ISSN 2186-7437

NII Shonan Meeting Report

No. 236

Provable Security for Trustworthy Embedded Systems against Physical Attacks - From Theory to Practice

Yuko Hara
Svetla Nikova

March 16–19, 2026

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

Provable Security for Trustworthy Embedded Systems against Physical Attacks - From Theory to Practice

Organizers:
Yuko Hara (CNRS, France)
Svetla Nikova (KU Leuven, Belgium)

March 16–19, 2026

In the Internet of Things (IoT) era, a number of embedded devices are considered as security-critical in a sense that they are handling confidential data. The communication between such devices can be protected using cryptographic techniques, such as the AES and RSA primitives. However, these devices can be attacked via implementation means (a combination of active/passive and local/remote approaches). Representative examples are Differential Power Analysis, which is to observe the device's power consumption (developed by Kocher et al. in 1999), and Differential Fault Analysis, which is to inject faults in the device's computation for figuring out the cryptographic secrets (developed by Biham and Shamir in 1997). Since then, these threats have been well recognized and a whole new area of research has been established looking at both new physical attacks and efficient countermeasures to resist the attacks.

The first Shonan Meeting focusing on hardware security (No.028), which was still an evolving research field back then, was held in 2014 to discuss how designers can construct secure hardware in a systematic way. In the past 10 years after the Shonan Meeting (No.028), we have been observing drastic advancements and evolutions in the information processing infrastructure to lead the current IoT technology – CMOS downscaling, the prevalence of system-level (higher abstraction level) design methods, further advancement of cryptographic theory, artificial intelligence (AI) technologies, etc. Because of them, there are now a variety of attack targets/applications (not only cryptographic primitives but also (for example) AI-based applications), countermeasures in hardware and software (and combinations of both), and attack techniques, for which the secure implementations are much more complicated. For such a multidisciplinary research field, it was not easy to have tight interactions between different research communities. Therefore, this meeting encouraged researchers in both - theory and practice - to tightly interact and exchange insights and work together towards practical methods for secure embedded systems implementations. Specifically, this meeting aims to discuss the following scientific questions:

- Development of practical adversary models reflecting well real physical attacks and making it possible to prove the security of countermeasures in these models

- Design of efficient and effective provably secure countermeasures against side-channel attacks (SCA), fault attacks and combined attacks both in hardware and software
- The new threats when IoT and AI technologies are advanced and diversified, for example in terms of the attack targets/applications and methods.

To encourage these discussions, we invited leading researchers and practitioners with a variety of background in terms of theory/practice, hardware/software, countermeasure/attack, etc.

Background and introduction

This meeting was organized to address emerging challenges in ensuring the security of embedded systems against physical attacks. Advances in semiconductor technologies, such as chiplet integration and backside processing, are creating new attack surfaces. At the same time, side-channel and fault injection techniques are becoming increasingly sophisticated.

The meeting aimed to bridge theory, circuit-level design, and system-level implementation to better understand and mitigate these threats. It brought together researchers from hardware security theory, circuit design, and system architecture, including experts in provable security and information-theoretic analysis, to foster interdisciplinary discussion. Particular emphasis was placed on integrating pre-silicon modeling with post-silicon countermeasures. The goal was to advance a unified understanding of practical and provable security.

Overview of the meeting

The meeting consisted of technical presentations and brainstorming sessions on physical attack models and countermeasures. Topics included side-channel leakage, fault injection, and emerging attack surfaces such as the silicon substrate backside. Several presentations highlighted how advanced packaging technologies can both introduce vulnerabilities and enable defenses. Discussions emphasized combining pre-silicon simulation with post-silicon protection strategies, as well as the need for security-aware design tools and evaluation frameworks. Case studies showed how modeling can predict leakage and fault susceptibility. The meeting also included discussions on AI and its possible impact on hardware security research. Overall, the meeting fostered productive exchanges and strengthened collaboration toward secure system design.

Overview of Talks

Si Substrate Backside of ICs as Surfaces and Countermeasures for Fault Attacks

Makoto Nagata, Kobe University, Japan

IC chip's backside, more precisely, the backside surface of a silicon substrate where ICs are manufactured on its frontside, provides open spaces for performance improvements as well as for adversarial security attacks, however, potentially leads to trade offs between performance and security. An attacker leverages Si substrate backside to scan side channel leakages from and also to inject intentional faults to a crypto processor, through a variety of physical interactions with electromagnetic, electrical, thermal and luminescent medias. This talk discusses the usage of Si backside for and against physical attacks. The IC chip backside will be experimentally explored for physical attacks, captured in system-wide simulation models for pre-Si verification, and also investigated for potential countermeasures with advanced packaging technologies.

Provably Secure, Practically Broken: Advanced Physical Attacks Beyond the Model

Shahin Tajik, Worcester Polytechnic Institute, USA

Side-channel and fault security models typically assume that an adversary's capabilities are bounded within a clock cycle, for example, limited to a fixed number of probes, restricted in the number of injectable faults, or unable to reposition probes between cycles. These bounded models underpin many formal proofs of side-channel and fault resistance. In practice, such assumptions have long appeared reasonable. The number and mobility of probes in conventional attacks are constrained by the physical complexity of the measurement setup. Moreover, tampering with the system clock to halt or sufficiently slow execution, so that probes can be moved between cycles, has historically been considered impractical. In this talk, we argue that these foundational assumptions no longer hold in the presence of recently developed advanced physical attacks. We first review static backscattering side-channel techniques, such as laser logic-state imaging and impedance analysis, as well as frequency-selective glitching attacks. These methods effectively provide an unlimited number of contactless probes for both state extraction and localized fault injection. We then demonstrate, for the first time, that the system clock can be indirectly halted, enabling full-circuit state extraction and fundamentally violating bounded-adversary models. Based on these findings, we discuss why we should shift our focus from impending increasingly powerful physical attacks to designing hardware that detects and responds to them at the physical level.

When Provable MPC Meets Physical Leakage: Side-Channel Security of Privacy-Preserving Computation

Fatemeh Ganji, Worcester Polytechnic Institute, USA

Secure multiparty computation (MPC) allows mutually distrustful parties

to compute on private data without revealing anything beyond the intended output. While MPC now benefits from decades of theoretical development and increasingly mature implementations, the security of real systems depends not only on protocol design but also on the security of the underlying implementation. In practice, software and hardware implementations of MPC can leak sensitive information through timing, power, and electromagnetic side channels, thereby creating an attack surface not captured by standard protocol-level security guarantees.

This talk focuses on side-channel attacks against MPC implementations and on the challenges of securing privacy-preserving computation in realistic deployment settings. By revisiting MPC from an implementation-security perspective, the talk aims to connect provable cryptographic security with the physical realities of deployed systems and to motivate further research on trustworthy privacy-preserving computation.

The AI Frontier in Hardware Security: Automated Vulnerability Discovery vs. Generative Attacks

Jeyavijayan (JV) Rajendran, Texas A&M University, USA

The integration of Artificial Intelligence into the hardware design cycle has introduced a paradigm shift in the adversarial landscape. This talk explores the "dual-use" nature of Large Language Models (LLMs) and Machine Learning in hardware security. We first examine the defensive frontier: using AI-driven frameworks to automate the detection of subtle security vulnerabilities in RTL and HLS designs that manual inspection and traditional CAD tools often miss. Conversely, we pivot to the adversarial perspective, demonstrating how AI can be leveraged to synthesize stealthy hardware Trojans or automate side-channel analysis with unprecedented efficiency. By analyzing the efficacy of AI-based "Red Teaming" against modern silicon, we will discuss the urgent need for robust, AI-resilient hardware primitives and the future of self-healing hardware architectures.

RTL Simulation for Masked Software Verification

Quentin Meunier, Sorbonne University, France

Masking is a countermeasure against Side-Channel Attacks (SCA) that aims to ensure that intermediate computations in an algorithm have secret-independent distributions through the use of random variables. This theoretically prevents SCAs, as power consumption is directly linked to the values manipulated by the program or hardware device. Designing a masking scheme is often non-trivial, and a critical need for automated masking verification has emerged to ensure the correctness of masked implementations, both in hardware and in software. In the first part of the presentation, I will introduce work on masking schemes verification, with a focus on the tool VerifMSI. Then, based on the observation that software implementations proven secure can still lead to observable secret leakages, the considered intermediate computations have progressively shifted closer to the hardware level. I will then present a work called Armistice, in which the datapath of a Cortex-M3 core was modeled to capture all relevant

internal expressions that need to be verified when running masked software. Finally, I will present how this work was extended to take as input the HDL of the processor core, resulting in a tool called aLeakator. The latter combines traditional HDL simulation with symbolic expression-based simulation in order to verify the absence of secret leakage on all wires of the HDL model. Comparisons with measured power traces from Cortex-M3 and Cortex-M4 cores have been performed and demonstrate the relevance of the approach.

Should We Trust Complex and Heterogeneous SoCs?

Lilian Bossuet, University Jean Monnet, Saint-Etienne, France

Over the past two decades, the complexity and performance of heterogeneous system on chips (SoCs)–comprising multiple processor cores and programmable logic–has greatly increased. But what about their security? Can we trust them? These essential questions give rise to others, because these SoCs, which closely link software and hardware, create an unprecedented vulnerability space. Is the software separation provided by TrustZone technology sufficient? What can an attacker corrupt in these systems? Can the software part of these systems be attacked from the hardware part? Are there covert channels that bypass the secure isolation of these systems? Can physical attacks (SCA and FIA) targeting cryptographic functions (software or hardware) be carried out remotely on such systems? Drawing on research published over the past ten years, this presentation attempts to answer these questions and identify lessons that can be learned to enhance the security of complex, heterogeneous system on chips (SoCs).

When Energy Meets Security: Internal Energy-Based Attacks

Maria Méndez Real, Université Bretagne-Sud, France

Performance optimization has been widely exploited to compromise embedded systems security. After performance, energy optimization is one of the most important constraints and can also provide ways to disrupt the system, observe it, and extract information that would otherwise be secret. For the past decade, energy optimization techniques such as DVFS, shared between sensitive and potentially malicious applications, have been exploited to inject timing faults. More importantly, these attacks do not require physical access to the system and are carried out from software. More recently, specially designed circuits in FPGA logic have also been used to induce voltage drop attacks to inject faults that can even propagate to processor in heterogeneous CPU+FPGA systems. In this presentation, we will introduce these attacks and discuss the challenges for mitigation.

Evaluation of Hardware Masked Circuits, From Conservative Models to Heuristic Approaches

Amir Moradi, Technical University of Darmstadt, Germany

Realization of masking in hardware has been always a challenging issue.

Conceptual flaws as well as engineering mistakes would potentially lead to delivering insecure implementations. Hence, evaluation of such designs prior to fabrication seems to be essential, while being solely dependent on experimental analyses does not necessarily always lead to secure circuits. This talk will give an overview of the relevant challenges, known solutions, their difficulties and shortcomings, and most importantly open problems which need attention in the near future.

Trace-Efficient Transformer Models for Side-Channel Analysis of Masked AES

Elif Bilge Kavun, Barkhausen Institut & Dresden University of Technology, Germany

Deep learning has become a powerful tool for side-channel analysis (SCA), but many existing architectures struggle with masked implementations and trace misalignment. This talk presents a transformer-based late-fusion approach that jointly processes side-channel traces and plaintext information to predict masked AES Sbox values. The architecture combines a pretrained language-model pathway for plaintext with a trace embedding module and leverages attention to learn leakage relationships. Experiments on the ASCAD variable-key dataset show highly trace-efficient attacks, achieving key recovery with only nine traces in the aligned setting and maintaining robustness under significant desynchronization. The talk also briefly discusses earlier unpublished transformer-based SCA results obtained from the DPAv2 dataset (unprotected hardware implementation of AES), highlighting practical challenges and insights when applying deep learning-based SCA to real-world implementations.

New Sensors for Remote Power Attacks

Sri Parameswaran, The University of Sydney, Australia

Deep devastation is felt when privacy is breached, personal information is lost, or property is stolen. Now imagine what happens when all of this occurs at once, and the victim is unaware of it until much later. This is the reality: an increasing number of electronic devices are used as keys, wallets, and files. Security attacks targeting embedded systems illegally gain access to information or destroy information. Advanced Encryption Standard (AES) is used to protect many of these embedded systems. While mathematically proven to be quite secure, it is now well known that AES circuits and software implementations are vulnerable to side-channel attacks. Side-channel attacks exploit the observation of system properties (such as power consumption, electromagnetic emissions, etc.) while the system performs cryptographic operations. In this talk, Remote power attacks are described, and novel sensors are demonstrated, which enable stealthy deployment of attacks on Cloud-based FPGAs (such as the Amazon Cloud FPGA).

Towards Cost-Efficient Share-Reduced Masking: A Combined Countermeasure Against Side-Channel and Fault Attacks

Haruka Hirata, The University of Electro-Communications, Japan

Physical attacks such as side-channel analysis (SCA) and fault analysis (FA) have become realistic and practical threats to cryptographic implementations. Moreover, these attacks are often combined to extract secret keys. Therefore, designing countermeasures that simultaneously provide resistance against both SCA and FA is essential. However, existing combined countermeasures typically incur significant implementation overhead due to duplication-based constructions and high-order masking implementations. In this work, we present a cost-efficient design methodology for combined countermeasures based on share-reduced masking integrated with a multiplicative MAC tag for fault detection. Our design reduces the number of shares while preserving resistance against both information leakage and injected faults. We provide a general construction framework and demonstrate its applicability to an AES S-box implementation, showing improved cost-efficiency compared to existing approaches.

TBA

Francesco Regazzoni, University of Amsterdam, the Netherlands & Universita' della Svizzera Italiana, Switzerland

TBA.

Information-Theoretic Analysis of Side-Channel Attacks and Provable Security of Masking

Rei Ueno, Kyoto University, Japan

In evaluating side-channel attacks (SCAs), the success rate (SR) is a major theoretical metric to evaluate the capability of SCA, which would correspond to the adversarial advantage. In this talk, starting from an information-theoretic SR bound by de Chérisey et al. in TCHES 2019, we discuss two information-theoretic aspects of SCA. (1) Deep-learning based SCA (DL-SCA) is the strongest SCA in the current situation. Perceived information (PI) is an information-theoretical metric to evaluate SR of DL-SCA, yet it has not been proven that PI is an SR upper-bound of DL-SCA. We show that PI is not an upper-bound of SR, and present information-theoretical SR upper-bounds of DL-SCA by another information-theoretical metrics. (2) We discuss the security of major SCA countermeasure, masking. We develop a communication channel model representing the SCA on masked cryptographic implementation, and prove an SR upper-bound based on it. The bound is followed by a security proof of higher-order masking, meaning that the SR of SCA on d-th order masking decreases exponentially by an increase of d.

Active Side-Channel/TEMPEST Attacks Using Electromagnetic Waves

Yuichi Hayashi, Nara Institute of Science and Technology, Japan

Electromagnetic (EM) emanations from electronic devices have long been known to cause unintended information leakage, forming the basis of passive EM side-channel and TEMPEST attacks. In these attacks, adversaries observe EM emissions from a target device to infer sensitive information such as processed data or secret keys. In this talk, we first review the principles and representative examples of such passive EM attacks. We then introduce a newer class of threats: active EM attacks.

Bridging Theory and Hardware: Circuit Techniques to Support Masking Assumptions

Takeshi Sugawara, The University of Electro-Communications, Japan

Masking schemes proven secure under the probing model ensure resistance against side-channel attacks; however, the underlying physical assumptions are not automatically satisfied in practice. First, interactions between adjacent circuit elements (i.e., coupling) can degrade the effective protection order. Furthermore, although high-order masking forces attackers to exploit higher-order statistical moments, achieving a target security level (e.g., in terms of MTD) still requires a sufficiently low signal-to-noise ratio (SNR). In this work, we highlight the importance of circuit-level design techniques to satisfy these assumptions by presenting our implementation of the high-order masking scheme HPC2 using the Wave Dynamic Differential Logic (WDDL) technique. WDDL effectively reduces the data-dependent signal component in the SNR while mitigating the impact of coupling effects.

List of Participants

- Lilian Bossuet, University Jean Monnet, Saint-Etienne, France
- Fatemeh Ganji, Worcester Polytechnic Institute, USA
- Benedikt Gierlichs, KU Leuven, Belgium
- Yuichi Hayashi, Nara Institute of Science and Technology, Japan
- Haruka Hirata, The University of Electro-Communications, Japan
- Mike Hutter, University of the Bundeswehr Munich, Germany
- Elif Kavun, Barkhausen Institut & Dresden University of Technology, Germany
- Yang Li, The University of Electro-Communications, Japan
- Maria Méndez Real, Université Bretagne-Sud, Lab-STICC, France
- Quentin Meunier, Sorbonne University, France
- Amir Moradi, Technical University of Darmstadt, Germany
- Makoto Nagata, Kobe University, Japan
- Artemii Ovchinnikov, KU Leuven, Belgium
- Sri Parameswaran, The University of Sydney, Australia
- Jeyavijayan Rajendran, Texas A&M University, USA
- Francesco Regazzoni, University of Amsterdam, the Netherlands & Universita' della Svizzera Italiana, Switzerland
- Patrick Schaumont, Worcester Polytechnic Institute, USA
- Takeshi Sugawara, The University of Electro-Communications, Japan
- Shahin Tajik, Worcester Polytechnic Institute, USA
- Rei Ueno, Kyoto University, Japan
- Ingrid Verbauwhede, KU Leuven, Belgium

Meeting Schedule

Check-in Day: March 15 (Sun)

- 19:00–21:00 Welcome Reception

Day1: March 16 (Mon)

- 09:00–10:00 Opening Introduction + Fault Adversaries and Tools
- 10:00–10:30 Coffee Break
- 10:30–12:00 Session 1: Side-Channel Adversaries and Tools (1)
- 12:00–13:30 Lunch
- 13:30–15:00 Session 2: Real World Adversary Models (1)
- 15:00–15:30 Coffee Break
- 15:30–17:00 Session 3: Physical Attacks
- 18:00–21:00 Dinner & Networking

Day2: March 17 (Tue)

- 09:00–09:45 Session 4: Real World Adversary Models (2)
- 09:45–10:15 Coffee Break + Group Photo
- 10:15–12:00 Session 5: Brainstorming
- 12:00–13:30 Lunch
- 13:30–15:00 Session 6: Side-Channel Adversaries and Tools (2)
- 15:00–15:30 Coffee Break
- 15:30–17:00 Session 7: Fault-and-Side-Channel Combined Adversaries and Countermeasures
- 18:00–21:00 Dinner & Networking

Day3: March 18 (Wed)

- 09:00–10:45 Session 8: Discussion Side-Channel Adversaries and Fault Adversaries Tools
- 10:45–11:15 Coffee Break
- 11:15–12:00 Session 9: Side-Channel Adversaries and Countermeasures
- 12:00–13:30 Lunch
- 13:30–21:00 Excursion & Banquet

Day4: March 19 (Thu)

- 09:00–10:30 Session 10: Side-Channel Adversaries and Tools (3)
- 10:30–11:00 Coffee Break
- 11:00–12:00 Closing Discussion
- 12:00–13:30 Lunch

Summary of discussions

Si Substrate Backside of ICs as Surfaces and Countermeasures for Fault Attacks (Makoto Nagata): The discussion focused on the practicality and limitations of backside attacks as well as corresponding countermeasures. Backside thinning was examined as a realistic attack step, already used in laser fault injection, while its cost, time, and potential detectability through physical changes such as capacitance variation were highlighted. The mechanical difficulty introduced by additional bonded wafers was also discussed as a factor affecting attack feasibility.

The effectiveness of shielding-based countermeasures was critically analyzed. Backside metal shielding was discussed as limited, particularly against magnetic fields and in the presence of unavoidable structural gaps. Shielding was framed as a frequency-dependent problem, similar to Faraday cage design, rather than a binary protection. The feasibility of injecting not only pulse-based glitches but also sinusoidal or frequency-specific signals was also discussed, expanding the attack model beyond conventional assumptions.

Measurement and evaluation aspects were also examined. Trade-offs between correlation and mutual information as leakage metrics were discussed in terms of computational cost and generality. Probe design was identified as a key factor, involving trade-offs between sensitivity, spatial resolution, and locality, and requiring careful parameter tuning. Overall, the discussion highlighted both the feasibility of backside-based attacks and the complexity of accurately modeling, measuring, and mitigating them.

Provably Secure, Practically Broken: Advanced Physical Attacks Beyond the Model (Shahin Tajik): The discussion focused on the behavior of circuits under voltage reduction and its implications for data integrity. Operation in the brownout region was examined, highlighting that data may remain stable within certain voltage ranges but becomes unreliable under more extreme conditions. The feasibility of exploiting this region for controlled extraction was discussed, particularly through repeated clock manipulation and fine-grained execution control.

Limitations of existing countermeasures were also analyzed. Detection mechanisms such as brownout sensors were shown to be insufficient if their response depends on clock-driven logic, introducing a delay between detection and reaction. This gap allows attacks to bypass protection mechanisms. The fundamental reliance on digital control paths for responding to analog events was identified as a structural weakness.

Physical attack feasibility and practical constraints were further explored. Backside and electromagnetic probing techniques were discussed in terms of spatial resolution, measurement time, and stability issues. While technically feasible, such attacks require careful tuning and long acquisition times. Overall, the discussion emphasized that physical-layer effects can invalidate assumptions underlying provable security models.

When Provable MPC Meets Physical Leakage: Side-Channel Security of Privacy-Preserving Computation (Fatemeh Ganji): The discussion focused on the gap between protocol-level security guarantees of MPC and their physical implementations. While MPC provides provable security under ideal assumptions, implementation-level leakages such as timing variations and non-constant-time behavior can violate these assumptions. This highlights

the limitation of relying solely on protocol-level analysis without considering physical execution.

The interaction between distributed computation and side-channel leakage was also examined. Even when data is secret-shared, intermediate computations and execution patterns may expose information through side channels. Repeated executions and statistical analysis can amplify such leakage, enabling practical attacks despite theoretical protections. The role of microarchitectural effects and shared resources was identified as a key factor in this gap.

The discussion emphasized the need to integrate side-channel resistance into MPC implementations. Techniques such as constant-time execution and careful memory access control were identified as necessary but not sufficient. Overall, bridging the gap between theoretical security models and real-world implementations was highlighted as a central challenge for privacy-preserving computation.

The AI Frontier in Hardware Security: Automated Vulnerability Discovery vs. Generative Attacks (Jeyavijayan Rajendran): The discussion focused on the limitations of current evaluation methodologies based on static threat models. Evaluating defenses against predefined scenarios or random attacks was identified as insufficient, as real attackers adaptively target weaknesses. This highlighted the risk of a "false sense of security" when defenses are validated only within fixed assumptions, motivating the need for more realistic evaluation frameworks.

The role of AI in constructing adaptive adversaries was examined. Reinforcement learning was discussed as a means to model sequential attack processes, such as selecting trigger conditions for hardware Trojans or identifying effective fault injection points. AI-driven approaches demonstrated the ability to generate attacks that evade multiple defenses simultaneously, but their lack of interpretability was noted as a limitation in understanding and validating the resulting strategies.

Scalability and practical deployment were also key concerns. The importance of reducing the action space to make reinforcement learning tractable was emphasized, along with the feasibility of applying these techniques to larger designs. The possibility of evaluating multiple defenses in a unified framework using AI was also discussed. Overall, the discussion emphasized a shift toward adaptive, AI-driven adversarial evaluation and reconsideration of how security is assessed.

RTL Simulation for Masked Software Verification (Quentin Meunier): The discussion focused on the scope and applicability of RTL-based verification across both software and hardware. The approach was characterized as general, as it operates on expressions derived from execution or circuit behavior, avoiding the need for manual abstraction. Using synthesized RTL as input was highlighted as a way to improve reliability and eliminate inconsistencies between high-level descriptions and implementation.

The ability to capture microarchitectural leakage sources was a central topic. RTL-based analysis naturally includes effects such as multiplexers, datapath interactions, and pipeline behavior, which are often missed in higher-level models. Stability analysis was discussed as a key technique to reduce false positives by identifying signals that actually change between cycles. Differences between theoretical guarantees and practical leakage behavior were also highlighted, particularly in optimized versus manually secured implementations.

Scalability remains a major challenge. While first-order verification is practical, higher-order analysis leads to combinatorial explosion due to the large number of signals and cycles in processor-level designs. Portability across architectures also requires non-trivial effort when adapting the framework. Overall, the discussion emphasized trade-offs between accuracy, scalability, and practical deployment of RTL-based verification methods.

Should We Trust Complex and Heterogeneous SoCs? (Lilian Bossuet): The discussion focused on the effectiveness and limitations of on-chip sensing mechanisms used for remote attacks. Multiplier-based sensors and TDC-based designs were compared, with both achieving similar resolution through different physical principles such as carry propagation and delay chains. Implementation choices, including DSP-based versus CLB-based designs, were discussed in terms of efficiency and detectability.

The security implications of control asymmetries in heterogeneous SoCs were also examined. In particular, frequency and clock control mechanisms enable covert channels, with hardware-side control offering fewer restrictions than software-side control. This highlighted fundamental weaknesses in current isolation assumptions. The discussion also connected these issues to multi-tenant FPGA scenarios, where similar attack vectors may arise.

A central theme was the limitation of logical isolation in the presence of shared physical resources. The shared power delivery network was identified as a persistent channel for information leakage, regardless of higher-level protections. Overall, the discussion emphasized that increasing system complexity and resource sharing inherently weaken isolation guarantees in modern SoCs.

When Energy Meets Security: Internal Energy-Based Attacks (Maria Méndez Real): The discussion focused on the practicality and robustness of voltage-drop-based attacks and their detection. Detection mechanisms based on voltage monitoring were examined, highlighting challenges due to dependence on sensor placement, environmental conditions, and system variability. Variance-based metrics were identified as more robust than fixed thresholds for capturing anomalous behavior.

Architectural aspects of attack propagation were also discussed, particularly the role of shared power delivery networks. Voltage disturbances were shown to propagate across domains in current FPGA-based systems, enabling cross-component interference. The potential role of integrated voltage regulators was considered, but their effectiveness in isolating such effects remains uncertain. The timing precision of DVFS-based attacks and the need for reliable synchronization were also key points.

A fundamental limitation identified was the gap between detection and prevention. Faults can be injected before detection mechanisms react, making detection alone insufficient. Possible directions such as securing the energy interface or introducing stronger physical isolation were discussed, but raise challenges in cost and practicality. Overall, the discussion emphasized the need for system-level approaches that move beyond detection toward effective prevention.

Evaluation of Hardware Masked Circuits — From Conservative Models to Heuristic Approaches (Amir Moradi): The discussion focused on how transitions and previous signal values are incorporated into leakage analysis. Leakage was characterized as fundamentally arising from transitions, requiring both current and previous values to be considered. Practical implica-

tions include the need for additional registers or ensuring stable inputs across consecutive cycles, with registers acting as barriers to glitch propagation.

The distinguishability of glitch patterns and the role of circuit-level effects were also examined. Even seemingly similar transitions were considered distinguishable due to differences in routing delays and gate propagation characteristics, which accumulate across cascaded logic. Process variation was discussed as a factor that further increases distinguishability rather than masking leakage, supporting a conservative modeling approach.

Scalability and evaluation methodology emerged as key challenges. Instead of explicit formula expansion, probe conditions are propagated incrementally through the circuit, but the number of probes and required simulations still grows significantly. The evaluation process requires forward propagation through the circuit rather than backtracking from outputs. Overall, the discussion highlighted trade-offs between model accuracy, conservativeness, and scalability in leakage evaluation.

Brainstorming (1) – AI and the Future of Hardware Security Research: The discussion examined the role of AI in hardware security research, particularly what aspects of research are unlikely to be replaced. It was emphasized that experimental work requiring physical interaction—such as measurements and probe setup—remains essential, although future automation was also acknowledged. It was also pointed out that current AI lacks physics-informed reasoning and mainly extrapolates from existing data, making it less suited for generating fundamentally new insights. The distinction between incremental research and genuinely original contributions was highlighted, with concerns that current evaluation practices may favor work that is more easily automated.

The session also explored how AI could support research activities. Common suggestions included handling administrative tasks, assisting with coding and evaluation, accelerating simulations, and improving literature search. However, it was noted that while AI can summarize prior work, it often misses subtle flaws or insights that inspire new research. The discussion further addressed whether surveys remain valuable, concluding that interpretation and identification of open problems remain human-driven tasks.

Another major focus was the application of hardware security techniques to AI systems. Key assets requiring protection include training data, model weights, and system integrity. It was noted that many existing techniques—such as side-channel analysis, fault injection, and Trojan detection—can be applied to AI hardware, although GPUs and large-scale systems introduce new challenges. Additional topics included geolocation verification, self-exfiltration risks, and broader questions at the boundary between security and safety.

The discussion also considered future directions, including the development of security-aware EDA tools, automated verification frameworks, and design methodologies that reduce leakage sources. Ideas such as AI-assisted security design, IP protection in an AI-driven design flow, and hardware-level control mechanisms were proposed. Finally, open problems such as distinguishing attacks from noise in detection systems and concerns about whether the field is reaching a plateau were discussed, with the conclusion that AI should be leveraged to accelerate progress rather than viewed as a replacement for human researchers.

Trace-Efficient Transformer Models for Side-Channel Analysis of Masked AES (Elif Bilge Kavun): The discussion focused on the reliabil-

ity and reproducibility of deep-learning-based side-channel analysis, particularly when using transformer architectures. A key issue identified was the strong dependence on random seed initialization, which can lead to significant variability in results. Fixing the seed was recognized as essential for reproducibility, while differences across datasets and evaluation setups highlighted the difficulty of making fair comparisons between models.

The impact of data conditions and preprocessing choices was also examined. The role of masking information during profiling, as well as the influence of trace alignment and desynchronization, were discussed in terms of their effect on attack performance. Fixed shifts in desynchronization were shown to sometimes create favorable leakage conditions, and results depend strongly on the specific subset of traces used. This raised broader concerns about evaluation bias and the representativeness of commonly used benchmarks.

Scalability and generalization were also emphasized. Model performance was shown to depend on platform-specific characteristics such as noise and implementation style, making cross-platform generalization non-trivial. The need for systematic validation across datasets, including techniques such as cross-validation and testing under varying measurement conditions, was highlighted. Overall, the discussion underscored challenges in reproducibility, evaluation methodology, and generalization for DL-based side-channel analysis.

New Sensors for Remote Power Attacks (Sri Parameswaran): The discussion focused on the nature of measurements performed by on-chip sensors and their interpretation. Voltage fluctuations on the power delivery network were clarified as the primary observable, with power consumption inferred indirectly through these variations. The effectiveness of different sensor types was examined, including trade-offs in spatial coverage, aggregation, and detectability, particularly for ring oscillators and compact delay-based sensors.

The robustness of sensing under realistic operating conditions was also discussed. The presence of other active circuits introduces noise, but uncorrelated activity was shown to primarily increase the number of required traces rather than prevent attacks. Calibration mechanisms were identified as critical for maintaining sensitivity, especially for low-resolution sensors. The feasibility of extracting meaningful information from minimal (e.g., single-bit) measurements was highlighted, relying on statistical accumulation over many traces.

Scalability and broader applicability were key concerns. Translation to ASIC platforms was identified as a major challenge due to the lack of PVT-invariant delay elements comparable to FPGA primitives. The dual-use nature of sensors was also emphasized, as the same mechanisms can enable both attacks and countermeasures. Overall, the discussion highlighted trade-offs between stealthiness, robustness, and portability of sensor-based side-channel techniques.

Towards Cost-Efficient Share-Reduced Masking: A Combined Countermeasure Against Side-Channel and Fault Attacks (Haruka Hirata): The discussion focused on how share reduction affects fault detection in the combined masking and multiplicative MAC scheme. Reusing shares between the value and tag computations introduces a vulnerability where a single fault injected on a shared component propagates identically to both paths. As a result, fault detection can fail under specific conditions, making such attacks easier compared to standard duplicated implementations that require independent faults.

The probing security of the share-reduced constructions was also examined.

The 5-share implementation preserves second-order probing security despite the shared structure, while the 4-share version introduces leakage due to common internal variables between value and tag computations. This effectively reduces the security order, which is also reflected in leakage evaluation results showing first-order leakage in the reduced design.

The discussion emphasized the fundamental limits of share reduction. Maintaining a given security order requires a minimum number of independent shares, restricting how much reduction is possible without degrading security. It was also noted that the current evaluation considers side-channel and fault attacks separately, and that combined attack scenarios remain an open issue.

TBA — Combined Side-Channel and Fault Attack Countermeasures (Francesco Regazzoni): The discussion focused on concrete examples of combined attacks. These included using side-channel information to identify when and where to inject faults, and injecting faults to weaken masking (e.g., targeting the random number generator) before exploiting leakage. Statistical ineffective fault analysis (SIFA) was also discussed as a case where the attacker distinguishes whether a fault had an effect, rather than relying on faulty outputs.

The ordering of countermeasures was discussed through specific examples. Applying masking first exposes the randomness source as a target for fault injection, while applying fault detection first leads to issues when masking is applied to the fault-detection signal. It was pointed out that the random number generator is typically implemented as a separate module communicating over a bus, making both the generator and its interface potential targets. Error-detecting codes were also discussed, including cases where detecting faults can still enable information leakage when combined with side-channel observations.

The discussion also addressed system-level aspects beyond circuit-level countermeasures. Detecting attack preparation, such as connecting probes or modifying the power delivery path, was considered as an additional direction. Questions were raised about how systems should react after attack detection, including key deletion and the risk of denial-of-service. The limitations of sensor-based approaches and the lack of clear security models for them were also discussed, along with the role of attacker cost and scalability in evaluating practical threats.

Brainstorming (2) – Side-Channel and Fault Adversary Tools: The discussion focused on what kinds of tools are needed to evaluate and design secure systems. Two complementary roles of tools were identified: assisting designers in evaluating security from an adversarial perspective, and enabling efficient attack strategies that inform realistic threat models. A key point was that these two directions should feed into each other, leading to a systematic framework that connects threat models, evaluation methods, and design guidelines. It was also emphasized that tools should assist designers in handling complexity rather than replacing core expertise.

A major theme was the level of abstraction at which security tools should operate. While existing EDA flows span from high-level design to physical implementation, current security tools are typically limited to specific abstraction levels and are not integrated into these flows. It was pointed out that proving security at a higher abstraction level is necessary but not sufficient, as leakage may appear after synthesis or in the physical implementation. The gap between pre-silicon models and post-silicon behavior was repeatedly emphasized, particularly due to effects from PCB, package, and system-level integration that are not captured in current models.

Another central issue was the lack of appropriate models and metrics for security evaluation. Existing power models in EDA are designed for average power estimation and are insufficient for capturing data-dependent leakage. At higher abstraction levels, simple models can be used for early screening, but accurate evaluation requires detailed physical modeling, which remains difficult and inaccessible. The absence of standardized measurement methodologies was also highlighted, as experimental setups vary widely and no common baseline exists.

The discussion also addressed practical constraints from industry. There is a fundamental gap between provable security in academic models and threshold-based evaluation used in certification. Industry focuses on the total effort required to break a system rather than the formal attack order, making it difficult to compare different countermeasures. It was also noted that tools must integrate with existing EDA ecosystems to be adopted, and that interoperability and standard interfaces are critical challenges.

Finally, the role of AI and future directions were discussed. AI was considered useful for design space exploration and optimization across multiple parameters, but not suitable for providing security guarantees, which require rigorous and explainable methods. A promising direction is iterative workflows combining AI-generated designs with formal evaluation. Broader community-level needs were also identified, including shared platforms for tools and benchmarks, security specifications for system integration, and the incorporation of security as a design constraint alongside power, performance, and reliability.

Information-Theoretic Analysis of Side-Channel Attacks and Provable Security of Masking (Rei Ueno): The discussion focused on the relationship between success rate (SR) and information-theoretic metrics in side-channel analysis. It was clarified that minimizing cross-entropy in deep-learning-based SCA corresponds to an optimal distinguisher in the asymptotic sense, but does not uniquely define optimality. The de Chérisey inequality relating SR to mutual information was discussed as a fundamental result, while the practical difficulty of computing mutual information for high-dimensional traces was emphasized.

The validity of Perceived Information (PI) as an SR upper bound was critically examined. It was shown that PI depends on cross-entropy, which can be arbitrarily changed through inverse temperature scaling without affecting SR. This demonstrates that PI cannot serve as an upper bound. Effective Perceived Information (EPI) was introduced to remove this ambiguity by minimizing cross-entropy, providing a more stable estimate, although its formal validity as an SR bound remains unproven.

The discussion also covered Latent Perceived Information (LPI) and its advantages. By evaluating mutual information in the latent feature space of neural networks, LPI becomes computationally tractable and serves as a provable SR upper bound. The separation between feature extraction and classification was emphasized, with LPI enabling identification of inefficiencies in the classifier. The extension to masked implementations was also discussed, including frequency-domain analysis and tighter bounds on masking security based on leakage energy.

Active Side-Channel / TEMPEST Attacks Using Electromagnetic Waves (Yuichi Hayashi): The discussion focused on how leakage depends on frequency and measurement conditions. Frequency-domain analysis was dis-

cussed as a way to identify which frequency components carry information, with leakage strongly influenced by path and antenna characteristics. It was also noted that higher-speed devices shift leakage to higher frequencies, requiring more advanced measurement equipment.

The effectiveness and limitations of countermeasures were also examined. Pin-level filtering can suppress conducted emissions when leakage is concentrated at specific frequencies, but is less effective when signals spread across the PCB through coupling. Shielding was discussed as an effective countermeasure at the enclosure level, but connectors such as power and communication cables remain major leakage points. Grounding and impedance mismatches at connectors were identified as critical factors affecting leakage.

The discussion also addressed modeling assumptions and practical effects. While the presented model assumes linear behavior, real systems include nonlinear components that generate additional frequency components not present in the original signal. These nonlinear effects further complicate leakage behavior and make real-world EM side-channel analysis more complex than simplified models.

Bridging Theory and Hardware: Circuit Techniques to Support Masking Assumptions (Takeshi Sugawara):

The discussion focused on circuit-level issues arising from implementing masking assumptions in practice. In particular, replacing registers with self-retimers introduces data-dependent timing, since the enable signal is triggered by completion of the previous stage. This was identified as a potential leakage source, in contrast to clock-driven registers which provide data-independent timing. Possible mitigations include using delay chains to generate data-independent timing, at the cost of additional overhead.

Device- and technology-level constraints were also examined. Pass-transistor logic introduces challenges such as threshold voltage drop and limited drive strength, which become more severe in advanced process nodes. While the presented implementation works in 180nm technology, scaling to smaller nodes may require additional design techniques such as pull-up feedback, which impacts performance. The accumulation of timing variations when cascading multiple stages was also highlighted as a concern.

The discussion further addressed interpretation of experimental results and design methodology. Observed leakage in masked implementations raised questions about whether it originates from coupling effects or potential design issues, and a clear causality analysis remains open. The lack of scalable design flows was also emphasized, as the current approach relies on full-custom transistor-level design without EDA tool support. Extending the approach to larger designs such as full S-boxes or AES remains a significant challenge.

Summary of new findings

The meeting highlighted a growing gap between theoretical security guarantees and physical implementation realities. A key finding is that circuit- and system-level physical effects (e.g., coupling, PDN behavior, PCB/package interactions) can fundamentally undermine provable security models such as masking. Several discussions confirmed that protection order degradation occurs in practice due to unintended interactions beyond the abstraction level assumed in theory. Another important observation is that system-level deployment conditions significantly influence leakage behavior, challenging the validity of pre-silicon evaluations.

Significant progress was reported in attack methodologies and evaluation techniques. New compact on-chip sensors demonstrated that even minimal hardware resources (e.g., one-LUT sensors) are sufficient for effective remote side-channel attacks, while also revealing their dual-use nature as potential countermeasures (e.g., laser detection). In deep-learning-based SCA, new insights were obtained regarding reproducibility issues, particularly the critical role of random seed selection. Information-theoretic metrics such as LPI were introduced as more principled and tractable evaluation tools, while limitations of widely used metrics (e.g., perceived information) were clarified.

The meeting also revealed important trade-offs in countermeasure design. Share-reduction techniques demonstrated meaningful cost savings but introduced new vulnerabilities, particularly facilitating identical fault injection, highlighting an inherent trade-off between efficiency and robustness. Discussions on combined attacks showed that independent countermeasure design is insufficient, and that the ordering and interaction of countermeasures are critical design dimensions requiring holistic adversary models.

Finally, several new research directions emerged. Detecting attack preparation stages, rather than only reacting to attacks, was identified as a promising direction. Sensor-based approaches were recognized as effective but lacking formal security models and coverage guarantees. The importance of post-detection response strategies (e.g., graceful degradation versus key erasure) was emphasized as an underexplored area. Overall, the community recognized the necessity of bridging theory, hardware implementation, and system-level considerations, as well as the need for new EDA tools and models that incorporate security across multiple abstraction levels.

Identified issues and future directions

A key issue is the gap between theoretical security models and physical implementations. Assumptions used in proofs (e.g., independent leakage) are often violated due to coupling, PDN effects, and PCB/package interactions. Future work should develop models and evaluation methods that incorporate these physical effects across abstraction levels.

Another major challenge is the lack of unified adversary models. Current approaches treat attack vectors independently, while real attackers combine them. The interaction and ordering of countermeasures remain unclear. Future research should focus on holistic threat models and systematic evaluation of combined attacks.

Tool support and evaluation methodologies are also insufficient. Existing EDA tools and power models are not designed for security, and there are no standardized benchmarks or measurement protocols. Future directions include security-aware EDA integration and common evaluation frameworks.

Sensor-based detection raises additional issues. While promising, sensors lack formal security models and may be bypassed. Moreover, appropriate system responses after detection (e.g., avoiding denial-of-service) remain underexplored.

Finally, scalability and emerging platforms present challenges. Many techniques do not scale to complex systems or advanced technologies. Future work should address cross-layer security, scalable design methodologies, and effective use of AI for design exploration.