

Monitoring and Metering

The market is not an invention of capitalism. It has existed for centuries. It is an invention of civilization.

— Mikhail Gorbachev

12.1 Introduction

Many secure systems are concerned with monitoring and metering the environment. They go back a long way. James Watt, the inventor of the steam engine, licensed his patents using a sealed counter that measured the number of revolutions an engine had made; his inspectors read these from time to time and billed the licensee for royalties.

Electronic systems that use cryptography and tamper-resistance are rapidly displacing older mechanical systems, and also opening up all sorts of new applications. Ticketing is a huge application, from transport tickets through sports tickets to theatre tickets; my case study for ticketing is the meters used for utilities such as gas and electricity. Then I'll turn to vehicle systems; the most familiar of these may be taxi meters but I'll mainly discuss tachographs — devices used in Europe to record the speed and working hours of truck and coach drivers, and in the USA to record the comings and goings of bank trucks. My third case study is the electronic postage meter used to frank letters.

You will recall that in order to defeat a burglar alarm it is sufficient to make it appear unreliable. Such service-denial attacks can be tricky enough to deal with; meters add further subtleties.

When we discussed an alarm in a bank vault, we were largely concerned with attacks on communications (though sensor defeats also matter).

But many metering systems are much more exposed physically. A taxi driver (or owner) may want the meter to read more miles or more minutes than were actually worked, so he may manipulate its inputs or try to disrupt it so that it over-measures. With tachographs, it's the reverse: the truck driver usually wants to drive above the speed limit, or work dangerously long hours, so he wants the tachograph to ignore some of the driving. Utility consumers similarly have a motive to cause their meters to ignore some of the passing electricity. In these the attacker can either cause the device to make false readings, or simply to fail. There are also markets for bad people who can sell exploits, whether by forging tickets for electricity meters or selling devices that can be installed in vehicles to deceive a taxi meter or tachograph.

In many metering and vehicle monitoring systems (as indeed with nuclear verification) we are also concerned with evidence. An opponent could get an advantage either by manipulating communications (such as by replaying old messages) or by falsely claiming that someone else had done so. As for postal franking systems, it's not sufficient for the attacker to cause a failure (as then he can't post his letters) but the threat model has some interesting twists; the post office is mostly concerned with stopping wholesale fraud, such as crooked direct marketers who bribe postal employees to slip a truckload of mail into the system. It's thus directed internally more than externally.

Metering systems also have quite a lot in common with systems designed to enforce the copyright of software and other digital media, which I will discuss in a later chapter.

12.2 Prepayment Meters

Our first case study comes from prepayment metering. There are many systems where the user pays in one place for a token — whether a magic number, or a cardboard ticket with a magnetic strip, or even a rechargeable token such as a smartcard — and uses this stored value in some other place.

Examples include the stored-value cards that operate photocopiers in libraries, lift passes at ski resorts, and washing machine tokens in university halls of residence. Many transport tickets are similar — especially if the terminals which validate the tickets are mounted on buses or trains and so are not usually online.

The main protection goal in these systems is to prevent the stored-value tokens being duplicated or forged en masse. Duplicating a single subway ticket is not too hard, and repeating a magic number a second time is trivial. This can be made irrelevant if we make all the tokens unique and log their use at both ends. But things get more complicated when the device that accepts the token does not have a channel of communication back to the ticket issuer, so all the replay and forgery detection must be done offline — in a terminal that

is often vulnerable to physical attack. So if we simply encipher all our tokens using a universal master key, a villain could extract it from a stolen terminal and set up in businesses selling tokens.

There are also attacks on the server end of things. One neat attack on a vending card system used in the staff canteen of one of our local supermarkets exploited the fact that when a card was recharged, the vending machine first read the old amount, then asked for money, and then wrote the amended amount. The attack was to insert a card with some money in it, say £49, on top of a blank card. The top card would then be removed and a £1 coin inserted in the machine, which would duly write £50 to the blank card. This left the perpetrator with two cards, with a total value of £99. This kind of attack was supposed to be prevented by two levers that extended to grip the card in the machine. However, by cutting the corners off the top card, this precaution could easily be defeated (see Figure 12.1) [749]. This attack is interesting because no amount of encryption of the card contents will make any difference. Although it could in theory be stopped by keeping logs at both ends, they would have to be designed a bit more carefully than is usual.

But we mustn't get carried away with neat tricks like this, or we risk getting so involved with even more clever countermeasures that we fall prey to the Titanic Effect again by ignoring the system level issues. In most ticketing systems, petty fraud is easy. A free rider can jump the barrier at a subway station; an electricity meter can have a bypass switch wired across it; things like barcoded ski lift passes and parking lot tickets can be forged with a




Figure 12.1: Superposing two payment cards

scanner and printer. The goal is to prevent fraud becoming systematic. So petty fraud should be at least slightly inconvenient and — more importantly — there should be more serious mechanisms to prevent anyone forging tickets on a large enough scale to develop a black market that could affect your client's business.

The first example I'll discuss in detail is the prepayment electricity meter. I chose this because I was lucky enough to consult on a project to electrify three million households in South Africa (a central election pledge made by Nelson Mandela when he took power). This work is described in some detail in [59]. Most of the lessons learned apply directly to other ticketing systems.

12.2.1 Utility Metering

In a number of European countries, householders who can't get credit (because they are on welfare, have court judgements against them, or whatever) buy gas and electricity services using prepayment meters (Figure 2.2). In the old days they were coin-operated, but the costs of coin collection led vendors to develop token-based meters instead. The UK now has 3.6 million electricity meters and 2 million gas meters. In South Africa, development was particularly rapid because of a national priority project to electrify the townships; as many of the houses were informally constructed, and the owners did not even have addresses (let alone credit ratings), prepayment was the only way to go. There are now 5.5 million of these meters in use in South Africa, which has exported 1.5 million to other countries in Africa, Latin America and elsewhere.

The customer goes to a shop and buys a token, which may be a smartcard, or a disposable cardboard ticket with a magnetic strip, or even just a magic number. Of the UK's electricity meters, 2.4 million use smartcards¹ and 1.2 million use magnetic tickets. Most of South Africa's meters use a magic number. This is perhaps the most convenient for the customer, as no special vending apparatus is required: a ticket can be dispensed at a supermarket checkout, at an ATM, or even over the phone.

The token is really just a string of bits containing one or more instructions, encrypted using a key unique to the meter, which decodes them and acts on them. Most tokens say something like 'meter 12345 — dispense 50KWh of electricity!' The idea is that the meter will dispense the purchased amount and then interrupt the supply. Some tokens have engineering functions too. For example, if the power company charges different rates for the daytime and evening, the meter may have to know the relative prices and the times at which the tariffs change. Special tokens may be used to change these, and to change

¹1.6 million of these smartcards are repackaged in plastic keys; the other 0.8 are normal smartcards. The packaging may improve usability, especially in a darkened house, but does not affect the cryptographic security.



Figure 12.2: A prepayment electricity meter (Courtesy of Schlumberger)

keys. The meters that use smartcards are able to report consumption patterns, tampering attempts and so on back to the power company; however the magnetic-ticket and magic-number meters do not have such a back channel.

The manufacture of these meters has become big business. Growth in the third world is strong: prepayment metering was the only way the government in South Africa could meet its election pledge to electrify millions of homes quickly. In the developed world, the main impetus for prepayment metering is reducing administrative costs. Electric utilities find that billing systems can devour 20 percent of retail customer revenue, when you add up the costs of meter reading, billing, credit control, bad debts and so on. Prepayment systems typically cost under 10 percent: the shop that sells the tokens gets five percent, while the meters and the infrastructure cost about the same again.

12.2.2 How the System Works

The security requirements for prepayment meters seem straightforward. Tokens should not be easy to forge, while genuine tokens should not work in the wrong meter, or in the right meter twice. One strategy is to make tokens tamper-resistant, by using smartcard chips of some kind or another;

the alternative is to tie each token to a unique meter, so that someone can't use the same magic number in two different meters, and also make each token unique using serial numbers or random numbers, so that the same token can't be used twice in the same meter. But it has taken a surprising amount of field experience to develop the idea into a robust system.

The meter needs a cryptographic key to authenticate its instructions from the vending station. The original system had a single vending machine for each neighbourhood, usually located in a local store. The machine has a vend key K_V which acts as the master key for a neighborhood and derives the device key when needed by encrypting the meter ID under the vend key:

$$K_{ID} = \{ID\}_{K_V}$$

This is the same key diversification technique described for parking lot access devices in Chapter 3. Diversifying the vend key K_V to a group of meter keys K_{ID} provides a very simple solution where all the tokens are bought locally. However, once the system rolled out, we found that real life was less straightforward. In Britain, deregulation of the electricity industry led to a multitude of electricity companies who buy power from generators and sell it onward to households through a common infrastructure, so metering systems must support multiple power companies with different tariff structures. In South Africa, many people commute long distances from townships or homelands to their places of work, so they are never at home during business hours and want to buy tickets where they work. So we had to support multiple retailers, by letting customers register at an out-of-area vending station. This meant protocols to send a customer meter key from the vending station that 'owns' the meter to another station, and to pass sales data in the opposite direction for balancing and settlement, somewhat like in ATM networks. The most recent development (2007) is online vending; a customer can buy a magic number over the Internet or via their mobile phone from a central token server. This server can deal directly with four million customers and also about 10,000 online vend points such as ATMs.

Statistical balancing is used to detect what are euphemistically known as *non-technical losses*, that is, theft of power through meter tampering or unauthorized direct connections to mains cables. The mechanism is to compare the readings on a feeder meter, which might supply 30 houses, with token sales to those houses. This turns out to be harder than it looks. Customers hoard tickets, meter readers lie about the date when they read the meter, and many other things go wrong. Vending statistics are also used in conventional balancing systems, like those discussed in Chapter 10.

There have been a few cases where vending machines were stolen and used to sell tokens in competition with the utility. These 'ghost vendors' are extremely difficult to track down, and the early ones generally stayed in business until the keys in all the meters were changed. The countermeasure has

been to get the vending machines to maintain a credit balance in the tamper-resistant security processor that also protects vend keys and foreign meter keys. The balance is decremented with each sale and only credited again when cash is banked with the local operating company; the company then sends a magic number that reloads the vending machine with credit. The operating company in turn has to account to the next level up in the distribution network, and so on. So here we have an accounting system enforced by a value counter at the point of sale, rather than just by ledger data kept on servers in a vault. Subversion of value counters can in theory be picked up by statistical and balancing checks at higher layers.

This distribution of security state is seen in other applications too, such as in some smartcard-based electronic purse schemes. However, the strategic direction for power vending is now believed to be centralisation. Now that the communications infrastructure is more dependable, many of the original 1200 vending machines will be replaced by online vending points that get their tokens in real time from the central service. (In banking, too, there is a move away from offline operation as communications get more dependable.)

So what can go wrong?

12.2.3 What Goes Wrong

Service denial remains an important issue. Where there is no return channel from the meter to the vending station, the only evidence of how much electricity has been sold resides in the vending equipment itself. The agents who operate the vending machines are typically small shopkeepers or other township entrepreneurs who have little capital so are allowed to sell electricity on credit. In some cases, agents who couldn't pay the electricity bill to the operating company at the end of the month just dumped their equipment and claimed that it had been stolen. This is manageable with small agents, but when an organization such as a local government is allowed to sell large amounts of electricity through multiple outlets, there is definitely an exposure. A lot of the complexity was needed to deal with untrustworthy (and mutually mistrustful) principals.

As with burglar alarms, environmental robustness is critical. Apart from the huge range of temperatures (as variable in South Africa as in the continental United States) many areas have severe thunderstorms: the meter is in effect a microprocessor with a 3-kilometer lightning conductor attached.

When meters were destroyed by lightning, the customers complained and got credit for the value they said was still unused. So their next step was to poke live mains wires into the meter to try to emulate the effects of the lightning. It turned out that one make of meter would give unlimited credit if a particular part of the circuitry (that lay under the token slot) was destroyed. So service denial attacks worked well enough to become popular.

It was to get worse. The most expensive security failure in the program came when kids in Soweto observed that when there was a brown-out — a fall in voltage from 220 to 180 volts — then a particular make of meter went to maximum credit. Soon kids were throwing steel chains over the 11KV feeders and crediting all the meters in the neighborhood. This was the fault of a simple bug in the meter ROM, which wasn't picked up because brown-out testing hadn't been specified. The underlying problem was that developed-country environmental standards were inadequate for use in Africa and had to be rewritten. The effect on the business was that 100,000 meters had to be pulled out and re-ROMmed; the responsible company almost went bust.

There were numerous other bugs. One make of meter didn't vend a specified quantity of electricity, but so much worth of electricity at such-and-such a rate. It turned out that the tariff could be set to a minute amount by vending staff, so that it would operate almost for ever. Another allowed refunds, but a copy of the refunded token could still be used (blacklisting the serial numbers of refunded tokens in subsequent token commands is hard, as tokens are hoarded and used out of order). Another only remembered only the last token serial number entered, so by alternately entering duplicates of two tokens it could be charged up indefinitely.

As with cash machines, the real security breaches resulted from bugs and blunders, which could be quite obscure, but were discovered by accident and exploited in quite opportunistic ways. These exploits were sometimes on a large scale, costing millions to fix.

Other lessons learned, which we wrote up in [59], were:

- prepayment may be cheap so long as you control the marketing channel, but when you try to make it even cheaper by selling prepayment tokens through third parties (such as banks and supermarkets) it can rapidly become expensive, complicated and risky. This is largely because of the security engineering problems created by mutual mistrust between the various organizations involved;
- changes to a business process can be very expensive if they affect the security infrastructure. For example, the requirement to sell meter tokens at distant shops, to support commuters, was not anticipated and was costly to implement;
- recycle technology if you can, as it's likely to have fewer bugs than something designed on a blank sheet of paper. Much of what we needed for prepayment metering was borrowed from the world of cash machines;
- use multiple experts. One expert alone can not usually span all the issues, and even the best will miss things;

- no matter what is done, small mistakes with large consequences will still creep in. So you absolutely need prolonged field testing. This is where many errors and impracticalities will first make themselves known.

Meters are a good case study for ticketing. Transport ticketing, theater ticketing and sports ticketing may be larger applications, but I don't know of any serious and publicly available studies of their failure modes. In general the end systems — such as the meters or turnstiles — are fairly soft, so the main concern is to prevent large scale fraud. This means paying a lot of attention to the intermediate servers such as vending machines, and hardening them to ensure they will resist manipulation and tampering. In the case of London transport tickets, deregulation of the railways led to reports of problems with train companies manipulating ticket sales by booking them at stations where they got a larger percentage of the takings, and clearly if you're designing a system that shares revenue between multiple vendors, you should try to think of how arbitrage opportunities can be minimised. One still does what one economically can to prevent the end users developing systematic attacks on the end systems that are too hard to detect.

I'll now look at a class of applications where there are severe and prolonged attacks on end systems which must therefore be made much more tamper resistant than electricity meters. The threat model includes sensor manipulation, service denial, accounting fiddles, procedural defeats and the corruption of operating staff. This exemplary field of study is vehicle monitoring systems.

12.3 Taxi Meters, Tachographs and Truck Speed Limiters

A number of systems are used to monitor and control vehicles. The most familiar is probably the odometer in your car. When buying a used car you'll be concerned that the car has been *clocked*, that is, had its indicated mileage reduced. As odometers become digital, clocking is becoming a type of computer fraud; a conviction has already been reported [274]. A related problem is *chipping*, that is, replacing or reprogramming the engine controller. This can be done for two basic reasons. First, the engine controller acts as the server for the remote key-entry systems that protect most modern cars from theft, as described in Chapter 3; so if you want to steal a car without stealing the key, the engine controller is the natural target (you might replace the controller in the street, or else tow the car and replace or reprogram the controller later). Second, people reprogram their cars' engine controllers to make them go faster, and the manufacturers dislike this because of the increased warranty claims from burned-out engines. So they try to make the controllers more tamper resistant, or at least tamper-evident.

This fascinating arms race is described in [426]. Some vehicles now keep logs that are uploaded to the manufacturer during servicing. General Motors started equipping some vehicles with black boxes to record crash data in 1990. By the time the logging became public in 1999, some six million vehicles had been instrumented, and the disclosure caused protests from privacy activists [1282]. Indeed, there's now a whole conference, ESCAR, devoted to electronic security in cars.

There are a number of monitoring systems separate from those provided by the engine manufacturer, and the most familiar may be the taxi meter. A taxi driver has an incentive to manipulate the meter to show more miles travelled (or minutes waited) if he can get away with it. There are various other kinds of 'black box' used to record the movement of vehicles from aircraft through fishing vessels to armored bank trucks, and their operators have differing levels of motive for tampering with them. A recent development is the black boxes supplied by insurers who sell 'pay-as-you-drive' insurance to young and high-risk drivers; these boxes contain satellite navigation devices that let the insurer charge a couple of pennies a mile for driving along a country road in the afternoon but a couple of dollars a mile for evening driving in an inner city [1264]. It's conceivable that within a few years this will be the only type of insurance available to many youngsters; if the dangerous drivers flock to any flat-rate contracts still on offer, they may become unaffordable. In that case, any young man who wants to impress girls by driving around town on a Saturday night will have a strong incentive to beat the black box.

12.3.1 The Tachograph

The case study I'm going to use here is the tachograph. These are devices used to monitor truck drivers' speed and working hours; they have recently been the subject of a huge experiment in Europe, in which old analogue devices are being replaced by digital ones. This gives us some interesting data on how such equipment works, and can fail; and it's an example of how a move to digital technology didn't necessarily make things better. It contains useful warnings for engineers trying to modernise systems that do analogue monitoring in hostile environments.

Vehicle accidents resulting from a driver falling asleep at the wheel cause several times more accidents than drunkenness (20 percent versus 3 percent of accidents in the UK, for example). Accidents involving trucks are more likely to lead to fatal injuries because of the truck's mass. So most countries regulate truck drivers' working hours. While these laws are enforced in the USA using weigh stations and drivers' log books, countries in Europe use tachographs that record a 24-hour history of the vehicle's speed. Until 2005–6, this was recorded on a circular waxed paper chart (Figure 12.3); since then,




Figure 12.3: A tachograph chart

digital tachographs have been introduced and the two systems are currently running side-by-side². Eventually the analogue systems will be phased out; in the meantime they provide an interesting study of the relative strengths and weaknesses of analogue and digital systems. First let's look at the old analogue system, which is still used in most trucks on Europe's roads.

The chart is loaded into the tachograph, which is part of the vehicle's speedometer/odometer unit. It turns slowly on a turntable inside the instrument and a speed history is inscribed by a fine stylus connected to the speedometer. With some exceptions that needn't concern us, it is an offence to drive a truck in Europe unless you have a tachograph; if it's analogue you must have a chart installed, and have written on it your starting time and location. You must also keep several days' charts with you to establish that you've complied with the relevant driving hours regulations (typically 8.5 hours per day with rules for rest breaks per day and rest days per week). If it's digital you have to have a driver card plugged into it; the card and the vehicle unit both keep records.

²Vehicles registered since August 2004 in the UK have had to have digital systems fitted, cards have been issued since June 2005 and since August 2006 the use of digital systems in new vehicles has been mandatory; the dates vary slightly for other EU countries.

European law also restricts trucks to 100 Km/h (62 mph) on freeways and less on other roads. This is enforced not just by police speed traps and the tachograph record, but directly by a speed limiter that is also driven by the tachograph. Tachograph charts are also used to investigate other offences, such as unlicensed toxic waste dumping, and by fleet operators to detect fuel theft. So there are plenty reasons why a truck driver might want to fiddle his tachograph. Indeed, it's a general principle in security engineering that one shouldn't aggregate targets. So NATO rules prohibit money or other valuables being carried in a container for classified information — you don't want someone who set out to steal your regiment's payroll getting away with your spy satellite photographs too. Forcing a truck driver to defeat his tachograph in order to circumvent his speed limiter, and vice versa, was a serious design error — but one that's now too entrenched to change easily.

Most of what we have to say applies just as well to taxi meters and other monitoring devices. While the truck driver wants his vehicle to appear to have gone less distance, the taxi driver wants the opposite. This has little effect on the actual tampering techniques.

12.3.2 What Goes Wrong

According to a 1998 survey of 1060 convictions of drivers and operators [45], the offences were distributed as follows.

12.3.2.1 How Most Tachograph Manipulation Is Done

About 70% of offences that result in conviction do not involve tampering but exploit procedural weaknesses. For example, a company with premises in Dundee and Southampton should have four drivers in order to operate one vehicle per day in each direction, as the distance is about 500 miles and the journey takes about 10 hours — which is illegal for a single driver to do every day. The standard fiddle is to have two drivers who meet at an intermediate point such as Penrith, change trucks, and insert new paper charts into the tachographs. So the driver who had come from Southampton now returns home with the vehicle from Dundee. When stopped and asked for his charts, he shows the current chart from Penrith to Southampton, the previous day's for Southampton to Penrith, the day before's for Penrith to Southampton, and so on. In this way he can give the false impression that he spent every other night in Penrith and was thus legal. This (widespread) practice, of swapping vehicles halfway through the working day, is called *ghosting*. It's even harder to detect in mainland Europe, where a driver might be operating out of a depot in France on Monday, in Belgium on Tuesday and in Holland on Wednesday.

Simpler frauds include setting the clock wrongly, pretending that a hitch-hiker is a relief driver, and recording the start point as a village with a very common name — such as ‘Milton’ in England or ‘La Hoya’ in Spain. If stopped, the driver can claim he started from a nearby Milton or La Hoya.

Such tricks often involve collusion between the driver and the operator. When the operator is ordered to produce charts and supporting documents such as pay records, weigh station slips and ferry tickets, his office may well conveniently burn down. (It’s remarkable how many truck companies operate out of small cheap wooden sheds that are located a safe distance from the trucks in their yard.)

12.3.2.2 Tampering with the Supply

The next largest category of fraud, amounting to about 20% of the total, involves tampering with the supply to the tachograph instrument, including interference with the power and impulse supply, cables and seals.

Old-fashioned tachographs used a rotating wire cable — as did the speedometers in cars up until the early 1980s — that was hard to fiddle with; if you jammed the truck’s odometer it was quite likely that you’d shear off the cable. More recent analogue tachographs are ‘electronic’, in that they use electric cables rather than rotating wire. The input comes from a sensor in the gearbox, which sends electrical impulses as the prop shaft rotates. This has made fiddling much easier! A common attack is to unscrew the sensor about a tenth of an inch, which causes the impulses to cease, as if the vehicle were stationary. To prevent this, sensors are fixed in place with a wire and lead seal. Fitters are bribed to wrap the wire anticlockwise rather than clockwise, which causes it to loosen rather than break when the sensor is unscrewed. The fact that seals are issued to workshops rather than to individual fitters complicates prosecution.

But most of the fiddles are much simpler still. Drivers short out the cable or replace the tachograph fuse with a blown one. (One manufacturer tried to stop this trick by putting the truck’s anti-lock braking system on the same fuse. Many drivers preferred to get home sooner than to drive a safe vehicle.) Again, there is evidence of a power supply interruption on the chart in Figure 12.3: around 11 A.M., there are several places where the speed indicated in the outside trace goes suddenly from zero to over 100 km/h. These indicate power interruptions, except where there’s also a discontinuity in the distance trace. There, the unit was open.

12.3.2.3 Tampering with the Instrument

The third category of fraud is tampering with the tachograph unit itself. This amounts for some 6% of offences, but declined through the 1990s as tampering

with digital communications is much easier than tampering with a rotating wire cable used to be. The typical offence in this category is miscalibration, usually done in cahoots with the fitter but sometimes by the driver defeating the seal on the device.

12.3.2.4 High-Tech Attacks

The state of the tampering art at the time of the 1998 survey was the equipment in Figure 12.4. The plastic cylinder on the left of the photo is marked 'Voltage Regulator — Made in Japan' and is certainly not a voltage regulator. (It actually appears to be made in Italy.) It is spliced into the tachograph cable and controlled by the driver using the remote control key fob. A first press causes the indicated speed to drop by 10%, a second press causes a drop of 20%, a third press causes it to fall to zero, and a fourth causes the device to return to proper operation.

This kind of device amounted for under 1% of convictions but its use is believed to be much more widespread. It's extremely hard to find as it can be hidden at many different places in the truck's cable harness. Police officers who stop a speeding truck equipped with such a device, and can't find it, have difficulty getting a conviction: the sealed and apparently correctly calibrated tachograph contradicts the evidence from their radar or camera.




Figure 12.4: A tachograph with an interruptor controlled by the driver using a radio key fob. (Courtesy of Hampshire Constabulary, England)

12.3.3 The Digital Tachograph Project

The countermeasures taken against tachograph manipulation vary by country. In Britain, trucks are stopped at the roadside for random checks by vehicle inspectors, and particularly suspect trucks may be shadowed across the country. In the Netherlands, inspectors prefer to descend on a trucking company and go through their delivery documents, drivers' timesheets, fuel records etc. In Italy, data from the toll booths on the freeways are used to prosecute drivers who've averaged more than the speed limit (you can often see trucks parked just in front of Italian toll booths). But such measures are only partially effective, and drivers can arbitrage between the differing control regimes. For example, a truck driver operating between France and Holland can keep his documents at a depot in France where the Dutch vehicle inspectors can't get at them.

So the European Union took the initiative to design a unified electronic tachograph system to replace the existing paper-based charts with smartcards. Each driver can now get a 'driver card' that contains a record of his driving hours over the last 28 days. Every new vehicle has a vehicle unit that can hold a year's history. There are two further types of credential: workshop cards used by mechanics to calibrate devices, and control cards used by law enforcement officers to read them out at the roadside. In 1998, I was hired by the UK Department of Transport to look at the new scheme and try to figure out what would go wrong. After talking to a wide range of people from policemen and vehicle inspectors to tachograph vendors and accident investigators, I wrote a report [45]. I revisited the field in 2007 when writing the second edition of this book; it was simultaneously pleasing and depressing to find that I'd mostly predicted the problems correctly. However a few interesting new twists also emerged.

The most substantial objection raised to the project was that it was not clear how going digital will help combat the procedural frauds that make up 70% of the current total. Indeed, our pair of drivers 'ghosting' between Dundee and Southampton will have their lives made even easier. It will take maybe ten years — the lifetime of a truck — to change over to the new system and meantime a crooked company can run one new digital truck and one old analogue one. Each driver will now have one chart and one card, with five hours a day on each, rather than two charts which they might accidentally mix up when stopped. This has turned out to be well-founded. In the UK, it's now estimated that 20% of the vehicle fleet has digital tachographs — somewhat more than would be expected — which suggests that operators have been installing digital devices before they need to as they're easier to fiddle. (So, in the short term at least, the equipment vendors appear to be profiting from the poor design of the system: in the medium term they may well profit even more, if European governments decide on yet another technology change.)

Another objection was that enforcement would be made harder by the loss of detailed speed and driving hours information. Back in 1998, the Germans had wanted the driver card to be a memory device so it could contain detailed records; the French insisted on a smartcard, as they're proud of their smartcard industry. So the driver card has only 32K of memory, and can only contain a limited number of alarm events. (Indeed until a late change in the regulations it didn't contain data on rest periods at all.) The choice of a smartcard rather than a memory card was probably the most critical wrong decision in the whole programme.

12.3.3.1 System Level Problems

The response to this problem varies by country. Germany has gone for an infrastructure of fleet management systems that accept digital tachograph data, digitized versions of the analog data from the existing paper charts, fuel data, delivery data and even payroll, and reconcile them all to provide not just management information for the trucking company but surveillance data for the police. Britain has something similar, although it's up to the police to decide which companies to inspect; unless they do so, data on driving infringements is only available to the employer. Germany has also introduced a system of road pricing for heavy goods vehicles that gives further inputs into fleet management systems.

Britain thought initially of road pricing, with tachograph data correlated with GPS location sensors in the trucks, and of using the country's network of automatic number plate reader (ANPR) cameras. The nationwide road-charging plan has become bogged down, as initial plans involved road charging for cars too and drew heavy resistance from motoring organisations. So in the UK ANPR has become the main means of complementary surveillance. It was initially installed around London to make IRA bombing attacks harder, and has now been extended nationwide. It was initially justified on the basis of detecting car tax evaders, but has turned out to be useful in many other policing tasks. We see ANPR data adduced in more and more prosecutions, for everything from terrorism down to burglary. 'Denying criminals the use of the roads' has now become an element in UK police doctrine. In the case of drivers' hours enforcement, the strategy is to verify a sample of logged journeys against the ANPR database; where discrepancies are found, the company's operations are then scrutinised more closely.

However, disagreements about privacy issues and about national economic interests have prevented any EU-wide standardization. It's up to individual countries whether they require truck companies to download and analyze the data from their trucks. And even among countries that require this, fleet management systems aren't a panacea, because of arbitrage. For example, the German police are much more vigorous at enforcing drivers' hours regulations

than their Italian counterparts. So, under the analogue system, an Italian driver who normally doesn't bother to put a chart in his machine will do so while driving over the Alps. Meanwhile, the driver of the German truck going the other way takes his chart out. The net effect is that all drivers in a given country are subject to the same level of law enforcement. But if the driving data get regularly uploaded from the Italian driver's card and kept on a PC at a truck company in Rome then they'll be subject to Italian levels of enforcement (or even less if the Italian police decide they don't care about accidents in Germany). The fix to this was extraterritoriality; an Italian truck driver stopped in Germany can be prosecuted there if he can't show satisfactory records of his driving in Italy for the week before he crossed the border.

12.3.3.2 Other Problems

So the move from analogue to digital isn't always an improvement. As well as the lower tamper-resistance of electronic versus mechanical signalling, and the system level problem that the location of the security state can't be tackled in a uniform way, there are several further interesting problems with tachographs being digital.

A curious problem for the policy folks is that digital tachographs have for the first time caused digital signatures to turn up in court in large numbers. For many years, security researchers have been writing academic papers with punchlines like 'the judge then raises X to the power Y , finds it's equal to Z , and sends Bob to jail'. The reality is very different. Apparently judges find digital signatures too 'difficult' as they're all in hex. The police, always eager to please, have resolved the problem by applying standard procedures for 'securing' digital evidence. When they raid a dodgy trucking company, they image the PC's disk drive and take copies on DVDs that are sealed in evidence bags. One gets given to the defence and one kept for appeal. The paper logs documenting the copying are available for Their Worships to inspect. Everyone's happy, and truckers duly get fined as before.

From the operational viewpoint, the main emerging problem is that many drivers have more than one driver card. This is an offence everywhere but that doesn't stop it! One source of cards is to borrow them from drivers who use them only occasionally — for example because they usually drive analogue trucks, or trucks under 3.5 tonnes. Another is that many drivers have more than one address; the Jean Moulin of Toulouse may also be Jean Moulin of Antwerp. A database, 'Tachonet', was set up to try to catch duplicate applications across European countries but it doesn't seem to work very well. For example, drivers may forget their middle name in one of their countries of residence.

Second, there will be new kinds of service denial attacks (as well as the traditional ones involving gearbox sensors, fuses and so on). A truck driver

can easily destroy his smartcard by feeding it with mains electricity (in fact, even a truck's 24 volts will do fine). Under the regulations he is allowed to drive for 15 days while waiting for a replacement. As static electricity destroys maybe 1% of cards a year anyway, it is hard to prosecute drivers for doing this occasionally. Similar card-destruction attacks have been perpetrated on bank smartcard systems in order to force a merchant back into less robust fallback modes of operation.

Third, I mentioned that the loss of detailed, redundant data on the tachograph chart makes enforcement harder. At present, experienced vehicle inspectors have a 'feel' for when a chart isn't right, but the analogue trace is replaced by a binary signal saying either that the driver infringed the regulations or that he didn't. This spills over into other enforcement tasks; analogue charts were often used to collect evidence of illegal toxic waste dumping, for example, as the recorded speed patterns often give a knowledgeable inspectors a good idea of the truck's route.

Next, some of the cards in the system (notably the workshop cards used to set up the instruments, and the control cards used by police and vehicle inspectors) are very powerful. They can be used to erase evidence of wrongdoing. For example, if you use a workshop card to wind back the clock in a vehicle unit from 10th July to 8th July, then the entries for July 9th and 10th become unreadable. (Of course the vendors should have implemented a proper append-only file system, but they had only 32Kb smartcards to work with not 32Mb memory cards.) Some countries have therefore gone to great lengths to minimise the number of workshop cards that fall into bad hands. In the UK, for example, truck mechanics have to pass a criminal records check to get one; yet this isn't foolproof as it's often companies that get convicted, and the wealthy owners of crooked truck-maintenance firms just set up new firms. There's no company licensing scheme, and although wrongdoers can be blacklisted from acting as directors of licensed firms, crooks just hide behind nominee directors.

Various technical attacks are possible. When assessing the security of the proposed design in the late 1990s, I was concerned that villains might physically reverse-engineer a card, extracting its master key and enabling a powerful workshop or police card to be forged. Since then, tamper-resistance has got better, so attacks are more expensive; and my 1998 critique helped move the design from shared-key to public-key crypto, limiting the damage from a single card compromise. But the most recent attacks on smartcard systems are definitely a concern. Consider for example relay attacks, in which a bogus card, connected to a bogus reader using mobile phones, enables a smartcard in country A to appear to be in country B [401]. In Chapter 10, I discussed this as a means of bank fraud. In the tachograph world, its implications are different. If any police card, or workshop card, can be used to erase evidence of a crime, then what's to stop a corrupt mechanic or policeman in Sicily or in Romania

from using his card to destroy evidence in London or in Berlin? This really is arbitrage with a vengeance. It's no longer enough for British coppers or the German Polizei to be honest, if officials from less well governed countries can indulge in telepresence. Perhaps we'll need region coding for policemen just as we have for DVDs.

This helps illustrate that key management is, as always, difficult. This is a pervasive problem with security systems in vehicles — not just tachographs and taxi meters, but even such simple devices as card door locks and the PIN codes used to protect car radios against theft. If the garage must always be able to override the security mechanisms, and a third of garage mechanics have criminal records, then what sort of protection are you buying? (In my own experience, PIN-protected radios are just a protection racket for the garages — you end up paying a garage £20 to tell you the PIN after you get a flat battery.)

12.3.3.3 *The Resurrecting Duckling*

In the late 1990s, a European Union regulation decreed that, in order to frustrate the use of interruptors of the kind shown in Figure 12.4 above, all digital tachographs had to encrypt the pulse train from the gearbox sensor to the vehicle unit. As both of these devices contain a microcontroller, and the data rate is fairly low, this shouldn't in theory have been a problem. But how on earth could we distribute the keys? If we just set up a hotline that garages could call, it is likely to be abused. There's a long history of fitters conspiring with truck drivers to defeat the system, and of garage staff abusing helplines to get unlocking data for stolen cars and even PIN codes for stolen radios.

One solution is given by the *resurrecting duckling* security policy model. This is named after the fact that a duckling emerging from its egg will recognize as its mother the first moving object it sees that makes a sound: this is called imprinting. Similarly, a 'newborn' vehicle unit, just removed from the shrink wrap, can recognize as its owner the first gearbox sensor that sends it a secret key. The sensor does this on power-up. As soon as this key is received, the vehicle unit is no longer a newborn and will stay faithful to the gearbox sensor for the rest of its 'life'. If the sensor fails and has to be replaced, a workshop card can be used to 'kill' the vehicle unit's key store and resurrect it as a newborn, whereupon it can imprint on the new sensor. Each act of resurrection is indelibly logged in the vehicle unit to make abuse harder. (This at least was the theory — the implementation fell somewhat short in that in one unit the error code for sensor rekeying is the same as the error code for a power outage.)

The resurrecting duckling model of key management was originally developed to deal with the secure imprinting of a digital thermometer or other piece of medical equipment to a doctor's PDA or a bedside monitor. It can

also be used to imprint consumer electronics to a remote control in such a way as to make it more difficult for a thief who steals the device but not the controller to make use of it [1218].

Another possible application is weapon security. Many of the police officers who are shot dead on duty are killed with their own guns, so there has been some interest in safety mechanisms. One approach is to design the gun so it will fire only when within a foot or so of a signet ring the officer wears. The problem is managing the relationship between rings and guns, and a possible solution is to let the gun imprint on any ring, but with a delay of a minute or so. This is not a big deal for the policeman signing a gun out of the armory, but is a problem for the crook who snatches it. (One may assume that if the policeman can't either overpower the crook or run for it within a minute, then he's a goner in any case.) Such mechanisms might also mitigate the effects of battlefield capture of military weapons, for which passwords are often unacceptable [175].

However, one last problem with the idea of a secure sensor has emerged in the last two years, since digital tachographs started shipping. The folks in Italy who brought you the interruptor now have a new product. This is a black box containing electromagnets and electronics to simulate a gearbox. The errant truck driver unscrews his gearbox sensor and places it in this virtual gearbox. The box comes with its own cable and a sensor that he plugs into his actual gearbox. The system now operates as before; on command it will either relay impulses faithfully, or discard them, or filter some of them out. The dodgy pulse-train arrives at the tachograph as before, but this time beautifully encrypted using triple-DES. Secure sensing is harder than it looks!

12.4 Postage Meters

My third case history of metering is the postage meter. Postage stamps were introduced in Britain 1840 by Sir Rowland Hill to simplify charging for post, and developed into a special currency that could be used for certain purposes, from paying for postage to paying certain taxes and topping up the value of postal money orders. Bulk users of the postal system started to find stamps unsatisfactory by the late 19th century, and the postage meter was invented in 1889 by Josef Baumann. Its first commercial use was in Norway in 1903; in the USA Arthur Pitney and Walter Bowes had a meter approved for use in 1920 and built a large business on it. Early postal meters were analogue, and would print a stamp (known as an indicium) on a letter, or on a tape to stick on a parcel. The indicium had a date so that old indicia couldn't be peeled off and reused. Each meter had a mechanical value counter, protected by a physical seal; every so often you'd take your meter into the post office to be read and

reset. Fraud prevention relied on users taking their mail to the local post office, which knew them; the clerk could check the date and the meter serial number.

In 1979, Pitney Bowes introduced a 'reset-by-phone' service, which enabled firms to buy an extra \$500 worth of credit over the phone; the implementation involved a mechanical one-time pad, with the meter containing a tape with successive recharge codes [328]. In 1981, this was upgraded to a DES-based system that enabled a meter to be recharged with any sum of money. The recharge codes were calculated in part from the value counter — so if the firm lied about how much postage they'd used, they couldn't recharge the device. However, these meters still produced inked indicia.

In 1990, José Pastor of Pitney Bowes suggested replacing stamps and indicia with digital marks protected by digital signatures [1007]. This caught the attention of the U.S. Postal Service, which started a program to investigate whether cryptography could help produce better postage meters. One concern was whether the availability of color scanners and copiers would make stamps and indicia too easy to forge. A threat analysis done for them by Doug Tygar, Bennett Yee and Nevin Heintze revealed that the big problem was not so much the forging or copying of stamps, or even tampering with meters to get extra postage. It was bulk mailers corrupting Postal Service employees so as to insert truckloads of junk mail into the system without paying for them [1265]. As a bulk mailer who was fiddling his meter seriously would risk arousing the suspicion of alert staff, there was a temptation to cut them in on the deal; and then it was natural to forge a meter plate whose inducting post office was elsewhere. By 1990 U.S. Postal service losses were in nine figures, and through the 1990s there were a number of high-profile convictions of bulk mailers who had manipulated their meters, and got away with millions of dollars of free postage [190].

This led to a development programme to produce a design based on digital signatures, generated by tamper-resistant processors in the postage meters, that were developed from Pastor's ideas into an open standard available to multiple competing manufacturers. The basic idea is that the indicium, which is machine-readable, contains both the sender and recipient postal codes, the meter number, the date, the postage rate, the amount of postage ever sold by the meter and the amount of credit remaining in it, all protected with a digital signature. The private signature key is kept in the meter's processor while its corresponding public signature verification key is kept in a Postal Service directory, indexed by the meter serial number. In this way, postal inspectors can sample mail in bulk at sorting offices, checking that each item is not only franked but on a logical route from its ostensible source to its destination.

The USA introduced the technology in 2000, with Germany next in 2004 and Canada in 2006; other countries are moving to follow suit. By 2006, the USA had about 450,000 digital meters out of a total market of 1.6 million, and it's

expected that digital devices will have taken over completely by 2012. Also, by 2006, all U.S. postal facilities had the scanners needed to read the new indicia, of which an example is illustrated in Figure 12.5 below.

Such indicia can be produced by postage meters that are drop-in replacements for the old-fashioned devices; you weigh a letter, frank it, and get billed at the end of the month. You don't have to take the meter in to be read though, as that can be done over the Internet for a credit meter, while if you buy a prepayment meter you replenish it by phoning a call center and buying a magic number with your credit card. This works in much the same way as the prepayment electricity meters discussed earlier in this chapter. The tamper-resistance is used to implement what's in effect prepaid digital cash (or preauthorized digital credit) that's kept in the meter on the customer's premises.

Indicia can also be produced without any special equipment locally; you can buy a stamp over the Internet by simply specifying the sender and destination postal codes. This facility, 'online postage', is aimed at small businesses and people working from home who don't send enough mail for it to be worth their while buying a meter. Both metered and online postage are cheaper than stamps to distribute. Also, it becomes possible to manage the system much better, by tracking volumes and profitability of mail down to local level. This matters as many countries' postal systems are deregulated and open up to competition. And despite predictions that email would displace physical mail, post volumes continue to grow by 1% a year, and it turns out that households with broadband actually receive more post.




Figure 12.5: One of the new formats for U.S. postal meters (courtesy of Symbol Technologies)

So, all told, digital postal meters offer more flexibility to both users and postal services than the old analogue ones. But what about security?

One way of looking at postage meters is that they're a slight extension of the utility metering model. There is a tamper-resistant processor, either in the meter itself, or attached to a web server in the case of online postage; this has a value counter and a crypto key. It dispenses value, by creating indicia in the requested denominations, until the value counter is exhausted. It then requires replenishment via a cryptographically-protected message from a control unit higher up in the chain. There are some additional features in each case. Many postage meters include a 'Clark-Wilson' feature whereby the value counter actually consist of two counters, an Ascending Register (AR) containing the total value ever dispensed by the meter, and a Descending Register (DR) indicating the remaining credit. The balancing control is $AR + DR = TS$, the 'total setting', that is, the total of all the sales made by or authorised for that device. If the balance fails, the meter locks up and can only be accessed by inspectors.

An alternative way of looking at postage meters is that they constitute a distributed mint. In Sir Rowland Hill's original British postage system, the penny black stamps were printed at the mint; and postage engineers to this day refer to the step of generating an indicium as 'minting', whether it's done in a tamper-resistant processor in a meter or in an online server. These systems can provide a lot of useful experience for anyone wanting to design novel payment and e-money protocols.

The full threat model includes stolen postage meters, meters that have been tampered with to provide free postage, genuine meters used by unauthorised people, mail pieces with indicia of insufficient value to cover the weight and service class, and straightforward copies of valid indicia. Various sampling and other tests are used to control these risks. Subtleties include how you deal with features like certified mail and reply mail. There are also national differences on matters ranging from which authentication algorithms are used to what sort of usage data the meters have to upload back to the postal service.

One interesting development is that, as operators get real experience of digital postal metering, the industry is moving away from the initial design of using digital signatures to one of using message authentication codes. Signatures appealed because they were elegant; but in real life, signature verification is expensive, and has also turned out to be unnecessary. Equipment at major sorting offices must process thousands of mail pieces a minute, and even using low-exponent RSA, this entails a lot of computation. One argument for signatures was that indicia could be verified even when central servers are offline; but in real operations, postal services usually verify indicia as an offline batch operation. This means that forged mail pieces go through initially and are only intercepted once a pattern of abuse emerges. Once the verification is done offline, this can just as easily be MAC verification as signature verification. (The central servers have hardware security modules

with master keys that were diversified to a MAC key in each meter.) It turns out that only two digits of the MAC are needed, as this is enough to detect any systematic abuse before it becomes significant [328].

The most recent optimisation is for the postal service not to do any cryptography at all, but to contract it out to the meter vendors. This in turn means that indicia are verified only in the home postal system, as overseas systems will often use different vendors. (So if you want to bribe a postal employee to let a few tons of junk mail into the system, the place to do it is now at a border crossing.) The upshot of the move away from public-key cryptography is a diversity of architectures, and sometimes multiple architectures in one country. Canada, for example, uses both signatures and MACs on its indicia.

How stuff actually breaks in real life is — as always — instructive. In the German post office's 'Stampit' scheme, a user buys 'smart pdf' files that contact the post office to say they're being printed, without any interaction with the user or her software. If the paper jams, or the printer is out of toner, then tough. So users arrange to photocopy the stamp, or to copy it to a file from which it can be printed again if need be. The UK system has learnt from this: although a stamp is grey-listed when a user PC phones home and says it's been printed, the grey doesn't turn to black until the stamp appears at the sorting office. The difference in syntax is subtle: the German system tried to stop you printing the stamp more than once, while the British system more realistically tries to stop you using it more than once [592].

All told, moving to digital postal meters involves a nontrivial investment, but enables much better control than was possible in the old days, when postal inspectors had to refer to paper records of mechanical meter readings. The hardware tamper-resistance also facilitates prepayment business models that extend the service's scope to many more customers and that also improve a service's cash flow and credit control. Unlike the case of digital tachographs, digital postal meters appear to be a success story.

12.5 Summary

Many security systems are concerned one way or another with monitoring or metering some aspect of the environment. They range from utility meters to taxi meters, tachographs, and postal meters. We'll come across further metering and payment systems in later chapters, ranging from the mechanisms used to stop printer cartridges working once they have printed a certain number of pages, to prepay scratch cards for mobile phone use, which may be the world's largest application-specific payment mechanism.

Many monitoring, metering and payment systems are being redesigned as the world moves from analogue to digital technology. Some of the redesigns

are a success, and others aren't. The new digital prepayment electricity meters have been a success, and are being introduced around the developing world as an enabling technology that lets utility companies sell power to people who don't even have addresses, let alone credit ratings. Digital tachographs have been much less impressive; they just do what the old analogue systems did, but less well. Our third example, postage meters, appear to be a success.

As with burglar alarms, the protection of these systems is tied up with dependability; if you're designing such a thing, you have to think long and hard about what sort of service denial attacks are possible on system components. Key management can be an issue, especially in low cost widely distributed systems where a central key management facility can't be justified or an adequate base of trustworthy personnel doesn't exist. Systems may have to deal with numerous mutually suspicious parties, and must often be implemented on the cheapest possible microcontrollers. Many of them are routinely in the hands of the opponent. And there are all sorts of application-level subtleties that had better be understood if you want your design to succeed.

Research Problems

We're gradually acquiring a set of worked examples of secure embedded metering, thanks to the kinds of systems described here. We don't yet have a really general set of tools for building them. At the component level, we have crypto algorithms as seen in Chapter 5, protocols as described in Chapter 3, security policies like Clark-Wilson which I described in Chapter 10, and tamper resistant processors, which I'll describe later. However we don't have many concepts, abstractions, or middle that help us pull these components together into larger building blocks. Although the mechanisms (and products) developed for automatic teller machine networks can be adapted (and are), much of the design work has to be redone and the end result often has vulnerabilities. Top level standards for ways in which crypto and other mechanisms can be built into a range of monitoring and ticketing systems might save engineers a lot of effort. Meanwhile we have to rely on case histories like those presented here. Metering applications are particularly useful because of the pervasive mutual mistrust caused not just by competing commercial entities but by the presence of dishonest staff at every level, as well as dishonest customers; and the fact that most of the equipment is in the custody of the attackers.

Again, there are questions for the security economist, and the business school researchers. Why did some digitisations of existing metering systems work well (utilities, postage) while others were much less impressive (tachographs)? Why were some disruptive, in that new entrants successfully challenged the previous incumbent suppliers, while in other cases (such as postage) the

existing suppliers managed the transition to better digital systems and largely saw off competition from dotcom startups?

Further Reading

Prepayment electricity meters are described in [59]. Tachographs are written up in [45]; other papers relevant to transport appear in the annual ESCAR conference on electronic security in cars. The early work on postal meters is in [1265] and the U.S. regulations can be found in [894]. However by far the most detailed exposition of postage meter security is in a book written by Gerrit Bleumer, a scientist at Francotyp-Postalia that took a leading role in the program [190].