The Sorcerer's Apprentice Guide to Fault Attacks

HAGAI BAR-EL, HAMID CHOUKRI, DAVID NACCACHE, MICHAEL TUNSTALL, AND CLAIRE WHELAN

Invited Paper

The effect of faults on electronic systems has been studied since the 1970s when it was noticed that radioactive particles caused errors in chips. This led to further research on the effect of charged particles on silicon, motivated by the aerospace industry, which was becoming concerned about the effect of faults in airborne electronic systems. Since then various mechanisms for fault creation and propagation have been discovered and researched. This paper covers the various methods that can be used to induce faults in semiconductors and exploit such errors maliciously. Several examples of attacks stemming from the exploiting of faults are explained. Finally a series of countermeasures to thwart these attacks are described.

Keywords*—Fault attacks, glitch attacks, side-channel attacks, smart cards.*

I. Introduction

One of the first examples of faults being injected into a chip was accidental. It was noticed that radioactive particles produced by elements naturally present in packaging material [1] caused faults in chips. Specifically, uranium-235, uranium-238, and thorium-230 residues present in the packaging decay to lead-206 while releasing α particles. These particles create a charge in sensitive chip areas, causing bits to flip. While these elements were only present in two or three parts per million, this concentration was sufficient to affect chip behavior. Subsequent research included studying and simulating the effects of cosmic rays on semiconductors [2]. Cosmic rays are very weak at ground level due to the earth's atmosphere, but their effect becomes more pronounced in the upper atmosphere and outer space. This problem is further compounded by the fact that the more RAM a computer has, the higher the chance of a fault occurring. This has provoked a great deal of research by organizations such as NASA and Boeing. Most of the work on fault resistance was motivated by this vulnerability to charged particles. Considerable engineering endeavors were devoted to the "hardening" of electronic devices designed to operate in harsh environments. This has mainly been done using simulators to model circuits and study the effect of randomly induced faults. Various fault induction methods have since been discovered but all have in common similar effects on chips. One such example is the use of a laser to imitate the effect of charged particles [3]. The different faults that can be produced have been characterized to enable the design of suitable protections. The first attack that used a fault to derive secret information [4] targeted the RSA public-key cryptosystem. Basically, by introducing a fault into one of the primes, the modulus can be exposed and as a result compromise the RSA system. This led to similar attacks on other cryptographic algorithms. The countermeasures that can be used to thwart fault attacks had already been largely defined and successfully deployed.

This survey is organized as follows. In Section II the various methods of fault injection and their effects are described. We then turn to theoretical (Section III) and practical (Section IV) attacks. Finally, countermeasures are described in Section V.

II. Methods of Fault Injection

The most common fault injection techniques are as follows.

1) *Variations in supply voltage* during execution may cause a processor to misinterpret or skip instructions. This method is widely researched and practiced behind closed doors by the smart-card industry but does not often appear in the open literature.
2) *Variations in the external clock* may cause data misread (the circuit tries to read a value from the data bus before the memory had time to latch out the asked value) or an instruction miss (the circuit starts executing instruction

Manuscript received August 4, 2004; revised December 20, 2004. The work of C. Whelan is supported by the Irish Research Council (IRCSET).

H. Bar-El is with Discretix Technologies Ltd., Rehovot 76574, Israel (e-mail: hagai.bar-el@discretix.com).

H. Choukri is with IXL Laboratory, Bordeaux 1 University, Talence Cedex F-33405, France (e-mail: h.choukri@voila.fr).

D. Naccache and M. Tunstall are with the Information Security Group, Royal Holloway, University of London, Egham TW20 0EX, U.K. (e-mail: david.naccache@rhul.ac.uk; m.j.tunstall@rhul.co.uk).

C. Whelan is with the School of Computing, Dublin City University, Dublin 9, Ireland (e-mail: cwhelan@computing.dcu.ie).

Digital Object Identifier 10.1109/JPROC.2005.862424

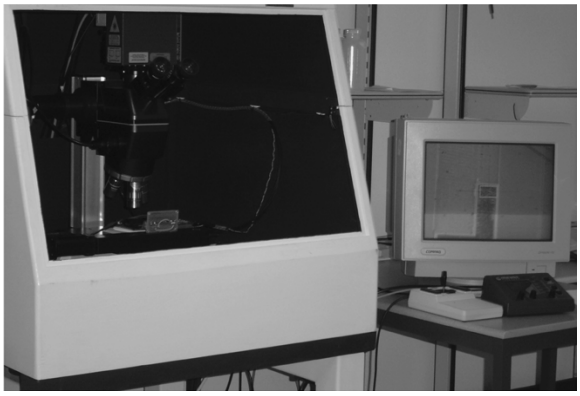

Fig. 1. Laser fault injection equipment.

$n+1$ before the microprocessor finished executing instruction n).

3) *Temperature*: circuit manufacturers define upper and lower temperature thresholds within which their circuits will function correctly. The goal here is to vary temperature using an alcoholic cooler until the chip exceeds the threshold's bounds. When conducting temperature attacks on smart cards (never documented in the open literature to the authors' knowledge) two effects can be obtained: the random modification of RAM cells due to heating and the exploitation of the fact that read and write temperature thresholds do not coincide in most nonvolatile memories (NVMs). By tuning the chip's temperature to a value where write operations work but reads do not or the other way around a number of attacks can be mounted (components are classified into three temperature vulnerability classes the description of which is beyond the scope of this survey).
4) *White light*: All electric circuits are sensitive to light due to photoelectric effects. The current induced by photons can be used to induce faults if a circuit is exposed to intense light for a brief period. This can be used as an inexpensive means of fault induction [5].
5) *Laser* can reproduce a wide variety of faults and can be used to simulate [3] faults induced by particle accelerators [6], [7]. The effect produced is similar to white light but the advantage of a laser over white light is directionality that allows to precisely target a small circuit area. Examples of laser fault injection equipment is shown in Figs. 1 and 2.
6) *X-rays and ion beams* can also used as fault sources (although less common). These have the advantage of allowing the implementation of fault attacks without necessarily depackaging the chip.

A. The Different Types of Faults

Electronic circuits can be subject to two classes of faults: provisional (transient) and destructive (permanent) faults. In a provisional fault, silicon is locally ionized so as to induce a current that, when strong enough, is falsely interpreted by the circuit as an internal signal. As ionization ceases so does the induced current (and the resulting faulty signal) and the chip recovers its normal behavior. By opposition, destructive faults, created by purposely inflicted defects to the chip's structure, have a permanent effect. Once inflicted, such destructions will affect the chip's behavior permanently.

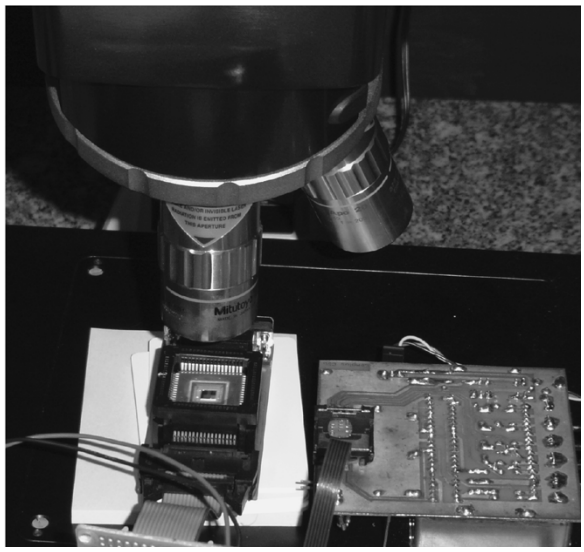

Fig. 2. Laser fault injection equipment (inner view).

1) Provisional Faults (Taxonomy): Provisional faults have reversible effects and the circuit will recover its original behavior after the system is reset or when the fault's stimulus ceases.

- *Single-event upsets* (SEUs) are flips in a cell's logical state to a complementary state. The transition can be temporary, if the fault is produced in a dynamic system, or permanent if it appears in a static system. SEU was first noticed during a space mission in 1975 [8], [9] and stimulated research into the mechanisms by which faults could be created in chips. SEUs can also manifest themselves as a variation in an analogue signal such as the supply voltage or the clock signal.
- *Multiple-event upsets* (MEUs) are a generalization of SEUs. The fault consists of several SEUs occurring simultaneously. A high integration density is a risk factor that can provide conditions favorable to the genesis of MEUs.
- *Dose rate faults* [10] are due to several particles whose individual effect is negligible but whose cumulative effect generates a sufficient disturbance for a fault to appear.

2) Destructive Faults (Taxonomy):

- *Single-event burnout faults* (SEBs) are due to a parasitic thyristor being formed in the MOS power transistors [11], [12]. This can cause thermal runaway in the circuit causing its destruction.
- *Single-event snap back faults* (SESs) [13] are due to the self-sustained current by the parasitic bipolar transistor in MOS transistor channel N. This type of fault is not likely to occur in devices with a low supply voltage.

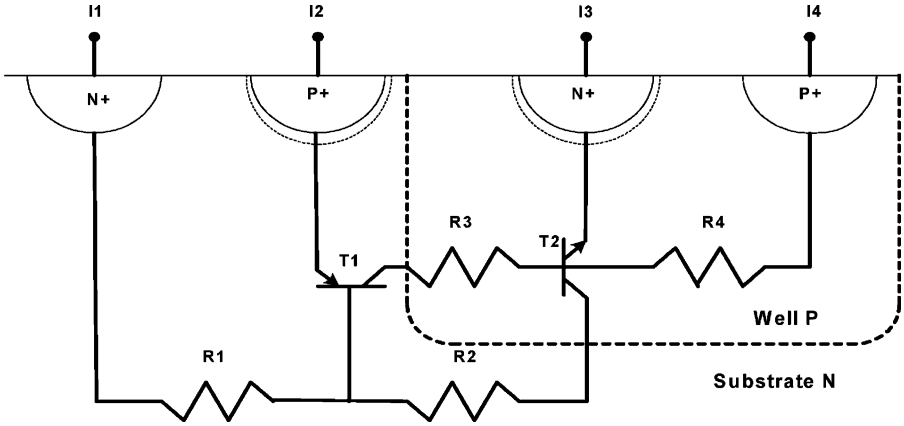

Fig. 3. Single event latch-up—parasitic transistors T1 and T2.

- *Single-event latch-up faults* (SELs) [6], [14]depicted in Fig. 3, are propagated in an electronic circuit by the creation of a self-sustained current with the releasing of PNPN parasitic bipolar transistors in CMOS technology. This can potentially destroy the circuit.
- *Total dose rate faults* [15] are due to a progressive degradation of the electronic circuit subsequent to exposure to an environment that can cause defects in the circuit [16].

When using fault injection as an attack strategy provisional faults are the method of choice. These allow for faults under numerous experimental conditions to be attempted until the desired effect is achieved. As a side-bonus the system remains functional after the attack's completion. By opposition, a destructive fault would (usually) render the target unusable and will necessitate the manufacturing of a clone.

III. Fault Attacks in Theory

The first academic fault attack paper [4] proposed a number of methods for attacking public key algorithms. One attack focused on an implementation of RSA using the Chinese Remainder Theorem (CRT). The attack is very simple as it only requires one fault to be inserted in order to factor the RSA modulus. Basically the attack works as follows.

A. Fault Attack on RSA Signature With Chinese Reminder Theorem

Let $N = p \times q$, where p and q are two large prime numbers. Let $m \in \mathbb{Z}_N^*$ be the message to be signed, d the private key and s the RSA signature. We denote by a and b the precomputed values required for use in the CRT, such that

$$\begin{cases} a \equiv 1 \pmod{p} \\ a \equiv 0 \pmod{q} \end{cases} \text{and} \begin{cases} b \equiv 0 \pmod{p} \\ b \equiv 1 \pmod{q} \end{cases}$$

and define

$$\begin{aligned} d_p &= d \pmod{p-1} \\ d_q &= d \pmod{q-1}. \end{aligned}$$

Using repeated squaring calculate

$$\begin{aligned} s_p &= m^{d_p} \pmod{p} \\ s_q &= m^{d_q} \pmod{q}. \end{aligned}$$

The RSA signature s is then obtained by the linear combination $s = a \times s_p + b \times s_q \pmod{N}$

The attack is based on being able to obtain two signatures of the same message, where one signature is correct and the other faulty. By "faulty" we mean that a fault injected during the computation corrupted either the computation of s_p or s_q.

Let $\hat{s} = a \times s_p + b \times \hat{s}_q \pmod{N}$. be the faulty signature (we arbitrarily assume that the error occurred during the computation of s_q but the attack works just as well when s_p is corrupted). Subtraction yields

$$\begin{aligned} \Delta &\equiv s - \hat{s} \\ &\equiv (a \times s_p + b \times s_q) - (a \times s_p + b \times \hat{s}_q) \\ &\equiv b(s_q - \hat{s}_q) \pmod{N} \end{aligned}$$

Hence, as $b \equiv 0 \pmod{p}$ and $b \equiv 1 \pmod{q}$ it follows that $\Delta \equiv 0 \pmod{p}$ (but $\Delta \not\equiv 0 \pmod{q}$) meaning that Δ is a multiple of p (but not of q). Hence, a GCD calculation gives the secret factors of N: $GCD(\Delta \pmod{N}, N) = p$ and $q = N/p$.

In summary, all that is required to break RSA is one correct signature and one faulty one. This attack will be successful regardless of the type or number of faults injected during the process provided that all faults affect the computation of s_p or (mutually exclusive or!) s_q.

This attack was extended in [34] to show that it is not necessary to generate a correct signature to achieve this attack. The faulty signature can be compared to the message, as $GCD(\hat{s}^e - m \pmod{N}, N) = p$, where e is the public verification exponent.

Although initially theoretical, this attack (implemented in [17]) stimulated the genesis of a variety of fault attacks against a wide gamut of cryptographic algorithms. The following subsections describe some more of these attacks.

B. Fault Attack on RSA Signature Without CRT

Suppose that one bit in the binary representation of the secret key d flips from from one to zero or vice versa, and that this faulty bit position is randomly located. An attacker arbitrarily chooses a plaintext m and computes the signature s. While the signature is being generated, the attacker generates a fault such that the one bit of the private exponent is changed resulting in a faulty signature $\hat{s}$. Assuming that the ith bit flips to its complement, then the division of the generated $\hat{s}$ by s, as described in [18], gives

$$\frac{\hat{s}}{s} \equiv m^{\hat{s}-s} \equiv \begin{cases} m^{2^i} \pmod N & \text{if } i\text{th bit of } d = 0, \\ \frac{1}{m^{2^i}} \pmod N & \text{if } i\text{th bit of } d = 1. \end{cases}$$

In [35] it was shown that this can be optimized by raising the formula to the power of the public exponent e, giving

$$\frac{\hat{s^e}}{m} \equiv \begin{cases} (m^e)^{2^i} \pmod N & \text{if } i\text{th bit of } d = 0, \\ \frac{1}{(m^e)^{2^i}} \pmod N & \text{if } i\text{th bit of } d = 1 \end{cases}$$

where it is not necessary to know the correct signature but just the signature with a fault.

An attacker can compare $\hat{s}^e/m$ to all the possible signatures that can be created by one bit fault errors in d. This can be done quickly as only $2 \log N$ signatures are possible. This process is repeated until enough information is obtained on d to derive all the key elements.

It has been proposed that this attack can be extended to when more than one bit is changed [18]. However, this attack works if and only if one bit is changed. If, for example, two bits (j and k) are changed, then the signature will be interpreted using the following relationship:

$$\frac{\hat{s}}{s} \equiv m^{\hat{s}-s} \equiv m^{\pm 2^j \pm 2^k}.$$

The $\pm$ depending on how the bit has been changed. Another fault on two bits (l and n) will yield the same value for $\hat{s}$ where:

$$\pm 2^j \pm 2^k = \pm 2^l \pm 2^n$$

A variant of this attack can be applied to discrete logarithm based public key cryptosystems such as DSA, this is described in [18].

C. Fault Attacks on Key Transfer or NVM

In this scenario [19], a fault is injected during the transfer of secret data from one memory component to another. Although the attack is applicable to any algorithm, let us assume that a DES key is being transferred from EEPROM to RAM in a smart card. If we change the value of parts of the key to some fixed value (for example, one byte at a time), it becomes possible to derive the secret key.

Table 1
The Biham–Shamir Attack

Input	DES Key	Output
$M \rightarrow$	K_0 =XX XX XX XX XX XX XX XX	$\rightarrow C_0$
$M \rightarrow$	K_1 =XX XX XX XX XX XX XX 00	$\rightarrow C_1$
$M \rightarrow$	K_2 =XX XX XX XX XX XX 00 00	$\rightarrow C_2$
$M \rightarrow$	K_3 =XX XX XX XX XX 00 00 00	$\rightarrow C_3$
$M \rightarrow$	K_4 =XX XX XX XX 00 00 00 00	$\rightarrow C_4$
$M \rightarrow$	K_5 =XX XX XX 00 00 00 00 00	$\rightarrow C_5$
$M \rightarrow$	K_6 =XX XX 00 00 00 00 00 00	$\rightarrow C_6$
$M \rightarrow$	K_7 =XX 00 00 00 00 00 00 00	$\rightarrow C_7$

We DES-encrypt a message M to obtain a faultless ciphertext C_0. Then, during the key transfer from EEPROM to RAM, one key byte is changed to a fixed known value (00 in our example). The resulting C_1 is recorded and the process is repeated by forcing two bytes to a fixed value, then three bytes, and so on. This continues until the whole key but one byte has been set, byte by byte, to the fixed value.

This procedure shown in Table 1, where C_i represents the ciphertext of an unknown key with i bytes set to a fixed value. Once this data has been collected, it can be used to derive the DES key.

Let K_n represent the original DES key with n bytes replaced with known values. To find K_7 the 128 different possible values for the first byte of the DES key are tried until one produces the ciphertext C_7.[1] After this K_6 can be found by searching through the 128 different possible, values for the second byte, as the first byte will be known. Finding the entire key will require a search through a key space of 1024 different keys. This attack can also be used when unknown data is manipulated by an known algorithm. The prerequisite for doing so is the ability to rekey the device (running the unknown algorithm) with keys of our choosing. In which case the exhaustive search phase can be performed on the attacked device itself.

Historical note: An attack similar to [19] was discovered and documented (but never published) in 1994. The code was that of a smart-card operating system where a special file contained DES keys saved in records. This OS featured two commands: `erase` i, a command that erases the ith key record and `encrypt` i, M a command that outputs the ciphertext of the message M using the key contained in the ith record. While invisible for the user, the OS was using the convention that all-zero keys are free records (an `encrypt` command on a zero (erased) record would return an error). The attack here was exploiting the fact that EEPROM could only be erased by 32-block units. In other words, upon an `erase`, the OS would erase twice four bytes. The attack consisted of encrypting a message with an unknown key and then instructing the OS to erase this key but cutting power just after the first 32-bit block's deletion. The card will then contain a

[1] Although a byte is changed, only 128 different values are possible, as the least significant bit is a parity bit.

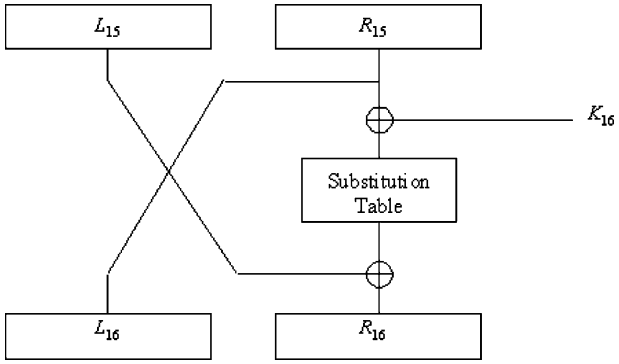

Fig. 4. Simplified DES last round model.

56-bit key which rightmost half is zeroed (which is not interpreted by the OS as an empty key record!). An encryption with this key followed by two 2^{28} exhaustive search campaigns would have eventually revealed the key.

Since that date, OSs associate a security bit σ to each key. When a user instructs to delete a key, the σ bit is erased first, thereby recording the information that the key cannot be used anymore for cryptographic operations. Only then will the OS undertake the task of erasing the key's actual bits. Upon reset, the OS ascertains that all $\sigma = 0$ keys contain zero bytes if any nonzero $\sigma = 0$ keys are found, the OS simply resumes the deletion of their bits.

D. Fault Attacks on DES

DES is a 16-round secret key algorithm based on a Feistel structure. This attack targets DES' 15th round. We use a simplified description of the last round (Fig. 4) to explain what happens when the 15th round does not execute properly.[2]

The output of the last round can be expressed as:

$$\begin{aligned} R_{16} &= S(R_{15} \oplus K_{16}) \oplus L_{15} \\ &= S(L_{16} \oplus K_{16}) \oplus L_{15}. \end{aligned}$$

If a fault occurs during the execution of the 15th round, i.e., R_{15} is changed into a faulty $\hat{R}_{15}$, then

$$\begin{aligned} \hat{R}_{16} &= S(\hat{R}_{15} \oplus K_{16}) \oplus L_{15} \\ &= S(\hat{L}_{16} \oplus K_{16}) \oplus L_{15}. \end{aligned}$$

If we XOR R_{16} and $\hat{R}_{16}$, we get

$$\begin{aligned} R_{16} \oplus \hat{R}_{16} &= S(L_{16} \oplus K_{16}) \oplus L_{15} \oplus S(\hat{L}_{16} \oplus K_{16}) \oplus L_{15} \\ &= S(L_{16} \oplus K_{16}) \oplus S(\hat{L}_{16} \oplus K_{16}). \end{aligned}$$

This gives a relationship where only the value of the 16th subkey (K_{16}) is unknown; all the other variables being given directly as an output of the DES. For each substitution table used in the last DES round this relationship will be true. An exhaustive search of the 64 possible values that validate this equation can be conducted for each of the six bits corresponding to the input of each substitution table. This will give approximately 2^{18} different hypotheses for the last subkey leading to a final exhaustive search through 2^{26} DES keys to find the whole key. In practice, it is simplest to conduct the attack several times either at different positions in the 15th round or with a varying message. When the lists of possible hypotheses are generated the actual subkey will show up in the intersection of all the sets of hypotheses. If the difference between the two output values for a given substitution table (R_{16} and $\hat{R}_{16}$) is zero, then all the possible values of K_{15} for that substitution table will be valid. This means that it is advantageous to induce a fault as early as possible in the 15th round so that the effect of the fault spreads over as many different substitution tables in the 16th round as possible.

[2]In Fig. 4 bit permutations were removed, as these do not fundamentally change theory although they somewhat complicate explanation.

E. Other Fault Attacks—Further Reading

While the bibliography on the matter would be too voluminous to overview exhaustively, the authors attract the reader's attention to a more powerful attack [20] applicable to all secret key algorithms. Several authors have published other fault attacks on DES [33] and other algorithms such as AES [21], [22] and RC5 [23]. The details of these are beyond the scope of this paper and are presented as further reading.

IV. Some Experimental Fault Attacks

In a glitch attack, the attacker deliberately generates a malfunction that causes one or more flip-flops to transition into a wrong state. The aim is usually to replace a single critical machine instruction with an almost arbitrary one. Glitches can also aim to corrupt data values as information is transferred between registers and memory [24]. There are three main techniques for creating fairly reliable malfunctions that affect only a very small number of machine cycles in smart-card processors. These are clock signal transients, power supply transients, and external electrical field transients. All three were successfully experimentally implemented. Particularly interesting instructions that an attacker might want to target with glitches are conditional jumps or the test instructions preceding them. They create a window of vulnerability in the processing stages of many security applications that often allow the attacker to bypass sophisticated cryptographic barriers by simply preventing the execution of the code that detects that an authentication attempt was unsuccessful. Instruction glitches can also be used to extend the runtime of loops—for instance, in serial port output routines—to see more of the memory after output buffer, or reduce the runtime of loops, thereby transforming an iterated block-cipher into an easy to break single-round variant [24]. Clock-signal glitches are currently the simplest and most practical ones. They temporarily increase the clock frequency for one or more half cycles, such that some flip-flops sample their input before the new state has reached them. Power analysis was used by this survey's authors to monitor how far a program has progressed and launch a fault as the power profile of a specific instruction was recognized. This in turn can be used to determine when, for example, a branch instruction

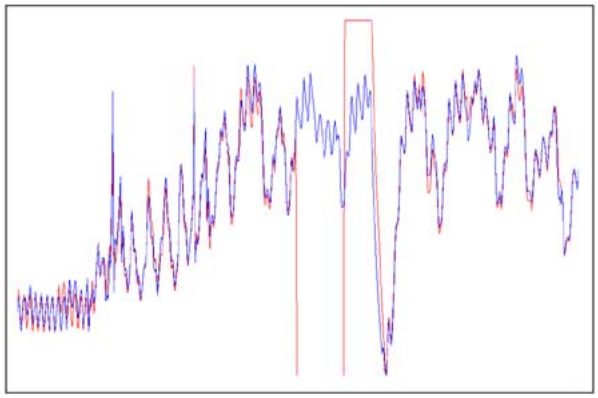

Fig. 5. Instruction-only glitch attack.

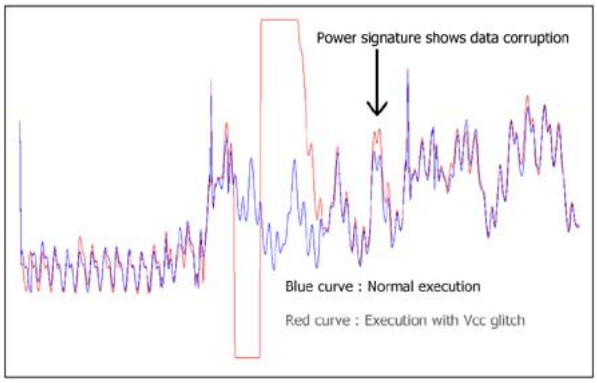

Fig. 6. Instruction and data glitch attack.

is about to be taken. A more rapid clock cycle at this point (a clock glitch) may provide insufficient time for the processor to write the jump address to the program counter, thereby annulling the branch operation [25]. A similar clock-glitch attack is also presented in [23]. Because of the different number of gate delays in various signal paths and the varying parameters of the circuits on the chip, this affects only some signals, and by varying the precise timing and duration of the glitch, the CPU can be fooled to execute a number of completely different, wrong instructions. These will vary from one instance of the chip to another, but can be found by a systematic search using specialized hardware.

The following figures illustrate different effects that glitches can have. In this experiment power was dropped from $V_{\rm cc}$ to 0 V during a few nanoseconds. By carefully playing with the glitch's parameters (duration, falling edge, amplitude etc.) two types of behavior were obtained.

- Under a first set of conditions (Fig. 5), the processor just skipped a number of instructions and resumed normal execution several microseconds after the glitch. This fault allows the selective execution of instructions in a program.
- Under a second set of conditions, not only does the processor skip instructions, but the value of data manipulated by the processor is also modified in a precise manner. This is visually reflected in the power curves of Fig. 6.

It should be noted that a third set of conditions was tested in this experiment. Although the results are not shown here, the outcome was that the value of data could be corrupted while the interpretation of instructions was left unchanged.

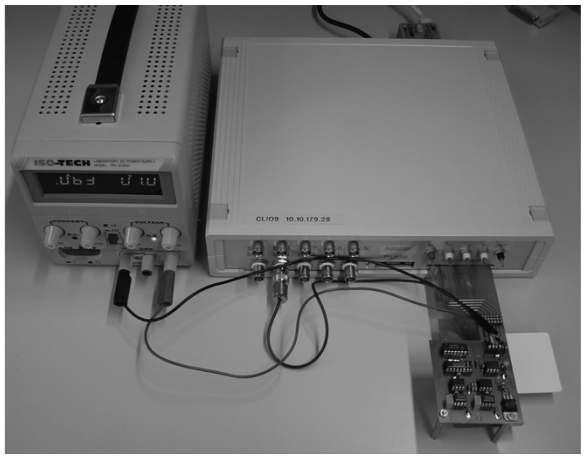

Fig. 7. Glitch fault attack board with CLIO reader.

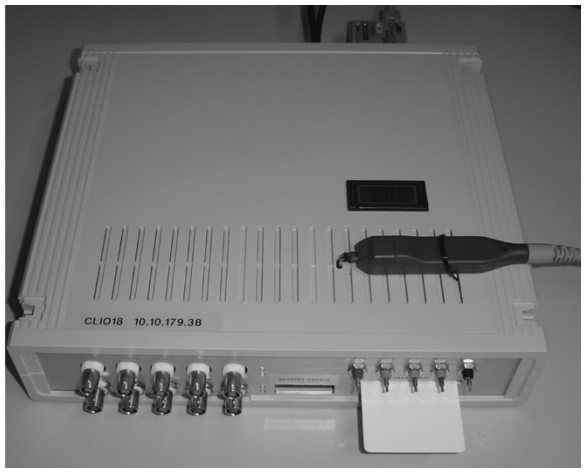

Fig. 8. A modified CLIO reader.

Figs. 7 and 8 show glitch injection electronics used in mounting these attacks. The board shown in Fig. 7 was developed to perform glitch attacks. The board accepts a signal from a CLIO reader instructing the board to apply a lower voltage to the $V_{\rm cc}$ for the duration of that signal. The levels of voltage that are applied during the glitch are controlled via potentiometers configured with a screwdriver. A similar setup was used to modify the clock sent to the card for short periods of time. This type of setup is very inexpensive but requires that all the possible glitch configurations be tested by hand.

Fig. 8 shows a modified CLIO reader that can be used to inject a glitch at a specific point during a command. This setup can be configured via the network to allow for a large number of glitch configurations to be tested when searching for vulnerabilities in new chips. This is more expensive than building a simple electronic board but can be automated so that more glitch configurations can be tested.

Glitch attacks have been reported against a number of cryptographic systems. We will describe here a few such attacks in further detail.

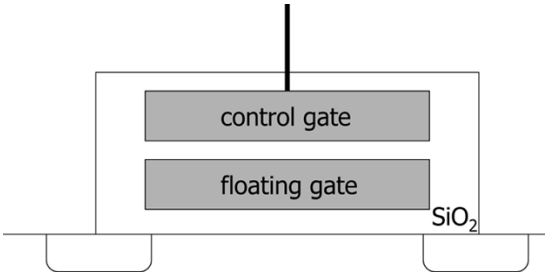

Fig. 9. EEPROM.

A. Glitch Attack on RSA

The GCD attack presented in Section III was implemented by [17] and others. We also refer the reader to [18] and [26] which report clock-glitch attacks against RSA and DES.

B. Glitch Attack on DES

When we can cause an instruction of our choice to fail, then there are several fairly straightforward ways to attack DES. We can remove one of the 8-bit XOR operations that are used to combine the round keys with the inputs to the S-boxes from the last two rounds of the algorithm, and repeat this for each of these key bytes in turn. The erroneous ciphertext outputs that we receive as a result of this will each differ from the genuine ciphertext in the output of usually two, and sometimes three, S-boxes. Using the techniques of differential cryptanalysis, we obtain about five bits of information about the eight key bits that were not XORed as a result of the induced fault. So, for example, six ciphertexts with faulty last rounds should leak about 30 key bits, leaving an easy brute-force search [23]. An even faster attack brutally reduces the number of DES rounds to one or two by corrupting the appropriate loop variable or conditional jump. As a conclusion, unprotected DES can be compromised in a variety of ways with somewhere between one and ten faulty ciphertexts.

C. Glitch Attack on EEPROM

EEPROM stores information as charges in the gate insulator of a MOSFET; charge is stored on the floating gate of a MOS transistor and the control gate is used to program the transistor, as shown in Fig. 9. EEPROM transfers electrons by Fowler–Nordheim tunnelling and program/erase operations are carried out by electrons tunnelling through the thin oxide. Control gate voltage is high for programming while for erasure the control gate is grounded and the drain voltage is raised. To read information from a cell, the cell's static voltage is compared to a reference detection voltage V_{det} (usually $V_{\mathrm{det}} = V_{\mathrm{cc}}/2$). Consequently, if programming is done under the lowest tolerable voltage, a lesser amount of particles will be forced into the cell. Then, if during reading V_{cc} is increased to the highest value tolerated by the circuit V_{det} is artificially boosted and, hence, data will be read as zero regardless its actual value. To attack an n byte key one can simply subject the circuit to $n-1$ power glitches to obtain the encryption of a known plaintext under a vulnerable key of the form

$$00\ 00\ \ldots\ 00\ 00\ \mathtt{XX}\ 00\ 00\ \ldots\ 00\ 00.$$

The attacker will then move the glitch's position to successively scan the entire key. This attack was implemented in the late 1990s.

D. Analogous Laser Attack on a Data Bus

In a specific smart card chip, a laser impact on the data bus during information transfer has the effect of reading the value 255 (0xFF) regardless the transferred information's actual value. The attack described in the previous subsection could hence be directly readapted in a laser laboratory.

E. The Java Sandbox

The Java sandbox is an environment in which applets are run without direct access to the computer's resources, the idea being that an applet need not be trusted as it is incapable of running malicious code. The most common example of Java programs being used is on the Internet, where an applet is downloaded and executed on a PC to achieve a given effect on the webpage being observed. A relatively recent paper [27] describes a fault attack on a PC forcing the Java Virtual Machine to execute arbitrary code. This was done by using a spotlight to heat up the PC's RAM to the point where a fault (in this case a bit flip) occurs. In this case a special applet was loaded into the computer's memory and the RAM heated up to the point where some bits would change their value. The expected fault was that the address of a function a called by the applet would have one bit changed, so that the address called was $a \pm 2^i$, where $0 \leq i < 32$ (the computer's word size). The programmer arranges to have a function present at that address that will return a variable of a type that is not expected by the calling function, for example an integer to a pointer. This can then be used to read/write to arbitrary addresses in the computers memory. One of the possible uses of such a fault would be to change fields in the Java runtime system's security manager to grant the applet illegal rights.

V. Countermeasures

Since the identification of faults as a problem in electronic systems several hardening methods were deployed. These solutions help circuits to avoid, detect, and/or correct faults. Hardware and software countermeasures will be overviewed separately for the sake of clarity.

A. Hardware Countermeasures

Hardware protections are implemented by the chip manufacturer and can be further subdivided into two categories: *active* and *passive* protections.

1) Active Protections:

- *Light detectors* detect changes in the gradient of light.
- *Supply voltage detectors* react to abrupt variations in the applied potential and continuously ascertain that voltage is within the circuit's tolerance thresholds.

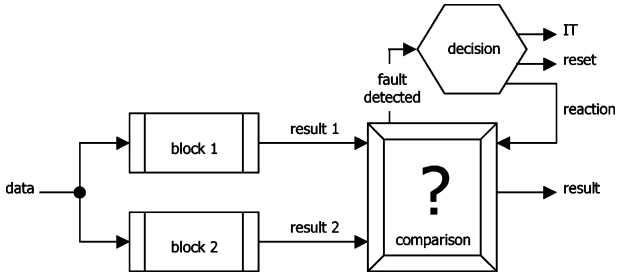

Fig. 10. Simple duplication with comparison.

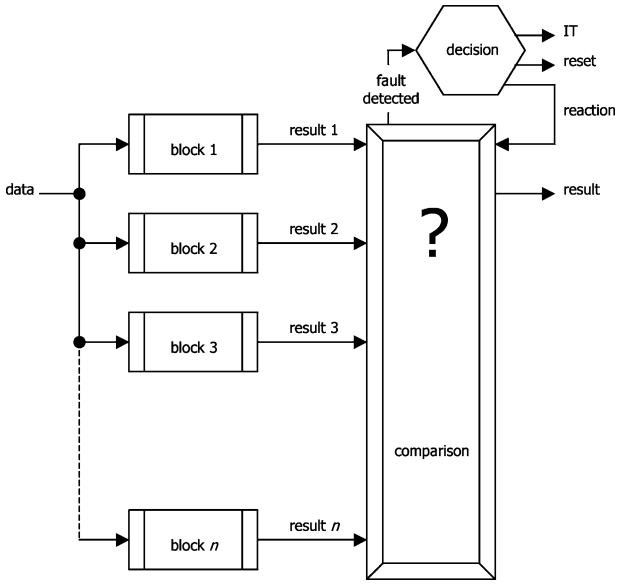

Fig. 11. Multiple duplication with comparison.

- *Frequency detectors* impose an interval of operation outside which the electronic circuit will reset itself.
- *Active shields* are metal meshes that cover the entire chip and has data passing continuously in them. If there is a disconnection or modification of this mesh, the chip will not operate anymore. This is primarily a countermeasure against probing, although it helps protecting against fault injection, as it makes the location of specific blocks in a circuit harder.
- Hardware redundancy:
  1) *Simple duplication with comparison (SDC)*, illustrated in Fig. 10, is the duplication of hardware blocks followed by a test by a comparator. When the two blocks' results do not match, an alert signal is transmitted to a decision block. Two types of reaction can be implemented: a hardware reset or the activation of an interruption that triggers dedicated countermeasures. SDC protects against single focused errors and only permits their detection. A feedback signal is usually triggered to stop all outgoing data flows.
  2) *Multiple duplication with comparison (MDC)*, illustrated in Fig. 11, is where each hardware block is duplicated at least thrice. The comparator detects any mismatch between results and transmits the alert signal to the decision block. As previously, two types of reaction can be implemented, a hardware reset or the activation of an interruption, the difference with SDC being the possibility to correct the fault through a majority vote and correct the outgoing signal.

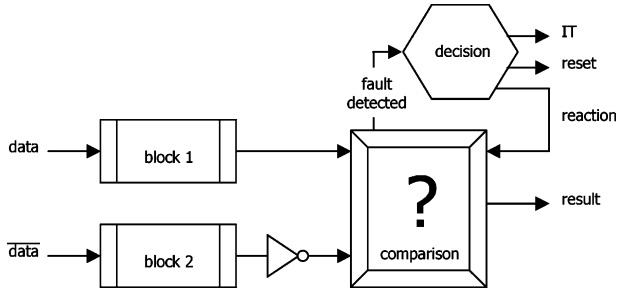

Fig. 12. Simple duplication with complementary redundancy.

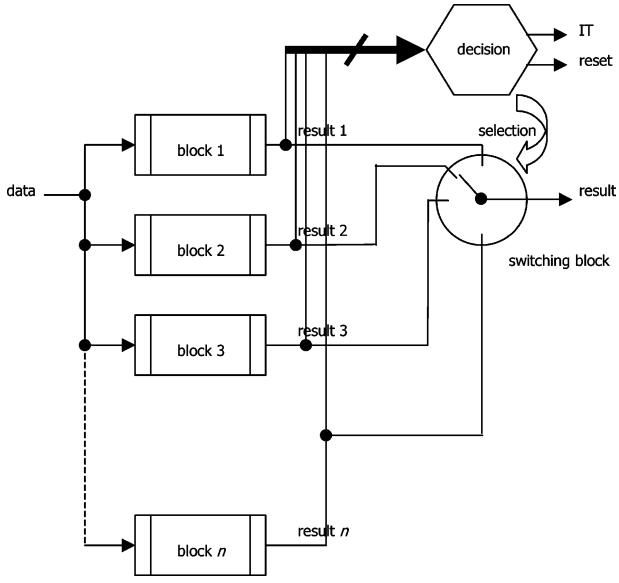

Fig. 13. Dynamic duplication.

  3) *Simple duplication with complementary redundancy (SDCR)*, illustrated in Fig. 12, is based on the same principles as SDC but the two blocks store complemented data. When the result of the two blocks match, the comparison block transmits an alert to the system that triggers a hardware reset or an interrupt. SDCR protects against multiple focused errors, since it is difficult to inject two different errors with complementary effects, but (just as SDC) SDCR only permits error detection.
  4) *Dynamic duplication*, illustrated in Fig. 13, consists of multiple redundancies with a decision module, commanding a data switch upon fault detection. The vote block is a switch, which transmits the correct

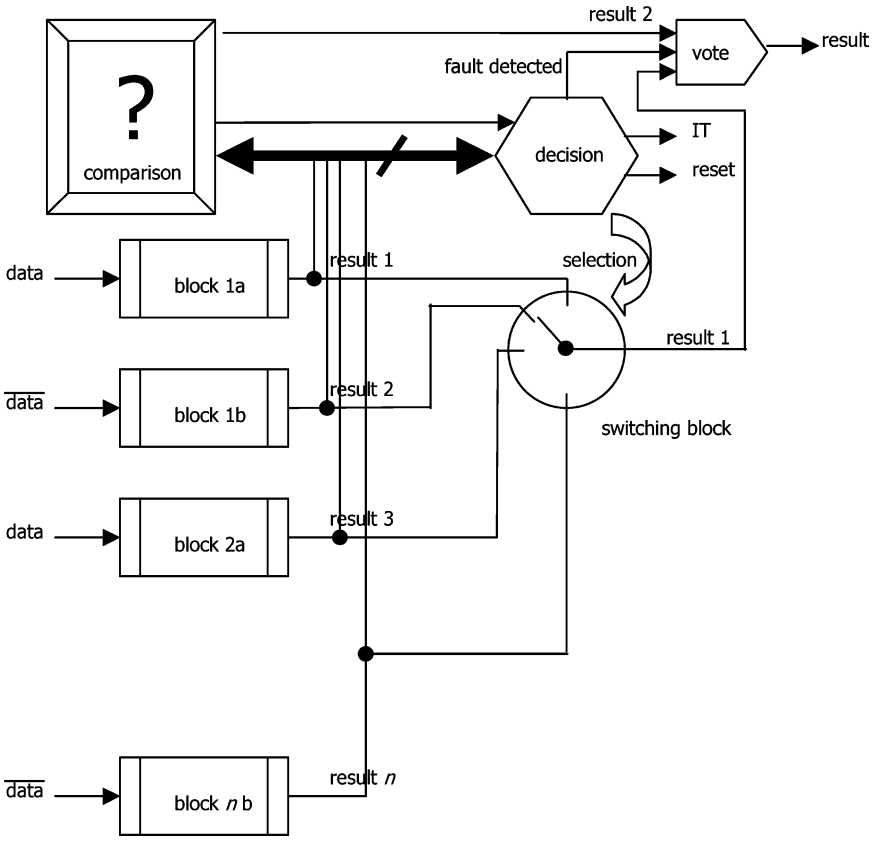

Fig. 14. Hybrid duplication.

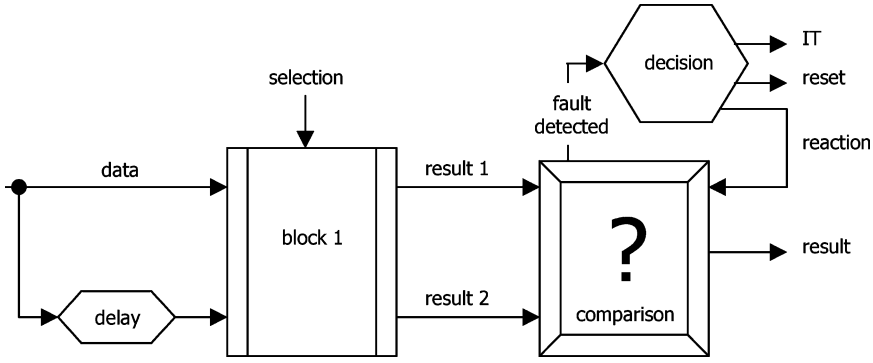

Fig. 15. Simple time redundancy with comparison.

result as instructed by the comparator. Corrupted blocks are disabled and their results discarded. This type of implementation permits detection and subsequent reaction to the detected error [28].

5) *Hybrid duplication*, illustrated in Fig. 14, is a combination of multiple duplications with complementary redundancy and dynamic duplication. This protects against single and multiple focused faults, as it is very difficult to inject multiple faults with complementary effects.

- Protection using time redundancy:
  1) *Simple time redundancy with comparison (STRC)*, illustrated in Fig. 15, consists of processing each operation twice and comparing results [29]. This protects against single and multiple time synchronized errors, but is only capable of detecting faults. Reaction is limited to the discarding of the corrupted results.
  2) *Multiple time redundancy with comparison* (Fig. 16) is based on the principle used by STRC, but the result is processed more than twice. This detects, reacts, and possibly corrects single and multiple faults.
  3) *Recomputing with swapped operands* (Fig. 17) consists of recomputing results with the operands' little endian and big endian bits swapped. The result is reswapped and compared to detect potential faults. This type of protection has the advantage of desynchronizing two different processes and makes fault attacks very difficult. This countermeasure protects against single and multiple time synchronized errors.

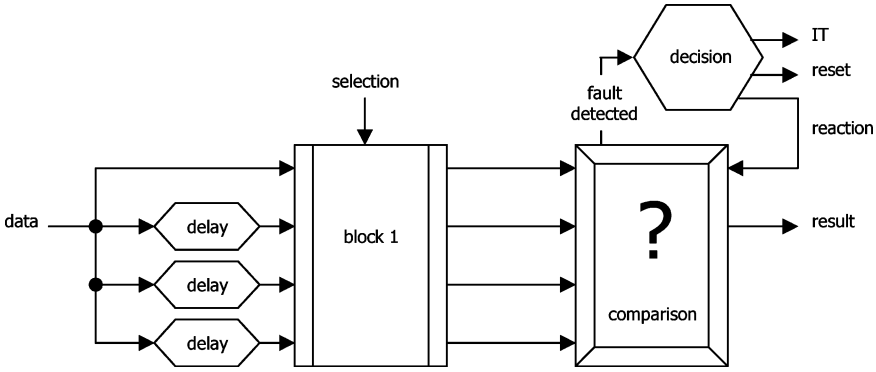

Fig. 16. Multiple time redundancy with comparison.

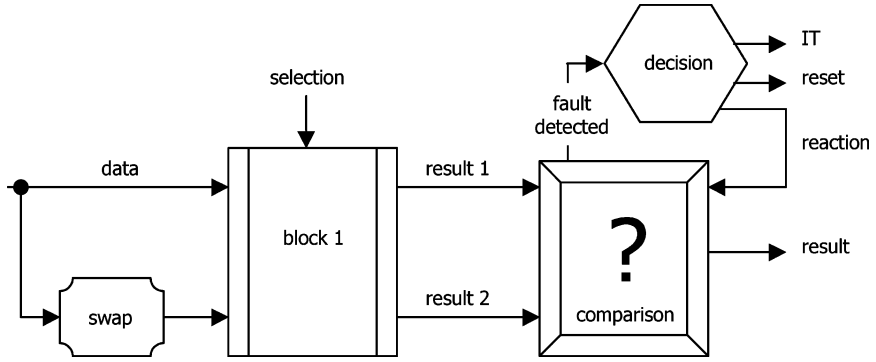

Fig. 17. Recomputing with swapped operand.

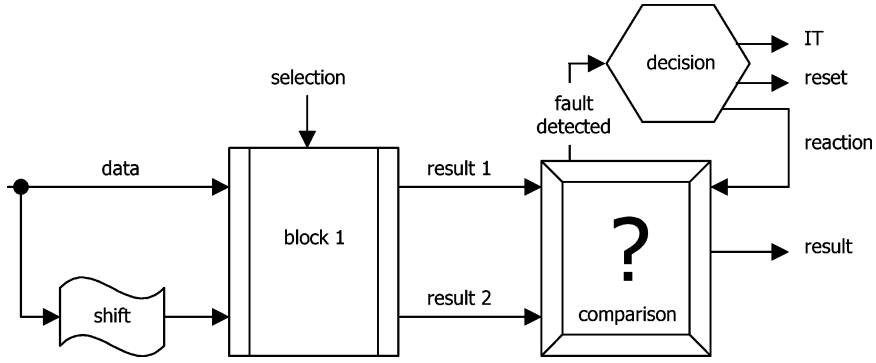

Fig. 18. Recomputing with shifted operand.

4) *Recomputing with shifted operands* (Fig. 18) [30]: operations are recomputed by shifting the operands by a given number of bits. The result is shifted backward and compared to the original one.
5) *Recomputing with duplication with comparison* (Fig. 19) is a combination of time redundancy and hardware redundancy. This protects against single, multiple, and time-synchronized faults, but the time penalty and the increase in block size limit this countermeasure's use.

- *Protection by redundancy mechanisms* such as Hamming codes [31], hardwired checksums, and error correction codes are also used to avoid or detect faults [32], the typical example being checksums attached to each machine word in RAM or EEPROM to ensure integrity.

2) Passive Protections: The second class of hardware protection mechanisms consists of *passive protections* that increase the difficulty of successfully attacking a device. These protections can be self-activated or managed by the device's programmer.

- Mechanisms that introduce *dummy random cycles* during code processing.
- *Bus and memory encryption*. Let h be a hardwired keyed permutation and f a simple hardwired block-cipher. Upon power-on, the chip generates an ephemeral key k. When the microprocessor wishes to write the value m at RAM address i, the system stores $v = f_k(m, i)$ at address $h_k(i)$. When the microprocessor requires the contents of address i, the system recomputes $h_k(i)$, fetches v from address $h_k(i)$, decrypts $m = f_k^{-1}(v, i)$, and hands m to the microprocessor. This makes laser or

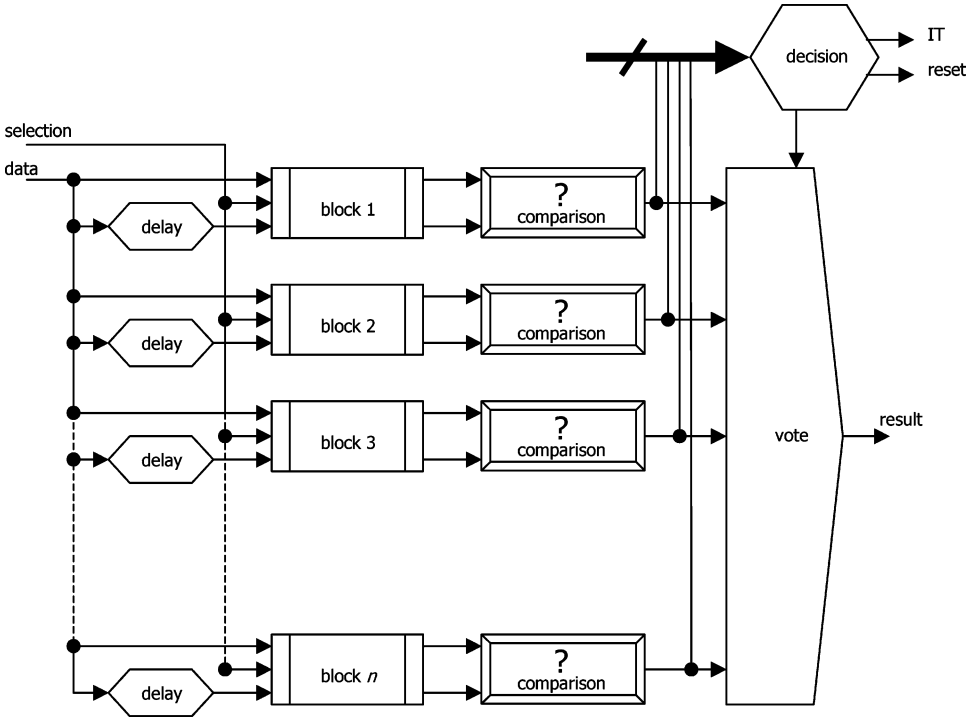

Fig. 19. Recomputing with duplication with comparison.

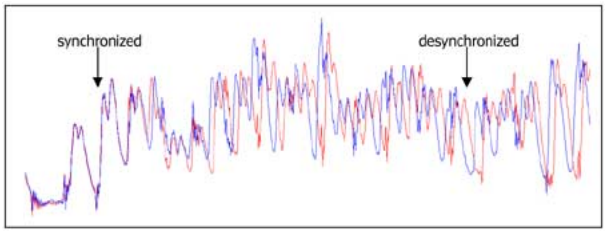

Fig. 20. Unstable internal frequency generation reflected in power consumption.

glitch targeting of a specific memory cell useless as successive computations with *identical* data use *different* memory cells.

- *Passive shield*: a full metal layer covers some sensitive chip parts, which makes light or electromagnetic beam attacks more difficult as the shield needs to be removed before the attack can proceed. This also allows to contain information leakage through electromagnetic radiations (i.e., thwart some side-channel attacks).
- *Unstable internal frequency generators* protect against attacks that need to be synchronized with a certain event, as events occur at different moments in different executions. An example of such a deterrent is depicted in Fig. 20.

B. Software Countermeasures

Software countermeasures are implemented when hardware countermeasures are insufficient or as cautious protection against future attack techniques that might defeat present-generation hardware countermeasures. The advantage of software countermeasures is that they do not increase the hardware block size, although they do impact the protected functions' execution time.

- *Checksums* can be implemented in software. This is often complementary to hardware checksums, as software CRCs can be applied to buffers of data (sometimes fragmented over various physical addresses) rather than machine words.
- *Execution randomization*: If the order in which operations in an algorithm are executed is randomized, it becomes difficult to predict what the machine is doing at any given cycle. For most fault attacks, this countermeasure will only slow down a determined adversary, as eventually a fault will hit the desired instruction. This will, however, thwart attacks that require faults in specific places or in a specific order, such as the transferring of secret data attack described previously.
- *Variable redundancy* is nothing but SDC in software.
- *Execution redundancy* is the repeating of algorithms and comparing the results to verify that the correct result is generated. As SDCR, redundancy is more secure if the second calculation is different than the first (for example, its inverse[3]) so that two identical faults cannot be used at different times.
- *Ratification counters and baits*: baits are small ($<$ 10 byte) code fragments that perform an operation and test its result. A typical bait writes, reads and compares data, performs XORs, additions, multiplications, and other operations whose results can be easily checked. When a bait detects an error, it increments an NVM counter, and when this counter exceeds a tolerance limit (usually three), the card ceased to function.

In theory, all data redundancy method used in hardware can be implemented in software. The problem then becomes execution time rather than block size. As some of the pro-

[3]Encrypt-decrypt, sign-verify, etc.

posed hardware designs become extremely time consuming when imitated by software.

VI. Conclusion

Various methods for creating faults were presented. Practical applications of these attacks were presented. These applications included attacks on keys and symmetric and asymmetric cryptosystems. Finally, hardware and software countermeasures were overviewed. Unfortunately, these countermeasures never come for free and impact the cost of the system being developed. Also, the resulting system will be slower and may feature an increased block size. There will always be a tradeoff between cost, efficiency, and security, and it will be a judgment call by designers, developers, and users to choose which of these requirements best suit their needs. There is still much work to be done in this area with the ultimate goal being an optimal balance between security, efficiency and cost.

The attacks described were implemented on chips which did not contain hardware countermeasures specifically designed to prevent fault attacks. They were part of an effort to develop suitable countermeasures knowing that the hardware was vulnerable to fault injection. This is due to the delay between such countermeasures being implemented and actually appearing as silicon that can be used for a product. There are currently European projects underway, such as [36], which aim to characterize the effects of fault injection and evaluate the efficieny of current hardware countermeasures.

References

[1] T. May and M. Woods, "A new physical mechanism for soft errors in dynamic memories," in *Proc. 16th Int. Reliability Physics Symp.* Apr. 1978.

[2] J. Ziegler, "Effect of cosmic rays on computer memories," *Science,* vol. 206, pp. 776–788, 1979.

[3] D. H. Habing, "The use of lasers to simulate radiation-induced transients in semiconductor devices and circuits," *IEEE Trans. Nucl. Sci.,* vol. 39, pp. 1647–1653, 1992.

[4] D. Boneh, R. DeMillo, and R. Lipton, "On the importance of checking cryptographic protocols for faults," *J. Cryptol.* vol. 14, no. 2, pp. 101–119, 2001.

[5] R. Anderson and S. Skoroboatov, "Optical fault induction attacks," in *Cryptographic Hardware and Embedded Systems (CHES 2002)*. Heidelberg, Germany: Springer-Verlag, 2002, vol. 2523, Lecture Notes on Computer Science, pp. 2–12.

[6] P. Fouillat, Contribution à l'étude de l'interaction entre un faisceau laser et un milieu semiconducteur Applications à l'étude du Latchup et à l'analyze d'états logiques dans les circuits intégrés en technologie CMOS, Thèse de doctorat de l'université Bordeaux I, 1990.

[7] V. Pouget, "Simulation expérimentale par impulsions laser ultracourtes des effets des radiations ionisantes sur les circuits intégrés," Thèse de doctorat de l'Université de Bordeaux I, Bordeaux, France, 2000.

[8] T. J. O'Gorman, "The effect of cosmic rays on soft error rate of a DRAM at ground level," *IEEE Trans. Electron Devices,* vol. 41, no. 4, pp. 553–557, Apr. 1994.

[9] J. C. Pickel and J. T. Blandford, Jr., "Cosmic ray induced errors in MOS memory circuits," *IEEE Trans. Nucl. Sci.,* vol. NS-25, pp. 1166–1171, 1978.

[10] R. Koga, M. D. Looper, S. D. Pinkerton, W. J. Stapor, and P. T. McDonald, "Low dose rate proton irradiation of quartz crystal resonators," *IEEE Trans. Nucl. Sci.,* vol. 43, no. 6, pp. 3174–3181, Dec. 1996.

[11] S. Kuboyama, S. Matsuda, T. Kanno, and T. Ishii, "Mechanism for single-event burnout of power MOSFETs and its characterization technique," *IEEE Trans. Nucl. Sci.,* vol. 39, no. 6, pp. 1698–1703, Dec. 1992.

[12] E. G. Stassinopoulos, G. J. Brucker, P. Calvel, A. Baiget, C. Peyrotte, and R. Gaillard, "Charge generation by heavy ions in power MOSFETs, burnout space predictions and dynamic SEB sensitivity," *IEEE Trans. Nucl. Sci.,* vol. 39, no. 6, pp. 1704–1711, Dec. 1992.

[13] R. Koga and W. A. Kolasinski, "Heavy ion induced snapback in CMOS devices," *IEEE Trans. Nucl. Sci.,* vol. 36, no. 6, pp. 2367–2374, Dec. 1989.

[14] L. Adams, E. J. Daly, and R. Harboe-Sorensen *et al.*, "A verified proton induced latchup in space," *IEEE Trans. Nucl. Sci.,* vol. 39, no. 6, pp. 1804–1808, Dec. 1992.

[15] P. Cazenave, P. Fouillat, X. Montagner, H. Barnaby, R. D. Schrimpf, L. Bonora, J. P. David, A. Touboul, M.-C. Calvet, and P. Calvel, "Total dose effects on gate controlled lateral PNP bipolar junction transistors," *IEEE Trans. Nucl. Sci.,* vol. 45, no. 6, pp. 2577–2583, Dec. 1998.

[16] B. G. Rax, C. I. Lee, A. H. Johnston, and C. E. Barnes, "Total dose and proton damage in optocouplers," *IEEE Trans. Nucl. Sci.,* vol. 43, no. 6, pp. 3167–3173, Dec. 1996.

[17] C. Aumüller, P. Bier, P. Hofreiter, W. Fischer, and J.-P. Seifert, "Fault attacks on RSA with CRT: Concrete results and practical countermeasures" , Cryptology ePrint Archive: Rep. 2002/073 [Online]. Available: http://www.iacr.org

[18] F. Bao, R. H. Deng, Y. Han, A. Jeng, A. D. Narasimhalu, and T. Ngair, "Breaking public key cryptosystems on tamper resistant devices in the presence of transient faults," in *Security Protocols*. Heidelberg, Germany: Springer-Verlag, 1997, vol. 1361, Lecture Notes in Computer Science, pp. 115–124.

[19] E. Biham and A. Shamir, "Differential fault analysis of secret key cryptosystems," in *Advances in Cryptology—CRYPTO '97*. Heidelberg, Germany: Springer-Verlag, 1997, vol. 1294, Lecture Notes in Computer Science, pp. 513–525.

[20] G. Piret and J. J. Quisquater, "A differential fault attack technique against SPN structure, with application to the AES and KHAZAD," in *Cryptographic Hardware and Embedded Systems (CHES 2003)*. Heidelberg, Germany: Springer-Verlag, 2003, vol. 2779, Lecture Notes in Computer Science, pp. 77–88.

[21] C. Giraud, "DFA on AES" , Cryptology ePrint Archive: Rep. 2003/008 [Online]. Available: http://www.iacr.org

[22] P. Dusart, G. Letourneux, and O. Vivolo, "Differential fault analysis on A.E.S." , Cryptology ePrint Archive: Rep. 2003/010 [Online]. Available: http://www.iacr.org

[23] R. Anderson and M. Kuhn, "Low cost attacks on tamper resistant devices," in *Security Protocols*. Heidelberg, Germany: Springer-Verlag, 1997, vol. 1361, Lecture Notes in Computer Science, pp. 125–136.

[24] O. Kommerling and M. Kuhn, "Design principles for tamper resistant smartcard processors," in *Proc. USENIX Workshop on Smartcard Technology* 1999, pp. 9–20.

[25] S. Moore, R. Anderson, and M. Kuhn, "Improving smartcard security using self-timed circuit technology," in *IEEE Int. Symp. Ansychnronous Circuits and Systems* 2002, pp. 120–126.

[26] O. Grabbe, "Smartcards and private currencies" [Online]. Available: http://www.aci/net/kalliste/smartcards.htm

[27] S. Govindavajhala and A. W. Appel, "Using memory errors to attack a virtual machine," in *Proc. 2003 IEEE Symp. Security and Privacy* pp. 154–165.

[28] J. Losq, "Influence of fault detection and switching mechanisms on reliability of stand-by systems," in *Digest 5th Int. Symp Fault-Tolerant Computing* 1975, pp. 81–86.

[29] L. Anghel and M. Nicolaidis, "Cost reduction and evaluation of a temporary faults detecting technique," in *Proc. Design, Automation and Test in Europe (DATE '00)* 2000, pp. 591–597.

[30] J. H. Patel and L. Y. Fung, "Concurrent error detection in ALUs by recomputing with shifted operands," *IEEE Trans. Comput.,* vol. C-31, pp. 589–595, 1982.

[31] F. Lima, E. Costa, L. Carro, M. Lubaszewski, R. Reis, S. Rezgui, and R. Velazco, "Designing and testing a radiation hardened 8051-like micro-controller," presented at the 3rd Military and Aerospace Applications of Programmable Devices and Technologies Int. Conf., Laurel, MD, 2000.

[32] M. Pflanz, K. Walther, C. Galke, and H. T. Vierhaus, "On-line detection and correction in storage elements with cross-parity check," in *Proc. 8th IEEE Int. On-Line Testing Workshop (IOLTW'02)* pp. 69–73.
[33] L. Hemme, "A differential fault attack against early rounds of (Triple-)DES," in *Cryptographic Hardware and Embedded Systems (CHES 2004)*. Heidelberg, Germany: Springer-Verlag, 2004, vol. 3156, Lecture Notes in Computer Science, pp. 254–267.
[34] M. Joye, A. Lenstra, and J.-J. Quisquater, "Chinese remaindering cryptosystems in the presence of faults," *J. Cryptol.,* vol. 12, no. 4, pp. 241–245, 1999.
[35] M. Joye, J.-J. Quisquater, F. Bao, and R. H. Deng, "RSA-type signatures in the presence of transient faults," in *Cryptography and Coding*, M. Darnell, Ed. Heidelberg, Germany: Springer-Verlag, 1997, vol. 1355, Lecture Notes in Computer Science, pp. 155–160.
[36] Duracell, "Durcissement aux Attaques par fautes de Cellules pour circuits sécuritaires," [Online]. Available: http://www.telecom.gouv.fr/rnrt/rnrt/projets/res_02_9.htm

David Naccache is currently a member of the Information Security Group at Royal Holloway, University of London, Egham, U.K. He was previously with Philips where he designed, broke, and coded various smart card security solutions and Thomson Consumer Electronics where he worked on Videocrypt's Pay-TV Module. He has published more than 50 papers in cryptography and security and holds 60 patents in the field. His primary areas of interest are whatever can be attacked or protected and, in particular, public-key cryptography, side-channel attacks and mobile code security. He has served on more than 30 program committees of specialized conferences such as ACM CSS, Eurocrypt, or Crypto and cosupervised a dozen of Ph.D.s in the area in France and abroad. A few years ago he cosigned an attack that caused the withdrawal of the ISO 9796-1 standard and the modification of ISO 9796-2 as well as an attack against PKCS#1 v1.5 encryption which was upgraded to 2.0. He spends much of his time coding in assembly, C, Java, doing maths, and training his subordinates.

Under his management, Dr. Naccache's security group won the RSA Security Industry Award, was awarded over 20 Common Criteria and FIPS certificates, and signed more than 250 security publications in refereed journals and specialized conferences.

Hagai Bar-El is an information security analyst specializing in security analysis of cryptographic applications and secure system design. Previously, he was involved in the security design and evaluation of applications and systems from his own office, hbarel.com. He currently serves as the information security analyst at Discretix Technologies Ltd., Rehovot, Israel, a leading company that develops hardware and software-based embedded security solutions.

Michael Tunstall is currently working toward the Ph.D. degree at Royal Holloway, University of London, Egham, U.K., supervised by C. Mitchell.

His current research interests are based on applied cryptography, specifically side-channel attacks, fault attacks, and developing efficient countermeasures.

Hamid Choukri is a smart card security expert and security project leader. His current research interests focus on the development of secure operating systems for smartcards, fault attacks and dedicated countermeasures.

Claire Whelan is currently working toward the Ph.D. degree at Dublin City University (DCU), Dublin, Ireland.

Sh is currently a member of the Dependable Systems Research Group at DCU. She is jointly supervised by Dr. M. Scott at DCU and Dr. D. Naccache. Her current research interests are in the field of side-channel attacks and specifically power analysis and fault analysis of smart card systems.