CryptArchi 2019

17th International Workshop on Cryptographic Architectures
Embedded in Logic Devices

Průhonice, Czech Republic.

June 23rd - 26th 2019

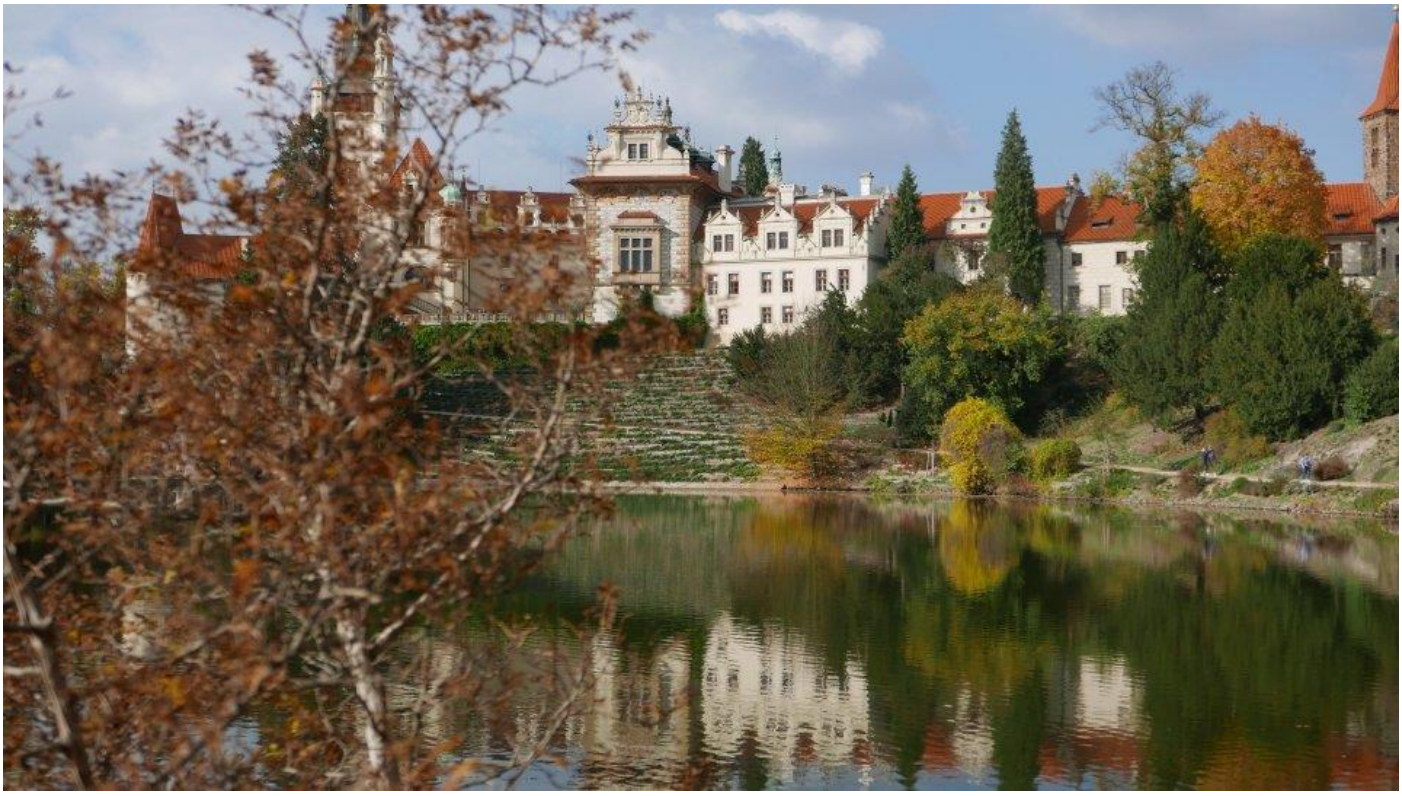

SCIENTIFICAL PROGRAM

MONDAY June 24th 2019

9:00 – 10:30 **Lightweight crypto / IoT** **Chair: V. Fischer**

Cedric Marchand, Ecole Centrale de Lyon **1**
Internet of things security: review from communication to sensor

Julien Francq, Airbus Defence & Space - CyberSecurity **2**
Design Exploration of the NIST LWC Competition Lilliput-AE

Etienne Tehrani, Telecom ParisTech **3**
Acceleration of Lightweight Block Ciphers on Microprocessors

11:00 – 12:30 **Session II: Masking and other countermeasures** **Chair: R. Zimmermann**

Vojtech Miskovsky, Czech Technical University in Prague **4**
Area-efficient fault-tolerant architectures exploiting masking scheme randomness

Wei Cheng, Telecom ParisTech **5**
Optimal Codes for Inner Product Masking

Tania Richmond, Inria - IRISA, Rennes **6**
Security aspects at the compilation level

14:30 – 16:00 **Session III: Microcontroller security** **Chair: K. Gaj**

Jan Riha, Czech Technical University in Prague **7**
Implementation and Effectiveness Evaluation of the VeraGreg Scheme on a Low-Cost Microcontroller

Alexander Schaub, Telecom ParisTech **8**
STAnalyzer: A simple static analysis tool for detecting cache-timing leakages

Petr Socha, Czech Technical University in Prague **9**
Toolkit for side-channel analysis: SICAK

16:30 – 18:00 **Session IV: RNG security** **Chair: M. Peter**

Marco Bucci, Infineon Technologies AG, Austria **10**
Offline and online testability of Random Number Generators

Markus Dichtl, Germany **11**
How (not) to end up with dependent random bits

Maciej Skorski, Dell, Austria **12**
Minimax Study of Bias Correctors

TUESDAY June 25th 2019

9:00 – 10:30 **Session V: Deep learning and other attacks** **Chair: F. Regazzoni**

Francis Olivier, Thales-DIS (formerly Gemalto) **13**
Deep Learning versus Template Attacks: experimental comparison

Damien Robissout, LabHC St-Etienne **14**
Improved Deep-Learning Side-Channel Attacks using Normalization layers

Martin Jurecek, Czech Technical University in Prague **15**
Cryptanalysis of the A5/1 Using Power Analysis

11:00 – 12:30 **Session VI: Hardware implementations** **Chair: M. Bucci**

Milos Grujic, imec-COSIC, KU Leuven, Belgium **16**
A Multimode Ring Oscillator based TRNG for FPGAs

Bertrand Cambou, Northern Arizona University, USA **17**
Replacing error correction by key fragmentation and search engines To generate error-free cryptographic keys from PUFs

Michal Andrzejczak, Military University of Technology in Warsaw, Poland **18**
Lattice sieving acceleration in FPGAs

14:30 – 16:00 **Session VII: Security challenges** **Chair: M. Novotny**

Kris Gaj, George Mason University, USA **19**
Toward Efficient and Fair Software/Hardware Codesign and Benchmarking of Candidates in Round 2 of the NIST PQC Standardization Process

G. Richard Newell, Microchip Technology (FPGA Business Unit), USA **20**
Survey of Notable Security-Enhancing Activities in the RISC-V Universe

Francesco Regazzoni, ALaRI - USI, Lugano **21**
Security Challenges in Cyber-Physical Systems

16:30 – 18:00 **Session VIII: Physical attacks** **Chair: J. Francq**

Ugo Mureddu, STMicroelectronics France **22**
Transient Effect Ring Oscillators Leak Too

Brice Colombier, LabHC St-Etienne **23**
Laser-induced Single-bit Faults in Flash Memory: Instructions Corruption on a 32-bit Microcontroller

Jan Belohoubek, Czech Technical University in Prague **24**
Modulated CMOS Static Power is Data Dependendent and Observable

ABSTRACTS

1- **Cedric Marchand,** Ecole Centrale de Lyon

Internet of things security: review from communication to sensor

The Internet of Things (IoT) is today a well-known ecosystem, continuously growing at an impressive speed (from about 500 million connected things in 2016 to more than 6.5 billion in 2018), where small and smart objects interact through communicating networks. At this rate, there will be 6 times more connected things than human on earth by 2025. Consequently, the amount of collected and transmitted data is becoming enormous and depending on the application, these data will be sensitive and require protection. However, it is important to remember that the very first constraint in the IoT context is the energy consumption of the devices. Thus, it is a great challenge to combine security and energy consumption in this context.
Concerning security, an effort has been done to identify risks and challenges resulting from a lack of security but also in order enhance the security of communication protocols used in the IoT. Another solution is to add security features externally using either hardware accelerators or hardware secure elements to encrypt collected data just before the communication. However, these solutions require important area and power consumption overhead. Thus, it is necessary to find new ultra-lightweight solutions that make it possible to bring the security as close as possible to sensor in order to enhance the security in the IoT context.
In this presentation, we will first present some statistic and public concern about IoT security. Then, we will review security solutions proposed to increase IoT security from communication to sensor. Finally, we will discuss about non-volatile computing perspectives to bring security at the sensor closest possible point and how this will enhance the security of sensor node in the IoT context.

2- **Julien Francq, Airbus Defence & Space - CyberSecurity**

Design Exploration of the NIST LWC Competition Lilliput-AE

Around 60 candidates will be competitors of the new upcoming NIST Lightweight Cryptographic Standardization Process. Among them, Lilliput-AE is a candidate which has serious advantages from security and performance point of view.
It has been already shown in its specification package that Lilliput-AE performs very well on software on 8-bit (e.g., ATMega 128) and 16-bit (e.g., MSP430) platforms since it has comparable or smaller execution time than the two final members of CAESAR (Competition for Authenticated Encryption: Security, Applicability, and Robustness) lightweight portfolio: Ascon and Acorn.
This talk will detail implementation results on FPGAs and ASICs for 3 architectures: straightforward, serial and thresholded. Comparisons will also be done with Ascon and Acorn.
This will show that Lilliput-AE is also well suited for hardware constrained environments.

3- **Etienne Tehrani,** Telecom ParisTech

Acceleration of Lightweight Block Ciphers on Microprocessors

Cryptography is a key element to the development of secure communication in embedded environment such as within or between connected cars. In such constrained devices standard cryptographic algorithms have been considered too costly which lead to the emergence of specific Lightweight Block Ciphers (LBC). The lack of

standards alongside industry's desire to use uniquely tweaked LBC calls for a generic and efficient implementation of those algorithms. Microprocessors are a part of most of these embedded systems which allows them to implement any of these algorithms but not efficiently way as it lacks specific instructions. For instance, the RiscV is an open source ISA which can be used in these microprocessors and is currently being enhanced by research through extensions. In this work we propose the study of this ISA and the development of an extension for efficient implementation of LBC.
From the state of the art we have selected some LBC based on the following criteria: at least a 128-bit key for security and a 64-bit block size to limit the necessary resources. In order to identify useful extensions, we first identified which parts of LBC are slow when implemented in pure software, and how common they are in state of the art LBC. We only studied the datapath of the cipher as we considered the key scheduling to be part of preprocessing. We used a software implementation of each of the studied algorithms to isolate the costly parts of the ciphers. The computation time was evaluated in number of RiscV assembly language instructions.
Studied LBC algorithms exhibit 3 main computation steps:
• The key addition which is a simple XOR and doesn't require additional instructions
• 8 or 16 4x4 Sbox (common for LBC) which can be implemented as LUT and can be accelerated thanks to the addition of a specific (SIMD) LUT instruction
• The diffusion is generally not trivial to implement in pure software and as it can be quite different from one algorithm to the other it is not obvious to provide a unique extension to implement it.

4- **Vojtech Miskovsky,** Czech Technical University in Prague

Area-efficient fault-tolerant architectures exploiting masking scheme randomness

Electronic systems become more and more important part of everyday life including safety-critical spheres like transportation or even medical devices. These systems need to fulfill strict dependability properties. To ensure correct operation of such devices, we need to make them fault-tolerant. This is usually achieved using common redundancy schemes like triple-modular redundancy (TMR). These schemes are simple to implement, but introduce high overhead. In case of TMR, we need to replicate the design three times and place majority voters on the outputs.
Since these systems are usually connected to some network, their activity and communication should be encrypted. Nevertheless, side-channel analysis pose a threat even to modern cryptographical algorithms. Luckily, many side-channel countermeasures based on hiding or masking exist to protect the device against such an attack. As these countermeasures can introduce very high area and/or power overhead, the total resources of a SCA-protected and fault-tolerant cryptographical module can be unbearable.
In our work, we deal with this issue and we present area-efficient fault-tolerant schemes exploiting randomness introduced in masking-based SCA countermeasure. Using our approach, it is for example possible to achieve properties similar to TMR (tolerance to single module failure) using only two redundant modules.

5- **Wei Cheng,** Telecom ParisTech

Optimal Codes for Inner Product Masking

Masking is the most popular countermeasure to protect cryptographic implementations against side-channel analysis, since it is provable secure and can be deployed at algorithm level. To strengthen the original Boolean masking scheme, several works have suggested to use more complicated schemes with high algebraic complexity, like affine masking and polynomial masking. Therefore, the Inner Product Masking (IPM) was proposed to be a better alternative with its intrinsic algebraic complexity. In this work, we express the security order of generalized IPM schemes from the viewpoint of coding theory, which allows us to optimize it.

Specifically, we highlight first that the IPM scheme is not optimal by showing different security order in byte- and bit-level, respectively. In particular, this result confirms the previous observations made by Balasch et al. at EUROCRYPT' 15 and at ASIACRYPT' 17 and Poussier et al. at CARDIS' 17 regarding the parameters effect in IPM. More importantly, we characterize this parameter effect by linking the side-channel resistance of IPM to the concept of minimum distance and one coefficient in weight enumeration polynomial of a linear code. The closed-form expression is proposed for depicting the connection, also allows us to systemically choose optimal codes for IPM. As the last contribution, we present the optimal linear code in several scenarios for IPM with two and three shares. The experiments are in perfect accordance with our theoretic analysis and finely demonstrate the optimality of the codes chosen by our method. Our results also present a solid explanation on parameters effect found by Balasch et al. and Poussier et al.

6- **Tania Richmond,** Inria - IRISA, Rennes

Security aspects at the compilation level

We are involved in the EU Horizon 2020 TeamPlay project. The aim of TeamPlay is to add non-functional properties such as time, energy and security as first class citizens in programs, targeting multicore heterogeneous platforms (e.g. mobile applications, IoT). At high level, we consider non-functional properties and their expressiveness. We transfer these high level properties down to the low level using the TeamPlay toolchain. At low level, we analyze side-channel vulnerabilities caused by time and power differences and furthermore automatically apply compiler-level countermeasures like equalization or noise addition. To illustrate this work, we rely on the modular exponentiation of RSA.

7- **Jan Riha,** Czech Technical University in Prague

Implementation and Effectiveness Evaluation of the VeraGreg Scheme on a Low-Cost Microcontroller

Homomorphic encryption is an effective way of securing data privacy while maintaining the possibility to process the data. The VeraGreg framework, unlike other existing homomorphic cryptosystem allows for verification of computation that was done with the encrypted data.
This work deals with an implementation of the VeraGreg framework and its effectiveness comparison with a naïve scheme based on symmetric encryption. Secure microcontroller CE1302 was chosen as the implementation platform. A new library for multiprecision integer arithmetic was created as well as the first published implementation of Paillier cryptosystem using hardware RSA accelerator.
The VeraGreg framework is 200 times slower compared to the naive scheme and occupies one third more space in the program memory, so it is not a suitable alternative to symmetric cryptosystems. On the other hand, it provides privacy to the user while allowing computations with the encrypted data, and verifying that is has not been manipulated during the computation.

8- **Alexander Schaub,** Telecom ParisTech

STAnalyzer: A simple static analysis tool for detecting cache-timing leakages

Cache-timing attacks are a class of side-channel attacks that target software implementations of cryptographic algorithms. If the cache-access pattern of the implementation depends on sensitive information, then a cache-timing attack can retrieve this information, which can potentially lead to a secret-key recovery. Implementations which branch on conditions depending on sensitive information, or that access memory locations whose address

depend on sensitive information, are potentially vulnerable to such attacks.

This paper presents an algorithm for verifying that a program, implemented in the C language, is free from cache-timing leakages. It consists in computing the dependencies of all the variables used in the program, and listing all sensible values that leak due to branching and memory accesses. An implementation of this algorithm, STAnalyzer, is also provided. It allows to flag sensitive values, and those are tracked across computations, function calls, etc. Therefore, only leakages of sensitive values are reported. Because the algorithm runs directly on an abstract syntaxic tree (AST) of the C program, the output is straightforward to interpret: dependencies between C variables are reported, as well as the stack of function calls and instructions that lead to the leakage of sensitive values.

9- **Petr Socha,** Czech Technical University in Prague

Toolkit for side-channel analysis: SICAK

Side-channel cryptanalysis pose a serious threat to many modern cryptographic systems. Typical scenario of a side-channel attack consists of an active phase, where data are acquired, and of an analytical phase, where the data get examined and evaluated.
This work presents a software toolkit which includes support for both phases of the side-channel attack. The toolkit consists of non-interactive text-based utilities with modular plug-in architecture. The measurement utility supports different oscilloscopes, target interfaces and measurement scenarios. The evaluation utilities include support for the test vector leakage assessment and the CPA attack. Different approaches to the algorithmical evaluation of the attack are implemented in order to extract the cipher key. The visualisation utility allows for the visual examination of the attack results by the user.
The toolkit aims to be multiplatform and it is written using C/C++ with performance in mind. Time-demanding operations (such as the statistical analysis) are accelerated using OpenMP and OpenCL for an efficient computation on both CPU and GPU devices.

10- **Marco Bucci,** Infineon Technologies AG, Austria

Offline and online testability of Random Number Generators

In this work, the offline and online testability of Random Number Generators (RNG's) and entropy sources is discussed. It is pointed out that, despite this is commonly accepted, the evaluation method used for pseudo-RNG's is not applicable to RNG's where the focus must be on the actual entropy. It is shown that the min-entropy, which is proposed as a quality figure for entropy sources, is not really a conservative estimation of the actual entropy but, in most of the cases, just a misleading quantity which can also result in an over estimation. On the other side, it is shown that, provided that the sequence under test features a relatively short memory, the actual Shannon entropy can be straightforwardly and correctly evaluated. Surprisingly, if the digital post-processing of an RNG is properly designed, this method can be also applied after post-processing, thus assessing whether or not the RNG features (a practically) maximal entropy.
Finally, feasibility, costs and robustness of online tests is compared vs the usage of redundant (i.e. multiple) entropy sources. It is shown that redundancy provides a better security coverage while, online tests, can even be exploited for mounting attacks.

11 **Markus Dichtl,** Germany

How (not) to end up with dependent random bits

Sampling a jittering ring oscillator with a D-flip-flop in order to generate random bits seems to be a completely trivial operation. However, this talk provides experimental evidence from an Artix-7 FPGA that the bits sampled strongly depend on the bit stored in the flip-flop prior to the sampling. In a ring oscillator based true random number generator, the bit stored in the flip-flop usually is the previous random bit. As a consequence, subsequent random bits can be dependent.
This phenomenon was studied on ring oscillators of different lengths and in different phase situations. Additionally, the sampling of the xor of signals from multiple ring oscillators was evaluated.
The talk discusses possible reasons for this problem and why it has not been observed previously. Furthermore, several approaches to overcome the problem are suggested.

12- **Maciej Skorski,** Dell, Austria

Minimax Study of Bias Correctors

Bias correctors are deterministic functions used to reduce the bias in physical random number generators; the most famous example is the xor function. In this talk we discuss how to construct such correctors under some prior assumptions on the input bias (for example, that the bias is bounded). We will see how to find the optimal construction by solving a min-max optimization over boolean functions.
It turns out that the xor corrector is optimal when the bias is sufficiently small, but - interestingly - can be improved for certain bias values.
We will also overview related results on the xor corrector, such as bias formulas due to Lacharme and Davies.

13- **Francis Olivier,** Thales-DIS (formerly Gemalto)

Deep Learning versus Template Attacks: experimental comparison

This study compares the experimental results of Template Attacks (TA) and Deep Learning (DL) techniques called Multi Layer Perceptron (MLP) and Convolutional Neural Network (CNN), concurrently in front of classical use cases often encountered in the side-channel analysis of cryptographic devices (restricted to SK). The starting point regards their comparative effectiveness against masked encryption which appears as intrinsically vulnerable. Surprisingly TA improved with Principal Components Analysis (PCA) and normalization, honorably makes the grade versus the latest DL methods which demand more calculation power. Another result is that both approaches face high difficulties against static targets such as secret data transfers or key schedule. The explanation of these observations resides in cross-matching. Beyond masking, the effects of other protections like jittering, shuffling and coding size are also tested. At the end of the day the benefit of DL techniques, stands in the better resistance of CNN to misalignment.

14- **Damien Robissout,** LabHC St-Etienne

Improved Deep-Learning Side-Channel Attacks using Normalization layers

Recent papers used deep neural networks to improve profiled side-channel attacks. The networks were efficient even in the presence of countermeasures such as masking and de-synchronization. Nevertheless, tuning networks to perform specific tasks is difficult and there is still room for improvement. One of such improvements

takes the form of normalization. Batch normalization is a well-known technique in the machine learning community to prevent the network from overfitting. In our work, using the ASCAD database, we test the addition of batch normalization layers and compare the results against the existing networks. Our experimental results cleary show that we significantly improve the attacks by reducing the number of traces needed to retrieve the key to about 1000 no matter the de-synchronization.

15- **Martin Jurecek,** Czech Technical University in Prague

SPA and DPA attack on the A5/1 Stream Cipher

In our contribution we describe cryptanalysis of the A5/1 stream cipher based on power analysis where we utilize the fact that the power consumption while clocking 3 LFSRs is different than when clocking 2 LFSRs.
Main part of our contribution is the presentation of a key recovery algorithm that uses information gathered by a simple power analysis (SPA) attack. We also discuss possible use of an existing differential power analysis (DPA) attack for providing input to the key recovery. The SPA attack was demonstrated on a simple prototype implementation of A5/1 on an 8-bit microcontroller.
The attack does not require (nor use) any knowledge of the keystream. For key recovery, we assume the SPA provides correct sequence of how many registers were shifted in each clock of the initialization phase (100 empty clocks). DPA also provides information which registers were shifted. The key recovery attack has a 100% success rate and requires minimal storage.
An average time complexity of our attack based on SPA is 2^(33.27) where the computation unit is a resolution of system of linear equations over the Z_2. Recovering the secret key using information from DPA has a constant complexity.

16- **Milos Grujic,** imec-COSIC, KU Leuven, Belgium

A Multimode Ring Oscillator based TRNG for FPGAs

True random number generators (TRNGs) are essential cryptographic components. Due to ever increasing ubiquity of FPGAs in modern embedded security systems, there exists a growing demand for fully-digital TRNGs that are suitable for FPGAs. As one possible solution to this problem, the TRNGs based on time-to-digital conversion are proposed. TDC based TRNGs rely on the concept of sampling the unstable signal of the ring oscillator with very high precision by using fast delay-chains.
In this work, we present a novel delay-chain based TRNG design with conservative security evaluation. Our design aims to minimize the effects of the unwanted noise sources and pseudo-randomness. Unlike previous DC-TRNG designs, our TRNG extracts significant entropy from the jittery pulse of the multimode ring oscillator. Additionally, we improved the design of digitization part to further boost the TRNG entropy rate. Our design is accompanied with a stochastic model which takes into account disparities and non-linearities in the structure of FPGAs. Further, we discuss influence of the correlated noises and the switched-mode power supply (SMPS) on the performance and security of the proposed TRNG.

17- **Bertrand Cambou,** Northern Arizona University, USA

Replacing error correction by key fragmentation and search engines To generate error-free cryptographic keys from PUFs

Physical Unclonable Functions (PUFs), the fingerprints of microelectronic components, are subject to aging, temperature drifts, electromagnetic interactions, and various environmental effects. Typically, this results in 2-

10% error rates between the initial readings of the PUFs that are stored as references, and the responses generated from these PUFs. Error correcting (ECC) algorithms need to correct all errors, i.e. 100%, to generate usable cryptographic keys from PUFs, single-bit mismatches being unacceptable. ECC algorithms use helper data from the transmitting party, and iterative methods such as fuzzy extractors at the receiving party, which consume computing power, and could thereby leak information to the opponents. The response based cryptographic method (RBC) eliminates the need to use error correction at the client level, as it generates cryptographic keys directly from the un-corrected responses of the PUFs. This technology relies on the implementation of an efficient search engine, driven by the secure server, interacting with the network of client devices, which finds the uncorrected PUF responses, rather than correcting them. The matching algorithms start from the original PUF challenges, and exploit the knowledge of the cipher texts generated by the transmitting client devices, which have access to uncorrected PUF responses as cryptographic keys.
In this paper, we are presenting an optimization of the RBC algorithm by fragmenting the keys generated from the uncorrected responses. The experimental work is based on window 10 PCs powered by Intel I7 quad core processors, able to process the Advanced Encryption Standard (AES) in five microseconds. The client devices are the WiFire development boards from Microchip with 200Mhz 32-bit RISCs from ARM. The PUFs are generated from commercial 32Kbyte SRAMs produced by Cypress Semiconductor. Below 1% challenge-response-pair (CRP) error rates, the RBC algorithm is fast enough to be able to handle 256-bit long keys, and find in a few seconds the uncorrected keys generated by the client device, and its SRAM PUF. In order to enhance effectiveness at higher error rates, we propose the fragmentation of the keys generated by the PUF into sub-keys, also 256-bit long, which are padded with known random numbers. In a fragmentation by two, the first sub-key is generated by keeping the first 128 bits of the key generated by the PUF, filled with 128 bits that do not contain errors. The second sub-key is generated from the last 128 bits of the PUF. Statistically, the two sub-keys are showing error rates that are half of those of the full keys, and they are faster to find by the RBC engine.
Both modelling and experimental data demonstrate that the RBC with fragmentation can be effective on any PUFs, regardless of the CRP error rates. It is always desirable to minimize the level of fragmentation to reduce computing power at the client level. The suggested method has the potential of eliminating the need to use error correcting methods, as well as helper data, and thereby simplifies PUF-based cryptographic protocols, and enhances security.

18- **Michal Andrzejczak,** Military University of Technology in Warsaw, Poland

Lattice sieving acceleration in FPGAs

Lattice-based cryptography is a promising option for secure communication in post-quantum era. Recently, a significant effort has been put into improving algorithms for solving lattice problems, such as the Shortest Vector Problem, especially using an algorithm called lattice sieving. The aforementioned algorithm is used for the cryptoanalysis of lattice-based schemes and as a result also for proper security parameters selection. Recent improvements and acceleration for lattice-sieving have been accomplished primarily using progress in the underlying math. In this talk we present the first reported attempt at speeding up lattice sieving algorithms with FPGAs in various scenarios, in particular, by using software/hardware codesign approach.

19- **Kris Gaj,** George Mason University, USA

Toward Efficient and Fair Software/Hardware Codesign and Benchmarking of Candidates in Round 2 of the NIST PQC Standardization Process

Post-Quantum Cryptography (PQC) refers to a new class of cryptographic algorithms that are resistant against all known attacks using quantum computers, but at the same time can be implemented by themselves using traditional computing platforms, such as microprocessors, microcontrollers, Field Programmable Gate Arrays

(FPGAs), and Application Specific Integrated Circuits (ASICs). PQC is a cryptographic community's response to the emerging threat of full-scale quantum computers, expected to be developed within the next decade or two. The main goal of PQC is to replace the existing public-key cryptography standards, based on RSA and Elliptic Curve Cryptography, which seem to be the most vulnerable to quantum computing and impossible to defend using traditional approaches, such as gradually increasing key sizes.
In order to initiate a timely transition to a new class of cryptographic schemes, in December 2016, NIST issued an official "Call for Proposals and Request for Nominations for Public-Key Post-Quantum Cryptographic Algorithms," commencing the NIST PQC standardization process. The number of submissions qualified to Round 1, started in December 2017, reached 69. In January 2019, based on the results of the initial security analysis and preliminary software benchmarking, 26 submissions were qualified by NIST to Round 2. These submissions included multiple public-key encryption, key encapsulation mechanism (KEM), and digital signature schemes, often more than one per a single submission.
Traditionally, hardware benchmarking played a major role in all recent cryptographic standardization efforts, such as AES, eSTREAM, SHA-3, and CAESAR contests. Unfortunately, this trend is not likely to be sustained in case of the NIST PQC standardization process, by simply following the old practices and hardware benchmarking approaches. In many respects, PQC schemes are dramatically different from those evaluated in previous cryptographic contests, and new challenges call for new substantially different solutions.
During the past contests, software and hardware benchmarking were conducted separately, by different groups of experts, equipped with different knowledge and tools. For PQC algorithms, this approach is hard to maintain. These algorithms are simply too complex and too different from the current state-of-theart to permit the development of optimized purely hardware implementations of a significant percentage of the remaining candidates by a single group within the time frame imposed by the NIST evaluation process (about 12-18 months per single round).
At the same time, there is little if any consensus, regarding basic design choices. In 16 months since the start of the PQC Round 1 (or before), only a few purely hardware implementations of Round 1 candidates were announced and even fewer were made open source. These implementations used different Application Programming Interfaces (APIs), targeted different platforms, and aimed at different optimization targets from high-speed to low-area. No conclusions regarding ranking of these algorithms in terms of their performance in hardware can be reached based on such divergent efforts.
In this talk, we propose a new approach to systematic benchmarking of candidates in cryptographic contests, based on the development and experimental measurements of their software/hardware codesigns. This approach is particularly applicable to the current stage of the NIST PQC standardization process, where a large number and high complexity of the evaluated algorithms makes the traditional hardware benchmarking practically infeasible. We propose and justify the choice of a suitable platform and design methodology. We demonstrate the validity of our approach by applying it to 7 Key Encapsulation Mechanisms (KEMs), representing 5 NIST Round 2 PQC candidates.
The obtained results indicate a potential for very substantial speed-ups vs. purely software implementations, ranging between 5 and 187 for encapsulation and between 15 and 444 for decapsulation. These speed-ups depend primarily on the percentage of the software execution time taken by functions offloaded to hardware (rather than the amount of acceleration itself). Ranking of the investigated candidates is affected, but not dramatically changed, by hardware acceleration.
At the same time, it should be noted that our current study cannot be used to predict the performance and ranking of the investigated candidates when implemented entirely in hardware. Such implementations can further benefit from elimination of the communication overhead between a processor and a hardware accelerator. They may also take advantage of an ability to parallelize some additional operations, left in software in the current study. As a result, more effort, by multiple groups, is needed to determine and realize the most efficient and fair software/hardware partitioning schemes, and to extend our study to the remaining Round 2 PQC candidates.

20- G. Richard Newell, Microchip Technology (FPGA Business Unit), USA

Survey of Notable Security-Enhancing Activities in the RISC-V Universe

The RISC-V open source Instruction Set Architecture (ISA) has quickly become the premier vehicle for CPU security research and many new commercial CPU implementations. The non-profit RISC-V Foundation, with over

235 member organizations representing industry and academia world-wide, has made security a top priority, elevating the Security Standing Committee as the only single-topic committee to report directly to the Foundation's Board of Directors. There are two security-focused technical task groups dedicated to developing ISA extensions, for Trusted Execution Environments and Cryptography, respectively, with influence in task groups working on Formal Specifications, Debug, the Privilege Specification, and other RISC-V task groups having a security impact. The RISC-V ISA has been selected by the DARPA SSITH program for all performers to use for its cyber security designs, and will be one of just two ISAs chosen by the DARPA AISS program for its security research. Together, these two DARPA programs are investing approximately $100M (US) in RISC-V HW-based security. The free and open RISC-V ISA has quickly become the de facto vehicle for CPU security research in academia. This talk will provide a necessarily brief survey of notable RISC-V security activities that are heralding in a new age in CPU security.

21- **Francesco Regazzoni,** ALaRI - USI, Lugano

Security Challenges in Cyber-Physical Systems

Cyber-Physical Systems (CPSs) tightly integrate cyber components (typically computation and communication elements) with physical components, such as sensors and actuators. These systems are often used in safety-critical applications, such as autonomous driving or medical devices, and are often used to manage and control our critical infrastructure, including smart grids and transportation systems. CPSs used in these application must withstand error and failures. These failures can be "natural" or caused by environmental conditions, but they can even be the outcome of a deliberated attack to the system. The capability of attackers have been demonstrated in different scenarios, including manipulation of industrial implant or hacking cars.
As embedded systems, CPSs are vulnerable to the cyber-attacks such as malware injection and to physical attacks such as power analysis and fault injection attacks. However, the presence of a physical component opens novel possibilities to adversaries. Side channels, for instance, can be used to extract design files from 3D printers while fabricating objects and manipulations of the design tools can be used to increase the leak of information or to alter the quality of a product.
Counteracting these threats is of utmost importance for the safe deployment of the applications relying on CPSs, including autonomous driving and industry 4.0. If on the one side there is a solid knowledge regarding cyber-attacks and physical attacks targeting the cyber part of a CPS, very little is know about attacks targeting the physical part and the possible countermeasures. Addressing this issue, this talk summarizes the main security and reliability challenges specific to CPS, discussing the main threats, the most relevant approaches to counteract them, and highlighting novel research directions.

22- **Ugo Mureddu,** STMicroelectronics France

Transient Effect Ring Oscillators Leak Too

Up to now, the transient effect ring oscillator (TERO) seemed to be a better building block for PUFs than a standard ring oscillator, since it was thought to be immune to electromagnetic analysis. Here,we report for the first time that TERO PUFs are in fact vulnerable to
electromagnetic analysis too. First, we propose a spectral model of a TERO cell output, showing how to fit it to experimental data obtained with the help of a spectrum analyser to recover the number of oscillations of a TERO cell. We then extend it to two TERO cells oscillating simultaneously, and show how this ability can be used to fully clone a TERO PUF. These results should help designers to better plan for
susceptibility of TERO PUFs to electromagnetic analysis in their future designs.

23- **Brice Colombier,** LabHC St-Etienne

Laser-induced Single-bit Faults in Flash Memory: Instructions Corruption on a 32-bit Microcontroller

Physical attacks are a known threat posed against secure embedded systems. Notable among these is laser fault injection, which is often considered as the most effective fault injection technique. Indeed, laser fault injection provides a high spatial accuracy, which enables an attacker to induce bit-level faults. However, experience gained from attacking 8-bit targets might not be relevant on more advanced micro-architectures, and these attacks become increasingly challenging on 32-bit microcontrollers. In this article, we show that the flash memory area of a 32-bit microcontroller is sensitive to laser fault injection. These faults occur during the instruction fetch process, hence the stored value remains unaltered. After a thorough characterisation of the induced faults and the associated fault model, we provide detailed examples of bit-level corruption of instructions and demonstrate practical applications in compromising the security of real-life codes. Based on these experimental results, we formulate a hypothesis about the underlying micro-architectural features that explain the observed fault model.

24- **Jan Belohoubek,** Czech Technical University in Prague

Modulated CMOS Static Power is Data Dependendent and Observable

As digital devices penetrate into many areas important for the present society, it is important to analyze even potential threats to mitigate device vulnerability during the lifetime of a digital device.
Skorobogatov has shown, that it is possible to obtain data stored in a register by using the invasive methods (the chip decapsulation and the laser beam). The disadvantage is, that the probing method is strongly limited by the transistor size and it is not applicable to recent technologies.
To overcome the size limitation, we analyzed the data dependency of the static power of CMOS combinational logic. We found that the static power (of a relative large combinational logic) modulated by the laser beam may decrease the entropy of the processed data. This is achieved by correlating the measured current consumption (induced by a laser beam) with the power model reflecting the data dependency of the laser-induced current. Moreover, the results have shown, that in certain cases, it is possible to obtain the processed data directly. When targeting sufficiently large combinational logic with the special structure, namely majority voter, the processed bit leaks.


UNIVERSITÉ
DE LYON


UNIVERSITÉ
JEAN MONNET
SAINT-ÉTIENNE


GDR
Sécurité Informatique


FACULTY
OF INFORMATION
TECHNOLOGY
CTU IN PRAGUE


cnrs


LABORATOIRE
HUBERT CURIEN
UMR • CNRS • 5516 • SAINT-ETIENNE