Computer Science Tripos part 2 — Security

Lecture 5 – Electrical Engineering Aspects

Ross J. Anderson

In this lecture, I will talk about some of the aspects of computer security that involve topics more commonly thought of as part of electrical engineering, such as signal processing, Tempest and the reverse engineering of silicon chips.

1 Information Hiding

We saw in the last lecture that in order to prevent selective service denial attacks, we may wish to make communications not merely confidential but anonymous. There are other applications, such as espionage, where we want our opponent to remain unaware that there is any traffic at all. The techniques for achieving such covertness of communication are generically referred to as steganography, which is Greek for 'hidden writing'. In fact, Herodotus tells us that in ancient Greece, messages were occasionally sent by tattooing them on a slave's head, allowing the hair to grow back, and then sending this slave as a messenger.

A more widespread method was invisible ink. During the middle ages, people discovered that if a message were written in milk, urine or a number of other fluids, the writing would remain invisible until the letter was heated, at which time it would become visible as a brown trace. The invisible ink could either be used to write the text itself, or to mark letters in a newspaper or book. The development of chemistry, and the spy mania in Europe at the turn of the century, led to further development, with various combinations of ink and developer in use; however, chemistry led during WW1 to the invention of universal developers which diminished the usefulness of this technology.

Also during WW1, the Germans developed the microdot, which provided the next two generations of spies with the tools of their trade. Documents were subjected to two or three successive photographic reductions, until the image could be cut out and placed over the full stop in a cover text such as a book. Microdots have been used into modern times, and throughout the cold war one hallmark of the KGB agent (for example) was the microdot reader concealed inside a tin of talcum powder.

Prearranged messages can also be used in low-bandwidth applications. For example, during WW2, transatlantic cable traffic was censored in Bermuda, and suspicious messages were often paraphrased to try and evoke a response. Once, a message from Germany to the USA saying 'Aunt Matilda is dead' was changed to 'Aunt Matilda is deceased', which prompted a panicky reply 'say again, is Aunt Matilda dead or deceased?'

1.1 Low-probability-of-intercept communications

More modern techniques include burst transmission systems, in which the spy or special forces unit commander composes his report, which is then compressed and transmitted at very high speed. British special forces have been reported using such equipment to send reports via satellite in the Falklands and the Gulf; the idea is that so long as the transmission is kept to under a second, it is hard for the other side to get a fix on the transmitter.

A similar technique is used in meteor scatter radio. This relies on the fact that billions of micrometeorites strike the Earth's atmosphere each day, and each will typically leave an ionisation trail many kilometres long, which persists for a third of a second or so and provides a temporary transmission path between a mother station and a 'footprint' which might be a hundred kilometres long and a few kilometres wide. The mother station transmits continuously, and whenever one of the daughter stations is within a footprint, it starts to send packets of data at high speed, which the mother acknowledges.

With the low power levels used in covert operations – and in civilian applications such as monitoring rainfall in remote parts of the third world – one can achieve an average data rate of about 50 bps, with an average delay of about 5 minutes and a range of 500–1500 miles. With the high power levels used in overt military applications, average data rates can rise into the tens of kilobits per second.

Other technical measures used to increase communication covertness include spread spectrum radio, in which the datastream is combined with a very fast pseudorandom keystream in order to increase the bandwidth. Under the right circumstances, the existence of the transmission may be undetectable to receivers which do not possess the key material needed to correlate and 'lock-on' to the keystream. Also, there is significant research being done in the use of optical systems such as infra-red lasers. Finally, most of the top-end bugs (radio microphones) in use – the kind that cost £10,000+ and are used by governments – are frequency hoppers, burst or both, in order to be harder to detect.

1.2 Law enforcement

The main technique used by criminals to evade surveillance by law enforcement officials is to hide their communications in innocuous traffic. In this application, encryption could be unwise as even if the cipher could not be broken, its existence would mark the traffic as worthy of attention. Few people realise that encryption can often make traffic less secure.

The NSA reckons that the drug cartels are its most formidable nongovernmental opponents, and their standard modus operandi is to hijack telephones and telephone accounts. In New York, for example, crooks get telephone credit card numbers by looking over the shoulders of people at phone boxes, and then use these to make calls overseas (these numbers are also used in 'call-sell' operations in which immigrants are charged a few dollars to phone home – the call-sell traffic providing extra cover for the narcotrafficantes). Another modus operandi is to observe the calling id of an analogue cellular telephone and reprogram another phone to steal this id; there is even a black market in telephones which assume a different identity for each outgoing call. In Paris, the modus operandi is to get a cordless telephone handset and drive around the smart suburbs until you get a dial tone. Near Newcastle, drug dealers tapped into a neighbour's phone line in order to make calls free of surveillance (and free of charge).

The spread of digital GSM phones, which use a proper authentication protocol to make cloning harder, caused crooks to move the attack upstream and buy phones using stolen credit cards. The most alarming recent development, from the police point of view, is the rapid spread of pre-paid mobiles, which have already been used for crimes such as stalking and extortion. In France, it is already mandatory that buyers of prepaid phones show an identity card to the dealer.

Even for intelligence agencies, the main problem is real-time traffic selection, rather than the subsequent cryptanalysis, translation, evaluation and distribution of intercepts. But technology is coming to their rescue. In the old days, many business telephone calls were protected by virtue of their being made from one large switchboard to another, and a sigint agency would have to be desperate to sift through all of a large company's telephone traffic. But now with direct inward dialing, personalised numbers – and of course electronic mail – the intelligence collection opportunities are multiplying rapidly.

Secure covert communications are possible if the communicating parties are prepared to pay a high enough price in bandwidth efficiency. For example, A CIA camera used a keystream generator to select a small number of pixels in an image; a covert ciphertext was embedded in the least significant bits of those pixels whose correlation with their neighbours was neither too high nor too low. We have produced demo software in which ciphertexts can be hidden in MP3 audio files. The possibilities are enormous.

1.3 Copyright marking

Over the last few years, interest in steganography has spread from the defence and intelligence communities to academia (and even Hollywood) with a surge of research on copyright marking. The problem is how to embed a hidden copyright message in digital objects (such as music, videos and software) which will resist intelligent attempts to remove it.

In the early 1980's, IBM successfully litigated against companies who had pirated parts of the code in its personal computer ROMs by using 'software birthmarks' such as the order in which registers were pushed and popped. Nowadays most code is compiled rather than hand assembled, so this is less useful. We might therefore tweak our compiler, or write an object code mangler to manipulate the executable file in ways that do not impair its functioning. However, an adversary could also write a mangler and hope to remove the marks we put in. Who would win?

When we are dealing with images, things become more complicated still, as an image may be subjected to lossy compression. Simply tweaking the least significant bits of a subset of pixels chosen by a secret key will not be enough, as such a mark would be erased by compression. So how do we create copyright marks that will survive any amount of mangling that is likely to leave the image valuable, and yet remain invisible to the naked eye? This is a hard problem, and the subject of current work. Most proposed solutions involve massaging the components of the image's discrete cosine transform in some way or another. So far, most of them can be broken (in the sense of preventing the marks being detected) by applying carefully chosen sub-perceptual distortions such as stretching, cropping and nonlinear filtering. Similarly, many audio marking schemes can be defeated by suitable filtering tricks. See the security group's web page on information hiding for more details.

2 Tempest

'Tempest' attacks involve reconstructing the screen image (or other useful information) from the stray radio frequency emissions from monitors and other components. In addition to passively listening in to whatever emissions are available, the attacker can also be active: it is possible to excite a terminal's keyboard cable with RF of an appropriate frequency and read the key codes which are modulated in the return signal.

In his book 'Spycatcher', former MI5 scientist Peter Wright recounts the origin of this technology. In 1960, Britain was negotiating to join the European Economic Community, and the Prime Minister was worried that the French president De Gaulle would block Britain's entry. He therefore asked the intelligence services to find out what the French position was. GCHQ tried to break the French diplomatic cipher system, but without success. However, Wright and his assistant Tony Sale (now the curator at Bletchley) noticed that the enciphered traffic carried a faint secondary signal, and constructed equipment to recover it. It turned out to be the plaintext, which somehow leaked through the cipher machine. The trick was then used against the Russians in East Berlin, but was betrayed by the double agent George Blake.

Electromagnetic surveillance is often referred to as 'Tempest' which appears to have been an acronym for 'transient electromagnetic pulse emission standards'; the cover story was that this referred to shielding military equipment against the electromagnetic pulse from a nuclear explosion. In modern usage, 'Tempest' refers to the passive interception of stray information-bearing radio frequency emissions from computer and communications equipment, while the kind of electromagnetic leakage observed by Wright is called 'Hijack'. There is a third category, 'Nonstop', which refers to secondary RF emissions; for example, if a cipher machine is placed next to a radio transmitter, a wire in it may pick up the signal and re-radiate it, modulated with the information that they are carrying (e.g., because of the nonlinearity of a transistor which terminates the wire). For this reason, it is against US military regulations to use a mobile phone within 5 metres of a cipher machine.

The open community became aware of Tempest in 1985 with the publication by Wim van Eck from the Netherlands PTT of a paper describing snooping on a PC monitor. This sent a wave of fear through the computer security community, but there was no realistic solution on offer. Shielding, whether of rooms or of equipment, is expensive. More recently, we have developed software techniques which enable us to reduce the amount of energy radiated by the VDU, by filtering out the top one third or so of the frequency spectrum from the graphics we display on it.

Although the Tempest threat affects all kinds of electronic information processing equipment, from computers through fax machines to digital telephones, it can be understood most simply in the context of the stray RF emissions from a vdu. In such a device, the beam current is modulated with a signal consisting of 0's, 1's and perhaps greyscale values as well. These constitute a square wave whose fundamental frequency is the dot rate $2\pi\omega$, and which therefore has a Fourier expansion of the form

$$I = \sum_{n=0}^{K} a_n \sin(n\omega t) \tag{1}$$

With a square wave, one expects that $a_n = O(\frac{1}{n})$, but as higher frequencies radiate better, the received signals should have $a'_n = O(1)$. Given that the dot rate is typically about 30MHz in a PC monitor, this means that we might expect a series of harmonics at 30MHz, 60MHz, 90MHz and so on up to a cut-off point of perhaps 900MHz at the Nyqvist frequency $2\pi nK$.

However, in a real system the power spectrum is not so simple, as parasitic resonances amplify particular harmonics at the expense of others. It is not uncommon to find that a typical PC's vdu cable, which is about 1 metre long, gives rise to a strong harmonic at about 150MHz, with other usable signals at (say) 75MHz, 300MHz and 450 MHz. These harmonics can be measured in a Faraday cage using a spectrum analyser. One will usually see other peaks in the spectrum which do not carry much useful information; these may be generated in the power supply and in the system unit, with the latter being stronger in more recent PCs due to the higher clock frequency.

We have often been asked whether organisations with a large number of identical PCs gain any safety in numbers. The answer is no. Polarisation and the variation in signal strength are only two means to discriminate between nearby machines; in addition, manufacturing variations give similar PCs quite different RF characteristics. It is also no defence to switch to LCD screens or to use wire mesh screen covers; in order to suppress the signal, one has to completely shield either the machine or the environment. Both are too expensive in most environments.

As the power dissipated by a vdu will typically be in the dozens of watts, and a radio signal of this strength can be received at a distance of dozens or even hundreds of kilometers, there is clearly a potential threat. Empirical investigation shows that detecting vdu signals is easy at short range, but is generally quite difficult at longer ranges except in a quiet RF environment. In town, the background noise level varies a lot during the working day, and signals which are clearly visible from ten metres at 0830 may have disappeared completely by 1000. Empirically, we have also found that filtering out the high frequencies from display graphics significantly reduces the strength of the signal. This may be the pragmatic approach where the signal strength is already marginal; if it forces the attacker to be inside the building rather than across the street, it may be enough.

The flip side of filtered fonts is the 'Tempest virus'. This uses the fact that while one can make text that is visible to the user much less visible to the man in the grey van, one can also do the reverse and encode information in a screen image which is hard for the user to see but nonetheless radiates strongly. A virus with such a payload could steal private information such as a PGP key and then radiate it unobtrusively.

However, the mainstream Tempest threat model remains one of attacks on governments by other governments. The attacks are labour intensive, and thus expensive, and require close proximity to the target, which might typically be an embassy. In the diplomatic community, the common solution is to use shielded rooms to house classified documents and hold secret meetings. Intelligence services may use completely shielded buildings, and some military units rely on remoteness. Such precautions also provide some protection against bugs of the radio microphone type.

3 Tamper Resistance

There are a number of computing devices that we may wish to make tamper resistant, such as smartcards, the cryptographic hardware used by banks to manage customer PINs and ATM keys, the TCB of multilevel secure systems evaluated to B3 and above, and any objects that we wish to protect against reverse engineering.

This raises the question of how tamper resistance is to be measured. There is a shortage of effective standards, but the classification used by IBM is as follows:

1. Grade one tamper resistance should protect against casual attack by unskilled individuals. An example is the hardware security modules used in the banking industry, which typically have lid switches that interrupt power to the key memory when the device is opened. Such devices can be easily defeated by a service engineer, who would bypass a switch at one visit and then obtain key material at his next visit;

2. Grade two tamper resistance should protect against attack by knowledgable individuals. IBM's 47xx series cryptographic units are aimed at this level; they contain a crypto chip, a processor, a battery and memory in an epoxy module that contains many fine copper wires. Most penetration techniques would cut at least one of these wires, resulting in the keys being overwritten;

3. Grade three tamper resistance should protect against attacks by funded organisations. Examples of systems designed to provide this level of protection include the Clipper chip, which used a special silicon technology in which ROM bits are implemented as fusible links, and the key memory in nuclear command and control systems which is typically destroyed by a shaped charge in the event of tampering.

The advance of silicon technology is posing serious problems for the less aggressive forms of tamper resistance, as more effective nondestructive test methods are constantly developed by the chip testing community. The attached article, "Tamper Resistance – A Cautionary Note" surveys the state of the art in 1996. The main lesson is that tamper resistance is a lot harder than most people think. Where an opponent can get unsupervised access to a piece of equipment, it is prudent to assume that all secrets in it will eventually be read out.

In addition, deleted data can often be recovered from both magnetic and silicon storage using appropriate tools. Magnetic force microscopes, which are used to test disk drives, can often recover at least one previous generation of stored data from them, and data that have been kept on the disk for a long time may never be properly removed in normal operation. (The head alignment may have shifted slightly, leaving a trace of that the write head can no longer overwrite.)

Both static and dynamic RAM are also vulnerable. Data can be 'burned in' to SRAM in the sense that storing the same data for a long period of time will make this the preferred power-up state. Even volatile data may be recovered, as it can take some time for charge to dissipate, and cryogenic temperatures can extend this time to minutes or even hours. With DRAM, the thin oxide which forms the storage capacitor dielectric is stressed by an applied field and ionic contaminants such as sodium migrate; the effect is to change the cell's threshold. Some DRAM chips have (often undocumented) test modes in which a low voltage can be applied to measure this, and these can be used to read out data that one might have thought lost.

For these reasons, the secure destruction of old cryptographic equipment and of computer equipment that has contained highly classified data is an important lifecycle management consideration. If one cannot rely on eventual secure destruction, then one ought to assume that an opponent will get unsupervised access to the device, and design it accordingly.

Designers' reactions to these threats vary. Military organisations tend to go for defence in depth, in which a government's master keys are hidden in a tamper resistant chip, inside a safe with a self-destruct charge, which in turn is in a vault in a closely guarded building in a large military base. Banks take similar although less elaborate precautions; their cryptographic hardware will typically be kept in the main computer room, and anyone who could get access to that could in principle alter his account balance directly from the console rather than having to break any cryptographic security mechanisms.

The approach of academic researchers in recent years has been to try to use mathematics instead of metal. It is possible, for example, to distribute one's security critical processing between seven servers in such a way that any four of them can perform a transaction, while up to two may be subverted without harm – in the sense that the attack will be detected, and the faulty processors can be removed from the group and replaced with others by the cooperation of the five processors that are still secure. We will spend the next few lectures discussing the cryptographic techniques which make solutions like this possible.

4 Further Reading

Donald Schilling's *'Meteor Burst Communications — Theory and Practice'* is the standard reference work on that subject, and describes the US Air Force's network in Alaska. Dixon's *Spread Spectrum Systems with Commercial Applications'* is the standard work in that field.

The early history of Tempest is in *'Spycatcher'*, by Peter Wright; the most cited paper on the subject is W van Eck, "Electromagnetic Radiation from Video Display Units: An Eavesdropping Risk?" in *Computers and Security* **4** no 4 (Dec 1985) pp 269 - 286. The reference article on recovering deleted data is "Secure Deletion of Data from Magnetic and Solid-State Memory", Peter Gutman, in *Usenix Security 96* pp 77–89.

5 Problems

1. Alice wants to send secret messages to Bob by email. However, all her traffic is read by her jealous husband, Willie. Willie has engaged you to assist him in proving Alice's infidelity, so that he can get an advantageous divorce settlement. We have already discussed a number of ways in which a ciphertext can be hidden in a covertext – for example, in nonprinting

characters. Describe as many of the signalling channels in email as you can, and estimate their bandwidths.

2. How would you go about hiding a message in a .gif file? In a jpeg file? In an mpeg datastream? Estimate the covert bandwidth in each case.

3. Would you expect a fax machine to be more vulnerable to Tempest attack when it is transmitting or when it is receiving?

4. Measure the redundancy of Intel object code, and estimate how much of it could be used for hiding information.

5. Assume that any smartcard system can be broken at a cost of £10,000 for the first example of any given mask to be fielded and thereafter £500 to extract keys from any specific instance of the card (these are approximately reasonable price estimates for 1996). What are the effects of this on the use of smartcards in (a) SIM cards for GSM phones (b) memory cards for pay telephones (c) subscriber cards for pay-TV (d) bank cards for online transactions (e) bank cards for offline transactions?

6. Estimate the signal strength that can be received by an attacker who lays down a capacitive pickup line on top of a bus line in a smartcard chip. Assume the bus line is 1 micron wide and 2 mm long, and separated from the probe line by 0.2 microns of silicon nitride passivation. Assume a 3.5MHz clock with 0.6 micron technology giving a rise time of 8 ns at 3.5V. Estimate also the load that this coupling would impose on the bus drivers. Would it be greater or less than the load of an extra inverter?