Simple Photonic Emission Analysis of AES

Photonic side channel analysis for the rest of us

Alexander Schlösser*,[1], Dmitry Nedospasov*,[2], Juliane Krämer[2],
Susanna Orlic[1], Jean-Pierre Seifert[2]

[1]Optical Technologies, Technische Universität Berlin, Germany
{schloesser,orlic}@opttech.tu-berlin.de
[2]Security in Telecommunications, Technische Universität Berlin, Germany
{dmitry,juliane,jpseifert}@sec.t-labs.tu-berlin.de
* Equal contribution

Abstract. This work presents a novel low-cost optoelectronic setup for time- and spatially resolved analysis of photonic emissions and a corresponding methodology, Simple Photonic Emission Analysis (SPEA). Observing the backside of ICs, the system captures extremly weak photoemissions from switching transistors and relates them to program running in the chip. SPEA utilizes both spatial and temporal information about these emissions to perform side channel analysis of ICs. We successfully performed SPEA of a proof-of-concept AES implementation and were able to recover the full AES secret key by monitoring accesses to the S-Box. This attack directly exploits the side channel leakage of a single transistor and requires no additional data processing. The system costs and the necessary time for an attack are comparable to power analysis techniques. The presented approach significantly reduces the amount of effort required to perform attacks based on photonic emission analysis and allows AES key recovery in a relevant amount of time.

Keywords: Photonic side channel, emission analysis, optical, temporal analysis, spatial analysis, AES, full key recovery

1 Introduction

Most side channel attacks focus on system-wide information leakage. However, photonic side channels also allow selective in-depth analysis of specific parts of the hardware. Leakage can even be extracted from single transistors within an integrated circuit (IC). This selectivity has important implications for side channel analysis. Potentially, signals can be captured that consist entirely of leakage and are not impeded by side effects originating from the rest of the system.

Photonic side channel analysis can be considered far more powerful than other side channel analysis techniques in use today. Recovering side channel leakage across large areas of an IC or logic becomes unnecessary. Instead, by

investing some time and effort into spatial photonic analysis of an IC's layout, potential weaknesses of the implementation can be efficiently identified and exploited. In this fashion, attacks targeting single transistors become reality. Moreover, targeting specific elements of a chip's logic results in significantly better Signal-to-Noise-Ratios (SNR). Subsequent analysis of such signals can be as simple as a binary evaluation of the traces, akin to Simple Power Analysis (SPA).

The main contributions of this paper are as follows:

A low-cost photonic emission analysis system. The system is a low-cost solution to capture photonic emissions of ICs. At approximately the price of a mid-range oscilloscope, it is specifically tailored for photonic side channel analysis. This system also cuts down on measurement times when compared to other state-of-the-art photonic emission analysis methods, see Section 2.4.

A novel methodology: Simple Photonic Emission Analysis. With this methodology we are able to recover signals that consist entirely of side channel leakage. By carefully identifying potential targets, we can exploit the spatial separation of logic circuits to eliminate noise within the measurement, circumvent countermeasures and eliminate other potential side effects stemming from the rest of the cryptosystem. Signals recovered by this methodology require little to no additional analysis or post-processing and the key can be recovered directly by simply observing the traces akin to SPA.

Results of a successful SPEA of AES. Using SPEA in combination with our photonic system, we were able to correctly recover the full secret key of a Proof-of-Concept (PoC) AES-128 implementation running on a common microcontroller, the ATmega328P. The process technology of the ATmega328P, approximately 250 nm, is the technology used in most smartcard deployments today. We exploited the photonic side channel leakage of the row address decoders to monitor accesses to the AES S-Box and were able to recover the full AES secret key. This attack works in software and hardware, and even in the presence of hardware countermeasures, such as memory scrambling and encryption.

Advantages and limitations of our system and methodology. In addition to presenting the system, the methodology and practical results, we present an overview of several potential countermeasures to this kind of attack. Also, we explain how such an attack could be extended to AES implementations using compressed tables and to alternative hardware implementations, including other volatile and non-volatile memories.

Organization. The rest of this work is structured as follows: In Section 2 we present additional background information on photonic emissions in CMOS, detection techniques, the AES algorithm and related work. Section 3 describes the optoelectronic system used in this work. In Section 4 we detail our attack against a PoC AES implementation. Section 5 presents additional considerations for our system and methodology and also includes several potential countermeasures. Finally, we conclude in Section 6.

2 Background

2.1 Photonic Emissions in CMOS

In CMOS technology, carriers gain kinetic energy in a transistor's conductive channel as they are accelerated by the source-drain electric field. At the drain edge of the channel where the field is most intense, this energy is released in radiative transitions, generating photons [21]. This hot-carrier luminescence is dominant in n-type transistors due to the higher mobility of electrons as compared to holes. Consequently, optical emissions of CMOS logic show a data-dependent behaviour similar, but not equal to power consumption. The photon generation rate is proportional to the supply voltage and the transistor switching frequency.

Since multiple interconnect layers of modern IC designs prevent generated photons from escaping the IC on the frontside, hot-carrier luminescence is best observed from the backside. In this case emitted photons have to pass through the silicon substrate, which is absorptive for wavelengths shorter than the bandgap energy, leaving only very few Near-Infrared (NIR) photons for analysis. To reduce absorption the substrate can be mechanically thinned with standard backside polishing machines or alternative techniques [12].

2.2 Detection Techniques

Detection of hot-carrier luminescence from the backside, i.e., single near-infrared photons, must overcome two issues. Firstly, charge coupled devices (CCD) exhibit high spatial resolution, but only allow slow frame rates; single pixel detectors like Photo Multiplier Tubes (PMT), Avalanche Photo Diodes (APD) and Superconducting Single Photon Detectors (SSPD) offer picosecond timing resolution, but only for one small detection area. Secondly, readily available and affordable Si-based detectors only cover a fraction of the relevant NIR spectral range. Thus, for efficient photonic emission analysis more complex and expensive solutions, such as InGaAs-based detectors, are necessary. This is especially true for analysis of modern ICs with small feature sizes, as the emission spectrum shifts further to the infrared with decreasing transistor gate length.

One of the most complex detector technologies in use today is Picosecond Image Circuit Analysis (PICA), which is based on gated multi-channel plates with NIR-sensitive cathode materials. It was developed explicitly for failure analysis of semiconductors [20, 10]. PICA delivers both spatial and temporal resolution, but offers only very limited NIR-sensitivity. Additionally, integrated PICA systems have a starting price of around one million Euros in 2012. In contrast to our optoelectronic system, it is unlikely that these systems will ever become a commodity.

2.3 The AES Algorithm

The Advanced Encryption Standard (AES) is a secret key encryption algorithm based on the Rijndael cipher [4]. AES has a fixed block size of 128 bits and

operates on a 4×4 matrix of bytes, named the state. Depending on the length of the key, which is 128, 192, or 256 bits, the cipher is termed AES-128, AES-192, or AES-256. The algorithm is specified as a number of identical rounds (except for the last one) that transform the input plaintext into the ciphertext. AES consists of 10, 12 and 14 rounds for 128-, 192- and 256-bit keys, respectively.

Since our attack exploits the leakage obtained during the beginning of the first round of AES, we present only the two operations that are executed until then, namely `AddRoundKey` and `SubBytes`. In the `AddRoundKey` step, each byte of the state is combined with a byte of the round key using the exclusive or operation ($\oplus$). The round key is derived from the original secret key using Rijndael's key schedule; each roundkey is the same size as the state, i.e., 128 bits. The first `AddRoundKey` operation uses the original secret key, or the first 128 bits of the secret key for AES-192 and AES-256, respectively. In the `SubBytes` step, each byte in the state is replaced with its entry in a fixed 8-bit lookup table, denoted the S-Box. This is the only operation that provides non-linearity in the algorithm.

2.4 Related Work

In the failure analysis community hot-carrier luminescence has been used primarily to characterize implementation and manufacturing faults and defects [6, 15]. In this field the technology of choice to perform backside analysis is PICA [1] and superconducting single photon detectors [17]. Both technologies are able to capture photonic emission with high performance in their respective field, but carry the downside of immense cost and complexity. One of the first uses of photonic emissions in CMOS in a security application was presented in [7], where the authors utilize PICA to spatially recover information about exclusive or operations related to the `AddRoundKey` operation of AES. Employing PICA in this manner, led to enormous acquisition times. This is especially true considering the size of the executed code. It took the authors 12 hours to recover a single potential key byte [7]. In the same time our system recovers all 16 bytes of the 128-bit AES key twice. In [16] low-cost equipment was used to capture photonic emissions via backside analysis and gain basic information about the operations executed on an IC. Even though the author presented low-cost solutions to both spatially and temporally resolved photonic emission analysis, no attacks using temporal information were demonstrated. Most recently an integrated PICA system and laser stimulation techniques were used to attack a DES implementation on an FPGA [5]. The authors showed that the optical side channel can be used for differential analysis and partly recovered the secret key using temporally resolved measurements. As the authors noted, the use of equipment valued at more than two million Euros does not make such analysis particularly relevant. Additionally, the analysis strongly relied on a specific implementation of DES in which inputs were zeroed. The results required differential analysis and full key recovery was not presented.

In the field of electromagnetic side channel analysis, location-dependent leakage was successfully exploited in an attack on an elliptic curve scalar multiplica-

tion implementation on an FPGA using a near-field EM probe [9]. The authors scanned the die surface and collected EM traces at every point. They demonstrated that location-dependent leakage can be used in a template attack and countermeasures against system-wide leakage thus can be circumvented.

Side channels based on monitoring memory accesses have also been researched in the field of cache attacks. In 2004, Bernstein conducted a known-plaintext memory-access timing attack on the OpenSSL AES implementation that uses precomputed tables [2]. The mathematical analysis of our attack is very similar.

3 Experimental Setup

To increase the relevance of semi-invasive optical vulnerability analyses and attacks, the experimental system was constructed with off-the shelf components and employs readily available technical solutions. As neither Si- nor InGaAs-based detectors can deliver both spatial and temporal resolution in the NIR range for an affordable price, our system combines the inherent advantages of both detector technologies in an integrated system. The overall complexity and cost of this system is considerably lower than common semiconductor failure analysis and even power analysis systems, as its price is comparable to that of a mid-range oscilloscope.

3.1 Hardware

The experimental setup consists of two detectors optically and electrically connected to the Device Under Test (DUT) via a custom-built near-infrared microscope and an FPGA-based controller, see Figure 1(a). The DUT is soldered onto a custom printed circuit board and mounted on lateral travel stages. The microscope itself uses finite conjugate reflection type objectives with high numerical aperture and gold plated mirrors to achieve maximum throughput and a spectrally flat transmission curve. After passing the objective, the hot-carrier luminescence spectrum is split by a dichroic mirror and each part directed to the relevant detector. A single Si-CCD serves as the primary detector and captures NIR photons below the silicon bandgap energy. Its mega-pixel deep depletion sensor is back illuminated and thermoelectrically cooled to ensure optimal NIR sensitivity and low dark current rates. This detector can take dark-field reflected-light as well as emission images through the substrate silicon with a diffraction-limited spatial resolution below 1 μm. The acquisition time necessary for adequate emission images ranges from a few seconds to many minutes. It depends strongly on the supply voltage of the DUT, the switching frequency of the transistors under observation and the substrate thickness. Optimized software implementations can increase the execution loop frequency and thereby the switching frequency, which often reduces acquisition times to seconds [11].

The secondary detector is a single InGaAs/InP Avalanche Photo Diode (APD) commonly found in telecom applications (Telcordia GR-468-CORE). It is operated in Geiger mode and thermoelectrically cooled. Increased dark current

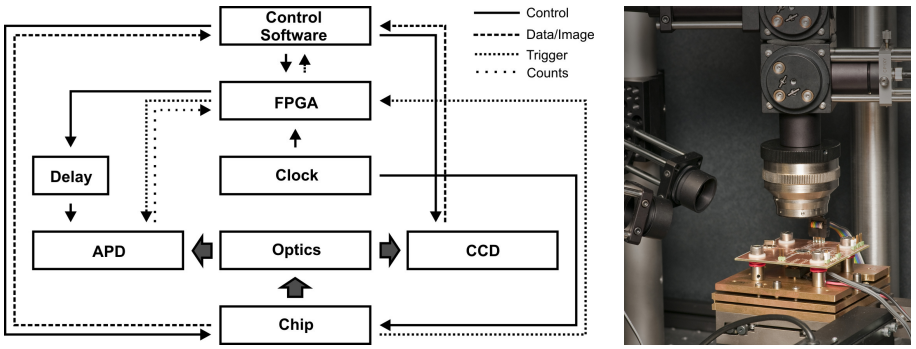

(a) Optical Emission Analysis Setup (b) Microscope and DUT

Fig. 1. The NIR microscope connects the DUT to the detectors. The two detectors are controlled via an FPGA-based controller, which handles gate synchronization and delay control as well as time-to-amplitude conversion and multichannel counting, see Figure 1(a). The DUT is mounted upside down on a custom printed circuit board underneath the microscope objective, see Figure 1(b).

and afterpulsing, common to InGaAs/InP-APDs, are reduced by gated operation and extensive quenching circuits. The diode is coupled to the microscope via an optical fiber. The fiber's aperture can be freely positioned in the image plane and its object plane size varied by changing the position and magnification of the fiber coupler. Areas of interest, identified in an emission image, can thus be selected for temporal analysis with high spatial selectivity. Because of its spectral sensitivity above 1 μm, unlike the CCD, this detector does not require a thinned DUT substrate, as silicon is transparent in this spectral range. Hence, if spatial orientation relative to the IC's layout can be obtained by other means, substrate thinning can be omitted completely. This is especially true when applying our methodologies across multiple samples of an identical IC, as only a single sample has to be prepared.

In gated operation the APD is rendered sensitive only for a short window in time, the detection gate, in every signal cycle. To reconstruct the complete signal temporally, the detection gate has to be synchronized and shifted relative to the signal with every signal cycle iteration, similar to a sampling oscilloscope. Provided the gate delay can be controlled with high resolution, the time resolution and inversely the measurement time depends only on the minimal gate width. To implement this detection scheme we use an FPGA-based controller phase-locked to the DUT clock. As the DUT executes the target program code, the phase-locked FPGA digitally delays and triggers the APD detection gate. Detection events are sent back to the FPGA and counted in the corresponding time bins. An additional analog delay can be employed for fine delay control. The absolute time resolution of our system is jitter-limited to approximately 1 ns.

The measurement time to reconstruct the extremely weak photoemission signals can be immense: Hundreds of thousands of samples may be necessary

to achieve an adequate SNR. To drastically reduce the measurement times, the FPGA triggers hundreds of APD detection gates per execution of the IC. This results in interleaved measurements.

For our practical evaluation the DUT board consisted of an inversely soldered ATmega328P supplied with 5 V operating voltage, a 16 MHz quartz oscillator as well as decoupling capacitors and an I/O header for external communications and programming. To shorten the emission image acquisition times, the backside of the ATmega328P was mechanically polished with an automated backside sample preparation machine. The remaining substrate was approximately 25 μm thick.

3.2 Software

The PoC implementation running on the ATmega328P microcontroller consisted of a software AES implementation. To increase the frequency of the execution, only the first `AddRoundKey` operation and part of the first round of the AES algorithm were computed on the chip after which the input was reset and the measurement restarted. Specifically, the implementation computed only the first `AddRoundKey` and `SubBytes` operations, see Section 2.3. Most notably, the AES S-Box, see Table 2 in the appendix, was implemented in the microcontroller's data memory, i.e. the SRAM. Both stack- and heap-based AES implementations were tested and resulted in offsets within SRAM as described in Section 4.1.

4 Practical Results

This section details our practical results in which we applied SPEA to mount a practical attack against a Proof-of Concept (PoC) AES implementation.

4.1 Monitoring SRAM Access

Address logic is implemented similarly across all platforms and memories. Hence, it is a particularly interesting and relevant target. In our implementation the S-Box is contained within the SRAM, which led us to consider possible side channels that exist within this memory. Specifically, we considered how SRAM is accessed in general and how the S-Box would be accessed in the PoC AES implementation. Memory is structured in rows and columns. In the case of the ATmega328P, each 512 kilobyte SRAM bank is made up of 64 rows and 64 columns. Thus, each row stores a total of 64 bits, or 8 bytes. An SRAM cell read begins with assertion of the word line [13]. When any cell within a row of memory is accessed, the word line for the entire row is asserted. This row select signal is driven by an inverter that is part of the row decode logic [22].

We used the Si-CCD detector and techniques introduced in [11] to analyze the emissions of these memory accesses for different addresses and values. With spatial resolution, accesses between even two adjacent rows of SRAM can be clearly differentiated, see Figure 2(a) and 2(b). The same kind of analysis can also be applied to the column addresses.

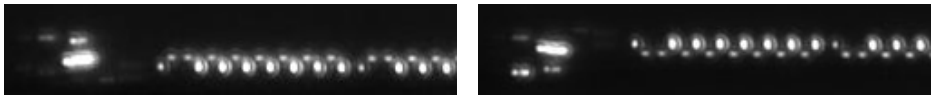

(a) Access to 0x300 (b) Access to 0x308

Fig. 2. 120 s emission images of memory accesses to two adjacent memory rows obtained with the Si-CCD detector.

By studying the emission images of our PoC AES implementation we identified the S-Box within memory, see Figure 3 and Table 3 in the appendix. The emissions of the row drivers are clearly visible to the left of the individual memory lines. The emission image Figure 3 reveals that the S-Box spans 33 rows of the SRAM and not 32 as expected. Since the S-Box was implemented as an array of bytes within data memory, its address depends on how many other variables are allocated within memory. It is therefore unlikely to be aligned to the beginning of a memory row. During our experiments, the S-Box always exhibited an offset unless the address was set explicitly.

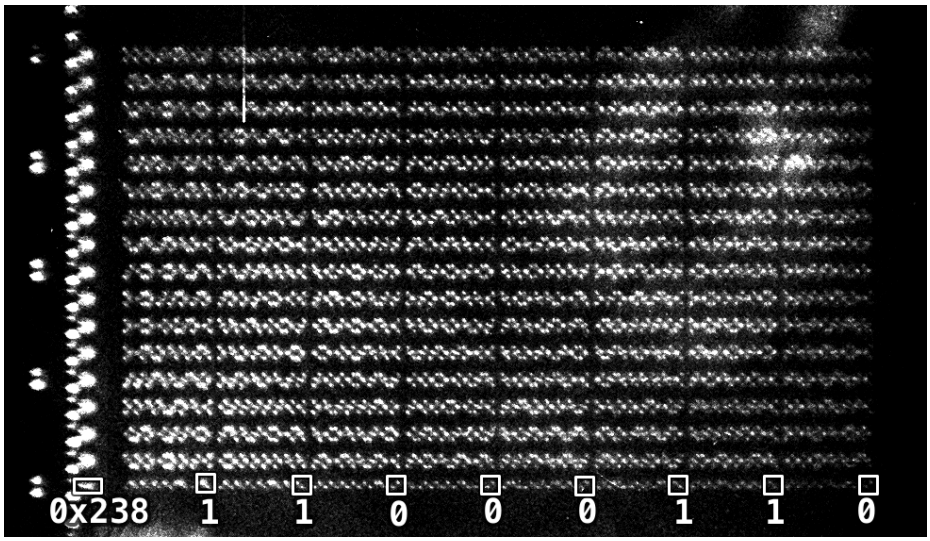

Fig. 3. Optical emission image of the S-Box in memory. The 256 bytes of the S-Box are located from `0x23F` to `0x33E`, see Table 3 in the appendix. The address `0x23F` is the seventh byte of the `0x238` SRAM line, i.e. the S-Box has an offset of 7 bytes. The emissions of the row drivers are clearly visible to the left of the memory bank. The image allows direct readout of the bit-values of the stored data. The first byte for example, as shown in the overlay, corresponds to $01100011_2 = 63_{16}$, the first value of the AES S-Box.

4.2 Key Recovery Using the Photonic Side Channel

By observing time-resolved access patterns to a specific row within the S-Box, the set of key candidates can be greatly reduced. In this attack on the first round of AES, we observe accesses to a single fixed row r over 256 input messages, in which all bytes of the input are equal. That is, for every text byte $p_i = i$, $i \in \{0, \ldots, 255\}$, there is an input message m_i that consists of 16 concatenated p_i's. Thus, after the exclusive or of the first `AddRoundKey` operation, every element of the S-Box is accessed exactly once for every processed byte $b \in \{1, \ldots, 16\}$ in the first AES round. However, we only measure accesses to rows, i.e., we use the access patterns to a given row r to eliminate certain key candidates. If potential offsets are also taken into consideration, the number of key candidates that can be identified using this attack are shown in Table 1.

Offset	remaining candidates per key byte		unresolved bits of the whole key	
	r = 1 or r = 33	$r \in \{2, ..., 32\}$	r = 1 or r = 33	$r \in \{2, ..., 32\}$
0	8	8	48	48
1	**1**	**1**	**0**	**0**
2	2	2	16	16
3	**1**	**1**	**0**	**0**
4	4	8	32	48
5	**1**	**1**	**0**	**0**
6	2	2	16	16
7	**1**	**1**	**0**	**0**

Table 1. Number of candidates per key byte and unresolved bits of the whole key, depending on the offset and row. For an offset of 0, there are only 32 rows.

Note, uneven offsets always result in unique access patterns, allowing for full key recovery directly.

An attack against our PoC implementation was thus implemented in three parts: (1) an offline precomputation of potential access patterns, (2) an online measurement of emissions over all possible input bytes, and (3) an evaluation resulting in a reduced set of key candidates.

Offline precomputation. For the first round of AES and independent of the value b, to determine which row of the S-Box is accessed by a plaintext byte p_i, key byte $k_j = j, j \in \{0, \ldots, 255\}$, and for a given offset $o \in \mathbb{N}^{\geq 0}$ and width of the rows $w \in \mathbb{N}^{\geq 1}$, we introduce the following function $row : \{0, \ldots, 255\} \times \{0, \ldots, 255\} \to \mathbb{N}$:

$$row(p_i, k_j) = \lfloor (p_i \oplus k_j) + o)/w \rfloor + 1$$

Using this function, we can generate the array A, storing the sets of plaintexts accessing row $r_l = l, l \in \{1, \ldots, \lfloor (256 + o)/w \rfloor + 1\}$ for key byte k_j:

$$A(r_l, k_j) = \{p_i | r_l = row(p_i, k_j)\}$$

Time-resolved emission measurement. Expecting input-dependent accesses to the memory rows, we set our APD detector to measure the photonic emissions from the row driver inverter corresponding to a fixed memory row with S-Box elements. Emission traces for all 256 inputs were captured, with an acquisition time of 90 seconds each. A row access resulted in a clearly defined photon detection peak with a high SNR, see Figure 4. Empirically defining a threshold level l determines if a fixed row r of the S-Box was accessed within the measurement. For a period of time t_b in which the b-th byte is processed, and for the measured photonic emission intensity I_r of a row r:

$$M_r(t_b) = \{p_i | max(I_r(p_i, t_b)) > l\}$$

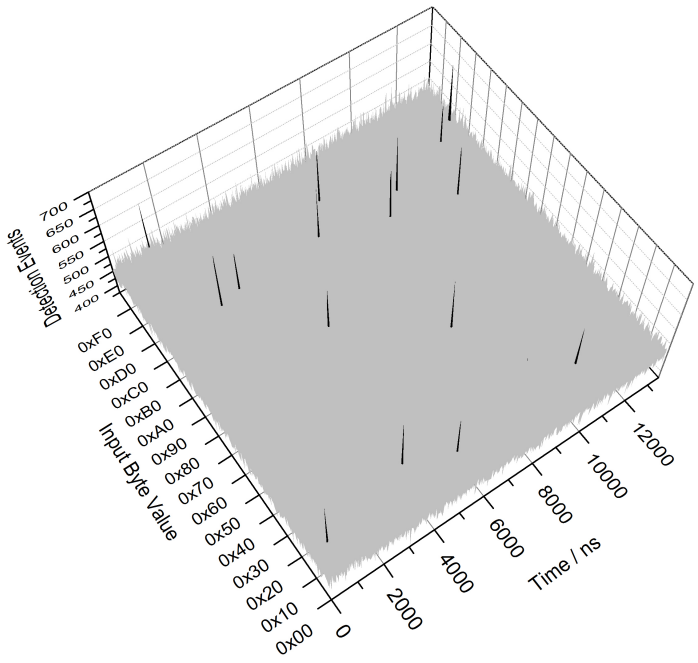

Fig. 4. Time-resolved photonic emission traces of the PoC implementation over all possible input bytes. All 16 peaks, corresponding to the bytes of the AES-128 key, are clearly identifiable. Specifically, results are shown for the following key, `[0x10, 0xF1, 0xB3, 0xB7, 0x1E, 0x81, 0x12, 0xBA, 0xD1, 0x56, 0xAD, 0xBB, 0x17, 0xA2, 0xCA, 0xD5]`. Note the SNR, comparing peaks to the noise floor of traces that are not accessing the SRAM. A two dimensional plot, Figure 5, can be found in the Appendix.

Full key recovery. Finally, we can identify the resulting set of candidates for the b-th key byte by analyzing, which key bytes could have caused the measured accesses, using the precomputed array A:

$$K(b) = \{k_j | A(r, k_j) = M_r(t_b)\}$$

For an S-Box with an uneven offset this results in a complete key recovery. In the case of an even offset a maximum of eight candidates per key byte remain. This set can be easily reduced by a second set of measurements selecting another row or cross-correlation with measurements from column decoder logic.

The measurement time to fully capture the photonic emission signal over time and all possible inputs, as seen in Figures 4 and 5, amounted to just over six hours. However, since accesses to the S-Box occur at constant points in time in every execution, these are the only points that need to be observed in subsequent measurements. This cuts down the necessary measurement time immensely. After initial analysis, any subsequent attacks on identical implementations can therefore recover the key in less than 45 minutes.

5 Discussion

The effectiveness and relevance of the presented attack depends on and therefore highlights the importance of preliminary spatial analysis. It demonstrates how the measurement time can be considerably reduced, while greatly boosting the SNR, by targeting the leakage of specific transistors directly. This basic methodology can be applied to many different attack vectors. Effectively, every transistor exhibiting data-dependent behaviour becomes a potential target.

The initial spatial analysis is necessary to allow for at least a basic understanding of the chip's functionality and the identification of potential points of interest. In our approach we use a Si-CCD, which operates with very few impeding photons at the edge of the spectral range to which silicon is sensitive. This low cost approach requires DUT substrate thinning. However, more expensive InGaAs-cameras can also be used, which are sensitive to photons above 1 μm wavelength. In this spectral range silicon is transparent and substrate thinning therefore becomes unnecessary. It is worth noting that many modern security ICs, such as smartcards, have a far thinner substrate than common general purpose microcontrollers. As a result, many security ICs do not have to be thinned at all for semi-invasive optical backside analysis. In contrast, modern ICs generate less and less photonic emissions due to lower supply voltage. However, recent works have demonstrated that emission analysis of modern ICs is possible [18, 19], especially if NIR-sensitive detectors are employed. An example of such a device is the InGaAs/InP-APD used in this work.

In the presented attack, full key recovery was achieved for an S-Box with an uneven offset. For a fully-aligned S-Box implemented in 8-byte-rows of memory, the set of potential key candidates can still be greatly reduced to approximately 2^{48}, see Table 1. However, by performing additional temporal measurements of alternative points of leakage on the chip, the set of potential key candidates can

nevertheless be reduced to a single key candidate. It is feasible, for example, to exploit the leakage of the column decode logic and reveal the exact address of memory accesses using cross-correlation. SPEA can potentially be extended to any addressable memory. It can also easily be adopted to AES-192 and AES-256 and to alternative implementations, which use compressed tables. If additional measurements can be captured in a reasonable amount of time, the attack could even be implemented as an unknown-plaintext attack, as demonstrated for a cache timing attack by [8].

It is worth noting, that many industry standard countermeasures, in fact, do not prevent photonic emission attacks at all. Shields and meshes generally only protect against attacks from the frontside. Memory encryption and scrambling [14] may protect against probing attacks, but have no effect on the optical emissions. For SPEA, the memory access patterns would be unaffected. However, memory encryption would make the initial spatial analysis more cumbersome. It would obfuscate the values in memory, preventing the memory from being read out optically. Memory scrambling would, on the other hand, potentially make the attack far easier. Since the goal of memory scrambling is to obfuscate the layout of memory in terms of addresses, it increases the likelihood that a single S-Box element may be isolated in a row of memory. If the positions of memory are spatially obfuscated, accessing a single line would reveal all of its elements.

Nevertheless, several potential hardware and software countermeasures to our attack do exist. Specifically, hardware and software delays, masking and dummy rounds can make such attacks vastly more difficult. Any form of randomization forces longer measurement iteration times, thus greatly increasing trace acquisition times. The effects of such countermeasures are identical to the effects on power analysis, as described in [3] and could be minimized with more advanced, e.g. differential, signal analysis techniques. On the other hand, at least in this specific attack randomization techniques could be thwarted by employing APD detection gates long enough to encompass any and all randomization clock cycles. Since accesses to the S-Box occur at vastly different points in time, the resulting temporal resolution would still be sufficient to yield all S-Box accesses.

It has also been argued, that shrinking structure sizes will eventually defeat optical emission analysis and thus photonic side channels. As already mentioned, recent works show that shrinking feature sizes do not eliminate optical emission [17–19] and photonic detection techniques continue to improve rapidly. Also, in practice structure sizes only apply to the smallest structures on an IC. In every IC there are plenty data-dependent transistors that have far larger channels than that of the smallest logic on the chip. One such example is the very row driver exploited in this work, which is sized up to cope with the large capacitances of SRAM memory rows.

As a countermeasure that prohibits any optical emission attacks we would like to propose an active shield or mesh on the backside of the IC. While a single metal layer on the backside can trap all photons generated within the chip, active integrity checks prevent the removal of such a layer. Further countermeasures are under development with our partners.

6 Conclusion

Optoelectronic systems are currently the only systems capable of recovering leakage from single transistors directly. Despite this fact, the photonic side channel attacks presented so far failed to utilize both temporal and spatial information in the attacks [5, 7, 16]. By combining temporal measurements of an InGaAs/InP-APD detector with the spatial resolution of a Si-CCD, potential targets are identified quickly and easily. Also, unprecedented levels in terms of SNR and therefore information leakage over acquisition time are achieved. The optoelectronic system employed in this work outperforms many state of the art optoelectronic systems, for a fraction of the cost of the system utilized in [5]. The necessary acquisition time for the attack presented is comparable to power analysis attacks.

In our *Simple Photonic Emission Analysis* of AES, we initially evaluated emission images and subsequently performed temporal measurements. As a result we were able to mount a successful attack that utilizes both spatial and temporal information retrieved from photonic emissions. Provided the spatial information from the preliminary evaluation of emission images, we were able to focus the temporal measurements directly on a single transistor of the IC. The information leakage was so high that no additional analysis was necessary to recover all 16 bytes of the AES-128 secret key, akin to SPA and for a similarly low price in terms of equipment. The optoelectronic system employed in this work costs approximately the same as a mid-range oscilloscope, yet because it is spatially selective, offers vastly better characteristics in terms of leakage SNR. To the best of our knowledge this is the first work to combine temporal and spatial photonic side channel analysis. It demonstrates that even the leakage of a single transistor can be exploited directly to recover the full AES secret key.

Acknowledgements. The authors acknowledge support by the German Federal Ministry of Education and Research in the project PhotonDA through grant number 01IS10029A and the Helmholtz Research School on Security Technologies. Also, the authors would like to thank our project partners at NXP Semiconductors Germany for their insight and cooperation, the Semiconductor Devices research group at TU Berlin for sample preparation and our colleagues Enrico Dietz, Sven Frohmann, Collin Mulliner and Cristoph Bayer for helpful discussions and feedback.

References

1. Bascoul, G., Perdu, P., Benigni, A., Dudit, S., Celi, G., Lewis, D.: Time Resolved Imaging: From logical states to events, a new and efficient pattern matching method for VLSI analysis. Microelectronics Reliability 51(9-11), 1640–1645 (2011), `http://dx.doi.org/10.1016/j.microrel.2011.06.043`
2. Bernstein, D.: Cache-timing attacks on AES (2004), `http://cr.yp.to/papers.html\#cachetiming`

3. Clavier, C., Coron, J.S., Dabbous, N.: Differential Power Analysis in the Presence of Hardware Countermeasures. In: Koç, C., Paar, C. (eds.) Cryptographic Hardware and Embedded Systems — CHES 2000, Lecture Notes in Computer Science, vol. 1965, pp. 13–48. Springer Berlin / Heidelberg (2000), http://dx.doi.org/10.1007/3-540-44499-8_20
4. Daemen, J., Rijmen, V.: The design of Rijndael: AES – the Advanced Encryption Standard. Springer Berlin / Heidelberg (2002)
5. Di-Battista, J., Courrege, J.C., Rouzeyre, B., Torres, L., Perdu, P.: When Failure Analysis Meets Side-Channel Attacks. In: Mangard, S., Standaert, F.X. (eds.) Cryptographic Hardware and Embedded Systems — CHES 2010, Lecture Notes in Computer Science, vol. 6225, pp. 188–202. Springer Berlin/Heidelberg (2011), http://dx.doi.org/10.1007/978-3-642-15031-9_13
6. Egger, P., Grutzner, M., Burmer, C., Dudkiewicz, F.: Application of time resolved emission techniques within the failure analysis flow. Microelectronics Reliability 47(9-11), 1545–1549 (2007), http://dx.doi.org/10.1016/j.microrel.2007.07.067
7. Ferrigno, J., Hlaváč, M.: When AES blinks: introducing optical side channel. Information Security, IET 2(3), 94 –98 (2008), http://dx.doi.org/10.1049/iet-ifs:20080038
8. Gullasch, D., Bangerter, E., Krenn, S.: Cache games – bringing access-based cache attacks on aes to practice. In: Security and Privacy, 2011 IEEE Symposium on. pp. 490 –505 (2011), http://dx.doi.org/10.1109/SP.2011.22
9. Heyszl, J., Mangard, S., Heinz, B., Stumpf, F., Sigl, G.: Localized Electromagnetic Analysis of Cryptographic Implementations. In: Dunkelman, O. (ed.) Topics in Cryptology — CT-RSA 2012, Lecture Notes in Computer Science, vol. 7178, pp. 231–244. Springer Berlin / Heidelberg (2012), http://dx.doi.org/10.1007/978-3-642-27954-6_15
10. Kash, J., Tsang, J.: Dynamic internal testing of CMOS circuits using hot luminescence. Electron Device Letters, IEEE 18(7), 330 –332 (1997), http://dx.doi.org/10.1109/55.596927
11. Nedospasov, D., Schlösser, A., Seifert, J., Orlic, S.: Functional integrated circuit analysis. In: Hardware-Oriented Security and Trust (HOST), 2012 IEEE International Symposium on (2012)
12. Nohl, K., Evans, D., Starbug, S.: Reverse-engineering a cryptographic RFID tag. 17th USENIX Security Symposium pp. 185–193 (2008), http://www.usenix.org/event/sec08/tech/full_papers/nohl/nohl_html/
13. Rabaey, J.M., Chandrakasan, A.: Digital Integrated Circuits. A Design Prespective, Pearson Education, second edn. (2003)
14. Rankl, W., Effing, W.: Smart Card Handbook. Wiley, fourth edn. (2010)
15. Selmi, L., Mastrapasqua, M., Boulin, D., Bude, J., Pavesi, M., Sangiorgi, E., Pinto, M.: Verification of electron distributions in silicon by means of hot carrier luminescence measurements. Electron Devices, IEEE Transactions on 45(4), 802 –808 (1998), http://dx.doi.org/10.1109/16.662779
16. Skorobogatov, S.: Using Optical Emission Analysis for Estimating Contribution to Power Analysis. In: Fault Diagnosis and Tolerance in Cryptography (FDTC), 2009 Workshop on. pp. 111 –119 (2009), http://dx.doi.org/10.1109/FDTC.2009.39
17. Song, P., Stellari, F., Huott, B., Wagner, O., Srinivasan, U., Chan, Y., Rizzolo, R., Nam, H., Eckhardt, J., McNamara, T., Tong, C.L., Weger, A., McManus, M.: An advanced optical diagnostic technique of IBM z990 eServer microprocessor pp. 9 pp. –1235 (2005), http://dx.doi.org/10.1109/TEST.2005.1584091

18. Tosi, A., Stellari, F., Pigozzi, A., Marchesi, G., Zappa, F., Heights, Y.: A Challenge For Emission Based Testing And Diagnostics. Reliability physics pp. 595–601 (2006), http://dx.doi.org/10.1109/RELPHY.2006.251284
19. Tsang, J.C., Fischetti, M.V.: Why hot carrier emission based timing probes will work for 50 nm, 1V CMOS technologies. Microelectronics Reliability pp. 1465–1470 (2001), http://dx.doi.org/10.1016/S0026-2714(01)00194-9
20. Tsang, J.C., Kash, J.A., Vallett, D.P.: Picosecond imaging circuit analysis. IBM Journal of Research and Development 44(4), 583 –603 (2000), http://dx.doi.org/10.1147/rd.444.0583
21. Villa, S., Lacaita, A., Pacelli, a.: Photon emission from hot electrons in silicon. Physical Review B 52(15), 10993–10999 (1995), http://www.dx.doi.org/10.1103/PhysRevB.52.10993
22. Weste, N.H.E., Harris, D.: CMOS VLSI Design: A Circuits and Systems Perspective. Addison Wesley, fourth edn. (2010)

Appendix

	0	1	2	3	4	5	6	7	8	9	a	b	c	d	e	f
0	63	7c	77	7b	f2	6b	6f	c5	30	01	67	2b	fe	d7	ab	76
1	ca	82	c9	7d	fa	59	47	f0	ad	d4	a2	af	9c	a4	72	c0
2	b7	fd	93	26	36	3f	f7	cc	34	a5	e5	f1	71	d8	31	15
3	04	c7	23	c3	18	96	05	9a	07	12	80	e2	eb	27	b2	75
4	09	83	2c	1a	1b	6e	5a	a0	52	3b	d6	b3	29	e3	2f	84
5	53	d1	00	ed	20	fc	b1	5b	6a	cb	be	39	4a	4c	58	cf
6	d0	ef	aa	fb	43	4d	33	85	45	f9	02	7f	50	3c	9f	a8
7	51	a3	40	8f	92	9d	38	f5	bc	b6	da	21	10	ff	f3	d2
8	cd	0c	13	ec	5f	97	44	17	c4	a7	7e	3d	64	5d	19	73
9	60	81	4f	dc	22	2a	90	88	46	ee	b8	14	de	5e	0b	db
a	e0	32	3a	0a	49	06	24	5c	c2	d3	ac	62	91	95	e4	79
b	e7	c8	37	6d	8d	d5	4e	a9	6c	56	f4	ea	65	7a	ae	08
c	ba	78	25	2e	1c	a6	b4	c6	e8	dd	74	1f	4b	bd	8b	8a
d	70	3e	b5	66	48	03	f6	0e	61	35	57	b9	86	c1	1d	9e
e	e1	f8	98	11	69	d9	8e	94	9b	1e	87	e9	ce	55	28	df
f	8c	a1	89	0d	bf	e6	42	68	41	99	2d	0f	b0	54	bb	16

Table 2. The AES S-Box, used during `SubBytes` operation, in hexadecimal representation. The first 4 bits of the input determine the row, the last 4 bits determine the column.

	0	1	2	3	4	5	6	7
0x338	f9 (99)	fa (2d)	fb (0f)	fc (b0)	fd (54)	fe (bb)	ff (16)	
0x330	f1 (a1)	f2 (89)	f3 (0d)	f4 (bf)	f5 (e6)	f6 (42)	f7 (68)	f8 (41)
0x328	e9 (1e)	ea (87)	eb (e9)	ec (ce)	ed (55)	ee (28)	ef (df)	f0 (8c)
0x320	e1 (f8)	e2 (98)	e3 (11)	e4 (69)	e5 (d9)	e6 (8e)	e7 (94)	e8 (9b)
0x318	d9 (35)	da (57)	db (b9)	dc (86)	dd (c1)	de (1d)	df (9e)	e0 (e1)
0x310	d1 (3e)	d2 (b5)	d3 (66)	d4 (48)	d5 (03)	d6 (f6)	d7 (0e)	d8 (61)
0x308	c9 (dd)	ca (74)	cb (1f)	cc (4b)	cd (bd)	ce (8b)	cf (8a)	d0 (70)
0x300	c1 (78)	c2 (25)	c3 (2e)	c4 (1c)	c5 (a6)	c6 (b4)	c7 (c6)	c8 (e8)
0x2F8	b9 (56)	ba (f4)	bb (ea)	bc (65)	bd (7a)	be (ae)	bf (08)	c0 (ba)
0x2F0	b1 (c8)	b2 (37)	b3 (6d)	b4 (8d)	b5 (d5)	b6 (4e)	b7 (a9)	b8 (6c)
0x2E8	a9 (d3)	aa (ac)	ab (62)	ac (91)	ad (95)	ae (e4)	af (79)	b0 (e7)
0x2E0	a1 (32)	a2 (3a)	a3 (0a)	a4 (49)	a5 (06)	a6 (24)	a7 (5c)	a8 (c2)
0x2D8	99 (ee)	9a (b8)	9b (14)	9c (de)	9d (5e)	9e (0b)	9f (db)	a0 (e0)
0x2D0	91 (81)	92 (4f)	93 (dc)	94 (22)	95 (2a)	96 (90)	97 (88)	98 (46)
0x2C8	89 (a7)	8a (7e)	8b (3d)	8c (64)	8d (5d)	8e (19)	8f (73)	90 (60)
0x2C0	81 (0c)	82 (13)	83 (ec)	84 (5f)	85 (97)	86 (44)	87 (17)	88 (c4)
0x2B8	79 (b6)	7a (da)	7b (21)	7c (10)	7d (ff)	7e (f3)	7f (d2)	80 (cd)
0x2B0	71 (a3)	72 (40)	73 (8f)	74 (92)	75 (9d)	76 (38)	77 (f5)	78 (bc)
0x2A8	69 (f9)	6a (02)	6b (7f)	6c (50)	6d (3c)	6e (9f)	6f (a8)	70 (51)
0x2A0	61 (ef)	62 (aa)	63 (fb)	64 (43)	65 (4d)	66 (33)	67 (85)	68 (45)
0x298	59 (cb)	5a (be)	5b (39)	5c (4a)	5d (4c)	5e (58)	5f (cf)	60 (d0)
0x290	51 (d1)	52 (00)	53 (ed)	54 (20)	55 (fc)	56 (b1)	57 (5b)	58 (6a)
0x288	49 (3b)	4a (d6)	4b (b3)	4c (29)	4d (e3)	4e (2f)	4f (84)	50 (53)
0x280	41 (83)	42 (2c)	43 (1a)	44 (1b)	45 (6e)	46 (5a)	47 (a0)	48 (52)
0x278	39 (12)	3a (80)	3b (e2)	3c (eb)	3d (27)	3e (b2)	3f (75)	40 (09)
0x270	31 (c7)	32 (23)	33 (c3)	34 (18)	35 (96)	36 (05)	37 (9a)	38 (07)
0x268	29 (a5)	2a (e5)	2b (f1)	2c (71)	2d (d8)	2e (31)	2f (15)	30 (04)
0x260	21 (fd)	22 (93)	23 (26)	24 (36)	25 (3f)	26 (f7)	27 (cc)	28 (34)
0x258	19 (d4)	1a (a2)	1b (af)	1c (9c)	1d (a4)	1e (72)	1f (c0)	20 (b7)
0x250	11 (82)	12 (c9)	13 (7d)	14 (fa)	15 (59)	16 (47)	17 (f0)	18 (ad)
0x248	09 (01)	0a (67)	0b (2b)	0c (fe)	0d (d7)	0e (ab)	0f (76)	10 (ca)
0x240	01 (7c)	02 (77)	03 (7b)	04 (f2)	05 (6b)	06 (6f)	07 (c5)	08 (30)
0x238								00 (63)

Table 3. AES S-Box in 32×8 implementation with an offset of 7. The sum of row and column index yields the entry's index, each entry denotes the corresponding input and output value.

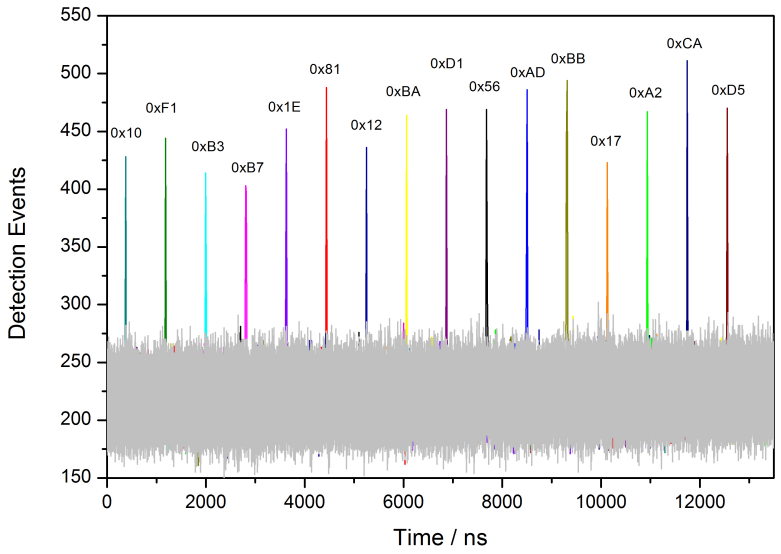

Fig. 5. Time-resolved photonic emission traces of the PoC implementation over all possible input bytes as a 2D representation. All 16 peaks, corresponding to the 16 Bytes of the AES key, are clearly identifiable and show a width of 20 ns. This corresponds to the employed detection gate width.

Offset	unresolved bits of the whole AES-192 key		unresolved bits of the whole AES-256 key	
	r = 1 or r = 33	$r \in \{2, ..., 32\}$	r = 1 or r = 33	$r \in \{2, ..., 32\}$
0	72	72	96	96
1	0	0	0	0
2	24	24	32	32
3	0	0	0	0
4	48	72	64	96
5	0	0	0	0
6	24	24	32	32
7	0	0	0	0

Table 4. Amount of unresolved bits of the complete key for AES-192 and AES-256, depending on the offset and row. For an offset of 0, there are only 32 rows. Attacking AES-192 or AES-256 requires measuring the photonic emissions of `SubBytes` during the first two rounds.