

Physical Protection

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

— Francis Bacon

*The greatest of faults, I should say,
is to be conscious of none.*

— Thomas Carlyle

11.1 Introduction

Most security engineers nowadays are largely concerned with electronic systems, but there are several reasons why physical protection cannot be entirely neglected. First, if you're advising on a company's overall risk management strategy, then walls and locks are a factor. Second, as it's easier to teach someone with an electrical engineering/computer science background the basics of physical security than the other way round, interactions between physical and logical protection will be up to the systems person to manage. Third, you will often be asked for your opinion on your client's installations — which will often have been installed by local contractors who are well known to your client but have rather narrow horizons as far as system issues are concerned. You'll need to be able to give informed, but diplomatic, answers. Fourth, many security mechanisms can be defeated if a bad man has physical access to them, whether at the factory, or during shipment, or before installation. Fifth, many locks have recently been completely compromised by 'bumping', an easy covert-entry technique; their manufacturers (even those selling 'high-security' devices) seemed to be unaware of vulnerabilities that enable their products to be quickly bypassed. Finally, your client's hosting

centres will usually be its most hardened facilities, and will be the responsibility of the systems managers who will most often seek your advice.

Much of physical security is just common sense, but there are some non-obvious twists and there have been significant recent advances in technology, notably in lock-picking and other forms of covert entry. There are ideas from criminology and architecture on how you can reduce the incidence of crime around your facilities. And perhaps most importantly, there are burglar alarms — which have a number of interesting system aspects.

For example, in order 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 book-keeping systems whose goal is preserving 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 gets to know (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. Now, historically, about 90% of computer security research was about confidentiality, about 9% about authenticity and 1% about availability. But actual attacks — and companies' expenditures — tend to be the other way round: more is spent on availability than on authenticity and confidentiality combined. And it's alarm systems, above all else, that can teach us about availability.

11.2 Threats and Barriers

Physical protection is no different at heart from computer security: you perform a threat analysis, then design a system that involves equipment and procedures, then test it. The system itself typically has a number of elements:

Deter–detect–alarm–delay–respond

A facility can deter intruders using hard methods such as guards and razor-wire fences, or softer methods such as being inconspicuous. It will then have one or more layers of barriers and sensors whose job is to keep out casual intruders, detect deliberate intruders, and make it difficult for them to defeat your security too quickly. This defense-in-depth will be complemented by an alarm system designed to bring a response to the scene in time. The barriers will have doors in them for authorized staff to go in and out; this means some kind of entry control system that could be anything from metal keys to biometric scanners. Finally, these measures will be supported by operational

controls. How do you cope, for example, with your facility manager having his family taken hostage by villains?

As I noted earlier, one of the ways in which you get your staff to accept dual controls and integrate them into their work culture is that these controls protect them, as well as protecting the assets. Unless the operational aspects of security are embedded in the firm's culture, they won't work well, and this applies to physical security as much as to the computer variety. It's also vital to get unified operational security across the physical, business and information domains: there's little point in spending \$10m to protect a vault containing \$100m of diamonds if a bad man can sneak a false delivery order into your system, and send a DHL van to pick up the diamonds from reception. That is another reason why, as the information security guy, you have to pay attention to the physical side too or you won't get joined-up protection.

11.2.1 Threat Model

An important design consideration is the level of skill, equipment and motivation that the attacker might have. Movies like 'Entrapment' might be good entertainment, but don't give a realistic view of the world of theft. As we have seen in one context after another, 'security' isn't a scalar. It doesn't make sense to ask 'Is device X secure?' without a context: 'secure against whom and in what environment?'

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 [118].

- *Derek* is a 19-year old addict. He's looking for a low-risk opportunity to steal something he can sell for his next fix.
- *Charlie* is a 40-year old inadequate with seven convictions for burglary. He's 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 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's done so far is stripping \$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 himself came third out of 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. (It's true that many terrorists these days aren't even as skilled as Charlie, but it would not be prudent to design a nuclear power station on the assumption that Charlie would be the highest grade of attacker.)

While the sociologists focus on Derek, the criminologists on Charlie and the military on 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. (I'll talk about them in a later chapter.) But the physical defenses of banks and computer rooms tend to be designed with someone like Bruno in mind. (Whether this is rational, or an overplay, will depend on the business your client is in.)

11.2.2 Deterrence

The first consideration is whether you can prevent bad people ever trying to break in. It's a good idea to make your asset anonymous and inconspicuous if you can. It might be a nondescript building in the suburbs; in somewhere like Hong Kong, with astronomical property prices, it might be half a floor of an undistinguished skyscraper.

Location matters; some neighbourhoods have much less crime than others. Part of this has to do with whether other property nearby is protected vigorously, and how easy it is for a crook to tell which properties are protected. If some owners just install visible alarms, they may redistribute crime to their neighbours; but invisible alarms that get criminals caught rather than just sent next door can have strongly positive externalities. For example, Ian Ayres and Steven Levitt studied the effect on auto thefts of Lojack, a radio tag that's embedded invisibly in cars and lets the police find them if they're stolen. In towns where a lot of cars have Lojack, car thieves are caught quickly and 'chop-shops' that break up stolen cars for parts are closed down. Ayres and Levitt found that although a motorist who installs Lojack pays about \$100 a year, the social benefit from his doing this — the reduced car crime suffered by others — is \$1500 [100]. One implication is that good alarm services may be undersupplied by the free market, as many people will free-ride off their neighbours: only rich people, or people with newer cars, or who are

particularly loss-averse, will install alarms. The same principle applies to real estate; an upper-class neighbourhood in which a fair number of houses have high-grade alarms that quietly call the police is a dangerous place for a burglar to work.

However, that is by no means all. Since the 1960s, there has arisen a substantial literature on using environmental design to deflect and deter threats. Much of this evolved in the context of low-income housing, as criminologists and architects learned which designs made crime more or less likely. In 1961, Elizabeth Wood urged architects to improve the visibility of apartment units by residents, and create communal spaces where people would gather and keep apartment entrances in view, thus fostering social surveillance; areas that are out of sight are more vulnerable [1355]. In 1972, Oscar Newman developed this into the concept of 'Defensible Space': buildings should be designed 'to release the latent sense of territoriality and community' of residents [968]. Small courtyards are better than large parks, as intruders are more likely to be identified, and residents are more likely to challenge them. At the same time, Ray Jeffery developed a model that is based on psychology rather than sociology and thus takes account of the wide differences between individual offenders; it is reflected in our four 'model' villains. Intruders are not all the same, and not all rational [1079].

Jeffery's 'Crime Prevention Through Environmental Design' has been influential and challenges a number of old-fashioned ideas about deterrence. Old timers liked bright security lights; but they create glare, and pools of shadow in which villains can lurk. It's better to have a civilised front, with windows overlooking sidewalks and car parks. In the old days, cyclone fences with barbed wire were thought to be a good thing; but they communicate an absence of personal control. A communal area with picnic seating, in which activities happen frequently, has a greater deterrent effect. Trees also help, as they make shared areas feel safer (perhaps a throwback to an ancestral environment where grassland with some trees helped us see predators coming and take refuge from them). Access matters too; defensible spaces should have single egress points, so that potential intruders are afraid of being trapped. It's been found, for example, that CCTV cameras only deter crime in facilities such as car parks where there's a single exit [527]. There are also many tricks developed over the years, from using passing vehicles to enhance site visibility to planting low thorn bushes under windows. Advice on these can be found in the more modern standards such as [229].

Another influential idea is the broken windows theory of George Kelling and Catherine Coles [700]. They noted that if a building has a broken window that's not repaired, then soon vandals will break more, and perhaps squatters or drug dealers will move in; if litter is left on a sidewalk then eventually people will start dumping their trash there. The moral is that problems should be fixed when they're still small. Kelling was hired as a consultant to help

New York clean up its vandalised subways, and inspired the zero-tolerance policing movement of police chief William Bratton, who cracked down on public drinkers, squeegee men and other nuisances. Both petty crime and serious crime in New York fell sharply. Criminologists still arguing about whether the fall was due to zero tolerance, or to other simultaneous changes such as demographics [787] and right-to-carry laws [814].

A related set of ideas can be found in the situational crime prevention theory of Ronald Clarke. This builds on the work of Jeffery and Newman, and is broader than just property crime; it proposes a number of principles for reducing crime generally by increasing the risks and effort, reducing the rewards and provocations, and removing excuses. Its focus is largely on designing crime out of products and out of the routines of everyday life; it's pragmatic and driven by applications rather than drawing on theories of psychology and sociology [298]. It involves detailed study of specific threats; for example, car theft is considered to be a number of different problems, such as joyriding by juveniles, theft to get home at night, theft of parts, and theft by professional gangs of cars for dismantling or sale abroad — and these threats can be countered by quite different measures. Such empirical studies are often criticised by criminologists who have a sociology background as lacking 'theory', but are gaining influence and are not far from what security engineers do. Many of the mechanisms discussed in this book fit easily within a framework of application-level opportunity reduction.

This framework naturally accommodates the extension of environmental controls to other topics when needed. Thus, for example, if you're planning on anonymity of your premises as a defence against targeted attack, you have to think about how you limit the number of people who know that the basement of your Norwich sales office actually contains your main hosting centre. This brings in internal control, culture and even information security policy. Governments often use multilevel policies for this; there may be a rule that the location of all public-sector hosting centres is 'Restricted'. Even in a commercial firm that doesn't burden itself with all the overhead of multilevel security, some of the ideas I discussed in that context in Chapter 8 may be useful.

11.2.3 Walls and Barriers

Anyway, once you've decided what environmental features you'll use to deter Derek or Charlie from trying to break into your site, and how you make it harder for Bruno to find out which of your sites he should break into, you then have the problem of designing the physical barriers.

The first task is figure out what you're really trying to protect. In the old days, banks used to go to great lengths to make life really tough for robbers, but this has its limits: a robber can always threaten to shoot a customer. So

by a generation ago, the philosophy had shifted to 'give him all the cash he can see'. This philosophy has spread to the rest of retail. In 1997, Starbucks reviewed physical security following an incident in which three employees were shot dead in a bungled robbery. They decided to move the safes from the manager's office to the front of the store, and made these safes highly visible not just to staff, customers and passers-by, but also to the control room via CCTV. A side-benefit was improved customer service. The new design was tested at a number of U.S. locations, where increased sales and loss reductions gave a good return on investment [341]. Indeed, I notice that young people increasingly leave their car keys by the front door at home; if someone breaks into your house in order to steal a car, do you really want to engage them in hand-to-hand combat?

Second, having settled your protection goals, you have to decide what security perimeters or boundaries there will be for what purposes, and where they'll be located. A growth industry recently has been the provision of vehicle traps to prevent car bombs being brought close to iconic terrorist targets. However a common failing is to focus on rare but 'exciting' threats at the expense of mundane ones. It's common to find buildings with stout walls but whose roofs are easy to penetrate, for example; perhaps a terrorist would blow himself up at your main gate to no effect, but an environmental protester could cripple your fab and cost you hundreds of millions in lost production by climbing on the roof, cutting a hole and dropping some burning newspaper.

For this reason, organisations such as NIST, the Builders' Hardware Manufacturers' Association, Underwriters' Laboratories, and their equivalents in other countries have a plethora of test results and standards for walls, roofs, safes and so on. The basic idea is to assess how long a barrier will resist an attacker who has certain resources — typically hand tools or power tools. Normal building materials don't offer much delay at all; a man can get through a cavity brick wall in less than a minute using a sledgehammer, and regardless of how good a lock you put on your front door, a police unit raiding your house will typically break the door off its hinges with a battering-ram. So could a robber. Thus for many years the designers of data centres, bank vaults and the like have favoured reinforced concrete walls, floors and roofs, with steel doorframes. Of course, if the bad guys can work undisturbed all weekend, then even eight inches of concrete won't keep them out.

There's a further problem in that the organisations that certify locks, safes and vaults often place unrealistic constraints on the tools available to an attacker. The lock on your car steering wheel is certified to resist a man putting his weight on it; car thieves just use a scaffolding pole, which gives them enough leverage to break it. The typical bank vault is certified to resist attack for ten minutes, yet your local Fire Department can get in there in two minutes using an abrasive wheel. And if the bad guys have access to proper explosives

such as shaped charges, they can get through almost anything in seconds. Another issue is the thermic lance, or burning bar, which will cut through most barrier materials quickly: safe engineers use them to get into a vault whose combination has been lost. Robbers can get them too. So barriers can't be seen in isolation. You have to evaluate them in the context of assumptions about the threats, and about the intrusion detection and response on which you can rely.

11.2.4 Mechanical Locks

The locksmithing industry has been seriously upset in the last couple of years by a couple of developments that have exposed the vulnerability of many low-cost mechanical locks.

The first of these is *bumping*. This technique enables many locks to be opened quickly and without damage by unskilled people using tools that are now readily available. Its main target is the pin-tumbler lock originally patented by Linus Yale in 1860 (see Figure 11.1). This was actually used in ancient Egypt, but Yale rediscovered it and it's often known as a 'Yale lock', although many firms make versions nowadays.

These locks have a cylindrical plug set inside a shell, and prevented from rotating by a number of *pin stacks*. Each stack usually consists of two or three pins, one on top of the other. The *bottom pin* or *key pin* makes direct contact with the key; behind it is a spring-loaded *top pin* or *driver pin* that forces the bottom pin as far down as possible in the keyway. When the correct key is inserted, the gaps between the top pin and the bottom pin align with the edge of the plug, creating a *shear line*; the plug can now be turned. A typical house or office lock might have five or six pins each of which could have the gap in ten different positions, giving a theoretical key diversity of 10^5 or 10^6 possible *key*




Figure 11.1: A cutaway pin-tumbler lock (Courtesy of Marc Weber Tobias)

differs. The actual number will be less because of mechanical tolerances and key-cutting restrictions.

It had been known for years that such locks can be picked, given special tools. You can find details in the MIT Lock Picking Manual [1258] or in treatises such as that by Marc Weber Tobias [1253]: the basic idea is that you twist the plug slightly using a tension wrench, and then manipulate the pins with a lockpick until they all line up along the shear line. Such techniques are used by intelligence agencies, locksmiths and high-grade crooks; but they take a lot of practice, and it's unlawful to possess the tools in many jurisdictions (for the laws in the USA, see [1255]). Until recently, lockpicking was generally thought to be a threat only to high-value targets where covert entry was of particular value to an attacker, such as investment banks and embassies.

The new discovery is that an attacker can insert a specially made *bump key* each of whose teeth is set at the lowest pin position and whose shoulder is slightly rounded. (Such keys are also known as '999' keys as all the teeth are at the lowest position, or *bitting*, namely number 9.) He can then place the key under slight torsion with his fingertips and tap the key head with a rubber mallet. The shock causes the pins to bounce upwards; the applied torsion causes them to stick as the spring pushes them back down, but with the gap at the cylinder edge. The net effect is that with a few taps of the mallet, the lock can be opened.

This trick had been known for years, but recently became much more effective because of better tools and techniques. It was publicised by a 2005 white paper written by Barry Wels and Rop Gonggrijp of The Open Organization Of Lockpickers (TOOOL), a Dutch 'lock sports' group (as the pastime of amateur locksmithing is starting to be known [1337]). TV coverage spread the message to a wide audience. There followed a technical analysis by lock expert Marc Weber Tobias [1254]; in his view, the main threat from bumping is that it deskills lockpicking. The consequences are potentially serious. It's been found, for example, that the locks in U.S. mailboxes can be opened easily, as can the pin-tumbler locks with 70% of the U.S. domestic market. The Dutch paper, and the subsequent publicity, have kicked off an arms race, with vendors producing more complex designs and amateur locksmiths reporting bumping attacks on many of them.

Until recently, locks from Medeco were thought to be unpickable (as well as being certified as such), and the company had a dominant position in the high-security lock market. Medeco uses secondary keying not in the form of a sidebar but in the angle at which cuts are made in the key. In this 'biaxial' system, angled cuts rotate the pins to engage sliders. In 2005, Medeco introduced the m3 which also has a simple sidebar in the form of a slider cut into the side of the key. In 2007, Tobias reported an attack on the m3 and biaxial locks, using a bent paperclip to set the slider and then a combination of bumping and picking to rotate the plug [1256].

What can a householder do? As an experiment, I replaced my own front door lock. The only high-security product I could find in a store within an hour's drive turned out to be a rebranded Mul-T-Lock device from Israel. It took two attempts to install, jamming the first time; it then took about a week for family members to learn to use the more complex deadbolt, which can easily fail open if operated carelessly. And the next time we were visited by someone with an intelligence background, he remarked that in the UK only drug dealers fitted such locks; so if the police ever pass by, I might end up on their database as a suspected pusher. This dubious improvement to my home security cost me \$200 as opposed to under \$20 for a standard product; and as in practice a burglar could always break a window, our actual protection still depends more on our location and our dogs than on any hardware. Indeed, Yochanan Shachmurove and colleagues surveyed the residents of Greenwich, Connecticut, and built a model of how domestic burglaries varied as a function of the precautions taken; locks and deadbolts had essentially no effect, as there were always alternative means of entry such as windows. The most effective deterrents were alarms and visible signs of occupancy such as cars in the drive [1154].

The situation for commercial firms is slightly better (but not much). The usual standards for high-security locks in the USA, UL 437 and ANSI 156.30, specify resistance to picking and drilling, but not to bumping; and although pick-resistant locks are generally more difficult to bump, this is no guarantee. Knowledge does exist about which lock designs resist bumping, but you have to look for it. (Tobias' paper, and www.toool.org, are good starting points.) UL has just recently taken up the issue of bumping and has formed a task force to determine whether this method of attack should be included in their testing of high security locks. BHMA/ANSI are also looking at the issue.

Purchasers therefore face a lemons market — as one might suspect anyway from the glossiness and lack of technical content of many lock vendors' marketing literature. And even expensive pick-resistant locks are often poorly installed by builders or OEMs; when I once had to break into a cryptographic processor with a really expensive lock, I found it could be levered open easily as the lock turned a cam that was made of soft metal. Indeed a recent security alert by Tobias disclosed that one of the most popular high security deadbolts could be mechanically bypassed by sliding a narrow screwdriver down the keyway, catching the bolt at the end and turning it, even without defeating the extensive security protections within the lock. This design had existed for more than twenty years and the vulnerability was unknown to the manufacturer before the disclosure. Many high security installations employ this or similar hardware.

The second recent class of problems are *master key attacks*. These have also been known to locksmiths for some time but have recently been improved and published, in this case by Matt Blaze. Master key systems are designed so that

in addition to the individual key for each door in a building, there can be a top-level master key that opens them all — say, for use by the cleaners. More complex schemes are common; in our building, for example, I can open my students' doors while the system administrators and cleaners can open mine. In pin-tumbler locks, such schemes are implemented by having extra cuts in some of the pin stacks. Thus instead of having a top pin and a bottom pin with a single cut between them, some of the pin stacks will have a middle pin as well.

The master-key attack is to search for the extra cuts one at a time. Suppose my key bitting is 557346, and the master key for my corridor is 232346. I make a key with the bitting 157346, and try it in the lock. It doesn't work. I then file the first position down to 257346. As 2 is a valid bitting for the first pin, this opens the lock, and as it's different from my user bitting of 5, I know it is the master key bitting for that pin. I will have to try on average four bittings for each pin, and if three pins are master-keyed then I will have a master key after about twelve tests. So master keying allows much greater convenience not just to the building occupants but also to the burglar. This is really important, as most large commercial premises use master keying. There are master-keying systems that resist this attack — for example, the Austrian lockmaker Evva has a system involving magnets embedded in metal keys which are much harder to duplicate. But most fielded systems appear vulnerable.

Another thing to worry about is, as always, revocation. Keyholders leave, and may become hostile. They may have made a copy of their key, and sell it to an attacker. Mechanical locks are easy to change singly but locking systems generally cope very poorly with revocation. Master-key attacks are important here, and so is bumping. Indeed, many expensive, pick-resistant locks actually make the problem worse. They often depend on a secondary keying mechanism such as a sidebar: the keys look like two normal pin-tumbler keys welded together, as in Figure 11.2. The sidebar is often the same for all the locks in the building (master-keyed systems generally require common sidebars in locks that share master keys). So if a bad man can get hold of a genuine key belonging to one of my students, he may be able to turn it into a bump key that will open my door, and indeed every door in the building, as in Figure 11.3. This may not be a problem in normal commercial premises, but it definitely is for banks, bullion dealers and wholesale jewelers where



Figure 11.2: Key for a sidebar lock




Figure 11.3: Bump key for a sidebar lock

attackers might spend two years planning a raid. Indeed, if such a facility had a master-keying system using sidebar locks, and a staff member were even suspected of having leaked a key, the prudent course of action would be to replace every single lock.

The combined effect of bumping, bad deadbolts, master-key attacks and other recent discoveries might be summarised as follows. Within a few years — as the tools and knowledge spread — a career criminal like Charlie will be able to open almost any house lock quickly and without leaving any forensic trace, while more professional attackers like Bruno and Abdurrahman will be able to open the locks in most commercial premises too. House locks may not matter all that much, as Charlie will just go through the window anyway; but the vulnerability of most mechanical locks in commercial premises could have much more complex and serious implications. If your responsibilities include the physical protection of computer or other assets, it's time to start thinking about them.

11.2.5 Electronic Locks

The difficulty of revocation is just one reason why electronic locks are starting to gain market share. They have been around for a long time — hotels have been using card locks since the 1970s. There's an enormous diversity of product offerings, using all sorts of mechanisms from contactless smartcards through PIN pads to biometrics. Many of these can be bypassed in various ways, and most of the chapters of this book can be applied in one way or another to their design, evaluation and assurance. There are also some electromechanical locks that combine mechanical and electronic (or magnetic) components; some of these we just don't know how to attack short of physical destruction. But, from the viewpoint of a company using locks to protect sensitive premises, the big problem is not so much the locks themselves but how you hook up dozens or hundreds of locks in a building. Think of a research laboratory some of whose rooms contain valuable inventions that haven't been patented yet, or a law firm where the offices might contain highly sensitive documents on forthcoming takeovers. Here you worry about insiders as well as outsiders.

In the long run, buildings may become aware of who is where, using multiple sensors, and integrate physical with logical access control. Knowing who went through which door in real time enables interesting security policies to be enforced; for example, if classified material is being handled, you can sound an alarm if there's anyone in the room without the right clearance. Buildings can monitor objects as well as people; in an experiment at our lab, both people and devices carried active badges for location tracking [1318]. Electronic systems can be fully, or almost always, online, making revocation

easy. As well as enforcing security policy, smart buildings could provide other benefits, such as saving energy by turning lights off and by tailoring airconditioning to the presence of occupants. But we're not there yet.

One practical problem, as we found when we built our new lab building, is that only a few firms sell turnkey entry control systems. We initially wanted to have biometric entry control based on iris scanners, as they were invented by one of my faculty colleagues, John Daugman. But back in 2000, we couldn't do that. The vendors' protocols didn't support the kit and we didn't have the time and the people to build our own entry control system from scratch. (And if a computer science department can't do that, the average customer has no chance.) We learned that the existing vendors operate just as other systems houses do: they make their money from lockin (in the economic, rather than locksmithing sense). However, the systems you buy can be extraordinarily troublesome, dysfunctional and expensive. You end up paying \$2000 for a door lock that cost maybe \$10 to manufacture, because of proprietary cabling systems and card designs. The main limit to the lockin is the cost of ripping and replacing the whole system — hence the vendors' love of proprietary cabling.

Our lab is now moving to a more open system based on standard contactless smartcards that are available from multiple vendors. The experience has taught us that an entry control system should be managed like any other computer system purchase, with very careful attention to maintenance costs, standards, extensibility, and total cost of ownership. We are keen to get a system we can install and maintain ourselves, and that allows us to specify security policies at a decent level of abstraction. (Our old system just has a matrix specifying which key opens which lock.) We are just starting to see the sort of components that will make decent systems integration possible — such as reasonably-priced door locks that run off the building's standard Ethernet. In short, the locksmithing industry is ripe for competition and modernisation. It's going to go digital, like most other industries.

It reminds me of the long conflict between phone companies and computer companies. The phone companies had their solid, established ways of doing things and assumed they could dictate to the computer industry how data would be sent along their lines. They lost; computer firms were nimbler and more entrepreneurial, and understood the technology better. I expect the same will happen with locks. Within ten years, commercial entry control systems will just be computer systems, albeit with some specialised peripherals. They will be run by your systems administrator rather than by the retired policeman who now controls your site guards. They will finally integrate with environmental controls, personnel systems and alarms, making the smart building practical. The entry control industry will resist for a while, and use all the complex government and insurance certification requirements that have accreted over the years, just as the phone companies used their own regulators

to try to strangle almost every innovation from the early data networks to VOIP. And just as computer firms have had to learn about dependability as they got into online systems, so there are other dependability lessons to be learned when doing physical security. This brings us to the most automated and sophisticated aspect of physical security, namely alarms.

11.3 Alarms

Alarms are used to deal with much more than burglary. Their applications range from monitoring freezer temperatures in supermarkets (so staff don't 'accidentally' switch off freezer cabinets in the hope of being given food to take home), right through to improvised explosive devices in Iraq and elsewhere that are sometimes booby-trapped. However, it's convenient to discuss them in the context of burglary and of protecting rooms where computer equipment and other assets are kept. Alarms also give us a good grounding in the wider problem of service denial attacks, which dominate the business of electronic warfare and are a problem elsewhere too.

Standards and requirements for alarms vary between countries and between different types of risk. You will normally 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 used to protect bank computer rooms.

An alarm in a bank vault is very well protected from tampering (at least by outsiders), so is a rather simple case. In order to look at the problem more generally, I'll consider the task of designing an alarm system for an art gallery. This is more interesting, because attackers can come in during the day as members of the public and get up to mischief. We'll imagine that the attacker is Bruno — the educated professional art thief. 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 removes a dozen roofing tiles and lowers himself down a rope so as not to activate the pressure mats under the carpet. He grabs the painting, climbs back out without touching the floor, and probably sells 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. Let's work through the threat scenarios systematically.

11.3.1 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 which you can loop round protected objects and which will alarm if the cable is stretched or relaxed by less than 100nm — a ten-thousandth of a millimeter. Isn't modern science marvellous? So the naive art gallery owner will buy a few feet of this magic cable, glue it to the back of his prize 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 will be on his motorbike. By the time the cops arrive twelve 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 realise that unless you can positively account for all the people who've entered the premises during the day, it may be prudent to take some precautions against the 'stay-behind' villain — even if this is only an inspection tour after the gallery has closed. So 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 [1170, 1174]. But I'm still 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 even just by helpful people holding the door open.

In addition, 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 below, this leads to still simpler attacks on most systems.

So we mustn't think of the alarm mechanism in isolation. As I mentioned above, a physical protection system has several steps: *deter* — *detect* — *alarm* — *delay* — *respond*, and the emphasis will vary from one application to another. If our opponent is Derek or Charlie, we will mostly be concerned with

deterrence. At the sort of targets Abdurrahman's 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. So you have to look carefully at the circumstances, and decide whether the bigger problem is with detection, with delay or with response.

11.3.2 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 using 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, maybe with movement detectors, to alarm automatically or 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 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). A standard, if slightly dated, reference on sensor installation is [283].

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: 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 intruder detection for a nuclear power station in a snow zone see [118]).

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)

So 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 retires a 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, so 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 responders 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 three 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.

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. A common strategy is to have remote video surveillance as a second line of defense, so the customer's premises can be inspected by the alarm company's dispatcher; and many police forces prioritize alarms confirmed by such means [661]. But even online video links are not a panacea. The attacker can disable the lighting, or start a fire. He can set off alarms in other buildings in the same street. 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 so the vibration sensors don't work so well, the criminal community learns new tricks, and meanwhile the sentries become complacent.

For this reason, sites with a serious physical protection requirement typically have several concentric perimeters. The traditional approach was an outer fence to keep out drunks, wildlife and other low-grade intruders; then 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 regulations laid down by the International Atomic Energy Agency for sites that hold more than 15g of plutonium are an instructive read [640]. A modern hosting centre might follow the same strategy; it may be in a nondescript building whose walls keep out the drunks and the rats, but with a more serious internal walls and sensors protecting the machine room.

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

Anyway, the combination of sensors and physical barriers which you select and install are still less than half the story.

11.3.3 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 backup power. A less obvious interaction is with fire alarms and firefighting.

How to steal a painting (4)

Bruno visits the gallery as a tourist and leaves a smoke grenade on a timer. It 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 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. At nuclear reactors, there's typically a security rule that if a bomb is discovered, the site's locked down, with no-one allowed in or out; and a fire safety rule that in the event of a blaze, much of the staff have to be evacuated (plus perhaps some of the local population too). This raises the interesting question of which rule prevails should a bomb ever go off. And there are fire precautions that can only be used if there are effective means of keeping out innocent intruders. Many computer rooms have automatic fire extinguishers, and since fears over the ozone layer 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.

11.3.4 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 coerced into creating a vulnerability (attacks on really high-value targets such as bank cash processing centres and diamond exchanges commonly involve insiders). 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. High-value sites that take seriously the possibility of suborned insiders insist that alarm maintenance and testing be done by two people rather than one; another edge case is the prison system, where attacks on sensors, cabling and indeed the very fabric of the building are so frequent that a continuing program of test and inspection is essential. It can be useful to ask yourself, 'How would I do this differently if half my staff were convicts on day release?'

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 [483].

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

How to steal a painting (6)

Bruno phones up his rival gallery claiming to be from the security company that handles their alarms. He says that they're updating their computers so could

they please tell him the serial number on their alarm controller unit? An office junior helpfully does so — not realising 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 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 a computer that mimics it, is known as 'spoofing'. There have been reports for many years of 'black boxes' that spoof various alarm controllers. As early as 1981, thieves made off with \$1.5 million in jade statuary and gold jewelry imported from China, driving the importer into bankruptcy. The alarm system protecting its warehouse in Hackensack, New Jersey, was cut off. Normally that would trigger an alarm at a security company, but the burglars attached a homemade electronic device to an external cable to insure continuous voltage [581].

With the better 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 periodic random challenges to the controller 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 crypto algorithm may be weak, or its key may be too short (whether because of incompetence or export regulations). Even if not, Bruno might be able 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 for remote maintenance, with a default password that most users never change. Another is making the crypto key equal to the device serial number. As well as 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 and hides a few hundred yards away in the bushes. He counts the number of men in blue uniforms who arrive,

and the number who 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. He now knows he has several hours to work.

This is more or less the standard way to attack a bank vault, and it's also been used on computer installations. The *modus operandi* can vary from simply reversing a truck into the phone company's kerbside 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 [581]. In another, an Oklahoma deputy sheriff cut the phone lines to 50,000 homes in Tulsa before burgling a narcotics warehouse [1275]. In a third, a villain blew up a telephone exchange, interrupting service to dozens of shops in London's 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 future they might not involve explosives but a software-based distributed denial-of-service attack on network facilities, as computers and communications converge. 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 alarms all over New York, creating an effect similar to that of a hurricane or a power cut but at a time convenient for the crooks. Another possibility might be to run a service-denial attack against the alarm company's control centre.

An angle which seriously concerns insurers is that phone company staff might be bribed to create false alarms. So 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 — who carry most of the alarm signal traffic — prefer to concentrate it in exchanges, which makes targeted service denial attacks easier. The police are also generally suspicious of anonymous communications. These tensions are discussed in [957].

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 500g of plutonium or 2Kg of U-235 must have their alarm control center and response force on the premises [640].

Where the asset you're protecting isn't a vault but a hosting center, the network is also critical to your operations. There's little point in having eight-inch concrete walls and roofs if the single fibre connecting you to the net runs through a kerbside junction box. You'll want two buried fibres going to two different telcos — and do you want them to be using switches and routers from different vendors? Even so, the simplest way for a knowledgeable opponent to take out a hosting centre is usually to cut its communications. That's why firms have two, three or even four centres. But it would still only take four, six or eight holes in the ground to close down your operations. Who wants to dig, who knows where to, and would you detect them in time?

Finally, it's worth bearing in mind 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 arise during the day as well as at night.

11.3.5 Lessons Learned

The reader might still ask why a book that's essentially about security in computer systems should spend several pages describing walls, locks and alarm systems. There are many reasons.

- Dealing with service denial attacks is the hardest part of many secure system designs. 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 you'd most like to rely on, but also the ones 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 — the receiver operating characteristic — is a pervasive problem in security engineering.

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

- There are some lessons we can learn from the alarm business. For example, the U.S. Transportation Security Administration inserts false alarms into airport luggage to ensure that screeners stay alert; there are X-ray machines whose software inserts an image of a gun or bomb about once per shift, and there are also penetration teams who insert real objects into real suitcases. This still doesn't work very well — a 2006 report showed that 75% of the threats got through at Los Angeles and 60% at O'Hare where the screening is done once per checkpoint per shift. But it may be fixable: at San Francisco, where screeners work for a private company and are tested several times per shift, only 20% of threats get through [492].
- 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, you can't just leave the technical aspects of a security engineering project to specialist subcontractors, as critical stuff will always fall down between the cracks.

As well as 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 respond by destroying all their cryptographic key material.

11.4 Summary

Like it or not, security engineers have to deal with physical protection as well as with computers and cipher systems. Indeed, just as the confluence of computers and telecomms saw computer-industry standards and methods of working displace the old phone company ways of doing things, so the increasing automation of physical protection systems will bring the world of barriers, locks and alarms within our orbit. Future buildings are likely to have much more integrated entry controls, alarms and system security. Their management will be the job of systems administrators rather than retired policemen.

In this chapter, I highlighted a few things worth noting. First, environmental protection matters; things like architecture, landscaping and lighting can make a difference, and quite a lot is known about them.

Second, physical locks are not as secure as you might think. Recent developments in covert entry technology have led to wide publication of attacks that compromise the most widely-used mechanical locks and even the most

widely-used high-security locks. The bump keys and other tools needed for such attacks on many locks are easily available online.

Third, there's quite a lot to learn from the one aspect of physical security that is already fairly well automated, namely alarms. Alarms provide us with a good example of a system whose security policy hinges on availability rather than on confidentiality or integrity. They can give us some useful insights when dealing with service-denial attacks in other contexts.

Research Problems

At the strategic level, the confluence of physical security and systems security is bound to throw up all sorts of new problems. I expect that novel research challenges will be found by those who first explore the information / physical security boundary in new applications. From the viewpoint of security economics, I'm eager to see whether the locksmithing industry will be disrupted by its collision with digital systems, or whether the incumbents will manage to adapt. I suspect it will be disputed — but what does this teach us about the strategies existing industries should adopt as the world goes digital?

At the technical level, we will probably need better middleware, in the sense of mechanisms for specifying and implementing policy engines that can manage both physical and other forms of protection. And as for low-level mechanisms, we could do with better tools to manage keys in embedded systems. As one engineer from Philips put it to me, will the smart building mean that I have to perform a security protocol every time I change a lightbulb?

Further Reading

The best all round reference I know of on alarm systems is [118] while the system issues are discussed succinctly in [957]. Resources for specific countries are often available through trade societies such as the American Society for Industrial Security [25], and though the local insurance industry; many countries have a not-for-profit body such as Underwriters' Laboratories [1268] in the USA, and schemes to certify products, installations or both. For progress on lock bumping and related topics, I'd monitor troublemakers like the Toool group, Marc Weber Tobias, and Matt Blaze; Matt has also written on safecracking [186]. Research papers on the latest sensor technologies appear at the IEEE Carnahan conferences [643]. Finally, the systems used to monitor compliance with nuclear arms control treaties are written up in [1171].