

The Temperature Side Channel and Heating Fault Attacks

Michael Hutter¹(✉) and Jörn-Marc Schmidt²

¹ Institute for Applied Information Processing and Communications (IAIK),
Graz University of Technology, Inffeldgasse 16a, 8010 Graz, Austria

Michael.Hutter@iaik.tugraz.at

² Secunet Security Networks AG,
Mergenthalerallee 77, 65760 Eschborn, Germany

joern-marc.schmidt@secunet.com

Abstract. In this paper, we present practical results of data leakages of CMOS devices via the temperature side channel—a side channel that has been widely cited in literature but not well characterized yet. We investigate the leakage of processed data by passively measuring the dissipated heat of the devices. The temperature leakage is thereby linearly correlated with the power leakage model but is limited by the physical properties of thermal conductivity and capacitance. We further present heating faults by operating the devices beyond their specified temperature ratings. The efficiency of this kind of attack is shown by a practical attack on an RSA implementation. Finally, we introduce data remanence attacks on AVR microcontrollers that exploit the Negative Bias Temperature Instability (NBTI) property of internal SRAM cells. We show how to recover parts of the internal memory and present first results on an ATmega162. The work encourages the awareness of temperature-based attacks that are known for years now but not well described in literature. It also serves as a starting point for further research investigations.

Keywords: Temperature · Side channels · Fault injection · Negative Bias Temperature Instability · AVR · Smart cards

1 Introduction

It has been known since the late 1990s that implementations of cryptographic algorithms leak information from different side channels. The first paper that demonstrates the exploitation of a side channel was published by P. C. Kocher [24] in 1996. He highlighted that implementations might provide timing characteristics that leak information of private keys. Attacks are therefore able to extract the keys by simply measuring the runtime of the implemented algorithm. Three years later, he introduced the power-consumption side channel

J.-M. Schmidt – This work was done while the author was with Graz University of Technology.

together with B. Jun and J. Jaffe in [25]. They observed that key material can be also extracted from the power consumption of cryptographic devices. Since then, many researchers started to investigate the properties of these leakages on different platforms and devices. They also proposed to exploit other powerful side channels like electromagnetic (EM) emanation [1, 13, 16, 31]. Up to now, the power and the EM side-channels have been widely established and used in academia and industry precisely because of simplicity, low-cost, and efficiency compared to other existing side channels.

Other lesser known and rather more exotic side channels (but not necessarily less powerful) are, for example, acoustic or optical emissions. Acoustic side channels have been first introduced by A. Shamir and E. Tromer [39] in 2004. They extended their work recently in 2013 [17] and provided a wide range of possible acoustic attacks, e.g., on GnuPG's RSA implementation. Heat causes mechanical stress which produces acoustic noise. This noise contains information about the power usage of CPUs and thus information about the processed data. A very related attack was also presented by D. Asonov and R. Agrawal [4] who exploited the fact that PC keyboards emanate different sounds that can be recognized at a distance. Improvements of the latter attack were reported by L. Zhuan et al. [43] in 2009. Optical emissions, in contrast, were investigated by J. Ferrigno and M. Hlaváč [15] as well as A. Schlösser et al. [35] who targeted an AES implementation. S. Skorobogatov [38] used a low-cost CCD camera to analyze the leakage of photons emitted from SRAM, EEPROM, and Flash memories.

The temperature side-channel has been often cited in literature [6, 8, 10, 11, 22, 23, 30, 36, 42]. However, most of the publications only mention the existence and possibility to exploit this channel but without providing further investigations. In particular, H. Bar-El stated in [6] that temperature attacks on smart cards are “never documented in the open literature to the author’s knowledge”. The only publication that pinpoints the existence of the temperature side-channel is due to J. Brouchier et al. [10, 11] from 2009. They showed that a cooling fan can carry information about the processed data indirectly through the dissipated temperature of a CPU. Within an experiment, they demonstrated how to extract bits from a secret password or possible RSA key (by assuming a low-frequency leak of the bits though, i.e., leaking a bit per three minutes). Furthermore, they emphasized that IP cores integrated in FPGAs might leak information to other IP cores in the system via the temperature side channel.

There are a few more papers on *active* temperature attacks, i.e., attacks that actively tamper the environmental temperature of a device (cooling or heating). Most of them demonstrate the efficiency of low-temperature attacks, e.g., as reported by S. Skorobogatov [40] and D. Samyde et al. [34] in 2002. They showed that by cooling down SRAM devices up to -50°C , they were able to *freeze* the data and to recover the content of the memory even after seconds after power-down (by exploiting the data retention property of SRAM cells). The same idea was used by T. Müller et al. [29] who presented a tool called FROST¹.

¹ FROST stands for *Forensic Recovery of Scrambled Telephones*.

The tool is able to recover the RAM content of modern Android smart phones similar to *cold boot attacks* on PCs [21]. High-temperature attacks, in contrast, have been investigated by J.-J. Quisquater and D. Samyde [32] who observed memory errors after hours of extensive heating. Similar results have been reported by S. Govindavajhala and A. Appel [19] who were able to induce errors into memories using a 50 W spotlight clip-on lamp. By heating an IBM JVM to 100 °C, they were able to inject faults with a probability of 71.4 % before their machine crashed.

In this paper, we describe a set of temperature-related attacks on common AVR and PIC 8-bit microcontrollers. There are three main contributions listed in the following:

1. We first characterize the “temperature side-channel” by presenting results of data leakages of AVR and PIC microcontrollers. We investigate the leakages and identify the linear relationship between heat radiation and circuit activity. It shows that the analyzed devices leak the Hamming weight of the processed data via the (low-frequency) temperature side channel.
2. We conduct high-temperature fault attacks on AVRs by operating the devices beyond their specified temperature ratings (>150 °C). A practical attack is shown on an RSA implementation where we successfully extracted the used private key.
3. Finally, we exploit the physical property of data remanence attacks on AVRs. By extensive heating, constant data like the private key gets burned in memory that can be recovered even after years. We identify permanent as well as transient NBTI degradation components and were able to fully recover 65 % of the entire memory of an ATmega162.

The rest of the paper is structured as follows. In Sect. 2, we characterize the temperature leakage of AVR microcontrollers. Section 3 presents results of heating fault attacks on RSA. In Sect. 4, we describe data remanence attacks. A discussion of the results is given in Sect. 5.

2 Temperature Leakage Characterization

In this section, we aim to characterize the temperature side channel by analyzing the leakage of an 8-bit ATmega162 AVR microcontroller [5]. This family of microcontrollers is widely used in embedded systems such as industrial automation, control, or in smart cards, e.g., integrated in the Funcard, ATmega Card, M2, KNOT, and Titanium programmable smart cards. First, we describe the setup to measure the temperature dissipation of these devices. Afterwards, we characterize the side channel in terms of its physical limits and possible exploitation efficiency.

2.1 The Setup to Measure the Temperature

Our setup is very similar to setups that are used to perform power-analysis attacks. Instead of measuring the power consumption of our target device, we




Fig. 1. Schematic view of the used setup to exploit the temperature side channel.

measure the dissipating temperature using a PT100 sensor element. The PT100 is a very common thermometer applied in various industry products. It measures the resistance of a platinum element having a resistance of $100\ \Omega$ at $0\ ^\circ\text{C}$. Next to the PT100, we use an AD693 amplifier (voltage to current converter) that provides a pre-calibrated Resistance Temperature Detector (RTD) interface allowing accurate measurements in the temperature range of 0 to $+104\ ^\circ\text{C}$. In that configuration the output current span is 4 to $20\ \text{mA}$ which was measured by calculating the voltage drop over a $390\ \Omega$ resistor in series to the power supply. We used a standard $1\ \text{GHz}$ digital oscilloscope for that purpose and connected it to a PC that runs Matlab for controlling the measurement process. Figure 1 shows the schematic of the used setup.

In order to allow an accurate characterization of the temperature leakage, we decided to decapsulate the chip from the rear side and to measure the temperature dissipation directly on the surface of the silicon substrate. Figure 2 shows the decapsulated ATmega162 and the PT100 touching the rear side of the chip die. The PT100 has been surrounded by a thermal-conductance paste ($4\text{--}10\ \text{W}/(\text{mK})$) to allow a stable and accurate sensing of the temperature. Note that for the targeted device we had to remove the copper plate which is located below the chip die and the plastic package as similarly done by [36, 41]. Subsequently, we polished the substrate but we did not apply any additional thinning procedures. This can be simply done by using a skew driver.

In general, silicon substrate has a good thermal conductivity which is much higher than the conductivity of the surrounding die package. So it is advantageous to decapsulate the chip but this is not necessarily required. We also performed the experiments without decapsulating the chip and measured the temperature on the surface of the package (which corresponds to a non-invasive attack). We observed that the leakage is slightly weaker but strong enough to obtain similar results. Figure 3 shows a picture of the overall setup including a controller board that is used to communicate with the ATmega162 over a serial-communication interface.

2.2 Temperature Analysis

In order to characterize the leakage of our targeted device, we measured the temperature dissipation of various processed intermediate bytes and used a long




Fig. 2. Rear-side decapsulation of an ATmega162 to provide direct contact to the silicon substrate for the PT100.




Fig. 3. Side-channel measurement setup including controller board, ATmega162, and temperature-sensing circuit.

acquisition window to evaluate the impact of thermal conductivity and capacitance. As a target operation, we made use of a MOV instruction and moved all possible values of one input byte (i.e., 256) to 24 internal registers (the remaining 8 registers are used for loop indices and other temporary data). These MOV instructions were executed in a loop where the loop index and execution duration had been configurable by the PC. The temperature dissipation was then measured for a period of 20 s where in the first 10 s only zero values were moved to the registers and in the last 10 s, the current input-byte value was written to all registers². For each input-byte value, we measured 100 traces and averaged them to reduce noise. Figure 4 zooms into the acquisition window showing the most interesting 10 s during the transient phase. It shows that in the first part of the traces, the temperature is equal for all inputs (the temperature is decreasing because of the fast acquisition runs and the higher temperature dissipation of previously measured traces). At the point when the actual input byte is written, the temperature increases depending on the Hamming weight of the processed value.

Figure 5 shows the temperature dissipation of all input-byte values at the time right before the end of the acquisition window, i.e., after about 18 s. It clearly shows that the temperature corresponds to the Hamming weight of the processed intermediate values ranging from 26.6 to about 26.8 °C. The temperature therefore linearly correlates with the power model of the device which was known and characterized using power analysis. If the circuit activity is high (meaning if there occur many bit transitions), the dissipated temperature increases. If the circuit activity is low due to less or no bit transitions, the temperature decreases. This, however, holds true not only for dynamic power consumption (caused by charging and discharging capacitances) but also for static power consumption

² We set all registers to zero before writing of new values to guarantee the transitions of all bits (avoiding Hamming-distance leaks).




Fig. 4. Slow temperature increase of all Hamming weights that are processed by the ATmega162.

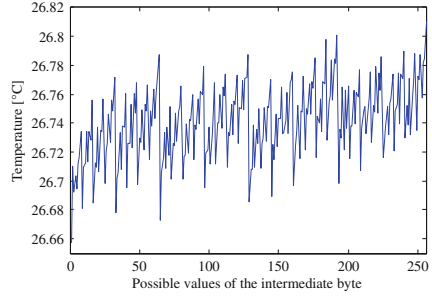


Fig. 5. The ATmega162 leaks the Hamming weight of all 256 possible intermediate values through the temperature.

(caused through sub-threshold and leakage current). Both components cause an averaged DC increase in temperature.

From Fig. 4 it is also observable that the temperature increases and decreases very slowly. While the circuit activity is constantly high or low, the temperature increases or decreases by about 0.3°C over a period of 10 s. This has its reason in the following facts. First, the temperature variation is limited by the physical property of thermal conductivity. The heat flow from the die (causing thermal power) to the sensing element can be seen as an RC network including resistance and capacitances for the junction (i.e., chip die), the package/case, heat sink, and ambient air. This RC network has the typical property that it consists of large (thermal) capacitances that causes the network to behave like a low-pass filter. The cut-off frequency is thereby very low (typically between some Hz and tens of kHz [2]). This means also that high frequency leakages ($>1\text{ MHz}$), which usually appear in CMOS devices, will not be easily exploitable from that side channel because the information will be largely filtered by the RC network. The measured temperature signal at the sensor element therefore contains all superimposed and integrated signal components of the power consumption. As a second reason, the used temperature sensors have a certain response time and acquisition resolution. In our experiment, we used a PT100 that has a thermal response time of 100 ms and a resolution of 0.01°C .

It shows that the temperature side-channel has a very low bandwidth limiting practical attacks. In the following, we discuss possible attack scenarios for low-frequency temperature leakages:

1. An attack exploiting the temperature side-channel is possible in case the leakage of the data is present over a period of several milliseconds or seconds. If an application repeatedly checks a password, for example in a loop, enough information is available even in limited frequency bands that allows low-bandwidth attacks [10, 11].
2. Many RSA implementations involve operations that take a long time, e.g., modular exponentiations. These operations create signals in a low frequency

band that can be revealed by low-bandwidth acoustic attacks as recently shown by D. Genkin, A. Shamir, and E. Tromer [17]. These low-bandwidth signals can be also extracted from the temperature side channel.

3. A very powerful attack which is not well investigated yet is the exploitation of static power-consumption leakages. Most of side-channel based analyses are exploiting the dynamic power-consumption which is the main contributor to the total power consumption of electronic devices. With shrinking CMOS technology, static leakages become more significant. Temperature attacks exploiting the static power-consumption benefit from less strict timing constraints because the leakage is *statically* available over an infinite period of time. A. Moradi recently demonstrated successful power-analysis attacks exploiting the static leakage of FPGAs in [28]. Other works also characterized and exploited static leakages of CMOS devices, for example, in [18, 26].

3 Exploiting Heating Faults on AVR_s

In this section, we intentionally operate a target device beyond its maximum temperature ratings in order to produce exploitable faults due to extensive heating. Each electronic device specifies a certain temperature range where the correct operation is guaranteed by the manufacturer. If these limits are exceeded by external influences, data might get modified that is stored in memories or processed by the CPU. Faulty cryptographic operations can then be exploited in attacks to reveal the secret key [6, 7, 23, 36].

In the following experiments, we used the same ATmega162 as used in the previous experiments. To prove the practicability of our attack, we implemented RSA, induced heating faults, and successfully extracted the private key used during encryption of data, cf. Bellcore attack [9].

Setup and RSA Implementation. We used a low-cost laboratory heating plate from Schott instruments (SLK 1) to heat-up and induce faults in an ATmega162. The microcontroller has been placed directly on top of the hot-plate surface, lying top-side down to allow a good heat transfer. The temperature of the internal IC has then been measured by calculating the mean of two PT100 sensors to be more accurate. One PT100 has been placed on the rear side of the ATmega162, the other PT100 has been placed directly on the surface of the heating plate. Both PT100 are connected to an oscilloscope similar to the setup described in the previous section. Figure 6 shows the setup.

We connected and used only six mandatory pins of the ATmega162: power supply (VCC and ground), serial communication (RX and TX), clock signal, and reset. For these connections, we used exposed wires to avoid any contact to the hot plate and the melting of solder³ during long-time heat exposure. As a controlling device, we used an FPGA board (Spartan-3) that is connected with the measurement PC.

³ The temperature melting point of Sn63/Pb37 lead solder, which is commonly used for electrical soldering, is 456 K (183 °C).




Fig. 6. Heating plate with two PT100 sensors measuring the rear-side and front-side temperature of an ATmega162.




Fig. 7. Distribution of fault occurrence between 150 and 160 °C. Mean fault-induction temperature is 154.4 °C.

We decided to target an RSA implementation that implements the Chinese Remainder Theorem (CRT). This attack is very simple and well documented in literature [6, 9, 36]. Only one single fault during the computation can reveal the secret RSA primes p and q . The attack and evaluation of faulty computations has been performed using Sage [33]. A description of the attack is given in Appendix A.1.

3.1 Heating-up the Target

The used heating plate provides ten heating stages that can be adjusted manually going up to 1,000 K, i.e., about 727 °C. By manual adjustment, we were able to heat up the device under test up to 150 °C within about 5 min.

In a first observation, we identified that the ATmega162 does not respond to requests anymore if the heating temperature is higher than around 160 °C. Note that this is much higher than the *operating temperature* maximum ratings given in the device specification [5], i.e., between -55 °C and $+125$ °C and it is approximately as high as the given maximum *storage temperature* rating of the device which is 150 °C. As a key observation, we identified that the device starts producing faults between a certain heating window of about 152 and 158 °C. During this window, the probability is high that the device outputs an incorrect result because of an induced fault during the computation of the implemented algorithm.

In order to quantify this behavior, we performed several measurements on that device taking about 70 min. We performed an RSA decryption operation every 650 ms and kept the heating temperature between 150 and 160 °C by manually adjusting the heating plate. As a result, we got 100 faults where 31 of this set have been exploitable, i.e., the attack revealed one of the used prime moduli. In the other cases, the fault was induced during I/O communication or other parts of the computation such that the difference of the faulty and a correct signature output was coprime to the RSA modulus and thus was not factorizable. In addition to these outcomes, we identified that 16 out of the 31 faulty

computations, revealed the prime modulus p and 15 revealed the prime modulus q . Hence, the probability that a fault reveals p or q is expected to be about 50%—a result that was expected since both p and q have the nearly the same size and the modular exponentiation with these primes take the same amount of time and therefore provide the same fault-induction window. Furthermore, 23 out of the 31 faulty computations have been unique and different whereas 7 faults have been repeated, i.e., the RSA decryption yielded an output that was already obtained in a previous measurement. For the latter case, we can therefore assume that either some internal memory locations or internal logic parts are more sensitive to heating than others, thus causing the same RSA outputs. Figure 7 shows the frequency of the induced faults. Most of the faults occurred between 152 and 158 °C having a mean fault-induction temperature of about 154.4 °C.

We made the same experiment also using other ATmega162 devices (new once) in order to verify our results. First, it showed that the mean fault-induction temperature for each device is slightly different and varies per device. Second, it showed that the number of faults is different per device (e.g., we got 182 faults within 30 min for another device) which is likely because the temperature has to be adjusted manually in our setup and varies per measurement. But the attack succeeded for all devices within less than 30 min.

4 Data Remanence Attacks on AVR_s

In this section, we characterize the property of data remanence effects of the internal SRAM cells of an ATmega162. It is known that data which is stored in the same location after each power-up of a device (such as a secret key that is loaded from program memory/flash to RAM) leaves a permanent mark that can be recovered as, for example, detailed by P. Gutmann in [20]. He observed that data that is stored in SRAM or DRAM for a long period of time remembers the value when powered-up again even after years. This effect has been practically exploited in an attack by R. Anderson and M. Kuhn [3] who were able to recover 90–95 % of a DES master key used by an old bank security module from the late 1980s. C. Cakir et al. [12] have characterized the data remanence effect on newer 65 nm CMOS RAMs which was considered to be not that efficient because of the newer SRAM structures. However, they were able to recover about 18 % of the entire SRAM content (in fact, 82 % have been recovered correctly out of 22 % predictable bits).

The SRAM data remanence effect can be explained as follows. SRAM cells that are exposed to extensive heating are subject to accelerated aging where internal transistor parameters get changed. This effect is known under Negative Bias Temperature Instability (NBTI) and has been first observed in the late 1960s. Since then, many researchers identified that this effect decreases parameters such as speed, drive current, and noise margins. In fact, NBTI occurs when transistors are stressed with negative gate voltages at elevated temperatures, e.g., during burn-in stress. Then, the (absolute) threshold voltages of the

transistors increase which change the preferred power-up state of SRAM cells. Thus, transistors get “weaker” and tend to a certain bit value after power-up. Note that NBTI has been basically observed for both PMOS and NMOS transistors while PMOS transistors are more affected [37].

The study of M. Ershov et al. [14] showed that there exist two NBTI degradation components: one component that remains after burn-in stress (permanent damage), and another component (transient damage) that recovers within a certain amount of time (seconds up to days). In the following, we identify these two degradation components in practical experiments on the ATmega162. We used the same setup as it was used in the previous section to heat up the device.

Before we started the experiments, we determined the preferred power-up values of the ATmega162 in order to evaluate the effect of burn-in stress. A small program was written that allows reading and writing of SRAM content over the serial interface. A byte array of 6,144 bits (out of the available 8,192 bits) was used for testing. After we disconnected all I/O connections from the device for some milliseconds⁴, which has been accomplished with our flexible FPGA controller board, we read out the SRAM content. This has been done 100 times to average noise. It showed that 3,101 bits (50.47 %) are powered-up to the value *one* and 3,043 bits (49.53 %) are set to *zero* on average, i.e., there was almost no bias in the distributions.

4.1 Permanent Data Remanence Effects After Burn-In Stress

We programmed the internal SRAM of the ATmega162 with normally distributed random data. We set 3,072 bits to 0 and 3,072 bits to 1 at random locations⁵. A first test after a power-up reset showed that the probability of guessing the bits correctly was around 50 % as expected. After this test, we exposed the device to extensive burn-in stress to accelerate aging effects. The stress conditions were a high ambient temperature of about 100 °C and an over-voltage supply of 5.5 Volts. We applied the stress over a period of 36 h and read out the power-up SRAM content afterwards. The read out values are then compared with the initial values.

It showed that the number of bits that are one or zero got biased due to NBTI degradation. There were 3,210 bits (52.24 %) set to one and 2,934 bits (47.75 %) set to zero after the burn-in stress. Furthermore, 919 bits of the total memory changed their state (i.e., 15 %): 405 bits moved from 1 to 0 and 514 bits moved from 0 to 1. From these 405 bits there were 393 bits that moved to the correct value zero (97.04 %) and 489 from 514 bits moved to the correct value one (95.14 %).

⁴ We disconnected not only the power supply but also the RS232 interface and the clock signal to guarantee that the device (and SRAM respectively) is completely unconnected and not powered by I/O interfaces. Note also that we used hardware relays to actually disconnect all connections.

⁵ We do not assume the knowledge of “preferred power-up values” before burn-in stress to guarantee a realistic attacking scenario.

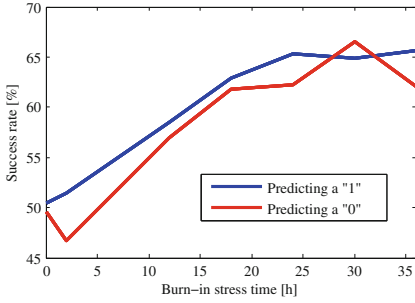


Fig. 8. Probability of predicting a SRAM bit correctly increases from 50 % to about 65 % after burn-in stress.




Fig. 9. Bit-value distribution during heating and cooling (transient data remanence effect).

Note that in our analysis, we were only able to identify half of the *unstable* SRAM cells. This is because from the set of unstable SRAM cells, statistically only half of the bits changed to the correct value while the other half already held the correct value. So, in our experiment, we can assume about 30 % of unstable SRAM cells and 70 % of cells that are *stable*, i.e., they did not change at least during our burn-in stress of 36 h. Now, by guessing 50 % of all stable cells correctly and by assuming a high probability of guessing all unstable bits correctly, we are able to successfully predict 65 % of the entire SRAM memory. This nicely corresponds to our practical results where we achieved a success rate of about 63 %, i.e., 3,842 bits have been predicted correctly as shown in Fig. 8. It shows that there is a huge increase during the first 20 h. After that, the probability keeps nearly constant. We also performed a measurement one week after the burn-in stress and identified no changes in the success rate.

4.2 Transient Data Remanence Effects During Burn-In Stress

We also characterized the transient data remanence effect during the burn-in stress process. For this experiment, we used a new ATmega162 and read out the SRAM content every 4 s while the device was heated up. Figure 9 shows the bit-value distribution during the burn-in stress and right after it. In the first 115 s, the device is heated up to 170 °C. After that, the heating plate was turned off to cool down the device. Between 105 and 125 s, the device was overheated which produced a significant jump in the probability distribution.

The experiment showed that the effect of NBTI degradation of PMOS and NMOS transistors in SRAM cells is also observable in the transient heating phase. The heating temporarily enforces the transistors to change the state. Note that this effect is only transient and the cells regenerate after cooling. The number of zero values increases during heating while the number of ones decreases accordingly. In our experiment, 820 bits changed their state to either 0 or 1. From these 820 bits, 257 bits were the same that also changed during

the long-term burn-in stress, i.e., we could identify 31.3% of all bits that are apparently unstable.

How to Exploit NBTI Degradation? Many implementations of cryptographic algorithms store the secret key in non-volatile memory and load it into SRAM when needed. This key is then loaded always at the same memory location in SRAM because the program code or hardware implementation is usually given and not changeable. Hence, the key value gets “burned” into SRAM within a period of time, e.g., some weeks or months. The key can then be extracted by the following data-remanence attacks:

1. If an attacker is in possession of several implementations that store the same secret key, the attacker can recover parts of the internal SRAM of each implementation and can then average the obtained results to reveal all bits of the memory with high probability. This attack assumes that each implementation reveals the content of different SRAM cells.
2. If an attacker is in possession of only one implementation (or several implementations using different keys), he/she can first apply a burn-in stress test over several hours in order to artificially accelerate aging. As a second step, he/she can read out the preferred SRAM content. In order to identify useful bits, the attacker can mount a transient remanence attack to reveal *unstable* SRAM cells which potentially changed their state during the burn-in stress and which contain useful information about the secret key. All other cells are then considered as stable in this attack and provide no useful information. Finally, all recovered bits are used in partial-key recovery attacks.

5 Discussion and Further Research Suggestions

In the following, we discuss further research questions arisen by this work:

- The presented attacks have been performed on microcontrollers that were fabricated in rather old process technologies. Further research has to be done to evaluate the impact of thermal attacks on ICs with newer CMOS technologies.
- The temperature side channel provides a low-bandwidth characteristic. High-frequency leakages (e.g., containing data-dependent signals in the MHz or even GHz scale) can therefore not be directly exploited due to the low-pass filter characteristics. However, there exist implementations that use long operations like exponentiations in RSA that can create low-frequency signatures that are however exploitable (as demonstrated in [17, 39] using acoustic signals in the kHz range).
- In order to validate our measurement setup, we also performed the measurements using a $0\ \Omega$ resistor instead of a PT100 element. This is done because it naturally raises the question if the measured data really corresponds to the dissipated temperature or if it is caused by other side-channel sources, e.g., EM modulated signals. Using the resistor, however, we were not able to identify any data-dependent signals so that we can exclude any signal interferences or the coupling of EM signals.

- We also characterized the temperature leakage of an 8-bit PIC16F84 microcontroller from Microchip Technology. We obtained similar results and could identify temperature-dependent processing of data. Details are given in Appendix A.2.
- Heat penetrates through different materials. Thus, the thermal conductivity of CMOS devices might be exploited in attacks where the heat conducts through EM shielding countermeasures (metal plates, mesh of power lines, sensors, etc.).
- *How does the temperature affect power-analysis attacks?* This question has been answered by A. Vijaykumar [42] in her master’s thesis. She evaluated temperature variation effects on Differential Power Analysis (DPA) attacks. By targeting KeeLoq and DES, she showed that the efficiency of DPA decreases with increased temperature due to decreasing power variances.
- By heating or cooling CMOS devices, the characteristics not only of memory but also of logic changes. Thermal attacks might even be used to circumvent countermeasure implementations, e.g., by increasing/decreasing the threshold voltages of watchdog implementations.
- The suitability of temperature attacks that indirectly exploit the leakage of *static* power consumption has to be investigated in future. Static power consumption is becoming more and more important as CMOS technology is shrinking. A static leak of an intermediate value of a cryptographic implementation creates a DC offset signal in the baseband that can be exploited in attacks that are provided with very low-bandwidth leakages from side channels such as the temperature or acoustic sound.

Acknowledgements. The work has been supported by the European Commission through the ICT program under contract ICT-SEC-2009-5-258754 (Tamper Resistant Sensor Node - TAMPRES), by the Austrian Science Fund (FWF) under the grant number TRP251-N23 (Realizing a Secure Internet of Things - ReSIT), and the European Cooperation in Science and Technology (COST) Action IC1204 (Trustworthy Manufacturing and Utilization of Secure Devices - TRUDEVICE).

A Appendix

A.1 Attacking CRT-RSA Using Faults

In the following, we consider an implementation of an RSA decryption that uses the Chinese Remainder Theorem (CRT) to speed up the computation. In our scenario, an adversary is able to supply the card with an input that is encrypted using textbook RSA and receives the decrypted message from the card. Further, the adversary is able to disturb the computation of this decryption and receives the result of this faulted computation. In order to describe how an adversary can benefit from this scenario to factor the modulus and thus compute the secret decryption key, we denote $n = pq$ an RSA modulus, where p and q are two large prime numbers. Let d be the private key and $e = d^{-1} \bmod \varphi(n)$ the corresponding public exponent. Furthermore, $z = \text{CRT}(x, y)$ denotes the CRT

recombination of the value $z \in \mathbf{Z}_n$ from values x, y of the subgroups \mathbf{Z}_p and \mathbf{Z}_q where

$$\text{CRT}(x, y) = xc_p + yc_q \bmod n$$

with $c_p = q(q^{-1} \bmod p)$ and $c_q = p(p^{-1} \bmod q)$ [27].

The usage of the CRT in this scenario allows computing two exponentiations in smaller sub-groups compared to a single exponentiations modulo n :

$$S \equiv \text{CRT}((m^d \bmod p), (m^d \bmod q)) \bmod n.$$

The first fault attack that takes advantage of injecting a random fault Δ in this scenario was presented by Boneh et al. [9]. The fault Δ causes the device to output a value \tilde{S} instead of S :

$$\begin{aligned} \tilde{S} &\equiv \text{CRT}((m \bmod p)^d, (m \bmod q)^d + \Delta) \bmod n \\ &\equiv m^d + \Delta p(p^{-1} \bmod q) \bmod n. \end{aligned}$$

If an adversary gets hold of both a faulty \tilde{S} and a correct signature S , the modulus n can be easily factorized by calculating $p = \gcd(\tilde{S} - S, n)$.

A.2 Temperature Leakage of a PIC16F84

We also investigated the leakage of a PIC16F84 microcontroller. We used the same measurement setup as described in Sect. 2 and measured the temperature on the decapsulated rear-side of the chip using a PT100 element. Instead of a MOV operation, we target an ADD instruction that adds either 0x00 or 0xFF to all internal registers that are previously initialized with zero. We measured 500 traces and averaged them to reduce noise.

Figure 10 shows the result where a zero value was written continuously over a period of 10 s. The value 0xFF is written afterwards for another 10 s. It shows an increase of temperature in the second half of the acquisition window. No leakage occurs in the first half of the trace. In Fig. 11, the result is shown when 0xFF is




Fig. 10. Leakage of 0xFF in the second half of the acquisition window. No leakage during the first 10 s.




Fig. 11. Leakage of 0xFF in the first half of the acquisition window. The mean temperature decreases afterwards.

written during the first 10s, and zero is written afterwards. There, it shows that the temperature slowly increases, similarly to the second half of Fig. 10. After 10s, the temperature is decreasing again.

References

1. Agrawal, D., Archambeault, B., Rao, J.R., Rohatgi, P.: The EM side-channel(s). In: Kaliski Jr, B.S., Koç, Ç.K., Paar, C. (eds.) CHES 2002. LNCS, vol. 2523, pp. 29–45. Springer, Heidelberg (2003)
2. Altet, J., Rubio, A., Schaub, E., Dilhaire, S., Claeys, W.: Thermal coupling in integrated circuits: application to thermal testing. *IEEE J. Solid-State Circ.* **36**(1), 81–91 (2001)
3. Anderson, R.J., Kuhn, M.G.: Low cost attacks on tamper resistant devices. In: Christianson, B., Lomas, M., Crispo, B., Roe, M. (eds.) Security Protocols 1997. LNCS, vol. 1361, pp. 125–136. Springer, Heidelberg (1998)
4. Asonov, D., Agrawal, R.: Keyboard acoustic emanations. In: IEEE Symposium on Security and Privacy, pp. 3–11 (2004)
5. Atmel Corporation.: ATmega 162/v Datasheet (2003)
6. Bar-El, H., Choukri, H., Naccache, D., Tunstall, M., Whelan, C.: The Sorcerer’s apprentice guide to fault attacks. *Cryptology ePrint Archive*. Report 2004/100 (2004). <http://eprint.iacr.org/>
7. Barenghi, A., Bertoni, G., Parrinello, E., Pelosi, G.: Low voltage fault attacks on the RSA cryptosystem. In: Workshop on Fault Diagnosis and Tolerance in Cryptography - FDTC 2009, pp. 23–31, Lausanne, Switzerland, 2009. Proceedings (2009)
8. Barenghi, A., Breveglieri, L., Koren, I., Naccache, D.: Fault injection attacks on cryptographic devices: theory, practice and countermeasures. *Proc. IEEE* **100**(11), 3056–3076 (2012)
9. Boneh, D., DeMillo, R.A., Lipton, R.J.: On the importance of checking cryptographic protocols for faults (extended abstract). In: Fumy, W. (ed.) EUROCRYPT 1997. LNCS, vol. 1233, pp. 37–51. Springer, Heidelberg (1997)
10. Bouchier, J., Dabbous, N., Kean, T., Marsh, C., Naccache, D.: Thermocommunication. *ePrint* (2009)
11. Bouchier, J., Kean, T., Marsh, C., Naccache, D.: Temperature attacks. *IEEE Secur. Priv.* **7**(2), 79–82 (2009)
12. Cakir, C., Bhargava, M., Mai, K.: 6 T SRAM and 3 T DRAM data retention and remanence characterization in 65 nm bulk CMOS. In: Custom Integrated Circuits Conference - CICC 2012, pp. 1–4, San Jose, USA, 9–12 Sept 2012
13. Carluccio, D., Lemke, K., Paar, C.: Electromagnetic side channel analysis of a contactless smart card: first results. In: Oswald, E. (ed.) Workshop on RFID and Lightweight Crypto (RFIDSec05), pp. 44–51, Graz, Austria, 13–15 July 2005
14. Ershov, M., Saxena, S., Karbasi, H., Winters, S., Minehane, S., Babcock, J., Lindley, R., Clifton, P., Redford, M., Shibkov, A.: Dynamic recovery of negative bias temperature instability in p-type metal-oxide-semiconductor field-effect transistors. *Appl. Phys. Lett.* **83**(8), 1647–1649 (2003)
15. Ferrigno, J., Hlaváč, M.: When AES blinks: introducing optical side channel. *IET Inf. Secur.* **2**(3), 94–98 (2008)
16. Gandolfi, K., Mourtel, C., Olivier, F.: Electromagnetic analysis: concrete results. In: Koç, Ç.K., Naccache, D., Paar, C. (eds.) CHES 2001. LNCS, vol. 2162, pp. 251–261. Springer, Heidelberg (2001)

17. Genkin, D., Shamir, A., Tromer, E.: RSA key extraction via low-bandwidth acoustic cryptanalysis. ePrint, Dec 2013
18. Giogetti, J., Scotti, G., Simonetti, A., Trifiletti, A.: Analysis of data dependence of leakage current in CMOS cryptographic hardware. In: Proceedings of the 17th ACM Great Lakes Symposium on VLSI, pp. 78–83, Stresa-Lago Maggiore, Italy. ACM, 11–13 Mar 2007
19. Govindavajhala, S., Appel, A.W.: Using memory errors to attack a virtual machine. In: Proceedings of the 2003 IEEE Symposium on Security and Privacy, pp. 154–165 (2003)
20. Gutmann, P.: Data remanence in semiconductor devices. In : USENIX 2001 - Proceedings of the 10th Conference on USENIX Security Symposium, Washington, DC, USA, Berkeley, CA, USA, 2001. USENIX Association, 13–17 Aug 2001
21. Halderman, J., Schoen, S.D., Heninger, N., Clarkson, W., Paul, W., Calandrino, J.A., Feldman, A.J., Appelbaum, J., Felten, E.W.: Lest we remember: cold boot attacks on encryption keys. In: 17th USENIX Security Symposium, pp. 45–60, San Jose, CA, July 2008
22. Hutter, M., Schmidt, J.-M., Plos, T.: RFID and its vulnerability to faults. In: Oswald, E., Rohatgi, P. (eds.) CHES 2008. LNCS, vol. 5154, pp. 363–379. Springer, Heidelberg (2008)
23. Karaklajić, D., Schmidt, J.-M., Verbauwhede, I.: Hardware designers guide to fault attacks. In: IEEE Transactions on Very Large Scale Integration (VLSI) Systems, pp. 1–12 (2012)
24. Kocher, P.C.: Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: Koblitz, N. (ed.) CRYPTO 1996. LNCS, vol. 1109, pp. 104–113. Springer, Heidelberg (1996)
25. Kocher, P.C., Jaffe, J., Jun, B.: Differential power analysis. In: Wiener, M. (ed.) CRYPTO 1999. LNCS, vol. 1666, pp. 388–397. Springer, Heidelberg (1999)
26. Lin, L., Burleson, W.: Leakage-based differential power analysis (LDPA) on sub-90 nm CMOS cryptosystems. In: ISCAS 2008 - IEEE International Symposium on Circuits and Systems, pp. 252–255, Seattle, USA, 18–21 May 2008
27. Menezes, A.J., van Oorschot, P.C., Vanstone, S.A.: Handbook of Applied Cryptography. Series on Discrete Mathematics and Its Applications. CRC Press, Boca Raton (1997). ISBN 0-8493-8523-7. <http://www.cacr.math.uwaterloo.ca/hac/>
28. Moradi, A.: Side-channel leakage through static power - should we care about in practice? ePrint, Jan 2014
29. Müller, T., Spreitzenbarth, M.: FROST. In: Jacobson, M., Locasto, M., Mohassel, P., Safavi-Naini, R. (eds.) ACNS 2013. LNCS, vol. 7954, pp. 373–388. Springer, Heidelberg (2013)
30. Otto, M.: Fault attacks and countermeasures. Ph.D. thesis, Universität Paderborn (2005)
31. Quisquater, J.-J., Samyde, D.: A new tool for non-intrusive analysis of smart cards based on electro-magnetic emissions, the SEMA and DEMA methods. Presented at the rump session of EUROCRYPT 2000 (2000)
32. Quisquater, J.-J., Samyde, D.: Eddy current for magnetic analysis with active sensor. In: Proceedings of the 3rd International Conference on Research in SmartCards (E-Smart'02), pp. 185–194, Nice, France. UCL, Sept 2002
33. SageMath.: Sage: open source mathematics software system (2013). <http://sagemath.org>
34. Samyde, D., Skorobogatov, S.P., Anderson, R.J., Quisquater, J.-J.: On a new way to read data from memory. In: IEEE Security in Storage Workshop (SISW02), pp. 65–69. IEEE Computer Society (2002)

35. Schlösser, A., Nedospasov, D., Krämer, J., Orlic, S., Seifert, J.-P.: Simple photonic emission analysis of AES. In: Prouff, E., Schaumont, P. (eds.) CHES 2012. LNCS, vol. 7428, pp. 41–57. Springer, Heidelberg (2012)
36. Schmidt, J.-M., Hutter, M.: Optical and EM fault-attacks on CRT-based RSA: concrete results. In: Posch, K.C., Wolkerstorfer, J. (eds.) Proceedings of Austrochip 2007, pp. 61–67, Graz, Austria. Verlag der Technischen Universität Graz, 11 Oct 2007. ISBN 978-3-902465-87-0
37. Schroder, D.K.: Negative bias temperature instability: what do we understand? *J. Microelectr. Reliab.* **47**(6), 841–852 (2006)
38. Skorobogatov, S.: Using optical emission analysis for estimating contribution to power consumption. In: Fault Diagnosis and Tolerance in Cryptography (FDTC) (2009)
39. Shamir, A., Tromer, E.: Acoustic cryptanalysis - on nosy people and noisy machines. <http://www.wisdom.weizmann.ac.il/~tromer/acoustic/>. Preliminary proof-of-concept presentation
40. Skorobogatov, S.: Low temperature data remanence in static RAM. Technical report, University of Cambridge Computer Laboratory, June 2002
41. Skorobogatov, S.P.: Semi-invasive attacks - a new approach to hardware security analysis. Ph.D. thesis, University of Cambridge - Computer Laboratory (2005). <http://www.cl.cam.ac.uk/TechReports/>
42. Vijaykumar, A.: DPA resistance of cryptographic circuits considering temperature and process variations. Master's thesis, University of Cincinnati, Engineering and Applied Science: Computer Engineering, July 2012
43. Zhuang, L., Zhou, F., Tyga, J.D.: Keyboard acoustic emanations revisited. *ACM Trans. Inf. Syst. Secur.* **13**(1), 373–382 (2009)