Appears in
Proceedings of the 11th European Conference on Antennas and Propagation (EuCAP), April 2017.

Path Loss Prediction for Electromagnetic Side-Channel Signals

Alenka Zajic, Milos Prvulovic, and Derrick Chu
Georgia Institute of Technology, Atlanta, GA 30332 USA

***Abstract*— This paper investigates propagation mechanisms that EM side-channel signals experience at different frequencies and proposes models for near-field and far-field propagation of side-channel signals. The near-field propagation is modelled as a field created by an electric monopole (Hertzian dipole) and a magnetic dipole, where the received power is collected using only magnetic components of the EM field. This model resulted in excellent match with measured data. Furthermore, this paper investigates unintentionally modulated side-channel signals. The propagation of EM side-channel signals was modelled using free-space propagation model which resulted in excellent match with measured data. In both cases we have observed that signal can be received at several meters from the side-channel source. The proposed models are the first step in understanding propagation mechanisms of EM side-channel signals and how to predict the distance at which they can be received.**

I. Introduction

Side-channel attacks use information gained or leaked from the physical implementation of a system to extract sensitive information, such as cryptographic keys. This includes measuring the time and power that system needs to perform computation (e.g. encryption) [1], [2], [3], leaked electromagnetic (EM) radiation [4], [5], sounds (acoustic "noise") [6], [7], and temperature variation produced during computation e.g. [8].

Among analog-signal attacks, the EM-based ones are particularly interesting: they can be mounted from a distance and can exploit sub-channels (at different frequencies and modulations) that leak somewhat different information [9]. Cryptographic EM side-channel attacks are more difficult for systems with high clock rates (laptop/desktop/server) [10], [11], [12], mainly because attacks often require signal sampling rates that match or exceed the victim device's clock rate. Even then, EM covert channel transmission has been demonstrated for such systems even in the presence of significant countermeasures [11], and side-channel cryptographic attacks have also been demonstrated [10].

Recent work also shows that analog signals can be used to learn more about a program's behavior. For instance, current (power) fluctuations were used to identify webpages during browsing [13] and even find anomalies in software activity [14], [15]. Our recent results show that differences between different instructions can be measured in EM analog signals across different devices (e.g. desktops, laptops, FPGAs) [12], [16], [17]. We can also identify which aspects of program activity modulate which EM-emanated signals [18]) and we have shown that the (approximate) number of times a specific loop or path in the program is exercised can be found by recording their EM signals in a known (training) execution and then counting matches in the EM signal collected during another (profiling) program execution [19], [20]. All these results indicate that execution monitoring through analog (and in particular EM) signals is possible.

This work has been supported, in part, by NSF grants 1563991 and 1318934, AFOSR grant FA9550-14-1-0223, and DARPA LADS contract FA8650-16-C-7620. The views and findings in this paper are those of the authors and do not necessarily reflect the views of NSF, AFOSR, or DARPA.

Often, EM side-channels attacks are not perceived as a serious security threat because of assumption that an attacker has to be very close to a device of interest. On the other hand, our work in [11] has shown that EM side-channel signals can be received several meters away, even through a wall. A natural question is: how far EM side-channel emanations can propagate? The goal of this paper is to investigate propagation mechanisms that EM side-channel signals experience at different frequencies and to model received power as a function of a distance. To achieve this goal, we use SAVAT [12], an exposure quantification metric, which does not present or imply a specific side-channel attack, but instead provides direct quantitative feedback to programmers and hardware designers about which instructions (or combination of instructions) have the greatest potential to create side-channel vulnerabilities. Additionally, we have shown in [11] that this method can be used to transmit covert-channel information.

We first measure signal energy directly created by SAVAT execution at several different distances. This signal is created at relatively low frequencies (e.g. hundreds of kHz) and all measurements are performed in the near-field. We have proposed to model the near-field side-channel field as a field created by an electric monopole (Hertzian dipole) and a magnetic dipole, where we can receive only magnetic components of the EM field. This model resulted in excellent match with measured data. Furthermore, we have investigated unintentionally modulated SAVAT signals. In this case, we have observed that the propagation of EM side-channel signals follows well free space propagation model. This is not a surprising result because at modulated frequencies, (e.g. 1GHz), received signals are in the antenna far-field. The proposed models are the first step in understanding propagation mechanisms of EM side-channel signals and how to predict the distance at which they can be received.

The rest of this paper reviews an exposure quantification metric SAVAT (Section II), reviews unintentionally modulated side-channel signals (Section III), details our path loss modeling and measurements (Section IV) and presents conclusions (Section V).

II. A Method for Measuring Very Low EM Side-Channel Energy from Processor Instructions

In [12] we have proposed a method for measuring EM side-channel energy from processor instructions. Our analysis assumes that an attacker has access to a program's source or executable code, and can observe EM emanations from the victim's system while this program is running. The attacker attempts to extract sensitive information by recording EM emanations, using them to infer which instructions are executed, and then infers sensitive data from knowledge of the executed instructions. For example, suppose an attacker can isolate (in the recorded EM signal) the time offset of a single branch instruction in the program, and suppose that this branch instruction is taken or not taken depending on a sensitive data bit. The attacker observes the side channel signal for a time period immediately following the branch. The taken/not-taken outcome of most branches results in executing different instructions after that branch, which would result in different signals, this signal difference may enable attacks (such as DPA [21]) that determine whether the branch was taken (and therefore the value of the sensitive bit).

The most direct approach to quantifying the EM emanations from side-channel signal created by executing instruction A vs. executing instruction B is to measure the EM emanations while instruction A is active, measure the EM emanations while instruction B is active, and then take the difference between these two signals. This approach is impractical in high performance systems for several reasons. First, equipment capable of measuring the low amplitude $a(t)$ and $b(t)$ signals at greater than 10G samples/sec (as required to test a processor using a GHz clock) is prohibitively expensive or non-existent. Second, complex processors heavily optimize the scheduling and execution of instructions, so determining the times where the test instructions A or B are actually active would be problematic. Third, some other instructions must be present around A and B to make the measurement practical (to trigger the measurement, setup the registers and memory used by instructions A and B, etc.), and so noise and other unrelated components of the received signal obfuscate the signal components created by the A and B instructions themselves.

To overcome these problems, we force the system to generate controllable emanations by executing the A and B instructions in a way that minimizes the effect of all other unrelated system activities, and then measure the leaked side-channel energy. We produce these controllable emanations by choosing a repetition period T_{alt} and then create a benchmark containing a `for` loop such that the first half of the loop does many repetitions of activity A and the second half does many repetitions of activity B. The microbenchmark in Figure 1 implements this idea by executing A and B instructions n_{inst} times (denoted as `n_inst` in Figure 1) in each iteration of the outer loop. Lines 2 through 7 execute n_{inst} instances of the A instruction, and then lines 8 through 13 execute the same number of instances of the B instruction. Thus lines 2 through 13 represent one A/B alternation, and this alternation is repeated (line 1) until the measurement of the side-channel signal is complete. It is critical to note that the value of T_{alt} is controlled directly by varying n_{inst}. For example, increasing n_{inst} increases the time required to execute one iteration of the outer loop (T_{alt}). The value of T_{alt} can be directly measured using counters available through processor instructions (e.g. the x86 `rdtsc` instruction) or the operating system (e.g. the Windows API `QueryPerformanceCounter()` function). We can then select the n_{inst} value that produces the desired alternation frequency ($f_{\text{alt}} = 1/T_{\text{alt}}$).

```
1  while(1){
2    // Do some instances of the A instruction
3    for(i=0;i<n_inst;i++){
4      ptr1=(ptr1&~mask1)|((ptr1+offset)&mask1);
5      // The A-instruction, e.g. a load
6      value=*ptr1;
7    }
8    // Do some instances of the B instruction
9    for(i=0;i<n_inst;i++){
10     ptr2=(ptr2&~mask2)|((ptr2+offset)&mask2);
11     // The B-instruction, e.g. a store
12     *ptr2=value;
13   }
14 }
```

Fig. 1. The A/B alternation pseudo-code.

These microbenchmarks create EM emanations as shown in Figure 2. Intuitively we expect differences between the A and B instructions to appear at the frequency $f_{\text{alt}} = 1/T_{\text{alt}}$ where T_{alt} is the time required to execute one iteration of the outer loop in Figure 1.

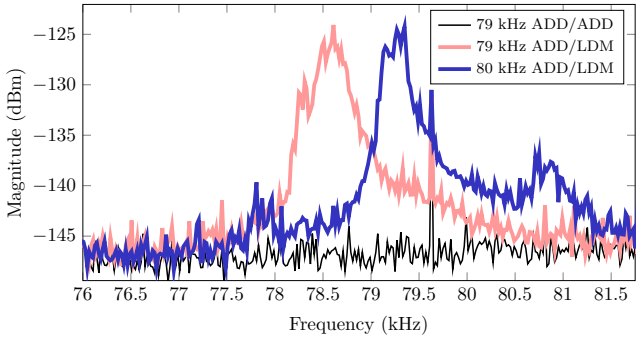

Fig. 2. Power spectrum of ADD/LDM instruction pair at 79 kHz and 80 kHz.

Figure 2 shows that we can choose alternation period T_{alt}, allowing us to avoid parts of the spectrum where other signals might be present. These spectra shows the ADD/LDM instruction pair (integer addition vs an off-chip memory load) with 79 and 80 kHz alternation frequencies, and also the ADD/ADD "pair". The ADD/ADD spectrum illustrates that, when A and B instructions are identical, the EM emanations for A and B activity are effectively the same, resulting in no signal at the alternation frequency. The 79 kHz and 80 kHz ADD/LDM spectra show broad peaks. These peaks are clearly not caused by other unrelated signals (such as nearby switching power supplies, CRT or LCD monitors, or other cabling) because the signal is only present when the A and B instructions differ (i.e. there is no signal for ADD/ADD),

and because the observed peak follows the intended alternation frequency. The generated signals are not perfectly concentrated at the intended f_{alt} because 1) f_{alt} cannot be controlled perfectly in a real system and 2) T_{alt}, i.e. the time to execute one iteration of the outer loop in Figure 1, varies slightly in complex processors and systems, resulting in the dispersion of power around the alternation frequency. The path loss measurements are performed by recording the peak of the signal at the alternation frequency.

III. Unintentional AM Carriers in Computer Systems

Amplitude modulation (AM) is well-studied and is used in numerous communication systems. Traditional communications rely on carefully designed transmitters and thoroughly regulated allocation of the frequency spectrum to optimize communication. On the other hand, unintentional AM signals in computer systems are generated by many possible "transmitters." A memory clock signal, for example, may act as a carrier. A clock signal creates periodic currents at the clock frequency f_c, and these currents flow through power and signal wires, generating a strong EM field. When the memory is active, more current is drawn by the clock, and less current is drawn when the memory is less active. If we alternate between high memory activity and low memory activity with a frequency f_{alt}, the amplitude of the carrier at f_c is modulated creating signals at $f_c \pm f_{alt}$. Figure 3 illustrates how SAVAT activity is modulated onto voltage regulator clock.

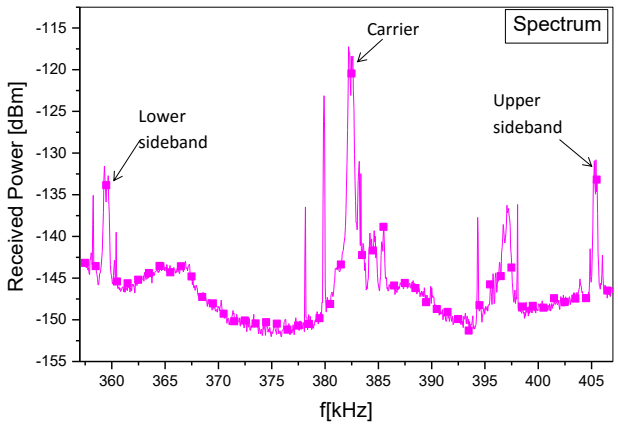

Fig. 3. A measured spectrum at a carrier frequency at 382 kHz produced by a voltage regulator clock and a lower and upper sidebands around 359 kHz and 405 kHz, respectively, produced by SAVAT activity.

IV. Path Loss Measurements and Modeling of EM Side-Channel Signals

In this section, we investigate near-field and far-field propagation mechanisms that govern EM side-channel signals and present models that characterize them. Furthermore, we compare proposed models with measured data to verify the validity of the models.

A. Path Loss Measurements and Modeling of Direct EM Side-Channel Signals

To predict path loss of direct EM side-channel signals, we start by measuring received signal power directly created by SAVAT executions at several different distances. Our measurements use the A/B alternation microbenchmark in Figure 1 to measure SAVAT for several NIOS instructions in Figure 4 including loads and stores that go to different levels of the cache/memory hierarchy, simple (ADD and SUB) and more complex (MUL and DIV) integer arithmetic, and the "No Instruction" case where the appropriate line in our alternation code (Line 6 or 12 in Figure 1) is simply left empty. For each pair of instructions A and B, we run the A/B microbenchmark and record the maximum of the spectrum in the vicinity of the alternation frequency. The SAVAT benchmarks ran on a NIOS II soft processor implemented on a DE1 Cyclone II FPGA board, with no memory management or operating system. No other logic was active on the FPGA.

	Instruction	**Description**
LDM	ldw r21, 0(r21)	Load from main memory
LDL1	ldw r21, 0(r21)	Load from L1 cache
ADD	addi r22,r22,173	Add imm to reg
SUB	subi r22,r22,173	Sub imm from reg
MUL	muli r22,r22,173	Integer multiplication
DIV	div r22,r22,r22	Integer division
NOI		No instruction

Fig. 4. NIOS instructions for our DE1 FPGA A/B SAVAT measurements.

A probe's type, position, and orientation affect the strength of the emanations it receives. A small "sniffer" probe placed a few millimeters above components picks up signals from only the components near the probe, but receives these signals very strongly. On the other hand, placing a probe with a larger effective area far away ($>$ 2 meters) will pick up signals from all the parts of the system, but is often not sensitive enough to pick up the weakest signals. To allow us to pick up emanations from all the parts of the system while at the same time being close enough to pick up the weakest signals tested, we settled on a compromise: a medium sized multiple turn loop (16cm^2 loop area, 20 turns) above the processor as shown Figure 5 (left). For our measurements the loop was rotated in all three directions to collect magnetic field in x, y, and z direction, and from those measurements total magnetic field is calculated and used for verification against the model. The power across the loop probe was measured using a spectrum analyzer (Agilent MXA N9020A) with a resolution bandwidth of 1 Hz to minimize the effects of variation in unrelated signals and noise.

Since we are measuring near-field signals using a magnetic loop probe, we have expected that the magnetic field will decay as $1/r^3$ and that the emanations source can be modelled as a magnetic dipole. However, our measurements did not match this model. Hence, we modelled the side-channel field as a field created by an electric monopole (Hertzian dipole) and a magnetic dipole, where we can receive only magnetic components of the EM field. Hence, the received power can

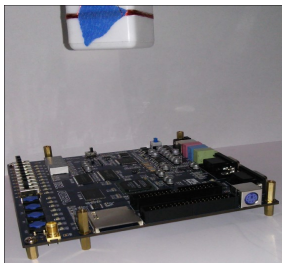

Fig. 5. Measurement setup for near-field measurements (left) and measurement setup for far-field measurements (right).

be modelled as

$$P_{rx}(H) \approx < |H|^2 > = P_{rx0} \left(\frac{1}{(kr^2)^2} + \frac{1}{(kr^3)^2} \right), \quad (1)$$

where $k = 2\pi/\lambda$ is the wavenumber, P_rx0 is a reference received power that corresponds to power measurements at 0.25 meters, and r is the distance between the antenna and the system. One of the main challenges in predicting propagation loss for EM side-channel signals is the fact that the transmit power and transmit "antenna" gain are unknown. To overcome this problem, we perform the measurements at "zero" distance to capture all losses that signal accumulates by exiting the electronics and reaching receive antenna. Often, this distance is not exactly zero because the computer casing, thickness of the motherboard, size of the probe, etc. can add significant distance between the transmitter and receiver. In our case, that distance was 0.25 m. After estimating transmit power $P_{tx}{}'$, we can use this model to predict the propagation distance.

Figure 6 compares modelled and measured received power for several representative instructions at 215 kHz. We observe that the received power of on-chip pairs of instructions (e.g. LDL1/DIV and LDL1/MUL) decays at the same rate as the received power on-chip/off-chip instruction pairs (e.g. LDL1/LDM) but that off-chip/on-chip signals weaker. Similar agreement between theoretical and measured SAVAT was found at 70 kHz and 150 kHz.

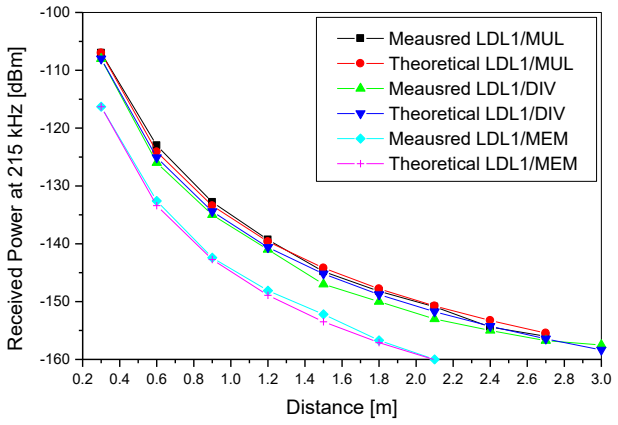

Fig. 6. A measured and modeled received power at 215 kHz produced by SAVAT activity.

B. Path Loss Measurements and Modeling of Unintentionally Modulated EM Side-Channel Signals

To predict path loss of indirect EM side-channel signals, we measure received signal power of unintentional carrier, and upper and lower sideband signals created by SAVAT executions at several different distances. Similarly as for direct emanations, here we use the A/B alternation microbenchmark in Figure 1 to measure SAVAT for several NIOS instructions in Figure 4 including loads and stores that go to different levels of the cache/memory hierarchy, simple (ADD and SUB) and more complex (MUL and DIV) integer arithmetic, and the "No Instruction" case. For each pair of instructions A and B, we run the A/B microbenchmark and record the maximum of the spectrum in the vicinity of the alternation frequency. The SAVAT benchmarks ran on a NIOS II soft processor implemented on a DE1 Cyclone II FPGA board, with no memory management or operating system. No other logic was active on the FPGA.

For far-field measurements we use horn antenna with frequency range of 1 GHz to 18 GHz with gain of 9 dBi in the frequency range of interest for this paper. We measure two harmonics of processor clock frequency, one at 1.083 GHz and the second one at 2.583 GHz as well as their sidebands created by SAVAT activity. The power across the horn antenna was measured using a spectrum analyzer (Agilent MXA N9020A) as shown in Figure 5 (right).

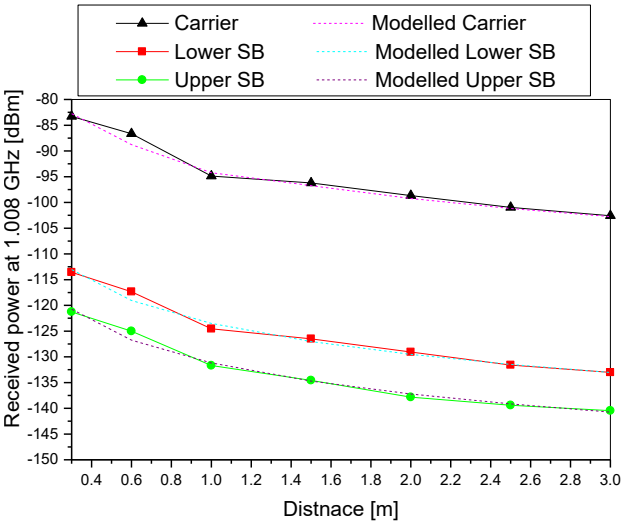

Fig. 7. A measured and modeled received power of processor clock harmonic at 1.083 GHz and AM modulated sidebands produced by SAVAT activity.

Here, our assumption that propagation loss can be modelled using Friis formula was adequate and the measurements were able to match the model. Hence, the received power is modelled as

$$P_{rx} = P_{rx0} \frac{1}{(kr)^2}, \quad (2)$$

where $k = 4\pi/\lambda$, P_{rx0} is a reference received power that corresponds to power measurements at 0.25 meters, and r is the distance between the antenna and the system. One of the

main challenges in predicting propagation loss for EM side-channel signals is the fact that the transmit power and transmit "antenna" gain are unknown. After estimating transmit power $P_{tx}{}'$, we can use this model to predict the propagation distance.

Figures 7 and 8 compare modelled and measured received power for LDL1/LDM instructions at 1.008 GHz and 2.583 GHz. The results show that the received power on-chip/off-chip instruction pairs (e.g. LDL2/LDM) is 20-30 dB weaker then the harmonic of the processor clock signal, but still about 20-30 dB above noise floor of the measurement system. We can observe that the signals can be detected up to 3 m away.

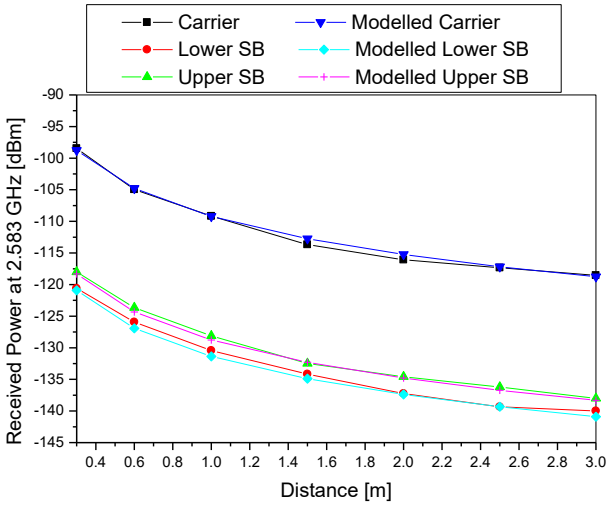

Fig. 8. A measured and modeled received power of processor clock harmonic at 2.583 GHz and AM modulated sidebands produced by SAVAT activity.

V. Conclusions

This paper investigated propagation mechanisms that EM side-channel signals experience at different frequencies and proposed models for near-field and far-field propagation of side-channel signals. The near-field propagation is modelled as a field created by an electric monopole (Hertzian dipole) and a magnetic dipole, where the received power is collected using only magnetic components of the EM field. This model resulted in excellent match with measured data. Furthermore, this paper investigates unintentionally modulated side-channel signals. The propagation of EM side-channel signals was modelled using free-space propagation model which resulted in excellent match with measured data. In both cases we have observed that signal can be received at several meters from the side-channel source. The proposed models are the first step in understanding propagation mechanisms of EM side-channel signals and how to predict the distance at which they can be received.

References

[1] A. G. Bayrak, F. Regazzoni, P. Brisk, F.-X. Standaert, and P. Ienne, "A first step towards automatic application of power analysis countermeasures," in *Proceedings of the 48th Design Automation Conference (DAC)*, 2011.

[2] D. Boneh and D. Brumley, "Remote Timing Attacks are Practical," in *Proceedings of the USENIX Security Symposium*, 2003.

[3] S. Chari, C. S. Jutla, J. R. Rao, and P. Rohatgi, "Towards sound countermeasures to counteract power-analysis attacks," in *Proceedings of CRYPTO'99, Springer, Lecture Notes in computer science*, pp. 398–412, 1999.

[4] D. Agrawal, B. Archambeult, J. R. Rao, and P. Rohatgi, "The EM side-channel(s)," in *Proceedings of Cryptographic Hardware and Embedded Systems - CHES 2002*, pp. 29–45, 2002.

[5] M. G. Khun, "Compromising emanations: eavesdropping risks of computer displays," *The complete unofficial TEMPEST web page: http://www.eskimo.com/˜joelm/tempest.html*, 2003.

[6] D. Genkin, A. Shamir, and E. Tromer, "Rsa key extraction via low-bandwidth acoustic cryptanalysis," in *Advances in Cryptology CRYPTO 2014* (J. Garay and R. Gennaro, eds.), vol. 8616 of *Lecture Notes in Computer Science*, pp. 444–461, Springer Berlin Heidelberg, 2014.

[7] S. Chari, J. R. Rao, and P. Rohatgi, "Template attacks," in *Proceedings of Cryptographic Hardware and Embedded Systems - CHES 2002*, pp. 13–28, 2002.

[8] M. Hutter and J.-M. Schmidt, "The temperature side channel and heating fault attacks," in *Smart Card Research and Advanced Applications* (A. Francillon and P. Rohatgi, eds.), vol. 8419 of *Lecture Notes in Computer Science*, pp. 219–235, Springer International Publishing, 2014.

[9] D. Agrawal, B. Archambeult, S. Chari, and J. R. Rao, "Advances in side-channel cryptanalysis electromagnetic analysis and template attacks," in *RSA laboratories cryptobytes*, pp. 20–32, 2003.

[10] D. Genkin, L. Pachmanov, I. Pipman, and E. Tromer, "Stealing keys from pcs using a radio: Cheap electromagnetic attacks on windowed exponentiation," in *Cryptographic Hardware and Embedded Systems – CHES 2015* (T. Gneysu and H. Handschuh, eds.), vol. 9293 of *Lecture Notes in Computer Science*, pp. 207–228, Springer Berlin Heidelberg, 2015.

[11] A. Zajic and M. Prvulovic, "Experimental demonstration of electromagnetic information leakage from modern processor-memory systems," *Electromagnetic Compatibility, IEEE Transactions on*, vol. 56, pp. 885–893, Aug 2014.

[12] R. Callan, A. Zajic, and M. Prvulovic, "A Practical Methodology for Measuring the Side-Channel Signal Available to the Attacker for Instruction-Level Events," in *Proceedings of the 47th International Symposium on Microarchitecture (MICRO)*, 2014.

[13] S. Clark, H. Mustafa, B. Ransford, J. Sorber, K. Fu, and W. Xu, "Current events: Identifying webpages by tapping the electrical outlet," in *Computer Security ESORICS 2013* (J. Crampton, S. Jajodia, and K. Mayes, eds.), vol. 8134 of *Lecture Notes in Computer Science*, pp. 700–717, Springer Berlin Heidelberg, 2013.

[14] C. Aguayo Gonzlez and J. Reed, "Power fingerprinting in sdr integrity assessment for security and regulatory compliance," *Analog Integrated Circuits and Signal Processing*, vol. 69, no. 2-3, pp. 307–327, 2011.

[15] S. S. Clark, B. Ransford, A. Rahmati, S. Guineau, J. Sorber, W. Xu, and K. Fu, "Wattsupdoc: Power side channels to nonintrusively discover untargeted malware on embedded medical devices," in *Presented as part of the 2013 USENIX Workshop on Health Information Technologies*, (Berkeley, CA), USENIX, 2013.

[16] R. Callan, A. Zajic, and M. Prvulovic, "FASE: Finding Amplitude-modulated Side-channel Emanations," in *42nd International Symposium on Computer Architecture (ISCA)*, 2015.

[17] R. Callan, N. Popovic, A. Daruna, E. Pollmann, A. Zajic, and M. Prvulovic, "Comparison of electromagnetic side-channel energy available to the attacker from different computer systems," in *Electromagnetic Compatibility (EMC), 2015 IEEE International Symposium on*, pp. 219–223, Aug 2015.

[18] R. Callan, N. Basta, A. Zajic, and M. Prvulovic, "A new approach for measuring electromagnetic side-channel energy available to the attacker in modern processor-memory systems," in *Proceedings of the 9th European Conference on Antennas and Propagation (EuCAP)*, 2015.

[19] N. Sehatbakhsh, A. Nazari, A. Zajic, and M. Prvulovic, "Spectral profiling: Observer-effect-free profiling by monitoring em emanations," in *The 49th Annual IEEE/ACM International Symposium on Microarchitecture, 2016*, 2016.

[20] R. Callan, F. Behrang, A. Zajic, M. Prvulovic, and A. Orso, "Zero-overhead profiling via em emanations," in *The International Symposium on Software Testing and Analysis 2016*, 2016.

[21] P. Kocher, J. Jaffe, and B. Jun, "Differential power analysis: leaking secrets," in *Proceedings of CRYPTO'99, Springer, Lecture notes in computer science*, pp. 388–397, 1999.