Next-Generation Anti-Tamper Envelopes forCyber Physical Defense Systems

- Extended Abstract -

Vincent Immler[1], Johannes Obermaier[1], Martin König[2], Matthias Hiller[1], and Georg Sigl[1,3]
[1]Fraunhofer Institute for Applied and Integrated Security (AISEC)
Parkring 4, 85748 Garching b. München, GERMANY
[2]Fraunhofer Research Institution for Microsystems and Solid State Technologies (EMFT)
Hansastraße 27d, 80686 München, GERMANY
[3]Chair of Security in Information Technology, Technical University of Munich
Arcisstraße 21, 80333 München, GERMANY

{forename.surname}@{aisec,emft}.fraunhofer.de

To assure successful missions, protecting the full stack of defense systems is mandatory, i.e., both hardware and software must be secured to provide the intended functions, prevent attacks, and continue to operate within a reasonable scope while under attack. As these systems often store classified data or otherwise sensitive information such as Intellectual Property (IP), this necessitates the use of cryptographic routines to encrypt and authenticate data while in transit and at rest. Consequently, proper handling and storage of Critical Security Parameters (CSPs), such as cryptographic key material is essential for device security. As a last resort this also entails deliberate destruction of sensitive information, i.e., zeroization of CSPs upon detection of a security breach which renders the device inoperable, if permissible according to mission objectives and required by the concept of operations (CONOPS).

Once a system is deployed in the field, rogue personnel and external adversaries can gain control over the device. Restricting malicious physical access in a hostile environment is therefore a complex task. If not counteracted, it is possible to carry out a wide range of physical attacks to obtain information from the device, e.g., by probing data lines on the Printed Circuit Board (PCB) or reading memory contents, either while powered-off or during runtime. More specialized attacks include side-channel analysis, fault injection, and delayering combined with optical analysis to extract stored cryptographic keys [3, 9]. Correspondingly, countermeasures are available that only protect against specific types of each of these attacks which makes the protection against the full range of potential attacks difficult.

In contrast to single-chip devices such as smartcards that can be protected in silicon, multiple-chip systems on a PCB can only be sufficiently secured by an additional physical security boundary that separates the secure and insecure domains of a system. Unfortunately, the rapid development of cyber physical systems that operate physically unattended outperforms the growth of generic solutions to mitigate the risk of physical attacks. Hence, developers seek solutions to seamlessly integrate security without impeding system's functionality or its development, especially for low to mid-range volume products that make a full-custom approach impractical. Typically, such a system comprises multiple chips on a PCB where only a fraction of the Integrated Circuits (IC) provides physical security countermeasures of varying degree. Therefore, an additional layer of protection is required to prevent any type of useful physical access and thereby reduce the risk of the aforementioned specialized attacks that could otherwise compromise device security.

The challenge in securing such systems was recognized by the National Institute of Standards and Technology

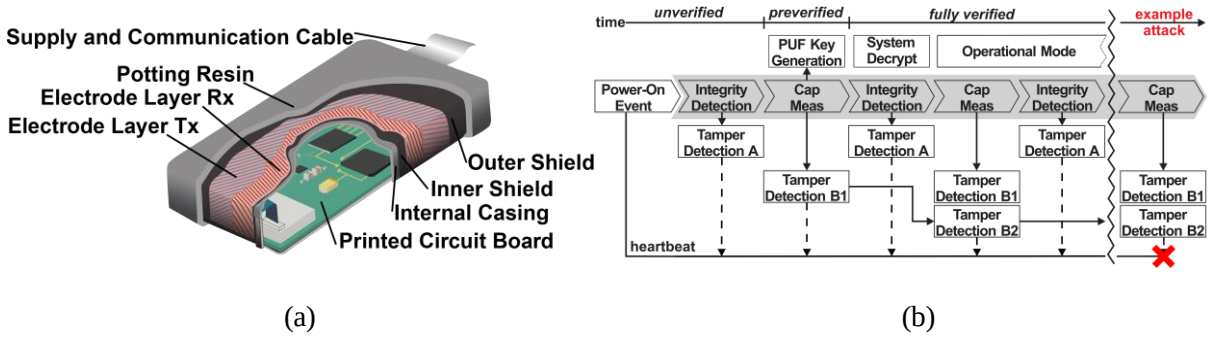

Figure 1: System overview: Figure 1a Conceptual view of a device enclosed by our envelope. Figure 1b Power-on sequence for a protected device employing our solution.

(NIST) and led to a government/ industry work group to provide defined levels of security which have been published as "FIPS 140: Security Requirements for Cryptographic Modules". For the most secure level, a "tamper detection/ response envelope with tamper response and zeroization circuitry" is required, i.e., an envelope that encloses the system and protects it from physical tampering by detecting intruders and other types of adversarial operating conditions. While FIPS 140-2 targets cryptographic modules only, the same principles apply to cyber physical systems that store sensitive information and are at risk of being physically attacked.

One previous approach to meet these requirements and that has been available commercially is the envelope by "GORE" which is made of a flexible polymer and contains a dense mesh of carbon-ink printed tracks [1]. Attempts to physically penetrate the mesh are very likely to destroy its tracks and result in open circuits. A continuous measurement from inside the system detects these open circuits and triggers an alarm that causes the zeroization of CSPs (i.e., tamper-detection and response). However, a battery is required for this monitoring mechanism whenever the supplementing carrier system of the protected module is powered off. Additionally, the CSPs must be stored in a volatile Battery-Backed Random-Access Memory (BBRAM) to enable instantaneous zeroization. Furthermore, a potting material is applied to conceal the wrapped module to further increase resistance towards attacks.

While conceptually providing a high level of security, this approach has significant practical drawbacks: adding a battery to the system increases bulk and weight, it lowers its robustness with regard to the device's operating temperature range, and prohibits prolonged storage. When the battery is fully discharged, the CSP are lost and physical integrity can no longer be guaranteed [2]. Moreover, storing CSPs in a BBRAM leaves room for the zeroization circuit to fail. However, storing a key in a non-volatile memory is also not an option as zeroization would be too slow and its contents can also be extracted while the system is powered off [10]. Alternatively, Physical Unclonable Functions (PUFs) can be used for key storage [4]. Once the device is running, this security primitive derives a cryptographic key from the device's inherent manufacturing variations, e.g., from the unique fingerprint-like start-up patterns of uninitialized SRAM. As long as the device is powered off, extracting these parameters is assumed to be difficult.

However, since most PUFs are implemented in an IC, it is impossible to use them for aftermarket protection of Commercial-Off-The-Shelf (COTS) components. Furthermore, silicon-based PUFs typically do not have the property of tamper-evidence [5], i.e., once powered on, they cannot verify if an attack was carried out on the system while powered off. Even worse, they are incapable of detecting online attacks that extract values during runtime [6] from, e.g., the data bus of the system, as they are typically just a component in a System-on-Chip.

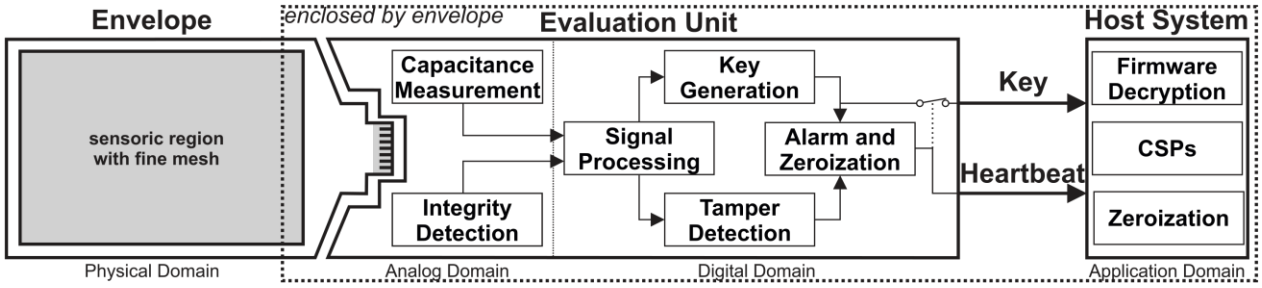

Figure 2: Tamper-resistant architecture based on envelope and evaluation unit.

To overcome these limitations, we present a novel, batteryless tamper-resistant envelope as shown in Figure 1a. After power-up, it checks its structural integrity (Tamper Detection A) and continues to verify itself similar to a tamper-evident PUF, i.e., if the system has not been tampered with, the correct key is derived from the unique physical properties of the envelope and the system's data is self-authenticated and decrypted. Afterwards, multiple mechanisms start in parallel to ensure continuous protection while the system is running, e.g., to check for the range of values (Tamper Detection B1) and their rate of change (Tamper Detection B2). This is illustrated in Figure 1b and exceeds previous concepts. It is therefore a solution towards meeting security standards, such as FIPS 140-2 Level 4, without a battery for the security mechanism. We consider this a next-generation anti-tamper envelope as it combines the properties of PUFs and tamper-respondent envelopes without the need of a battery-backed monitoring circuit.

To achieve these properties, our envelope contains an advanced mesh implementation, capable of not only detecting short and open circuits to ensure integrity of the mesh, but also of measuring the capacitances between tracks that are subject to manufacturing variation. These capacitive measurements are the basis to implement the tamper-evident PUF and enable a dual-approach with more sensitive integrity checks and secret key derivation. Hence, recovery of the key is only possible from inside the system as long as the envelope has not been tampered with. The challenge in successfully implementing this is, enclosing the PCB in a large-scale physical object while only using small-scale intrinsic variations for the PUF-based key derivation to make their extraction by an attacker improbable. Furthermore, a wider range of physical attacks must be taken into account that previously have been outside the scope of battery-backed approaches, as their security mechanism is never powered off. In this extended abstract, we refer to several conceptual and practical considerations of our design and how they relate to common security standards, briefly present its various components, and demonstrate its feasibility. Moreover, we explain how to incorporate this solution in a majority of cyber physical systems.

Our design [12] is based on an architecture that comprises three building blocks: the envelope, an evaluation unit, and the protected host system as illustrated in Figure 2. The envelope is manufactured using a custom thin-film technology which causes manufacturing-intrinsic variations in the envelope's mesh structure. These minuscule variations are then extracted by the measurement circuit [13] which is part of the evaluation unit. Subsequent processing of the data up to the PUF key generation is done by the evaluation unit or the host system, depending on the chosen device architecture which in our case, for development purposes, is based on a dedicated microcontroller running a customized version of FreeRTOS [11].

To have a fixed design goal, we limit the capabilities of the adversary to holes with a diameter of 0.3 mm/12 mil based on requirements of standards for security certification. While creating smaller holes might be possible, practically making use of them is difficult considering the shaft diameter of commonly available micro-probing

needles and related tools to get inside the envelope-wrapped enclosure, as depicted in Figure 1a. It is evident that this does not cover the full range of possible attacks. Thwarting additional attacks on a conceptual level was done in [12] and is beyond the scope of this article.

The envelope's capacitive structure is created from two layers of 16 Tx and 16 Rx electrodes that are used as capacitive sensors. These electrode layers are enclosed by a shield to provide a defined boundary condition for the capacitive measurement. A stochastic model is used to support the design process of envelopes of varying size. In order to extract local variations and ignore a coarse-grained global bias [14], the measurement focuses on exclusive Tx electrode pairs in close vicinity that are measured differentially against all Rx electrodes. In our proof-of-concept [12], this results in a total of $(16/2)\cdot 16 = 128$ differential capacitive nodes that are considered as the entropy source for the PUF key generation.

As part of preliminary testing, we currently use a custom discrete measurement circuit [13]. Its basic operating principle is to use two antiphasic excitation signals for each Tx electrode pair while the other Tx electrodes remain inactive, thereby creating an in-situ differential capacitance towards the Rx electrodes inside the envelope. This approach results in a current on the Rx electrodes representing the differential capacitance which is then further processed by analog circuitry before being sampled, filtered, and evaluated by an STM32 microcontroller. The obtained full-scale range is ±73 fF at a theoretical digital resolution of $\Delta M = 7.3$ aF which is however limited by circuit noise of $\sigma N = 0.2$ fF when the envelope is connected (without additional oversampling). Performing a single differential measurement can be completed in 0.6 ms. Since it can be parallelized on the Rx side for each TX pair, this results in only $(16/2)\cdot 0.6$ ms = 4.8 ms for the overall envelope.

The obtained data is then further processed to compensate for environmental influence which includes an equi-distant quantization scheme [7] and error-correcting code [8]. This results in an entropy of approx. 2.5 bit per differential capacitive node, i.e., a total of $128\cdot 2.5$ bit = 320 bit are extracted from the envelope prior to the error-correcting code with a per-symbol-error rate of less than 0.1% over the temperature range of -20 °C to +60 °C. The specifics of a suitable quantization and error-correcting code are described in [7, 8]. The thusly obtained entropy can be conveniently turned into a 128 bit cryptographic key with a failure rate of less than 10-6. We analyzed the quantized data using common PUF metrics, namely uniqueness and robustness to confirm the required properties, i.e., the envelopes indeed provide a unique data pattern that sufficiently differs from other envelopes while at the same time, providing robustness.

To also verify the tamper-evident properties of our enclosure, we attacked one of the envelopes using a 0.3 mm drill. Its design guarantees that at least one Tx and Rx electrode is destroyed by such an attack. Hence, detecting intrusions can already be done independently of the PUF-properties while not considering attempted repairs. Our experiments confirm that ≥ 80 bit of entropy are inherently destroyed by such a drilling attempt without dedicated zeroization mechanism. This observation is based on the significant change in the affected differential capacitance for each of the measurement nodes as direct result of the attack. This exceeds the threshold the error-correcting code can tolerate and therefore causes the key derivation to fail upon device start-up. Please note, attacks during runtime cause a heartbeat signal to stop that in turn triggers the active zeroization, as sketched in Figure 1b. Therefore, the attacker cannot recover any useful data from the system without either facing a significant computational effort when the device is powered-off or risking active detection during runtime. According to our understanding this should fulfill the definition of "zeroization" in the FIPS 140-2 standard which states: "a method of erasing electronically stored data, cryptographic keys, and CSPs by altering or deleting the contents of the data storage to prevent recovery of the data".

Additional tests conducted in the range of -20 °C to +60 °C confirm the chosen design rationale and already exceed the ambient temperature range of, e.g., the IBM 4765 PCIe crypto coprocessor which is only +10 °C to

+35 °C [2]. Further optimizing each part of the design and continued testing is an ongoing effort, i.e., these results are from an early stage of this work and should be considered preliminary. Several optimizations are currently under investigation: optimizing the material properties, transferring the discrete measurement circuit into an IC, developing more advanced error-correcting codes.

Putting the previous statements into perspective, we introduce a holistic approach to protect defense systems from the ground-up. Our approach should allow the replacement of previously used battery-backed approaches. Both envelope and security architecture have been developed to meet the highest levels of the existing security standards. We point out that the necessary concepts are generic and could also be implemented using other enclosure technologies. Our tests provide initial evidence that this concept fulfills the desired criteria in terms of tamper-resistance and operating range.

REFERENCES

[1] P. Isaacs, M. J. Fisher, K. Cuthbert and T. Morris Jr, "Tamper Proof, Tamper Evident Encryption Technology," in Pan Pacific Symposium, SMTA, 2013.

[2] IBM, "IBM 4765 Cryptographic Coprocessor Security Module," 2012.

[3] S. P. Skorobogatov, "Semi-invasive attacks -- A new approach to hardware," University of Cambridge, Computer Laboratory, Tech. Rep. UCAM-CL-TR-630, Apr. 2005.

[4] C. Herder, M. Yu, F. Koushanfar and S. Devadas, "Physical Unclonable Functions and Applications," Proceedings of the IEEE, vol. 102, 2014.

[5] R. Maes and I. Verbauwhede, "A discussion on the Properties of Physically Unclonable Functions," in TRUST, 2010.

[6] C. Helfmeier, D. Nedospasov, C. Tarnovsky, J. S. Krissler, C. Boit and J.-P. Seifert, "Breaking and entering through the silicon," in ACM CCS, 2013.

[7] V. Immler, M. Henning, L. K. Kürzinger and G. Sigl, "Practical aspects of quantization and tamper-sensitivity for physically obfuscated keys," in Workshop on Cryptography and Security in Computing Systems, ACM, 2016, pp. 12-18.

[8] V. Immler, M. Hiller, Q. Liu, A. Lenz and A.Wachter-Zeh, "Variable-Length Bit Mapping and Errror-Correcting Codes for Higher-Order Alphabet PUFs," Cryptology ePrint Archive, Report 2017/911, 2017. [Online]. Available: https://eprint.iacr.org/2017/911.pdf.

[9] S. H. Weingart, "Physical Security Devices for Computer Subsystems: A Survey of Attacks and Defense," in CHES, 2000.

[10] D. Samyde, S. Skorobogatov, R. Anderson and J. J. Quisquater, "On a new way to read data from memory," in First International IEEE Security in Storage Workshop., 2002.

[11] J. Obermaier, F. Hauschild, M. Hiller, and G. Sigl, "An Embedded Key Management System for PUF-based Security Enclosures," in 7th Mediterranean Conference on Embedded Computing (MECO), 2018.

[12] V. Immler, J. Obermaier, M. König, M. Hiller, and G. Sigl, "B-TREPID: Batteryless Tamper-Resistant Envelope with a PUF and Integrity Detection," in IEEE International Symposium on Hardware Oriented Security and Trust (HOST), 2018.

[13] J. Obermaier, V. Immler, M. Hiller, and G. Sigl, "A Measurement System for Capacitive PUF-Based Security Enclosures," in 55th ACM/EDAC/IEEE Design Automation Conference (DAC), June 2018

[14] V. Immler, M. Hiller, J. Obermaier, and G. Sigl, "Take a moment and have some t: Hypothesis testing on raw PUF data," in IEEE International Symposium on Hardware Oriented Security and Trust (HOST), 2017.