

Monitoring Systems

For if a man watch too long, it is odds he will fall asleep.

—FRANCIS BACON

10.1 Introduction

A significant number of secure systems are concerned with monitoring the environment. The most obvious example is the burglar alarm. Then there are meters for measuring consumption of utilities such as gas and electricity. At the top end of the scale, there are systems used to verify nuclear nonproliferation treaties, where a number of sensors (seismometers, closed-circuit TV, and so on) are emplaced in a state's nuclear facilities by the International Atomic Energy Authority (IAEA) to create an immediate, indelible, and remote log of all movements of fissile substances. There are also vehicle systems, such as missile telemetry, taximeters, and tachographs (devices used in Europe to record the speed and working hours of truck and bus drivers).

These have a number of interesting features in common. For example, to defeat a burglar alarm it is sufficient to make it stop working, or—in many cases—to persuade its operators that it has become unreliable. This raises the spectre of *denial of service attacks*, which are increasingly important yet often difficult to deal with.

Just as we have seen military messaging systems designed to enforce confidentiality, and bookkeeping systems whose goal is to preserve record authenticity, monitoring applications give us the classic example of systems designed to be dependably available. If there is a burglar in my bank vault, then I do not care very much who else finds out (so I'm not worried about confidentiality) or who it was who told me (so authenticity isn't a major concern); but I do care very much that an attempt to tell me is not thwarted.

An alarm in a bank vault is very well protected from tampering (at least by outsiders), so it provides the simplest case study. We are largely concerned with attacks on communications (though sensor defeats are also a worry). But many other monitoring systems are very exposed physically. Utility meters are usually on the premises of the

consumer, who has a motive to cause them to make incorrect readings. Much the same goes with taximeters: the taxi driver (or owner) may want the meter to read more miles or more minutes than were actually worked. With tachographs, it's the reverse. The truck driver usually wants to drive above the speed limit, or work dangerously long hours, so both types of attack are found. The driver can either cause the tachograph to fail, or to make false readings of time and distance. These devices, too, are very exposed to tampering. In both metering and monitoring systems (and especially with nuclear verification) we are also concerned with evidence. An opponent could get an advantage not just by manipulating communications (such as by replaying old messages) but by falsely claiming that someone else had done so.

Monitoring systems are also important because they 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. They also provide a gentle introduction to the wider problem of service denial attacks, which dominate the business of electronic warfare, and are starting to be of grave concern to electronic commerce.

10.2 Alarms

Alarms are used to deal with much more than burglary. Their applications range from monitoring freezer temperatures in supermarkets (to stop staff “accidentally” switching off freezer cabinets in the hope of being given spoiling food to take home) right through to improvised explosive devices that are booby-trapped to deter the bomb disposal squad. However, it's convenient to discuss them in the context of burglary and of protecting rooms where computer equipment is kept.

Standards and requirements for alarms vary between countries and between different types of risk. Normally, you will use a local specialist firm for this kind of work; but as a security engineer, you must be aware of the issues. Alarms often affect larger system designs in my own professional practice, this has ranged from the alarms built into automatic teller machines through the evaluation of the security of the communications used by an alarm system for large risks such as wholesale jewelers, to continually staffed systems such as those used to protect bank computer rooms.

It's easier to teach someone with an electrical engineering/computer science background the basics of physical security than the other way round. Therefore, interactions between physical and logical protection will be up to the systems person to deal with. You are also likely to be asked for your opinion on your client's installations—which will often have been designed and installed by local contractors who may have established links with your clients but rather narrow horizons as far as system issues are concerned.

10.2.1 Threat Model

An important design consideration is the level of skill, equipment, and determination that the attacker might have. Movies such as *Entrapment* might be good entertainment, but they don't give a realistic view of the world of theft. In the absence of an “international standard burglar,” the nearest I know to a working classification is one developed by a U.S. Army expert [74].

Derek is a 19-year-old addict. He's looking for a low-risk opportunity to steal something such as a video recorder to fund his next fix.

Charlie is a 40-year-old inadequate with seven convictions for burglary. He has spent seventeen of the last twenty-five years in prison. Although not very intelligent, he is cunning and experienced; he has picked up a lot of "lore" during his spells "inside." He steals from small shops and prosperous-looking suburban houses, taking whatever he thinks he can sell to local fences.

Bruno is a "gentleman criminal." His business is mostly stealing art. As a cover, he runs a small art gallery. He has a (forged) university degree in art history on the wall, and one conviction for robbery eighteen years ago. After two years in jail, he changed his name and moved to a different part of the country. He has done occasional "black bag" jobs for intelligence agencies, who know his past. He'd like to get into computer crime, but the most he has done so far is to strip \$100,000 worth of memory chips from a university's PCs back in the mid-1990s, when there was a memory famine.

Abdurrahman heads a cell of a dozen militants, most with military training. They have infantry weapons and explosives, with PhD-grade technical support provided by a disreputable country. Abdurrahman was third in a class of 280 at the military academy of that country, but was not promoted because he's from the wrong ethnic group. He thinks of himself as a good man rather than a bad man. His mission is to steal plutonium.

So Derek is unskilled; Charlie is skilled; Bruno is highly skilled, and may have the help of an unskilled insider such as a cleaner; while Abdurrahman is not only highly skilled but has substantial resources. He may even have the help of a technician or other skilled insider who has been suborned.

The sociologists are interested in Derek, the criminologists in Charlie, and the military in Abdurrahman; our concern is mainly with Bruno. He isn't the highest available grade of "civilian criminal" (that distinction probably goes to the bent bankers and lawyers who launder money for drug gangs, whom I'll get to later). But in countries without a terrorism problem, the physical defenses of computer rooms tend to be designed with someone like Bruno in mind. (Whether this is rational, or an overplay, will depend on the kind of business your client is in.)

The common view of Bruno is that he organizes cunning attacks on alarm systems, having spent days poring over the building plans in the local town hall. You probably read about this kind of crime several times a year in the papers.

HOW TO STEAL A PAINTING (1)

A Picasso is stolen from a gallery, with supposedly state-of-the-art alarm systems, by a thief who removed a dozen roofing tiles and lowered himself down a rope so as not to activate the pressure mats under the carpet. He grabbed the painting, climbed back out without touching the floor, and probably sold the thing for a quarter of a million dollars to a wealthy cocaine dealer.

The press loves this kind of stuff, and it does happen from time to time. Reality is both simpler and stranger.

10.2.2 How Not to Protect a Painting

A common mistake when designing alarm systems is to be captivated by the latest sensor technology. There's a lot of impressive stuff on the market, such as a fiber optic cable that you can loop around protected objects and that will sense if the cable is stretched or relaxed by less than a thousandth of a millimeter. The naive art gallery owner will buy a few feet of this magic cable, glue it to the back of the Picasso and connect it to an alarm company.

HOW TO STEAL A PAINTING (2)

Bruno's attack is to visit as a tourist and hide in a broom cupboard. At one in the morning, he emerges, snatches the painting and heads for the fire exit. Off goes the alarm, but so what in less than a minute, Bruno is on his motorcycle. By the time the cops arrive 12 minutes later, he has vanished.

This sort of theft is much more likely than a bosun's chair through the roof. It's often easy because alarms are rarely integrated well with building entry controls. Many designers don't realize that where they can't positively account for all the people who enter the premises during the day, it may be prudent to take some precautions against the "stay-behind" villain—even if this means only an inspection tour after the gallery has closed. Serious physical security means serious controls on people. In fact, the first recorded use of the RSA cryptosystem—in 1978—was not to encrypt communications but to provide digital signatures on credentials used by staff to get past the entry barrier to a plutonium reactor at Idaho Falls. The credentials contained data such as body weight and hand geometry [701, 705]. But I continue to be amazed by the ease with which building entry controls are defeated at most secure sites I visit—whether by mildly technical means, such as sitting on somebody else's shoulders to go through an entry booth, or by helpful people holding the door open.

Moreover, the alarm response process often hasn't been thought through carefully. (The *Titanic effect* of over-reliance on the latest gee-whiz technology often blinds people to common sense.) As we'll see shortly, this leads to still simpler attacks on most systems.

So we mustn't think of the alarm mechanism in isolation. A physical security system has a number of elements:

Deter – detect – alarm – delay – respond

The emphasis will vary from one application to another. If our opponent is Derek or Charlie, we will be concerned primarily with deterrence. At the sort of targets Abdurrahman is interested in, an attack will almost certainly be detected; the main problem is to delay him long enough for the Marines to arrive. Bruno is the most interesting case as we won't have the military budget to spend on keeping him out, and there are many more premises whose defenders worry about Bruno than about Abdurrahman. Depending on the circumstances, they might have a problem with detection, and also with the response.

10.2.3 Sensor Defeats

Burglar alarms use a wide range of *sensors*, including

- Vibration detectors, to sense fence disturbance, footsteps, breaking glass, or other attacks on buildings or perimeters
- Switches on doors and windows
- Passive infrared devices to detect body heat
- Motion detectors that use ultrasonics or microwave
- Invisible barriers of microwave or infrared beams
- Pressure pads under the carpet, which in extreme cases may extend to instrumenting the entire floor with pressure transducers under each tile
- Video cameras, perhaps with movement detectors, to alarm automatically or to provide a live video feed to a monitoring center
- Movement sensors on equipment, ranging from simple tie-down cables through seismometers to loops of optical fiber.

Most of these sensors can be circumvented one way or another. Fence-disturbance sensors can be defeated by vaulting the fence; motion sensors by moving very slowly; door and window switches by breaking through a wall. Designing a good combination of sensors comes down to skill and experience (with the latter not always guaranteeing the former).

The main problem is limiting the number of false alarms. Ultrasonics don't perform well near moving air such as central heating inlets, while vibration detectors can be rendered useless by traffic. Severe weather, such as lightning, will trigger most systems, and a hurricane can increase the number of calls per day on a town's police force from dozens to thousands. In some places, even normal weather can make protection difficult. Protecting a site where the intruder might be able to ski over your sensors (and even over your fence) is an interesting challenge for the security engineer. (For an instructive worked example of the design of intruder detection systems for a nuclear power station in a snow zone see [74]).

But regardless of whether you're in Alaska or Arizona, the principal dilemma is that the closer you get to the object being protected, the more tightly you can control the environment, and so the lower the achievable false alarm rate. Conversely, at the perimeter, it's hard to keep the false alarm rate down. But to delay an intruder long enough for the guards to get there, the outer perimeter is exactly where you need reliable sensors.

HOW TO STEAL A PAINTING (3)

Bruno's next attack is to wait for a dark and stormy night. He sets off the alarm somehow, taking care not to get caught on CCTV or otherwise leave any hard evidence that the alarm was a real one. He retreats few hundred yards and hides in the bushes. The guards come out and find nothing. He waits half an hour and sets off the alarm again. This time the guards don't bother, and in he goes.

False alarms—whether induced deliberately or not—are the bane of the industry. They provide a direct denial-of-service attack on the alarm response force. Experience from the world of electronic warfare is that a false alarm rate of greater than about 15% degrades the performance of radar operators; and most intruder alarm response forces are operating well above this threshold. Deliberately induced false alarms are especially effective against sites that don't have round-the-clock guards. Many police forces have a policy that, after a certain number of false alarms from a given site (typically two to five in a year), they will no longer send a squad car there until the alarm company, or another keyholder, has been there to check.

In addition to service denial issues, false alarms degrade systems in other ways. The rate at which they are caused by environmental stimuli, such as weather conditions and traffic noise, limits the sensitivity of the sensors that can usefully be deployed. Also, the very success of the alarm industry has greatly increased the total number of alarms, and thus decreased police tolerance of false alarms. So many people install remote video surveillance, so the customer's premises can be inspected by the alarm company's dispatcher. Many police forces prioritize alarms confirmed by such means [417].

But even online video links are not a panacea. The attacker can disable the lighting, start a fire, or set off alarms in other buildings nearby. The failure of a telephone exchange, as a result of a flood or hurricane, may well lead to opportunistic looting.

After environmental constraints such as traffic and weather, Bruno's next ally is time. Vegetation grows into the path of sensor beams; fences become slack, and the vibration sensors don't work so well; and the criminal community learns new tricks. Meanwhile, the sentries become complacent.

For this reason, sites with a serious physical protection requirement typically have several concentric perimeters. The outer fence keeps out drunks, wildlife, and other low-grade intruders; next there may be level grass with buried sensors, then an inner fence with an infrared barrier, and finally a building of sufficiently massive construction to delay the bad guys until the cavalry gets there. (The international regulations laid down by the IAEA for sites that hold more than 15g of plutonium are an instructive read [409].)

At most sites, this kind of protection isn't possible; it is too expensive. And even if you have loads of money, you may be in a city such as Hong Kong where real estate's in really short supply, and like it or not, your bank computer room will just be a floor of an office building that you'll have to protect as best you can.

In any case, the combination of sensors and physical barriers you select and install are still less than half the story.

10.2.4 Feature Interactions

Intruder alarms and barriers interact in a number of ways with other services. The most obvious of these is electricity. A power cut will leave many sites dark and unprotected, so a serious alarm installation needs batteries or other backup power supplies. A less obvious interaction is with fire alarms and firefighting.

HOW TO STEAL A PAINTING (4)

Bruno again visits the gallery as a tourist and leaves a smoke grenade on a timer. This goes off at one in the morning and sets off the fire alarm, which in turn causes the burglar alarm to ignore signals from its passive infrared sensors. (If it doesn't, the alarm dispatcher will probably ignore them anyway as he concentrates on getting the fire trucks to the scene). Bruno smashes his way in through a fire exit and grabs the Picasso. He'll probably manage to escape in the general chaos, but if he doesn't, he has a cunning plan: to claim he was a public-spirited bystander who saw the fire and risked his life to save the town's priceless cultural heritage. The police might not believe him, but they'll have a hard time prosecuting him.

The interaction between fire and intrusion works in a number of ways. Some fire precautions can be used only if there are effective barriers to keep out innocent intruders. Many computer rooms have automatic fire extinguishers, and since fears over global warming made Halon unavailable, this means carbon dioxide flooding. A CO₂ dump is lethal to untrained personnel. Getting out of a room on the air you have in your lungs is much harder than it looks when visibility drops to a few inches and you are disoriented by the terrible shrieking noise of the dump. A malfunctioning intruder alarm that let a drunk into your computer room, where he lit up a cigarette and was promptly executed by your fire extinguisher, might raise a few chuckles among the anti-smoking militants but is unlikely to make your lawyers very happy.

In any case, the most severe feature interactions are between alarm and communication systems.

10.2.5 Attacks on Communications

A sophisticated attacker is at least as likely to attack the communications as the sensors. Sometimes, this will mean the cabling between the sensors and the alarm controller.

HOW TO STEAL A PAINTING (5)

Bruno goes into an art gallery and, while the staff are distracted, he cuts the wire from a window switch. He goes back that evening and helps himself.

It's also quite possible that one of your staff, or a cleaner, will be bribed, seduced, or otherwise coerced into creating a vulnerability (especially if you're dealing with Abdurrahman rather than Bruno). So frequent operational testing is a good idea, along with sensor overlap, means to detect equipment substitution (such as seals), strict configuration management, and tamper-resistant cabling. (Serious sites insist that alarm maintenance and testing be done by two people rather than one.)

The old-fashioned way of protecting the communications between the alarm sensors and the controller was physical: lay multiple wires to each sensor and bury them in concrete, or use armored gas-pressurized cables. The more modern way is to encrypt

the communications. An example is Argus, a system originally developed for nuclear labs, which uses DES encryption to protect sensor links [303].

But the more typical attack on communications is to go for the link between the alarm controller and the security company that provides or organizes the response force.

HOW TO STEAL A PAINTING (6)

Bruno calls his rival gallery claiming to be from the security company that handles its alarms. He says that he's updating his computers and he needs the serial number on their alarm controller unit. An office junior helpfully gives it to him, not realizing that the serial number on the box is also the cryptographic key that secures the communications. Bruno buys an identical controller for \$200 and, after spending half an hour learning how to use an EEPROM programmer, he has a functionally identical unit, which he splices into his rival's phone line. This continues to report "all's well" even when it isn't.

Substituting bogus alarm equipment, or using a computer that mimics it, is known as *spoofing*. There have been many reports of 'black boxes' that spoof the older or less well-designed alarm controllers. In one case, thieves made off with \$1.5 million in jade statuary and gold jewelry imported from China, a theft which drove the importer into bankruptcy. The alarm system protecting its warehouse in Hackensack, New Jersey, was cut off. Normally, that would have triggered an alarm at a security company, but the burglars had attached a homemade electronic device to an external cable to ensure a continuous "all's well" signal [371].

With modern systems, either the alarm controller in the vault sends a cryptographic pseudorandom sequence to the alarm company, which will assume the worst if it's interrupted, or the alarm company sends the controller periodic random challenges, which are encrypted and returned, just as with IFF.

However, the design is often faulty, having been done by engineers with no training in security protocols. The cryptographic algorithm may be primitive, or its key may be too short (whether because of incompetence or export regulations). It may well be possible for Bruno to record the pseudorandom sequence and replay it slightly more slowly, so that by early Monday morning he might have accumulated five minutes of "slack" to cover a lightning raid.

An even more frequent cause of failure is the gross design blunder. One typical example is having a dial-up modem port that allows remote maintenance, with a default password that many users never change; another is making the crypto key equal to the device serial number. Besides being vulnerable to social engineering, the serial number often appears in the purchase order, invoice, and other paperwork, which lots of people get to see. (In general, it's a good idea to buy your alarm controller for cash. This also makes it less likely that you'll get one that's been "spiked." But big firms often have difficulty doing this.)

By now you've probably decided not to go into the art gallery business. But I've saved the best for last. Here is the most powerful attack on burglar alarm systems. It's a variant on 3, but rather than targeting the sensors, it goes for the communications.

HOW TO STEAL A PAINTING (7)

Bruno cuts the telephone line to his rival's gallery, then hides a few hundred yards away in the bushes. He counts the number of blue uniforms that arrive, and the number that depart. If the two numbers are equal, then it's a fair guess the custodian has said, 'Oh bother, we'll fix it in the morning,' or words to that effect. Bruno now knows he has several hours to work.

This is more or less the standard way to attack a bank vault, and it has also been used on computer installations. The modus operandi can vary, from simply reversing a truck into the phone company's curbside junction box to more sophisticated attempts to cause multiple simultaneous alarms in different premises and thus swamp the local police force. (This is why it's so much more powerful than just rattling the fence.)

In one case, thieves in New Jersey cut three main telephone cables, knocking out phones and alarm apparatus in three police stations and thousands of homes and businesses in the Hackensack Meadowlands. They used this opportunity to steal Lucien Piccard wristwatches from the American distributor, with a value of \$2.1 million wholesale and perhaps \$8 million retail [371]. In another, an Oklahoma deputy sheriff cut the phone lines to 50,000 homes in Tulsa before burgling a narcotics warehouse [762]. In a third, a villain set off a bomb at the telephone exchange in Holborn, London, interrupting service to dozens of shops in the Hatton Garden jewelry quarter. Blanket service denial attacks of this kind, which saturate the response force's capacity, are the burglarious equivalent of a nuclear strike.

In the future as computers and communications converge these attacks might not involve explosives but a software-based distributed denial-of-service attack on network facilities. Rather than causing all the alarms to go off in a neighborhood (which could be protected to some extent by swamping it with police), it might be possible to set off several thousand random alarms all over New York, creating an effect similar to that of a hurricane or a power outage, but at a time convenient for the crooks.

An angle that seriously concerns insurers is that phone company staff might be bribed to create false alarms. Insurance companies would prefer it if alarm communications consisted of anonymous packets, which most of the phone company's staff could not relate to any particular alarm. This would make targeted service denial attacks harder. But phone companies—which carry most of the alarm signal traffic—prefer to concentrate it in exchanges, which makes targeted service denial attacks easier. (These tensions are discussed in [586].)

For these reasons, the rule in the London insurance market (which does most of the world's major reinsurance business) is that alarm controllers in places insured for over £20 million must have two independent means of communication. One option is a leased line and a packet radio service. Another is a radio system with two antennas, each of which will send an alarm if the other is tampered with.¹ In the nuclear world, IAEA regulations stipulate that sites containing more than 500 g of plutonium or 2 Kg

¹ I used to wonder, back in the days when I was a banker, whether two bad men who practiced a bit could cut both cables simultaneously. I concluded that the threat wasn't worth worrying about for bank branches with a mere \$100,000 or so in the vault. Our large cash processing centers were staffed 24/7, so the threat model there focused on dishonest insiders, hostage taking, and so on.

of U-235 must have their alarm control center and response force on the premises [409].

Finally, although physical security isn't a main topic of this book, it's worth noting that many physical security incidents arise from angry people coming into the workplace—whether spouses, former employees, or customers. Alarm systems should be able to cope with incidents that occur during the day as well as at night.

10.2.6 Lessons Learned

You might be wondering why a book that's essentially about security in computer systems should spend several pages describing burglar alarm systems. There are many reasons.

Dealing with service denial attacks is the hardest part of many secure system designs. And, as the bad guys come to understand system-level vulnerabilities, it's also often the most important. Intruder alarms give us one of the largest available bodies of applicable knowledge and experience.

The lesson that one must look at the overall system—from intrusion through detection, alarm, delay and response—is widely applicable, yet increasingly hard to follow in general-purpose distributed systems.

The observation that the outermost perimeter defenses are the ones that we'd most like to rely on, but also those on which the least reliance can be placed, is also quite general.

The trade-off between the missed alarm rate and the false alarm rate is a pervasive problem in security engineering.

There are some subtleties though where we can learn from the alarm business. For example, some U.S. airport X-ray machines use *false alarm insertion* to ensure that alarm systems and personnel stay effective: they insert an image of a gun or bomb about once per shift. Staff are graded continually on their error rates.

Failure to understand the threat model—designing for Charlie and hoping to keep out Bruno—causes many real-life failures. It's necessary to know what actually goes wrong, not just what crime writers think goes wrong.

And, finally, we can't just leave the technical aspects of a security engineering project to specialist subcontractors, as critical stuff will always fall between the cracks.

In addition to these system-level lessons, there are a number of other applications where the experience of the burglar alarm industry is relevant. I already mentioned improvised explosive devices; in a later chapter, I'll discuss tamper-resistant processors that are designed to detect attempts to dismantle them and destroy all their cryptographic key material by way of an alarm response.

10.3 Prepayment Meters

Our next case study comes from prepayment metering. In many systems, 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 postal franking machines, the stored value cards that operate photocopiers in libraries, lift passes at ski resorts, and washing machine tokens in university residence halls. Many transport tickets are similar—especially if the terminals that 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; in this case, all the replay and forgery detection must be done offline, on a terminal that is often vulnerable to physical attack. So if we simply enciphered all our tokens using a universal master key, we might expect that a villain would extract this key from a stolen terminal, then set up as a token vendor in competition with us.

There are also attacks on the server end of things. One neat attack on a vending card system used in the staff cafeteria 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 10.1) [479]. This attack is interesting because no amount of encryption of the card contents would make any difference. Although it could, in theory, be stopped by keeping logs at both ends, the design would not be trivial.

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; while barcoded ski lift passes, parking lot tickets, and the like can be forged with a scanner and printer. The goal is to prevent fraud becoming systematic. 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 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 more than 2.5 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 [39].) Most of the lessons learned apply directly to other ticketing systems.




Figure 10.1 Superposing two payment cards.

10.3.1 Utility Metering

In a number of European countries, householders who can't get credit (because they are on welfare, have court judgments against them, or whatever) buy gas and electricity services using prepayment meters. In the old days, they were coin-operated, but the costs of coin collection led vendors to develop token-based meters instead. The customer goes to a shop and buys a token, which may be a smartcard, a disposable cardboard ticket with a magnetic strip, or even just a magic number. A magic number is often the most convenient, as no special vending apparatus is required: a ticket can be dispensed at a supermarket checkout, or even over the phone. U.S. readers may be used to replenishing a postal meter by phoning a call center and buying a magic number with a credit card: the magic number replenishes the meter. This is exactly the same kind of system as a prepayment utility meter.

The token should be thought of as 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 read something like, "meter 12345, dispense 50 kWh of electricity!" but some have maintenance functions, too (see Figure 10.2). The idea is that the meter will dispense the purchased amount and then interrupt the supply.

The manufacture of these meters has become big business. Britain has about a million electricity meters using two proprietary schemes, and some six hundred thousand gas meters using smartcards. Prepaid electricity meters have been installed in a number of other countries, including Brazil, India, Namibia, and the Ivory Coast. Growth in the Third World is strong because the customers may not even have addresses, let alone credit ratings. This was the case in South Africa: prepayment metering was the only

way the government could meet its election pledge to electrify millions of homes quickly. In the developed world, the main impetus for metering is reducing administrative costs. Electric utilities find that billing systems can devour 20 percent of retail customer revenue in urban areas, while prepayment systems typically cost under 10 percent.



Figure 10.2 A prepayment electricity meter (courtesy of Schlumberger).

10.3.2 How the System Works

The security requirements for a prepayment meter system seem fairly straightforward. Tokens should not be easy to forge, and genuine tokens should not work in the wrong meter, or in the right meter twice. Tokens should either be tamper-resistant (which is expensive) or unique (which can be done fairly easily using serial numbers and cryptography). 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 typical system 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 2. 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. It's often less straightforward than this. In South Africa, many people commute long distances from townships or homelands to their places of work, so are never at home during business hours and prefer to buy tickets where they work. So they can register at an out-of-area vending station, where there is a security protocol to send their meter key to this vending station from the vending station that "owns" the meter. Sales data then get passed in the opposite direction for balancing and settlement. These mechanisms are very much like those developed for ATM networks.

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 9.

The vending machines themselves maintain a credit balance. They rely on tamper-resistant security processors to keep the vendor from extracting vend keys and foreign meter keys, or interfering with the credit balance. The balance is decremented with each sale, and only credited again when cash is banked with the local operating company. This company in turn has to account to the next level up in the distribution network, and so on. Here we have an example of an accounting system partially 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 something we may see a lot more of; for example, it's the model used by the Mondex electronic purse scheme promoted by Mastercard.

So what can go wrong?

10.3.3 What Goes Wrong

Service denial remains an important issue. As 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 and so are allowed to sell electricity on credit. In some cases, agents just dumped their equipment, then claimed that it got 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 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 (which lay under the token slot) was destroyed. Thus, service denial attacks worked well enough to become popular. (They could become a serious problem if banks fielded offline electronic purse smartcards that don't do full balancing, but rely instead on value counters plus statistical balancing. When a customer complains that a card has stopped working, all the bank can do is either refund the amount the customer claims was on the card, or tell him or her to get lost.)

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—a particular make of meter went to maximum credit. Soon kids were throwing steel chains over the 11 KV 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 brownout testing hadn't been specified. In fact, developed country environmental standards were inadequate 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, rather so much worth of electricity at such-and-such a rate. It turned out that the tariff could be set to a tiny amount by vending staff, so that it would operate almost forever. Another make 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 meter remembered the last token serial number entered, and 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 were the following.

Prepayment may be cheap as 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 other than at local shops, to support commuters, was not anticipated and was costly to implement.

Recycle technology if possible, 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 even sports ticketing may be larger applications, but I don't know of any publicly available studies of their failure modes. In many cases, the end systems—such as the meters or turnstiles—are fairly soft, so our 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. One still does what one economically can to prevent people developing efficient systematic attacks on the end systems that are too hard to detect.

We'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.

10.4 Taximeters, 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 worried whether 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 [170].

The next most familiar may be the taximeter. 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. Starting in 1990, for example, General Motors equipped 6 million vehicles with black boxes to record crash data. This could be a bonanza for trial lawyers; there are also privacy aspects, as the existence of the boxes only became public in 1999 [768]. (I'll discuss these issues in Chapter 21.)

The case study we're going to use here is the tachograph. A driver falling asleep at the wheel is the cause of several times more accidents than drunkenness (20 percent versus 3 percent of accidents in Britain, for example). An accidents involving a truck is 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 United States using weigh stations, countries in Europe use devices called *tachographs*, which record a 24-hour history of the vehicle's speed on a circular waxed paper chart (see Figure 10.3).

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; there are three styli which record, the speed (the outside trace), whether the driver was working or resting (the middle trace), and the distance travelled (the inner trace—each tick being

10 km). With some exceptions, which needn't concern us here, it is an offense to drive a truck in Europe unless you have a tachograph 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). Some tachographs have extra needles to record some environmental variable: examples include the flashing lights of emergency vehicles, the temperature of refrigerated trucks, and whether the doors of armored trucks are open or closed. (It is for such applications that tachographs are most widely used in North America.)




Figure 10.3 A tachograph chart.

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 which is also driven by the tachograph. Tachograph charts are also used to investigate other offenses, such as unlicensed toxic waste dumping, and by fleet operators to detect fuel theft. Clearly, there are plenty reasons why a truck driver might want to fiddle his tachograph.²

² It's a general principle in security engineering that one shouldn't aggregate targets. Thus, 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 also getting away with your spy satellite photographs. Forcing a truck driver to defeat his or her tachograph to circumvent the speed limiter, and vice versa, was a serious design error—but one that's now too entrenched to change easily.

The EU is in the process of moving from paper-based to smartcard-based systems, which makes the issue highly topical. As with any security engineering task, we first need to know what actually goes wrong. Most of what I have to say applies equally well to taximeters 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.

10.4.1 What Goes Wrong

According to a 1998 survey of 1,060 convictions of drivers and operators [31], the offenses were distributed as follows.

10.4.1.1 How Most Tachograph Manipulation Is Done

About 70% of offenses 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 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. 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 the driver 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 hitchhiker 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. (The chart in Figure 10.3 shows several violations of this type. For example, the start point is listed as “B’HAM” which could be Birmingham or Buckingham, and the clock was wound back from 14.30 to 14.00, as can be seen from the overlapping traces.)

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 at a safe distance from the trucks in their yard.)

10.4.1.2 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.

When old-fashioned tachographs used a rotating wire cable—as did the speedometers in cars up until the early 1980s—it was hard to fiddle with. For example, if you jammed the truck’s odometer it was quite likely that you’d shear off the cable. Electronic tachographs have made fiddling much easier. They get their input from a sensor in the gearbox, which sends electrical impulses as the prop shaft rotates. A common attack is to unscrew the sensor about a tenth of an inch. This 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.

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 antilock 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 10.3: 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.

10.4.1.3 Tampering with the Instrument

The third category of fraud is tampering with the tachograph unit itself. This amounts for some 6% of offenses, but is in decline with modern equipment, because tampering with digital communications is so much easier than tampering with a rotating wire cable used to be. The typical offense in this category is miscalibration, usually done in cahoots with the fitter, but sometimes by the driver defeating the seal on the device.

10.4.1.4 High-Tech Attacks

The state of the tampering art is the equipment in Figure 10.4. 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 0, and a fourth causes the device to return to proper operation.

This kind of device amounts 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. The next step in the arms race is the use by the police of electronic warfare techniques to detect and neutralize these “interruptors”—after that, no doubt, the bad guys will start using cryptography to secure the communications from the key fob.

10.4.2 Countermeasures

The countermeasures taken against tachograph manipulation vary by country. In Britain, trucks are stopped at the roadside for random checks by vehicle inspectors; particularly suspect trucks may be shadowed across the country. In the Netherlands, enforcement focuses on inspectors descending on a trucking company and going through their delivery documents, drivers' timesheets, fuel records, and the like. In Italy, data from the toll booths on the freeways are used to prosecute drivers who've averaged more than the speed limit (this is why 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 inspectors can't get at them.




Figure 10.4 A tachograph with an interruptor controlled by the driver using a radio key fob (courtesy of Hampshire Constabulary, England).

10.4.2.1 Tachosmart

So the European Union is taking an initiative to design a unified electronic tachograph system, called Tachosmart, which will replace the existing paper-based charts with smartcards. Each driver will have a “driver card” that will, in effect, be the truck driver's license and contain a record of his driving hours over the last 28 days. Each vehicle will have a vehicle unit with a year's history. Special types of smartcard will be used by mechanics to calibrate devices, and by law enforcement officers to read them out at the roadside.

The most substantial objection to the move to smartcards is that it's not clear how it 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; in the meantime, they can run one truck with an old chart system and the other with the new card system. 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.

10.4.2.2 System Level Problems

The response to this problem varies by country. Germany wants an infrastructure of fleet management systems that will accept digital tachograph data, digitized versions of the analogue 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. The idea, as with some mid-1990s proposals for the regulation of cryptography, is that large companies would be trusted to run their own fleet management systems, while small ones would have to use a licensed bureau

Britain doesn't have as large a share of the existing bureau business as Germany does, so British proposals have included integrating tachograph systems either with GPS location sensors in the trucks or with an existing system of automatic number plate readers. (This was first deployed around London to make IRA bombing attacks harder and has now been extended nationwide to detect car tax evaders.)

However, disagreements about privacy issues and about national economic interests have prevented any EU-wide standardization. It's going to be up to individual countries whether they require truck companies to download and analyze the data from their trucks.

Even if everyone does this, it won't be a panacea, because of arbitrage. At present, the German police are much more vigorous at enforcing drivers' hours regulations than their Italian counterparts. So 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). It's easy to see that this will cause downward pressure on enforcement.

10.4.2.3 Other Problems

The move from analogue devices to digital isn't always an improvement. In addition to 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 other interesting problems with tachographs being digital.

First, the loss of detailed, redundant data on the tachograph chart will make enforcement harder. At present, experienced vehicle inspectors have a "feel" for when a chart isn't right; but once the analogue trace is replaced by a binary signal, which says either that the driver complied with the regulations or that he didn't, they have little else to go on (especially if the truck company's HQ with the supporting paperwork is

in another jurisdiction). The new digital system is less likely to degrade gracefully under attack than its analogue predecessor.

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; and under the regulations, he will be allowed to drive for 15 days while waiting for a replacement. As static electricity destroys maybe 1 percent of cards a year anyway, it would be hard to prosecute drivers for doing this. Similar card-destruction attacks have been perpetrated on bank smart-card systems in France and elsewhere, to force systems back into less robust fallback modes of operation.

Third, some of the cards in the system (notably the workshop and calibration cards used to set up the instruments) are very powerful. They can be used to erase evidence of wrongdoing and to restore a tachograph to a virgin state. A black market in them is likely, and they may become valuable enough for it to be worth someone's while to forge them. As a result of this problem, plus some other technical concerns, the Tachosmart system is being redesigned to use public key cryptography rather than universal master secrets in the cards and vehicle units.

A particularly difficult problem turns out to be key management. This is a general problem with security systems involving vehicles—not just tachographs and similar devices such as taximeters, but even such simple devices as car-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 secure system do you think you can build?

10.4.2.4 The Resurrecting Duckling

A recent EU directive stated that, in order to frustrate the use of interruptors of the kind shown in Figure 10.4, 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 PIN codes for stolen car 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, will 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, there is a procedure whereby the vehicle unit can be ‘killed’ and resurrected 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.

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 con-

sumer 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 [731].

Another possible application is weapons security. Many of the police officers who are shot dead on duty are killed with their own guns, so there is now a lot of interest in safety mechanisms. One approach is to design the gun to fire only when within a foot or so of a signet ring that 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 after a delay of a minute or so. This is not a big deal for police officers signing the gun out of the armory, but is a problem for the crook who snatches it. (One may assume that if a 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 [106].

10.5 Summary

Many security systems are concerned one way or another with monitoring some aspect of the environment. They range from ordinary domestic burglar alarms through utility meters to taximeters, tachographs, and even a number of systems critically concerned with nuclear safety.

The protection of these systems is most often more concerned with preventing attacks that involve denial of service, such as swamping communications, overwhelming sensors with noise, or doing other things which, directly or indirectly, decrease the amount of trust that the system owners place in it. Service denial attacks may be augmented, or complemented, with various kinds of data manipulation. Key management can be an issue, especially in low-cost widely distributed systems where a central key management facility can't be justified or trustworthy field personnel don't exist. Systems may have to deal with numerous mutually suspicious parties, and must often be implemented on the cheapest possible microcontrollers. Finally, many of them are routinely in the hands of the enemy.

I've illustrated the problems of this exacting environment with three case studies—burglar alarms, utility meters, and vehicle tachographs—which may be instructive now that denial of service attacks on the Internet such as SYN floods and DDoS have become a major issue.

Research Problems

We don't yet have a really general set of tools to manage keys in embedded systems. Although the mechanisms (and products) developed for automatic teller machine networks can be (and are) adapted, much of the design work has to be redone; the result often has security vulnerabilities (I'll discuss this in Chapter 14, which deals with the special processors used for this purpose).

Although we have some industry standards (such as CANBUS, which is used for communications between vehicle systems), we don't have any top-level standards for ways in which cryptography and other mechanisms, such as anonymity and balancing,

can be built into a range of monitoring and ticketing systems. Such standards could save a lot of engineers a lot of effort.

Further Reading

The best all-round reference I know of on alarm systems is [74]; the system issues are discussed succinctly in [586]. Resources for specific countries are often available through trade societies, such as the American Society for Industrial Security [14], and through the local insurance industry; many countries have a not-for-profit body such as Underwriters' Laboratories [756] in the United States, and schemes to certify products, installations, or both. Research papers on the latest sensor technologies appear at the IEEE Carnahan conferences [399].

Prepayment electricity meters are described in [39], and a rather similar application—postal metering machines—in [753]. Tachographs, including the Tachosmart project, are written up in [31]. Finally, the systems used to monitor compliance with nuclear arms control treaties are discussed in [702].