CHAPTER

15

Emission Security

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

—WILLIAM SHAKESPEARE, KING HENRY V, ACT IV

15.1 Introduction

Emission security, or *Emsec*, refers to preventing a system being attacked 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. 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.

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 (at least, once the research needed to design the attack had been carried out), 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. There has also been much speculation about disruptive electromagnetic attacks, whereby, for example, a terrorist group uses 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 preventing random system disruption happening as a result of problems with *electromagnetic compatibility* (EMC) and *radio frequency interference* (RFI). If you fly regularly, you've no doubt heard the captain say something like, "All electronic devices must be switched off now, and not switched on again until I turn off the seat belt sign 10 minutes after takeoff." This problem is worsening as everything becomes electronic and clock frequencies go up. And how do you do as the captain says when more and more devices are designed to be "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 into the gigahertz range, all these problems—RFI/EMC, Emsec and various electronic warfare threats—are set to get worse.

15.2 History

"Crosstalk" between telephone wires was a problem known to the pioneers of telephony in the nineteenth century, with their two-wire circuits stacked on tiers of crosstrees on supporting poles. One way of dealing with it was to use "transpositions," whereby 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 1884–1885 [569].

The first appearance of compromising emanations in warfare seems to date to 1914. Field telephone wires were laid to connect the troops with their headquarters, and these often ran for miles, parallel to enemy trenches that were only a few hundred yards away. These wires used a single-core insulated cable and 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 3,000 yards of the front. When the United States joined the war, the techniques were passed on. More information can be found in [542, 569].

During the World War II, radio engineering saw advances in radar, passive direction finding, and low-probability-of-intercept techniques, which I'll discuss in the next chapter. 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 that is supposed to support public broadcast services. Its use has since expanded to satellite and cable TV operators, who use detector vans to find pirate decoders. Some people in the computer security community were also aware that information could leak from cross-coupling and stray RF (see, for example, [259, 791]).

The intelligence community also started to exploit RF effects. In 1960, the British prime minister ordered surveillance on the French embassy in the course of negotiations about joining the European Economic Community. Scientists from his domestic intelligence agency, MI5, 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 [814]. 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 [259]. The revelation that Tempest attacks were not just feasible, but could be mounted with simple equipment that could be built at home, 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 observed in [43] that many smartcards could be broken by inserting transients, or *glitches*, in their power or clock lines (this attack wasn't discovered by us, but by pay-TV hackers). Paul Kocher also showed that many common implementations of cryptosystems could be broken by making precise measurements of the time taken [466]. In 1998, Kuhn and I published a paper showing that many of the compromising emanations from a PC could be made better, or worse, by appropriate software measures [478]. 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—which we'll discuss in detail in Section 15.4.1.2 below [467]. In 2000, David Samyde and Jean-Jacques Quisquater demonstrated that similar attacks could be carried out by bringing small electromagnetic field sensors close to the card's surface [668].

15.3 Technical Surveillance and Countermeasures

Before getting carried away with high-tech toys such as Tempest monitoring receivers, we need to stop and think about bugs. The simplest and most widespread attacks that use the electromagnetic spectrum are not those that exploit some unintended design 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.

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, which 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 a few days or weeks.
- At the next step up are devices that draw their power from the mains, a telephone cable or some other external electricity supply. This means that they can last indefinitely once positioned. 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. Still others, used by the U.K. police in recent gangland surveillance cases, look like electrical adaptors, but actually contain a microphone, a radio transmitter, and a TV camera. Others monitor data—for example, there is a Trojan computer keyboard with bugging hardware contained in the cable connector.
- Many modern bugs use off-the-shelf mobile radio technology. They can be seen as slightly modified cellphone handsets which go off-hook silently when called.
- One exotic device, on show at the NSA Museum in Fort Meade, 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 the office. Right up to the end of the Cold War, embassies in Moscow were regularly irradiated with microwaves, so presumably variants of the technique continued to be used.
- 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 upward 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.

A number of countermeasures can give a fair amount of protection against such attacks.

- 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 have the effect of rectifying incident radio frequency signals. The device broadcasts a weak radio signal, and listens for harmonics of this signal. 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. An interesting 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 se-

curity 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.

- A number of *surveillance receivers* are 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 3 GHz, 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 it is often impractical. It can be cheaper to have a shielded internal room for particularly sensitive meetings; and there are vendors who sell prefabricated rooms with acoustic and electromagnetic shielding for just this purpose.
- 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 [55]. This is very expensive, and in many cases impractical. A second-best 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 it was recently reported that Britain's counterintelligence service had to tear down and rebuild a large part of a new headquarters building, at a cost of about $50 million, after an employee of one of the building contractors was found to have past associations with the Provisional IRA.

The traditional tension here is between technological defenses, which can be very 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 greater scope for attacks involving equipment that's already in place rather than stuff that has to emplaced for the purpose. For example:

- The risks associated with telephones are much higher 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 (ISDN) phones have a facility built into them that allows this (it's said that some repressive countries make this feature 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 have a radio LAN connection. An attacker might infect the device with a virus that listens to conversations in the room, compresses them, encrypts them, and emails 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.

15.4 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 her part to create. Broadly speaking, there are two categories. 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. Still, it's a reasonable working classification most of the time.

15.4.1 Leakage through Power and Signal Cables

Since the nineteenth century, engineers have been aware that high-frequency signals leak everywhere, and that careful measures are needed to stop them causing problems; as 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.

15.4.1.1 Red/Black Separation

Red equipment (carrying confidential data such as plaintext) has to be isolated by filters and shields from *black* equipment (which 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 [660].

So properly shielded equipment tends to be available only in small quantities, and made specifically for defense markets. This makes it extremely expensive. And 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.

15.4.1.2 Power Analysis

Often, people aren't aware of the need to filter signals until an exploit is found. A recent, and 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 chokes, capacitors and so on. 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 simply by 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 10 Ω 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 15.1.

Different instructions have quite different power-consumption profiles, and, as you can see in the figure, 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 14.5). 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 [547]. 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 [549]. EEPROM reads and writes can give even more substantial signals.

The effect of this leakage is that an attacker who understands how a cipher is implemented (for example, as a result of probing out the card software and disassembling it) 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 dismantle a card, understand its contents, and design the attack, a very large number of cards may be compromised at little marginal cost.

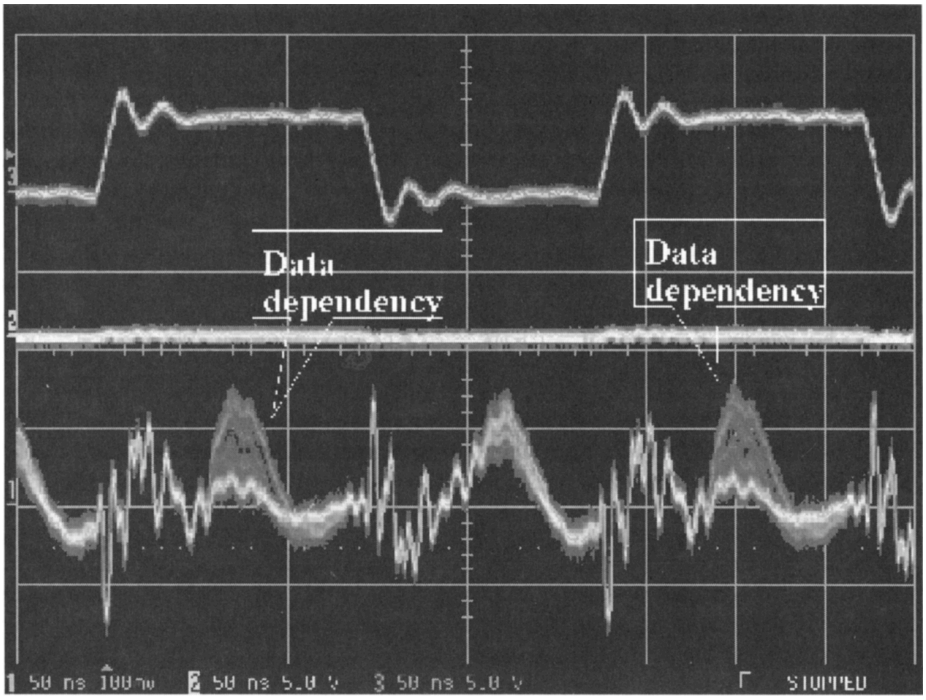

Figure 15.1 Superimposed power consumption traces of a Siemens SLE44 smartcard, showing data dependency (courtesy of Andy Santoso). The upper half of the screen shows the clock signal; the lower shows the power consumption.

The threat posed to smartcards by power analysis was brought forcefully to the industry's attention in 1998 with the development by Paul Kocher of a specific signal-processing technique to extract the key bits used in a block cipher, such as DES, from a collection of power curves, without knowing the implementation details of the card software. This technique, *differential power analysis*, works as follows [467].

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 on the market in 1998 were claimed to be vulnerable [467].) Furthermore, even attackers without access to probing equipment could mount attacks cheaply and quickly.

This discovery was widely publicized, and held up the deployment of smartcards while people worked on defenses. In some cases, protocol-level defenses are possible; one can design protocols that update the key with every few encryptions, and thus prevent the attacker getting enough data (some point-of-sale terminals are designed this way). 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. For example, at each round of DES, one might look up the eight S-boxes in a random order. However, all this achieves is that instead of one large spike in the differential trace, one gets eight spikes each with an eighth the amplitude; the attacker only has to collect some more power curves.

The defenses now being fielded against power analysis are hardware-based. One of the common 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. The next generation of cards may use more robust defenses, such as potting capacitors with the smartcard chip to enable the supply voltage to be properly decoupled, or using silicon design techniques such as dual-rail encoding where the current drawn is independent of the data being processed. Yet another approach is to use self-timed logic, which uses no clock. At the time of writing, this is an area of active research.

15.4.2 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. Similar phenomena were noted by many people, some of whom used the noise 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 that 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 [259, 478]. Contrary to popular belief, LCD displays are also generally easy for the eavesdropper.

Other researchers quickly established the possibility of remote snooping on everything from fax machines through shielded RS-232 cables to Ethernet [719, 230]. A few companies sprang up to sell "jammers," but these are hard to implement properly [60], as they can interfere with TV and other services. Military Tempest-shielded equipment 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 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 [239]. However, Tempest precautions remain a rarity in commerce and in nondefense-sector industry.[1]

Meanwhile, with the end of the Cold War, military budgets were cut, and often there was no alternative to using commercial off-the-shelf equipment; there was no longer the money to develop systems exclusively for government use. Government organizations in NATO countries have switched to a *zone* model of Emsec protection, whereby the most sensitive equipment is kept in the rooms furthest from the facility perimeter, and shielding is reserved for the most sensitive systems (such as national intelligence) or where the threat is highest (such as in overseas embassies). Nonetheless, the bill for Tempest protection in NATO government agencies comes to over a billion dollars a year.

A lower-cost protection technology, called Soft Tempest, has emerged and been deployed in some commercial products (such as the email encryption package PGP) [478]. Soft Tempest uses software techniques to filter, mask, or render incomprehensible the information bearing electromagnetic emanations from a computer system.

Markus Kuhn and I 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 15.2 and 15.3).

This turns out to have an almost imperceptible effect on the screen contents as seen by the user. Figures 15.4 and 15.5 display photographs of the screen with the two video signals from Figures 15.2 and 15.3.

The difference in the emitted RF is dramatic, as illustrated in the photographs in Figures 15.6 and 15.7. These show the potentially compromising emanations, as seen by a Tempest monitoring receiver.

The level of protection that Soft Tempest techniques can provide for VDUs is only on the order of 10–20dB, but this translates to a difference of a zone—which, in an organization the size of a government, can give a considerable cost saving [45].

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

[1] Just as I got the copyedited manuscript of this book back from Wiley for checking. I heard, for the first time, a believable report of a commercial Tempest attack. Apparently, one financial institution was spied on by a private investigator retained by a rival. But the big picture remains military.

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.

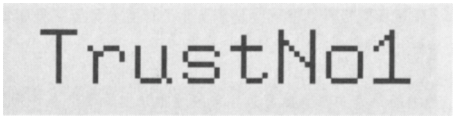

Figure 15.2 Normal text.

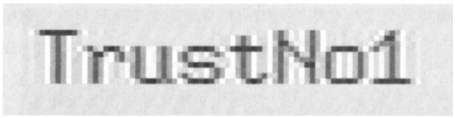

Figure 15.3 Same text, low-pass filtered.

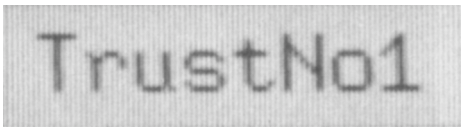

Figure 15.4 Screenshot, normal text.

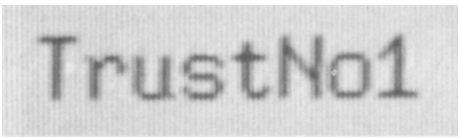

Figure 15.5 Screenshot, filtered text.

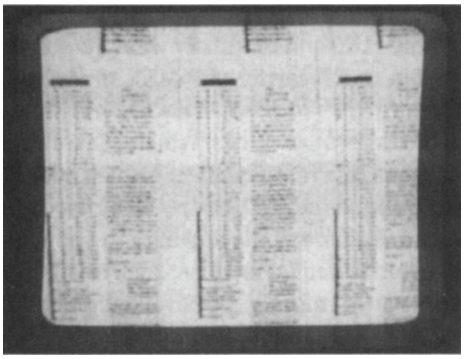

Figure 15.6 Page of normal text.

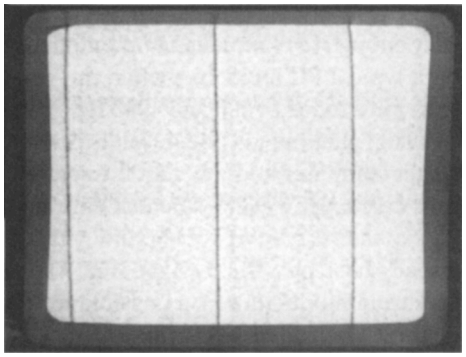

Figure 15.7 Page of filtered text.

15.5 Active Attacks

But 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 to 100 yards. To prevent it, one must also encrypt the signal from the keyboard to the PC [478].

15.5.1 Tempest Viruses

There are quite a few other active attacks possible on various systems. The phenomenon observed with our school computer in 1972—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 15.8 and 15.9 show what the screen on a typical PC looks like when the video signal is an RF carrier at 2 MHz, modulated with pure tones of 300 and 1200 Hz.

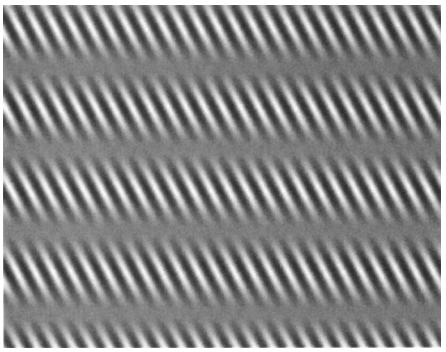

Figure 15.8 A 300 Hz broadcast signal.

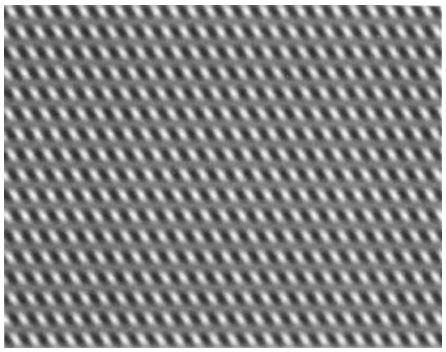

Figure 15.9 A 1200 Hz broadcast signal.

Using phenomena like this, it is possible to write a *Tempest virus*, which 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 on the Net. With more sophisticated techniques, such as spread-spectrum modulation, it's possible for the attacker with more expensive equipment to get much better ranges [478].

Some of these methods may already have been 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 [464]). Material recently declassified by the NSA in response to a FOIA request [542, 420] 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 has been shielded and Tempest-certified up to a certain frequency (say, 1 GHz) by irradiating it through the ventilation slots using microwaves of a much higher frequency (say 10 GHz), at which these slots become transparent [478].

The possibility of attacks using malicious code is one reason why Tempest testing may involve 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 equipment and thus simulate the worst case attack in which the opponent has used a software exploit to take over the device [108].

15.5.2 Nonstop

Another class of active methods, called *Nonstop* by the U.S. military [55], is the exploitation of RF emanations that are accidentally induced by nearby radio transmitters and other RF sources. If equipment that is processing sensitive data is used near a mo-

bile phone, the phone's transmitter may induce in it currents 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 because 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.

15.5.3 Glitching

Active Emsec threats are also significant in the smartcard world, where perhaps the best known is the *glitch attack* [43]. Here, as I mentioned above, 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 triggered a reset only 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. By introducing glitches at suitable times, the attacker could step over jump instructions, and thus bypass access controls.

15.5.4 Differential Fault Analysis

Even where the attacker does not know the card's software in detail, glitch attacks can still be a threat. It had been noticed that a number of public key cryptographic algorithms would break if a random error could be induced [126]. For example, when doing an RSA signature, the secret computation $S = h(m)^d \pmod{pq}$ is typically carried out mod p, then mod q; the results are then combined. However, if the card returns a defective signature S which is correct modulo p but incorrect modulo q, then we will have:

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

which breaks the system at once. These attacks can be implemented easily if the card isn't protected against glitches; they can also be extended to many symmetric algorithms and protocols [103].

15.5.5 Combination Attacks

Other attacks use a combination of active and passive methods. I mentioned in passing in Part 1 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. This involved writing a byte to EEPROM, so the current consumed by the card rose measurably as the capacitors in the EEPROM voltage multiplier circuit were charged up. On noticing this, the attacker could simply reset the card and try the next candidate PIN.

15.5.6 Commercial Exploitation

Not all Emsec attacks are conducted in the context of covert military surveillance or laboratory attacks on tamper-resistant devices. I already mentioned the TV detector vans used in Britain to catch TV license defaulters and the customers of pay-TV pirates. 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 alarms privacy advocates [728]. The same equipment has been sold to car dealers, mall operators, and radio stations.

15.5.7 Defenses

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

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 in Section 15.5.5 is prevented in more modern implementations by decrementing the counter, soliciting the PIN, then increasing the counter again if it's correct. Differential 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 [470].

Current defenses against glitch attacks are not entirely foolproof, and extensive device testing is highly advisable. New technologies, such as the use of self-timed logic, may improve things by providing a high level of protection against both active and passive threats. In the meantime, if you have to write a smartcard application, attacks based on glitching merit careful consideration.

15.6 How Serious Are Emsec Attacks?

Technical surveillance and its countermeasures 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 others 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 that 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 that were initially developed in the government sector but then start being important in the design of commercial products.

15.6.1 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 an extremely hard problem. Shielding all electronic equipment (except that used for deception) will be part of the solution. In less threatening environments, the use of hardware Tempest shielding is more doubtful.

Despite the hype with which the Tempest industry maintained itself during the Cold War, there is growing scepticism about whether any actual Tempest attacks had ever been mounted by foreign agents, though 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; and that the East German Stasi were found to have maps of suitable parking places for Tempest vans in West German towns. But I've not found anything that can be nailed down to a reliable source, and having been driven around an English town looking for Tempest signals, I can testify that launching such attacks is much harder in practice than it might seem in theory. Governments now tend to be much more relaxed about Tempest risks than 10 years ago.

15.6.2 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. Blocking these attacks turns out to be difficult, and doing it properly will involve a further generation of hardware design.

And what about the future?

The "nonsecurity" 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 that were previously analogue or mechanical, are making electromagnetic compatibility a steadily harder and yet more pressing problem. Different industry groups, manage a host of incompatible standards 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 [455].

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—were, until recently, an expensive military curiosity [482], 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, radio LAN basestations, and to 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 might just as easily use them for Tempest attacks.

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

15.7 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, or from 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

The security industry badly needs a comprehensive set of emission security standards for commercial use. Military standards are classified, and the RFI/EMC standards are fragmented and contradictory, so a new and unified approach is overdue.

Further Reading

There is a shortage of open literature on Emsec. The classic van Eck article [259] is still worth a read; and the only book on computer security (until this one) to have a chapter on the subject is by Deborah Russell and G.T. Gangemi [660]. Our recent work on Soft Tempest, Teapot, and related topics can be found in [478]. For power analysis, see the papers by Paul Kocher [467], and by Tom Messergues, Ezzy Dabish and Robert Sloan [547]; more papers are appearing regularly. Finally, Joel McNamara runs a comprehensive unofficial Tempest Web site at [542].