The EM Side–Channel(s)

Dakshi Agrawal, Bruce Archambeault, Josyula R. Rao, and Pankaj Rohatgi

IBM T.J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598
{agrawal,barch,jrrao,rohatgi}@us.ibm.com

Abstract. We present results of a systematic investigation of leakage of compromising information via electromagnetic (EM) emanations from CMOS devices. These emanations are shown to consist of a multiplicity of signals, each leaking somewhat different information about the underlying computation. We show that not only can EM emanations be used to attack cryptographic devices where the power side–channel is unavailable, they can even be used to break power analysis countermeasures.

1 Introduction

Side–channel cryptanalysis has been used successfully to attack many cryptographic implementations [7,8]. Most published literature on side–channels deals with attacks based on timing or power. With the recent declassification of portions of the TEMPEST documents [5], and other recent results [9,6], an awareness of the potential of the EM side–channel is developing. However, some basic questions remain unanswered. For instance, what are the causes and types of EM emanations? How does information leaked via EM emanations compare with leakages from other side–channels? What new devices and implementations are vulnerable to EM side–channel attacks? Can the EM side–channel overcome countermeasures designed to provide protection against other side–channel attacks? With questions such as these in mind, we conducted a systematic investigation of EM side–channel leakage from CMOS devices. In this paper, we address each of these basic questions.

In Section 2, we discuss the causes and types of various EM signals and describe the equipment required to capture and extract these signals. In addition to the direct emanations, EM signals consist of several compromising signals which are unintentional and are found in unexpected places. For instance, researchers have thus far missed the faint, but far more compromising amplitude modulated EM signals present even in the power line.

Section 3 presents experimental results illustrating various types of emanations and Section 4 provides a qualitative comparison of information leakages from EM and power. These results are very instructive. One crucial observation is that even a *single* EM sensor can easily pick up multiple compromising signals of different types, strengths and information content. Moreover, significant amount of compromising information is to be found in very low energy signals.

B.S. Kaliski Jr. et al. (Eds.): CHES 2002, LNCS 2523, pp. 29–45, 2003.

It is therefore critical that signals be separated early in the acquisition process to avoid loss of these low energy signals due to precision limits of signal capturing equipment. A very effective way to achieve such a separation is to exploit unintentionally modulated carriers at higher frequencies where there is less interference and noise rather than focusing on direct emanations in the baseband where the large amount of interference and noise may require techniques such as chip decapsulation and use of carefully positioned micro–antenna [9,6].

Using EM, we launched attacks such as simple and differential electromagnetic attacks (SEMA and DEMA [9]) on straight–forward implementations of DES, RSA and COMP128 on smart cards, cryptographic tokens and SSL accelerators. While the EM side–channel remains the most viable avenue for attacking cryptographic devices where the power side–channel is unavailable, an important question is whether the EM side–channel provides any other advantage when the power side–channel is available. In Section 5, we answer this in the affirmative. We outline an approach that breaks some fielded systems with power analysis countermeasures. The approach is based on the observation that most devices have classes of "bad instructions" whose leakage in some EM side–channel far exceeds the corresponding leakage in the power side–channel and works against two major classes of power analysis countermeasures [8,2,4]. We illustrate this approach by attacking a test implementation[1] of the secret–sharing countermeasures of [2,4]. This approach works in many cases even when the code is unknown.

Despite their effectiveness, our low–cost attacks provide only a glimpse of what is possible: combining leakages from multiple EM signals could yield substantially better attacks. Furthermore, developing countermeasures requires a methodology to assess the net information leakage from all the EM signals realistically available to an adversary. Our work on these aspects of the EM side–channel(s) is described in more detail in [1].

2 EM Emanations and Acquisition

This section describes the origin and types of various compromising EM signals that we have observed[2] and the equipment and techniques to extract them.

2.1 Origin of EM Emanations

EM emanations arise as a consequence of current flows within the control, I/O, data processing or other parts of a device. These flows and resulting emanations may be *intentional* or *unintentional*. Each current carrying component of the device not only produces its own emanations based on its physical and electrical characteristics but also affects the emanations from other components due to coupling and circuit geometry.

[1] To avoid disclosing weaknesses of commercially deployed systems.

[2] While there is an obvious overlap with the declassified TEMPEST documents (NACSIM 5000) [5], we only describe what we have verified in our investigations.

An attacker is typically interested in emanations resulting from data processing operations. In CMOS devices, ideally, current only flows when there is a change in the logic state of a device and this logic state change is controlled by a "square–wave" shaped clock. These currents result in compromising emanations, sometimes, in unintended ways. Such emanations carry information about the currents flowing and hence the events occurring during each clock cycle. Since each active component of the device produces and induces various types of emanations, these emanations provide multiple views of events unfolding within the device at each clock cycle. This is in sharp contrast to the power side–channel where only a single aggregated view of net current inflow is available thus, explaining why the EM side–channel(s) are much more powerful.

2.2 Types of EM Emanations

There are two broad categories of EM emanations:

1. Direct Emanations: These result from *intentional* current flows. Many of these consist of short bursts of current with sharp rising edges resulting in emanations observable over a wide frequency band. Often, components at higher frequencies are more useful to the attacker due to noise and interference prevalent in the lower bands. In complex circuits, isolating direct emanations may require use of tiny field probes positioned very close to the signal source and/or special filters to minimize interference: getting good results may require decapsulating the chip packaging [6,9].

2. Unintentional Emanations: Increased miniaturization and complexity of modern CMOS devices results in electrical and electromagnetic coupling between components in close proximity. Small couplings, typically ignored by circuit designers, provide a rich source of compromising emanations. These emanations manifest themselves as *modulations* of carrier signals generated, present or "introduced" within the device. One strong source of carrier signals is the ubiquitous harmonic–rich "square–wave" clock signal[3]. Other sources include communication related signals. Ways in which modulation occurs include:

a. Amplitude Modulation: Non–linear coupling between a carrier signal and a data signal results in the generation and emanation of an Amplitude Modulated (AM) signal. The data signal can be extracted via AM demodulation using a receiver tuned to the carrier frequency.

b. Angle Modulation: Coupling of circuits also results in Angle Modulated Signals (FM or Phase modulation). For instance, while signal generation circuits should ideally be decoupled from data processing circuits, this is rarely achieved in practice. For example, if these circuits draw upon a limited energy source the generated signal will often be angle modulated by the data signal. The data signal is recoverable by angle demodulation of the generated signal.

[3] Theoretically a symmetric, square clock signal consists of the fundamental frequency and all the odd harmonics with progressively diminishing strengths. In practice, the clock signal is always imperfect.

Exploiting unintentional emanations can be much more effective that trying to work with direct emanations. Some modulated carriers have substantially better propagation than direct emanations. This enables attacks to be carried out without resorting to invasive techniques and even attacks that can be performed at a distance. None of the attacks described in this paper require any invasive techniques or fine grained positioning of probes. Secondly, careful field probe positioning cannot separate two sources of direct emanations in close proximity, while such sources may be easily separable due to their differing interaction with the carriers present in the vicinity.

2.3 Propagation and Capture of EM Signals

EM signals propagate via radiation and conduction, often by a complex combination of both. Thus two classes of sensors are required to capture the signals that emerge. The most effective method for capturing radiated signals is to place near field probes as close as possible to the device or at least in the near field, i.e., no more that a wavelength away. Some of these emanations can also be captured at much larger distances using standard antennas. In our experiments, the most effective near field probes are those made of a small plate of a highly conducting metal like silver or copper attached to a coaxial cable. In the far field, we used biconical and log–periodic wide–band antennas as well as hand–crafted narrow–band, high gain Yagi antennas. Conductive emanations consist of faint currents found on all conductive surfaces or lines attached to the device possibly riding on top of stronger, intentional currents within the same conductors. Capturing these emanations requires current probes. The quality of the received signal improves if the equipment is shielded from interfering EM emanations in the band of interest, though the shielding does not have to be elaborate.

The emanations received by the sensor have to be further processed to extract compromising information. For direct emanations, filters may suffice. For unintentional emanations, which manifest themselves as modulations of carrier signals, a receiver/demodulator is required. For experimental work, a wide bandwidth, wideband tunable receiver such as the R–1550 Receiver from Dynamic Sciences and the 8617 Receiver from Watkins–Johnson is convenient. A cheaper alternative is to use wide–band radio receivers such as the ICOM 7000/8500 which have intermediate frequency outputs and to then perform the demodulating functionality in software. An even cheaper approach is to construct the receiver using commonly available low noise electronic components. At some stage, the signal has to be digitized using digital scope/sampling card as done for power analysis attacks. Equipment such as spectrum analyzers are also useful for quickly identifying carriers and potentially useful emanations. A useful rule–of–thumb is to expect strong carriers at odd harmonics of the clock.

3 Experimental Results

We describe experiments that illustrate the various types and nature of EM emanations.

Experiment 1: Direct Near–Field Emanations: We programmed a recently deployed smart card, called smartcard A (to protect vendor identity[4]), to enter a 13 cycle infinite loop using the externally supplied 3.68MHz clock. A near–field probe (a small metal plate attached to a co–axial cable) was placed near the chip at the back of smart card. After wide–band amplification, 500K sample points (representing approx 284 iterations of the loop) were captured using an 8–bit, 500MHz digital scope. In the time domain, the baseband *direct emanations* signal (band centered at 0MHz), looked like a differentiated form of the external clock and provided *no* visual indication of a loop execution. In the frequency domain, the signal received by the probe consists of the signal of interest, i.e., a periodic signal corresponding to a loop iteration at 283KHz (3.68MHz/13), other signals from the chip and its vicinity such as the clock (periodic with freq 3.68MHz) and aperiodic noise. Capturing the received signal with a limited resolution scope further introduces quantization noise. Figure 1 plots the magnitude[5] of the FFT of the captured baseband signal against the frequency in KHz over the 0–200 MHz band. The large spikes below 100 MHz are the high energy harmonics of the clock signal and tiny spikes sprinkled between them are other types of direct and unintentional emanations which are of interest. Very little signal is noticeable above 125 MHz because these signals have lower strengths and have been overwhelmed by quantization noise. In the linear scale used in Figure 1, the loop execution is not apparent. On a log (base 10) scale, zooming into the region from 0 to 20MHz, as shown in Figure 2, the signal of interest at 283KHz and its harmonics can be seen interspersed between the clock signal and its harmonics. Note that the use of a large time window, i.e., 284 iterations of the loop, helps in detecting this periodic signal since aperiodic noise from the chipcard, environment and quantization gets reduced due to averaging. Since the direct emanations are at least an order of magnitude smaller than interfering signals, exploiting them in the presence of quantization noise, is quite challenging and has been addressed by [6,9]. Our approach focuses on the much easier task of exploiting *unintentional emanations*.

Experiment 2: Unintentional Near–Field AM Emanations: We use the same setup as in Experiment 1, but with the output of the probe connected to an AM receiver, tuned to the 41'st clock harmonic at 150.88 MHz with a band of 50MHz. The demodulated output was sampled with a 12–bit 100MHz scope[6] and 100K sample points representing approximately 284 loop iterations were collected. Figure 3 plots the magnitude of the FFT of this signal against the frequency in KHz. Notice that even in this *linear* scale plot, the signal of interest, i.e., the 283KHz signal corresponding to the loop and its harmonics, is clearly

[4] Smartcard A is 6805–based, uses 0.6 micron triple metal technology with an optional variable internal clock as one defense against DPA.

[5] In all figures, signal magnitudes should be treated as relative quantities: we don't track the absolute values as the signals typically undergo analog processing before being captured by an 8/12–bit scope. The scope sensitivity is set so that the 8/12–bit dynamic range is fully utilized.

[6] The lower bandwidth allows the use of a lower sampling rate with higher precision.

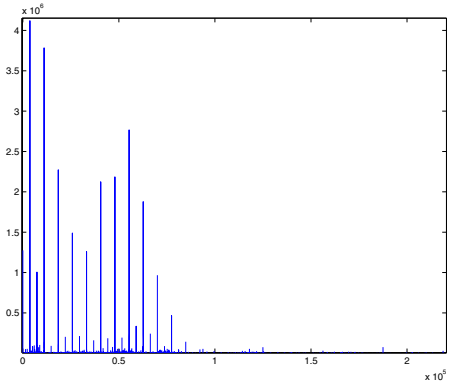

Fig. 1. FFT of baseband signal from Experiment 1 with Smartcard A

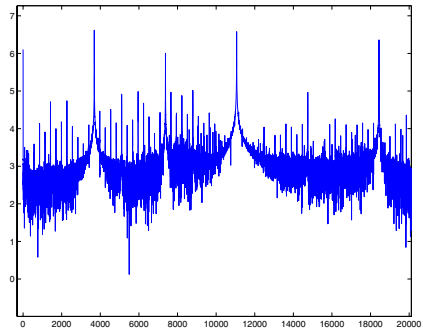

Fig. 2. Log of FFT in the region 0–20MHz from Experiment 1 with Smartcard A

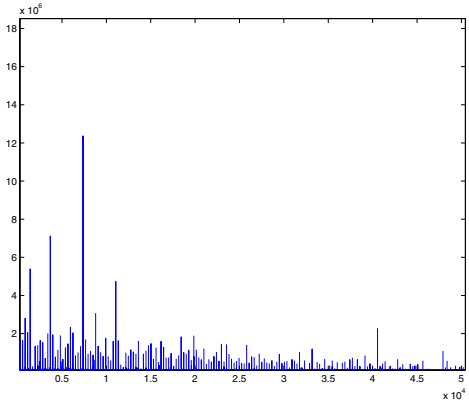

Fig. 3. FFT of demodulated signal (150.88 MHz carrier, 50Mz band) in Experiment 2 with Smartcard A

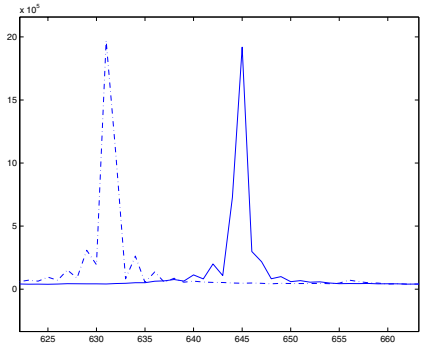

Fig. 4. Two FFTs showing loop frequency differences (LSB 0 and 1) for smartcard A

visible among the clock harmonics. The loop structure is also clearly visible in the time domain. Notice that these greatly improved results were obtained using the same sensor setting as in Experiment 1, and with the same number of loop iterations. Note that we are also operating in a part of the spectrum which showed hardly any signal according to Figure 1; since the signals in this band were overwhelmed by the quantization noise in that experiment.

Experiment 3: Unintentional Near/Far–Field Angle Modulated Emanations: Next we enabled the variable internal clock DPA protection mechanism in Smartcard A and kept everything else the same. One of the instructions in the 13–cycle loop was to load a user supplied byte B from RAM to accumulator. We experimented with different values of the byte B and made the following surprising observation: the average frequency of the 13–byte loop was dependent on the least significant bit (LSB) of B but not on other bits. This is shown in Figure 4, where the magnitude of FFT of the EM output for two different cases is plotted against the frequency in KHz. The first case (shown by a broken line) shows the loop frequency with the $LSB(B) = 1$ and in the second case (shown by a solid line) the loop frequency when the $LSB(B) = 0$. In the first case, the loop runs slower. This is due to coupling between the LSB and the circuitry generating the internal clock. Although the clock frequency itself varies frequently, when there is a 1 bit on the LSB line, the intrinsic variation is biased towards slowing down the clock for a couple of subsequent cycles. We speculate that this is due to the clock circuitry drawing energy from the same source as some other circuitry affected by the LSB. Thus, angle demodulation, e.g., FM demodulation, turns out to be a good avenue for attacking smartcard A using LSB based hypothesis. This effectively transforms a countermeasure into a liability! Another advantage of such an attack is that it can be performed at a distance in the far field since the clock signal is quite strong.

Experiment 4: Unintentional Far–Field AM Emanations: We examined emanations from an Intel–based server containing a commercial, PCI bus based SSL accelerator S[7]. We programmed the server to repeatedly invoke S to perform

[7] S is rated to perform 200, 1024-bit CRT based RSA private key ops/s.

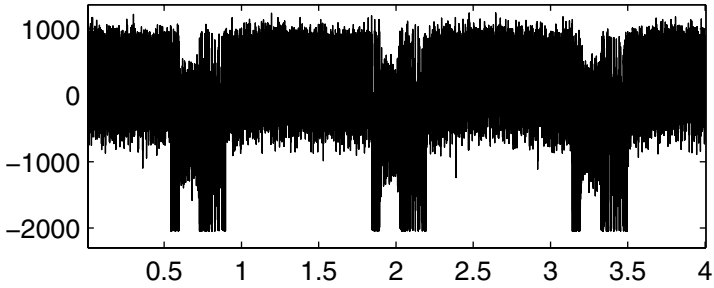

Fig. 5. EM Signal from SSL Accelerator S

a 2048 bit exponentiation with a single–nibble exponent. Several AM modulated carriers (at multiples of the 33MHz PCI clock) containing compromising information propagated to distances upto forty feet. Figure 5 plots a signal (amplitude vs. time in ms) captured by a log–periodic antenna 15 feet away using the 299MHz carrier and 1MHz bandwidth. Three invocations of S are clearly visible as bands where the amplitude goes below -1000. At this resolution, the macro structures of the exponentiation are already visible. At higher resolutions, there is enough information to enable the new class of template attacks [3].

Experiment 5: Conductive Emanations: Conductive emanations appear at unexpected places and are easy to overlook. In fact, if researchers experimenting with power analysis attacks re–analyze the raw signals from their current probes, they will discover that apart from the relatively low frequency, high amplitude power consumption signal, there are faint higher frequency AM modulated carriers representing conductive EM emanations from the device, since the power line is also a conductor. Figure 6 plots one such EM signal (amplitude vs time in 10ns units) extracted from the power line by AM demodulating one such carrier while a smart card (which we call smartcard B[8]) executes 3 rounds of DES. These rounds are clearly visible in the signal.

4 Information Leakage across EM Spectrum

In this section, we provide experimental evidence to reinforce a central theme of this paper, i.e., the output of even a single wide–band EM sensor logically consists of multiple EM signals each carrying qualitatively different compromising information and in some cases, EM leakages can be substantially superior to the power consumption signal.

While the presence of certain types of EM signals (e.g., angle modulated carriers, intermodulated carriers etc) are device dependent, our experiments show that invariably, AM carriers at clock harmonics are a rich and easily accessible source of compromising signals. For smart cards, since the fundamental

[8] Smartcard B is a 6805–based, 0.7micron, double metal technology card with inbuilt noise generators.

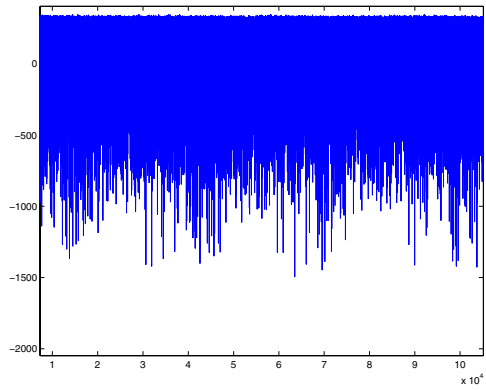

Fig. 6. EM Signal on Power Line for 3 rounds of DES on smartcard B

frequency is low, the intermediate harmonics are usually the best. Lower harmonics suffer from excessive noise and interference and higher harmonics tend to have extremely low signal strength[9].

We now examine the leakage of information from four types of signals obtained from smartcard B when it performed DES in software. No power analysis countermeasures, except for the internal noise generators, were enabled on the card. The smart card ran on the 3.68MHz external clock. Three of these signals were obtained by AM demodulating the output of a near field probe placed as in Experiment 1, at three different carrier frequencies (50MHz bands around 188MHz, 224.5MHz and 262MHz). The fourth signal was the power consumption signal. All signals were collected by a 12–bit, 100MHz digital scope.

It is well known that plotting the results of a differential side channel attack launched against a bit value used in a computation is a good way to assess the leakage of the bit [8]. This is because the plot is essentially the difference between the average of all signals in which the bit is 1 and the average of all signals in which the bit is 0, plotted against time. At points in the computation where this bit is not involved or where the bit is involved but information about it does not leak in the side–channel, the value of the difference is small and not noticeable. At points where the bit is used in the computation *and* this information leaks in the signal, this difference is likely to be large or noticeable.

Figures 7, 8, 9, and 10 show the results of a differential side–channel attack on an S–box output bit in the first cycle of the DES implementation, using the four different signals. Figures 7, 8 and 9, are for the EM signals and Figure 10 is for the power signal. All figures are aligned in time. In all figures, the X–axis shows elapsed time in 10ns units and the Y–axis shows the difference in the averages of signals with bit=0 and bit=1 for 2000 invocations of DES with random inputs. Even at this resolution, it is clear that the leakage results are qualitatively different from each other. There are some gross similarities between the EM leakages in Figures 7 and 8 and between the EM leakage in Figure 9

[9] This is because clock edges are not very sharp in practice.

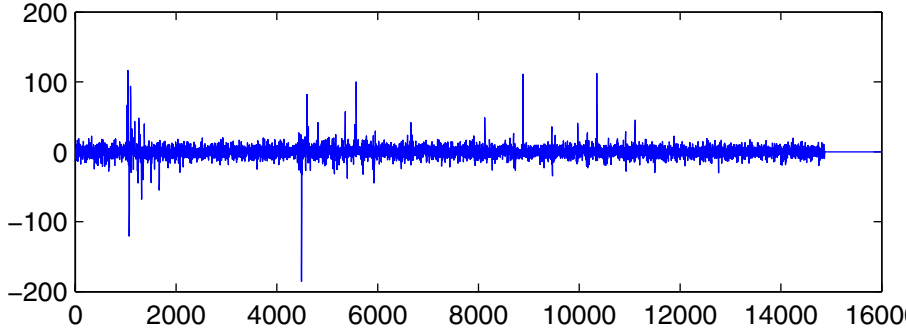

Fig. 7. DEMA attack on DES on smartcard B using the 224.5 MHz carrier

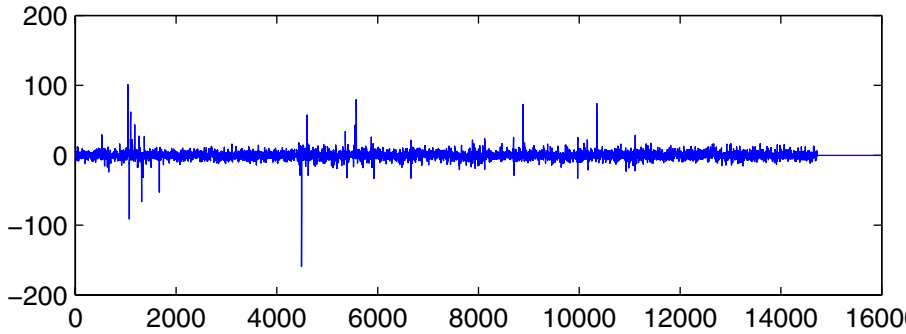

Fig. 8. DEMA attack on DES on smartcard B using the 262MHz carrier

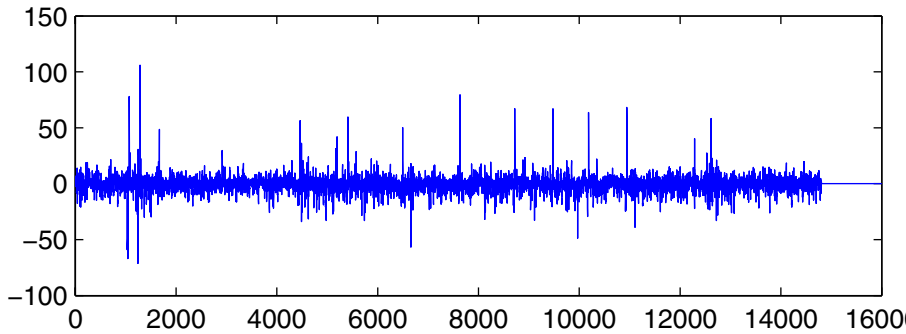

Fig. 9. DEMA attack on DES on smartcard B using the 188MHz carrier

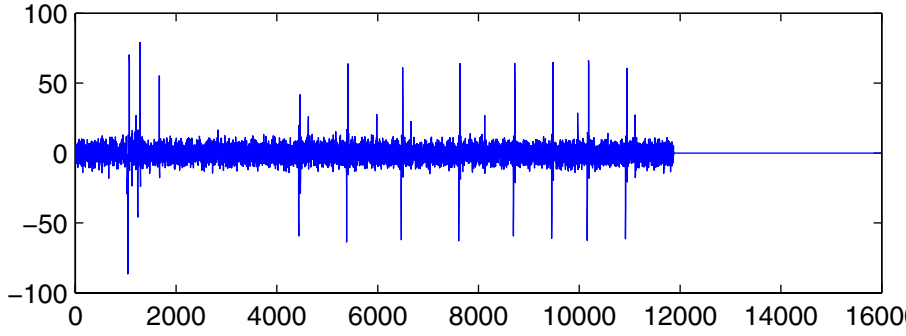

Fig. 10. DPA attack on DES on smartcard B

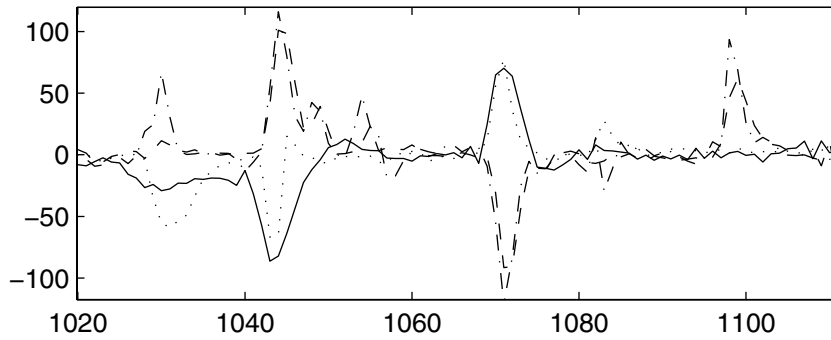

Fig. 11. Comparison of DEMA/DPA Leakages at region 1020–1110

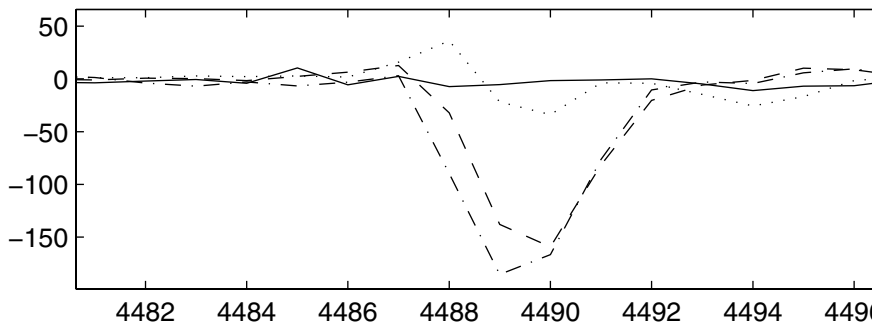

Fig. 12. Comparison of DEMA/DPA Leakages at region 4482–4496

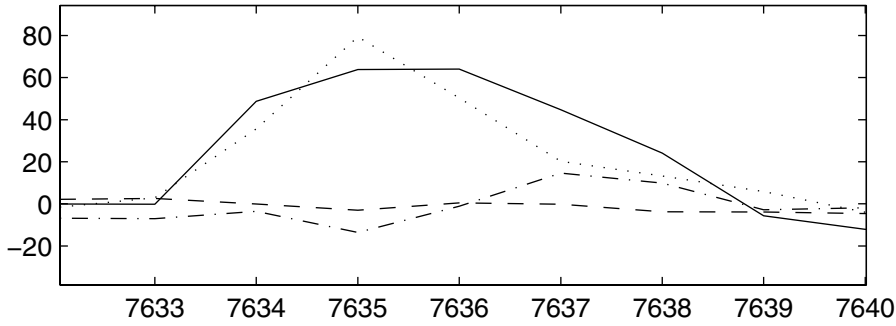

Fig. 13. Comparison of DEMA/DPA Leakages at region 7633–7640

and the power leakage in Figure 10. These leakages can be compared by plotting them together. Figures 11, 12, 13 show some of the regions in such a plot. Each leakage is plotted in a different line–style, with the power leakage being a solid line and the 3 EM leakages plotted in different broken–line styles (188MHz with a dotted line, 224.5MHz with a dashed line and 262MHz with alternate dot and dashes). It is clear from these figures that even though the signals fall into two gross classes at the macro level, there are significant differences even between signals within a class at a cycle level (see Figure 11). Moreover, there are leakages which appear in EM signals (and sometimes excessively so), which do not appear in the power signal (see Figure 12). Such leakages are due to what we will later term as a "bad" instruction. There are also leakages which are large in power, but low in some (but not all) EM signals (see Figure 13).

5 The Power of the EM Side–Channel(s)

Using low–cost EM equipment, which can collect only one signal at time, we have experimented with a wide variety of cryptographic equipment and computing peripherals. We could easily launch attacks such as simple and differential electromagnetic attacks (SEMA and DEMA [9]) on straight–forward implementations of DES, RSA and COMP128 on smart cards, cryptographic tokens and SSL accelerators. While these attacks are interesting, this does not justify why EM side–channel(s) should be used in preference to others. In some cases, e.g., attacking an SSL accelerator from a distance, the only strong side–channel available is EM. We now show that the EM side–channel is extremely useful even in cases where the power side–channel is available, i.e., the EM side channel can be used to break power analysis resistant implementations.

In [8], a suggested countermeasure to power analysis is to use only those instructions whose power leakage is not excessive and to refresh sensitive information, such as a key, after each invocation in a non–linear fashion. This forces the adversary to extract a substantial portion of the key from a *single invocation* since incomplete key information does not help in subsequent invocations. Another class of countermeasures is based on splitting all sensitive information

into shares [2,4]. The basic idea is that uncertainty in the information about each share is exponentially magnified in proportion to the number of shares.

5.1 Bad Instructions Defeat Power Analysis Countermeasures

The key to breaking both classes of countermeasures is to identify instructions, that we term *bad* instructions, which leak much more information in some EM signals as compared to the power signal. If bad instructions are used in power–analysis resistant implementations, the leakage assumptions made the implementation become invalid.

For all chip cards that we examined, there were several bad instructions. In our investigations, we did not find any instruction that leaked in the power side–channel but did not leak in some EM side–channel. This can happen if all critical parts of a chipcard are well–shielded but the power signal is not. We feel that this is unlikely since a designer who shields EM emanations so well is also likely to protect against power signal leakages.

For example, the bit test instruction is very useful for implementing algorithms, such as DES, which involve bit level permutations. For example, it can be used for key expansion and P–permutation. The value of the tested bit is known to have low power leakage characteristics on many smart cards. This is because the power signal is dominated by the larger currents needed to drive bus lines as opposed to the smaller currents within a CPU performing a bit–test. Thus, it is likely to be present in some power analysis resistant implementations.

However, this bit test instruction turned out to be a bad instruction for smartcard B. When the internal noise generators had been turned off, we observed that it leaked information about the tested bit from even a *single* signal sample in the EM side–channel but not in the power side–channel. This is illustrated in Figures 14 and 15 where the amplitudes of two EM signals are plotted against time (in 10ns units). In both figures, the data was collected by a 12–bit, 100MHz scope after demodulating at the 262MHz carrier. Figure 14 shows two EM signals in which the bits tested are both 0: this is seen as a low value in both the signals at the point 18915. Figure 15 shows two EM signals in which one of the bits tested is 0 and the other is 1: this is seen as a low value in one of the signals and a high value in the other at point 18780. These points correspond to the cycle where the value of the bit is tested.

Even with noise generators enabled, it was possible to classify the bit value correctly with high probability by using only a few samples (20–30). We experimentally verified that no such differences were to be found at the corresponding cycle for the power signals. Even after statistical tests involving thousands of power samples, there are no differences at this cycle although they show up at other cycles (such as the point where the byte containing the bit is loaded).

If the bit test instruction was used for implementing permutations in a power analysis resistant implementation of DES, with noise generators off, a SEMA attack would be sufficient to extract the DES key regardless of which class of countermeasures [8,2,4] was used. However, if noise was enabled, then the countermeasure of [8] may still remain immune. However, as we now show, higher order statistical attacks would still defeat the countermeasures of [2,4].

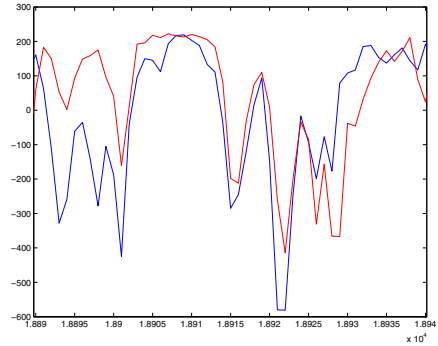

Fig. 14. Two EM Signals where tested bits are 0 (seen as low values at 18915)

Higher Order EM Attacks on Secret–Sharing. The secret–sharing based DPA countermeasure chooses a value for the number of shares based on leakage characteristics and the desired level of resistance against higher order power analysis attacks [2,4], in terms of the number of samples required to break the implementation. If a leakage is superior in an EM signal, then the number of samples for the corresponding higher order EM attack can be substantially lower. The task of an adversary attempting this higher order attack may be complicated by the fact that the code could be unknown. We now outline a general technique to perform higher order EM attacks exploiting bad instructions which can work even when the code is unknown.

Attacks on Unknown Code. Assume a chipcard containing an unknown k–way secret–sharing based DPA protected code for a known algorithm. Further assume that "bad" instructions have already been identified and some of these instructions are used to manipulate shares. These, of course, are necessary conditions for EM attacks to be more effective than power attacks. Let us also assume that it is possible to use signal processing to remove execution sequence and variable clock randomization that has been added as countermeasures to complicate alignment of signals and that each signal can be realigned into a canonical execution sequence[10].

The value of k is usually small. For simplicity, assume that k is 2: the attack generalizes for slightly larger k. Fix a reasonable limit L on the number of EM samples that can be collected. We now show that if k is small and if with knowledge of the code we could have broken the protected code using L samples, then this attack can break the unknown protected code with $O(L)$ samples.

In case of a two–way split, a first step is to identify the two locations where the shares of an algorithmic quantity are being manipulated using bad instructions. If code–execution randomization can be effectively neutralized, then this can be done for many algorithms. Knowing the algorithm, one can provide two different inputs such that the value of the variable is different for these inputs while most

[10] We have found this to be quite feasible, especially since canonicalization with a reasonable probability of correctness suffices.

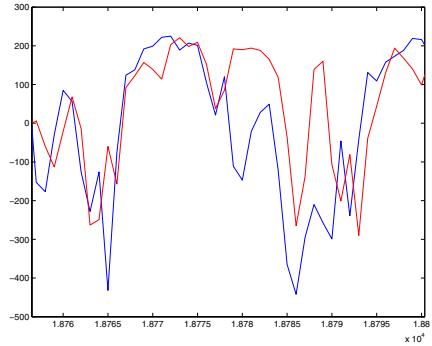

Fig. 15. Two EM Signals where tested bits are 0 and 1 (low and high values at 18780)

of the other variables are the same within the window of interest. For example, in DES if algorithmic quantity is an S–box output, one could choose two inputs which differ only on 1 bit so that only that S–box output is affected.

Take L EM samples for each of these two different inputs. If the exact locations where the two shares were manipulated was known, then there is second order statistic, S, that can be applied to the signal at these two locations to distinguish between the two different inputs, thus enabling hypothesis testing.

Without location information, one can only assume that the two locations are an integral number, D, of clock cycles apart. So the strategy is to compute the statistic S for each point on the signal with respect to a corresponding point D cycles away. This is done for both sets of inputs for all reasonable values of D. If the shares are not manipulated at distance D, then the values of the statistic S at all points will be similar for the two inputs. However, for the right value of D, there will be a significant difference in S exactly at the point where the first share is manipulated and thus the exact location of each share is revealed.

An optimization is to choose the two inputs so that multiple algorithmic variables are different. Then the above exercise will yield candidate locations for the shares for all these variables. Once these locations are identified, second (or higher) order attacks can be applied as if the code were known.

To validate this approach, we implemented a two–way XOR–based secret sharing scheme for bits on smartcard B with noise generators on. The sample code split the input bits into pairs of shares and tested the values of the bits of the shares using the bit test instruction. We confirmed that DPA and DEMA on input bits did not work. In the implementation, the shares of one of the input bits were tested 40 cycles apart. Section 5.1 shows that when a bit is 1, the signal at the bit test instruction is high and when the bit is 0, the signal is low. For a 2–way bit split using an XOR–scheme, the shares of a 0 bit will be (0, 0) or (1, 1) with equal probability and the shares of a 1 bit would be (0, 1) or (1, 0) with equal probability. This suggests that a good statistic S is the correlation coefficient between the corresponding signal points where the shares of bits are being tested. S will be positive when the bit is 0 and negative when the bit is 1.

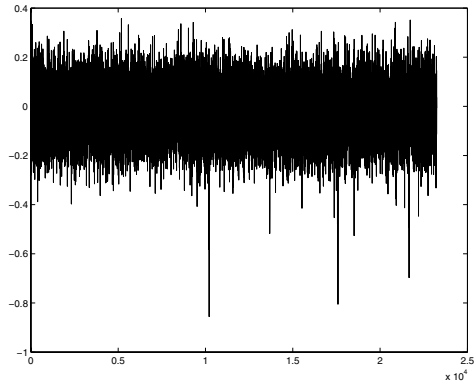

Fig. 16. Difference in correlation statistics for $D = 40$, $L = 500$

We experimented with $L = 500$, for two different inputs, which differed in exactly three bits. Figure 16 shows the difference in the statistic S when the distance D is 40, plotted against elapsed time in 10ns units. The three significant negative peaks were confirmed to be at exactly the points where the first shares of the three bits (that differ) were being manipulated. In fact this attack even worked when $L = 200$. No peaks were seen when D differed from 40. The same experiment when repeated for $D = 40$ for five thousand power signals did not work showing that higher order DPA does not work with five thousand signals.

6 Conclusion and Further Work

This paper, together with other recent work [9,6,1], lays the foundations for a theory of EM leakages during computation in CMOS devices. While a significant amount of information had been publicly available on EM leakages, that work mostly dealt with leakages from displays and other peripherals[10].

Our paper highlights a key aspect of the nature of EM leakage, i.e., the presence of multiple, unintentional, information–bearing signals within this side–channel. In addition, this paper also demonstrates why EM side–channel(s) are so useful: multiple signals with differing leakage characteristics enable a variety of attacks, including attacks against implementations secure against power analysis.

Despite their effectiveness, the single–channel attacks described in this paper provide only a glimpse of what is possible. Combining leakages from multiple EM channels using techniques from Signal Detection and Estimation Theory yield substantially stronger attacks. The existence of such multi-channel attacks highlights a pressing need for models and techniques to assess the net information leakage from all the EM signals realistically available to an adversary. Preliminary results on these aspects of the EM side–channel(s) is described in more detail in [1].

7 Countermeasures

Due to the presence of several unexpected EM leakages, a comprehensive EM vulnerability assessment has to be an integral part of any effort to develop countermeasures against EM attacks on specific implementations. Such countermeasures fall into two broad categories: *signal strength reduction* and *signal information reduction*. Techniques for signal strength reduction include circuit redesign to reduce egregious unintentional emanations and the use of shielding and physically secured zones to reduce the strength of compromising signals available to an adversary relative to ambient thermal noise. Techniques for signal information reduction rely on the use of randomization and/or frequent key refreshing within the computation [7,8,2,4] so as to substantially reduce the effectiveness of statistical attacks using the available signals.

Acknowledgments. This paper has greatly benefitted from the advice of anonymous CHES referees whose comments helped in selecting aspects of our work on the EM side–channel to create a more focussed paper. We would like to thank Helmut Scherzer for key components that enabled this work, in particular, his help with experimental setup, smart card programming and data collection and analysis tools. We would also like to thank Elaine and Charles Palmer for their encouragement and useful comments.

References

1. D. Agrawal, B. Archambeault, J. R. Rao, and P. Rohatgi, EM Side–Channel(s): Attacks and Assessment Methodologies, http://www.research.ibm.com/intsec.
2. S. Chari, C. S. Jutla, J. R. Rao and P. Rohatgi. Towards Sound Countermeasures to Counteract Power–Analysis Attacks. Proc CRYPTO '99, LNCS 1666, pp 398–412.
3. S. Chari, J. R. Rao and P. Rohatgi. Template Attacks, Proc CHES '02.
4. L. Goubin and J. Patarin. DES and Differential Power Analysis. Proc CHES '99, LNCS 1717, pp 158–172.
5. NSA Tempest Series `http://cryptome.org/#NSA--TS`.
6. K. Gandolfi, C. Mourtel and F. Olivier. Electromagnetic Attacks: Concrete Results. Proc CHES '01, LNCS 2162, pp 251–261.
7. P. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS and Other Systems. Proc CRYPTO '96, LNCS 1109, pp 104–113.
8. P. Kocher, J. Jaffe and B. Jun. Differential Power Analysis: Leaking Secrets. Proc CRYPTO '99, LNCS 1666, pp 388–397.
9. J.–J. Quisquater and D. Samyde. ElectroMagnetic Analysis (EMA): Measures and Counter-Measures for Smart Cards. In Smart Card Programming and Security (E-smart 2001), LNCS 2140, pp. 200–210.
10. The complete unofficial TEMPEST web page, http://www.eskimo.com/~joelm/tempest.html.