Physical Security Devices for Computer Subsystems: A Survey of Attacks and Defenses

Steve H. Weingart

Secure Systems and Smart Card Group
IBM Thomas J. Watson Research Center, Hawthorne, NY
weingart@us.ibm.com

Abstract. As the value of data on computing systems increases and operating systems become more secure, physical attacks on computing systems to steal or modify assets become more likely. This technology requires constant review and improvement, just as other competitive technologies need review to stay at the leading edge.

This paper describes known physical attacks, ranging from simple attacks that require little skill or resource, to complex attacks that require trained, technical people and considerable resources. Physical security methods to deter or prevent these attacks are presented. The intent is to match protection methods with the attack methods in terms of complexity and cost. In this way cost effective protection can be produced across a wide range of systems and needs.

Specific technical mechanisms now in use are shown, as well as mechanisms proposed for future use. Common design problems and solutions are discussed with consideration for manufacturing.

1 Introduction

Traditionally the term 'physical security' has been used to describe protection of material assets from fire, water damage, theft, or similar perils. However, recent concerns in computer security have caused physical security to take on a new meaning: Technologies used to safeguard information against physical attack.

In this new sense, physical security is a barrier placed around a computing system to deter unauthorized physical access to the computing system itself. This concept is complementary to logical security, the mechanisms by which operating systems and other software prevent unauthorized access to data. Both physical and logical security are complementary to environmental security. Environmental security is the protection the system receives by virtue of location such as guards, cameras, badge readers, access policies, etc. The reason for separating physical and environmental security is partly due to the change in the nature of the assets being protected. In the past the assets to be protected were nominally physical items: cash, jewelry, bonds, etc. Now the assets are often information, which can be stolen without being physically removed from where they are kept. If information can be seen, it can simply be copied. This information can be anything from a spreadsheet work file to cryptographic keys. It may

Ç.K. Koç and C. Paar (Eds.): CHES 2000, LNCS 1965, pp. 302–317, 2000.

be reasonable for an individual to have access to a location (environmental security) and not to have access to the information stored on a computing system in that environment (physical security).

Physical security is also becoming more important because computing systems are moving out of environmentally secure computer rooms and into less environmentally secure offices and homes. At the same time, the value of the data on these computing systems is increasing. Logical security has also been improved so that a physical attack may become more easily performed than a logical attack [1]. We can see that the motivation to attack computing systems is increasing because the rewards for doing so are increasing.

For physical security to be effective the following criteria must be met: in the event of an attack, there should be a low probability of success and a high probability of detection either during the attack, or subsequent to penetration [17].

It is possible to build physical security systems to protect sensitive data [12,5,6,15] These systems can make unauthorized access to the data difficult, as a bank vault makes stealing cash a daunting task (tamper resistant). They can trigger mechanisms to thwart the attack, much like an alarm system (tamper responding). They can make an attempted attack apparent so that subsequent inspection will show an attack had been attempted (tamper evident).

Classification systems have been proposed that evaluate computing systems according to criteria that measure the difficulty of mounting a successful attack [16,8]. Requiring additional documentation, testing, and quality assurance further ensures increasing degrees of security. Continued work has lead to the advancement of standards [9], these standards are becoming accepted since trying to do one's own evaluation is a daunting task and the standards are being rigorously and publicly evaluated.

Physical security technology is a relatively recent addition to computing system design. This paper attempts to describe and catalog the currently known design and implementation techniques. Effort is made to differentiate between simple methods, which are applicable in areas of low criticality vs. the sophisticated methods required for protecting very critical data.

2 Kinds of Physical Security

A number of physical security methods are currently in use. This is a new field in the commercial market and is still being developed. The US government has been working on this problem for over 25 years but the results remain classified. The ways and means described here are not an exhaustive list, nor are they represented as ultimate methods. Development is continuing in protection methods and it is proceeding in attack methods. Any evaluation of appropriateness of a physical security system is time dependent and must be repeated periodically. For example, the FIPS 140 standard [9] is to be re-evaluated at five-year intervals.

2.1 Tamper Resistant

Tamper resistant systems take the bank vault approach. This type of system is typified by the outer case design of an automated teller machine. Thick steel or other robust materials are utilized to slow down the attack by requiring tools and great effort to breach the system. This type of system can be used in many environments and sometimes has the advantage of being so physically heavy (as in automated teller machines), that it resists theft by sheer weight. However recent thefts of automated teller machines by thieves using towing chains and four-wheel drive vehicles may indicate that ATMs are no longer sufficiently tamper resistant. A system that is only tamper resistant has the disadvantage that the owner may not be aware of the loss until the break-in is discovered. That may be never, if the attacker did a 'neat' job and replaced any material that had been removed.

Tamper resistant physical security is usually the easiest to apply. Steel cases and locks are well-known technology and are easily manufactured. Weight and bulk can be a problem or benefit, depending on the application.

Complexity or size can be another variety of tamper resistance. Single chip implementations of secure devices have a certain level of physical security due to the small size of the features and complexity involved in the determination of which part of a circuit performs which function. However this advantage is rapidly being lost as the equipment and skills needed to work with semiconductor devices at the microscopic level are becoming commonly available at many universities and technology centers.

Tamper responding systems use the burglar alarm approach. The defense is the detection of the intrusion, followed by a response to protect the asset. In the case of attended systems, the response may consist of sounding an alarm. Erasure or destruction of secret data is sometimes employed to prevent theft in the case of isolated systems which cannot depend on outside response. Tamper responding systems do not depend on robust construction or weight to guard an asset. Therefore, they are good for portable systems or other systems where size and bulk are a disadvantage.

Tamper evident systems are designed to ensure that if a break-in occurs, evidence of the break-in is left behind. This is usually accomplished by chemical or chemical/mechanical means, such as a white paint that 'bleeds' red when cut or scratched, or tape or seals that show evidence of removal. This approach can be very sensitive to even the smallest of penetrations. Frangible (brittle, breakable) covers or seals are other methods available using current technology.

These systems are not designed to prevent an attack or to respond to the indication that one is in progress. Their job is to ensure that the fact of a break-in will remain known and can be ascertained at a later time. An audit policy must exist, and be adhered to, for a tamper evident system to be effective. Otherwise it may not be known if, or when, the system was breached. If no one looks for the evidence of tampering, that evidence will never be found.

Some Additional Physical Security Considerations: Some of the properties of specific methods of physical security were discussed with the introduction of each type. Here, some additional points are considered. One must examine each system to determine the correct protection.

Size and Weight: The size and weight implications of a potential physical security design must be considered in the light of the application. Thick steel would not be a good idea for a portable system. A lightweight system would not be effective for an automated teller machine, as it would allow the system to be carried away more easily.

Mixed and Layered Systems: In many cases a security system can be made substantially more secure by using more than one layer and more than one kind of system. For example, a typical safety deposit vault has steel walls, an alarm system, and a high quality vault lock. These methods might seem sufficient, but the individual safety deposit boxes have significant locks as well. The individual locks serve two purposes. They provide a second layer of general security by requiring an attacker to break into each box individually after breaking into the vault. The locks on the individual boxes also serve as an additional authorization/authentication process which requires an individual to possess the correct key to open the box.

Similarly, a layer of tamper evident security placed over a layer of tamper resistance or tamper response can prevent an attack, which might be attempted over a period of days. A regular audit may turn up indications of tampering before the system is fully breached and allow additional measures to be taken before the attack is completed.

Multiple layers of security also make the attack more difficult in general. The requirement for two different kinds of tools, skills, etc., may not make the two-layered system twice as difficult to attack, but it does increase the difficulty.

3 Physical Security Methods and Mechanisms

The following sections describe different methods of physical attack that may be attempted upon computing systems, as well as the defense mechanisms that can be useful in deterring or detecting such attacks.

Physical security can be broadly divided into two categories: high technology and low technology. Low technology concepts such as inserting desktop systems into external steel cases and using floppy drive cover locks are fairly well known and will not be discussed here. The high technology examples will explore existing and contemplated attack mechanisms, and the corresponding defense mechanisms that are being brought into commercial use now, or are being considered for the near future.

3.1 High Technology Attacks

This section deals with mechanisms that used to be considered unusual. The attacks described in this section, and the defenses described in the following section, far exceed the typical levels of skills and resources available to the common attacker. However, the skill level of the common attacker is increasing. These attacks and defenses are presented to meet the requirements of markets such as banking. However as data value increases, as is occurring now with the rise of Internet commerce, these defensive techniques should become a standard part of common business practice. These techniques have are now required to meet certain government requirements [9]. The business community is also beginning to embrace these standards as a means of assurance.

Probe Attacks: The purpose of a probe attack is to directly attach conductors to the circuit(s) being protected so that information can be obtained from, and/or changes injected into, the system under attack.

Passive Probes: These are common oscilloscope or logic analyzer probes. They may be used to watch and record information contained in circuits. When used with a logic analyzer, a trigger condition may be set such that the attacker waits for a predetermined event and then begins recording.

The term passive probe is somewhat of a misnomer in that so-called passive probes may be terminated in active circuitry, which gives them very high input impedance. This may prevent their detection by, or interference with, the circuit being attacked.

Active or Injector Probes: Active probes are generally used in conjunction with passive probes. Using a pattern generator or similar device, these probes can inject signals or information into an active system.

Pico-Probes: Pico-probes can be used in either of the capacities described above. Pico-probes are very tiny and are used to directly probe the surfaces of integrated circuits.

Energy Probes: Energy probes can be electron beams, ion beams, or focused beams of light. Depending on the technology being attacked, energy probes can read or write the contents of semiconductor storage, or change control signals. Ion beam deposition has been used to successfully reconnect fuse links, to return product level smart cards to their debug-state where the output of key registers, etc., was permitted.

Machining Methods: The purpose of machining is to cut or remove material. In this context, a cover or potting material is machined to access circuitry beneath the potting or cover. Once the covering is removed, a probe attack as described above can proceed.

If the system is protected by physical security, the intent is to perform the machining operation without tripping a sensor or leaving evidence[1]. After the covering material is removed, the sensor is then disabled or bypassed so that a probing attack may proceed. If the system is protected by a tamper evident system, there may be an attempt to cover the evidence after the attack is complete.

The list of machining methods include chemical and energy methods of material removal, as well as traditional machining methods.

Manual Material Removal: Manual material removal is commonly referred to as the 'brain surgery' attack. In this scenario an attacker using a knife, or other tool, attempts to remove material from a potted or sealed container while stopping short of tripping a sensor. This attack is much more effective than might be thought. If the attacker is dexterous and has good hand-eye coordination, extremely delicate work can be accomplished.

Mechanical Machining: This method removes much material, very precisely, in the shortest time. Its disadvantages lie in the fact that there is little or no feedback. This frequently causes cuts that are too deep. If the cutter is conductive, it may be detected by the tamper detector.

Water Machining: Water machining is a very precise method for material removal. The 'cutter' can be non-conductive (if the water is pure), does not dull, and is very effective for all but very soft materials. Its chief disadvantage is that water machining equipment is typically very large. However, in situations where cost and size are a concern, but time is not, a directed slow, steady, drip of water will effectively cut through many materials given sufficient time.

Laser Machining: This technique has many of the same advantages as water. One disadvantage of laser machining is that the process may generate a great deal of heat. The laser must be tuned for the material of interest, e.g. EXCIMER (U.V.) lasers are excellent for ablating organic materials (such as epoxy).

Chemical Machining: Almost any material can be dissolved. Jet Etch[2] and similar commercial tools are very good for removing coatings and potting material cleanly. These techniques work by using a high-pressure, very precise spray of a solvent or acid to dissolve away the material. The solvent or acid may be heated to increase effectiveness. The main disadvantage is the potentially high conductivity of highly ionic cutting liquids, which may cause short circuits.

[1] If the data has an extremely short duration of value, or the audit period is excessively long, there may be no effort to cover the evidence.

[2] Jet Etch is a commercial product commonly used for removal of semiconductor surface coatings for analysis.

Shaped Charge Technology: Shaped charge technology has become commonly available to the degree where that charge precision welding and cutting sample kits are available to universities to promote the technology. These techniques have the advantages of being very accurate and being extremely fast. The penetration speed can approach 25,000 ft/sec. At these hypersonic speeds, a package can be penetrated and circuits disabled before they can respond. For example, a memory zeroing circuit can be disabled before the energy can be removed from the memory. This could give the attacker from a few seconds to a minute to finish entering a package and to reapply power to the memory before its contents decay.

TEMPEST: This is a passive attack. Electromagnetic emanations from a computer, or other electronic device, can be detected at a distance and decoded to determine contents or behavior. The distance can be many hundreds to a thousand feet or more. Power supply current profiles can also be measured to determine circuit activity.

Most information on TEMPEST is government classified in the interests of national security. However it is well known, and has been demonstrated, that a video display or serial communication line can be tapped at distances of hundreds of feet. Recently more aspects of TEMPEST technology have been independently invented/discovered in the commercial sector. Smart cards have been successfully attacked by means of studying their power supply current [10,4], and others [11] have developed new approaches to using this method.

Energy Attacks: These attacks are both of the contact and non-contact variety. However even the non-contact attacks usually require close access to the system.

Radiation Imprinting: By irradiating CMOS RAM in the X-Ray band (and possibly other bands), the contents can be 'burned in' such that power down or over-write will fail to erase the contents.

The basic imprinting attack uses radiation to imprint the CMOS RAM used to store cryptographic keys or other secret data, then the unit is physically breached without regard for power down or rewrite mechanisms. The RAM may then be read at leisure.

Temperature Imprinting: CMOS RAM will retain its contents with the power removed for seconds to hours when the temperature of the RAM is lowered. This effect starts at just below freezing. Over-writing will erase the contents.

High Voltage Imprinting: By 'spiking' CMOS RAM with short duration, high-voltage pulses, it may be possible to imprint the contents in a manner similar to radiation imprinting. This technique has not been verified by the author.

High or Low Voltage: By changing Vcc to abnormally high or low values, erratic behavior may be induced in many circuits. The erratic behaviour may include the processor misinterpreting instructions, erase or over-write circuitry failing, or memory retaining its data when not desired.

Clock Glitching: By lengthening or shortening the clock pulses to a clocked circuit such as a microprocessor, it's operation can be subverted. Instructions or tests can be skipped or generally erratic operation can be induced [2].

Circuit Disruption: This area has not yet been studied in depth by the author, however it is known that strong electromagnetic interference may cause disruption in noise-diode type random number generators and computing circuits.

Electron Beam Read/Write: The electron beam of a conventional scanning electron microscope can be used to read, and possibly write, individual bits in an EPROM, EEPROM, or RAM. To do this the surface of the chip must be exposed first, usually via chemical machining. This is a very powerful attack once the chip is exposed since buried, normally non-readable, keys and secrets can possibly be stolen and/or modified.

IR LASER Read/Write: Silicon is transparent at IR frequencies. Because of this, it is possible to read and write storage cells in a computing device by using an IR LASER directed through the bulk silicon side of the chip. By going through the bulk side there is no need to jet etch or otherwise remove the device's passivation.

Imaging Technologies: Any of the current imaging technologies including X-Ray, tomography, ultrasound, etc. can all be used to visualize the contents of a sealed or potted package. This can assist the attacker by pinpointing areas of vulnerability, identifying printed circuit card layout, showing part placement, and possibly identifying specific parts.

3.2 High Technology Defenses

The detection methods below fall into three categories: preventing intrusion, detecting intrusion, detection of noninvasive energy attacks (cold, radiation, etc.). After detection, there are various methods of response. Each method must be examined when choosing the design point. For example, a design that calls for a low temperature sensor must take into account the temperatures which the unit could be exposed to while in transport.

Tamper Resistant: This is basic bank vault technology. For example, an automated teller machine required a one inch thick mild steel case which enclosed another one-inch thick cash box [3]. These types of systems also resist theft by means of bulk. Another approach is to attach the device to the tamper barrier so firmly that the attempt to separate the layers, or to penetrate the protection, results in the destruction of the protected device.

Hard Barriers: Steel, brick, ceramics, etc., can all be used as effective barriers. As noted above, this may also help to inhibit theft.

Single Chip Coatings: This technique is used to prevent attack on the single chip level (e.g. pico-probing). The surface of the chip may not be probed with the coating in place and these coatings are applied so that removal will damage the chip beyond reclamation. This is a very complex topic as new chemistry is constantly being developed.

Insulator Based Substrates: To prevent an attacker avoiding a protective coating by using an IR LASER technique, the bulk silicon must be replaced with a material that is not transparent at useful frequencies. Silicon/Metal Oxide (SiMOX), Silicon-on- Sapphire (SOS), or other silicon-on-insulator technologies, combined with advanced passivation represent the highest level of passive, single chip, protection. One must still carefully evaluate the possibility of using surface grinding techniques to thin the substrate to the point of transparency.

Special Semiconductor Topographies: To prevent scanning electron microscope or pico-probing attacks, even in the presence of chemical machining or other techniques that can remove coatings, a chip can be designed so as not to expose critical structures without removing active layers of the device.

Tamper Evident: Tamper evident systems are not designed to prevent attack or entry into the protected area. They are designed such that entry *will* leave evidence to be discovered during physical audit.

Brittle Packages: The device is sealed in a package that is made of ceramic, glass, or another frangible material. If an attempt is made to enter the package, it cracks or shatters, leaving evidence.

Crazed Aluminum: The package is made from aluminum or other similar material, which has been heated (usually above 1000 degrees F.) and quenched. This heat treating causes a myriad of shallow, web-like cracks to appear on the surface. These cracks, like a fingerprint, are unique to each piece. The case can be photographed and subsequently audited using the photograph and optical comparison devices.

Polished Packages: Similar to crazed aluminum the package is inspected for changes in surface appearance. In this case any mark at all represents an attempted breach.

Bleeding Paint: Again, the surface quality is the auditable characteristic. Paint of one color is mixed with micro-balloons containing paint of a contrasting color. If the surface is marred, the other color "bleeds' onto the surface.

Holographic Tape: The surface of tape, with a very firm adhesive, is printed with a holographic image similar to the kind used on credit cards. This kind of tape is moderately difficult to forge, and it is constructed so that attempts to remove it will damage it (the tape may be scored to promote tearing when removal is attempted). This is good for checking to see if doors or covers have been illicitly opened. Recently there have been several incidents of holograhic seals being counterfeited.

Tamper Responding Sensor Technology: Tamper sensors cover a wide variety of devices, like the tamper evident devices above. Each type of sensor is designed to detect a particular type of intrusion. Like the example above of the automated teller machine and its steel case, certain designs are better suited for particular environments than others.

Voltage Sensors: Voltage sensors are useful in almost any design that requires proper power delivery for correct operation. Both high and low voltage can be a deliberate or accidental attack. To guarantee correct operation of circuits all power supplies should be monitored. Any excursion outside of nominal operating range should be considered an attack, and response should be engaged. References for monitors should be independent of power supply variations.

Probe Sensors: Probe sensors form a large family of active tamper barriers. Individual designs may feature tamper resistance or evidence, as well as tamper detection for additional security. Some designs are more or less costly, or heavy, or manufacturable, than others.

Wire Sensors: Thin wire wrapped around the package to be protected and then potted forms the intrusion sensor. Ideally the wire should have a high resistance so the wire can be used as a distributed resistance, so small changes can be detected as well as opens and shorts. If the wire is folded back over itself, or wound as multiple parallel strands, the sensitivity is increased because two adjacent wires may be electrically distant. So shorting two wires gives a larger signal than would two adjacent strands on a continuous wrap. The insulation on the wire should be as similar to the potting material as possible in both appearance and chemistry. This makes machining more difficult because no hints as to the whereabouts of the wire are given. Chemical attacks are made more difficult because of the difficulty of dissolving the potting without dissolving the insulation and causing shorts. It is also an advantage if the wire is made from a material which is difficult to attach to.

Printed Circuit Board Sensors: A sensor similar to the wire sensor above can be made for a much lower cost by printing the wiring onto a printed circuit board. However, the regular spacing of the lines and the usual copper conducting material give somewhat less security. This is due to the ease with which the conductors may be isolated, owing to the regularity of a rigid printed circuit board. Once a conductor is located, it is very easy to attach another wire to it for the purpose of giving the tamper detection circuitry false information. However, with good potting material and small lines, this design gives moderate security.

Flexible Printed Circuit Sensors: This design incorporates the best features of the previous two. The flexible surface helps break up the regularity of the surface planes. The lines can be made of silk-screened conductive paste, which allows high resistance. It is even better to use lines made from a conductively doped version of the same material used for final potting. The realm of package shapes is wider because the package can be “gift wrapped’ with the material, then potted. Also, the narrow screened lines will be much more difficult to find without breaking. Multiple layers can be used for additional security.

Stressed Glass Printed Circuit Sensors: Metal, or metal oxide, lines can be printed on glass, in a manner similar to a printed circuit board sensor. Contacts to the glass can be made using elastomeric ’Zebra’ connectors. Stressed glass can be obtained that is virtually impenetrable without the glass breaking. This method is very good for large flat surfaces, or possibly, for secure doors.

Stressed Glass with Piezo-Electric Sensor: Using the same glass as in the previous example, this sensor uses a piezo-electric element to signal the breakage of the glass. The force of stressed glass breaking is enough to induce a large signal from a piezo-electric device attached to the inside of the glass.

Piezo-Electric Sheet: Plastic piezo-electric sheets can be used as probe barriers. If an area protected by a piezo-electric sheet is probed or punctured, an electric charge is generated proportional to the force applied. This charge can be measured and used to activate tamper response circuitry. There are problems with this application because of sensitivity to pressure and vibration, both making the design too sensitive to environmental conditions, and potentially insensitive to slow puncture attacks.

Bulk Multiple Scattering: This sensor uses the scattering properties of coherent light through bulk materials to create a very sensitive probe sensor based on measuring the optical speckle pattern.

Motion Sensors: These sensors are typically used to sense motion in an area or box. They are often need to be used in pairs because each type can sometimes cause a false positive or can miss under unusual conditions. An infrared sensor can trip falsely when the first rays of the sun fall on the protected package through a window.

Ultra-sonic: Ultra-sonic sensors average a picture of the protected space via ultra-sonic projection and reflection. They can be very effective, but can have false positives due to air currents, etc.

Microwave: Similar to ultra-sonic, with the same strengths and weaknesses, but at a higher frequency. The material of the walls of the protected area have to be taken into account with this type of system since some non-metallic materials can be transparent at these frequencies. This can cause false positives due to activity outside of the protected region.

Infra-red: This sensor is not typically sensitive to air currents or the like, but these systems have been known to trip due to light (and heat) changes due to sunrise through windows when the averaging is too sensitive. They are most useful for detecting warm bodies, people, animals, etc. A tool at ambient temperature will probably not be noticed unless it was moved to suddenly block an infrared radiating source that the sensor already 'sees.'

Acceleration Sensors: These sensors are used to detect movement or vibration. Their primary uses are to prevent theft, and to detect drilling or hammering.

Solid State: This sensor detects a beam of light reflected by mirrors that are attached to flexible mounts, or a piezo-electric device and a small mass. They are quite sensitive and reliable.

Micro-switches: Micro-switch motion sensors use mercury or pendulums to detect motion. They are lower in cost than solid state devices, but are less sensitive, and are more prone to failures. However, a liquid mercury switch can be reliable and virtually without wear.

Radiation Sensors: Radiation sensors are used to detect attempts at radiation imprinting. These sensors are most important for remotely located systems which could be taken into a laboratory and attacked.

Flux Sensors: Flux sensors sense the real-time radiation intensity. The advantage of this type of circuit is that it can be very low cost. The disadvantage is that this sensor has no cumulative memory (total dose measurement). If the data is invariant over a long period of time, low levels of radiation (below the sense point) can imprint the data. Given the power and cost budget for typical physical security systems, integrating the flux reading is too costly. So a compromise must be struck as to flux level trip point vs. minimum time to imprinting.

Phototransistors can be very effective radiation flux sensors. The circuit is the same as is used for light measurement, however a higher gain is typically required. The typical problems with this circuit are that the sensitivity in the radiation band of interest is usually not specified by the manufacturer and must be determined by testing, and that the sensors tend to degrade with time and exposure to radiation.

Dosage Sensors: These sensors store the total radiation dose over time. Total dose is the best indicator of imprinting in CMOS SRAM. Unfortunately, at this time there are no available dosage sensors which are small, low cost, low power, and directly readable.

Temperature Sensors: Temperature sensors are well-known and readily available at all cost performance points.

Tamper Responding - Response Technology: The methods of tamper response technology discussed here are means of removing data from RAM circuits which presumably contain secret information. This is currently the most common method of storing such information because the retention is reliable and the erasure is reasonably so. If one were to use the highest level of technology available to attempt recovery of data that had been stored and then erased on almost *any* known media, there is little, outside of physical destruction, that can prevent recovery.

RAM Power Drop: This is the most straightforward method of data erasure. If aided by a crowbar circuit that supplies a very low impedance path from Vcc to ground, it is reliable if imprinting protection (temperature sensing and radiation sensing) has been employed. Since there is a tendency for RAM contents to imprint over time, any information that is to be stored in RAM for long periods should be regularly scrambled, inverted, or otherwise changed to prevent imprinting.

RAM Overwrite: This method has had the widest acceptance in government specifications [17], however in a catastrophic condition it is difficult to guarantee that reliable power will be available to operate the over-write circuit. The common method is to over-write some number of times with all 0's, then all 1's.

It would seem random or pseudo-random data would be more effective, but this has not been shown. It would also take even longer to complete the overwrite since the data would have to be generated.

Physical Destruction: This is the only method of data erasure that is completely reliable. Destruction can be accomplished with a minimum of overt violence. The occurrence would barely be detectable at the surface of a metal hybrid package. Nonetheless, this method is typically reserved for the most sensitive circumstances.

3.3 Operating Envelope Concept

One of the main problems encountered while implementing physically secure systems is the prevention of the class of attacks that cause erratic operation. This can occur when the operating point is pushed to the boundaries of the operating range. For example, running the circuit at either marginally high or low supply voltages may cause erratic operation of the circuit such that secret information could be leaked. If one considers the possibility of adjusting both temperature and voltage, the problem can become even more complex.

Manufacturers define the operating range of the components that they make, but often the specification is incomplete. It can be incomplete because no one ever intended the part to be used in some particular way, and the manufacturer, justifiably, doesn't want to deal with the problem. In general, designers can design circuits that stay within prescribed limits and the circuit functions properly. For example, if the circuit is run at too high a temperature while at too low a supply voltage, the condition is undefined. This may open the system to attack.

It is the physical security designer's responsibility to determine the safe operating envelope of the circuit under all conditions, and to provide safeguards to detect conditions outside of the acceptable operating envelope. If these conditions are detected the response circuitry must protect the secret data. This is the basic idea behind the environmental failure protection requirement in FIPS 140-1 [9]. If conditions leave the safe operating envelope in a non-catastrophic manner (e.g. Vcc drop during power down), the system should be stopped (or held reset). If conditions leave the safe operating envelope in a catastrophic manner (e.g. ambient temperature exceeding safe operating range), the critical data should be erased and the system should be prevented from operating.

Secure designs should also employ good engineering practice to prevent improper clock signals from reaching sensitive circuits by use of phased lock loops (PLLs), or similar techniques. Power analysis attacks should be prevented by designing in adequate power filtering to reduce information leakage.

4 A High Technology Physical Security Design Example

The following design began as a concept and has now been developed.

A small printed circuit board contains the microcomputer, cryptographic, and tamper detection/response circuitry. The circuitry on the card includes voltage, temperature and radiation sensors to protect the battery backed-up CMOS SRAM from becoming imprinted as well as circuitry to erase the SRAM by power down with a crowbar to ground on the SRAM power pin. Circuitry also guarantees that the contents of the rest of the system (key registers, microprocessor contents, etc.) are lost on tamper. Additional circuitry monitors the tamper detection screen which surrounds the entire assembly. The tamper detection screen is constructed of conductive organic lines on a polyester substrate. These lines are arranged in a configuration so that changes in the resistance of the lines caused by shorting, breaking, or otherwise damaging the lines are detectable. The assembly is then potted using an organic material similar in composition to the conductors, in a metal case which serves as an electrcal shield [8].

If the design is examined it can be seen that a number of attacks have been anticipated and guarded against. The voltage, temperature and radiation sensors ensure that cold and radiation attacks will not succeed in causing imprinting. The voltage sensors protect from imprinting and disruption. The SRAM power down and crowbar circuit will reliably erase the SRAM in the event of an attack. The SRAM devices used in the design have been tested to assure erasure when the power down circuit is activated.

The probe sensor is sensitive to very small probes and the potting makes machining very difficult because the uneven surface of the polyester under the hard potting would make cutting too deep quite likely. Even if the screen were reached successfully, the lines are very difficult to manipulate and would most likely be damaged in the attack attempt, triggering the power down. The metal case protects additionally from machining as well as acting as an electrical interference barrier (Faraday cage).

This design has also been tested, and has been found not susceptable, to power analysis attacks.

This design is representative of the commercial state of the art in physical security design, and has been validated at FIPS 140-1 level 4 overall [13,14].

5 Conclusions

Physical security devices like those described here are becoming desirable in areas where a technical means for ensuring data secrecy is required. As data values climb, the motivation for using physical means to extract data from computing systems is steadily increasing. System design must meet this growing need for protection.

As with any developing technology, the design point and performance must be constantly reviewed. The technology of potential adversaries, as well as the value of the data which motivates these individuals, is increasing. So the technology and quality of the protection must keep up with the skills of the attackers.

Acknowledgements

The author would like to thank several people for their contributions to this work. Years ago, Gus Simmons gave some critical pointers that led me to deeper work in physical security, I thank him for that! Steve R. White and Bill Arnold at the IBM T. J. Watson Research Center, and Glen Double from IBM Charlotte, were central to this work for many years. The other members of the IBM 4758 team at Watson, Elaine Palmer, Sean Smith, Ron Perez, Mark Lindemann and Joan Dyer have all helped and shared ideas. Huge thanks to my wife Pat, for patience, support and copy editing.

References

1. R. Anderson, M. Kuhn, 'Tamper Resistance - A Cautionary Note', The Second USENIX Workshop on Electronic Commerce Proceedings, Oakland, California, November 18-21, 1996, pp 1-11, ISBN 1-880446-83-96.
2. R. Anderson, M. Kuhn, 'Low Cost Attacks on Tamper Resistant Devices'
3. R. E. Anderson, 'Bank Security', Butterworth Publishers 1981, pp. 9l-93.
4. S. Chari, C.S. Jutla. J.R. Rao, and P.Rohatgi. 'A Cautionary note regarding evaluation of AES candidates on smart cards'. Proceedings of Second AES Conference, Rome, Mar 1999.
5. David Chaum, 'Concepts for Design of Tamper Responding Systems', Advances in Cryptology, Proceedings of Crypto '83 , Plenum Press 1984, pp.387-392.
6. Andrew 1. Clark, 'Physical Protection of Cryptographic Devices', presented at Eurocrypt '87, Amsterdam.
7. 'Department of Defense Trusted Computer System Evaluation Criteria', U. S. Department of Defense, 5200.28 STD
8. G. P. Double, 'Physical Security for Transaction Systems: A Design Methodology', IBM Technical Report, TR 83.227 IBM 1990.
9. 'Federal Information Processing Standard 140-1: General Security Requirements for Equipment Using the Data Encryption Standard', National Institute for Standards and Technology
10. P. Kocher, J. Jaffe and B. Jun. 'Introduction to Differential Power Analysis and Related Attacks.' Manuscript, Cryptography Research, Inc. 1998.
11. M. Kuhn and R. Anderson, 'Soft Tempest: Hidden Data Transmission Using Electromagnetic Emanations', Information Hiding 1998, LNCS 1525, pp. 124-142, 1998.
12. W. L. Price, 'Physical Security of Transaction Devices', NPL Technical Memo DITC 4/86, National Physical Laboratory, Jan, 1986.
13. S.W. Smith, S.H. Weingart, 'Building a High Performance, Programmable Secure Coprocessor.' Computer Networks (Special Issue on Computing Network Security). 31: 831-860. April 1999.
14. S.W. Smith, V. Austel, R. Perez, S. Weingart. 'Validating a High-Performance, Programmable Secure Coprocessor or, the World's First FIPS 140-1 Level 4.' 22nd National Information Systems Security Cconerence, October 199.
15. S. H. Weingart, 'Physical Security for the uABYSS System', Proceedings of IEEE Symposium on Security and Privacy 1987, IEEE Publications, pp. 52-58.
16. S. H. Weingart, S. White, W. Arnold, and G. Double, 'An Evaluation System for the Physical Security of Computing Systems', Proceedings of the Sixth Annual Computer Security Applications Conference 1990, IEEE Publications, pp. 232-243.
17. U. S. Federal Standard 1027 Telecommunications: General Security Requirements for Equipment Using the Data Encryption Standard.