four elements seem to be an acceptable compromise, with a gain of 8.6 dBi and a length $l \approx \lambda/2$. Further gain can be achieved by connecting a group of Yagi antennas together, and each doubling of their number will in practice increase the directional gain by 2.5–2.8 dB. In such a stack of Yagi antennas, the minimum distance between the dipoles has to be about $\lambda/(2 \sin \alpha/2)$ [74, p. 478], where $\alpha = 60^\circ$ (100°) is the 6-dB-width of the directional characteristic of the four-element dipole in horizontal (vertical) stacking direction [74, p. 450].

As a practical example, a 2×3 group of six Yagi antennas with four elements each tuned for the 350 MHz ($\lambda = 0.86$ m) center frequency at which a Toshiba 440CDX laptop provides the best compromising video signal would be $\lambda/2 = 0.43$ m long, $1.5 \cdot \lambda/(2 \cdot \sin 30^\circ) = 1.3$ m wide and $2 \cdot \lambda/(2 \cdot \sin 50^\circ) = 1.1$ m high. Such a structure could be hidden and handled quite easily behind a window or inside a small van with non-conducting upper body and

it would provide a gain of about $G_a = 16$ dBi. Doubling the reception frequency roughly quadruples the number of antenna elements that can fit into the same space, leading to 5–6 dB more gain.

5.2.5 Processing gain

Averaging is a practical and highly effective technique for increasing the signal-to-noise ratio of a periodic signal, such as that generated by the image-refresh circuitry in a video display system or a rapidly repeating software-controlled broadcast signal.

When X_i are independent random variables, then the variance of their sum will be the sum of their variances: $\text{Var}(\sum_i X_i) = \sum_i \text{Var}(X_i)$. The variance of a DC-free signal voltage is just the average power of this signal, multiplied with the load impedance. If we add two sine waves with a random phase relationship together, the expected power of the result is the sum of the powers of each input signal. However, if we add two sine waves of identical frequency and phase together, their voltages will add up.

Similarly, adding two recorded segments of independent noise together will double the power of the noise and increase its root-mean-square voltage by a factor of $\sqrt{2} = 3$ dB. On the other hand, adding two phase-aligned repetitions of the same waveform together will increase its voltage by a factor of $2 = 6$ dB (and will therefore quadruple its power). When two recorded signals contain both independent noise and a wanted phase-aligned signal, then adding the two together will increase the signal-to-noise ratio by 3 dB. The subsequent division by the number of added signals that completes the average calculation will not affect the SNR.

This can be generalized to a processing gain of

$$G_p = \sqrt{N} = 3 \text{ dB} \cdot \log_2 N = 10 \text{ dB} \cdot \log_{10} N \quad (5.14)$$

when N repetitions of a signal can be observed and added up with correct phase alignment.

How many frames of a video signal can be averaged in practice depends on a number of factors:

- When the screen content is stable for a time period T , then obviously up to $f_v T$ frames can be received, and T can range in practice, depending on user behavior, from a few seconds to many minutes or longer, limiting N to about 10^2 – 10^6 .
- Periodic averaging of a video signal can only be successful if the refresh frequency f_v can be determined with a relative error of less than $[2x_t y_t (N - 1)]^{-1}$, if we demand that pixel intervals in the first and last averaged frame overlap by at least half a pixel time t_p . The frequency of crystal oscillators used in graphics cards wanders out of such a tight frequency tolerance within a small number of seconds, limiting N to less than 10^3 for averaging based on a manually adjusted vertical sync signal generated in an independent oscillator.
- An alternative to reconstructing the sync signal and averaging in real-time is to record the receiver output and then search for peaks in the auto-correlation of this signal. This way, the precise repetition frequency can be determined in a more compute-intensive post-processing step (as shown in Fig. 4.3). The number of frames

will here be limited by the available acquisition memory. Storage oscilloscopes offer today 16–64 MB, which limits N to 10^1 – 10^2 frames. The available memory and processing power can be expected to grow further with Moore’s law. They are limited in purpose-built hardware only by the eavesdropper’s budget. With enough signal processing power available, this auto-correlation – which will only be evaluated for peaks near the expected frame time f_v^{-1} – could even be performed in real-time, leading not only to unlimited integration time but also to real-time monitoring of the result.

A practical periodic-averaging device for real-time monitoring of eavesdropped video signals just needs to have the memory to hold all samples from one single frame repetition. It can then implement an infinite-impulse-response (IIR) filter for each sample slot. Such pixel filters are used to reduce noise in some analog TV sets with digital frame buffer. If s_0 is a newly acquired sample value, and b the one stored in the corresponding buffer slot, then the digital signal processor will execute the operation

$$b := \kappa \cdot b + (1 - \kappa) \cdot s_0 \quad (5.15)$$

to mix the new sample with the accumulated weighted average of the previous ones, where $0 \leq \kappa < 1$ is a filter parameter. If s_1, s_2, \dots denote the previous samples for buffer slot b , then

$$b = (1 - \kappa) \sum_{i=0}^{\infty} s_i \kappa^i. \quad (5.16)$$

Assuming all s_i have a common variance σ_s^2 , then the variance of the buffer content will be

$$\sigma_b^2 = \sigma_s^2 \cdot (1 - \kappa)^2 \sum_{i=0}^{\infty} \kappa^{2i} = \sigma_s^2 \cdot \frac{(1 - \kappa)^2}{1 - \kappa^2} = \sigma_s^2 \cdot \frac{1 - \kappa}{1 + \kappa} \quad (5.17)$$

and therefore the processing gain will be

$$G_p = \sqrt{\frac{1 + \kappa}{1 - \kappa}} \quad (5.18)$$

with such an IIR filter buffer.

5.3 Suggested emission limits

As the periodic signals of video displays are the most easily demonstrated emission security threat of personal computers, I will focus on suitable limits for a possible test standard for this type of signal.

Figure 5.2 illustrates that for the reconstruction of readable text, a signal-to-noise ratio of at least 10 dB seems necessary. This requirement could perhaps – especially with much larger fonts – be reduced by a few dB when a symbol detector is used to automatically recognize characters, but the additional gain achievable here depends a lot on the font and will be significantly below the square-root of the number of pixels per character, as many characters such as *i, l, I, 1* or *o, c, e* differ only in a few pixels. A SNR of not more than 0 dB seems therefore to be a reasonable security requirement.




Figure 5.2: This figure visualizes what different signal-to-noise ratios look like for a text video signal.

A further consideration is the receiver bandwidth. We can expect an eavesdropper to work with a bandwidth B somewhere near the pixel frequency f_p of the video mode, as this is necessary to separate the impulses received from individual pixels and reconstruct the full video bandwidth. For larger fonts, text might remain intelligible with somewhat lower bandwidths.

In general, a higher bandwidth will lead to an improved signal-to-noise ratio, because the frequency components of an impulse are correlated and therefore received impulse voltages will grow linearly with B (20 dB for every ten-fold increase in bandwidth), whereas thermal noise and narrow-band background signals are usually not correlated and their voltage will therefore only grow proportional to \sqrt{B} . As a result, in the best case, the signal-to-noise ratio can grow proportional to \sqrt{B} . In practice, the ratio will grow somewhat less, because some of the unwanted background noise (for example emissions from other nearby computers) also has the form of broadband impulses. A reason for keeping the bandwidth small in an eavesdropping receiver is that this will make it more likely to find a quiet window in the radio spectrum that is not used continuously by powerful narrow-band transmitters such as radio and TV broadcast stations. Overall, a bandwidth in the order of $B = 50$ MHz would be a typical practical compromise.

Based on all these considerations, we can now bring together possible values for equations (5.1) and (5.2). A calculation like the following example illustrates how a rationale for the emission limits in a compromising emanations test standard for video signals could look like:

- We measure field strength in the laboratory tests at a distance of $\hat{d} = 1$ m.
- We assume our eavesdropper uses a directional Yagi antenna array like the one described in section 5.2.4 with $G_a = 16$ dBi.

- We assume that an attacker will not get closer than $d = 30$ m with this type of antenna, therefore $a_d = 30$ dB. In a quiet rural site, securing the area 30 m around a device should be feasible, whereas in an urban environment, space is typically more confined, but noise levels are also 10 dB higher, providing the same protection against attackers at 10 m distance. If the threat model includes attackers in nearby rooms in the same building ($d = 3$ m), the resulting test limits will have to be lowered 10 dB further.
- We want to ensure protection even for rooms whose attenuation by building materials is located in the lowest decile of the available statistics and therefore use $a_w = 5$ dB.
- We assume the attacker uses a receiver with a noise figure of $f_r = 10$ dB (the value given for the Dynamic Sciences R-1250) and an impulse bandwidth of $B = 50$ MHz.
- We assume that an attacker will in practice have difficulties with aligning the antenna, tuning to a suitable center frequency, and synchronizing to the exact frame rate if there is no visibly usable signal after averaging $N = 32$ frames (or equivalently using an IIR pixel filter with $\kappa = 0.94$) and therefore we assume a possible processing gain of $G_p = 15$ dB.
- From Fig. 5.1 we can see that for the HF and VHF frequency range of 3 MHz to 300 MHz the background noise level $E_{n,1 \text{ MHz}}$ remains above about 10 dB μ V/m, even at a quiet receiver site. Above 200 MHz thermal noise from the antenna itself becomes the limiting factor, increasing with 20 dB per decade. After adjusting the bandwidth, we get $E_{n,50 \text{ MHz}} = 27$ dB μ V/m.
- The reference from Section 5.2.3 suggests that HF/VHF noise on mains lines can be expected to be at least 30 dB above available thermal noise power, which is 10 dB μ V at $B = 50$ MHz. Therefore $U_{n,50 \text{ MHz}} = 40$ dB μ V seems like a plausible value to work with, perhaps 10 dB more for the range below 20 MHz.
- To estimate the attenuation that conducted emanations will encounter on their way to the receiver, we take the value $a_c = 10$ dB mentioned in the reference from Section 5.2.3 for two mains outlets on the same circuit.

When we require $S/N \leq 0$ dB and use all of the just listed parameters in (5.1), we end up with

$$\begin{aligned}
 \hat{E}_{50 \text{ MHz}} &\leq \frac{S/N \cdot a_d \cdot a_w \cdot E_{n,50 \text{ MHz}} \cdot f_r}{G_a \cdot G_p} \\
 &= 0 \text{ dB} + 30 \text{ dB} + 5 \text{ dB} + 27 \text{ dB}\mu\text{V/m} + 10 \text{ dB} - 16 \text{ dBi} - 15 \text{ dB} \\
 &= 41 \text{ dB}\mu\text{V/m} \\
 \hat{U}_{50 \text{ MHz}} &\leq \frac{S/N \cdot a_c \cdot U_{n,50 \text{ MHz}} \cdot f_r}{G_p} \\
 &= 0 \text{ dB} + 10 \text{ dB} + 40 \text{ dB}\mu\text{V/m} + 10 \text{ dB} - 15 \text{ dB} \\
 &= 45 \text{ dB}\mu\text{V}
 \end{aligned}$$

It would be desirable if the suggested limits were verifiable with off-the-shelf EMC measurement equipment, such as a normal spectrum analyzer. Unlike a sophisticated eavesdropper, a spectrum analyzer will not be able to utilize the processing gain offered by

periodic averaging. Therefore, eavesdroppers can still use signals that are not visible on a spectrum analyzer. To compensate for this, we have to bring during spectrum analyzer tests the antenna as close as possible to the equipment under test, that is $\hat{d} = 1$ m, even if that means that we might encounter some near-field effects.

Like the eavesdropper, the test procedure should work at as high a bandwidth as possible to make use of the fact that impulse voltage will grow proportional to B whereas noise voltage grows proportional with \sqrt{B} . A wide-band receiver suitable for video eavesdropping uses extra-high intermediate frequencies in order to provide large bandwidths of 50 MHz and more. Measurements at such bandwidths are not possible with commonly used spectrum analyzers, whose intermediate frequencies only allow a maximum impulse bandwidth of 5 or 1 MHz. The corresponding limits would be 20 or 34 dB lower:

$$\begin{aligned}\hat{E}_{5 \text{ MHz}} &= 21 \text{ dB}\mu\text{V/m} \\ \hat{U}_{5 \text{ MHz}} &= 25 \text{ dB}\mu\text{V} \\ \hat{E}_{1 \text{ MHz}} &= 7 \text{ dB}\mu\text{V/m} \\ \hat{U}_{1 \text{ MHz}} &= 11 \text{ dB}\mu\text{V}\end{aligned}$$

The specified limits have to be above the receiver noise floor of a spectrum analyzer to be verifiable. As an example, the HP E4402B spectrum analyzer with additional preamplifier has, according to manufacturer claims [107, p. 235], at a resolution bandwidth of 1 kHz a noise level of -133 dBm. This corresponds to $4 \text{ dB}\mu\text{V}$ at 1 MHz and $11 \text{ dB}\mu\text{V}$ at 5 MHz and leaves a distance to $\hat{U}_{1 \text{ MHz}}$ of 7 dB and to $\hat{U}_{5 \text{ MHz}}$ a more comfortable distance of 14 dB. Therefore, the suggested conducted emanation limits can be verified with the impulse bandwidth offered by normal spectrum analyzers.

Verification of the radiated limits is slightly more problematic. The antenna factor for a typical measurement antenna at frequencies up to 100 MHz (see Fig. 2.3) is not more than 10 dB. Therefore, the noise floor of the above spectrum analyzer corresponds to field strengths of $14 \text{ dB}\mu\text{V/m}$ at 1 MHz and $21 \text{ dB}\mu\text{V/m}$ at 5 MHz bandwidth. This makes spectrum analyzer verification of the electric field strength limits problematic at a bandwidth of 1 MHz, and just about feasible at 5 MHz, at least with passive antennas. Active antennas designed for use in anechoic chambers³ might therefore have to be used instead, which offer a 20 dB lower antenna factor and a sensitivity of $6 \text{ dB}\mu\text{V/m}$ at 1 MHz.

Growing antenna factors make the limit of $\hat{E}_{5 \text{ MHz}} = 21 \text{ dB}\mu\text{V/m}$ also problematic to verify for frequencies above 100 MHz, but the eavesdropper would experience thermal noise as well here at the assumed quiet receiver site. Therefore it seems acceptable to increase the limit proportional to the frequency, starting at 100 MHz, up to at least 1 GHz. For higher frequencies, the eavesdropper can use parabolic antennas with higher gain to overcome the thermal noise.

In order to compare the proposed limits with those defined in CISPR 22 Class B, we have to take into account the different bandwidths and antenna distances. To increase the impulse bandwidth from 120 kHz to 5 MHz, we have to raise the permitted field strength by 32 dB, in order to keep the equivalent spectral density constant. The radiated field limits have to be raised further by 20 dB to convert the measurement distance from 10 m to 1 m.

³e.g., Rohde & Schwarz AM524 low noise active antenna system

This way, we can compare the emission security test limits proposed here with the established EMC emission limits:

- Radiated VHF field strength levels have to be 61 dB lower than those allowed in CISPR 22. This corresponds to a reduction of the maximum feasible eavesdropping distance by a factor of 10^3 . In other words, the CISPR 22 EMC limits do not prevent devices from emitting pulses that could in principle be receivable kilometers away.
- Conducted VHF voltage limits have to be about 57 dB lower on mains cables and 71 dB lower on communications cables.

We can conclude from this that a shielded room with an attenuation of 60 dB across the HF/VHF/UHF frequency range for radiated and 70 dB for conducted common-mode emissions should provide adequate protection if all the equipment operated inside it complies to the CISPR 22 Class B limits, as can be expected from all currently available office equipment. It is worth stressing that this entire analysis concentrates on VHF video emanations. It should be applicable to similar high-speed digital signals such as those from system busses, but it does not take into account any low-frequency magnetic or acoustic emanations from electromechanic equipment, as they might have been a concern, for example, with some 20th century impact printers.

If the protection provided by the shielded room must be effective immediately outside the room at a very quiet site ($d = 0.3$ m), and not only in 30 m distance, then another 40 dB attenuation is required, leading to the 100 dB attenuation defined in the NSA specification for shielded enclosures [13].

The US military EMC standard MIL-STD-461E [92] provides in its requirement R102 for mobile Army and Navy equipment radiated electric limits field limits similar to those suggested here, namely measured at 1 m distance not more than $24 \text{ dB}\mu\text{V}/\text{m}$ from 2 to 100 MHz, and then increasing with 20 dB per decade up to 18 GHz. However, the measurement bandwidth is with 10 kHz up to 30 MHz, 100 kHz in the 30–1000 MHz range and 1 MHz for frequencies above 1 GHz comparable to what CISPR 22 uses, and therefore still 37 dB less sensitive for broadband impulse signals than the limits proposed here.

The different emission-control standards can also be compared via the spectral density of the strongest radiated impulse that they permit (measured at 100 MHz and 1 m distance). It has $68 \text{ dB}\mu\text{V}/(\text{m} \cdot \text{MHz})$ under CISPR 22 Class B, $44 \text{ dB}\mu\text{V}/(\text{m} \cdot \text{MHz})$ with MIL-STD-461E/R102, and $7 \text{ dB}\mu\text{V}/(\text{m} \cdot \text{MHz})$ for the emission-security limit proposed here.

Dropping the logarithmic scales and the dependence on a measurement distance, we can also compare the radiated impulse emission limits in terms of the peak effective isotropic radiated power permitted within a given bandwidth. For a 50 MHz wide band, this would be about 0.5 mW under CISPR 22 Class B, 2 μW under MIL-STD-461E/R102, and 0.3 nW under the limit proposed here. For comparison, recall that the peak EIRPs observed during clearly readable eavesdropping demonstrations at this bandwidth in the previous sections were in the range 10–240 nW. The 10 dB stricter limit to protect even against an eavesdropper in a neighbor room in an urban environment would be 30 pW.

The general measurement procedure and setup (use of a wood table over a ground plane, arrangement of cables and impedance stabilization networks, etc.) in an emission security standard could be adopted from existing EMC specifications such as CISPR 22 or MIL-STD-461E. Some changes that would have to be made include:

- CISPR 16-1 suggests that the ambient noise levels at a test site should be 6 dB below the measurement limits, for perfect results even 20 dB below. While CISPR 22 describes the use of an open area test site, this will hardly be feasible with the emission security limits proposed here. A well-shielded anechoic chamber will be required instead as a measurement site, to remain at least 6 dB below the measurement limits.
- CISPR 22 allows 14 dB higher conducted emissions on communications cables than on mains cables. This makes sense in an EMC standard, as radios are connected directly to mains but usually not to communications cables. In an emission security standard, both mains and communications cable must be assumed to be equally accessible to an eavesdropper, therefore no distinction in conducted voltage limits should be made here.
- The CISPR 22 limits on common-mode signals on communications cables are measured while data is being transmitted on these cables, as the intentional signals on these cables are the cause of the interference risk concern. An eavesdropper who has physical access to a communications cable can see the exchanged data anyway and the emission security concern is more about other signals that couple into the communications cable. It should therefore be necessary in an emission security test to transmit data over a communications cable only during radiated tests, but not during the conducted test on that cable.
- While CISPR 22 limits are for quasi-peak and averaging detectors, emission security tests should be performed with peak detectors, because rather than perceived annoyance, the separability from noise is the concern. The video bandwidth should not be limited in spectrum analyzer measurements.

The HF/VHF/UHF emission limits suggested here are justified based on video signal eavesdropping, but they are also likely to provide adequate protection against other forms of compromising emanations above about 5 MHz. It seems unlikely that other emanations will offer substantially higher processing gain than video signals, except perhaps carefully encoded intentional broadcasts by malicious software. In addition, the suggested limits are already very close to what seems technically feasible in the form of generic limits on spectral energy as it can be measured with EMC broadband antennas and spectrum analyzers in anechoic chambers.

An example of another type of eavesdroppable signal would be the laser-diode current in a laser printer. The printed page is usually not repeated, at least not with a repetition frequency that can be predicted by the eavesdropper. Nevertheless, the high vertical resolution of modern laser printers (0.02–0.1 mm) has the effect that successive lines are nearly enough identical, and averaging groups of $N = 16$ lines from a laser-printer signal together will still leave text readable while improving the signal-to-noise ratio.

In some types of devices, comparatively powerful emissions can come from sources that are unlikely to carry any compromising emanations. In the case of a CRT, this would be deflection coils and their driver circuits (at least as long as malicious software has no access to the choice of video timing parameters, otherwise they could be frequency modulated slightly to deliberately broadcast information covertly). The deflection signals can also help an eavesdropper to synchronize on a signal and although they carry no information themselves they could be of value in utilizing the video signal more easily. In the case of

a laser printer, this could be the motors that rotate the drum and scan mirror or feed the paper. Again, their signal carries no information directly, but it could theoretically be of use to the eavesdropper as it might announce, for example, that another page is about to be printed. When such concerns about the information leaked have been addressed, it is possible for an emission security standard to allow device designers the implementation of special emission-security test modes. In these modes, units such as motors and deflection coils are switched off by the controlling software, while all components that handle red signals (CRT e-beam, laser diode, etc.) and their entire signal path are operated with a worst-case test pattern. This way, the emission security measurement can apply the limits on the compromising signal sources only, without requiring unnecessarily expensive shielding for non-compromising sources.

Different measurement techniques can be used in order to apply the same limits to specific signals that are suspected of being emitted. In the case of a periodic signal such as the output of a video refresh circuit, it is possible to use a wide-band receiver, just as an eavesdropper would, together with a good storage oscilloscope triggered from the graphic adapter's vertical sync signal. Averaging $N = 1024$ frames (about 12 s of video signal) will lead to a processing gain of 30 dB. If all 1024 lines of the test image are identical and averaged as well, this will lead to another 30 dB gain, making the measurement in principle better than one that took place inside a shielded room with 60 dB attenuation. An RF function generator can be used as a reference source to determine whether the averaged signal is above the specified emission limits.

In the case of non-periodic digital signals, such as mainboard bus lines, cross-correlation measurements as outlined in Section 2.4 provide an alternative. In this case, attention has to be paid to ensure that the receiver bandwidth is larger than the signal frequency, such that the impulse shapes generated by a signal transition in the receiver's demodulator input do not overlap and interfere with each other. Signal-specific measurement techniques like this do not, of course, protect against eavesdroppers who find another signal, for example, from the memory bus of the graphics card.

A fully developed emission-security standard that permits to analyze individual emission sources separately might also decide to apply different limits to different classes of signals. Parallel interfaces, for instance, leak less information than serial interfaces, and periodic signals or those involving a large amount of redundancy can be detected and processed far more easily than compact and transient signals.

Rarely are the design and rationale of emission limit standards based solely on a purely scientific reasoning process. Compromises have to be made in practice between different protection levels, threat models, as well as the cost of conformance and the cost and repeatability of verification. The limits proposed in this chapter will hopefully serve as a starting-point for such a discussion.

Chapter 6

Optical eavesdropping of displays

The available open literature on compromising emanations from video displays has so far only considered the threat of information carried in the radio frequency bands (primarily 3 MHz–3 GHz). We must not forget, however, that the very purpose for which video display devices are designed is the emission of information suitable for human perception in the optical bands (790–385 THz frequency or 380–780 nm wavelength). In this section, I will demonstrate that optical high-frequency emissions of cathode-ray tube (CRT) computer monitors are a feasible new eavesdropping channel that deserves serious consideration, even if the eavesdropper has no direct line of sight to the target display surface.

6.1 Projective observation with telescopes

It has of course not escaped the attention of security experts in the past that any video display surface that is within a line of sight to an eavesdropper's hiding place could be read with the help of a telescope. Many organizations dealing with critical information have security policies concerning the orientation and visibility of documents, whiteboards, computer monitors, and keyboards relative to windows that connect to uncontrolled spaces such as nearby streets, parking lots, buildings, or tall structures.

With high-quality optics, the limiting factor for the angular resolution of a telescope is the diffraction at its aperture. For an aperture (diameter of the first lens or mirror) D , the achievable angular resolution as defined by the Rayleigh criterion is

$$\theta = \frac{1.22 \cdot \lambda}{D}, \quad (6.1)$$

where $\lambda \approx 500$ nm is the wavelength of light. Typical modern office computer displays have a pixel size $r = 0.25$ mm, for example in the form of the 320×240 mm display area on a 43 cm CRT, divided into 1280×1024 pixels. If the observer is located at distance d and her viewing direction differs by an angle α from a perpendicular view onto the display surface, she will see a single pixel under a viewing angle $\theta = \frac{r}{d} \cdot \cos \alpha$. She will therefore need a telescope with an aperture of at least

$$D = \frac{1.22 \cdot \lambda \cdot d}{r \cdot \cos \alpha}. \quad (6.2)$$

Therefore, a simple amateur astronomy telescope ($D = 300$ mm) will be sufficient for reading high-resolution computer display content for up to 60 m distance and $\alpha < 60^\circ$, even with very small font sizes.

6.2 Time-domain observation of diffuse CRT light

Surprisingly though, the direct projection of a video display surface into the image plane of a camera with a good telescope, as discussed in section 6.1, is not the only way in which optical emanations of cathode-ray tubes can be used to read the screen content at a distance. In the rest of this chapter, I will demonstrate a new observation technique that does not depend on a direct line of sight to the display surface. It uses the fact that the high-frequency amplitude variations of the light emitted by a CRT depend on the video signal. This allows an eavesdropper to reconstruct the screen content even without knowing from which point on the display surface the received photons originated.

The overall light from a CRT is a weighted average of the luminosity of the last few thousand pixels that the electron beam addressed. More precisely, the intensity $I(t)$ of the light emitted is equivalent to the (gamma corrected¹) video signal $v_\gamma(t)$ convolved with the impulse response $P(t)$ of the screen phosphor:

$$I(t) = \int_0^\infty v_\gamma(t - t') P(t') dt'. \quad (6.3)$$

So even if an observer can pick up only the average luminosity of a CRT surface, for example, by observing with a telescope the diffuse light reflected from nearby walls, furniture or similar objects, a low-pass filtered version of the video signal becomes accessible. Not even curtains, blinds or windows with etched or frosted glass surfaces – as are frequently used to block views into rooms – are necessarily an effective protection, as the average luminosity inside a room can still leak out.

As with radio frequency eavesdropping, we utilize the fact that displayed pixels are updated sequentially, and again the periodic nature of the process allows an attacker to reduce noise and address individual display units out of several in a room via periodic averaging.

The light emitted by a cathode-ray tube is generated when the electron beam hits a luminescent substance, called the *phosphor* (not to be confused with the chemical element phosphorous). The measurements described in the next section show that when the electron beam hits the phosphor of a bright pixel, the emitted light intensity increases faster than a single pixel period, and even though the overall afterglow of the phosphor lasts typically more than a thousand pixel times, a noticeable drop of light output happens within a few pixel times. As we will see, this preserves enough high-frequency content of the video signal in the emitted light to allow for the reconstruction of readable text.

¹The intensity of the light emitted by the phosphor is, up to a saturation limit, proportional to the electron beam current $i(t)$, which is typically related to the video-signal voltage $v(t)$ by a power-law relationship $i(t) \sim (v(t) - v_0)^\gamma$. The “gamma corrected” video signal $v_\gamma(t) \sim i(t)$ used here is strictly speaking not the actual video voltage supplied to a monitor via the video cable. It is a hypothetical voltage that is proportional to the beam current and $v_\gamma(t) = 1$ V represents the maximum intensity. This way, we can quantify the phosphor impulse response of a monitor without having to measure the beam current.

6.3 Characterization of phosphor decay times

The exact shape of the decay curve of the phosphors used in the CRT is the most important factor for the image quality that the eavesdropper can obtain. It is of interest for at least three reasons:

- It determines the frequency characteristic of the phosphor, which shows how much high-frequency content of the video signal will be attenuated before appearing in the emitted light.
- It determines the initial luminosity during the first pixel time, which is a characteristic parameter for estimating how strong the received signal will be against the shot noise due to background light.
- It is needed as a parameter for the deconvolution operation that the eavesdropper can use to reconstruct the original image.

Every bright pixel of a CRT surface is hit by an electron beam of typically up to 100 μA for time $t_p = f_p^{-1}$, and this refresh is repeated once each time interval f_v^{-1} , where f_p and f_v are the pixel-clock and vertical-deflection frequency, respectively. The beam electrons push other electrons in the phosphor material to higher energy levels. As they fall back into their original position, they emit stored energy in the form of photons. The time delay in this process causes an afterglow for several milliseconds after the electron beam has passed by.

The VGA CRT monitor that I used in the measurements described in this chapter is a “Dell D1025HE”. Its phosphor type is identified in the user manual simply as “P22”. This is an old and obsolete designation, referring to an entry in an early version of the *Electronic Industries Alliance (EIA)* phosphor type registry. It merely describes the entire class of phosphors designed for color TV applications. The more modern *Worldwide Type Designation System (WTDS)* for CRTs [112] calls the old P22 family of phosphors “XX” instead and distinguishes subclasses. The most recent EIA TEP116-C phosphor type registry [113] lists seven different color TV RGB phosphor-type triples designated XXA (P22 sulfide/silicate/phosphate), XXB (P22 all-sulfide), XXC (P22 sulfide/vanadate), XXD (P22 sulfide/oxy-sulfide), XXE (P22 sulfide/oxide), XXF (P22 sulfide/oxide modified) and XXG. In addition, it contains partial information on composition, emission spectrum, decay curves and color coordinates for at least 15 further RGB phosphor type triplets designated XBA, XCA, etc. that were developed for data-display applications and that differ somewhat from the TV standards in their color. Unfortunately, the original manufacturer of the tested monitor was not able to answer my question on which exact P22 variant was used.

CRT screen phosphors are usually based on the sulfides of zinc and cadmium or rare-earth oxysulfides and are activated by additions of dopant elements to determine the color. Most EIA registered XX and X phosphor type triplets use for the red phosphor yttrium oxysulfide doped with europium ($\text{Y}_2\text{O}_2\text{S}:\text{Eu}$, phosphor types “RJ”/“RH”), often blended with zinc phosphate doped with manganese ($\text{Zn}_3(\text{PO}_4)_2:\text{Mn}$, “RE”). The green phosphor is often zinc sulfide doped with copper ($\text{ZnS}:\text{Cu}$, “GF”) and sometimes also with aluminium and/or gold, or zinc silicate doped with manganese and silver ($\text{Zn}_2\text{SiO}_4:\text{Mn,Ag}$,

“GR”). The blue phosphor is usually zinc sulfide doped with silver (ZnS:Ag, “BM”) and in some cases also aluminium or gallium.

Like many physical decay processes (e.g., radio activity), the luminosity of a typical excited phosphorescent substance follows an exponential law of the form

$$I_e(t) = I_0 \cdot e^{-\frac{t}{\tau}} \quad (6.4)$$

where I_0 is the initial luminosity right after the excitation ceases and the time constant τ is the time in which the luminosity drops by a factor e ($= 2.718$). Such decays can be identified easily in a plot of the logarithm of the luminosity over time as a straight line. For

$$\tau = \frac{1}{2\pi f} \quad (6.5)$$

the above exponential decay is also the impulse response of a first-order Butterworth low-pass filter consisting of a single resistor and capacitor, with a 3-dB cut-off frequency f . As the phosphor decay can be seen as a low-pass filter applied to the video signal before we can receive it with a photosensor, describing the decay in terms of the cut-off frequency is perhaps more illustrative than the time constant.

Zinc-sulfide based phosphors show instead a power-law decay curve of the form

$$I_p(t) = \frac{I_0}{(t + \alpha)^\beta} \quad (6.6)$$

Such a decay behavior can be identified on a plot of the logarithm of the luminosity versus the logarithm of the time since excitation has ceased as an asymptotically straight line that flattens somewhat near $t = 0$. The condition $\beta > 1$ must be fulfilled, otherwise the integral

$$\int_0^\infty \frac{1}{(t + \alpha)^\beta} dt = \frac{\alpha^{1-\beta}}{\beta - 1} \quad (6.7)$$

which is proportional to the total number of photons emitted would not be positive and finite.

Since commonly used phosphors are mixtures of various substances and different excitation modes occur (resulting in various wavelengths), actual decay curves have to be modeled as the sum of several exponential and power-law curves.

The TEP116-C standard does contain decay curves for most phosphor types, but these curves are plotted on a linear time-scale extending over many milliseconds. These curves give no indication about the detailed decay during the first microsecond and they are therefore not suitable for estimating the frequency characteristic of the phosphors above 1 MHz. The decay curves published in TEP116-C were measured primarily to provide information about how the phosphor type might affect the perceived flicker caused by the frame refresh. Suitable fast decay curves, or even closed form approximations, were not available in the existing CRT phosphor literature. So I decided to perform my own measurements on a typical example monitor.

6.3.1 Instrumentation

We are primarily interested in the rapid decay within a time interval not much longer than t_p , therefore we need a very sensitive light sensor with ideally more than 100 MHz bandwidth or less than 5 ns rise and fall time.



Figure 6.1: The Hamamatsu H6780-01 photosensor module used for these experiments contains a photomultiplier tube together with a high-voltage power supply.

One fast light sensor type is the PIN photodiode in photoconductive mode, in which a reverse bias voltage is applied and the resulting current is measured. The PIN photodiode has an undoped “intrinsic” layer between the p- and n-doped regions (hence the name PIN). Compared to normal photodiodes, PIN diodes have reduced capacity and can be used with a higher bias voltage, which increases their response time. For example, a PIN diode with a “rise and fall time of about 20 μs ” was used in [114] to evaluate the luminance decay of the P31 phosphor in a CRT used in vision research.

Ultra-highspeed photodiodes are now available with down to 1 ns response time for applications such as optical Gbit/s communication links and laser range finding. However, their low sensitivity, typically 0.5 A/W, makes significant additional amplification necessary, which would lead to additional noise and further limit the bandwidth. Avalanche photodiodes (APDs) provide greater sensitivity (10^2 A/W) and are also available with 1 ns response times.

Photomultiplier tubes (PMTs) are evacuated electron tubes with a photocathode. Received photons can trigger the emission of electrons, which are then accelerated with high voltage and multiplied in a cascade of further electrodes. A single received photon results in an entire cloud of electrons hitting the anode, contributing to the measured current. Photomultiplier tubes have response times in the nanosecond range and their sensitivity can be easily adjusted over many orders of magnitude.

I decided to use the Hamamatsu H6780-01 photosensor module (Fig. 6.1), which can be operated with radiant sensitivity levels in the 10^1 – 10^5 A/W range [115]. It can therefore be used under quite a wide range of light conditions. This device contains in a small robust metal package a photomultiplier tube and a high-voltage generating circuit and can thus be operated conveniently from a 12 V lab power supply. A separately applied 0.25–0.90 V control voltage U_c adjusts the radiant sensitivity of the sensor to

$$1.5 \times 10^5 \text{ A/W} \cdot \left(\frac{U_c}{1 \text{ V}} \right)^{7.2}.$$

The radiant sensitivity is the quotient of the output current generated by the sensor and the radiant energy received on its 8 mm diameter (50 mm^2) aperture. When operated within the specified parameters, a photomultiplier is a highly linear light-controlled current source. To prevent damage to the sensor, attention has to be paid that the maximum allowed average output current of $100 \text{ }\mu\text{A}$ is not exceeded, by selecting the control voltage appropriately ([109] even recommends keeping the current below $1 \text{ }\mu\text{A}$). In my experiments, I monitored the output current with an oscilloscope. I adjusted the control voltage manually with a $10 \text{ k}\Omega$ precision potentiometer that is part of a voltage divider and monitored its value with a digital voltmeter.

According to the data sheet, the anode-current rise time of the H6780 photomultiplier module is 0.78 ns , an order of magnitude faster than the pixel time t_p of commonly used video modes. Since a photomultiplier just pumps electrons through its anode output, the fall time is determined by the RC low-pass filter formed by the combined capacitance C of the anode and the cable as well as the input impedance R of the measurement amplifier. A higher termination resistance R will lead to a higher input voltage at the amplifier, but also to a reduced fall time for the signal. Since the photomultiplier operates with up to a kilovolt acceleration voltage, it is able to generate output voltages that could damage an amplifier with too high a DC input impedance, which is another reason for choosing a low value for R . I therefore decided to terminate the $50 \text{ }\Omega$ coaxial cable that leaves the H6780 module with a $50 \text{ }\Omega$ resistor, such that reflections are eliminated and the 100 pF/m capacitance of the cable does not affect the rise and fall times of the signal significantly. The Tektronix TDS 7054 storage oscilloscope that I used for data acquisition has, in its most sensitive setting, a voltage resolution of 1 mV/div . This corresponds to an 8-bit A/D converter resolution of $40 \text{ }\mu\text{V}$ and so we get a current resolution of $0.8 \text{ }\mu\text{A}$ over the $50 \text{ }\Omega$ resistor. I therefore decided to connect the photomultiplier module directly to the oscilloscope input without adding an additional video amplifier, even though that meant that I had to operate the sensor with a peak output current not too far below the maximum average current $100 \text{ }\mu\text{A}$ to make good use of the A/D converter resolution. For comparison, the thermal noise through an $R = 50 \text{ }\Omega$ resistor at $T = 300 \text{ K}$ and $B = 100 \text{ MHz}$ bandwidth is $\sqrt{4kTRB} = 9 \text{ }\mu\text{V}$ and therefore still well below the quantization noise.

6.3.2 Measurement method

In order to characterize the phosphor response times of the Dell D1025HE, I used seven test images. The first three have just one single illuminated pixel in either full intensity red, green, or blue, the other four showed a 320-pixel-long line, in either full intensity red, green, blue, or white. The use of both a single pixel and an entire illuminated line allowed me to characterize both very fast (tens of nanoseconds) and slower (millisecond) impulse features in the impulse response of the phosphor, which would have been difficult with a single measurement given the limited dynamic range of the recording. The signal timing used in this test was the VESA $640 \times 480 @ 85 \text{ Hz}$ video mode, in which the electron beam traverses the 320 mm wide screen at 18 km/s . I set the monitor controls during the measurement to 100 % contrast, 50 % brightness and a color temperature of 6500 K and ensured that it had been powered up for at least 30 min.

According to [109, p. 253], the decay curves of zinc-sulfide based phosphors can vary significantly under different drive conditions and [111] specifies that the characterization of CRT phosphor decay times has to be performed with a standard beam current of

100 μA . However, this adjustment would require the connection of a microammeter directly to the high-voltage anode connection of the CRT, a measurement for which I lacked the necessary equipment. Therefore I decided to simply use the default brightness setting and the full intensity color combinations that are most frequently used for text display. The resulting luminosity measurement is therefore with respect to a known video signal voltage only, and not with respect to a known beam current.


I positioned the photosensor 0.25 m in front of the center of the screen surface, facing the screen in normal viewing direction parallel to its normal vector. This means that the sensor aperture of 50 mm^2 covered a solid angle of about 0.8 msr. I triggered the oscilloscope connected to the photosensor from the vertical sync signal on pin 12 of the feature connector of the driving VGA card and reduced the noise of the signal by averaging the signal over 256 frames with 16-bit arithmetic in the oscilloscope. Each frame was acquired with 8-bit resolution at a sampling rate of 5 GHz for the single-pixel signal and 125 MHz for the 320-pixel line. For each measurement, I adjusted and recorded the photomultiplier control voltage individually to maximize the dynamic range.

6.3.3 Results

Figure 6.2 shows the results of these measurements (continuous lines), both the entire recording as well as a zoom into the start region. In these diagrams, I converted the voltages recorded by the oscilloscope into the radiant intensity from the CRT, based on the known termination resistor, control voltage, radiant sensitivity and solid angle of the sensor, in order to eliminate the dependency of the resulting measurement on the specific details of the sensor setup. The radiant sensitivity used is the one given in the data sheet for 420 nm (blue) and can vary for up to a factor of two for other wavelengths. Because of this, and since I had no calibration source for radiant intensity, the absolute values in these diagrams should only be seen as estimates with (in particular for red and green) up to a factor two uncertainty. In the case of the 320-pixel line measurement, I was able to verify the relative accuracy of the radiant sensitivity deduced from the control voltage to be in the order of 1 % by comparing the sum of the red, green and blue signals with the white signal.

For further theoretical analysis, as well as for optimizing the processing of eavesdropped photonic signals for best readability, it is helpful to have a simple closed-form approximation of the phosphor impulse response. I attempted to use various linear system identification algorithms, such as the Steiglitz-McBride iteration [116], to find suitable IIR filter coefficients. I also tried manually selecting several first-order Butterworth filter cut-off frequencies, which were then combined by least-squares approximation to form a model candidate. None of these techniques led to satisfactory results, in particular not for the blue, green and white signals, whose long-term decay – as I realized only later – is dominated by the power-law decay typical for zinc-sulfide compounds. This cannot be approximated well as the sum of a small number of exponential decays with different time constants.

I eventually matched the parameters manually. I adjusted the parameters, until I found a close fit between the measured curves and the sum of the smallest possible number of exponential and power-law decay functions, which I compared on all the linear, logarithmic and double-logarithmic scales shown in Figs. 6.2, 6.3, and 6.4. I ended up with the following model functions for the impulse response of the three phosphors, which after

(a) Emission decay of a single pixel ($f_p = 36$ MHz)

(b) Emission decay of a 320-pixel line





Figure 6.2: This figure shows the measured emission decay of P22 CRT phosphors (continuous lines). The video signal (magenta line near bottom) is shown here delayed by 29 ns in order to compensate for delays in both electron tubes and the signal cables. The dashed lines representing the modeled decay follow the measured values so closely that they are hardly visible without linear or logarithmic (Figs. 6.3 and 6.4) magnification.

convolution with the 1 V video signal stimulus (delayed by 29 ns to compensate for delays in both electron tubes and the signal cables) lead to the excellently matching dashed lines in the above mentioned figures:

$$\begin{aligned}
 P_{P22R}(t) / \frac{\text{W}}{\text{V} \cdot \text{s} \cdot \text{sr}} &= 4 \times e^{-2\pi t \times 360 \text{ Hz}} + 1.75 \times e^{-2\pi t \times 1.6 \text{ kHz}} + \\
 &2 \times e^{-2\pi t \times 8 \text{ kHz}} + 2.25 \times e^{-2\pi t \times 25 \text{ kHz}} + \\
 &15 \times e^{-2\pi t \times 700 \text{ kHz}} + 29 \times e^{-2\pi t \times 7 \text{ MHz}}
 \end{aligned} \tag{6.8}$$

(a) Emission decay of a single pixel ($f_p = 36$ MHz)


(b) Emission decay of a 320-pixel line




Figure 6.3: This figure shows the curves from Fig. 6.2 on a logarithmic scale, to demonstrate the accuracy of the model formula over several orders of magnitude.

$$\begin{aligned}
 \frac{P_{P22G}(t)}{V \cdot s \cdot sr} &= 210 \times 10^{-6} \times \left(\frac{t + 5.5 \mu s}{1 s} \right)^{-1.1} + 37 \times e^{-2\pi t \times 150 \text{ kHz}} + \\
 &100 \times e^{-2\pi t \times 700 \text{ kHz}} + 90 \times e^{-2\pi t \times 5 \text{ MHz}}
 \end{aligned} \quad (6.9)$$

$$\begin{aligned}
 \frac{P_{P22B}(t)}{V \cdot s \cdot sr} &= 190 \times 10^{-6} \times \left(\frac{t + 5 \mu s}{1 s} \right)^{-1.11} + 75 \times e^{-2\pi t \times 100 \text{ kHz}} + \\
 &1000 \times e^{-2\pi t \times 1.1 \text{ MHz}} + 1100 \times e^{-2\pi t \times 4 \text{ MHz}}
 \end{aligned} \quad (6.10)$$

$$P_{P22} = P_{P22R} + P_{P22G} + P_{P22B} \quad (6.11)$$




Figure 6.4: On a double-logarithmic plot of the measured phosphor decay – same data as in Fig. 6.2 (b) – the power-law decay of the green and blue P22 phosphors shows up as a straight line after about 10^{-4} s.




Figure 6.5: The numeric model of the measured P22 phosphor impulse response is shown here, plotted on logarithmic time and intensity scales. The clearest signal can be expected from the blue phosphor, not only because it has the largest initial power, but also because its power drops by the largest amount in the first 100 ns.




Figure 6.6: This graph shows the same decay functions as Fig. 6.5, but on a normalized linear intensity scale. The dashed lines are the respective normalized integrals of the decay functions, which show at which time which fraction of the overall number of photons has already been emitted.

Using this closed form representation, we can now have a better look at various attributes of the phosphor decay. Figure 6.6 shows as continuous lines the impulse response curves on a logarithmic time scale. Their amplitudes have been normalized to $P(0) = 1$ in this representation, in order to make the curves better comparable, independent of the sensitivity of the photosensor to the various wavelengths. The dashed lines represent the integrals of the decay functions. They show which fraction of the totally emitted energy after stimulation ceased has already been given off at any point in time.

The red phosphor emits practically all of its stored energy within 1–2 ms, but it still has not lost a significant part of its energy after the first 10 μ s. The red phosphor also shows the fastest relative initial decay during the first 0.1 μ s, after which the blue phosphor catches up and remains at a lower relative luminosity until after about 2 ms the red phosphor has exhausted its energy entirely. Recall that the red phosphor decays purely exponentially.

The blue and green phosphors show a far more heavy-tailed behavior, thanks to the power-law component in their decay function. Even long after the stimulus, they still have not emitted all of their stored energy and, as a result, even an unaided human observer with fully adapted scotopic vision can notice an afterglow on a CRT screen in an otherwise completely dark room for several minutes. The integral of $P_{P22G}(t)$ shows that, in theory, even hours after the excitation some energy remains unreleased in this phosphor type. Although this measurement was not designed to estimate the here significant parameter β in equation (6.6) with good accuracy, this observation still leads to the interesting question




Figure 6.7: This figure shows the frequency characteristic of the three measured P22 phosphors and their combination to white. The modeled impulse response was truncated with a rectangular window to 3 ms and then a 2^{19} point FFT applied. The 0 Hz value has been normalized to equal 1 for all four curves.

whether CRT phosphors could leak confidential information not only via instantaneous compromising emanation, but also via data remanence, an information security risk that has so far only been studied in memory devices [53, 55]. Placed in a dark and possibly cooled chamber, a sensitive camera with long exposure time might detect the afterglow from recently displayed text even some time after the CRT was switched off, limited only by thermal noise.

All three phosphors show a very noticeable relative drop of radiant intensity in the first tenth of a microsecond. Figure 6.5 shows that of all three phosphors, the blue one has the by far largest drop in absolute radiant intensity in the first 100 ns ($-1500 \text{ W}/(\text{V} \cdot \text{s} \cdot \text{sr})$) and therefore will provide the strongest high-frequency signal, while the red phosphor has the smallest absolute drop ($-34 \text{ W}/(\text{V} \cdot \text{s} \cdot \text{sr})$). The Fourier transforms of the impulse response curves in Fig. 6.7 show that the blue phosphor applies to the video signal a low-pass filter in which, for example, a 10 MHz component is less than 40 dB more attenuated than a 1 kHz signal. Only for frequencies above about 5–10 MHz, the phosphors show the continuous 20 dB per decade roll-off typical for a first order low-pass filter.

6.4 Optical eavesdropping demonstration

Perhaps more interesting than a theoretical discussion of phosphor decay frequency characteristics is a visually convincing actual reconstruction of a displayed image from an

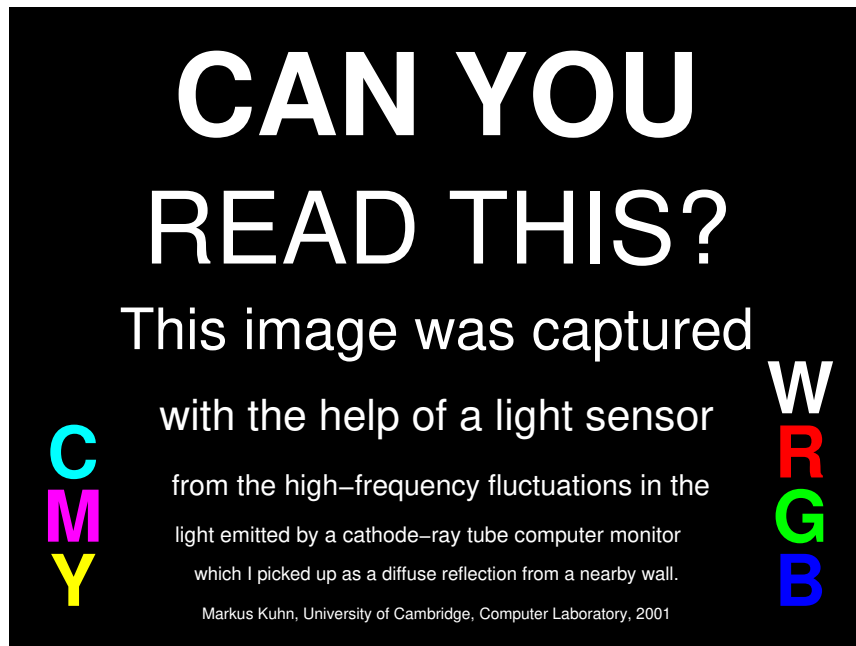


Figure 6.8: Testchart displayed on the target monitor (VESA $640 \times 480@85\text{Hz}$ video mode)




Figure 6.9: Unprocessed photomultiplier output after diffuse reflection from a wall

out-of-sight CRT surface.

In the following experiment, the same monitor (Dell D1025HE) faces a white office wall at about a meter distance. The photomultiplier is now located behind the monitor, facing the same wall at about 1.5 m distance. There is no direct line of sight between the sensor and the screen surface.

As the wall illuminated by the monitor covers a large solid angle, as seen from the photo-sensor, no additional optical elements, such as a focusing lens, are needed in this demonstration. The experiment was performed at night, with the room lights switched off;

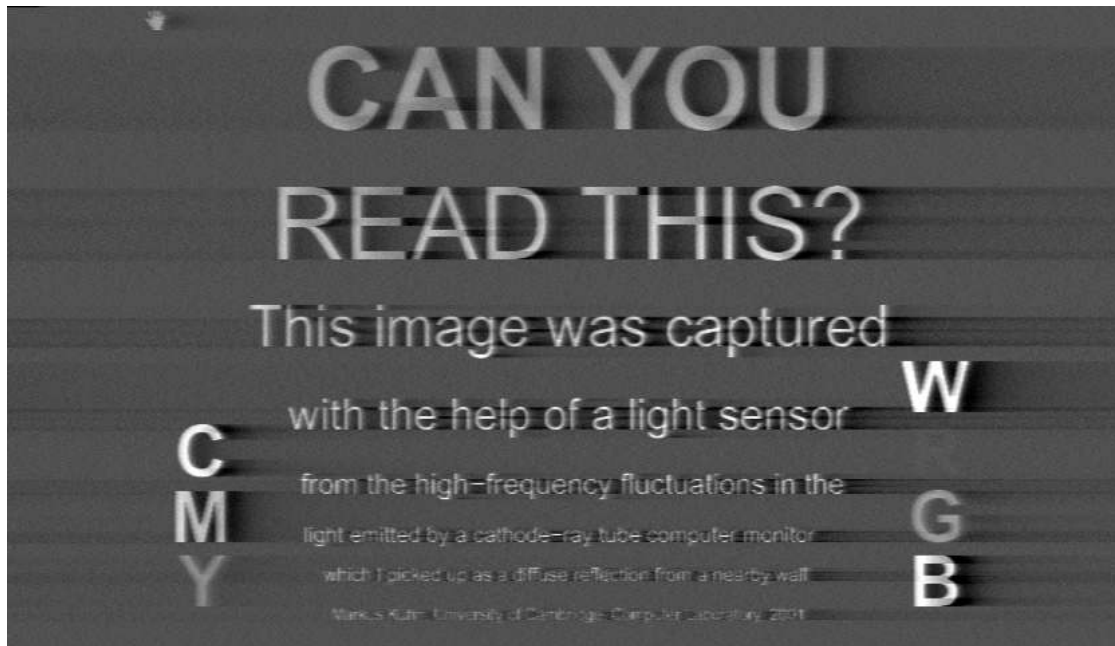


Figure 6.10: Signal from Fig. 6.9 with a 4 MHz Butterworth high-pass filter applied.

however, the room was far from completely dark, being illuminated by several computer displays and stray light from outside lamps.

Figure 6.8 shows a simple readability test chart with text in various sizes, which the monitor displayed as the target video signal during the test. The test text was displayed in full white on black, and additional test letters show all full intensity combinations of the three phosphor colors. I used the same video mode (VESA $640 \times 480@85\text{Hz}$) as in the previous experiment, and again I triggered the oscilloscope from a vertical sync signal provided by the graphics card and averaged over 256 frames (or equivalently 3 s), at a sampling frequency of 250 MHz. Quite acceptable readability of small text can also be achieved with lower sampling rates and numbers of averaged frames, as well as higher video modes.

Figure 6.9 shows the rasterized representation of the recorded and averaged photocurrent. The largest font sizes are readable, though the slow decay smears the luminosity of each white pixel along the path of the electron beam across the rest of the line and further. The gray values of all the rastered signals shown here were adjusted such that the values of the 0.1 % highest and 0.1 % lowest pixels in a histogram are all mapped to full white or full dark respectively, and the remaining values are linearly mapped to the corresponding points on a gray scale.

The raw photomultiplier current shown in Fig. 6.9 clearly has to be processed further in order to make text readable. Analog preprocessing has the advantage that it can significantly improve the signal-to-noise ratio (SNR) before any amplification and quantization steps limit the dynamic range of the signal to, for example, 8 bits or 48 dB in the case of the oscilloscope being used. With a digitized signal of too low quality, little further digital recovery is feasible, as the necessary high-pass filter step amplifies high-frequency noise further.

Figure 6.10 shows the digital simulation of a simple first-order Butterworth high-pass filter with a cut-off frequency of 4 MHz applied to the signal, which could be implemented quite




Figure 6.11: A much better image quality can be achieved by applying a matched filter with a frequency characteristic that is inverse to that of white shown in Fig. 6.7.

easily as just a resistor-capacitor combination. It leads to a dramatic improvement in text readability, though it still shows quite noticeable image distortions. These are due to the fact that this simple filter applies a 20 dB per decade roll-off from 4 to 0 MHz, whereas the frequency characteristic of the phosphors (Fig 6.7) is actually significantly flatter below 4 MHz.

A much better reconstruction can be obtained by deconvolution, that is with the help of a filter that has approximately the inverse phase and frequency characteristic of the phosphor. In order to generate the image of $\tilde{v}(t)$ in Fig. 6.11, I first sampled the model impulse response function $P_{P22}(t)$ with the same sampling frequency and number of samples as the recorded averaged luminosity signal $I(t)$ for a single frame. I then Fourier transformed both, divided the complex results and applied the inverse Fourier transform:

$$\tilde{v} = \mathcal{F}^{-1} \left\{ \frac{\mathcal{F}\{I\}}{\mathcal{F}\{P_{P22}\}} \right\} \quad (6.12)$$

No padding was necessary before performing the Fourier transform, since $I(t)$ is a periodic signal anyway, $P_{P22}(t)$ already drops close to zero near the beginning of the frame period, and the FFTW code [117] used to perform the calculation can also handle block sizes other than 2^n for the discrete Fourier transform quite efficiently.

The result of the deconvolution shows a significantly improved contrast, the smear along the electron beam path to the right of each illuminated pixel is reduced, and even the smallest font size of the test chart (with an H-height of 8 pixels or 4 mm) becomes readable. The restoration is slightly better for blue text than for green (and consequently white) text, and the high-frequency components of the red signal remain too weak for this sensor setup. Using a closed-form model for P_{P22} significantly adds to the numeric stability of the deconvolution and avoids division-by-zero problems.

6.5 Threat analysis

With the help of a phosphor decay curve like the one shown in Fig. 6.5, we can now estimate the signal strength that an eavesdropper can receive and what upper bound for the reception distance shot-noise imposes. For definitions of the radiometric and photometric quantities and SI units used here, see [128, 109].²

For the following order-of-magnitude estimates, we assume in the interest of simplicity that the screen, wall, and sensor surfaces involved are roughly parallel to each other and that the photons of interest travel perpendicular to these, otherwise the cosine of the relevant angles would have to be multiplied in as well. We also assume that the quantum efficiency of the photosensor – the probability that a photon passing through the aperture is registered – equals 1. Smaller values of real sensors can be taken into account by adjusting the sensor aperture A_r .

6.5.1 Direct observation

We first consider the case without diffuse reflection from a wall, where the eavesdropper can see the screen surface directly. This might allow projective observation with a telescope, but the result might not be satisfactory in situations with minor distortions such as aperture diffraction, atmospheric fluctuations or even a frosted glass window. Time-domain analysis of the received light could be of interest even where a line of sight is available.

Let $t_p = f_p^{-1}$ be the duration for which the electron beam illuminates a single pixel. The video voltage due to one pixel (full intensity: $V = 1$ V) will be

$$v_\gamma(t) = \begin{cases} V & \text{if } 0 < t \leq t_p \\ 0 & \text{otherwise} \end{cases} \quad (6.13)$$

and the resulting radiant intensity according to (6.3) is

$$I_p(t) = V \cdot \int_{t-t_p}^t P(t') dt'. \quad (6.14)$$

At distance d with receiver aperture area A_r , neglecting transmission delays and the directional characteristic of the emitter, the power received from the pixel is

$$P_p(t) = \frac{A_r}{d^2} \cdot I_p(t). \quad (6.15)$$

We approximate the detection process performed in the receiver by simply integrating the received pixel power over the pixel duration. The resulting energy collected per pixel is

$$Q_p = n \cdot \int_0^{t_p} P_p(t) dt \quad (6.16)$$

²In a nutshell: Luminous flux is measured in lumen (lm), which is the photometric equivalent of radiation power, weighted by the spectral sensitivity of the human eye, where 683 lm are per definition as bright as 1 W of (green) 540 THz light. In order-of-magnitude calculations, I will simply approximate 10^3 lm as 1 W. The steradian (sr) measures a solid angle (4π for the full sphere), candela (cd) is the same as lumen per steradian and measures the luminous intensity of a light source in a given direction, and lux (lx) is the same as lumen per square meter and measures the illuminance of a location. Commonly encountered illuminance levels cover ten orders of magnitude, from 10^5 lx for “direct sunlight” to 10^{-4} lx for “overcast night sky (no moon or light pollution)” [109, p. 16].

where n is the number of frame repetitions accumulated by periodic averaging. This is only a small fraction of the overall energy received from the pixel during its decay, but it approximates the amount of energy that can be separated from the contributions of neighbor pixels by high-pass filtering. At wavelength λ this energy corresponds to

$$N_p = \frac{Q_p \lambda}{hc} \quad (6.17)$$

photons per pixel ($hc = 1.986 \times 10^{-25}$ Jm).

We also have to consider background light as a noise source, both from other pixels of the observed CRT as well as any surrounding surfaces. The photon count per pixel duration from the background light can be estimated as

$$N_b = \frac{nt_p AA_r L_b \lambda}{hcd^2}, \quad (6.18)$$

where L_b is the average radiance and A is the area of the observed background surface.

The arrival of photons at a detector aperture is a Poisson process [118]. This means that when a random variable N describes the number of photons received per pixel and we expect $\mathcal{E}[N]$ photons on average then the standard deviation $\sqrt{\mathcal{E}[(N - \mathcal{E}[N])^2]}$ will be $\sqrt{\mathcal{E}[N]}$. This inevitable variability of the photon count is known as *shot noise*. As $N_b \gg N_p$, the background light determines the amount of shot noise against which the status of a single pixel has to be detected. This roughly becomes feasible when the signal-to-noise ratio is greater than one, that is

$$N_p > \sqrt{N_b} \quad (6.19)$$

or with $P(t) \approx P(0)$ for $0 \leq t \leq t_p$

$$\frac{nt_p^2 A_r V P(0) \lambda}{2hcd^2} > \sqrt{\frac{nt_p AA_r L_b \lambda}{hcd^2}}. \quad (6.20)$$

and therefore

$$\frac{A_r}{d^2} > \frac{4AhcL_b}{n\lambda V^2 t_p^3 P^2(0)}. \quad (6.21)$$

We can now fill this condition with some example parameters. Assuming a background luminance of 100 cd/m², as it is typical for a CRT and other bright surfaces in a well-lit office environment [109, 110], the corresponding background radiance will be in the order of not more than $L_b = 0.1$ W/(sr · m²), from which we mask off an observed area of $A = 0.2$ m². Together with other typical parameters such as $t_p = 20$ ns, $P(0) = 10^3$ W/(V · s · sr), $V = 1$ V, $\lambda = 500$ nm, and by averaging $n = 100$ frames, we get

$$\frac{A_r}{d^2} > 4 \times 10^{-5} \text{ sr}. \quad (6.22)$$

For example, a simple telescope with $A_r = 0.3$ m² could therefore theoretically receive a signal under these well-lit conditions up to 80 m away.

6.5.2 Indirect observation

We now consider an indirect observation in a dark environment, where the not directly visible CRT screen faces at distance d' a diffusely reflecting observable wall, which has a reflection factor $0 < \varrho < 1$. The radiant intensity (power per solid angle) $I_p(t)$ from a pixel will lead to an irradiance (incoming power per area)

$$E_p(t) = \frac{I_p(t)}{d'^2} \quad (6.23)$$

onto the wall and to a radiant exitance (outgoing power per area) of

$$M_p(t) = \varrho E_p(t). \quad (6.24)$$

For a uniformly diffusing (“Lambertian”) surface, we have to divide the radiant exitance by π [109] to obtain the corresponding radiance (power per solid angle per area)

$$L_p(t) = \frac{1}{\pi} M_p(t) \quad (6.25)$$

which leads us finally to the power

$$P_p(t) = \frac{AA_r}{d^2} \cdot L_p(t) \quad (6.26)$$

passing through the receiver aperture A_r , which is located at distance d from the observed wall area A . Using the same $P(t) \approx P(0)$ for $0 \leq t \leq t_p$ approximation as before, we can estimate the number

$$N_p = \frac{\varrho n t_p^2 A A_r V P(0) \lambda}{2\pi h c d^2 d'^2} \quad (6.27)$$

of photons received from a single pixel and compare it to the number

$$N_b = \frac{n t_p A A_r \varrho E_b \lambda}{\pi h c d^2} \quad (6.28)$$

of photons received from the background light, assuming the wall is exposed to an irradiance E_b . The signal to shot-noise ratio will again be of order unity under the condition $N_p > \sqrt{N_b}$, which leads to a receivability condition

$$\frac{A_r}{d^2} > \frac{4\pi E_b h c d'^4}{\varrho n \lambda t_p^3 A V^2 P^2(0)}. \quad (6.29)$$

Let’s again look at an example scenario. Assuming the observed monitor has a luminous intensity of $100 \text{ cd/m}^2 \times 240 \text{ mm} \times 320 \text{ mm} = 8 \text{ cd}$, a wall at a distance $d' = 2 \text{ m}$ would be exposed to an illuminance of in the order of 2 lx from the overall light given off by the monitor alone, which corresponds to the illuminance during “late twilight” [109] and is equivalent to an irradiance of in the order of $E_b = 1 \text{ mW/m}^2$. Using this with the same example parameters as before, as well as $A = 2 \text{ m}^2$ and $\varrho = 0.5$, we get

$$\frac{A_r}{d^2} > 1 \times 10^{-4} \text{ sr} \quad (6.30)$$

for this indirect observation under late twilight conditions. The $A_r = 0.3 \text{ m}^2$ mirror used as an example before could therefore receive a signal under these conditions up to 50 m away. This distance is proportional to $1/\sqrt{E_b}$, so, for example, under full daylight illuminance (10^4 lx), observation would already be infeasible a meter from the wall.

6.5.3 Observation of LEDs

It is worth noting that the very high pixel frequencies used by CRTs play a significant rôle in limiting the reception range. Optical displays with lower update frequencies could pose an eavesdropping risk, even if they do not offer the redundancy of a repetitive video signal. A practical example are devices with slow serial ports (10^4 – 10^5 bit/s), such as some modems, that feature light-emitting diode (LED) displays to indicate the logic level of data lines. Unless the displayed signal is distorted, for example, by a monostable multivibrator circuit that enforces a minimum LED-on period of at least a character time, an optical eavesdropper could manage to reconstruct transmitted data by monitoring the LED luminosity at a distance. A recent study found that of 39 tested communication devices, 14 emitted serial port-data in light from transmit/receive line status LEDs [119].

Another example would be software-controllable status LEDs, such as those connected to the keyboard and hard-disk controller of every PC. Malicious software could use these to covertly broadcast information in situations where this cannot be accomplished via normal network connections (e.g., due to “air gap” security or a mandatory access-control operating system).

Normal LEDs have a luminous intensity of 1–10 mcd, although super-bright variants with up to 100 mcd or more are available as well.

We can again estimate the expected number of photons N_p received from a single bit pulse of the LED, as well as the expected number N_b from the background illumination. For a sufficiently large N_b , we can approximate the distribution of the number N of photons received as a normal distribution

$$P\left(\frac{N - \mu}{\sigma} < x\right) \approx \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{y^2}{2}} dy \quad (6.31)$$

with the mean value

$$\mu = \begin{cases} N_b + N_p & \text{when LED on} \\ N_b & \text{when LED off} \end{cases} \quad (6.32)$$

and the standard deviation

$$\sigma = \sqrt{N_b}. \quad (6.33)$$

Assuming that transmitted bits 0 and 1 are equally likely, a matched filter detector [120] will count the photons N received per bit interval and compare the resulting number with the threshold $N_b + \frac{1}{2}N_p$ to decide whether the LED was on or not. The probability for a bit error due to shot noise will therefore be

$$p_{\text{BER}} = Q\left(\frac{N_p}{2\sqrt{N_b}}\right) \quad (6.34)$$

where

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^{\infty} e^{-\frac{y^2}{2}} dy = \frac{1}{2} - \frac{1}{2} \operatorname{erf}\left(\frac{x}{\sqrt{2}}\right) \approx \frac{e^{-\frac{x^2}{2}}}{x\sqrt{2\pi}} \quad \text{if } x > 3 \quad (6.35)$$

is the Gaussian error integral [120].

As a practical example, we consider a direct line of sight to a green ($\lambda = 565$ nm) LED with a luminous intensity of 7 mcd, which corresponds to a radiant intensity of roughly

$I_p = 10^{-5}$ W/sr. A telescope at distance d with aperture A_r will receive, over a single bit pulse time t_p , an expected number of photons

$$N_p = \frac{t_p A_r I_p \lambda}{h c d^2} \quad (6.36)$$

from the LED, plus an expected number of photons

$$N_b = \frac{t_p A A_r \rho E_b \lambda}{\pi h c d^2} \quad (6.37)$$

if the observed area A has a reflection factor ρ and is exposed to an ambient irradiance E_b . With example parameters $A_r = 0.3$ m², $d = 500$ m, $t_p = 10^{-5}$ s (100 kbit/s), $\rho = 1$, $A = 1$ cm² = 10^{-4} m² and $E_b = 1$ W/m² (roughly 10^3 lx, “overcast sky”), we end up with a lower bound for the bit error rate of 10^{-7} .

Finally an example where the same LED illuminates a wall at distance d' , of which the eavesdropper observes area A and collects from a single bit pulse an expected photon count

$$N_p = \frac{t_p A A_r \rho I_p \lambda}{\pi h c d^2 d'^2}, \quad (6.38)$$

whereas the photon count from the background illumination remains as in (6.37). Inserting example values of $A_r = 0.3$ m², $d = 50$ m, $t_p = 10^{-4}$ s (10 kbit/s), $\rho = 0.5$, $d' = 2$ m, $A = 2$ m² and $E_b = 1$ mW/m² (roughly 1 lx, “late twilight”), we end up with a lower bound for the bit error rate near 10^{-4} .

Figure 6.12 illustrates a possible detection and clock recovery algorithm for NRZ encoded binary data (as it appears on RS-232 serial port lines), which recovers the sampling clock signal if only the bitrate is known (or guessed correctly).

6.6 Receiver design considerations

The demonstration in Section 6.4 shows the image quality that an eavesdropper can achieve in principle under favorable conditions by using simple off-the-shelf instruments. It is just intended as a proof-of-concept lab experiment for the diffuse optical CRT eavesdropping risk and does not exploit a number of techniques for improving range and signal quality that could be used in purpose-built portable optical eavesdropping receivers.

The most important improvement is the use of a zoom telescope to capture more photons and provide for the exact selection of a target area with good signal-to-noise ratio. The image quality of the telescope needs to be only good enough to allow for the masking of an area of interest, usually with centimeter to decimeter resolution. This avoids the need for heavy high-precision mirrors like those used for astronomic imaging.

The ultimate performance limit is the amount of background light and associated shot noise that reaches the photosensor and an important design concern will be its reduction with the help of careful masking and filtering.

The data provided for “X” and “XX” screen phosphors in [113] as well as Fig. 6.13 show that the zinc-sulfide based blue and green phosphors have a bell-shaped spectral energy distribution centered mostly at 450 and 520 nm, respectively, with a standard deviation of roughly 20–30 nm. The red phosphors, on the other hand, typically have




Figure 6.12: This chart illustrates an algorithm for clock and data signal recovery from a NRZ binary signal (a) with added white Gaussian noise (b). We first convolve the received signal (b) with the pulse shape of a single bit (here a rectangular pulse) and obtain (c). Function (d) shows the distance of (c) from its mean, which we convolve with an impulse series with the same period as the bit length to get (e). The result has maxima at the edges of the original signal, which provides us with the clock signal for sampling (c). The sampled values (f) are then thresholded (g) and we have recovered the original bitstream (a) out of (b), knowing only its bit rate, but not its clock phase.

a spectrum consisting of several much narrower lines, usually between 600 and 700 nm, with a standard deviation of less than 5 nm. Color filters or a spectrograph can be used to separate the contributions from different phosphors to reconstruct color images and to provide for the application of phosphor-specific deconvolution parameters. Careful selection of filter frequencies can also be used to attenuate background light. While both the sun and incandescent lights have a relatively smooth spectrum in the optical band, this is not the case with some types of fluorescent lights commonly used in offices. As Fig. 6.13 also shows, fluorescent lamps emit strong mercury emission lines (405 nm, 436 nm, 546 nm, 615 nm), which can be attenuated with suitable filters or a spectrograph relative to CRT emissions, to improve the signal-to-noise ratio for an optical eavesdropper.

The phosphor decay curves shown in this chapter were measured with a sensor that is sensitive over the entire optical band. It might be worth investigating whether narrow-band sensors that record only a single spectral line observe a different decay curve. If this is the case, spectral bands with particularly low high-frequency attenuation could be selected by an optical eavesdropping receiver to improve the signal quality further, although a tradeoff will have to be made between optical bandwidth and shot noise.

If background light is generated directly from a 50 or 60 Hz power supply, it will be modulated with twice that frequency; fluorescent lights far more so than incandescent ones. Where the observed signal is repetitive, varying the receiver gain to be inversely proportional to the background light amplitude can further improve the SNR. A gated photomultiplier module could be used for that purpose.

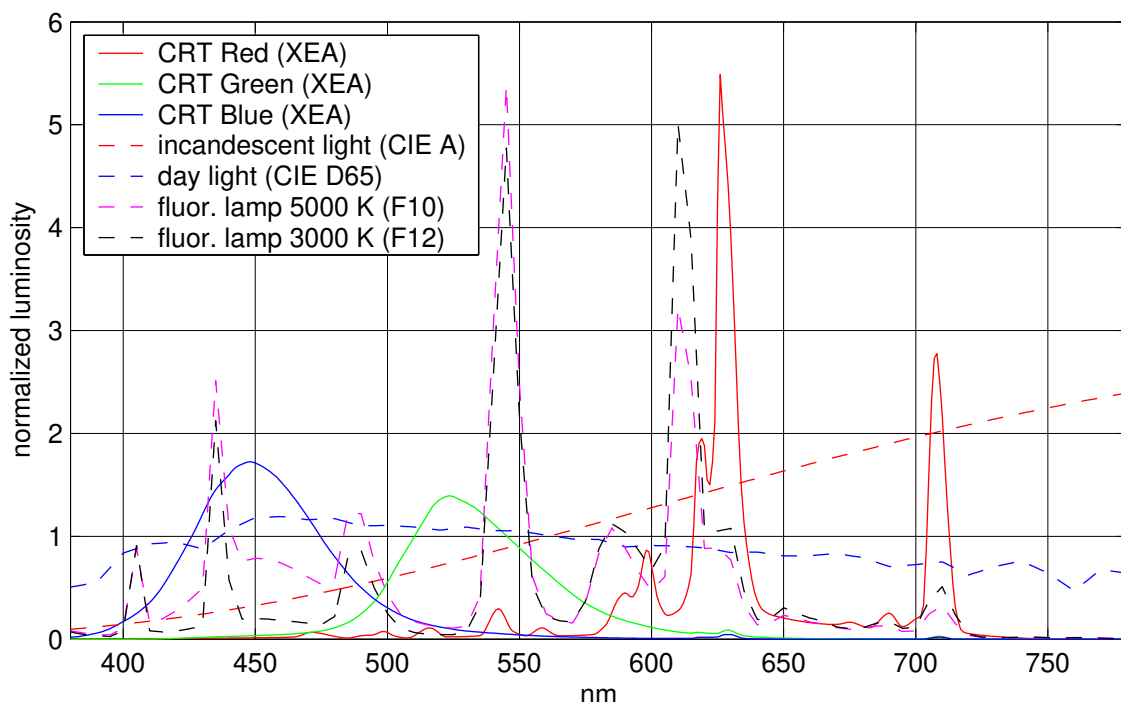


Figure 6.13: These spectra illustrate intensity distributions of emitted wavelengths for a typical CRT phosphor mixture (“XEA”, data from [113]), common types of tri-phosphor fluorescent office lights (data from [109, pp. 305–308]), as well as standard reference spectra for incandescent household lightbulbs (CIE illuminant A) and daylight (CIE illuminant D65) [121]. The illumination curves have been normalized for equal luminosity and the three CRT curves shown add up to an equally luminous white.

Analog preprocessing at the output of the photosensor could better approximate the optimal deconvolution filter than just the simple RC element simulated in the previous section. Digital processing would then only have to take care of any remaining inaccuracies of the analog stage.

6.7 Countermeasures

Once the nature of a new eavesdropping technique is understood, it is possible to suggest a range of countermeasures that, when combined and implemented properly, can significantly reduce the described risk.

Display surfaces as well as keyboards used for handling critical information should, of course, first of all be kept out of any line of sight to a potential eavesdropper. In addition, diffusely reflected stray light from cathode-ray tube displays should be treated as a potentially compromising emanation, especially when there is low background illumination and eavesdroppers can install large-aperture equipment within a few hundred meters. Rooms where a significant amount of the ambient light comes from displayed sensitive information should be shielded appropriately, for example, by avoiding windows.

Various measures for jamming diffuse optical emissions with good background illumination can be used. Background light should preferably be of a broadband nature (solar or high-temperature incandescent), or, in the case of fluorescent lights, be produced with

phosphors that have an emission spectrum similar to those used in CRTs. Many modern fluorescent lights are operated with a high-frequency current (≥ 20 kHz). These are preferable, as they have significantly reduced dark phases. In addition, individual lamps will not be phase synchronized.

Some monitors already include an ambient light sensor. These adjust the contrast and brightness automatically to the surrounding light conditions. It would be easy to extend this mechanism, such that a power-saving mode is activated when the ambient light levels fall below a secure jamming margin. Such a mechanism has not only security but also ecologic and ergonomic advantages. Less electric power would be wasted in dark and empty offices over night, if darkness acted as an additional power-saving trigger, and eye strain for users might be reduced, by discouraging work under bad background illumination.

The red phosphor in this demonstration experiment showed a significantly lower initial brightness and better high-frequency attenuation than the green and blue ones. Its emission lines are also in a part of the spectrum that is covered well by household light bulbs. In order to facilitate the selection of suitable CRT phosphors for information security applications, it would be helpful, if display tube and phosphor manufacturers, as well as phosphor-type registries, provided impulse-response information. This could be in the form of double-logarithmic diagrams, such as Fig. 6.5, that cover a time scale of 10^{-9} – 10^{-2} s, and perhaps even a closed-form approximation, along with a plot of the frequency-domain filter characteristic. An example for a suitable characteristic parameter of interest in the design of a security CRT might be the relative attenuation provided by a phosphor for beam currents with 100 Hz and 10 MHz frequency. It would also be helpful, if monitor manufacturers documented, which exact CRT and phosphor types they use, as well as which beam currents.

The need for special security CRTs is likely to be reduced significantly with the further proliferation of liquid crystal displays (LCDs). Their pixels react considerably slower than CRT phosphors and most types of flat-panel displays refresh all pixels in a line simultaneously. Both these factors suggest that this technology has a significantly reduced risk of leaking information about individual pixels in diffuse optical emanations.

Chapter 7

Review, outlook and conclusions

The major new effects demonstrated and insights gained in this investigation of compromising emanations from video displays are:

- The information displayed on a cathode-ray tube computer monitor can be reconstructed, in a readable form, by an eavesdropper from its distorted or even diffusely reflected light. This can be demonstrated using easily available components, such as a photomultiplier tube and a fast analog-to-digital converter, followed by numerical deconvolution with the help of a closed-form model of the screen phosphor's impulse response.
- Due to shot-noise limits, the eavesdropping from diffuse reflections of display light (both CRT and LED) is only applicable in relatively dark environments. Even then, it is limited to less than a few tens or hundreds of meters distance, which may still be of practical concern in some situations. Better eavesdropping distances, even under office-light conditions, become possible with a direct line of sight, which might include minor distortions, such as frosted glass, that would otherwise be deemed sufficient to frustrate projective observation.
- Very much like radio-frequency eavesdropping of video displays, the practical exploitation of compromising optical emanations will usually require specially designed equipment, expertise, and patience. However, it seems at least as powerful as the former, and organizations who have traditionally worried about compromising radio emanations should seriously consider this new set of eavesdropping techniques in their threat models.
- In contrast to conventional wisdom, flat-panel displays can pose a larger emission-security risk than CRTs. The Gbit/s digital interfaces used to connect them to the video controller use serial transmission formats that effectively modulate the video signal in ways which provide an eavesdropper with better reception quality at larger distances.
- A detailed understanding of the encoding algorithms and bit arrangement used in digital video links allows programmers fine-grained control over the emitted signal. This can be used for both attack and defense.

- 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.
- Modified glyph-rendering routines can significantly reduce the text readability of compromising emanations from CRT monitors. To achieve this, font anti-aliasing algorithms have to be modified to minimize the contrast of vertical edges. This can be accomplished, for example, by subsampling with a suitably chosen FIR filter.
- Combinations of foreground and background colors can be selected to reduce the readability of text in the compromising emanations of digital video interfaces. 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.
- The analog conditioning of the video signal on a video card, between the digital-to-analog converter and the VGA connector, significantly affects the compromising emanations of a CRT. Products differ widely in the amplitude of the harmonic frequencies they generate. The use of well-designed low-pass filters can significantly reduce the probability of a successful signal intercept.
- The steganographic emission of information in the video signal can be demonstrated by gamma-corrected amplitude modulation of an invisible dither pattern. The emitted signal can be designed such that it is readable even with TV-bandwidth receivers.
- The rigid timing of video signals makes it possible to use, with monospaced fonts, a very simple and effective radio character recognition technique for the automatic transcription of eavesdropped text.
- The principles of direct-sequence spread-spectrum modulation can be applied to reliably broadcast data and synchronization information via digital-video-link emanations.
- For a typical video-eavesdropping scenario, an effective emission-security standard for devices or shielded rooms should limit the spectral density of impulses to a level about a thousand times or 60 dB lower than the civilian radio-interference limit. This is still a factor 70 below the stricter US military RFI protection standard. None of the publicly available electromagnetic emanation test standards defines today a meaningful test to evaluate the emission security of a device.
- And finally, I provided brief introductions to instruments, measurement procedures and concepts that are of use for the investigation and quantitative description of compromising emanations. This includes broadband antennas, wide-range wide-band receivers, digital signal processing methods and techniques for the visualization of reconstructed video signals, along with the notions of impulse strength and spectral density.

Compromising emanations remain an area with many open questions. It is my impression that the field provides plenty of challenges and intellectual stimulation for researchers with an interest in hardware security, high-frequency technology, signal processing, and unconventional communication channels.

The following list highlights some ideas and minor observations that might warrant further investigation or could lead to future projects:

- The shot-noise calculations presented for optical emanations of CRTs involve a crude approximation of the signal-to-noise ratio available to an eavesdropper after deconvolution. A better estimate should be possible by taking a displayed image, convolving it using the phosphor impulse-response model presented, and using the shot-noise formulas to add a realistic amount of random noise to the result, in order to generate a simulated photocurrent, as it would be recorded in a given eavesdropping scenario. This could then undergo deconvolution, and the result should show far more accurately the readability that can be achieved for the reconstructed video signal.
- I have not yet performed any experiments with compromising emanations from mains power connections, therefore the limits for conducted emissions that I suggested should be seen as a very preliminary proposal that will need to be backed up with experimental data and practical experience.
- An AM receiver is not designed to reconstruct unmodulated digital signals accurately. It destroys the phase and polarity information that could help to distinguish rising and falling edges in digital signals. The non-linear detector it uses prevents the application of linear digital-signal-processing techniques, such as channel equalization, that could be used to reconstruct the originally transmitted waveform more accurately. Experiments could be performed with a quadrature-amplitude-modulation detector, which provides as output a two-dimensional vector signal. Such signals would be particularly easy to average and analyze, if it were possible to phase-lock the local oscillator to the clock frequency of the target device. If this is not feasible in the receiver hardware, equivalent transforms could be applied in software to the digitized IF output of the receiver before averaging.
- My eavesdropping demonstrations involved a time-consuming GPIB data-transfer from a storage oscilloscope into a PC, where the data was further processed, analyzed and formatted for display with MATLAB programs. A high-performance digital signal processing system could perform many of these steps far more conveniently in real time. Suitable hardware might soon be available for experiments with software radio architectures that digitize the filtered IF mixer output and perform all further demodulation steps in real-time on a normal workstation in software [122].
- I performed some very preliminary experiments towards eavesdropping of a laser printer (Apple LaserWriter II NTX, 300 dpi) and found that its laser diode scans the photosensitive drum with 555.6 lines per second and a pixel frequency of 1.865 MHz. The laser-diode drive signal has 1.2 V high edges with 5 ns raise time, which promises emanations at least up to 60 MHz. It was not difficult to receive AM tones from printed dither patterns in the shortwave bands in the immediate vicinity of the printer. More recent printers with four times the resolution have pixel frequencies

near 30 MHz and might have harmonics up to the lower UHF range, where far-field eavesdropping becomes more feasible.

- Rumors of eavesdropping attacks on electric typewriters by intelligence agencies in the 1970s, in particular on the IBM “golfball” variant, encouraged me to attempt a reconstruction of some early “Tempest” history. I obtained an *IBM Selectric II* typewriter, but found that this was a purely mechanical device, driven by a single continuously running motor. It is therefore not a suitable target for electromagnetic eavesdropping. The rumors might refer instead to golfball computer printers such as the IBM 1050 or IBM 2741, but unfortunately none were available for experimentation.
- In some devices, a digital state or signal can cause minor supply voltage variations. These can then cause a slight change of frequency in an oscillator on the same supply, which may lead to FM broadcast of data. I observed that a Logitech mouse emitted a 4.01490 MHz signal, which drops by 60 Hz whenever the built-in microcontroller wakes up during mouse movement. I also observed that a MAX233 RS-232 voltage converter chip uses a 125 kHz oscillator to operate an on-chip voltage multiplier. This frequency is also emitted as a series of switching pulses via the RS-232 outputs onto the serial line, and it changes by a few hertz with the current logic value.
- The R-1250 receiver that I used is a versatile piece of laboratory equipment, but its mass (40 kg) and price make it unlikely to be used in many real eavesdropping attacks. Receivers suitable for picking up compromising video emanations could be built with far simpler circuitry. Neither the wide frequency range nor the large number of bandwidths of a “Tempest” test receiver is required, nor is the frequency accuracy of a superheterodyne receiver with PLL local oscillator. A portable low-cost eavesdropping system could be built as a simple direct receiver, consisting of an antenna amplifier, several switchable band-pass filters, a rectifier, a low-pass filter and a video amplifier. The output would be connected to a fast digital storage oscilloscope card plugged into a portable PC, where software rasters the received signal in real-time on the display.
- Other routes to low-cost wide-band AM receivers could involve modifications to superheterodyne amateur scanners that have a sufficiently high first intermediate frequency (e.g., ≈ 780 MHz for the IC-R8500), or perhaps the use of commercial wide-band converters designed for satellite downlink applications. A creative eavesdropper is unlikely to be affected by export controls for “Tempest” laboratory receivers.
- Compromising emanations are just one electromagnetic aspect of hardware security. Interference immunity is a related area, including the protection against high-energy electromagnetic waveforms that could be used to destroy information processing equipment [123], interfere with its operation, trigger secondary compromising emanations, or transmit commands to preinstalled malicious software over an unconventional communications channel.
- Information theft through unconventional side-channels is unlikely to remain restricted to the electromagnetic, optical, and acoustic domain. For example, a mechanical side-channel could be exploited by installing pressure sensors underneath

the feet of a keyboard, or the table on which it rests. How the force vector of each keystroke will be split up among these anchor points will depend on the location of the respective key. Two or three attached transducers should suffice to distinguish and record keystroke sequences, an attack that may be of particular concern with trusted PIN-entry devices.

Are compromising emanations a practically relevant information security threat today? In realistic overall threat models of applications that depend on information security, compromising emanations no doubt play a relatively minor rôle at present. The vast majority of practical vulnerabilities can be exploited using comparatively simple and purely software-based techniques. This is likely to remain the case, as long as information security is only a secondary consideration in the design and selection of products, equally neglected by both product designers and end users.

Eavesdropping on unintended hardware emissions usually requires a physical presence close to the target. This can lead to significant cost and risk of discovery for the eavesdropper. Apart from a few historic examples and rumors, I am not aware of a single recent confirmed case of computer criminality or espionage, where there was clear evidence that compromising emanations have been exploited. If there are any cases, they are likely to be performed by persons who go to great lengths to keep their activities away from public disclosure. Compromising emanation attacks are therefore unlikely to become the future focus of the same sort of public attention enjoyed today by mostly untargeted remote vandalism such as Internet or WLAN break-ins and computer virus development.

Nevertheless, compromising emanations are an important field of research. Compared to the large number of minor and highly theoretical vulnerabilities of cryptographic primitives and protocols discussed in much of the current computer security literature, compromising emanations are a risk of practical interest that has so far remained mostly undocumented. Even the most basic concepts of compromising emanations are not even mentioned in textbooks at the moment, except perhaps for the occasional reference to classified military documents on the subject. Entire new classes of vulnerabilities, such as, for example, the diffuse optical emission risk demonstrated here, are still being discovered. Many others, such as RF emanations from devices that are illuminated by an external carrier wave, have been speculated about, but no experimental data has been published so far.

“Tempest-certified” stands for expensive low-volume production devices that undergo strict time-consuming individual emission tests, most likely from 100 Hz to over 1 GHz. Civilian EMC-tests on the other hand are only designed to keep complaints about radio interference in residential areas at a manageable level. They do not prevent severe interference in close proximity or critical environments, such as aircraft. There might be a market for mass-produced consumer electronic devices that conform to a, yet to be developed, new emission-control standard. It would only require type-certification and statistical compliance, so as to be applicable to mass-production devices, and would limit the emitted wide-band spectral density in the VHF/UHF bands to about 60 dB below what the current civilian EMC standard (CISPR 22) allows.

For the customer, products conforming to such a new standard would offer the benefit of safety from radio interference in very close proximity and sensitive environments (“safe for use in aircraft during takeoff and landing”) combined with a reassuring, but not

paranoid, level of emission security. This standard would not require expensive per-device testing, and it would probably also not require any particular shielding against low-frequency magnetic fields and power-line emissions. The latter two are far less likely to carry significant compromising emanations in modern products than they did at the time of low-speed electromechanic devices, when the original “Tempest” specifications were developed.

“My neighbor has had a new heart pacemaker fitted. Every time he makes love my garage doors open.”

— Bob Hope, 1975

Bibliography

- [1] Dan Tyler Moore, Martha Waller: Cloak & Cipher. George G. Harrap & Co. Ltd., London, 1965.
- [2] R.F.H. Nalder: The Royal Corps of Signals – A History of its Antecedents and Development (circa 1800–1955). Royal Signals Institution, London, 1958.
- [3] Arthur O. Bauer: Some aspects of military line communications as deployed by the German armed forces prior to 1945. The History of Military Communications, Proceedings of the Fifth Annual Colloquium, Centre for the History of Defence Electronics, Bournemouth University, 24 September 1999.
- [4] John Young: How Old is Tempest? Online response collection, February 2000. <http://cryptome.org/tempest-old.htm>
- [5] Peter Wright: Spycatcher – The Candid Autobiography of a Senior Intelligence Officer. William Heinemann Australia, 1987, ISBN 0-85561-098-0.
- [6] Electromagnetic Pulse (EMP) and Tempest Protection for Facilities. Engineer Pamphlet EP 1110-3-2, 469 pages, U.S. Army Corps of Engineers, Publications Depot, Hyattsville, December 31, 1990.
- [7] Deborah Russell, G. T. Gangemi Sr.: Computer Security Basics. Chapter 10: TEMPEST, O'Reilly & Associates, 1991, ISBN 0-937175-71-4.
- [8] National Security Telecommunications and Information Systems Security Advisory Memorandum NSTISSAM TEMPEST/1-92: Compromising Emanations Laboratory Test Requirements, Electromagnetics. National Security Agency, Fort George G. Meade, Maryland, 15 December 1992. Partially declassified transcript: <http://cryptome.org/nsa-tempest.htm>
- [9] NACSIM 5000: Tempest Fundamentals. National Security Agency, Fort George G. Meade, Maryland, February 1982. Partially declassified transcript: <http://cryptome.org/nacsim-5000.htm>
- [10] National Security Telecommunications and Information Systems Security Advisory Memorandum NSTISSAM TEMPEST/2-95: RED/BLACK Installation Guidance. National Security Agency, Fort George G. Meade, Maryland, 12 December 1995. Transcript: <http://cryptome.org/tempest-2-95.htm>
- [11] National COMSEC/EMSEC Information Memorandum NACSEM-5112: NONSTOP Evaluation Techniques. National Security Agency, Fort George G. Meade, Maryland, April 1975. Partially declassified transcript: <http://cryptome.org/nacsem-5112.htm>
- [12] National Security Telecommunications and Information Systems Security Instruction NSTISSI No. 7000: TEMPEST Countermeasures for Facilities. National Security Agency, Fort George G. Meade, Maryland, 29 November 1993. Partially declassified transcript: <http://cryptome.org/nstissi-7000.htm>

- [13] Specification NSA No. 94-106: Specification for Shielded Enclosures. National Security Agency, Fort George G. Meade, Maryland, 24 October 1994. Transcript: <http://cryptome.org/nsa-94-106.htm>
- [14] Tempest Glossary. National Communications Security Committee, Subcommittee on Compromising Emanations, NCSC 3. National Security Agency, Fort George G. Meade, Maryland, 30 March 1981. Partially declassified transcript: <http://cryptome.org/ncsc-3.htm>
- [15] A. J. Mauriello: Join a government program to unveil Tempest-spec mysteries. EDN, Vol. 28, No. 13, pp. 191–195, June 23, 1983.
- [16] 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.
- [17] Bloßstellende Abstrahlung [Compromising Emanations]. Faltblätter des BSI 12, German Information Security Agency, Bonn, 1996. <http://www.bsi.de/literat/faltbl/012-blab.htm>
- [18] Überkoppeln auf Leitungen [Cross-talk on cables]. BSI-Kurzinformationen, German Information Security Agency, Bonn, 2001. <http://www.bsi.de/literat/faltbl/uebkopl.htm>
- [19] Schutzmaßnahmen gegen illegales Abhören [Protection measures against illegal eavesdropping]. BSI-Kurzinformationen, German Information Security Agency, Bonn, 2003. <http://www.bsi.de/literat/faltbl/f25schutz.htm>
- [20] Joachim Engel: Bloßstellende Abstrahlung [Compromising emanations]. KES 93/6, pp. 67–68, 1993.
- [21] Ulrich Hecker: Das Zonenmodell: Maßnahme gegen bloßstellende Abstrahlung [The zone model: measure against compromising emanations]. KES 93/6, pp. 69–70, 1993.
- [22] Rainer Jung: Raumschirmungen und Schirmkabinen gegen Abstrahlung und Lauschangriffe [Room shielding and shielded cabin against emanations and eavesdropping attacks]. Part I: KES 94/5, pp. 69–70, Part II: KES 95/1, pp. 57–59.
- [23] Volker Fricke: Überkoppeln auf Leitungen [Cross-talk on cables]. KES 94/5, pp. 63–64.
- [24] Joachim Opfer, Reinhart Engelbart: Verfahren zum Nachweis von verzerrten und stark gestörten Digitalsignalen und Schaltungsanordnung zur Durchführung des Verfahrens [Method for the detection of distorted and strongly interfered digital signals and circuit arrangement for implementing this method]. German Patent DE 4301701 C1, Deutsches Patentamt, 5 May 1994.
- [25] Wolfgang Bitzer, Joachim Opfer: Schaltungsanordnung zum Messen der Korrelationsfunktion zwischen zwei vorgegebenen Signalen [Circuit arrangement for measuring the correlation function between two provided signals]. German Patent DE 3911155 C2, Deutsches Patentamt, 11 November 1993.
- [26] Harold Joseph Highland: Electromagnetic Radiation Revisited. Computers & Security, Vol. 5, pp. 85–93 and 181–184, 1986.

- [27] Kristian Beckman: Läckande Datorer – en information om ROS [Leaking Computers – information on compromising emanations]. Brochure, Brottförebyggande rådet (BRÅ) [National Council for Crime Prevention], Stockholm, Sweden, 1984. Cited in [26, 28].
- [28] Harold Joseph Highland: The Tempest over Leaking Computers. *Abacus*, Vol. 5, No. 2, pp. 10–18 and 53, 1998.
- [29] Wim van Eck: Electromagnetic Radiation from Video Display Units: An Eavesdropping Risk? *Computers & Security*, Vol. 4, pp. 269–286, 1985.
- [30] 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.
- [31] Gerd Schmidt, Michael Festerling: Entstehung, Nachweis und Vermeidung kompromittierender Strahlung [Origin, detection and avoidance of compromising radiation]. *Mess-Comp '92*, 6. Kongreßmesse für die industrielle Meßtechnik, Wiesbaden, 7–9 September 1992.
- [32] Peter Smulders: The Threat of Information Theft by Reception of Electromagnetic Radiation from RS-232 Cables. *Computers & Security*, Vol. 9, pp. 53–58, 1990.
- [33] secuDat 600 – Der Schutz vor den Datenpiraten. Product information, PPM Data Management AG, Zürich.
- [34] Hans-Georg Wolf, Invited talk at Chaos Communication Congress, Berlin, 29 December 1998.
- [35] Stefan Krempel: Lauschangriff am Geldautomaten [Eavesdropping attack on ATM]. *Spiegel Online*, 8 January 1999, Spiegel-Verlag, Germany. <http://www.spiegel.de/netzwelt/technologie/0,1518,13731,00.html>
- [36] Stefan Krempel: Konzerne im Visier – Abhörspezialist Hans-Georg Wolf über Lauschangriffe von Geheimdiensten auf Unternehmen [Interview with eavesdropping expert Hans-Georg Wolf]. c't 4/1999, p. 182, Verlag Heinz Heise, Germany, ISSN 0724-8679.
- [37] *Sicurezza Elettromagnetica nella Protezione dell'Informazione*, ATTI SEPI'88, Rome, Italy, 24–25 November 1988, Fondazione Ugo Bordoni.
- [38] *Symposium on Electromagnetic Security for Information Protection*, SEPI'91, Proceedings, Rome, Italy, 21–22 November 1991, Fondazione Ugo Bordoni.
- [39] Roland Briol: Emanation – How to keep your data confidential. In [38], pp. 225–234.
- [40] Henri Hodara: Secure Fiberoptic Communications. In [38], pp. 259–293.
- [41] B. Demoulin, L. Kone, C. Poudroux, P. Degauque: Electromagnetic Radiation of Shielded Data Transmission Lines. In [38], pp. 163–173.
- [42] Motohisa Kanda, A.R. Ondrejka: Ultra-Broadband and Nondispersive Sensor for the Measurement of Time-Domain Signals. In [38], pp. 65–78.
- [43] N.E. Köksaldı, S.S. Şeker, B. Sankur: Information Extraction from the Radiation of VDUs by Pattern Recognition Methods. EMC'98 Roma, International Symposium on Electromagnetic Compatibility, Roma, Italy, 14–18 September 1998, Vol. 2, pp. 678–683.

- [44] Identification cards – Integrated circuit(s) cards with contacts. International Standard ISO/IEC 7816, International Organization for Standardization, Geneva.
- [45] W. Rankl, W. Effing: Smart Card Handbook. John Wiley & Sons, 2004.
- [46] Security Requirements for Cryptographic Modules. FIPS PUB 140-1, Federal Information Processing Standards Publication, National Institute of Standards and Technology, U.S. Department of Commerce, 11 January 1994.
- [47] Access control system for the MAC/packet family: EUROCRYPT. European Standard EN 50094, European Committee for Standardization (CEN), Brussels, 1993.
- [48] Identification card systems – Intersector electronic purse. European Standard EN 1546, European Committee for Standardization (CEN), Brussels, 1999.
- [49] Oliver Kömmerling, Markus G. Kuhn: Design Principles for Tamper-Resistant Smartcard Processors. Proceedings of the USENIX Workshop on Smartcard Technology (Smartcard '99), Chicago, Illinois, USA, May 10–11, 1999, USENIX Association, pp. 9–20, ISBN 1-880446-34-0.
- [50] Ross J. Anderson, Markus G. Kuhn: Tamper Resistance – a Cautionary Note. The Second USENIX Workshop on Electronic Commerce Proceedings, Oakland, California, 18–21 November 1996, pp. 1–11, ISBN 1-880446-83-9.
- [51] Ross J. Anderson, Markus G. Kuhn: Low Cost Attacks on Tamper Resistant Devices. In M. Lomas et al. (ed.): Security Protocols, 5th International Workshop, Paris, France, 7–9 April 1997, Proceedings, Springer LNCS 1361, pp. 125–136, ISBN 3-540-64040-1.
- [52] Steve H. Weingart: Physical Security for the μ ABYSS System. In Proceedings of the 1987 IEEE Symposium on Security and Privacy, IEEE Computer Society Press, pp. 52–58.
- [53] A Guide to Understanding Data Remanence in Automated Information Systems. National Computer Security Center, Forest Green Book, NCSC-TG-025, September 1991. <http://www.radium.ncsc.mil/tpep/library/rainbow/>
- [54] Peter Gutmann: Data Remanence in Semiconductor Devices. Proceedings of the 10th USENIX Security Symposium, Washington, D.C., USA, 13–17 August 2001.
- [55] Sergei Skorobogatov: Low temperature data remanence in static RAM. Technical Report UCAM-CL-TR-536, University of Cambridge, Computer Laboratory, June 2002.
- [56] P. Kocher, J. Jaffe, B. Jun: Differential Power Analysis. In Michael Wiener (Ed.), Advances in Cryptology – CRYPTO'99, LNCS 1666, Springer, pp. 388–397, 1999.
- [57] Suresh Chari et al.: Towards Sound Approaches to Counteract Power-Analysis Attacks. Advances in Cryptology – CRYPTO'99, Proceedings, Lecture Notes in Computer Science 1666, Springer-Verlag, pp. 398–412, 1999.
- [58] T.S. Messerges, E.A. Dabbish, R.H. Sloan: Examining smart-card security under the threat of power analysis attacks. IEEE Transactions on Computers, Vol. 51, No. 5, May 2002, pp. 541–552.
- [59] Jean-Jacques Quisquater, David Samyde: ElectroMagnetic Analysis (EMA): Measures and Counter-Measures for Smart Cards. Smart Card Programming and Security (E-smart 2001), Cannes, France, LNCS 2140, September 2001, pp. 200–210.
- [60] K. Gandolfi, C. Moutrel, F. Olivier: Electromagnetic Analysis: Concrete Results. Cryptographic Hardware and Embedded Systems – CHES 2001, LNCS 2162, Springer, 2001, pp. 251–261.

- [61] Dakshi Agrawal, Bruce Archambeault, Josyula R. Rao, Pankaj Rohatgi: The EM Side-Channel(s). 4th International Workshop on Cryptographic Hardware and Embedded Systems, LNCS 2523, Springer, 2002, pp. 29–45.
- [62] Dakshi Agrawal, Josyula R. Rao, Pankaj Rohatgi: Multi-channel Attacks. 5th International Workshop on Cryptographic Hardware and Embedded Systems, LNCS 2779, Springer, 2003, pp. 2–16.
- [63] Suresh Chari, Josyula R. Rao, Pankaj Rohatgi: Template Attacks. 4th International Workshop on Cryptographic Hardware and Embedded Systems, LNCS 2523, Springer, 2002, pp. 13–28.
- [64] Egmont R. Koch: Zündstoff – Hacker mit Geheimauftrag. German TV report, Zweites Deutsches Fernsehen (ZDF), 24 April 1996, 22:15–23:00 CEST.
- [65] Egmont R. Koch, Jochen Sperber: Die Datenmafia – Computerspionage und neue Informationskartelle. Rowohlt, 1996, ISBN 3499602474.
- [66] Harry V. Martin: Federal Corruption – INSLAW. Napa Sentinel, 1991. <http://www.sonic.net/sentinel/gvcon7.html>
- [67] Richard L. Fricker: The INSLAW Octopus. Wired 1.01, March/April 1993. http://www.wired.com/wired/archive/1.01/inslaw_pr.html
- [68] Inslaw, the Continuing Caper. UNCLASSIFIED No. 37, Summer 1996. <http://www.io.com/~patrik/promis.htm>
- [69] Henning F. Harmuth: Transmission of Information by Orthogonal Functions. 2nd Ed., Springer, 1972, ISBN 3-540-05512-6.
- [70] NSA/CSS Regulation 90-6: Technical Security Program. National Security Agency – Central Security Service, Fort George G. Meade, Maryland, 12 pages, 31 May 199?. Partially declassified transcript: <http://cryptome.org/nsa-reg90-6.htm>
- [71] Joel McNamara: The Complete, Unofficial TEMPEST Information Page. Internet Web page, 1996–2003. <http://www.eskimo.com/~joelm/tempest.html>
- [72] 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.
- [73] 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.
- [74] Karl Rothammel, Alois Krischke: Rothammels Antennenbuch. Franckh-Kosmos, Stuttgart, 1995.
- [75] Ralph S. Carson: Radio Communication Concepts: Analog. Wiley, 1990.
- [76] Norbert Wiener: Extrapolation, Interpolation and Smoothing of Stationary Time Series with Engineering Applications, Wiley, New York, 1949.
- [77] William H. Press et al.: Numerical Recipes in C. Cambridge University Press, 1992.
- [78] Model R-1250 Wide Range Receiver Instruction Manual, document number IM 469750B, Dynamic Sciences Incorporated, March 1985, <http://www.dynamic-sciences.com/>

- [79] Test Receiver R&S FSET7, R&S FSET22, RF Preselector R&S FSET-Z2, R&S FSET-Z22. Data sheet, Rohde & Schwarz GmbH & Co. KG, München, Germany, <http://www.rohde-schwarz.com/>
- [80] Model R-1150-10A Portable Antenna Kit Instruction Manual, document number IM 470000, Dynamic Sciences Incorporated, Chatsworth, California, January 1988.
- [81] IEEE Standard for the Measurement of Impulse Strength and Impulse Bandwidth, ANSI/IEEE Std 376-1975.
- [82] Monitor Timing Specifications, Version 1.0, Revision 0.8, Video Electronics Standards Association (VESA), San Jose, California, September 17, 1998.
- [83] Michael Bach, Thomas Meigen, Hans Strasburger: Raster-scan cathode-ray tubes for vision research—limits of resolution in space, time and intensity and some solutions. *Spatial Vision*, Vol. 10, No. 4, pp. 403–414, 1997.
- [84] Stanley A. Klein, Q. James Hu, Thom Carney: The Adjacent Pixel Nonlinearity: Problems and Solutions. *Vision Research*, Vol. 36, No. 19, pp. 3167–3181, 1996.
- [85] Test methods for Visual Display Units – visual ergonomics and emission characteristics. Swedish National Board for Measurement and Testing, MPR 1990:8.
- [86] TCO'99 – Mandatory and recommended requirements for CRT-type Visual Display Units (VDUs). Swedish Confederation of Professional Employees (TCO), 1999. <http://www.tcodevelopment.com/>
- [87] Procedure for Measurement of Emissions of Electric and Magnetic Fields from VDUs from 5 Hz to 400 kHz. European Computer Manufacturers Association, Standard ECMA-172, June 1992. (also IEEE Std 1140-1994)
- [88] Council Directive 89/336/EEC of 3 May 1989 on the Approximation of the Laws of the Member States Relating to Electromagnetic Compatibility. *Official Journal of the European Community*, L 139, pp. 19–26, 1989-05-23.
- [89] Limits and Methods of Measurement of Radio Interference Characteristics of Information Technology Equipment. EN 55022, CENELEC, Brussels, 1987.
- [90] Specification for radio disturbance and immunity measuring apparatus and methods. CISPR 16, International Electrotechnical Commission (IEC), Geneva, 2000.
- [91] Information technology equipment – Radio disturbance characteristics – Limits and methods of measurement. CISPR 22, International Electrotechnical Commission (IEC), Geneva, 1997.
- [92] Requirements for the Control of Electromagnetic Interference Characteristics of Subsystems and Equipment. MIL-STD-461E, US Department of Defense, Interface Standard, 20 August 1999.
- [93] BSI-Handbuch für digitale Signaturen [German Information Security Agency Manual for Digital Signatures]. Version 1.1, Regulatory Authority for Telecommunications and Posts (RegTP) and German Information Security Agency (BSI), Germany, 18 November 1997. <http://www.bsi.bund.de/esig/basics/techbas/masskat/bsikat.pdf>
- [94] LVDS Transmitter 24-Bit Color Flat Panel Display (FPD) Link, National Semiconductor Cooperation, 1998. <http://www.national.com/pf/DS/DS90CF581.html>
- [95] Electrical characteristics of low voltage differential signaling (LVDS) interface circuits, ANSI/TIA/EIA-644, Electronic Industries Alliance, 1996.

- [96] PanelLink Technology, EMC Design Application Note – LCD Monitor & Notebook Applications. Revision B, Silicon Image Inc., April 1998. <http://www.siimage.com/documents/SiI-AN-0003-B.PDF>
- [97] VESA Plug and Display Standard. Version 1, Video Electronics Standards Association, 11 June 1997.
- [98] VESA Digital Flat Panel (DFP). Version 1, Video Electronics Standards Association, 14 February 1999.
- [99] Digital Visual Interface – DVI. Revision 1.0, Digital Display Working Group, April 1999. <http://www.ddwg.org/>
- [100] Scott Crosby, Ian Goldberg, Robert Johnson, Dawn Song, David Wagner: A Cryptanalysis of the High-Bandwidth Digital Content Protection System. ACM CCS-8 Workshop DRM 2001, 5 November 2001, Springer, LNCS 2320, pp. 192–200.
- [101] Radio noise. Recommendation ITU-R P.372-7, International Telecommunication Union, Geneva, 2001.
- [102] Propagation data and prediction methods for the planning of indoor radiocommunication systems and radio local area networks in the frequency range 900 MHz to 100 GHz. Recommendation ITU-R P.1238-2, International Telecommunication Union, Geneva, 2001.
- [103] Homayoun Hashemi: The Indoor Radio Propagation Channel. Proceedings of the IEEE, Vol. 81, No. 7, July 1993, pp. 943–968.
- [104] L. P. Rice: Radio Transmission into Buildings at 35 and 150 mc [MHz]. Bell System Technical Journal, Vol. 38, No. 1, January 1959, pp. 197–210.
- [105] Manfred Zimmermann, Klaus Dostert: A Multipath Model for the Powerline Channel. IEEE Transactions on Communications, Vol. 50, No. 4, April 2002, pp. 553–559.
- [106] D. Liu, E. Flint, B. Gaucher, Y. Kwark: Wide Band AC Power Line Characterization. IEEE Transactions on Consumer Electronics, Vol. 45, No. 4, November 1999, pp. 1087–1097.
- [107] Test & Measurement Catalog 2001, Agilent Technologies, USA.
- [108] Leon W. Couch: Digital and Analog Communication Systems. 2nd ed., Macmillan, New York, 1987.
- [109] Peter A. Keller: Electronic Display Measurement – Concepts, Techniques and Instrumentation. John Wiley & Sons, New York, 1997.
- [110] Dell D1025HE Color Monitor User’s Guide, ZF5368, April 1997.
- [111] Measurement of Phosphor Persistence of CRT Screens. Electronic Industries Alliance (EIA), Tube Electron Panel Advisory Council (TEPAC), Publication TEP105-14, Arlington, Virginia, April 1987.
- [112] Worldwide Type Designation System for TV Picture Tubes and Monitor Tubes. Electronic Industries Alliance (EIA), Tube Electron Panel Advisory Council (TEPAC), Publication TEP106-B, Arlington, Virginia, June 1988.
- [113] Optical Characteristics of Cathode-Ray Tube Screens. Electronic Industries Alliance (EIA), Tube Electron Panel Advisory Council (TEPAC), Publication TEP116-C, Arlington, Virginia, February 1993.

- [114] W. Wolf, H. Deubel: P31 phosphor persistence at photopic mean luminance level. *Spatial Vision*, Vol. 10, No. 4, pp. 323–333, 1997.
- [115] Photosensor Modules H5773/H5783/H6779/H6780/H5784 Series. Hamamatsu Photonics K.K., 2000. <http://www.hamamatsu.com/>
- [116] K. Steiglitz, L.E. McBride: A Technique for the Identification of Linear Systems. *IEEE Transactions on Automatic Control*, Vol. AC-10, pp. 461–464, 1965.
- [117] Matteo Frigo, Steven G. Johnson: FFTW: An Adaptive Software Architecture for the FFT, *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing*, Vol. 3, pp. 1381–1384, 1998. <http://www.fftw.org/>
- [118] Tudor E. Jenkins: *Optical Sensing Techniques and Signal Processing*. Prentice-Hall International, 1987.
- [119] Joe Loughry, David A. Umphress: Information Leakage from Optical Emanations. *ACM Transactions on Information Systems Security*, Vol. 5, No. 3, pp. 262–289, August 2002.
- [120] Rodger E. Ziemer, Roger L. Peterson: *Digital Communications and Spread Spectrum Systems*. Macmillan, New York, 1985.
- [121] CIE standard illuminants for colorimetry. International Standard ISO/CIE 10526, International Organization for Standardization, Geneva, 1999. See <http://www.cvr1.org/> for online tables.
- [122] Vanu Bose, Mike Ismert, Matt Welborn, John Gutttag: Virtual Radios. *IEEE Journal on Selected Areas in Communications*, Vol. 17, No. 4 (Special Issue on Software Radios), April 1999, pp. 591–590.
- [123] Carlo Kopp: An Introduction to the Technical and Operational Aspects of the Electromagnetic Bomb. Air Power Studies Centre, Paper No. 50, Royal Australian Air Force, November 1996.
- [124] James Clark Maxwell: *A Treatise on Electricity and Magnetism*. London, 1873.
- [125] Otto Zinke, Heinrich Brunswig: *Hochfrequenztechnik 1*. 6th edition, Springer, 2000.
- [126] Reinaldo Perez (ed.): *Handbook of Electromagnetic Compatibility*. Academic Press, 1995.
- [127] Quantities and units — Part 5: Electricity and magnetism. International Standard ISO 31-5, International Organization for Standardization, Geneva, 1992.
- [128] Quantities and units — Part 6: Light and related electromagnetic radiations. International Standard ISO 31-6, International Organization for Standardization, Geneva, 1992.
- [129] J. A. Stratton: *Electromagnetic Theory*. New York, London, 1941.

Appendix A

Electromagnetic fields

A.1 Maxwell's equations

The effects of electric and magnetic fields were first accurately investigated by Michael Faraday (1791–1867) and André Marie Ampère (1775–1836) and finally described in a unifying theory by James Clark Maxwell (1831–1879) in 1873 [124]. Some familiarity with Maxwell's theory is helpful for understanding not only the physical effects underlying electromagnetic emanations from electronic circuits but also the transducers used to measure these and the units in which they are quantified. This section provides only a quick review of the most important basics; readers interested in a more detailed introduction are referred to standard physics, electrodynamics and electromagnetic compatibility textbooks [125, 126]. The notational conventions used here follow the ISO 31 standard [127, 128].

In Maxwell's theory, the electromagnetic aspects of a portion of space and time are described by five vector fields \mathbf{D} , \mathbf{E} , \mathbf{B} , \mathbf{H} and \mathbf{J} , as well as four scalar fields ϱ , ε_r , μ_r , and σ . **Bold** letters designate vectors. The rôles of these fields are as follows:

ϱ The **charge density** represents the volume density of electric charge, that is the number of missing ($\varrho > 0$) or surplus ($\varrho < 0$) electrons per unit of space. The unit is 1 C/m^3 , where 1 coulomb ($1 \text{ C} = 1 \text{ As}$) represents the positive charge of 6.24146×10^{18} missing electrons.

\mathbf{D} The **electric flux density** or **displacement** is a vector quantity that if integrated over an area A (where $d\mathbf{A}$ denotes the outwards facing normal vector of area element dA) that encloses a volume V will give as a result the electric charge Q located in V , that is

$$\oint_A \mathbf{D} \cdot d\mathbf{A} = \int_V \varrho \cdot dV = Q, \quad (\text{A.1})$$

which according to Gauss' integral theorem is equivalent with

$$\text{div } \mathbf{D} = \varrho. \quad (\text{A.2})$$

The unit is 1 C/m^2 .

- E** The **electrical field strength** describes the charge-proportional force $\mathbf{F} = \mathbf{E} \cdot Q$ that acts on a small test charge Q placed at some point in space. The unit is $1 \text{ V/m} = 1 \text{ N/C}$. $|\mathbf{E}| = |\varphi_1 - \varphi_2|/d$ is also the homogeneous electrical field strength found in a capacitor with plate distance d and static plate potentials φ_1 and φ_2 . [Static electric fields with $\text{rot } \mathbf{E} = 0$ define a potential field φ with $\mathbf{E} = \mathbf{grad } \varphi$, which describes the charge-proportional potential energy of a small test charge at any point in space (unit: 1 V).]
- J** The **electric current density** describes the motion of charge as the electric current I flowing through area A , where $d\mathbf{A}$ denotes again the normal vector of area element dA facing in the direction of positive current:

$$\int_A \mathbf{J} \cdot d\mathbf{A} = I. \quad (\text{A.3})$$

The unit is 1 A/m^2 . Charge cannot be created or destroyed in a volume V enclosed by an area A (where $d\mathbf{A}$ denotes the outwards facing normal vector of area element dA), therefore

$$\oint_A \mathbf{J} \cdot d\mathbf{A} = - \int_V \frac{\partial \rho}{\partial t} dV \quad \Rightarrow \quad \text{div } \mathbf{J} = - \frac{\partial \rho}{\partial t}. \quad (\text{A.4})$$

- σ The **conductivity** describes for a conductive medium the current density generated by an electric field:

$$\mathbf{J} = \sigma \cdot \mathbf{E} \quad (\text{A.5})$$

The unit is 1 S/m .

- B** The **magnetic induction** or **magnetic flux density** describes the current-proportional force $d\mathbf{F}$ acting on a conductive volume dV with current density \mathbf{J}

$$d\mathbf{F} = dV \cdot \mathbf{J} \times \mathbf{B}$$

or alternatively, if we integrate over volume, it describes the force $\mathbf{F} = Q \cdot \mathbf{v} \times \mathbf{B}$ acting on a charge Q moving with speed \mathbf{v} . The unit is 1 tesla ($1 \text{ T} = 1 \text{ N/Am}$). There exists no magnetic charge, hence

$$\text{div } \mathbf{B} = 0. \quad (\text{A.6})$$

- H** The **magnetic field strength** is more an auxiliary mathematical construct than a directly measurable physical effect. It is a vector quantity defined such that if integrated over the boundary of an area A results in

$$\oint \mathbf{H} \cdot d\mathbf{s} = \int_A \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t} \cdot d\mathbf{A}, \quad (\text{A.7})$$

which is equivalent to

$$\text{rot } \mathbf{H} = \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t}. \quad (\text{A.8})$$

The unit is 1 A/m . $|\mathbf{H}| = I \cdot n/l$ is also the homogeneous magnetic field strength found inside a long coil with current I and n windings per length l along the axis of the coil.

S The Poynting vector

$$\mathbf{S} = \mathbf{E} \times \mathbf{H} \quad (\text{A.9})$$

represents the power flux density of an electromagnetic field. The unit is 1 W/m^2 .

ε_r The **relative permittivity** is a material constant that describes the factor by which the electric field strength \mathbf{E} is attenuated due to charge displacement in a dielectric substance. The value for vacuum is 1.

μ_r The **relative permeability** is a material constant that describes the factor by which the magnetic induction \mathbf{B} is changed due to the magnetization of a diamagnetic ($\mu_r < 1$) or paramagnetic ($\mu_r > 1$) material. The value for vacuum is 1.

To summarize, the \mathbf{E} and \mathbf{B} field together describe the forces that act on charged particles. The force caused by an \mathbf{E} field depends only on the charge, while the force caused by the \mathbf{B} field depends in addition on the velocity of the charge. The \mathbf{D} and \mathbf{H} field describe the origins of these fields, namely the charge density for the electric field and the electric current density for the magnetic field.

In order to relate the four fields describing the causes and effects of electromagnetism, we also need two constants:

- The **electric constant** or **permittivity of vacuum**

$$\varepsilon_0 = 8.854188 \times 10^{-12} \text{ C/Vm} \quad (\text{A.10})$$

is the quotient D/E between the surface density of charge D located on a flat surface and the electric field caused by this charge in the surrounding vacuum. In general:

$$\mathbf{D} = \varepsilon_0 \varepsilon_r \mathbf{E} \quad (\text{A.11})$$

- The **magnetic constant** or **permeability of vacuum**

$$\mu_0 = 4\pi \cdot 10^{-7} \text{ H/m} \quad (\text{A.12})$$

is the quotient of magnetic induction B and magnetic field strength H in vacuum. In general:

$$\mathbf{B} = \mu_0 \mu_r \mathbf{H} \quad (\text{A.13})$$

The term $c_0 = 1/\sqrt{\varepsilon_0 \mu_0} = 299\,792\,458 \text{ m/s}$ is the velocity of electromagnetic waves in vacuum and $c = c_0/\sqrt{\varepsilon_r \mu_r}$ is their velocity in an arbitrary medium.

Electric fields are not only generated by a charge distribution and magnetic fields not only by moving charges. Changes in either field type will generate an additional component of the other type.

We already included in the definition of the \mathbf{H} -field, how a variation in the electric flux density \mathbf{D} over time generates a curled magnetic field. Similarly, Faraday's law of induction describes how a changing **magnetic flux** $\Phi = \int_A \mathbf{B} \cdot d\mathbf{A}$ generates the curled electric field that a coil converts into a voltage:

$$\oint \mathbf{E} \cdot d\mathbf{s} = -\frac{d\Phi}{dt} \quad (\text{A.14})$$

which is equivalent to

$$\text{rot } \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}. \quad (\text{A.15})$$

This completes the definition of Maxwell's equations and the variables involved.

A.2 Quantities and units

The SI units for \mathbf{H} and \mathbf{E} field strength are A/m and V/m, respectively. In addition to these simple and well-defined SI base units, the field of communications engineering has evolved its very own set of allegedly more practical auxiliary notations for quantitative measurements of signals and noise. The quantities that communication engineers have to deal with often vary over many orders of magnitude, which makes it difficult to agree on a common SI prefix. In addition, the quotient of two quantities (e.g., amplification, attenuation) is usually of greater interest than the absolute difference. Thirdly, signal strengths can usually equally reasonably be expressed as either a field quantity (voltage, current, pressure, etc.) or as a power value, but the quadratic relationship between these two representations ($P = U^2/R = I^2R$) is rather inconvenient for mental arithmetic. (And finally, like navigators and typographers, communication engineers could not resist the temptation of distinguishing their trade by establishing magic special-purpose units, just to keep the uninitiated out.)

As a result, communication engineers have become highly accustomed to the use of relative logarithmic scales instead of absolute units. The unit **neper (Np)** is defined to denote the natural logarithm of the quotient of a field quantity F and a reference value F_0 , that is F has a level of $\ln F/F_0$ Np. Far more common in practice though is the unit **bel (B)**, which is defined as the decadic (base 10) logarithm of a power quantity P divided by a reference power P_0 , that is the level of P is $\log_{10} P/P_0$ B. The bel is usually used in the form of **decibel (dB)**, such that 1 B = 10 dB refers to a 10× increase of power, 20 dB stands for a 100× increase of power or equivalently a 10× increase in voltage or current, 30 dB for a 1000× increase of power, etc.

Common reference field and power levels are indicated in the form of a suffix to the unit abbreviation dB:

$$\begin{aligned} 0 \text{ dBW} &= 1 \text{ W} \\ 0 \text{ dBm} &= 1 \text{ mW} = -30 \text{ dBW} \\ 0 \text{ dB}\mu\text{V} &= 1 \mu\text{V} \\ 0 \text{ dBV/m} &= 1 \text{ V/m} \\ 0 \text{ dBA/m} &= 1 \text{ A/m} \\ 0 \text{ dB}_{\text{SPL}} &= 20 \mu\text{Pa} \quad (\text{sound pressure level}) \end{aligned}$$

The intensity of radio frequency signals, conducted by a cable, is commonly described in either dBm or dB μ V. The power P and voltage U can be converted into each other using the relation $P = U^2/R$ when both are considered in relation to a fixed load resistance R and when the voltage is root-mean-squared. In the case $R = 50 \Omega$, which is the standard impedance of RF measurement equipment and coaxial cables, we get $P = 0 \text{ dBm} = 10^{0/10} \text{ mW} \Rightarrow U = \sqrt{PR} = \sqrt{1 \text{ mW} \times 50 \Omega} \approx 0.2236 \text{ V} \approx 10^{107/20} \mu\text{V} = 107 \text{ dB}\mu\text{V}$.

A.3 Electromagnetic emanations

As a trivial example, we can calculate the magnetic field strength $|\mathbf{H}|$ at a point located at distance r from an infinitely long conductor with constant current I and no electric

field fluctuations. Due to the rotational symmetry of the setup, all points along the circle C of length $2\pi r$ that contains our point and is rotationally symmetric with regard to the conductor must have an equal field strength. Therefore, we can follow from

$$2\pi r |\mathbf{H}| = \oint_C \mathbf{H} \cdot d\mathbf{s} = \int_A \mathbf{J} + \frac{\partial \mathbf{D}}{\partial t} \cdot d\mathbf{A} = I \quad (\text{A.16})$$

that

$$|\mathbf{H}| = \frac{I}{2\pi r}. \quad (\text{A.17})$$

For numerically calculating the electromagnetic field generated by a given arbitrary current distribution, it is useful to introduce an auxiliary vector field \mathbf{A} called the **magnetic vector potential**. The lack of a magnetic charge ($\text{div } \mathbf{B} = 0$) allows us to represent the magnetic field such that

$$\mathbf{B} = \text{rot } \mathbf{A}. \quad (\text{A.18})$$

If the volume under consideration is not conductive ($\sigma = 0$), then both the magnetic and electric potential can be calculated from the generating time-dependent current distribution field $\mathbf{J}(t)$ and charge density field $\varrho(t)$ [125, p. 263],[129]:

$$\mathbf{A} = \frac{\mu_0 \mu_r}{4\pi} \int_V \frac{\mathbf{J}(t - \frac{r}{c})}{r} dV \quad (\text{A.19})$$

$$\varphi = \frac{1}{4\pi \epsilon_0 \epsilon_r} \int_V \frac{\varrho(t - \frac{r}{c})}{r} dV \quad (\text{A.20})$$

where r is the distance between dV and the current location. Electric and magnetic field strength can now be determined via

$$\mathbf{H} = \frac{1}{\mu_0 \mu_r} \text{rot } \mathbf{A} \quad (\text{A.21})$$

$$\mathbf{E} = -\text{grad } \varphi - \frac{\partial \mathbf{A}}{\partial t}. \quad (\text{A.22})$$

Very useful in practice are worst-case estimation formulae for the electric fields generated by currents in cables and other typical arrangements of conductors. For example, the following estimates derived in [126, pp. 178–187] are valid for cables of length l with two parallel conductors at distance d , where l is significantly shorter than the wavelengths appearing on the wire, such that the current can be assumed to be constant along the wires (Hertzian dipole).

Two types of currents on wire pairs have to be distinguished. A *differential-mode* current I_d is one that flows in both wires in opposite directions. This is the case, for example, if one of the two wires provides the ground return for the signal conducted in the other. A *common-mode* current I_c on the other hand is one that appears on each wire in the same direction. Common-mode currents occur if there is some alternative return path to the wire pair, for example a ground loop, or for higher frequencies even just a capacitance against ground. The electric field strength generated by both types of currents at distance D from their center in the plane of the wires is then:

$$|\mathbf{E}_d| = 1.316 \times 10^{-14} \frac{\text{V}}{\text{Hz}^2 \cdot \text{m}^2 \cdot \text{A}} \cdot \frac{f^2 \cdot l \cdot d \cdot |I_d|}{D} \quad (\text{A.23})$$

$$|\mathbf{E}_c| = 1.257 \times 10^{-6} \frac{\text{V}}{\text{Hz} \cdot \text{m} \cdot \text{A}} \cdot \frac{f \cdot l \cdot |I_c|}{D} \quad (\text{A.24})$$

Common-mode currents lead to much stronger fields than differential-mode currents of similar magnitude, because the fields generated by currents flowing in opposing directions in nearby wires will mainly cancel each other out. Common-mode current field strength grows linearly with the signal frequency (20 dB per decade), while differential-mode fields grow with the square of the frequency (40 dB per decade), as long as the wavelength is longer than the cable segment of interest. Differential-mode currents can be predicted quite well from transmission-line models, but common-mode currents in a cable usually have to be measured with a current probe clamped around the cable, as the parasitic capacitances against ground causing them are difficult to estimate [126]. A transmission-line technique aimed at reducing both differential and common-mode emissions are balanced twisted-pair lines. Twisting the wires reduces the effective size of the loop area $l \cdot d$, which is one factor in the field strength of differential-mode emissions. Driving the lines with voltages whose sum equals at all times the ground potential helps to minimize the common-mode current, as from both wires currents of equal magnitude but opposite direction will flow through the capacitance to ground.

A.4 Transmission lines and antennas

Figure A.1 shows a model circuit for an infinitely short cable segment of length dz . Here, functions $R(z)$, $L(z)$, $G(z)$ and $C(z)$ represent the total resistance, inductance, insulation conductivity and capacity in the cable from its start ($z = 0$) up to this segment. For moderate cable lengths, currents, and voltages, the resistance dR/dz and insulation conductivity dG/dz per cable length are negligibly small. For signals with a wavelength significantly longer than the cable, the relationship between voltage and current is primarily determined by the impedance of the load connected to the receiving end. For signals of higher frequency, however, the voltage changes at the transmitting end happen very quickly compared to the round-trip time, and the load connected to the far end loses its immediate influence on the current flowing into the cable. The current flowing into the transmission line depends, apart from the applied voltage, for signals of high frequency, primarily on the inductance dL/dz of the two conductors, and the capacitance dC/dz , which between them determine the impedance

$$Z = \sqrt{\frac{dR + j\omega dL}{dG + j\omega dC}} \approx \sqrt{\frac{dL}{dC}} \quad (\text{A.25})$$

of the transmission line. With the approximation $dR = dG = 0$, the transmission line impedance is independent of the signal frequency $\omega/2\pi$ and the cable length l .

The propagation of a voltage waveform $u(z, t)$ on a transmission line is described by

$$\frac{\partial^2 u}{\partial z^2} = \frac{dL dC}{dz^2} \left[\frac{\partial^2 u}{\partial t^2} + \left(\frac{dR}{dL} + \frac{dG}{dC} \right) \frac{\partial u}{\partial t} + \frac{dR dG}{dL dC} u \right] \approx \frac{dL dC}{dz^2} \frac{\partial^2 u}{\partial t^2}, \quad (\text{A.26})$$

the “telegrapher’s equation” [125, p. 54].

Solutions to the $dR = dG = 0$ simplification of this differential equation can be represented as the sum $u(z, t) = u_p(z, t) + u_r(z, t)$ of two waveforms that travel along the line with speeds

$$c \approx \sqrt{\frac{dL dC}{dz^2}} \quad (\text{A.27})$$




Figure A.1: Model circuit for a short segment of a transmission line.

and $-c$, respectively. The corresponding current on the line is $i(z, t) = i_p(z, t) - i_r(z, t) = u_p(z, t)/Z - u_r(z, t)/Z$.

If we apply to a cable a waveform (such as a short pulse) $u_p(0, t)$ at one end ($z = 0$), it will travel as a voltage

$$u_p(z, t) = u_p(0, t - z/c) \quad (\text{A.28})$$

and current

$$i_p(z, t) = \frac{u_p(z, t)}{Z} \quad (\text{A.29})$$

along the line until it reaches the far end. Let's assume the cable end ($z = l$) is terminated with a load of impedance Z_2 . If Z_2 differs from the impedance Z of the cable, then the current and voltage $i(l, t) = u(l, t)/Z_2$ that we can measure at the far end will also differ from the $i_p(l, t) = u_p(l, t)/Z$ propagated along the cable. The only way to accommodate the difference between these two voltages in the solution space of the telegrapher's equation is the addition of a second waveform $i_r(z, t) = u_r(z, t)/Z$ traveling in the opposite direction, such that the sum of both is consistent with the voltage over the terminating load:

$$Z_2 = \frac{u_p(l, t) + u_r(l, t)}{i_p(l, t) - i_r(l, t)} = Z \cdot \frac{u_p(l, t) + u_r(l, t)}{u_p(l, t) - u_r(l, t)} \Rightarrow \frac{u_r(l, t)}{u_p(l, t)} = \frac{Z_2 - Z}{Z_2 + Z} \quad (\text{A.30})$$

In other words, unless the impedance of what is connected to the end of the line matches the impedance of the line ($Z_2 = Z$), the signal will be reflected back, for example with unchanged voltage if the end is open ($Z_2 = \infty \Rightarrow u_r(0, t) = u_p(0, t - 2l/c)$) or inverted if the end is a short circuit ($Z_2 = 0 \Rightarrow u_r(0, t) = -u_p(0, t - 2l/c)$). In order to avoid signal distortions caused by such reflections at the end of transmission lines, every interface has to be designed with the same impedance as the interconnecting cable, which is by convention chosen to be 50Ω for RF laboratory equipment such as antennas, amplifiers, receivers, and oscilloscopes. In coaxial cables and connectors, this is achieved by choosing the inner mantle diameter D and the outer core diameter d such that $Z = (138 \Omega / \sqrt{\epsilon_r}) \cdot \log_{10} D/d = 50 \Omega$. For polyethylene, the most common insulator material in coaxial cables, we have $\epsilon_r = 2.3$, which results in a signal propagation speed of $c = c_0 / \sqrt{\epsilon_r} = 0.66 c_0$ [126, 74].

Antennas are interface devices that convert between the voltage and current signal propagated along a metallic transmission line and the electromagnetic field propagated through

free space. Free space has an impedance

$$Z_0 = E/H = \sqrt{\frac{\mu_0}{\varepsilon_0}} \approx 377 \, \Omega, \quad (\text{A.31})$$

that is, the peak electric field strength of a freely propagating electromagnetic wave measured in V/m is 377 times the peak magnetic field strength measured in A/m. At least, this is the case if we are sufficiently far away from the transmitting antenna, in its “far field”. In the “near field” of a transmitting antenna, the size of which is typically less than a wavelength, the magnitude relationship between magnetic and electric fields can differ significantly from $E/H = Z_0$.

One of the simplest narrow-band antennas for electric fields is the $\lambda/2$ dipole or free-space-resonant dipole. It consists of a metal wire, whose length is half the wavelength $\lambda = c_0/f$ of the intended transmission or reception frequency f . In a receiving dipole, the electric field in the surrounding space exercises a force on its electrons. This in turn leads to a voltage and current distribution over the dipole that can be propagated into a transmission line connected to the center end points of the dipole, whose two half segments are otherwise insulated from each other.

The dipole is also an LC resonator that filters a narrow frequency band out of the spectral composition of the surrounding electric field. For optimal frequency adjustment, the dipole has to be shortened by a factor $V \approx 0.90\text{--}0.95$ depending on the diameter of the wire [74, p. 57].

A receiving antenna in a field with an energy flux density of

$$S = \frac{E^2}{377 \, \Omega}, \quad (\text{A.32})$$

where E is the root mean square of the electric field magnitude, extracts from the field the signal power

$$P = SA_e, \quad (\text{A.33})$$

where A_e is the *effective aperture* of the antenna. For a $\lambda/2$ dipole, this value is

$$A_e = \frac{\lambda^2}{4\pi} \cdot G, \quad (\text{A.34})$$

where $G = 1.64$ is the directional gain of a dipole compared to a (hypothetic) isotropic antenna whose gain is identical in all directions.

The *effective length* l_e of an antenna is a quantity that characterizes the voltage

$$U = E \cdot l_e \quad (\text{A.35})$$

created at an open output connector of an antenna by an electric field strength E . Its value

$$l_e = \frac{\lambda}{\pi} \quad (\text{A.36})$$

for a $\lambda/2$ dipole is somewhat shorter than the actual length of the dipole due to the non-uniform distribution of current on the dipole. In order to extract most energy from a receiving antenna, the connected amplifier must have the same impedance as the antenna, in which case it will see as its input voltage

$$U = E \cdot \frac{\lambda}{2\pi}. \quad (\text{A.37})$$

which is half the antenna's open-circuit voltage. [74]

An effective antenna has to transform the voltage/current relationship (impedance) between the transmission line and free space, otherwise much of the signal energy would be reflected back, either into the cable in the case of a transmission antenna, or into the surrounding space in the case of a receiving antenna. The advantage of a $\lambda/2$ dipole is that its impedance at its resonance frequency is about 60Ω , which makes it suitable for connection to a 50Ω cable without an impedance matching transformer. Even though the near field and the current distribution can differ significantly when the same passive antenna is used for transmission or reception, the impedance and directional gain are identical in either case.

A.5 Time-domain characterization of antennas

A simple time-domain reflectometry experiment demonstrates the frequency dependency of the antenna impedance (Fig. A.3). In this experiment, an impulse generator produces a pulse that is shorter than 1 ns, which travels along a transmission line to an antenna. Part of the impulse energy will be emitted by the antenna as an electromagnetic wave, a small amount will be converted into heat and the rest will be reflected back into the cable. A digital storage oscilloscope can be used to record the impulse waveform on its way both to and from the antenna (Fig. A.3 (a) and (b)), and comparing the Fourier transforms of these signals allows us then to estimate the frequency-dependent impedance of the antenna compared to that of the transmission line (Fig. A.3 (c)–(e)).

The classic technique for determining the impedance of an antenna is to use a continuous waveform (CW) signal generator to send a sine wave $u_p(l, t) = \text{Re}(U_p \cdot e^{2\pi j f t})$ where $j^2 = -1$ to the antenna, via a transmission line of known impedance Z . From knowing both the primary and reflected voltages U_p and $U_r = U_p \frac{Z_a - Z}{Z_a + Z}$, or at least the reflection coefficient $r = U_r/U_p$, we can calculate the antenna impedance

$$Z_a = Z \cdot \frac{U_p + U_r}{U_p - U_r} = Z \cdot \frac{1 + \frac{U_r}{U_p}}{1 - \frac{U_r}{U_p}} = Z \cdot \frac{1 + r}{1 - r}. \quad (\text{A.38})$$

The problem of a continuous waveform in this experiment is that u_p and u_r will overlap and therefore the amplitudes U_p and U_r cannot be measured simply by tapping the transmission line at a single point. One solution is to tap the line at several points to observe the standing wave $u_p + u_r$ and determine the *standing wave ratio* (SWR)

$$s = \frac{|U_p| + |U_r|}{|U_p| - |U_r|} = \frac{1 + |r|}{1 - |r|} \quad (\text{A.39})$$

between the maximum and minimum of the peak voltage at various points on the line. Another approach is the use of a special transformer to separate u_p and u_r , such that their individual amplitudes can be determined. The third approach is to send a signal that is significantly shorter than the round-trip time of the cable, such that u_p and u_r are not non-zero simultaneously at some distance from the antenna. Short impulses carry energy distributed over a broad frequency spectrum, at least up to about $f = 1/(4T)$, where T is the width of the pulse (Fig. A.2). This way, the impedance of the antenna under test




Figure A.2: The amplitude spectrum of a rectangular pulse of duration T is $UT \left| \frac{\sin \pi f T}{\pi f T} \right|$. The spectrum of a pulse with smoother edges has a similar main lobe up to $1/T$, but higher lobes vanish more quickly.

can be measured for all frequencies of interest, by looking at just the reflection of a single pulse.

In the experimental setup used for the measurements in Fig. A.3, a Dynamic Sciences R-1160C impulse generator set to a pulse spectral density of $60 \text{ dB}\mu\text{V}/\text{MHz}$ injects a pulse into a 5 m long RG 213 50Ω coax cable, which leads to a BNC-T-connector on the input of a Tektronix TDS 7054 digital storage oscilloscope. The antenna that is to be tested is connected to the other end of this T-connector via a 15 m RG 213 coax cable (acting as a delay line). The oscilloscope input is configured to 50Ω impedance with DC coupling. This is necessary in order to achieve a constant input impedance over the entire frequency range. In the normal $1 \text{ M}\Omega$ high-impedance mode, the impedance of the oscilloscope (even if an available active probe head with only 1 pF input impedance were used) would drop from $1 \text{ M}\Omega$ at frequencies below 100 kHz to about 200Ω at 1 GHz .

Since the oscilloscope does not offer a high-impedance UHF input, an impedance mismatch and reflection at the T-connector is difficult to avoid and has to be compensated for numerically instead. The signal arriving via a 50Ω cable from the generator will see at the T-connector two 50Ω loads (oscilloscope and second cable) connected to it in parallel, which forms a 25Ω terminating load that will reflect $\frac{25-50}{25+50} = -\frac{1}{3}$ of the original pulse voltage back to the generator where it will be absorbed. The $-\frac{1}{3}u_p$ reflection reduces the pulse propagating into both the delay line and oscilloscope by that amount to $\frac{2}{3}u_p$.

This remaining pulse continues to propagate through the second cable towards the antenna. The antenna impedance $Z_a(f)$ does not match the 50Ω of the cable at all frequencies, therefore a pulse with an amplitude spectrum

$$\mathcal{F}(u_r)(f) = \frac{2}{3} \mathcal{F}(u_p)(f) \cdot \frac{Z_a(f) - 50 \Omega}{Z_a(f) + 50 \Omega} \quad (\text{A.40})$$

will be reflected back. The oscilloscope records that after another reduction by $\frac{2}{3}$, again due to the T-connector attenuation.

Figure A.3 (a) and (b) show the two pulses recorded by the oscilloscope, with the T-connector losses of $\frac{2}{3}$ and $(\frac{2}{3})^2$, respectively, compensated for, to show the voltages u_p and u_r as if the oscilloscope were not present. The Fourier transforms of the pulses in Fig A.3

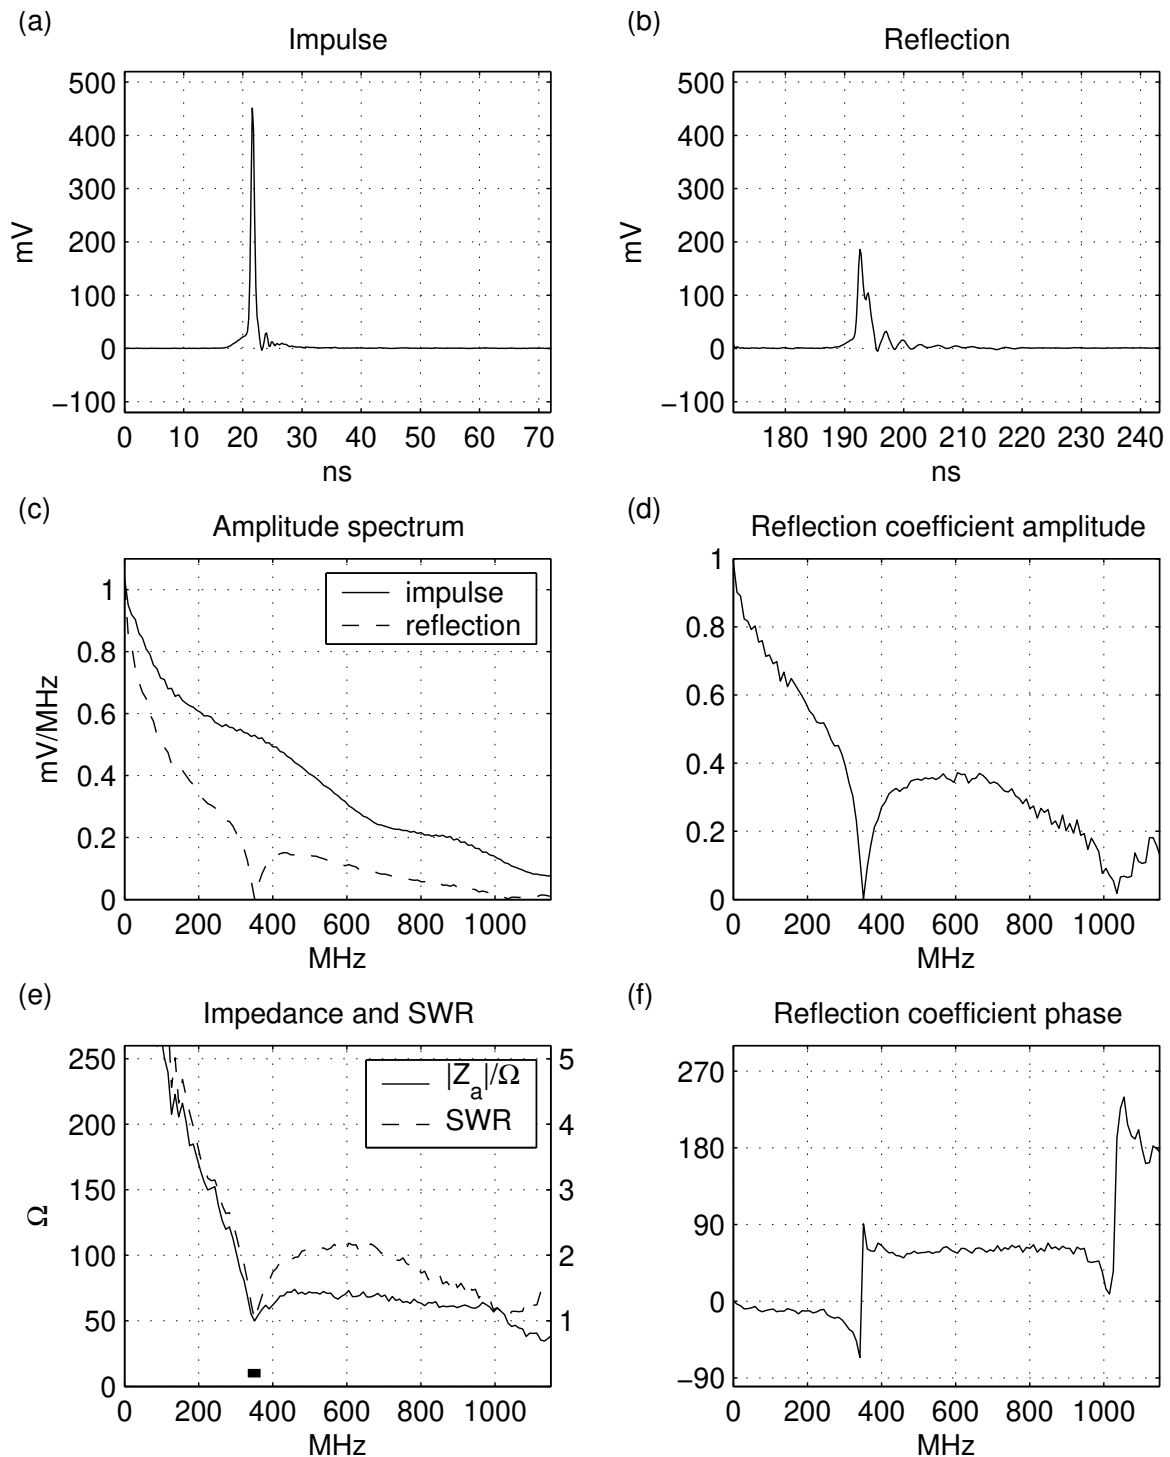


Figure A.3: Time-domain reflectometry impedance measurement of a 393 mm ($\lambda/2$ for 357 MHz) dipole: (a) shows a 60 dB μ V/MHz impulse sent to the antenna via 17 m of 50 Ω coax cable, (b) is the remaining pulse reflected back by the antenna into the cable, (c) shows the amplitude of the Fourier transform of both pulses, and (d) shows the normalized amplitude spectrum of the reflected signal relative to the original impulse. The dipole antenna completely absorbs the impulse energy at frequency 357 MHz and $3 \times 357 = 1071$ MHz. (e) shows the impedance that would cause such a reflection, and (f) shows the phase spectrum of the reflection.

(c) were calculated with a rectangular window, since the signals are broadband, transient, and vanish anyway at the window edges. The resulting amplitude spectra $A(f)$ are scaled such that with an $1 \text{ mV}_{\text{rms}}$ sine wave as input, we would get $\int_0^\infty A(f)df = 1 \text{ mV}$. The amplitude spectrum of u_p shows, for frequencies below the VHF band, a value of $1 \text{ mV/MHz} = 60 \text{ dB}\mu\text{V/MHz}$, which is exactly the level setting used on the impulse generator. The drop for higher frequencies is due to the oscilloscope amplifier having only a bandwidth of 500 MHz, combined with the fact that the spectrum of the produced pulse is not completely flat either. For frequencies over about 1.2 GHz, the attenuation is strong enough for the quantization noise of the 8-bit ADC in the oscilloscope to become noticeable in the plots of the amplitude and phase relationship of the pulses (Fig. A.3 (d) and (f)).

Figure A.3 (e) shows both the amplitude of the complex impedance, as well as the standing-wave ratio (SWR), of the antenna over the tested frequency range. Both were calculated according to equation A.38 from the pulse and reflection amplitude spectra in (c). The standing wave ratio s can be interpreted as the quotient of a real-valued (ohmic) antenna impedance Z'_a and the cable impedance Z that would cause a reflection of the same intensity as the one observed for this antenna. In other words, a termination resistor of $s \cdot 50 \Omega$ ($s \geq 1$) connected to a 50Ω transmission line causes a standing-wave ratio of s , and an ideally matched antenna has $s \approx 1$. The antenna tested here is a simple 393 mm dipole connected to a 2 m cable. With this length, it is designed to be a $\lambda/2$ dipole for about 357 MHz (a frequency optimized for a laptop eavesdropping experiment described in Section 4.1), and indeed at exactly that frequency the reflected pulse spectrum drops to almost zero, indicating that the SWR for 357 MHz signals of this dipole is about 1. The same dipole also seems to act as a $\frac{3}{2}\lambda$ dipole for $3 \times 357 \text{ MHz}$, as indicated by a second drop at that frequency.

I did not have access to equipment for verifying these impedance measurements with more conventional techniques (i.e., network analyzer) with better dynamic range, but nevertheless the results show nicely, at least in a qualitative way, the frequency dependency of the antenna impedance.

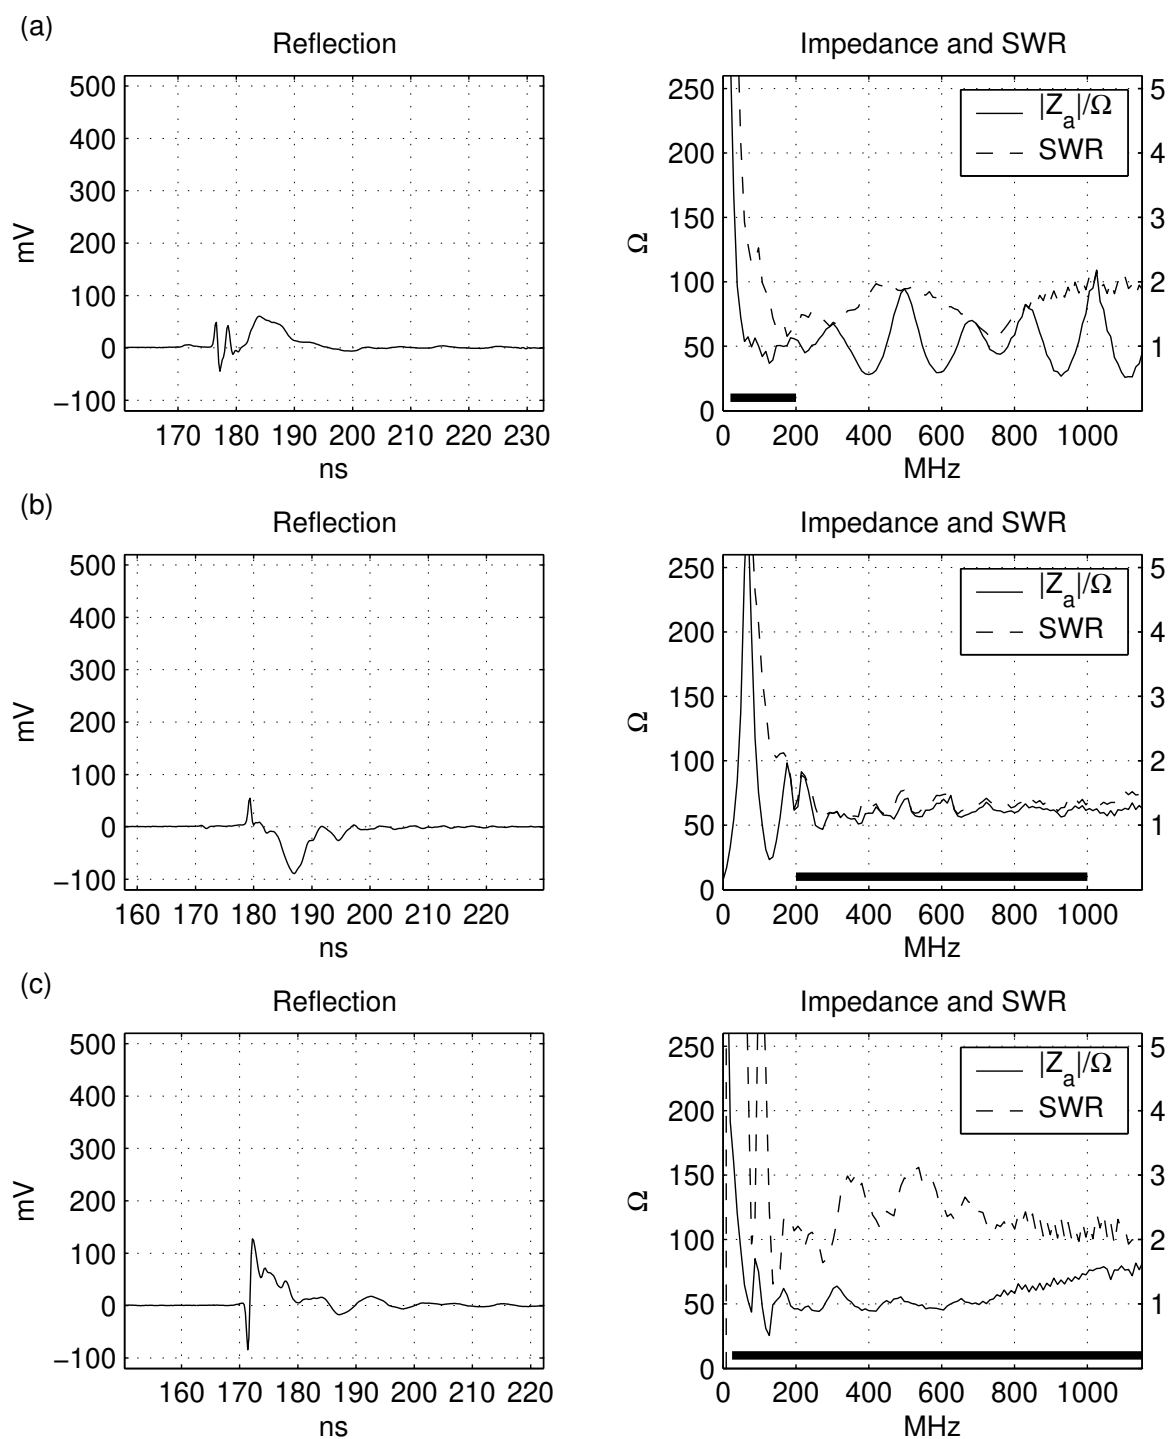


Figure A.4: Time-domain reflectometry impedance-match analysis of (a) the bi-conical 20–200 MHz and (b) log periodic 200–1000 MHz antenna in the Dynamic Sciences R-1150-10A Portable Antenna Kit, and (c) the low-cost Watson WBD-40 discone 25–1300 MHz scanner antenna.

Appendix B


Notes on experimental setups

This section collects a few practical tricks that I learned while designing some of the experiments. They may be of benefit for anyone planing to use similar setups and equipment.

B.1 Impedance-matched attenuators

RF measurement equipment commonly uses inputs and outputs with $50\ \Omega$ impedance, while video equipment manufacturers seem to prefer $75\ \Omega$. In addition to this incompatibility of impedance in connectors and cables, which can lead to distorting reflections if not handled correctly, it is also often necessary to attenuate a signal to match voltage levels.

So what we want is a simple passive 4-port resistor network that has impedances Z_1 and Z_2 on its two port pairs if the same impedances are connected to the respective ports. In addition, we want the output voltage on the second port pair $U_2 = A \cdot U_1$ to be attenuated by a factor A when voltage U_1 is applied to the first port pair. A normal voltage divider with just two resistors has not enough degrees of freedom to allow us to match all three design parameters Z_1 , Z_2 and A , therefore we have to use a π network consisting of three resistors R_1 , R_2 and R_3 instead:



The following three equations link the three design parameters with the three resistor values:

$$Z_1 = \frac{1}{\frac{1}{R_1} + \frac{1}{R_2 + \frac{1}{\frac{1}{R_3} + \frac{1}{Z_2}}}} \quad (\text{B.1})$$

$$Z_2 = \frac{1}{\frac{1}{R_3} + \frac{1}{R_2 + \frac{1}{\frac{1}{R_1} + \frac{1}{Z_1}}}} \quad (\text{B.2})$$

$$A = \frac{\frac{1}{\frac{1}{R_3} + \frac{1}{Z_2}}}{R_2 + \frac{1}{\frac{1}{R_1} + \frac{1}{Z_1}}} \quad (\text{B.3})$$

Solving for the resistor values, we get

$$R_1 = Z_1 \cdot \frac{\frac{Z_2}{Z_1} - A^2}{\frac{Z_2}{Z_1} + A^2 - 2A} \quad (\text{B.4})$$

$$R_2 = Z_1 \cdot \frac{\frac{Z_2}{Z_1} - A^2}{2A} \quad (\text{B.5})$$

$$R_3 = Z_2 \cdot \frac{\frac{Z_2}{Z_1} - A^2}{\frac{Z_2}{Z_1} + A^2 - 2A \cdot \frac{Z_2}{Z_1}} \quad (\text{B.6})$$

$$(\text{B.7})$$

and the resistor values are all non-negative if

$$0 \leq A < \begin{cases} 1 - \sqrt{1 - \frac{Z_2}{Z_1}} & \text{if } Z_1 \geq Z_2 \\ \frac{Z_2}{Z_1} \cdot \left(1 - \sqrt{1 - \frac{Z_1}{Z_2}}\right) & \text{if } Z_1 < Z_2 \end{cases} \quad (\text{B.8})$$

For example, to attenuate the 3 V signal at the $Z_1 = 50 \Omega$ output of a receiver to the 1 V required at the $Z_2 = 75 \Omega$ input of a monitor ($A = 1/3$), we have to use $R_1 = 73.5 \Omega$, $R_2 = 104.2 \Omega$ and $R_3 = 170.5 \Omega$.

B.2 Video sync signal generation

A multi-sync monitor connected to the AM output of the receiver is extremely helpful for real-time quality assessment of a received compromising video signal and the fine-tuning of parameters such as reception frequency, bandwidth, as well as antenna type and position. A deep-memory oscilloscope is an excellent post-processing tool once the signal has been found, but its one-dimensional display can visualize only a fraction of the signal recorded from an entire video frame, and software rasterization in a connected PC takes several seconds and is therefore too slow for manual fine tuning and signal searching.

Modern VGA computer monitors can synchronize to a wide and continuous range of vertical and horizontal deflection frequencies (e.g., 50–150 Hz and 30–95 kHz). The demodulator output of a receiver can be connected easily via an attenuator (see previous section) to one or more of the red/green/blue 0.7 V inputs of the monitor. However, an AM receiver does not recover the sync pulses of the eavesdropped video signal in a form suitable for driving the TTL H/V-sync inputs of a VGA monitor. The most practical solution is the use of an independent local oscillator to reconstruct sync signals with identical frequency.

I used a TTI TGA1230 30-MHz synthesized arbitrary-waveform generator. It features a 12-bit DAC output which I used for the generation of the horizontal sync signal. In addition, every sample of the stored waveform has an additional marker bit that is made available on a separate TTL/CMOS output port, which I utilized for the generation of the V-sync signal. If the total number of lines y_t , and thus the exact ratio between the horizontal and vertical sync frequencies, is known (as is the case thanks to the VESA standard modes [82] for most PCs), the use of a single function generator for the generation of both signals means that only the frequency drift of one single oscillator has to be controlled.

The TGA1230 has 64 kilosamples waveform memory, up to four separate waveforms can be linked together to a sequence, each single waveform can be repeated up to 2^{15} times, and the entire sequence can be repeated continuously. So I stored two waveforms, a normal line with a H-sync pulse at the end and a V-sync line that has in addition the marker bit set for all its samples. Both together are repeated y_t times and the sampling frequency divided by the number of samples per line is adjusted to match f_h as closely as possible.

Unfortunately, in arbitrary waveform mode, the sampling frequency of the TGA1230 can only be adjusted with 4 digits precision, which is several orders of magnitude less precise than what is required to have the $< 10^{-7}$ relative frequency error for the display of a stable image. I wrote the Perl program `gsync` that chooses the best approximation to the requested f_h and f_v that can be obtained from all feasible combinations of sampling frequency and the number of samples per line and programs the TGA1230 via the serial port appropriately. It turned out that the actual sampling frequency used was (as documented) only within one digit of the programmed frequency. In order to understand what exact frequency was generated for each setting, I reverse engineered the control algorithm for the PLL sampling frequency synthesizer of the TGA1230 (a Motorola MC145170D2 with a reference frequency $f_{\text{ref}} = 10.0000$ MHz divided by $R = 3333$ and a VCO frequency in the range 15–30 MHz divided by $N = 9999, \dots, 5003$). I found that a programmed sampling frequency $30.00 \text{ MHz} \geq f > 15.00 \text{ MHz}$ leads to a synthesized sampling frequency $f_s = f_{\text{ref}} \cdot N/3333$ where $N = \lfloor \frac{f-15 \text{ MHz}}{30 \text{ MHz}} \cdot 9999 \rfloor + 5000$; lower sampling frequencies are generated by division with the appropriate power of two.

Even with this insight into the operation of the PLL, the overall obtainable resolution for the generation of f_h was far from satisfactory and the reconstructed video signal on the eavesdropping monitor often rolled over horizontally within seconds. The solution I finally chose was the use of an additional TTI TG1010 programmable function generator. It can produce a square-wave signal up to 10 MHz with 7 digits resolution. I used it to produce a 9.990000 MHz signal that I can manually adjust with 1 Hz resolution, and I provided the TTL auxiliary output of this signal as the f_{ref} reference frequency signal to the TGA1230 generator, disabling its internal crystal oscillator. This way, I could generate a wide range of phase-locked VGA sync signals, while having control over the exact frequencies with the required seven digits of resolution, without designing a complete sync signal generator from scratch or purchasing additional hardware. The use of standard video graphics adapters turned out to be unsuitable, as their pixel-clock PLLs have an even lower resolution.

A useful extension of this setup would be blanking logic that switches off the video signal provided to the monitor near the sync pulses. A clamping circuit in the monitor depends on the video signal being at ground level shortly before or after each sync pulse, in order

to reconstruct the DC offset of the signal after a decoupling capacitor. If the supplied video signal is not blanked out during the horizontal beam flyback, the DC offset is mis-estimated by the monitor and the entire following line can appear darker as a result.

Appendix C

Glossary

AC	alternating current
ADC	analog-to-digital converter
AM	amplitude modulation
BFO	beat frequency oscillator
BW	bandwidth
CIE	Commission Internationale de l'Éclairage [International Commission on Illumination]
CISPR	Comité International Spécial des Perturbations Radioélectroniques [International Special Committee on Radio Interference]
CRT	cathode-ray tube
CW	continuous wave
DAC	digital-to-analog converter
dB	decibel
dBd	decibel gain relative to half-wavelength dipole
dBi	decibel gain relative to isotropic antenna
dBm	decibel relative to 1 mW
DC	direct current
DVI	Digital Visual Interface
EIA	Electronic Industries Alliance
EIRP	effective isotropic radiated power
EMC	electromagnetic compatibility
EUT	equipment under test
FIR	finite impulse response
FM	frequency modulation
FPD	flat-panel display
GCHQ	Government Communications Headquarters
GPIB	General Purpose Interface Bus (IEEE 488)
HF	high frequency (3–30 MHz)
IBW	impulse bandwidth
IF	intermediate frequency
IIR	infinite impulse response
ISO	International Organization for Standardization
LCD	liquid-crystal display
LED	light emitting diode
LO	local oscillator

LVDS	low voltage differential signaling (EIA-644)
NRZ	no return to zero
NSA	National Security Agency
PC	personal computer
PLL	phase-locked loop
PMT	photomultiplier tube
pp	peak to peak
RF	radio frequency
RFI	radio-frequency interference
rms	root mean square
SNR	signal-to-noise ratio
SWR	standing wave ratio
TMDS	transition minimized differential signaling
TTL	transistor-transistor logic
UHF	ultra high frequency (300–3000 MHz)
VCO	voltage controlled oscillator
VESA	Video Electronics Standards Association
VGA	video graphics adapter
VHF	very high frequency (30–300 MHz)

Index

- A/m, 146
- ADC, *see* analog-to-digital converter
- ADCS, *see* average depth of correct symbol
- air-gap security, 15
- AM, *see* amplitude modulation
- Ampère, André Marie, 143
- amplitude modulation, 24, 40, 57
- AMSG 720B, 12
- analog interface, 67
- analog-to-digital converter, 23
- Anderson, Ross, 4, 15
- anechoic chamber, 17, 100, 102
- angular resolution, 105
- antenna, 19–23, 149–151
 - active, 20, 21, 100
 - array, 95
 - bi-conical, 20, 21
 - broad-band, 19, 20
 - dipole, 19, 20, 150
 - directional gain, 17, 19
 - discone, 20, 21
 - effective aperture, 150
 - effective length, 23, 150
 - EMC measurement, 20, 23
 - factor, 23, 51, 100
 - gain, 90, 95, 150
 - horn, 21
 - impedance, 21, 151
 - isotropic, 91, 93, 95, 150
 - log-periodic, 20, 21, 47, 70
 - narrow-band, 19, 150
 - TV, 19
 - Yagi-Uda, 19, 95
- antenna-input voltage, 51
- anti-aliasing, 61, 63, 64
- aperture, 105, 120, 122
- ATM, 12
- attenuation, 64, 71, 91
 - building material, 71, 93, 99
 - free-space, 64, 93
 - mains network, 94
- attenuator, 26
 - impedance-matched, 44, 157
- audio tone, 38
- auto-correlation, 71
- automatic gain control, 26
- average depth of correct symbol, 55
- averaging, 46, 47, 54, 89, 96, 97, 121
- background light, 121, 125
- bandwidth, 24, 25, 28, 30, 31, 37, 45, 47, 51, 53,
54, 57, 68, 95, 98, 100
 - 3 dB, 28, 44
 - impulse, 28–32, 51, 90
- battle-field phones, 11
- beat-frequency oscillator, 26
- bel, 146
- BFO, *see* beat-frequency oscillator
- bit error rate, 123
- Briol, Roland, 12
- broadband signals, 19
- BSI, *see* German Information Security Agency
- building material, *see* attenuation
- bus lines, 103

- cable, 11, 12, 89, 148
- calibration, 27, 29, 51, 60
- candela, 120
- carrier frequency, 24
- cathode-ray tube, 57, 58, 67, 102, 105, 106
- character cell, 54
- character recognition, 56, 97
- charge density, 143
- CIA, 14
- cipher machine, 11
- CISPR, 87
- CISPR 16, 88
- CISPR 22, 87, 101
- clock recovery, 125
- common-mode current, 42, 70, 88, 147, 148
- compromising emanations, *see* emanations
- conducted interference, 89
- conductivity, 144
- conductor, 146
- continuous wave, 26
- contrast, 74, 75
- convolution, 27, 41, 62
- convolution theorem, 41, 42, 80
- correlation, 33
- countermeasures, 10, 11, 14, 37, 51, 60, 126
- covariance matrix, 34
- cover image, 58
- cross-correlation, 12, 35, 70, 76, 89, 103
- cross-talk, 44
- CRT, *see* cathode-ray tube
- cryptographic keys, 13, 83
- current probe, 148

- DAC, *see* digital-to-analog converter
data remanence, 13, 116
dB, 146
dB/m, 23
dBi, 95
dBm, 146
dB μ V, 146
dB μ V/m, 23, 146
dB μ V/MHz, 32
decay, 108, 110–112, 115
decibel, 146
deconvolution, 107, 119, 131
deflection coils, 102
demodulator, *see* detector
detector
 AM, 24, 25, 30, 35, 131
 BFO, 26
 FM, 26
 linear AM, 26
 logarithmic AM, 26
 non-tunable, 24
 peak, 102
 QAM, 131
 quasi-peak, 88, 89, 102
 tunable, 24
 vector, 131
DFP, *see* Digital Flat Panel
difference matrix, 56
differential power analysis, 14
differential-mode current, 42, 70, 147, 148
diffraction, 93, 105, 120
diffuse reflection, 106, 122
Digital Flat Panel standard, 77
digital interface, 67
digital oscilloscope, 26, 97, 110, 111, 158
digital rights management, 13
digital signal processing, 23, 89, 96, 131
Digital Visual Interface standard, 77, 82
digital-to-analog converter, 42, 44
dipole, *see* antenna, 68
display adapter, *see* graphics card
distance, 70, 98, 99, 101, 105, 120–122, 124, 129
dither pattern, 53, 57, 60
dithering, 58
driver, 35, 42
DVI, *see* Digital Visual Interface
Dynamic Sciences, 21–23, 25
- eavesdropping protection, 51
edges, 35
EIRP, 51, 70, 75, 101
electric current density, 144
electric field, 145
electric flux density, 143
electrical field strength, 144
electromagnetic compatibility, 16, 37, 86, 87
electromagnetic immunity, 132
electromagnetic interference, 70
electromechanic devices, 101, 134
electron beam, 38, 45, 53, 63
electron-beam current, 106
emanations
 acoustic, 12
 broadband, 19
 compromising, 9, 10, 12, 31, 85, 133
 low-frequency, 42, 101, 134
 mechanical, 132
 optical, 16, 105
 power-line, 89, 101, 131, 134
 software-controlled, 14, 15, 38, 60, 65, 75, 102, 123
 video, 42, 45, 57, 102
embedded image, 58, 60
embedded information, 57
EMC, *see* electromagnetic compatibility
emission limits, 11, 86, 100, 102
emission security, 11, 86, 89
EMSEC, *see* emission security
EN 55022, 87
energy flux density, 150
entropy, 55
error rate, 55
ESL 400, 4
- f_h , 38
 f_p , 38, 40
 f_v , 38
fall time, 35, 44
far field, 42, 150
Faraday, Michael, 143
FCC, 87
fiber-optic cables, 12
filter, 27, 29
 Butterworth, 108
 coefficients, 62
 color, 125
 finite-impulse-response, 62
 infinite-impulse-response, 97
 intermediate-frequency, 25–28
 low-pass, 61–63, 110
 matched, 123
 notch, 89
 red/black separation, 11
 Wiener, 24
firmware, 60
flat-panel displays, 68
fluorescent light, 126
flux density, 51
Fondazione Ugo Bordoni, 12
font, 54, 60
font renderer, 60
font size, 53, 97
Fourier analysis, 19, 28, 31, 41, 42, 80, 116
FPD-Link, 70, 75, 77
frame, 38, 45–48, 70, 96
frame buffer, 67, 82

- Freedom of Information Act, 11, 86
French embassy, 11
frequency domain, 41
frequency modulation, 132
Fresnel ellipsoid, 93
FS222, 11
FSET7/FSET22, 26
- gamma correction, 60, 106
GCHQ, 11
German army, 11
German Information Security Agency, 12, 36
glyphs, 60
golfball typewriter, 132
GPIB, 131
graphics card, 42, 43, 46, 67
- half-toning, 57
Hamming window, 62
hard-disk controller, 123
HDCP, 83
high-energy electromagnetic waveforms, 132
hinting, font, 60, 61, 63
horizontal deflection frequency, 38, 45, 46
horizontal resolution, 68
- IBM PC, 67
IBM Selectric typewriter, 132
illuminance, 120
impedance, 67, 110, 148, 150, 157
impulse, 21, 22, 24, 28, 31, 33, 98
 broadband, 98
 Dirac, 41
 generator, 25, 27, 32, 151
 rectangular, 31, 41, 42
 response, 27, 114, 119, 127
 series, 41
 source, 31
 spectral density, 32
 strength, 32
 voltage, 100
- indium-tungsten oxide, 68
information hiding, 57, 130
Inslaw, 14
intentional broadcast, 14, 15, 38, 75, 102, 123
inter-character interference, 54
interface standards, 68
intermediate frequency, 25–27
Internet, 133
irradiance, 122
ISO 31, 143
ITU-R P.372, 91
- jamming, 12, 51, 65, 75, 126
- key exchange protocol, 83
keyboard, 35, 105, 123, 126, 133
Koch, Egmont, 14
- Lambertian surface, 122
- laptops, 68, 75, 77
laser diode, 102
laser printer, 102, 131
least-significant bits, 64
LED, 123, 124
light, 106
light sensor, 127
line impedance stabilization network, 88
line of sight, 105, 126
liquid-crystal display, 68, 127
local oscillator, 24, 25
logarithmic scales, 146
low voltage differential signaling, 70
lumen, 120
luminous intensity, 120
lux, 120
LVDS, *see* low voltage differential signaling
- magnetic field, 145
magnetic field strength, 144, 146
magnetic induction, 144
magnetic stripe, 12
magnetic vector potential, 147
mains network, 89, 94
MATLAB, 131
MAX233, 132
Maxwell's equations, 19, 143–145
Maxwell, James Clark, 143
mechanical side-channel, 132
MI5, 11
MIL-STD-461E, 101
military, 10
monospaced fonts, 54
mouse, 132
multi-sync monitor, 158
- NACSIM 5000, 11
NACSIM 5100A, 11, 25
NAG1A, 11
National Security Agency, 9, 15, 25, 29, 55, 101
near field, 42, 150
near-field probe, 69
neper, 146
noise, 24, 51, 65, 91, 98
 atmospheric, 92
 factor, 90, 91
 figure, 91
 from background light, 121
 galactic, 92
 mains power, 94
 man-made, 92
 quantization, 110
 receiver, 99
 thermal, 110
non-linearity, 58
NRZ, 125
NSTISSAM TEMPEST/1-92, 11, 29, 55, 86
Nyquist, Harry, 40

- optical character recognition, 54
- optical eavesdropping, 105
- order statistic, 55
- oscilloscope, 45, 47
- overflow, 58

- P22 phosphor, 107, 112
- P&D, *see* Plug & Display
- PanelLink, 77
- password, 55
- pattern recognition, 54
- PCI, 13
- peak voltage, 63
- permeability, 145
- permittivity, 145
- phase-locked loop, 25, 45, 68, 132, 159
- phosphor, 106–112, 115, 119, 124, 127
- phosphor type registry, 107
- photo sensor, 109, 111
- photodiode, 109
- photomultiplier, 109, 110
- photons, 121
- PIN, 12, 133
- pixel, 37, 38, 40–42, 47, 53, 54, 63, 70, 105
 - pulse shape, 42–44
- pixel-clock frequency, 26, 37–39, 44, 47, 53, 67, 98
- plasma displays, 68
- Plug & Display standard, 77
- power, 146
- power analysis, 14
- Poynting vector, 145
- pressure sensor, 132
- primary color, 67
- printer, 12, 101–103, 131, 132
- PROMIS, 14, 16
- protection technique, 60
- protective color combination, 75
- pull-up resistor, 35
- pulse strength, 27
- pulse stretch, 26

- quantization noise, 40
- quasi-peak detector, *see* detector

- R-1250 receiver, 25, 27, 45, 47, 132
- radiant exitance, 122
- radiant intensity, 111, 120, 122
- radiant sensitivity, 110, 111
- radio transmitters, 52, 88, 89, 98
- RAMDAC, *see* digital-to-analog converter
- random sequence, 76
- raster graphic, 47
- raster scan, 37
- Rayleigh criterion, 105
- receiver, 24
 - AM, 24, 64, 131
 - broadband, 16
 - low-cost, 132
 - superheterodyne, 24, 26, 27
- red/black separation, 11
- redundancy, 42
- reflection, 149
- refresh, 45, 82
- return-to-zero, 42
- Riconosciuto, Michael J., 14, 15
- rise time, 35, 44, 110
- Rohde & Schwarz, 26, 100
- RS-232 cables, 12, 124, 132
- RZ, *see* return-to-zero

- sampling, 41, 61
- sampling theorem, 40
- secure, unconditionally, 33
- security policies, 105
- selective refresh, 82
- SEPI, 12
- serial line, 12, 124, 132
- serial links, 12
- shielding, 11, 16, 17, 37, 83, 101, 103
- shortwave, 40
- shot noise, 121, 123
- signal-to-noise ratio, 23, 40, 90, 97, 121
- slideback, 26
- smartcards, 13–14
- Smulders, Peter, 12
- software protection, 60
- software radio, 131
- solid angle, 111, 117, 120
- sound pressure level, 146
- spatial frequency, 57, 60
- spectral density, 32, 51
- spectrograph, 125
- spectrum, 24, 31, 43, 44, 60, 80
- spectrum analyzer, 28, 37, 43, 44, 51, 64, 99, 102
- spread-spectrum modulation, 15, 76
- Spycatcher, 11
- SRAM, 13
- standards, 68
 - emission-security, 86, 89, 97, 103, 133
 - ergonomic, 87
 - military EMC, 101
 - radio interference, 87, 101
- standing wave ratio, 151
- start bit, 76
- statistical independence, 33
- steganography, *see* information hiding
- Steiglitz-McBride iteration, 111
- steradian, 120
- sync pulses, 70
- sync signals, 67
- sync-pulse generator, 12, 45, 159
- synchronization, 76, 102

- tamper-resistant modules, 13–14
- TCO'92, 87
- Teapot, 15, 16, 76
- telegrapher's equation, 148

- telescope, 105, 106
- telex cable, 11
- Tempest, 10, 11, 16, 25, 29, 36, 55, 85, 86, 132–134
- test images, 110
- threat, 9, 12, 16, 86, 97, 133
- time domain, 41
- time-domain reflectometry, 151
- TMDS, 77, 79, 80, 82
- transistor-transistor logic, 67
- transmission format, 71
- transmission line, 70, 148
- trigger, 47
- Trojan Horse, 15
- TTL, *see* transistor-transistor logic
- TV set, 12, 16, 37, 42, 57
- twisted pair, 70, 77, 148
- typewriters, 132

- units, 146

- V/m, 146
- van Eck, Wim, 12, 16, 37, 42
- VAX, 15
- VDE, 87
- vertical deflection frequency, 38, 45–47
- VESA, 38, 39, 77
- VGA, *see* Video Graphics Adapter
- VGA connector, 44, 67
- video
 - cable, 42, 67
 - connector, 67
 - controller, 70
 - interface, 67, 70
 - mode, 38, 39, 46, 47, 68
 - signal, 37, 40–43, 51, 53, 60, 64, 71, 106
 - spectrum, 42
 - terminal, 42, 54, 67
 - timing, 39
- Video Graphics Adapter, 67
- virus, 15, 133

- wavelength, 42
- Wiener filter, 24
- WLAN, 133
- Wright, Peter, 11
- WTDS, 107

- x_d , 38
- x_t , 38
- X11 Window System, 47

- y_d , 38
- y_t , 38
- Yagi, *see* antenna
- Young, John, 11

- zinc sulfide, 107, 108