CHAPTER

17

Emission Security

The hum of either army stilly sounds,
That the fixed sentinels almost receive
The secret whispers of each others' watch;
Fire answers fire, and through their paly flames
Each battle sees the other's umbred face.

— William Shakespeare, King Henry V, Act IV

17.1 Introduction

Emission security, or Emsec, is about preventing attacks using *compromising emanations*, namely conducted or radiated electromagnetic signals. It has many aspects. Military organizations are greatly concerned with *Tempest* defenses, which prevent the stray RF emitted by computers and other electronic equipment from being picked up by an opponent and used to reconstruct the data being processed. Tempest has recently become an issue for electronic voting too, after a Dutch group found they could tell at a distance which party a voter had selected on a voting machine. The smartcard industry has been greatly exercised by *power analysis*, in which a computation being performed by a smartcard — such as a digital signature — is observed by measuring the current drawn by the CPU and the measurements used to reconstruct the key. These threats are closely related, and have a number of common countermeasures. Researchers have also discovered attacks that exploit stray optical, thermal and acoustic emanations from various kinds of equipment. Such techniques are also referred to as *side channel* attacks as the information is leaking through a channel other than those deliberately engineered for communication.

People often underestimate the importance of Emsec. However, it seems that the world's military organizations spent as much on it as on cryptography during the last quarter of the twentieth century. In the commercial world, the uptake of smartcards was materially set back in the last few years of that century by the realization that all the smartcards then on the market were extremely vulnerable to simple attacks which required the attacker only to trick the customer into using a specially adapted terminal that would analyze the current it drew during a small number of transactions. These attacks did not involve penetrating the card and thus might leave no trace. Once fielded, they were very much cheaper than probing attacks, and potentially allowed large-scale card-cloning attacks against an unsuspecting cardholder population.

Electromagnetic eavesdropping attacks have been demonstrated against other commercial systems, including automatic teller machines. They can interact with malware, in that rogue software can cause a computer to emit a stronger signal than it normally would, and even modulate the signal so as to get stolen data past a corporate firewall. There has also been alarm about disruptive electromagnetic attacks, in which a terrorist group might use a high-energy microwave source to destroy the computers in a target organization without killing people. (I'll discuss these in more detail in the chapter on electronic warfare.)

Both active and passive Emsec measures are closely related to *electromagnetic compatibility* (EMC) and *radio frequency interference* (RFI), which can disrupt systems accidentally. If you fly regularly, you'll be familiar with the captain saying something like 'All electronic devices must be switched off now, and not switched on again until I turn off the seat belt sign'. This problem is getting worse as everything becomes electronic and clock frequencies go up. And how do you obey the captain now that more and more devices are 'always on' — so that the 'off' switch only turns off the green tell-tale light?

As more and more everyday devices get hooked up to wireless networks, and as processor speeds head up into the gigahertz range, all these problems — RFI/EMC, Emsec and various electronic warfare threats — are set to get worse.

17.2 History

Crosstalk between telephone wires was well known to the pioneers of telephony in the 19th century, whose two-wire circuits were stacked on tiers of crosstrees on supporting poles. One way of dealing with it was to use 'transpositions', in which the wires were crossed over at intervals to make the circuit a twisted pair. This problem appears to have first come to the attention of the military during the British Army expedition to the Nile and Suakin in 1884–85 [923].

The first known exploitation of compromising emanations in warfare was in 1914. Field telephone wires were laid to connect the troops bogged down in the mud of Flanders with their headquarters, and these often ran for miles, parallel to enemy trenches that were only a few hundred yards away. A phone circuit was a single-core insulated cable, which used earth return in order to halve the weight and bulk of the cable. It was soon discovered that earth leakage caused a lot of crosstalk, including messages from the enemy side. Listening posts were quickly established and protective measures were introduced, including the use of twisted pair cable. By 1915, valve amplifiers had extended the earth leakage listening range to 100 yards for telephony and 300 yards for Morse code. It was found that the tangle of abandoned telegraph wire in no-man's land provided such a good communications channel, and leaked so much traffic to the Germans, that clearing it away become a task for which lives were spent. By 1916, earth return circuits had been abolished within 3000 yards of the front. When the USA joined the war, the techniques were passed on to them [869, 923].

The Second World War brought advances in radar, passive direction finding and low-probability-of-intercept techniques, which I'll discuss in the chapter on Electronic Warfare. By the 1960s, the stray RF leaking from the local oscillator signals in domestic television sets was being targeted by direction-finding equipment in 'TV detector vans' in Britain, where TV owners must pay an annual license fee to support public broadcast services. Some people in the computer security community were also aware that information could leak from cross-coupling and stray RF. The earliest published reference appears to be a 1970 Rand Corporation report written by Willis Ware [1319].

The intelligence community also started to exploit side channel attacks. During the Suez crisis in 1956, the British figured out the settings of the Egyptian embassy's Hagelin cipher machine using a phone bug. In 1960, after the Prime Minister ordered surveillance on the French embassy during negotiations about joining the European Economic Community, his security service's scientists noticed that the enciphered traffic from the embassy 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 [1363]. This is more common than one might suppose; there has been more than one case of a cipher machine broadcasting in clear on radio frequencies (though often there is reason to suspect that the vendor's government was aware of this).

During the 1970s, emission security became a highly classified topic and vanished from the open literature. It came back to public attention in 1985 when Wim van Eck, a Dutch researcher, published an article describing how he had managed to reconstruct the picture on a VDU at a distance using a modified TV set [408]. The revelation that Tempest attacks were not just

feasible, but could be mounted with simple home-built equipment, sent a shudder through the computer security industry.

Published research in emission security and related topics took off in the second half of the 1990s. In 1996 Markus Kuhn and I reported that many smartcards could be broken by inserting transients, or *glitches*, in their power or clock lines [68]. Paul Kocher also showed that many common implementations of cryptosystems could be broken by making precise measurements of the time taken [727]. In 1998 Kuhn and I showed that many of the compromising emanations from a PC could be made better, or worse, by appropriate software measures [753]. In 1998–9, Kocher showed that crypto keys used in smartcards could be recovered by appropriate processing of precise measurements of the current drawn by the card [728]. Although smartcard vendors had been aware of a possible problem from the late 1980s, Kocher's *differential power analysis* provided a simple and powerful signal processing technique for recovering data that completely overwhelmed the somewhat primitive defences that the industry had seen fit to provide.

In recent years, results have followed steadily. 2002 brought results on optical leakage: Markus Kuhn showed that a VDU's screen contents can be recovered optically, even from diffuse light reflected off room walls or the operator's face [750], while Joe Loughry and David Umphress also found serial port data in many of the LED status indicators on data serial lines [815]. In 2004, Dmitri Asonov and Rakesh Agrawal showed that the different keys on a keyboard made sufficiently different sounds that someone's typing could be picked up from acoustic emanations [91]; in 2005, Li Zhuang, Feng Zhou, and Doug Tygar improved this to use keyboard characteristics and text statistics to decipher a recording text typed for ten minutes on a random keyboard, to which there had been no previous access to train the recognition software [1376]. In 2006, Steven Murdoch showed that many computers reveal their CPU load via thermal leakage; clock skew is a function of ambient temperature, and can be measured remotely. He hypothesised that it might even be used to work out a target machine's latitude and longitude [914]. These results just seem to keep on coming.

17.3 Technical Surveillance and Countermeasures

Before getting carried away with high-tech toys such as Tempest monitoring receivers, we ought to stop and think about bugs. The simplest and most widespread attacks that use the electromagnetic spectrum are not those exploiting some unintended feature of innocuous equipment, but those in which a custom-designed device is introduced by the attacker.

No matter how well it is protected by encryption and access controls while in transit or storage, most highly confidential information originally comes

into being either as speech or as keystrokes on a PC. If it can be captured by the opponent at this stage, then no subsequent protective measures are likely to help very much.

So an extraordinary range of bugs is available on the market:

- At the low end, a few tens of dollars will buy a simple radio microphone that you can stick under a table when visiting the target. Battery life is the main constraint on these devices. They typically have a range of only a few hundred yards, and a lifetime of days to weeks.
- At the next step up are devices that draw their power from the mains, a telephone cable or some other external electricity supply, and so can last indefinitely once emplaced. Some are simple microphones, which can be installed quickly in cable ducting by an adversary who can get a few minutes alone in a room. Others are inserted from a neighboring building or apartment by drilling most of the way through a wall or floor. Yet others look like electrical adaptors but actually contain a microphone, a radio transmitter and a TV camera. Others monitor data — for example a Trojan computer keyboard with bugging hardware contained in the cable connector.
- Many modern bugs use off-the-shelf mobile phone technology. They can be seen as slightly modified cellphone handsets that go off-hook silently when called. This gives them worldwide range; whether they last more than a week or so depends on whether they can be connected to a power source when installed.
- One exotic device, on show at the NSA Museum in Fort Meade, Md., was presented to the U.S. ambassador in Moscow in 1946 by a class of schoolchildren. It was a wooden replica of the Great Seal of the United States, and the ambassador hung it on the wall of the office in his residence. In 1952, it was discovered to contain a resonant cavity that acted as a microphone when illuminated by microwaves from outside the building, and retransmitted the conversations that took place in his office. Right up to the end of the Cold War, embassies in Moscow were regularly irradiated with microwaves, so variants of the technique presumably remained in use.
- Laser microphones work by shining a laser beam at a reflective or partially reflective surface, such as a window pane, in the room where the target conversation is taking place. The sound waves modulate the reflected light, which can be picked up and decoded at a distance.
- High-end devices used today by governments, which can cost upwards of $10,000, use low-probability-of-intercept radio techniques such as frequency hopping and burst transmission. They can also be turned on and off remotely. These features can make them much harder to find.

- People constantly come up with creative new ideas. A recent one is the *jitterbug* which you put in a keyboard cable. It modulates keystroke data, such as passwords, into sub-perceptible keystroke delays. This means that a password you type can be more easily guessed by an attacker who wiretaps your connection, even if it's encrypted [1155].

A number of countermeasures can give a fair degree of protection against such attacks, provided they are used by skilled and experienced experts.

- The *nonlinear junction detector* is a device that can find hidden electronic equipment at close range. It works because the transistors, diodes and other nonlinear junctions in electronic equipment rectify incident RF signals. The nonlinear junction detector broadcasts a weak radio signal and listens for odd harmonics. It can detect unshielded electronics at a range of a few feet. However, if the bug has been planted in or near existing electronic equipment, then the nonlinear junction detector is not much help. There are also expensive bugs designed not to re-radiate at all. A variant was invented by the investigative journalist Duncan Campbell in the early 1970s to detect telephone taps: the amplifier used at that time by the security services re-radiated harmonics down the line. Following a raid on his house, the plans for this device were seized; it was then 'invented' in a government laboratory, and credited to a government scientist.
- There are a number of *surveillance receivers* on the market. The better ones sweep the radio spectrum from about 10 KHz to 3 GHz every few tens of seconds, and look for signals that can't be explained as broadcast, police, air traffic control and so on. (Above 3GHz, signals are so attenuated by building materials, and device antennas can be so directional, that general spectrum search is no longer as effective as nonlinear junction detectors and physical searching.) Contrary to popular belief, some low-probability-of-intercept techniques do not give complete protection. Direct sequence spread spectrum can be spotted from its power spectrum, and frequency hoppers will typically be observed at different frequencies on successive sweeps. Burst transmission does better. But the effectiveness of surveillance receivers is increasingly limited by the availability of bugs that use the same frequencies and protocols as legitimate mobile or cordless phones. Security conscious organizations can always try to forbid the use of mobiles, but this tends not to last long outside the military. For example, Britain's parliament forbade mobiles until 1997, but the rule was overturned when the government changed.
- Breaking the line of sight, such as by planting trees around your laboratory, can be effective against laser microphones but is often impractical.

- Some facilities at military organizations are placed in completely shielded buildings, or underground, so that even if bugs are introduced their signals can't be heard outside [87]. This is very expensive; but if you can confine your secret conversations to a single room, then there are vendors who sell prefabricated rooms with acoustic and electromagnetic shielding. Another option is to ensure that devices such as wire-line microphones aren't installed in the building when it's constructed, that there are frequent sweeps, and that untrusted visitors (and contractors such as cleaning staff) are kept out of the most sensitive areas. But this is harder than it looks. A new U.S. embassy building in Moscow had to be abandoned after large numbers of microphones were found in the structure, and Britain's counterintelligence service had to tear down and rebuild a large part of a new headquarters building, at a cost of about $50m, after an employee of one of the building contractors was found to have past associations with the Provisional IRA.

The tension here is between technological defenses, which can be effective but very expensive, and procedural controls, which are cheap but tedious.

All that said, technological developments are steadily making life easier for the bugger and harder for the defense. As more and more devices acquire intelligence and short-range radio or infrared communications — as 'things that think' become 'things that chatter' — there is ever more scope for attacks via equipment that's already in place rather than stuff that needs to emplaced for the purpose. For example:

- The risks associated with telephones are much more than many people would like to believe. More and more people use cordless phones for convenience, and forget that they're easy to eavesdrop. Phones can be doctored so that they'll go off-hook under remote control; some digital phones have such a facility already built into them (and it's said that some countries make this a condition of import licensing). Also, some makes of PBX can be reprogrammed to support this kind of surveillance.
- The typical laptop computer has a microphone that can be switched on under software control, and is increasingly likely to be online from time to time. An attacker can infect it with malware that listens to conversations in the room, compresses them, encrypts them and mails them back to its creator.
- The NSA banned Furby toys in its buildings, as the Furby remembers (and randomly repeats) things said in its presence.

But there are many more ways in which existing electronic equipment can be exploited by an adversary.

17.4 Passive Attacks

We'll first consider passive attacks, that is, attacks in which the opponent makes use of whatever electromagnetic signals are presented to him without any effort on his part to create them. I'll exclude optical signals for now, although light is electromagnetic; I'll discuss them along with acoustic attacks later.

Broadly speaking, there are two categories of electromagnetic attack. The signal can either be conducted over some kind of circuit (such as a power line or phone line), or it may be radiated as radio frequency energy. These two types of threat are referred to by the military as 'Hijack' and 'Tempest' respectively. They are not mutually exclusive; RF threats often have a conducted component. For example, radio signals emitted by a computer can be picked up by the mains power circuits and conducted into neighboring buildings. However it's a reasonable working classification most of the time.

17.4.1 Leakage Through Power and Signal Cables

Since the nineteenth century, engineers have been aware that high-frequency signals leak everywhere and that care is needed to stop them causing problems, and, as I noted, the leakage has been exploited for military purposes since in 1914. Conducted leakage of information can be largely suppressed by careful design, with power supplies and signal cables suitably filtered and suppressed. This makes up a significant part of the cost difference between otherwise comparable military and civilian electronics.

17.4.1.1 Red/Black Separation

Red equipment (carrying confidential data such as plaintext) has to be isolated by filters and shields from *black* equipment (that can send signals directly to the outside world). Equipment with both 'red' and 'black' connections, such as cipher machines, is particularly difficult to get right. It's made more expensive by the fact that the standards for emission security, such as the NACSIM 5100A that specifies the test requirements for Tempest protected equipment, and its NATO equivalent AMSG 720B, are classified [1098] (though they've leaked, as I'll discuss in section 17.4.2 later).

So properly shielded equipment tends to be available only in small quantities, and made specifically for defense markets. This makes it extremely expensive. However, the costs don't stop there. The operations room at an air base can have thousands of cables leading from it; filtering them all,

and imposing strict enough configuration management to preserve red/black separation, can cost millions.

17.4.1.2 Timing Analysis

In 1996, Paul Kocher showed that many implementations of public-key algorithms such as RSA and DSA leaked key information through the amount of time they took [727]. His idea was that when doing exponentiation, software typically steps through the secret exponent one bit at a time, and if the next bit is a one it does a multiply. This enables an opponent who knows the first b bits of an exponent to work out the $b+1$-st bit by observing a number of exponentiations. Attack and defence coevolved for a while, and many people thought their implementations were secure if they used the Chinese Remainder Theorem. But in 2003, David Brumley and Dan Boneh implemented a timing attack against Apache using OpenSSL, and showed how to extract the private key from a remote server by timing about a million decryptions [233]. Good implementations of public-key algorithms now use blinding to prevent such attacks (OpenSSL did offer blinding as an option, but Apache didn't use it).

John Kelsey, Bruce Schneier, David Wagner and Chris Hall pointed out in 1998 that block ciphers using large S-boxes, such as AES, could be vulnerable to timing attacks based on cache misses [704]. The attacker can verify guesses about the output of the first round of the cipher by predicting whether the guessed value would cause a cache miss on S-box lookup, and verifying this against observation. A number of researchers have improved this attack since then, and nowadays a naïve implementation of AES can be broken by observing a few hundred encryptions [999, 164, 994].

17.4.1.3 Power Analysis

Often people aren't aware of the need to filter signals until an exploit is found. A very important example comes from the discovery of power attacks on smartcards. As a smartcard is usually a single silicon chip in a very thin carrier, there is little scope for filtering the power supply using extra components such as chokes and capacitors — and given that a smartcard costs 50¢–$1 in bulk, while a capacitor would cost 10¢ to fit, it's rarely economic. The power supply may also be under the control of the enemy. If you use your bank smartcard to make a purchase in a Mafia-owned store, then the terminal might have extra electronics built into it to cheat you.

By the early 1990s, it appears to have been known to pay-TV hackers and to some government agencies that a lot of information could be gathered about the computations being performed in a smartcard by simply measuring the

current it drew from its power supply. This attack, known as *power analysis* or *rail noise analysis*, may involve as little as inserting a 10Ω resistor in the ground line and connecting a digital storage oscilloscope across it to observe fluctuations in the current drawn by the device. An example of the resulting power trace can be seen in Figure 17.1. This shows how a password can be extracted from a microcontroller by guessing it a byte at a time and looking for a different power trace when the correct byte is guessed.

Different instructions have quite different power consumption profiles, and, as you can see, the power consumption also depends on the data being processed. The main data-dependent contribution in many circumstances is from the bus driver transistors, which are quite large (see the top of Figure 16.6). Depending on the design, the current may vary by several hundred microamps over a period of several hundred nanoseconds for each bit of the bus whose state is changed [877]. Thus the Hamming weight of the difference between each data byte and the preceding byte on the bus (the *transition count*) is available to an attacker. In some devices, the Hamming weight of each data byte is available too [881]. EEPROM reads and writes can give even more substantial signals.

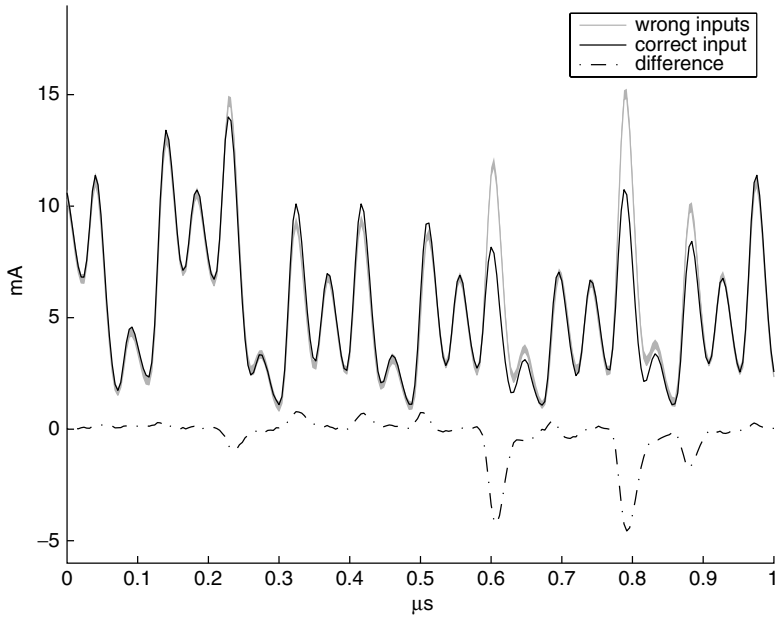

Figure 17.1: Plot of the current measured during 256 single attempts to guess the first byte of a service password stored in the microcontroller at the heart of a car immobilizer (courtesy of Markus Kuhn and Sergei Skorobogatov).

The effect of this leakage is not limited to password extraction. An attacker who understands (or guesses) how a cipher is implemented can obtain significant information about the card's secrets and in many cases deduce the value of the key in use. It is particularly significant because it is a noninvasive attack, and can be carried out by suitably modified terminal equipment on a smartcard carried by an unsuspecting customer. This means that once the attacker has taken the trouble to understand a card and design the attack, a very large number of cards may be compromised at little marginal cost.

The threat posed to smartcards by power analysis was brought forcefully to the industry's attention in 1998 with the development of an efficient signal processing technique to extract the key bits used in a block cipher such as DES from a collection of power curves, without knowing any implementation details of the card software. This technique, Paul Kocher's *differential power analysis*, works as follows [728].

The attacker first collects a number of curves (typically several hundred) by performing known transactions with the target card — transactions for which the encryption algorithm and either the plaintext or the ciphertext is known. She then guesses some of the internal state of the cipher. In the case of DES, each round of the cipher has eight table look-ups in which six bits of the current input is exclusive-or'ed with six bits of key material, and then used to look up a four-bit output from an S-box. So if it's the ciphertext to which the attacker has access, she will guess the six input bits to an S-box in the last round. The power curves are then sorted into two sets based on this guess and synchronized. Average curves are then computed and compared. The difference between the two average curves is called a *differential trace*.

The process is repeated for each of the 64 possible six-bit inputs to the target S-box. It is generally found that the correct input value — which separates the power curves into two sets each with a different S-box output value — will result in a differential trace with a noticeable peak. Wrong guesses of input values, however, generally result in randomly sorted curves and thus in a differential trace that looks like random noise. In this way, the six keybits which go to the S-box in question can be found, followed by the others used in the last round of the cipher. In the case of DES, this gives 48 of the 56 keybits, and the remainder can be found trivially by exhaustive search. If the cipher has many more keybits, then the attacker can unroll it a round at a time.

The effect is that, even if a card could be constructed that resisted probing attacks, it is likely to be vulnerable unless specific power analysis defenses are built in. (In fact, all smartcards then on the market appeared to be vulnerable [728].) Furthermore, even attackers without access to probing equipment could mount attacks cheaply and quickly.

This discovery got wide publicity and held up the deployment of smartcards while people worked on defenses. In some cases, protocol level defenses are possible; the EMV protocol for bank cards mandates (from version 4.1) that the

key used to compute the MAC on a transaction be a session key derived from an on-card master key by encrypting a counter. In this way, no two ciphertexts that are visible outside the card should ever be generated using the same key. But most existing protocols are too well entrenched to be changed radically. Another idea was to insert randomness into the way the cryptography was done. One (bad) idea was that, at each round of DES, one might look up the eight S-boxes in a random order; all this achieves is that instead of one large spike in the differential trace, one gets eight spikes each with an eighth the amplitude, so the attacker has merely to collect some more power curves. A better idea was to mask the computation by introducing some offsets in each round and recalculating the S-boxes to compensate for them. This way, the implementation of the cipher changes every time it's invoked.

The defenses now being fielded against power analysis in the better devices depend on special hardware. One of the market-leading cards has hardware that inserts a dummy operation about every 64 machine instructions; another has an internal clock that is only loosely coupled to the external one and that changes frequency about every 64 cycles. Neither of these is foolproof, as an attacker might use signal processing techniques to realign the power curves for averaging. Testing a device for DPA resistance is not straightforward; there is a discussion by Paul Kocher at [729].

There are many variants on power analysis. Attacks based on cache misses can be carried out by measuring power as well as the time taken to encrypt, as a miss activates a lot of circuitry to read nonvolatile memory; so you can't stop cache attacks on AES just by ensuring that each encryption takes a constant number of clock cycles. Another variant is to use different sensors: David Samyde and Jean-Jacques Quisquater created *electromagnetic analysis*, in which they move a tiny pickup coil over the surface of the chip to pick up local signals rather than relying simply on the device's power supply [1054]. The latest twist was invented by Sergei Skorobogatov, who uses a laser to illuminate a single target transistor in the device under test for half of the test runs [1185]. This gives access not just to a Hamming weight of a computation, but a single bit; even if the device is constructed using glue logic, the attacker can still target the sense amplifiers of memory structures.

17.4.2 Leakage Through RF Signals

When I first learned to program in 1972 at the Glasgow Schools' Computer Centre, we had an early IBM machine with a 1.5 MHz clock. A radio tuned to this frequency in the machine room would emit a loud whistle, which varied depending on the data being processed. This phenomenon was noted by many people, some of whom used it as a debugging aid. A school colleague of mine had a better idea: he wrote a set of subroutines of different lengths such that

by calling them in sequence, the computer could be made to play a tune. We didn't think of the security implications at the time.

Moving now to more modern equipment, all VDUs emit a weak TV signal — a VHF or UHF radio signal modulated with a distorted version of the image currently being displayed — unless they have been carefully designed not to. The video signal is available at a number of places in the equipment, notably in the beam current which is modulated with it. This signal contains many harmonics of the dot rate, some of which radiate better than others because cables and other components resonate at their wavelength. Given a suitable broadband receiver, these emissions can be picked up and reconstituted as video. The design of suitable equipment is discussed in [408, 753]. Contrary to popular belief, LCD displays are also generally easy for the eavesdropper; a typical laptop has a serial line going through the hinge from the system unit to the display and this carries the video signal (Figure 17.2).

Other researchers quickly established the possibility of remote snooping on everything from fax machines through shielded RS-232 cables to ethernet [365, 1196]. A few companies sprang up to sell 'jammers' but this is hard to do properly [98]: they can interfere with TV and other services. The military use 'Tempest' shielded equipment, but this has generally remained unavailable to the commercial sector. In any case, it is usually a generation out of date and five times as expensive as off-the-shelf PCs. The view taken in the banking industry was 'well, we don't do it to our competitors, so they probably don't do it to us, and we don't know where to get effective countermeasures anyway — so

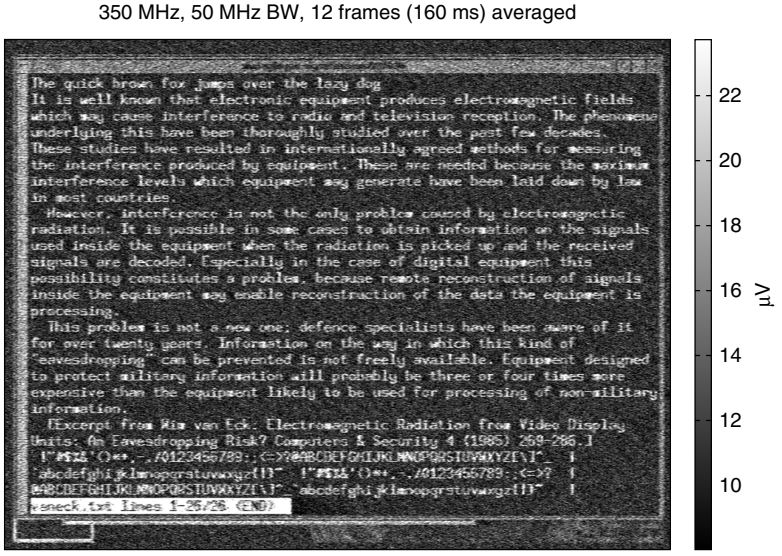

Figure 17.2: RF signal from a Toshiba laptop reconstructed several rooms away, through three plasterboard walls (courtesy of Markus Kuhn [752]).

put it in the ''too hard'' file'. This view got shaken somewhat in the late 1990s when Hans-Georg Wolf demonstrated a Tempest attack that could recover card and PIN data from a cash machine at a distance of eight meters [744].

Tempest precautions remain a rarity outside the defense sector, but one recent exception comes from the world of voting machines. In October 2006, a Dutch group opposed to electronic voting machines demonstrated that the voting machine used to collect 90% of the election ballots in the Netherlands could be eavesdropped from a distance of several tens of meters [541]. This has led to a Dutch government requirement that voting equipment be Tempest-tested to a level of 'Zone 1–12 dB'.

The *zone* system works as follows. Equipment certified as Zone 0 should not emit any signals that are exploitable at a distance of one meter; it should protect data from electronic eavesdropping even if the opponent is in the next room, and the wall is something flimsy like plasterboard. Zone 1 equipment should be safe from opponents at a distance of 20 meters, and thus the Dutch 'Zone 1–12 dB' criterion means that a voting machine should not leak any data on what vote was cast to an eavesdropper 5 meters away. Zone 2 and Zone 3 mean 120 and 1200 meters respectively. Technical details of zoning were briefly published in 2007, as [243]. (This document was then withdrawn, perhaps because the Americans objected to the Germans releasing it. However everything in it was already in the public domain except the zone limit curves, which are worst-case relative attenuations between distances of 20 m, 120 m and 1200 m from a small dipole or loop antenna, taking into account the difference between nearfield and farfield dropoff.)

The zone system has come into wide governmental use since the end of the Cold War, which slashed military budgets and forced government agencies to use commercial off-the-shelf equipment rather than developing hardware exclusively for their own use. Commercial off-the-shelf equipment tends to be zone 2 when tested, with some particularly noisy pieces of kit in zone 3. By knowing which equipment radiates what, you can keep most sensitive data on equipment furthest from the facility perimeter, and shield stuff only when you really have to. The most sensitive systems (such as national intelligence) and those exposed to the highest threats (such as in embassies overseas) are still either shielded, or kept in shielded rooms. Zoning has greatly cut the costs of emission security, but the overall bill in NATO government agencies comes to over a billion dollars a year.

Markus Kuhn and I developed a lower-cost protection technology, called 'Soft Tempest', which has been deployed in some products, from the email encryption package PGP to the latest Dutch election machines [753]. Soft Tempest uses software techniques to filter, mask or render incomprehensible the information bearing electromagnetic emanations from a computer system.

We discovered that most of the information bearing RF energy from a VDU was concentrated in the top of the spectrum, so filtering out this component

is a logical first step. We removed the top 30% of the Fourier transform of a standard font by convolving it with a suitable low-pass filter (see Figures 17.3 and 17.4).

Figure 17.3: Normal text

Figure 17.4: Text low-pass filtered

This turns out to have an almost imperceptible effect on the screen contents as seen by the user. Figures 17.5 and 17.6 display photographs of the screen with the two video signals from Figures 17.3 and 17.4.

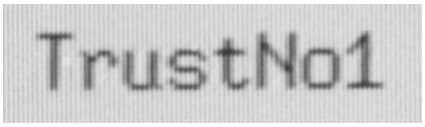

Figure 17.5: Screen, normal text

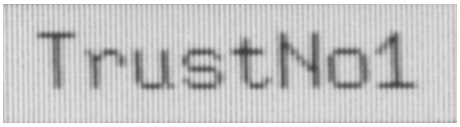

Figure 17.6: Screen, filtered text

However, the difference in the emitted RF is dramatic, as illustrated in the photographs in Figures 17.7 and 17.8. These show the potentially compromising emanations, as seen by a Tempest monitoring receiver.

Figure 17.7: Page of normal text

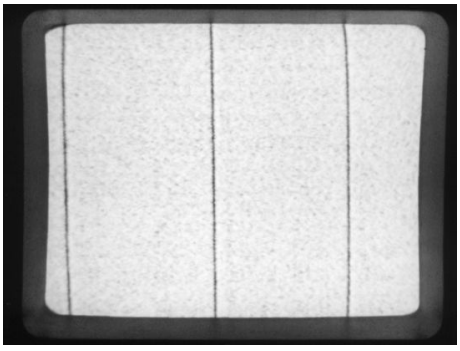

Figure 17.8: Page of filtered text

While the level of protection which Soft Tempest techniques can provide for VDUs is only of the order of 10–20 dB, this translates to a difference of a zone — which in an organization the size of a government, can save a lot of money [70].

There are other attacks that software tricks can block completely. For example, computer keyboards can be snooped on while the microcontroller goes through a loop that scans all the keys until it encounters one that is pressed. The currently pressed key is modulated on to the RF emissions from the keyboard. By encrypting the order in which the keys are scanned, this kind of attack can be completely blocked.

17.5 Active Attacks

However, it's not enough to simply encrypt a keyboard scan pattern to protect it, as the attacker can use active as well as passive techniques. Against a keyboard, the technique is to irradiate the cable with a radio wave at its resonant frequency. Thanks to the nonlinear junction effect, the keypress codes are modulated into the return signal which is reradiated by the cable. This can be picked up at a distance of 50–100 yards. To prevent it, one must also encrypt the signal from the keyboard to the PC [753].

17.5.1 Tempest Viruses

There are quite a few other active attacks possible on various systems. The phenomenon that we observed with the IBM 1401 — that a suitable program would cause a computer to play a tune on the radio, in effect turning it into a low-grade radio transmitter — is easy enough to reimplement on a modern PC. Figures 17.9 and 17.10 show what the screen on a PC looks like when the video signal is an RF carrier at 2 MHz, modulated with pure tones of 300 and 1200 Hz.

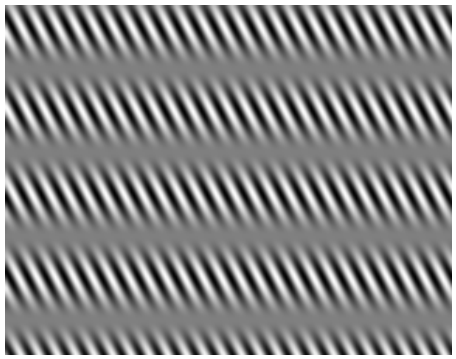

Figure 17.9: 300 Hz AM signal

Figure 17.10: 1200 Hz AM signal

Using phenomena like this, it is possible to write a *Tempest virus* that will infect a target computer and transmit the secret data it steals to a radio receiver hidden nearby. This can happen even if the machine is not connected to the net. The receiver need not be expensive; a short wave radio with a cassette recorder will do, and exploit code has already been published. With more sophisticated techniques, such as spread-spectrum modulation, it's possible for an attacker with more expensive equipment to get much better ranges [753].

Some of these methods were already known to the intelligence community. There have been reports of the CIA using software-based RF exploits in economic espionage against certain European countries (for example, in a TV documentary accompanying the release of [725]). Material recently declassified by the NSA in response to a FOIA request [869, 666] reveals the use of the codeword *Teapot* to refer 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'. A further example is to attack equipment that's been shielded and Tempest-certified up to a certain frequency (say, 3 GHz) by irradiating it through the ventilation slots using microwaves of a much higher frequency (say 10GHz) at which these slots become transparent [753].

The possibility of attacks using malicious code is one reason why Tempest testing involves not just listening passively to the emanations from the device under test, but injecting into it signals such as long linear feedback shift register sequences. These create a spread-spectrum signal which will likely be detectable outside the device and thus simulate the worst-case attack in which the opponent has used a software exploit to take over the device [177]. I understand that normal Tempest certification does not take account of the process gain that can be obtained by such techniques.

17.5.2 Nonstop

Another class of active methods, called *Nonstop* by the U.S. military [87], is the exploitation of RF emanations that are accidentally induced by nearby radio transmitters and other RF sources. If equipment processing sensitive data is used near a mobile phone, then the phone's transmitter may induce currents in the equipment that get modulated with sensitive data by the nonlinear junction effect and reradiated.

For this reason, it used to be forbidden to use a mobile phone within 5 meters of classified equipment. Nonstop attacks are also the main Emsec concern for ships and aircraft; here, an attacker who can get close enough to do a passive Tempest attack can probably do much more serious harm than eavesdropping, but as military ships and aircraft often carry very powerful radios and radars, one must be careful that their signals don't get modulated accidentally with something useful to the enemy.

17.5.3 Glitching

Active Emsec threats are also significant in the smartcard world, where perhaps the best known is the *glitch attack* [68]. Here, the opponent inserts transients into the power or clock supply to the card in the hope of inducing a useful error.

For example, one smartcard used in early banking applications had the feature that an unacceptably high clock frequency only triggered a reset after a number of cycles, so that transients would be less likely to cause false alarms. So it was possible to replace a single clock pulse with two much narrower pulses without causing an alarm that would reset the card. This reliably caused the processor to execute a NOP regardless of what instruction it was supposed to execute. This gives rise to a *selective code execution* attack that can be extraordinarily powerful. For example, the attacker can step over jump instructions and thus bypass access controls.

17.5.4 Differential Fault Analysis

Even where the attacker does not know the card's software in detail, glitch attacks can still be a knockout. Dan Boneh, Richard DeMillo and Richard Lipton noticed that a number of public key cryptographic algorithms break if a random error can be induced [206]. For example, when doing an RSA signature the secret computation $S = h(m)^d \pmod{pq}$ is carried out mod p, then mod q, and the results are then combined, as this is much faster. But if the card returns a defective signature S_p which is correct modulo p but incorrect modulo q, then we will have

$$p = \gcd(pq, S_p^e - h(m))$$

which breaks the system at once. These attacks can easily be implemented if the card isn't protected against glitches, and can also be easily extended to many symmetric algorithms and protocols. For example, Eli Biham and Adi Shamir pointed out that if we have the power to set a given bit of memory to zero (or one), and we know where in memory a key is kept, we can find out the key by just doing an encryption, zeroising the leading bit, doing another encryption and seeing if the result's different, then zeroising the next bit and so on [171]. Optical probing may be just the tool for this.

Our subsequent discovery of the power of optical probing means that such attacks can be implemented routinely by an attacker who can get access to the chip surface and identify the relevant memory structures [1111].

17.5.5 Combination Attacks

Other attacks use a combination of active and passive methods. I mentioned in passing in Part I a trick that could be used to find the PIN in a stolen

smartcard. Early card systems would ask the customer for a PIN and if it was incorrect they would decrement a retry counter. As this involved writing a byte to EEPROM, the current consumed by the card rose measurably as the capacitors in the circuit that boosts the supply voltage V_{cc} to the programming voltage V_{pp} were charged up. On noticing this, the attacker could simply reset the card and try the next candidate PIN.

The leading active-passive method at the time of writing is Sergei Skorobogatov's optically enhanced position-locked power analysis [1185], which uses a laser to partially ionise a target transistor while power analysis is carried out, and which I discussed in section 17.4 above. This can be extended to an active attack by increasing the laser power so as to make the target transistor conduct.

17.5.6 Commercial Exploitation

Not all Emsec attacks involve covert military surveillance or lab attacks on tamper-resistant devices. I already mentioned the TV detector vans used in Britain to catch TV license defaulters, and the uproar over voting machines in Holland. There are also marketing applications. U.S. venue operator SFX Entertainment monitors what customers are playing on their car radios as they drive into venue parking lots by picking up the stray RF from the radio's local oscillator. Although legal, this annoys privacy advocates [1212]. The same equipment has been sold to car dealers, mall operators and radio stations.

17.5.7 Defenses

The techniques that can be used to defend smartcards against active Emsec threats are similar to those used in the passive case, but not quite the same.

The use of timing randomness — jitter — is still useful, as a naive opponent might no longer know precisely when to insert the glitch. However, a clever opponent may well be able to analyze the power curve from the processor in real time and compare it against the code so as to spot the critical target instructions. In addition, fault attacks are hard to stop with jitter, as the precise location of the fault in the code is not usually critical.

In some cases, defensive programming is enough. For example, the PIN search described above in section 17.5.5 is prevented in more modern cards by decrementing the counter, soliciting the PIN, and then increasing the counter again if it's correct. Fault attacks on public key protocols can be made a lot harder if you just check the result.

Other systems use specific protective hardware, such as a circuit that integrates the card reset with the circuit that detects clock frequencies that are too high or too low. Normal resets involve halving the clock frequency for a few cycles, so an attacker who found some means of disabling the monitoring

function would quite likely find himself unable to reset the card at all on power up [733].

So we know in principle how to defend ourselves against glitch attacks (though extensive device testing is always advisable). Optical probing is harder; the sort of attack a card can face nowadays is from an opponent who puts it in a test rig, initiates a cryptographic operation, and fires a laser at the chip to cause a single device upset at a precisely measured time. With a motorised stage, hundreds of different targets can be tried per hour.

Colleagues and I at Cambridge designed and prototyped a defence technology that can resist such an onslaught, using dual-rail self-timed logic. In such logic, rather than signaling '1' by High and '0' by Low on a single line, we have two lines for each logic bit. We signal '1' by 'HighLow' and '0' by 'LowHigh'; the end of a logic operation is 'LowLow'. Such logic has been known for some years, and has the property that if the 'HighHigh' state someone enters the system, it tends to propagate across the chip, locking it up and rendering the device inoperable until it's reset. We made a virtue of this by defining 'HighHigh' to be the 'alarm' state, and redesigning the gates so that a single-event upset would cause an alarm to propagate. We also salted the chip with a number of alarm sensors. And because the logic is balanced, the power consumption is much less dependent on the data; the signals potentially exploitable by power analysis were reduced by about 20 dB [903]. A number of other researchers have recently started looking at such exotic design styles, and some have started to appear in products. Redundant design can be fairly expensive, costing at least three times as much as standard CMOS, but the cost per transistor may now have fallen to the point where it makes sense.

17.6 Optical, Acoustic and Thermal Side Channels

In recent years, there has been a stream of interesting new results on novel side-channel attacks. Have you ever looked across a city at night, and seen someone working late in their office, their face and shirt lit up by the diffuse reflected glow from thier computer monitor? Did you ever stop to wonder whether any information might be recovered from the glow? In 2002 Markus Kuhn showed that the answer was pretty well 'everything': he hooked up a high-performance photomultiplier tube to an oscilloscope, and found that the light from the blue and green phosphors used in common VDU tubes decays after a few microseconds. As a result, the diffuse reflected glow contains much of the screen information, encoded in the time domain. Thus, given a telescope, a photomultiplier tube and suitable image-processing software, it was possible to read the computer screen at which a banker was looking by decoding the light scattered from his face or his shirt [750].

The next headline was from Joe Loughry and David Umphress, who looked at the LED status indicators found on the data serial lines of PCs, modems, routers and other communications equipment. They found that a significant number of them were transmitting the serial data optically: 11 out of 12 modems tested, 2 out of 7 routers, and one data storage device. The designers were just driving the tell-tale light off the serial data line, without stopping to realise that the LED had sufficient bandwidth to transmit the data to a waiting telescope [815].

Acoustics came next. There had always been a 'folk rumour' that the spooks were able to tell what someone was typing on the old IBM Selectric typewriter by just recording the sound they made, and it had been reported that data could be recovered from the noise made by dot matrix printers [228]. In 2004, Dmitri Asonov and Rakesh Agrawal showed that the different keys on a keyboard made different enough sounds. They trained a neural network to recognise the clicks made by key presses on a target keyboard and concluded that someone's typing could be picked up from acoustic emanations with an error rate of only a few percent [91]. Now Dawn Song, David Wagner and XuQing Tian had also shown that SSH encrypted sessions leak a considerable amount of information as the keystrokes are sent in individual packets, the time-delays between which are visible to an attacker; they noted that this would enable an attacker about a factor of 50 advantage in guessing a password whose encrypted value he'd observed [1203].

In 2005, Li Zhuang, Feng Zhou, and Doug Tygar combined these threads to come up with an even more powerful attack. Given a recording of someone typing text in English for about ten minutes on an unknown keyboard, they recognised the individual keys, then used the inter-keypress times and the known statistics of English to figure out which key was which. Thus they could decode text from a recording of a keyboard to which they had never had access [1376].

In 2004, Eran Tromer and Adi Shamir took the acoustic analysis idea down to a much lower level: they showed that keys leak via the acoustic emanations from a PC, generated mostly at frequencies above 10KHz by capacitors on the motherboard [1263].

The latest development has been thermal covert channels. In 2006, Steven Murdoch discovered that a typical computer's clock skew, which can be measured remotely, showed diurnal variation, and realised this was a function of ambient temperature. His experiments showed that unless a machine's owner takes countermeasures, then anyone who can extract accurate timestamps from it can measure its CPU load; and this raises the question of whether an attacker can find where in the world a hidden machine is located. The longitude comes from the time zone, and the latitude (more slowly) from the seasons. So hiding behind an anonymity service such as Tor might not be as easy as it looks [914, 916].

17.7 How Serious are Emsec Attacks?

Technical surveillance and its countermeasures — bugs — are the most important aspect of Emsec, in both government and industry. They are likely to remain so. The range of bugs and other surveillance devices that can be bought easily is large and growing. The motivation for people to spy on their rivals, employees and lovers will continue. If anything, the move to a wired world will make electronic surveillance more important, and countermeasures will take up more of security budgets.

Those aspects of Emsec which concern equipment not designed for surveillance — Tempest, Teapot, Hijack, Nonstop and the various types of power and glitch attack — are set to become another of the many technologies which got their initial development in the government sector but which become important in the design of commercial products.

17.7.1 Governments

The Emsec threats to embassies in hostile countries are real. If your country is forced by the President of Lower Slobovia to place its embassy in the second floor of an office block whose first and third floors are occupied by the local secret police, then security is a hard problem. Shielding all electronic equipment (except that used for deception) will be part of the solution. It won't be all of it; your cleaning ladies will no doubt be in the pay of the Slobovian security forces and will helpfully loosen your equipment's Tempest gaskets, just as they change the batteries in the room bugs.

In less threatening environments, the cost-effectiveness of hardware Tempest shielding is more doubtful. Despite the hype with which the Tempest industry maintained itself during the Cold War, there is a growing scepticism about whether any actual Tempest attacks had ever been mounted by foreign agents in the USA. Anecdotes abound. It's said, for example, that the only known use of such surveillance techniques against U.S. interests in the whole of North America was by Canadian intelligence personnel, who overheard U.S. diplomats discussing the U.S. bottom line in grain sales to China.

There was a scandal in April 2007 when it emerged that Lockheed-Martin had ignored Tempest standards when installing equipment in U.S. Coast Guard vessels. Documents were left on the web site of the Coast Guard's Deepwater project and ended up an activist website, cryptome.org, which was closed down for a while. The documents tell a story not just of emission security defects — wrong cable types, violations of cable separation rules, incorrect grounding, missing filters, red/black violations, and so on — but of a more generally botched job. The ships also had hull cracks, outdoor radios that were not waterproof, a security CCTV installation that did not

provide the specified 360 degree coverage, and much more [338]. This led to a Congressional inquiry. The documents at least provide some insight into the otherwise classified Tempest and Nonstop accreditation procedures.

I must confess I might have some sympathy if Coast Guard personnel had simply placed a low priority on Tempest defences. Having been driven around an English town looking for Tempest signals, I can testify that doing such attacks is much harder in practice than it might seem in theory, and the kind of opponents we face nowadays are rather different from the old Soviet intelligence machine. Governments are rightly more relaxed about Tempest risks than twenty years ago.

17.7.2 Businesses

In the private sector, the reverse is the case. The discovery of fault attacks and then power attacks was a big deal for the smartcard industry, and held up for probably two years the deployment of smartcards in banking applications in those countries that hadn't already committed to them. The currently deployed devices are still not perfect; attacks are kept at bay by a mishmash of ad-hoc mechanisms whose effectiveness depends on there being capable designers who understand the problem and whose managers remain worried about it. As the 'DPA' scare recedes, and equipment becomes ever more complex, expect the residual vulnerabilities to reassert themselves. Building chip-scale devices that really block side-channel attacks is hard, and few customers are prepared to pay for it.

And what about the future?

The 'non-security' aspects of emission management, namely RFI/EMC, are becoming steadily more important. Ever higher clock speeds, plus the introduction of all sorts of wireless devices and networks and the proliferation of digital electronics into many devices which were previously analogue or mechanical, are making electromagnetic compatibility a steadily harder and yet more important problem. A host of incompatible standards are managed by different industry groups, many of which are rapidly becoming obsolete — for example, by not requiring testing above 1 GHz, or by assuming protection distances that are no longer reasonable [715].

On the 'security' side, attacks are likely to become easier. The advent of *software radios* — radios that digitize a signal at the intermediate frequency stage and do all the demodulation and subsequent processing in software — will be important. These were until recently an expensive military curiosity [761] but are now finding applications in places like cellular radio base stations. The next generation may be consumer devices, designed to function as GPS receivers, GSM phones, wireless LAN basestations, and support whatever other radio based services have been licensed locally — all with only a change in software.

Once people learn how to program them, they may well use them for Tempest attacks.

Finally, Emsec issues are not entirely divorced from electronic warfare. As society becomes more and more dependent on devices that are vulnerable to strong radio frequency signals — such as the high power microwaves generated by military radars — so the temptation to mount attacks will increase. I'll discuss high energy radio frequency attacks in the next chapter but one.

17.8 Summary

Emission security covers a whole range of threats in which the security of systems can be subverted by compromising emanations, whether from implanted bugs, from unintentional radio frequency or conducted electromagnetic leakage, to emanations that are induced in some way. Although originally a concern in the national intelligence community, Emsec is now a real issue for companies that build security products such as smartcards and cash machines. Many of these products can be defeated by observing stray RF or conducted signals. Protecting against such threats isn't as straightforward as it might seem.

Research Problems

We need a comprehensive set of emission security standards for commercial use. The military standards — NATO SDIP-27 and USA NSTISSAM — are classified, although they've leaked as described in section 17.4.2. RFI/EMC standards — the civilian IEC/CISPR 22 and the stricter MIL-STD-461E — were simply not designed to protect information. The recent panic in Holland about Tempest snooping on voting machines shows that standards are needed, so that equipment purchasers and vendors can take a view on whether they're needed in any given application.

Further Reading

There is a shortage of open literature on Emsec. The classic van Eck article [408] is still worth a read, and the only book on computer security (until this one) to have a chapter on the subject is Russell and Gangemi [1098]. Our work on Soft Tempest, Teapot and related topics can be found in [753]. For power analysis, see the papers by Kocher [728] and by Messergues et al. [877]; more papers appearing regularly at the CHES workshop. Joel McNamara runs an unofficial Tempest Web site at [869]. For timing and power analysis, the original papers by Paul Kocher and colleagues are the classic references [727, 728]; there's also a book by Stefan Mangard, Elisabeth Oswald and Thomas Popp [833].