Secure Physical Enclosures from Covers with Tamper-Resistance

Vincent Immler[1], Johannes Obermaier[1], Kuan Kuan Ng[2], Fei Xiang Ke[2], Jin Yu Lee[2], Yak Peng Lim[2], Wei Koon Oh[2], Keng Hoong Wee[2] and Georg Sigl[1,3]

[1] Fraunhofer Institute for Applied and Integrated Security (AISEC), Germany
forename.surname@aisec.fraunhofer.de
[2] DSO National Laboratories, Singapore
[3] Technical University Munich (TUM), Germany

Abstract. Ensuring physical security of multiple-chip embedded systems on a PCB is challenging, since the attacker can control the device in a hostile environment. To detect physical intruders as part of a layered approach to security, it is common to create a physical security boundary that is difficult to penetrate or remove, e.g., enclosures created from tamper-respondent envelopes or covers. Their physical integrity is usually checked by active sensing, i.e., a *battery*-backed circuit continuously monitors the enclosure. However, adoption is often hampered by the disadvantages of a battery and due to specialized equipment which is required to create the enclosure. In contrast, we present a *batteryless* tamper-resistant cover made from standard flexPCB technology, i.e., a commercially widespread, scalable, and proven technology. The cover comprises a fine mesh of electrodes and an evaluation unit underneath the cover checks their integrity by detecting short and open circuits. Additionally, it measures the capacitances between the electrodes of the mesh. Once its preliminary integrity is confirmed, a cryptographic key is derived from the capacitive measurements representing a PUF, to decrypt and authenticate sensitive data of the enclosed system. We demonstrate the feasibility of our concept, provide details on the layout, electrical properties of the cover, and explain the underlying security architecture. Practical results including statistics over a set of 115 flexPCB covers, physical attacks, and environmental testing support our design rationale. Hence, our work opens up a new direction of counteracting physical tampering without the need of batteries, while aiming at a physical security level comparable to FIPS 140-2 level 3.

Keywords: Tamper-resistance, Physical Unclonable Function (PUF), Secure Bootstrap, Security Standards, FIPS 140-2, Higher-Order Alphabet PUF (HOA PUF).

1 Introduction

Standards for security certification such as FIPS 140-2 [Nat02], PCI-HSM [Pay12], or certain protection profiles of Common Criteria (CC) [KLR08] demand a physical security boundary for higher certification levels to protect compliant devices against tampering. These boundaries aim at separating the secure and insecure domains of a system, thereby protecting the device against physical attacks, such as drilling, grinding, etching, or probing [Wei00]. They can be made from security covers, housings, envelopes, etc. and are typically required to secure multiple-chip embedded systems on Printed Circuit Boards (PCBs) [IMJFC13, ES05]. In contrast to single-chip devices such as smartcards that are protected in silicon, these countermeasures are of particular importance for high-performance cryptographic modules and other applications requiring anti-tamper mechanisms on a PCB

IACR Transactions on Cryptographic Hardware and Embedded Systems ISSN 2569-2925,
Vol. 2019, No. 1, pp. 51–96
DOI:10.13154/tches.v2019.i1.51-96

level. This type of generic countermeasure is designated to make a wide range of invasive, semi-invasive, and non-invasive attacks [Sko05] more difficult to perform, as they typically require direct hardware access which is hindered by the physical security boundary [OI18].

One of the formerly widespread but now discontinued approaches for FIPS 140-2 level 3 and level 4 protection is based on a cover [W.L07b] or envelope [IMJFC13, W.L07a] with a mesh that encloses the Module Under Protection (MUP). Attempts to 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 Critical Security Parameters (CSPs) such as cryptographic keys, i.e., tamper-detection and response. However, a battery is required for this monitoring mechanism whenever the supplementing carrier system is powered off. Additionally, the CSPs are stored in a *volatile* Battery-Backed Random-Access Memory (BBRAM) to enable instantaneous zeroization upon detection of a physical intruder. Unfortunately, the specifics of such countermeasures are typically not made public.

Moreover, since either cover or envelope are manufactured using specialized technology, there is a risk of single-source supplier problems, i.e., there is no free market in addition to trust issues. Furthermore, this approach has significant practical drawbacks: adding a battery to the system increases bulk and weight, it negatively impacts the device's operating temperature range, and prohibits prolonged storage. Once the battery is fully discharged, the CSPs are lost and physical integrity can no longer be guaranteed [IBM12]. Storing CSPs in a BBRAM also leaves room for the zeroization circuit to fail. However, storing a key in non-volatile memory is also not an option, as its contents can be extracted while the system is powered off [SSAQ02]. Alternatively, Physical Unclonable Functions (PUFs) can be used [GCDD02, HYKD14]. Once the device is running, this security primitive derives a cryptographic key from the device's inherent manufacturing variations. As long as the device is powered off, extracting these parameters is supposedly difficult.

Since most PUFs are implemented in an Integrated Circuit (IC), it is impossible to use them for aftermarket protection of *other* Commercial-Off-The-Shelf (COTS) components. Furthermore, silicon-based PUFs *typically* do not provide the property of tamper-evidence [MV10], i.e., once powered on, they cannot verify if an attack was carried out on *other* parts of the *system* while powered off. In general, without additional meshes, these silicon-based PUFs are incapable of detecting online attacks that extract values during runtime [HNT$^+$13], e.g., from the data bus of a System-on-Chip (SoC) by using probing needles [Sko17]. One of the exceptions detecting certain FPGA-level optical probing attacks during runtime is [TFL$^+$17]. Nevertheless, direct access to the PCB and its components would still be possible.

To overcome these issues, we present a batteryless tamper-resistant cover that encloses multiple chips on a PCB. This cover verifies its integrity after power-up similar to a tamper-evident PUF and continues providing protection for its enclosed components during runtime. It therefore exceeds the scope of silicon-based PUFs such as the SRAM-PUF that by design is neither tamper-evident nor able to protect other on-chip system components, e.g., a data bus, or off-chip system components such as power regulators and external memory chips from proximity-based physical attacks. These proximity-based attacks are for example Laser Fault Injection (LFI) or Localized Electromagnetic Analysis (EMA). As long as the enclosed system has not been tampered with, the correct cryptographic key is derived from the cover, the system's sensitive data is decrypted and authenticated, and multiple Tamper Detection (TD) mechanisms start to ensure continuous protection while the system is running. This advances previous concepts and is a step towards meeting security standards, such as FIPS 140-2 level 3 or higher, *without* a battery for the security mechanism of the enclosure which makes it possible to employ it in more applications.

To achieve this, our cover contains an advanced mesh concept to not only detect short and open circuits, but also to measure the capacitances between traces. This is the basis to implement the tamper-evident PUF and allows for a dual-approach with integrity checks *and* PUF-based secret key derivation. Hence, recovery of the key is only possible from inside the system as long as the cover has not been tampered with.

The challenge in successfully implementing this lies in the constraint of only using an easily accessible manufacturing technology and enclosing the PCB in a large-scale physical object while only using small-scale intrinsic variations for the PUF-based key derivation. This is necessary to make their extraction by an attacker improbable. Ideally, such an implementation provides the property of being read-proof, i.e., any attempt made from the outside to extract the physical parameters of the cover causes their immediate destruction [GLM$^+$04]. Furthermore, a wider range of physical attacks must be taken into account that have been outside the scope of battery-backed approaches as their security mechanism is never powered off. In the following, we present several conceptual and practical considerations of our design, its various components, and demonstrate its feasibility.

Contributions. We present the following conceptual and experimental contributions based on the exemplary implementation of our concepts using commercially available flexPCB technology for the enclosure:

- A security architecture that only relies on widespread commercial technology to implement a hybrid cover which combines the properties of traditional tamper-responding enclosures with PUFs and extends their concept

- A stochastic model of the contained PUF to estimate its entropy, support its design process, and emphasize the benefit of constructing a Higher-Order Alphabet (HOA) PUF from it, i.e., drawing *symbols* from its output instead of binary data

- Reconsidering the well-known PUF metric Uniqueness for higher-order alphabets, i.e., its interpretation needs to reflect the specifics of a symbol-based PUF output

- An advanced measurement circuit that combines differential *and* absolute capacitance measurements as complementary sensing mechanisms

- A proof of concept hardware implementation, including a detailed statistical evaluation of 115 flexPCB covers, exemplary physical attacks, and environmental tests

These contributions significantly extend [IOK$^+$18] and evaluate the effectiveness of such a system when using only commercially available manufacturing technology based on a cover which is in contrast to [IOK$^+$18] where a customized manufacturing process is used to create an envelope. Moreover, we provide a more thorough empirical study both in terms of statistics as well as performed attacks.

Outline. An overview on related work is provided in Section 2. Afterwards, we introduce our security architecture in Section 3. The cover as a crucial building block is discussed in Section 4 including its stochastic model. This is followed by the explanations of the evaluation unit, whereas its measurement chain is described in Section 5 and the subsequent PUF data processing in Section 6. The properties of the host system and the secure bootstrap process are addressed in Section 7. In Section 8, a case study of the proposed concept is carried out to investigate its feasibility. Finally, a conclusion is drawn in Section 9 based on the obtained results.

2 State of the Art

This section presents background information on two well-established security domains, namely battery-backed tamper-resistant enclosures and PUFs which are combined later on.

2.1 Tamper-Resistant Enclosures

A first overview of tamper-resistant enclosures can be drawn from the list of FIPS 140-2 validated modules provided by the National Institute of Standards and Technology (NIST). From this list, we selected only the modules with highest physical security levels, namely: the HP Atalla Cryptographic Subsystem (ACS) [Hew10, Hew09], the IBM 4765 PCIe crypto coprocessor [IBM12], and Ultra Electronics AEP's HSM Crypto Module [Ult17]. In short, all require a battery-backed monitoring mechanism. In the following, we focus on IBM's and HP's cryptographic module, since we had them at hand and could align our practical findings as detailed in [OI18] with information we found online.

The physical security offered by IBM's module is also described in part by Isaacs et al. [IMJFC13] and Abraham et al. [ADDS91]. It is based on a battery-backed envelope manufactured by GORE that encloses the system and thereby protects it from tampering. This envelope can be found in many other systems, too. It is made of a flexible polymer with a printed conductive mesh and a maximum distance of 300 µm between traces on its most sensitive layer (cf. Appendix A). Additionally, it is potted using an opaque resin with the following properties: difficult to penetrate or remove, either mechanically or using solvents. The mesh serves as a resistive sensor which is continuously evaluated. Its tracks are routed on four layers in a serpentine pattern with no visible gaps (cf. Figure 2a) such that penetrating the mesh is improbable without causing a detectable change which triggers the zeroization. Due to the envelope's wrapping technique, this mesh obstructs any possible angle of an attack. In addition to that, visibility of its tracks is limited upon optical inspection (e.g., X-rays) and a device-specific layout randomization further increases the difficulty of attacks. Additional sensors, e.g., to detect intruding light on the inside of the device, complement the security provided by the envelope.

This mechanism ensures security assuming that sensitive data is kept only in volatile memory and a continuous power supply is available for the BBRAM and monitoring circuit, even when the device itself is powered off. Therefore, the monitoring circuit must be armed at the factory and supplied by the battery throughout the product lifetime including its shipping. This is unfavorable since environmental conditions during transport often exceed those of the intended operating environment in terms of peak temperature, vibration, etc. The *actively running* battery-backed monitoring circuit is subject to these conditions and as a result is more likely to cause false alarms. Typically, these devices are therefore delivered by priority shipping in thermally insulated boxes with gel packs. After arrival, maintaining the battery creates an additional burden and costs [IBM12].

In contrast to the highly tailored material properties employed in IBM's solution is the approach in HP Atalla's module by far less complex [OI18]. It is based on two covers that are securely tightened together such that prying them open relieves the pressure onto the PCB which in turn opens several connections. This mechanism, in addition to a coarse-grained mesh in the covers with 1 mm track width and space is constantly monitored by a battery-backed circuit which is supplied by two batteries on the carrier board and *eight* on an expansion card. Unlike the GORE envelope, no layout randomization could be used since the mesh forms just one loop with a single input and output. It should be noted that this module is compliant to FIPS 140-2 level 3 [Hew10] and level 4 [Hew09] with the only apparent difference being that the latter provides Environmental Failure Protection (EFP), i.e., exceeding a certain temperature range also causes zeroization. Please note that the aforementioned approaches are all rather limited in their operating temperature range in comparison to IC-level ratings. For example, as operating temperature, IBM's 4765 PCIe crypto coprocessor is specified for +10 °C to +35 °C *only* [IBM12].

Another approach in this domain actively measures the difference in fringe-effect capacitances of the enclosure due to intruding objects [ES05] and claims compliance to FIPS 140-2 level 4. It also relies on a battery-backed mechanism and therefore suffers from similar limitations. The same applies to [BOU07] regardless of their different housing technology which is based on ceramic or plastic covers. Similar covers are often found in Point-of-Sales terminals. We conclude that while the technology specifics of these approaches vary, they all share the same underlying operating principle of using continuously powered sensing mechanisms and corresponding response mechanisms.

2.2 Physical Unclonable Functions (PUFs)

The aforementioned disadvantages could be solved if the device did not require a battery. This can partially be achieved by PUFs which offer a hardware-intrinsic key storage without the use of a dedicated memory for the key [GCDD02, HYKD14]. They make use of random variations of manufactured structures to derive an individual behavior for each device. In order to harness the PUF properties, these variations must be extracted. They are similar for each read-out of the same device but subject to noise and affected by environmental changes.

Secure key derivation with PUFs is a common use case [DGSV15]. During enrollment at the factory, the key is derived for the first time and discarded after helper data is created and stored, to enable later error-correction. During reconstruction in the field, helper data and the noisy PUF response are combined to derive the initial secret. PUFs can be integrated into IC designs, but they only offer limited tamper-resistance [HNT+13], especially if they are just a *component* in a larger system, e.g., a System-on-Chip (SoC). Therefore, we focus on non-silicon, *system-level* PUFs in the following, i.e., the PUFs are required to enclose a significant portion of the system, thereby obstructing physical access.

One such example is the Coating PUF [TSS+06a] that protects a whole IC by covering its top with a randomized coating material, which is measured to extract its unique properties and to derive a secret key. Reconstructing this key is infeasible if the coating has been damaged due to an attack. A similar approach using an optical "backscatter" PUF is presented in [EFK+12]. Both approaches do not address attacks during runtime. Furthermore, covering *every* IC of an embedded device with a coating requires a costly, fully customized sourcing of its components. Moreover, direct access to the PCB would still be possible and therefore simplify various attacks, e.g., voltage glitch or power side-channel attacks.

Based on the requirement to protect a system as a whole, Vai et al. present an optical waveguide coating PUF [SFIC14] with a corresponding system architecture in [VNK+15]. As the waveguide only covers the top of a PCB, its edges and bottom remain unprotected. In addition, generic shortcomings of backscatter based systems such as inhomogeneity of the "illumination" and relative shift of the PUF token to the sensor, e.g., due to vibration, were not addressed. Yet another aspect is protection of such a system during runtime which is vital to protect its keys that are temporarily stored in volatile memory. Therefore, implementing a runtime tamper detection that monitors the system after power-on is essential to detect possible tampering attempts which is not mentioned in [VNK+15].

3 System Architecture

This section introduces the components of our architecture as shown in Figure 1a and explains the basic PUF data processing steps as outlined in Figure 3. To protect a host system, e.g., a Hardware Security Module (HSM), two building blocks are required: a cover with a capacitive sensoric mesh enclosing the system and its corresponding evaluation unit which performs the PUF data processing. This architecture has been practically implemented in Figure 1b, based on the packaging concept shown in Figure 4. Please note, in this publication we focus on the top cover only to explain the fundamental concepts and properties. Despite its more challenging physical shape, its concepts are largely the same when compared to the bottom cover. Hence, when there is no need to specifically refer to either top or bottom cover, we mostly use the singular form of cover for simplicity reasons.

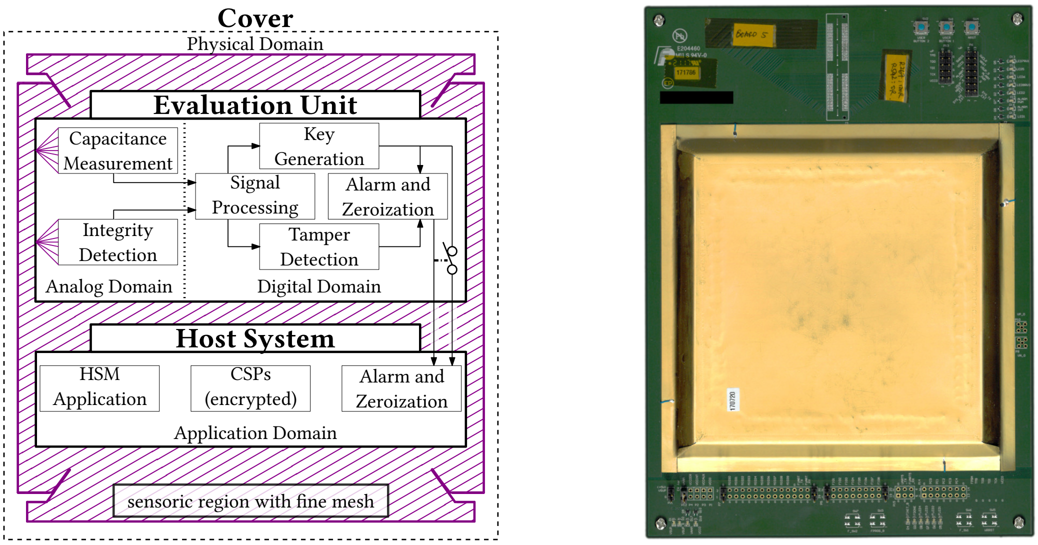

(a) System architecture with evaluation unit and enclosed host system.

(b) Test vehicle implementation with cover of size 140 mm × 140 mm.

Figure 1: System architecture and test vehicle with preliminary assembly (prior to potting).

3.1 Attacker Model

Our goal is *not* to present a solution for absolute security but new concepts that can be realized with moderate effort to achieve a *reasonable* level of security for multiple-chip modules and to support the design of even better follow-up solutions to ultimately pass certification without battery-backed mechanisms. Due to that and also to limit the complexity of this paper, we focus *only* on attempts to physically penetrate the cover, i.e., its mesh. More specifically, we assume penetrations to be at least 300 µm in diameter. Other attacks such as removing the cover are deemed impractical and result in severe damage to the cover, as explained when introducing the packaging concept in Section 4. As result of an attack above the given diameter, the system needs to be able to ensure that it becomes immediately inoperable and recovery of its sensitive data must be infeasible, e.g., once an attacker is detected, fuses are blown or similar steps taken to irrecoverably destroy the device. In the following, we briefly justify our reasoning for this diameter.

Standards for Security Certification. The Derived Test Requirement (DTR) A1 of [Pay13][1] demands a "Minimum width/separation (of active traces) of 6 mil" for an enclosure's mesh which translates to 300 µm based on geometrical considerations as illustrated in Figure 2a, i.e., if the track width is w then the detectable drill diameter is $2 \cdot w$. The same principles must be adhered to for other layouts, as shown in Figure 2b. Here, the detectable drill width is $3 \cdot w$, assuming equal width of tracks and space.

Please note that within the context of security certifications, just making a hole is not considered an attack [Joi15]. Instead, holes *and* subsequent attacks leading to *successful* exploitation of a system are rated on the scorecards. It is crucial that the determined attack potential (in points) is above a certain threshold to pass certification, i.e., there will always

[1] This document is officially available only under Non-Disclosure Agreement (NDA), nevertheless it can be found, e.g., on Baidu. It must not be confused with the *public* document "PCI PTS POI SR v4" on the Security Requirements (SR) of PCI PTS POI which does not include such detailed information.

be some attack possible, the only question is how much effort needs to be spent. Hence, we consider attempts of only making a hole as an evaluation-level analysis only without real-world significance. For practical exploitation, we assume that the underlying system has been designed such that either multiple smaller holes of 300 µm would be required or an increased drill diameter of $3 - 6$ mm for a single hole, e.g., to allow decapsulation of an IC which appears impractical through a 300 µm hole of several millimeters depth.

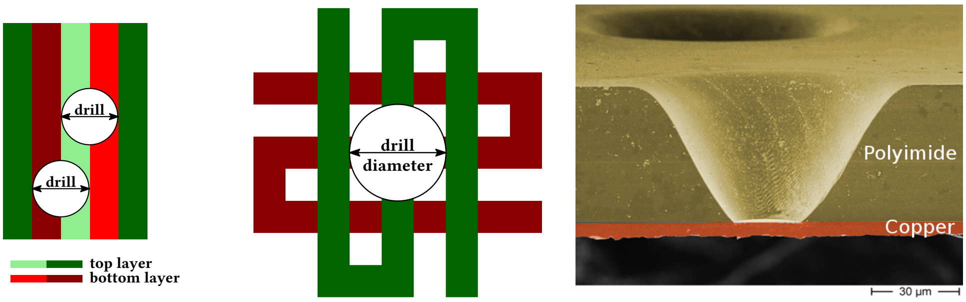

(a) Tracks with gapless design across layers.

(b) Layout variant with visible gaps (cf. Figure 6b).

(c) Laser drilling creating a cone-shaped hole in 50 µm polyimide.

Figure 2: Geometrical considerations of track width vs. mechanical and laser drilling.

Commercial Products. Another approach to limit the relevant diameter is to look at previous products and commercial brochures. According to our findings regarding the GORE envelope (cf. [OI18] and Appendix A), the smallest diameter to allow guaranteed detection should be around 300 µm as its track width and spacing is in the same range. Another security housing that was previously available [BOU07] was advertised to detect drills of 500 µm. Other solutions such as the one employed in HP Atalla's HSM have an even larger track width and spacing of 1 mm. Hence, to the best of the authors' knowledge, there are no commercial products offering a smaller mesh structure other than what is required to meet the 300 µm drill diameter. We point out that detecting smaller diameters is possible with advances in manufacturing technology to create smaller structures and that the concepts presented as part of this work scale accordingly.

Available Tools. Mostly the diameter of mechanical drills and *shaft*[2] diameter of micro-probing needles matters. While a micro-needle's tip is usually very small (~ 1 µm), it is also very short and not suited to reach far inside an enclosure. In contrast, the shaft is often several millimeters long but also much thicker, e.g., [GGB04b] offers tungsten needles with a copper shaft of 500 µm in diameter, i.e., a shaft already larger than the considered hole diameter. This shaft diameter does not account for a small gap around it, i.e., the hole itself would need to be slightly larger than 500 µm since a perfect alignment and insertion angle of 90° are difficult to achieve in practice.

Mechanical drills are easily available down to 100 µm as later illustrated in Figure 16c. However, as a rule of thumb [WTM08], a micro-drill's diameter versus its effective drill length – determined by its flute length – is a ratio of 1:15, e.g., a drill with 0.3 mm in diameter has an effective drill length of 4.5 mm at best. Therefore, such drills must be considered as part of the later security analysis in Section 8.4.

In contrast, we consider laser ablation or laser drilling not as a viable option as it tends to create cone shaped holes (depending on type of laser and material considered), i.e., the top hole will be typically larger than the bottom hole as illustrated in Figure 2c for a layer of 50 µm polyimide[3]. Considering the aspect ratio of hole diameter to material

[2] In some data sheets this part of a micro-probe is called shank instead of shaft.

[3] This figure is unrelated to the presented work and only intended to demonstrate technology aspects of laser drilling in polyimide as part of a regular manufacturing process (with courtesy of Fraunhofer EMFT)

thickness, in addition to the aspect ratio of top to bottom hole, it appears impractical to use laser drilling for the layer stack-up as presented in Table 1, i.e., even when assuming an idealized laser hole diameter to material thickness ratio of 1:1 (an assumption that is to our disadvantage) this still would create a hole of at least 243 µm. Moreover, the attacker would still need to penetrate further to ultimately gain access to the enclosed IC that performs the PUF key generation and additional checks for reasons of tamper-detection. Regarding chemical solvents for creating holes, we cannot make an educated statement as it would exceed our own expertise.

Conclusions on Attacker Model. Following these arguments, we are of the opinion that a 300 µm hole diameter is a reasonable choice for most practical applications and in accordance to current industry standards. As long as the enclosed system follows best design practices in this domain, such as routing all signal layers on the inner layers of a PCB, only using Ball Grid Array (BGA) components, buried vias, etc. it is difficult to foresee a successful exploitation with only few points on the score card of a security certification process, if it is possible at all when not deactivating some countermeasures for the evaluation process. Please note that such an enclosure is only one layer of a thorough Defense-in-Depth (DiP) concept. Therefore, we still require countermeasures at the appropriate level to counteract follow-up attacks such as LFI or EMA. Defeating these countermeasures would then in turn require more rework, requiring a larger degree of freedom to access the targeted IC which however is hindered by the enclosure, making the overall attack more complex to perform. Hence, the designated purpose of a physical security boundary of making such attacks more difficult would indeed be fulfilled. Other types of attacks are later briefly discussed in Section 8.4.3.

3.2 System Overview

The system is enclosed by a cover on the top and bottom of its PCB. Each cover contains capacitive sensors that act as a PUF and provide the basis for deriving a cryptographic key. After manufacturing the device, the key is derived for the first time to encrypt and authenticate sensitive data. The thusly protected data is stored in non-volatile memory, since an attacker can neither gain information from it nor change it in a useful way without damaging the cover, thereby destroying its key. During a later device start-up in the field, the system attempts to reconstruct the key and subsequently uses it to self-authenticate and decrypt sensitive data. Once the device is running, the same sensors that extracted the PUF properties of the cover now continuously monitor it. In case of an attack during runtime, an alarm is raised to trigger the zeroization of sensitive data which is temporarily stored as plaintext in volatile memory for processing it.

Cover. A flexPCB containing a mesh of fine conductive tracks is used to create the cover. The mesh represents a PUF to derive a cryptographic key by evaluating the capacitance measurements over its entire sensoric region. The cover's tracks are overlapping and represent electrodes which work as capacitive sensors. These tracks are subject to minuscule manufacturing variations in terms of surface roughness and physical dimension due to etching and related manufacturing processes [Bri04, WSL+15]. As a result, each overlap between electrodes represents a capacitance that cannot be accurately predetermined. Therefore, this concept relies on the intrinsic variation of a standard manufacturing process in contrast to artificially introduced randomness of, e.g., the Coating PUF [TSS+06a].

Evaluation Unit. This unit connects the cover to the host system. We refer to this as a separate unit primarily out of the reason for clarity of the explanations. In fact, it could be integrated into the host system which appears as the most secure but also least flexible approach in terms of development, as the process of incorporating it most likely results in changing the design of the host, too. Therefore, we implemented the evaluation unit in a dedicated microcontroller, controlling the PUF data processing concept, including:

- Analog domain: a single analog front-end that unifies distinct measurement concepts for the capacitance and integrity detection, i.e., they are sharing the same circuitry
- Digital domain: signal processing, PUF key generation, and runtime tamper detection logic including zeroization upon detection of physical intruders
- Data interface: to exchange information with the host, e.g., to serve as a decryption oracle, i.e., encrypted data is transferred to the evaluation unit and decrypted data is returned. Please note that this interface is within the physical security boundary.
- Heartbeat interface: with two independent alarm signals that are monitored by the host system during runtime to thwart “one-shot” intrusion attempts. This is for example, a Pulse-Width Modulated (PWM) signal with randomized frequency to which the host synchronizes and a static alarm signal which is active high.

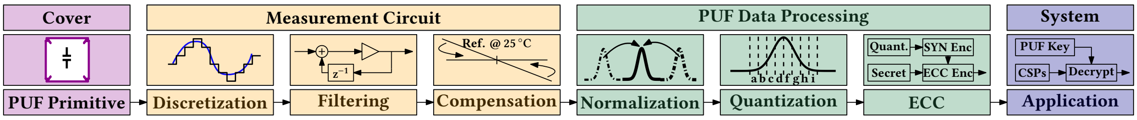

Figure 3: PUF data processing concept of the evaluation unit.

Host System. After each power-on, the host system synchronizes to the heartbeat signals and only then starts the interaction with the evaluation unit, e.g., to request the decryption of its firmware or additional CSPs using the key derived from the cover. Direct access to the key is denied to prevent software-based extraction. If the alarm signals indicate a tampering attempt, a zeroization is carried out. Following this generic approach, it is possible to implement a wide range of applications that may be unaware of their physically protected execution environment.

4 Physical Domain

In the following, we explain our packaging concept as shown in Figure 4. We assume that the PCB’s top is for active and its bottom for passive components. To protect them, we select a cover-based design, i.e., there is a top and bottom cover such that the majority of the surface which is exposed to an attacker is fully covered by the sensoric region contained in the covers. Both covers and their auxiliary mounting components such as the stiffener frame are attached to the PCB by at least two different mechanisms: firstly, by adhesives with high mechanical strength and good chemical resistance, secondly, by mechanical means such as screws. The covers themselves are additionally connected to the PCB using a secure seam which is beyond the scope of this publication and the simplified attacker model. Since the physical assembly of the covers is intertwined, removing them or prying them open without causing severe damage to one or the other is unlikely. To further harden the design and increase damage upon cover removal, we intend on using a conformal coating or potting resin for real-world designs which we omitted for this study.

As illustrated in Figure 4, there is sufficient space beneath the top cover to internally mount a heatsink to dissipate the heat. Moreover, the heatsink acts as an additional physical barrier once the attacker gets passed the cover itself. Since the distance between the top cover’s surface to the PCB is 7.4 mm, we assume that at least a drill diameter of 0.5 mm must be used for practical exploitation, i.e., a perfect attacker would know the best spot to attack, drill a hole to fully reach inside, decapsulate the area of the IC where the PUF data processing takes place, and extract its raw measurement data to reconstruct the PUF key. Such attacks must therefore be counteracted at the IC-level, too. However, in contrast to previous battery-backed solutions, it is no longer possible to only tamper with PCB-level tracks to defeat the security mechanism [OI18]. Instead, it is highly probable that the advanced evaluation logic at the IC-level must be attacked, too. This is a significant advantage of PUF-based enclosures over battery-backed approaches and their relatively crude but energy saving determination of the enclosure's physical integrity.

To complement the security provided by the covers, a vertical protection structure inside the PCB was designed to prevent attacks via its sides. Hence, any direct line of attack is obstructed either by the capacitive sensoric mesh or requires difficult angles to attack from which in turn are obstructed by the vertical protection structure. The packaging concept therefore already provides a comprehensive resistance towards attacks on a practical level.

Aside from the physical assembly which is designed to resist physical attacks, we still envision to use various other sensors, e.g., light, voltage, pressure contacts, and brittle components such as vias that easily get torn apart, to detect adversarial operating conditions upon power-on. An actual exploitation of the whole system therefore not only relies on defeating the tamper-resistant covers but also on successfully disabling additional layers of physical security on the inside which would require multiple holes to be made, thereby necessitating further damage to the covers and/or requiring a more advanced effort. Hence, an attacker will likely require more than one device to first design the best attack (identification) before attempting an actual attack (exploitation). This aspect is reflected on the scorecards during a certification process [Joi15].

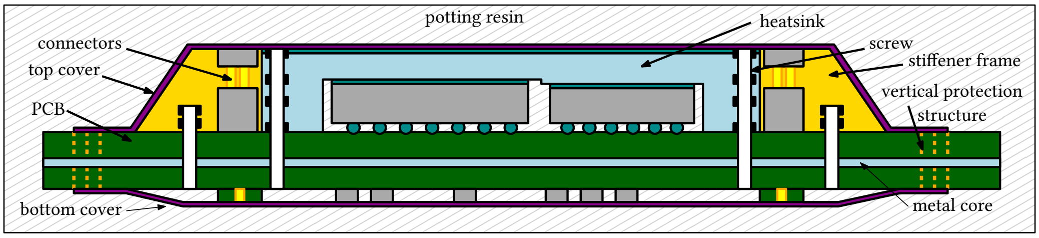

Figure 4: Packaging concept of a device enclosed by the proposed cover.

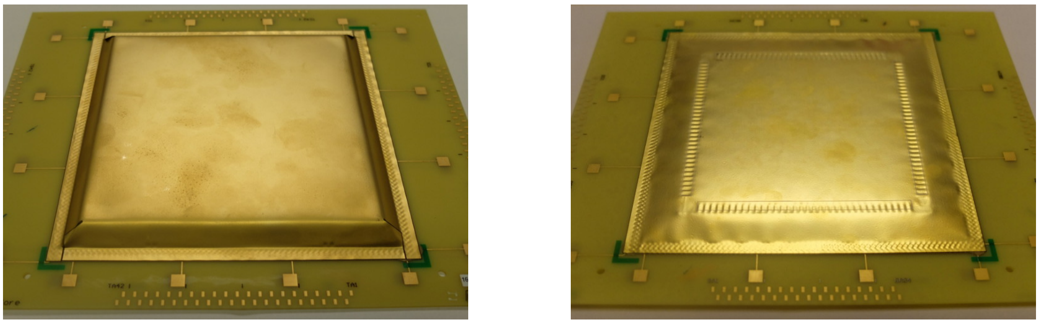

(a) Top cover.

(b) Bottom cover.

Figure 5: Assembly dummy at an early stage of the packaging concept.

4.1 Layer Stack-up of Cover

Designing a layer stack-up depends on the limitations of the manufacturing technology and the targeted sensor type. Thus far, tamper-respondent enclosures are primarily based on resistive sensors that are manufactured by a silk-screen printing process, i.e., fine tracks are printed on a flexible sheet and the resulting mesh is considered as resistors in the corresponding evaluation circuit. However, this has several disadvantages when compared to capacitive sensors, especially for devices that can be fully powered off, as the resistance of a track could be measured and replaced with a matched resistor which would result in a bypass difficult to detect. Moreover, resistive sensors only detect changes within their own tracks. In contrast, capacitive sensoric regions are conceptually less prone to bypassing of their tracks due to the small capacitances in the range of femtofarads. Furthermore, parasitic capacitances towards surrounding objects influence the measurement. Hence, not only are tracks considered part of the measurement but so are nearby layers and objects.

For the cover, we aim at a self-contained capacitive sensor to sense the intrinsic manufacturing variations of the mesh. This is achieved by implementing two layers of electrodes that are enclosed with a grounded shield to provide a defined boundary condition and prevent interference from the inside or outside. One layer of electrodes is named "Tx" while the other layer contains the corresponding "Rx" electrodes. As detailed in Section 5, the "Tx" electrodes are driven by an excitation signal and the "Rx" electrodes act as receivers. The capacitance between each Tx and Rx electrode is quantified as the "mutual capacitance", as noted in Table 1. Since the parasitic capacitance towards the shield is rather large compared to the mutual capacitance, partially removing or not grounding the shield already degrades the measurement up to the point that it no longer works. For connectivity to the measurement circuit, the cover requires an additional layer for connectors, resulting in a total of five conductive layers. For our implementation, we exemplarily use flexPCB technology which is a lithographic process and therefore allows a much smaller track width when compared to silk-screen printing.

Table 1: Layer stack-up of the flexPCB cover with overall thickness of 243 µm.

Layer	Height	Description	Comment
1	27 µm	Shield	Facing to environment
	52.5 µm	Bonding/Insulation	⇳ Parasitic capacitance C^{P}
2	24 µm	Tx electrodes	Driven electrodes
	12 µm	Polyimide substrate (carrier)	⇳ Mutual capacitance C^{M}
3	24 µm	Rx electrodes	Receiving electrodes
	52.5 µm	Bonding/Insulation	⇳ Parasitic capacitance C^{P}
4	12 µm	Shield	Facing inside (to PCB)
	12 µm	Polyimide substrate	
5	27 µm	Connectors and routing	

4.2 Sensor Design (Physical Layout)

The physical layout of the mesh follows the logical representation in Figure 6a, i.e., a matrix of overlapping electrodes to fill the whole cover with the intended sensor structure, thereby avoiding blind spots where attacks would go undetected. This has the advantage that if the cover is damaged in one spot, more than one capacitive sensor is destroyed, i.e., such an interconnected sensor arrangement makes physical attacks more easily detectable. To obtain a well-defined protection against a specific drill diameter, the sensor mesh for our study is manufactured with a structure size of 100 µm line and space as shown in Figure 6b. Consequently, any combination of "track-space-track" or vice-versa is guaranteed to be ≤ 300 µm and therefore provides protection against drilling attacks of such diameter as noted in Figure 2. Hence, countering smaller drills can be done with smaller structures.

Creating small structures increases the difficulty of attacks and improves manufacturing variations. Smaller tracks than 100 µm could already be created by flexPCB technology, as some manufacturers offer a track width down to 25 µm in aggressive lines. However, since the structure size is small and spans a significant area, manufacturing defects may occur that either result in short or open circuits, e.g., in case of over-etching. This reduces the yield and increases cost why we chose 100 µm as a start to exemplarily test the envisioned concept. At the time of device assembly, it is critical to verify that each cover is free of the aforementioned defects. This is considered a mesh with "full integrity" which provides assurance that the whole sensoric surface contributes to the PUF during enrollment.

To detect open circuits, the layout in Figure 6b allows checking the electrode's continuity by forming a loop, i.e., both *input* and *output* of an electrode are routed to the connector, denoted as Ri/Ro for Rx and Ti/To for Tx electrodes. To also check for short circuits, the electrodes are interleaved such that each neighboring track can be driven independently. Figure 6b shows the resulting advanced layout and its simplified equivalent circuit. It can easily be scaled to cover a larger area by increasing the number of sensor cells and/or the number of electrodes. For our test design, we selected a $(N_{\mathrm{Tx}},N_{\mathrm{Rx}}) = (16{,}16)$ configuration of electrodes, resulting in $N_{\mathrm{nodes}} = 16 \cdot 16 = 256$ sensor *nodes* with mutual capacitance $C^{\mathrm{M}} = C_{\mathrm{s}}$, whereas each node is created from approx. 1800 sensor *cells*, i.e., tiny squares with capacitance C_{c} that are the result of each Tx/Rx electrode track overlap.

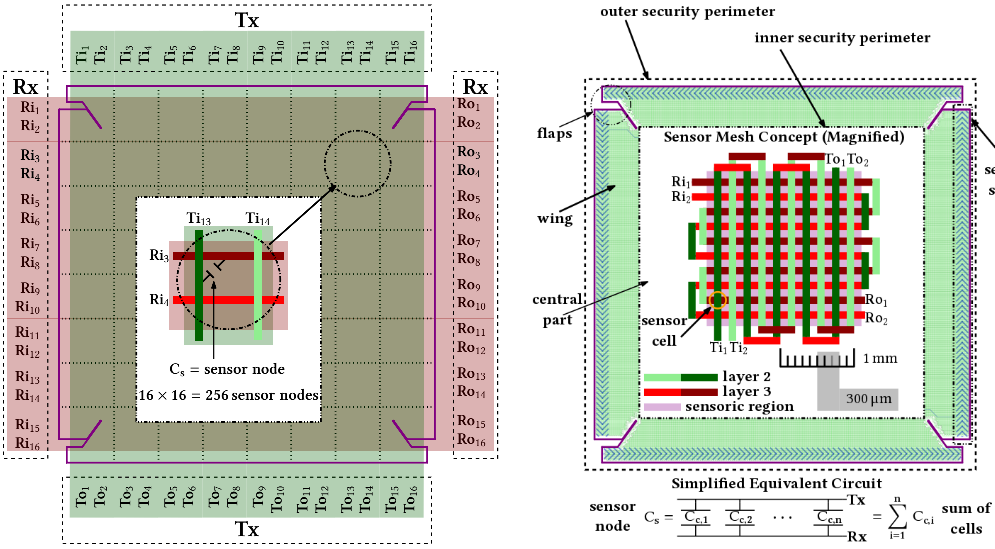
Tx
Ti1 Ti2 Ti3 Ti4 Ti5 Ti6 Ti7 Ti8 Ti9 Ti10 Ti11 Ti12 Ti13 Ti14 Ti15 Ti16
Rx
Ri1 Ri2 Ri3 Ri4 Ri5 Ri6 Ri7 Ri8 Ri9 Ri10 Ri11 Ri12 Ri13 Ri14 Ri15 Ri16
Rx
Ro1 Ro2 Ro3 Ro4 Ro5 Ro6 Ro7 Ro8 Ro9 Ro10 Ro11 Ro12 Ro13 Ro14 Ro15 Ro16
Ti13 Ti14
Ri3
Ri4
Cs = sensor node
16 × 16 = 256 sensor nodes
To1 To2 To3 To4 To5 To6 To7 To8 To9 To10 To11 To12 To13 To14 To15 To16
Tx
outer security perimeter
inner security perimeter
flaps
Sensor Mesh Concept (Magnified)
To1 To2
Ri1
Ri2
secure seam
wing
central part
sensor cell
Ro1
Ro2
Ti1 Ti2
1 mm
layer 2
layer 3
sensoric region
300 µm
Simplified Equivalent Circuit
sensor node
Cs = Cc,1 Cc,2 ··· Cc,n = ∑ Cc,i (i=1 to n)
Tx
Rx
sum of cells

(a) Logical layout.

(b) Physical layout and equivalent circuit.

Figure 6: Different representations and details of the top cover.

4.3 Stochastic Model of a Sensor Node

Based on the simplified equivalent circuit shown in Figure 6b, we analyze the capacitance C_{s} of a single sensor node. When neglecting the track resistance, each of the n track overlaps (sensor cells) between the electrodes is modeled as a tiny capacitor in parallel. This is a valid estimate based on our practical experience and leads to C_{s} being the sum over the capacitances $C_{\mathrm{c},i}$ with $i = 1,\ldots,n$.

In the following, we assume $C_{\mathrm{c},i} \sim \mathcal{N}(\mu_{\mathrm{c}}, \sigma_{\mathrm{c}}^2)$ as i.i.d. Recall that adding two Gaussian random variables results in a Gaussian distribution with the sum of means and sum of variances. Therefore $C_{\mathrm{s}} \sim \mathcal{N}(n \cdot \mu_{\mathrm{c}}, n \cdot \sigma_{\mathrm{c}}^2)$, i.e., $\mu_{\mathrm{s}} = n \cdot \mu_{\mathrm{c}}$ and $\sigma_{\mathrm{s}}^2 = n \cdot \sigma_{\mathrm{c}}^2$. According to the weak law of large numbers we then compute the respective means of the sensor cell

$$\overline{C_{\mathrm{c}}} = \frac{C_{\mathrm{s}}}{n}\ ,\ \overline{\mu_{\mathrm{c}}} = \frac{\mu_{\mathrm{s}}}{n}\ ,\ \overline{\sigma_{\mathrm{c}}^2} = \frac{\sigma_{\mathrm{s}}^2}{n} \tag{1}$$

and obtain equations that depend on n which is the number of parallel cells combined to a sensor node, i.e., $C_{\mathrm{s}} = n \cdot \overline{C_{\mathrm{c}}}$ which can be used as a guideline for subsequent designs to determine the number of electrodes, their number of overlaps, etc.

As a next step, the entropy of the thus far continuous Probability Distribution Function (PDF) of C_{s} needs to be estimated. To observe C_{s}, a measurement circuit is required. Hence, observing $\mathrm{PDF}(C_{\mathrm{s}})$ depends on the circuit's resolution Δ_{M}. As security objective, we target $\Delta_{\mathrm{M}} \leq \overline{C_{\mathrm{c}}}$, i.e., removing a single cell from the capacitance C_{s} of a sensor node would be detected with high probability. Based on the results of Section 8, we select $\Delta_{\mathrm{M}} = 1\,\mathrm{fF} \leq \overline{C_{\mathrm{c}}}$. Subsequent processing includes a quantization with bin size Δ_{Q} [IHKS16] as explained in Section 6.2. However, at this stage we are interested in the fundamental properties of the design only which is why we proceed with Δ_{M} instead. According to [CT06], the Shannon entropy H^{Δ} of a discretized Gaussian random variable is given by

$$\mathrm{H}^{\Delta} = \mathrm{ld}\left(\frac{\sigma_{\mathrm{s}}}{\Delta_{\mathrm{M}}} \cdot \sqrt{2\pi\mathrm{e}}\right) \tag{2}$$

To achieve, e.g., $\mathrm{H}^{\Delta} = 5\,\mathrm{bit}$ for the given Δ_{M}, we solve for σ_{s} which is 7.7 fF. This value can be verified empirically once a statistically relevant number of samples is available. In our case, σ_{s} closely matches the empirically determined deviation as presented in Section 8.2.

Other publications such as [WSL+15] and [Bri04] show that besides of local variation there is also global variation across manufacturing panels of PCBs. This would contradict our model and result in a capacitance gradient and global bias. To counteract this effect, we use a differential measurement as further detailed in Section 5. The basic idea is to interpret $C^{\mathrm{M}} = C^{\mathrm{N}} + C^{\mathrm{V}}$ with C^{N} being the nominal capacitance, i.e., $C^{\mathrm{N}} = \mu_{\mathrm{s}}$, and C^{V} as the variation in the range of, e.g., $[-3\sigma_{\mathrm{s}}; +3\sigma_{\mathrm{s}}]$. For an even h, the electrodes Tx_{h-1} and Tx_{h} are routed differentially. They form the Tx pair $(\mathrm{Tx}_{2k-1}, \mathrm{Tx}_{2k}),$ for $k \in \{1, 2, \ldots, N_{\mathrm{Tx}}/2\}$. All Rx are used as individual electrodes with $(\mathrm{Rx}_j),$ for $j \in \{1, 2, \ldots, N_{\mathrm{Rx}}\}$. Hence, the differential capacitance is $\gamma_{k,h} = C^{\mathrm{M}}_{(2k-1),j} - C^{\mathrm{M}}_{(2k),j} = C^{\mathrm{V}}_{(2k-1),j} - C^{\mathrm{V}}_{(2k),j}$.

Measuring $\gamma_{k,h}$ between two pairs of nodes in close vicinity isolates the local variation and minimizes global effects, as supported by our findings, e.g., in Figure 14l. We therefore argue that the independence of variables can indeed be assumed. Moreover, our results also confirm what would be expected from the central limit theorem, i.e., the sum of many independent cells combined to a node tends towards a normal distribution. For the entropy, let us also briefly consider the scenario of the differential measurement γ which reduces the number of sensor nodes to draw entropy from to $N_{\gamma,\mathrm{nodes}} = N_{\mathrm{nodes}}/2$. However, the resulting PDF of γ is $\mathcal{N}_{\gamma}(0, \sqrt{2} \cdot \sigma_{\mathrm{s}})$ and therefore Equation 2 can be rewritten as

$$\mathrm{H}^{\Delta}_{\gamma} = \mathrm{ld}\left(\sqrt{2}\right) + \mathrm{ld}\left(\frac{\sigma_{\mathrm{s}}}{\Delta_{\mathrm{M}}} \cdot \sqrt{2\pi\mathrm{e}}\right) \tag{3}$$

Hence, the theoretical entropy of the cover for this scenario is $16/2 \cdot 16 \cdot 5.5\,\mathrm{bit} = 704\,\mathrm{bit}$.

5 Measurement System

To address the previously outlined requirements, we leverage the solution presented in [OIHS18] to measure a differential capacitance and verify the electrode's integrity. Moreover, we extend the work presented in [OIHS18] and add another measurement mode to sense the absolute capacitance which completes the sensing capabilities of the system. This improved circuit extracts the differential capacitance variation C^{V} that is in the femtofarad range while effectively canceling out the orders of magnitude larger nominal capacitance C^{N} and ignoring the parasitic capacitances C^{P} as supported by our practical results in Section 8.

Additionally, the chosen approach handles a matrix of capacitors in parallel and is adjustable to various electrode configurations. Please note, the circuit concept was designed with respect to security, avoiding timing side-channels and emanations of, e.g., oscillator-based measurement circuits as reported in [MSSS11] that could put approaches such as [TSŠ+06b] at risk, as the unknown capacitance determines the oscillation frequency which can be observed from the outside. In the following, the concept of the circuit is explained while the implementation details are presented in Section 8.1 as part of our case study.

5.1 Differential Capacitance Measurement

As the targeted differential capacitance can be easily falsified by amplifier offsets and resistor mismatches, it is vital to prevent any such influence. In contrast to other methods that are prone to such influence, e.g., determining the capacitances and subsequently computing their difference in analog circuitry, we use a differential *excitation* to move the "computation" of the difference inside the cover as a kind of "in-situ measurement" which is briefly explained in the following. Moreover, our method shifts the measurement from time into the frequency domain which is then evaluated by digital signal processing and not susceptible to component imperfections, aging, etc. Hence, we employ analog circuitry only where absolutely necessary.

Figure 7 shows the basic concept. The digital domain is implemented on a microcontroller providing suitable analog peripherals. Its Digital-to-Analog Converter (DAC) generates two sine waves of equal amplitude at a frequency of 33.3 kHz with a relative phase of 180°. The following low-pass removes the DAC's sampling artifacts. The fully differential amplifier improves the signal's amplitude and relative phase match while being able to drive a larger capacitive load.

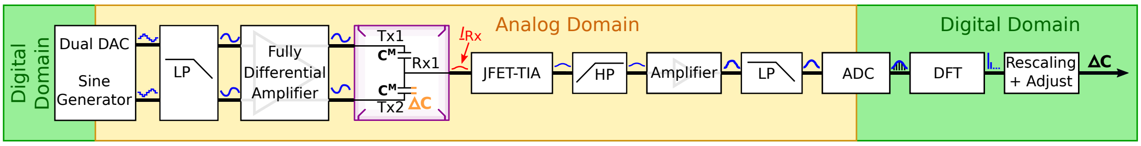

Figure 7: Measurement concept for the differential capacitance ΔC.

The resulting signal of amplitude V_{Tx} and angular frequency ω is connected to one Tx pair of the cover, e.g., Tx1 and Tx2, while the resulting complex current $\underline{I}_{\mathrm{Rx}}$ at *one* Rx electrode is observed. Both Tx electrodes are excited with an inversely phased signal of the same amplitude and frequency. This generates two cover-internal currents, proportional to the capacitances but inverted to each other. Therefore, C^{N} of the mutual capacitance C^{M} cancels itself out and only a current representing $\Delta C = C^{\mathrm{V}}$ remains:

$$\underline{I}_{\mathrm{Rx}} = j\omega V_{\mathrm{Tx}} \cdot C^{\mathrm{M}} + j\omega(-V_{\mathrm{Tx}}) \cdot (C^{\mathrm{M}} + \Delta C) = -j\omega V_{\mathrm{Tx}} \cdot \Delta C. \tag{4}$$

This current of less than a few nanoamperes is amplified and converted into a voltage by a JFET-based Transimpedance Amplifier (JFET-TIA) that provides a high gain at very low noise. Since the ΔC information is only contained in the small AC component of the signal, the DC component, caused by amplifier offsets, is completely removed by a high-pass filter. The subsequent second amplifier increases the signal's amplitude to match the Analog-to-Digital Converter's (ADC) dynamic range and comprises a low-pass filter to prevent aliasing during signal acquisition.

The microcontroller's ADC records eight periods of the signal which provides a reasonable trade-off between noise-reduction and measurement duration in our scenario. Our optimized digital signal processing chain computes a single-bin Discrete Fourier Transform (DFT) of the input signal by applying the Goertzel Algorithm [Goe58] which is possible due to the known excitation frequency. This algorithm requires very few resources and can be executed at a high performance using the system's Floating Point Unit (FPU). The complex result is converted to polar coordinates and evaluated for its magnitude, representing the absolute value $|\Delta C|$, while its phase provides $\texttt{sign}(\Delta C)$. Hence, $\gamma = |\Delta C| \cdot \texttt{sign}(\Delta C)$.

5.2 Absolute Capacitance Measurement

In addition to the measurement of a differential capacitance, we extend the system to support the acquisition of the absolute mutual capacitance C^{M}. As described in Section 4.3, it is assumed that the absolute capacitance is dominated by C^{N} which does not contain entropy. Thus, it should *not* be used directly as input to the PUF key generation. Nevertheless, C^{M} provides valuable information regarding the enclosure's physical integrity and should be considered in the overall process. Clearly, physically tampering with the electrodes while ensuring the same behavior of the differential *and* absolute capacitance measurement *at the same time* appears extremely challenging. Therefore, taking the absolute capacitance into consideration enables the detection of certain attacks that may exploit specific properties of the differential measurement, as explained in Section 8.4.

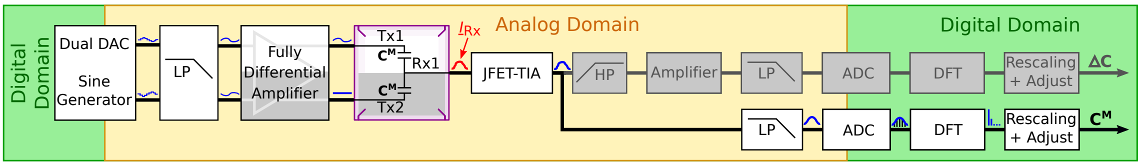

Figure 8: Measurement concept for the absolute capacitance measurement of C^{M}. The differential path in gray is disabled while reusing some of its component for the absolute capacitance measurement.

Implementing the absolute capacitance measurement is done at a nearly negligible area overhead as a major part of the circuit and digital signal processing of the differential measurement is reused. During an absolute capacitance measurement, shown in Figure 8, only a single differential amplifier output is enabled, its amplitude is reduced by G_{abs}, and the Rx signal is rerouted. For example, when measuring the mutual capacitance between Tx1 and Rx1, the Tx1 electrode is excited by a sine while the Tx2 electrode remains on a constant voltage level. This results in an Rx current directly proportional to the mutual capacitance, i.e.,

$$\underline{I}_{\mathrm{Rx}} = -j\omega \cdot G_{\mathrm{abs}} \cdot V_{\mathrm{Tx}} \cdot C^{\mathrm{M}}. \tag{5}$$

As the mutual capacitance is about three orders of magnitude larger than the variation this would overdrive the Rx circuitry. Therefore, the excitation sine is digitally attenuated by 20 dB, i.e., $G_{\mathrm{abs}} = 1/10$, and the 40 dB amplifier on the Rx side is skipped. The Rx signal is tapped early on within the processing chain and connected to a secondary channel of the ADC. Eventually, the same signal processing algorithms are applied in the digital domain. This yields the amplitude and phase of the complex signal which is then converted into C^{M}. While both outputs of the mutual capacitance measurement, i.e., amplitude and phase information, support the detection of attacks, we primarily focus on the magnitude for the remainder of the paper.

5.3 Integrity Verification of Sensor Mesh

The system additionally verifies the cover's integrity by reusing a majority of the capacitance measurement circuitry. Tx electrodes require little to no overhead, as the DAC can put a voltage on their Ti input and a comparator senses the voltage at all To outputs. Thereby, severed traces with no output at To, can be detected and shorted traces will show an output signal at *other* To outputs.

For Rx electrodes, the approach is different as applying a voltage signal to current-sensitive circuits overdrives them. Therefore, the opposite end of each Rx electrode is attached to a current source that injects a small current into a chosen Rx electrode. Subsequently, the JFET-TIA is used to sense the current and verify the electrodes' integrity. Therefore, this can be implemented at a negligible overhead while providing strong assurance of the electrodes' integrity. This exceeds previous PUF concepts as this allows to reliably distinguish between manufacturing defects and excessive variation. Another advantage is that the signal evaluation is done by a comparator as further detailed in Section 8.1. This results in a much faster integrity verification during runtime to detect rapid intrusion attempts during runtime.

6 PUF Data Processing

Several additional processing steps are required to yield a cryptographic key which is reliable, provides full entropy, and in addition to that offers the property of *tamper-sensitivity*, i.e., even small physical changes should result in a significant change of the PUF's output data.

6.1 Compensation and Normalization

The output of the previous stages is considered as raw differential capacitance data that must be adjusted to account for structural bias and environmental changes such as temperature drift. Removing structural bias is also called "normalization", as for example in [MVHV12]. Typically, this would require additional helper data to mitigate the effects of a structural bias. However, as seen later on in Figure 14h, the structural bias in our case is mostly in such a way that removing the mean of each Tx group also removes the structural bias, i.e., all Rx electrodes measured in parallel are subject to the same bias. Since a shift in these means is the predominant effect of temperature drift this serves as a simplified temperature compensating step, too. Hence, the values prior to the quantization are computed by the following equation

$$X_i = X_{k,h} = \gamma'_{k,h} - \left(\frac{1}{N_{\mathrm{Rx}}} \sum_{r=1}^{N_{\mathrm{Rx}}} \gamma'_{r,h}\right) \qquad h = 1, \ldots, N_{\mathrm{Tx}}/2 \quad \text{and} \quad k = 1, \ldots, N_{\mathrm{Rx}} \tag{6}$$

whereas $\gamma'_{k,h}$ is a representative of the previously obtained noisy differential capacitance. The output $X_{k,h}$ is created by subtracting each Tx group's mean. To simplify the notation, the result is reshaped to X_i with $i = 1, \ldots, k \cdot h$.

6.2 Quantization and Error-Correcting Code (ECC)

The previously compensated and normalized data is now further processed by an equidistant quantization [IHKS16]. This is an error-*reduction* technique to mitigate the remaining circuit noise σ_{N} that would otherwise cause frequent changes in the output data. Alternatives would have been, e.g., an equiprobable quantization as applied to the output of the Coating PUF [TSS+06c] which is typically based on a Gray code, as illustrated in

Figure 9b. However, the unequal width of equiprobable intervals causes helper data leakage in addition to an uneven tamper-sensitivity as explained in [IHKS16]. Other approaches to equiprobable quantization include [dGVL12, SAS17, BDHV07] using a partitioning scheme to avoid helper data leakage.

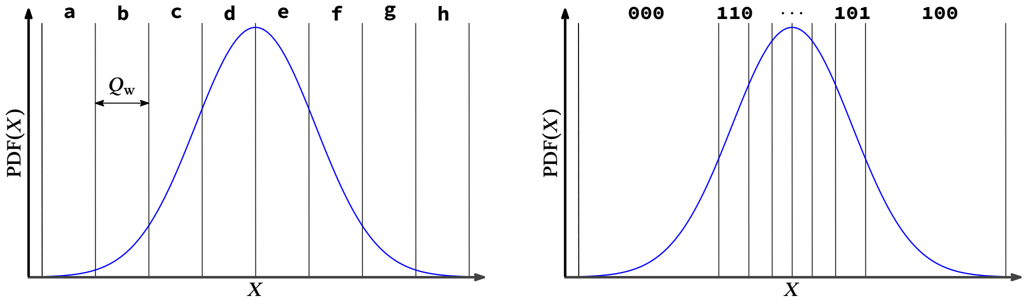

(a) Exemplary equidistant quantization.

(b) Exemplary equiprobable quantization.

Figure 9: Different quantization approaches with assignment of *symbols* for the equidistant quantization and a Gray code in case of equiprobable quantization.

However, two problems of equiprobable quantization remain. First of all, it is mandatory to precisely know the PDF in addition to its preferred symmetry. This is difficult for some practical scenarios, e.g., within the context of low volume manufacturing as it is typically the case for tamper-resistant enclosures. Secondly, the quantization error is mainly determined by the innermost intervals as illustrated in Figure 9b which either results in a relatively high error rate or in a diminished entropy when increasing the width of the two innermost intervals (assuming a relatively uniform noise level across the range of values).

In contrast, an equidistant quantization as illustrated in Figure 9a is relatively insensitive to, e.g., shifts of the PDF and also provides a constant quantization error probability across the range of values. It is therefore an attractive choice for practitioners at the downside of a biased PUF output *at the stage of quantization* which needs to be considered in subsequent processing steps. The equidistant quantization works as follows. The width Q_w of the quantization intervals is determined by $Q_\mathrm{w} = 2 \cdot y \cdot \sigma_N$ whereas y is a parameter of choice according to the required reliability. To obtain m-bit PUF responses, $\mathrm{PDF}(X)$ is divided into $L = 2^m$ intervals of the form $(\mu + l \cdot Q_\mathrm{w}, \mu + (l+1) \cdot Q_\mathrm{w}]$ where $l = -L/2, \ldots, -1, 0, 1, \ldots, L/2$. Aligning $l = 0$ and μ of the Gaussian distribution leads to the highest entropy output while it is slightly decreased by misalignment depending on the choice of y and the shift. However, due to symmetry reasons of the equidistant quantization this decrease is well-bounded and therefore a robust scheme.

Figure 9a exemplarily illustrates the quantization intervals for $L = 8$ and an optimal alignment. Each interval is represented by a symbol Q_l in $[0, L-1]$. As the measurement of the PUF values X'_i is non-ideal, i.e., affected by noise of the measurement process, values could move to a different interval compared to the time of enrollment. To additionally reduce such errors, the offsets between each value X_i and their corresponding interval center are stored as helper data W^*. By following this approach, the probability of a quantization error can be significantly reduced, e.g., by choosing $y = 3.29$ the symbol error-rate is at 0.1% for each node [IHKS16]. During PUF reconstruction, this value is then mapped to the PUF response R'_i, i.e., $(X'_i - W^*_i \in Q_{l_i} \rightarrow R'_i)$ for $i = 1, \ldots, N_{\gamma,\mathrm{nodes}}$.

To obtain a fully robust device, a subsequent error-correction scheme is still required. Since the underlying scenario of this work is similar to [IHL+17], we could make use of a variable-length bit mapping that represents the *symbols* from a *higher-order alphabet* by a variable-length bit string. This counteracts bias in the quantized PUF values and therefore mitigates large portions of the helper data leakage. Afterwards, the variable-length bit string is corrected by a specialized insertion/deletion error-correcting code.

However, due to difficulties in obtaining a constant-runtime for implementations of these codes, we are currently working on variants to continue operating on the symbols instead of applying the variable-length bit mapping. Hence, we consider this a Higher-Order Alphabet PUF as opposed to a Binary PUF. Please note, the combined approach of equidistant quantization and symbol-oriented post-processing is fundamentally different to other approaches in the domain of PUFs, as they typically only consider i.i.d. binary responses [JW99, DRS04, BGS$^+$08, YD10, MVHV12, HMSS12, HYS16]. Due to that, we also need to focus on the uniqueness and reliability in terms of symbols which has not been done beforehand, as further detailed in Section 8.3. This necessitates the extension of previous PUF works and requires a new perspective on how to interpret the PUF output. At the same time, this opens up the opportunity to assess the PUF data independently of the chosen bit mapping. Our own scheme does not fit the scope of this paper and will be presented at a later stage.

7 Application Domain

In the following, we briefly explain the secure boot process on a conceptual level. In addition to that, we describe our example application.

Boot process. The overall system's boot process is depicted in Figure 10. Immediately after power-up, two independent heartbeat signals are generated by the evaluation unit to which the host system synchronizes, in particular if the two units are two different ICs. This should prevent rapid "one-shot" attempts to directly interrupt the alarm later on. As a first line of defense, an integrity detection is carried out to verify if the electrodes contain any short or open circuits.

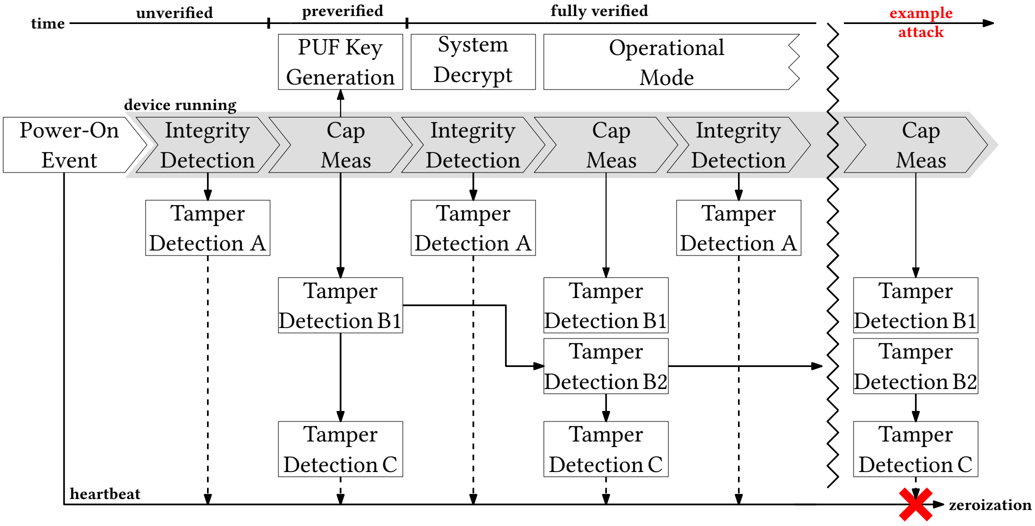

Figure 10: Secure boot process with intertwined security mechanisms, i.e., integrity detection and capacitance measurements. Attacks are detected shortly after power-on by either integrity detection or PUF key generation. Attacks during runtime are detected by the tamper detection (TD) mechanisms A, B1, B2 and C.

We name this Tamper Detection A (TD-A) which is then followed by a capacitive measurement. Both are continuously repeated during runtime, i.e., they take turns. The first differential capacitance measurement after power-up is considered a reference value and used for the PUF key reconstruction. Simultaneously, the *same* differential capacitance values are used to start another TD, termed TD-B1 and TD-B2. TD-B1 limits the valid range of each individual capacitance relative to its reference value, i.e., at $t = T_0$ boundaries for each sensor node are computed *once* based on the reference value $\pm p$, whereas p is a constant guard parameter. For each subsequent measurement, the then current capacitance value is checked against the computed boundaries: $|\gamma(t) - \gamma(T_0)| < p$. As additional precaution, TD-B2 limits the discrete rate of change, i.e., by computing $|(\gamma(t) - \gamma(t-1))| < q$, for a second security parameter q. Both parameters p and q must be tuned to the specific application profile of the device and are strongly related to the width Q_w of the equidistant quantization.

The output of the absolute capacitance measurement serves as input for TD-C. Here, zeroization is caused if any of the absolute capacitance values significantly deviates from the then-current mean of all absolute capacitance nodes. This approach is relatively insensitive to temperature drift in absolute capacitances as can be derived from Figure 26c. As later illustrated in Section 8.4, a deviation due to tampering can be assumed if the value is outside a ±15% range of the mean. Please note that tinkering with TD-A, TD-B1, TD-B2, and TD-C cannot be done easily without violating some of their properties.

By successfully generating the PUF key, the proper initialization of the TD-B mechanisms is ensured. Evaluating TD-C complements this approach. This PUF key could then be used to decrypt the firmware of the host or some of its CSPs. In our actual implementation, we combine the PUF key with IC-level roots-of-trust[4] to form a compound device identifier within the Device Identifier Composition Engine (DICE) framework for the secure boot process of the device. If either during power-up or runtime any of these checks fail, a tamper-event is caused that triggers the zeroization and stops the heartbeat signals. All mechanisms have been designed in an intertwined way to have a layered approach to security; individually disabling them is considered very challenging.

Firmware level. Following the ideas of [OHHS18], a custom firmware was developed for testing the concept. This is based on a security-enhanced fork of FreeRTOS that serves as operating system for the measurement setup and PUF data processing chain. Additionally, it implements an Embedded Key Management System (EKMS) which operates similar to the software of a Hardware Security Module (HSM). This system ensures real-time behavior of the measurement process while protecting and operating on sensitive data, i.e., PUF data and derived keys. The host system can request cryptographic operations to be performed on data using a handle to the key material. Thereby, the key material itself is not exposed and never leaves the measurement system. To achieve these goals, FreeRTOS has been extended with a secure syscall interface that allows a userspace task, e.g., the communication interface, to only execute well-defined operations. The Memory Protection Unit (MPU) provides hardened data protection such that an attacker cannot gain access to key material by taking over a single userspace task. Please note, the application domain was not the focus of this work. The given example is only intended to point out how such a PUF-based enclosure could be incorporated into a larger system.

8 Case Study

We present a case study that is based on the statistical evaluation of 115 top covers with a physical dimension of 140 mm × 140 mm and the test vehicle design as shown in Figure 1b and Figure 12. It is primarily based on an STM32F303 Cortex-M4F microcontroller running at 72 MHz for the evaluation unit. The cover design properties and the resulting

[4] This could be another tamper-evident PUF at the IC level or keys stored in Secure Non-Volatile Storage (SNVS) in COTS microcontrollers, i.e., the cover basically extends the physical trust domain of the IC to the whole enclosed area.

capacitive behavior are listed hereafter. Please note the significant difference in the order of magnitude between the capacitances.

- 16×16 electrodes resulting in 256 sensor nodes with $n = 1800$ sensor cells each
- $16/2 \times 16 = 128$ differential sensor nodes due to how the measurement circuit operates
- Parasitic capacitance: $C^{\mathrm{P}} \sim 1.8\,\mathrm{nF}$; mutual capacitance: $C^{\mathrm{M}} \sim 50\,\mathrm{pF}$; variation of differential capacitance: $C^{\mathrm{V}} < \pm 132\,\mathrm{fF}$; on average per-cell capacitance: $\overline{C_{\mathrm{c}}} = 50\,\mathrm{pF}/1800 = 27\,\mathrm{fF}$

Measurement Setup and Controlling Environmental Variables. For the statistical evaluation, each of the 115 top covers was measured with *the same circuit* and *in a temperature controlled room* such that influence by differing environmental temperature was insignificant, i.e., differences in the data is primarily rooted in the manufacturing variation of the covers. For all measurements, a regular work-bench power supply was used. As the measurement is based on an AC-measurement principle, variation in the DC-supply is not an issue within the tolerance of on-board regulators and the work-bench power supply. To ignore the influence of noise, each measurement of a cover was done at least 100 times in very short period of time such that averaging the result creates a virtually noise-free value. As additional precaution, we measured several covers using two different measurement boards, whereas the resulting data of the same covers were highly similar between the two boards with a correlation coefficient close to 1. Thus, the vast majority of entropy is extracted from the cover with negligible influence of the measurement circuit.

Determining Circuit Parameters. C^{P} is an important parameter to estimate the capacitive load on the amplifiers. For measuring C^{P}, all but one Tx *or* Rx electrode are grounded. A signal generator is connected to the remaining electrode and generates a sine wave of known frequency and amplitude. A high-precision ammeter is placed in between and measures the excitation current. Knowledge of current, voltage, and frequency allows to solve for the unknown parasitic capacitance C^{P}. We repeated this measurement for different voltages, frequencies, and electrodes to cross-verify the result.

Despite being capable of obtaining C^{M} on its own, the circuit behavior must be verified independently to ensure correctness of the measurement. To determine C^{M}, all but one Tx *and* one Rx electrode are grounded and a signal is applied to this one Tx electrode. In this experiment, the ammeter measures the current between the Rx electrode and ground such that the excitation terminal (Tx) is not the terminal of current measurement (Rx), i.e., a three-terminal measurement. Thereby, C^{P} is effectively ignored and only C^{M} remains which is then derived from the current. As this resulting value is rather small, i.e., few picofarads, several Rx electrodes were connected in parallel to validate the result in subsequent measurements. These manually obtained values match those from the circuit.

8.1 Details of Circuit Implementation

The implementation of the circuit is represented by the block diagram of Figure 11 and results in the physical design shown in Figure 12. Since it is based on COTS components, we tried to minimize the overall component count of the circuit while not specifically focusing on selecting the smallest package possible for convenience of research and development purposes, i.e., hand-soldering, debugging, etc. Still, the design results in a fully functional system with sufficient space left to accommodate an FPGA or similar components.

For the implementation on the Tx side, we make use of a single dual DAC output of the microcontroller that is multiplexed by a dual 1-to-8 MUX which feeds the inversely phased signals to the Tx electrode inputs. On the Rx side, the initial stage of processing has been implemented in parallel, i.e., 16 JFET-TIAs that convert the current signal to voltage. This voltage-signal is then multiplexed by a 16-to-1 MUX. Depending on the

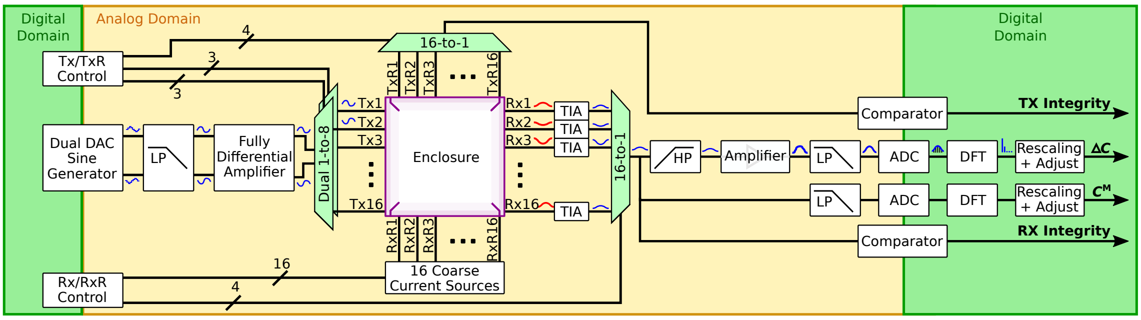

Figure 11: Block diagram of the entire measurement system for a 16 × 16 enclosure during differential capacitance measurement of the Tx1-2 group.

specific measurement mode, a different processing chain follows which is implemented sequentially. This is complemented by the components required for the integrity detection as included in Figure 11.

The chosen architecture represents a trade-off between signal robustness and the goal to reduce overall component count. In terms of absolute area, our implementation would clearly benefit from the development of an ASIC that integrates all previous analog components. This allows fully parallelizing the circuit on the Rx side such that they can be measured simultaneously. For the current circuit based on COTS components, performing a single differential measurement can be done in 270 µs for the analog part, i.e., excitation over eight periods and sampling. The digital signal processing requires another 258 µs, resulting in a total of less than 0.6 ms per node when performed sequentially.

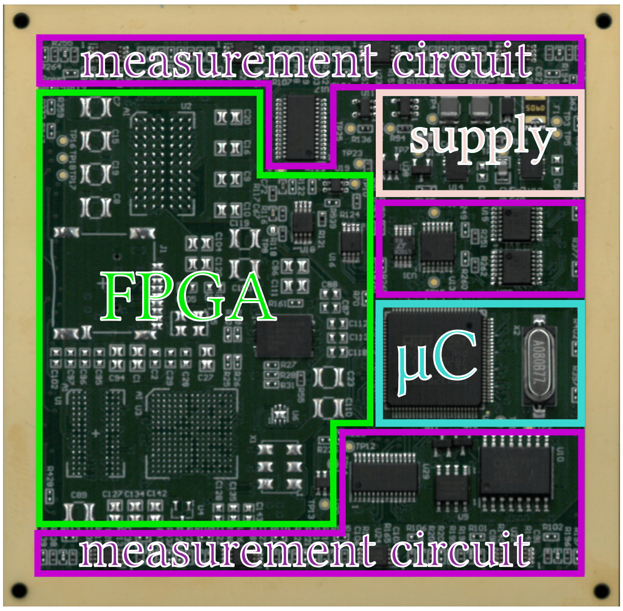

Figure 12: Top side of the test vehicle's inner security perimeter showing various system components. Please note, the components for the FPGA have not been populated.

Since the software part is performance-optimized with concurrent execution of code by peripheral units and double-buffering, *the signal processing is done in parallel to the next analog data acquisition*. Neglecting the overhead of context switches, measuring the whole cover differentially with the given circuit therefore only requires $128 \cdot 0.3\,\text{ms} = 38.4\,\text{ms}$ without oversampling. With additional oversampling, e.g., 10× or 20×, this increases to 384 ms or 768 ms. For the initial start-up of the device, this is an acceptable overhead for our use case to reduce the noise (cf. Figures 14a,14b,14c) and ensures optimized sensitivity towards the most advanced attacks that are likely to be carried out when the device is powered off. If the circuit would be fully parallelized on the Rx side, the differential acquisition could be done in less than $(16/2) \cdot 0.6\,\text{ms} = 4.8\,\text{ms}$ for the whole cover. Straightforward optimizations to speed up the measurement include but are not limited to increasing the excitation frequency and selecting a faster microcontroller such that we expect an improvement by a factor of 2× to 4× with moderate effort. Other performance figures of the circuit are:

- **Differential measurement:** full-scale range of ±132 fF at a theoretical digital resolution of $\Delta_{\text{M,Diff}} = 13.2\,\text{aF}$ (equivalent to 1 point) which is however limited by noise of $\sigma_{\text{N,Diff}} = 1.7\,\text{fF}$ without oversampling when the cover is connected.

- **Absolute measurement:** measurement range of approx. 0 pF to 100 pF at a theoretical digital resolution of $\Delta_{\text{M,Abs}} = 10\,\text{fF}$ (equivalent to 1 point) which is however limited by noise of $\sigma_{\text{N,Abs}} = 30\,\text{fF}$ when the cover is connected.

- **Power:** on average during a fully operational mode, the measurement system draws 132 mA on the analog rail and 43 mA on the digital rail at 3.3 V. This results in a power dissipation of 0.6 W. If needed, this can be lowered during run-time by, e.g., lowering the number of capacitance measurements per considered period.

8.2 Statistical Evaluation

In the following, let us consider basic statistics obtained from the measurement of 115 top covers. This is done for the differential measurement in Section 8.2.1. Afterwards, in Section 8.2.2, the corresponding results for absolute capacitance measurement are analyzed.

Exemplary Measurement Output. In Figure 13, an exemplary output of a *single* measured cover is shown. The output of the differential output is plotted in Figure 13a. Clearly visible is the random distribution of values in the range of −10 000 to +10 000 (in points). This is in contrast to Figure 13b which shows the output of the absolute capacitance measurement. The structural bias on a coarse-grained level becomes visible when zooming into the plot of the absolute capacitance measurement, more specifically, into the range of 4000 to 5000. This is best analyzed when considering the overall set of 115 covers as visualized in Figure 15d as part of the statistical evaluation of the absolute capacitance. It well reflects the expectation that directly neighboring electrodes have about the same nominal capacitance C^{N}.

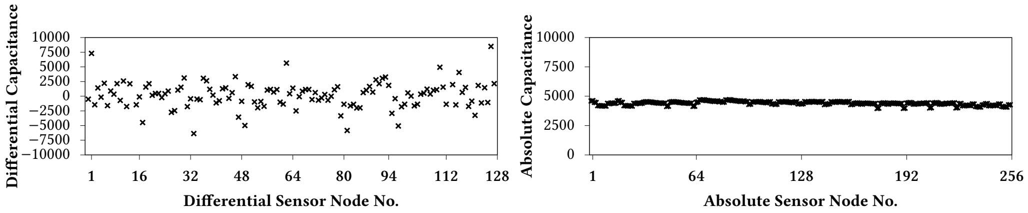

(a) Output of a differential measurement.

(b) Output of an absolute measurement.

Figure 13: Exemplary measurement output of a *single* cover to illustrate basic properties of the system. 200 samples over time were averaged to create a noise-free representation.

8.2.1 Statistical Evaluation of Differential Capacitance Measurement

The statistical evaluation of the differential measurement concentrates on the noise, the manufacturing variation, and the resulting entropy. This comprehensive evaluation strongly supports the chosen design rationale based on the provided data.

Measurement noise. In Figure 14a, the noise standard deviation $\sigma_{\mathrm{N,Diff}}$ of the differential measurement is plotted for each individual sensor node over the set of all 115 covers. Clearly visible is a mostly uniform behavior across the whole range of nodes and an expected value of $\bar{\sigma}_{\mathrm{N,Diff}} = 130$. Only Tx-group 6 (Tx11 and Tx12) shows a slightly degraded noise performance which may require further investigation. Without further adjusting the number of measurement periods, a direct oversampling of the values by a factor of 10 leads to the plot in Figure 14b with a reduced noise level of $\bar{\sigma}_{\mathrm{N},10} = 39$ which then requires 384 ms in our proof of concept implementation. Further increasing this to a 20× oversampling only reduces this to $\bar{\sigma}_{\mathrm{N},20} = 29$ at the cost of 768 ms (cf. Figure 14c). Even with this tremendous oversampling, resulting in an extremely low noise behavior, we would still be at an equal performance level compared to the solution of [VNK+15] whose authors state a measurement duration of 620 ms to 930 ms. To minimize the time for device start-up, we choose an oversampling of 10× while still reducing the noise.

The distribution of the occurring noise per node deviation (*not* of the noise itself which is Gaussian) is shown in Figure 14d and illustrates that the higher the noise is, the fewer occurrences are seen. Overall, this ensures a high level of confidence in the low noise behavior of the design which is essential for PUF-based tamper-evident applications. Of course, with a fully parallelized implementation of the circuit in an ASIC, both noise level and measurement duration are likely to be further improved.

Manufacturing variation. In Figure 14e, the device-specific standard deviation of the observed capacitance values is plotted with an expected value of $\bar{\sigma} = 2290$. To investigate the question whether there are "weak" spots of little deviation, we created Figure 14f which shows the standard deviation of the capacitance values per sensor node. This is of particular importance within the context of physical attacks, since we assume that the PUF entropy is spatially distributed. If this would not be the case, an attacker may characterize the PUF by gaining partial knowledge of its distribution from previously analyzed devices and then use this knowledge to attack that specific location of the cover where the standard deviation is the smallest, thereby minimizing the damage. As supported by the plot in Figure 14f is in terms of variation the differential measurement indeed a suitable approach to prevent such structural bias or imperfections, thereby avoiding the risk of the aforementioned attack scenario. There are only two nodes that appear to have a rather low manufacturing variation. However, as seen in Figure 14h this stems from the fact that the corresponding sensor nodes are affected by a structural bias in their expected value, causing some of the variation to hit the limit of the measurement range. This is an imperfection of the layout due to the irregular shape of the top cover and will be addressed in the next hardware revision.

To complement these tests, we used Welch's t-test as proposed in [IHOS17] to create Figure 14k and Figure 14l. As indicated by Figure 14h are the means across different Tx groups different. Figure 14k clearly supports that this difference is statistically relevant, i.e., the considered PDFs are distinguishable by their first statistical moment, indicating a structural bias that is present across different Tx groups. In contrast, Figure 14l only shows few comparison that exceed the threshold of $|t| > 4.5$, i.e., the differences in the variation of Figure 14f are not statistically relevant most of the time. As our data processing attempts to extract only the variation by removing first-order structural bias, this confirms the good PUF behavior at the stage of the raw data already.

Entropy (Global Analysis). Figure 14g shows the PDF of $\Delta C = \gamma$ and contains all sensor nodes from all covers. Its standard deviation σ is 2241 points which equals 29.58 fF. To compute the entropy, we apply an equidistant quantization [IHKS16]. Its bin size Δ_{Q} is chosen as multiples of the noise deviation $\sigma_{\mathrm{N,Diff}}$, thereby making the result more robust.

For $\Delta_{\mathrm{Q}} = 2 \cdot 3.29 \cdot \bar{\sigma}_{\mathrm{N,Diff}} \approx 11.3\,\mathrm{fF}$, the computed Shannon entropy yields 3.45 bit per node. With 10× oversampling, this changes to $\Delta_{\mathrm{Q}} = 2 \cdot 3.29 \cdot \bar{\sigma}_{\mathrm{N,Diff,10}} \approx 3.4\,\mathrm{fF}$ resulting in 5.2 bit per node. Hence, a total of $128 \cdot 5.2\,\mathrm{bit} \approx 665\,\mathrm{bit}$ can be expected from the PUF under ideal conditions. Using the given $y = 3.29$ for the quantization, we experience an average error rate of $\leq 0.1\%$ per differential sensor node at room temperature. For the full temperature range of −20 °C to +60 °C, the results are presented in Section 8.5.

Entropy (Spatial Analysis). To further investigate inter-dependencies of neighboring nodes from an information-theoretic point of view, we developed an extension of the Context Tree Weighting (CTW) method [WST95, ISS+06] which we call Spatial CTW (or SCTW in short). Due to how we interpret the PUF output, this spatial extension is based on q-ary symbols as opposed to bits. Hence, the differences to the classical CTW are: instead of considering the successive bit of a context does our approach operate on the successive higher-order alphabet symbols, in addition, we consider a context comprising all nodes within a certain spatial radius around the targeted node, whereas a radius of 1 corresponds to a tree depth of 8 and a radius of 3 to a tree depth of 48. This analysis can be interpreted as follows: if an attacker would be able to destroy one node only and obtain all values of the surrounding nodes, what is the remaining conditional entropy left to reconstruct the single destroyed node. In our case, for a total of 32 quantization intervals, the obtained results of the SCTW analysis were 3.1 bit for a radius of 3, the same for a radius of 2, and 3.7 bit for a radius of 1, i.e., lower than the Shannon entropy, indicating a minor degradation in entropy due to inter-dependency of values. Still, the results support the properties of the overall design. Details of this approach will be published at a later stage.

8.2.2 Statistical Evaluation of Absolute Capacitance Measurement

For the sake of completeness, we include the statistical properties of the absolute capacitance measurement. While they are by far less critical for the contained PUF, they are nevertheless important for the overall design to provide consistency with our assumptions regarding the differential measurement. The statistical evaluation of the absolute capacitance measurement is done on the same data set of 115 flexPCB covers. In Figure 15a, the noise standard deviation per node is shown. Clearly visible is that the noise of the absolute capacitance measurement only has a minor impact on the data acquisition, i.e., $\bar{\sigma}_{\mathrm{N,Abs}} = 3$ which is equivalent to ±30 fF.

To analyze the absolute capacitance variation, we provide Figures 15b and 15c that show a per-device average absolute capacitance varying in the range of 40 pF to 50 pF while the per-node standard deviation is approx. at 4 pF. Few outliers are observed that are attributed to bending the flaps which induces mechanical stress resulting in miniature cracks in the copper tracks, as the bending radius is rather tight. In Figure 15d is the per-node mean of the capacitance shown. While there is a distinct pattern, it is also visible that data points occur in pairs, i.e., directly neighboring absolute capacitance nodes indeed have a highly similar nominal capacitance. This supports our previous arguments regarding the differential measurement and chosen pairwise electrode layout.

8.3 PUF Properties – Uniqueness and Reliability

Originally defined by Maiti et. al. in [MCMS10] and [MGS11], the metrics *Uniqueness* and *Reliability* have become the de facto standard for PUF publications. While various other metrics have been proposed, they can still be considered as a starting point to assess the fundamental PUF properties, i.e., if their values sufficiently differ from each other and if they can be reconstructed reliably. While we recommend to always complement them with additional tests, we nevertheless focus on these two most common metrics with regard to higher-order alphabets which is owed to their popularity. In the following, the PUF responses R'_i are considered for the Uniqueness and Reliability, as this represents the

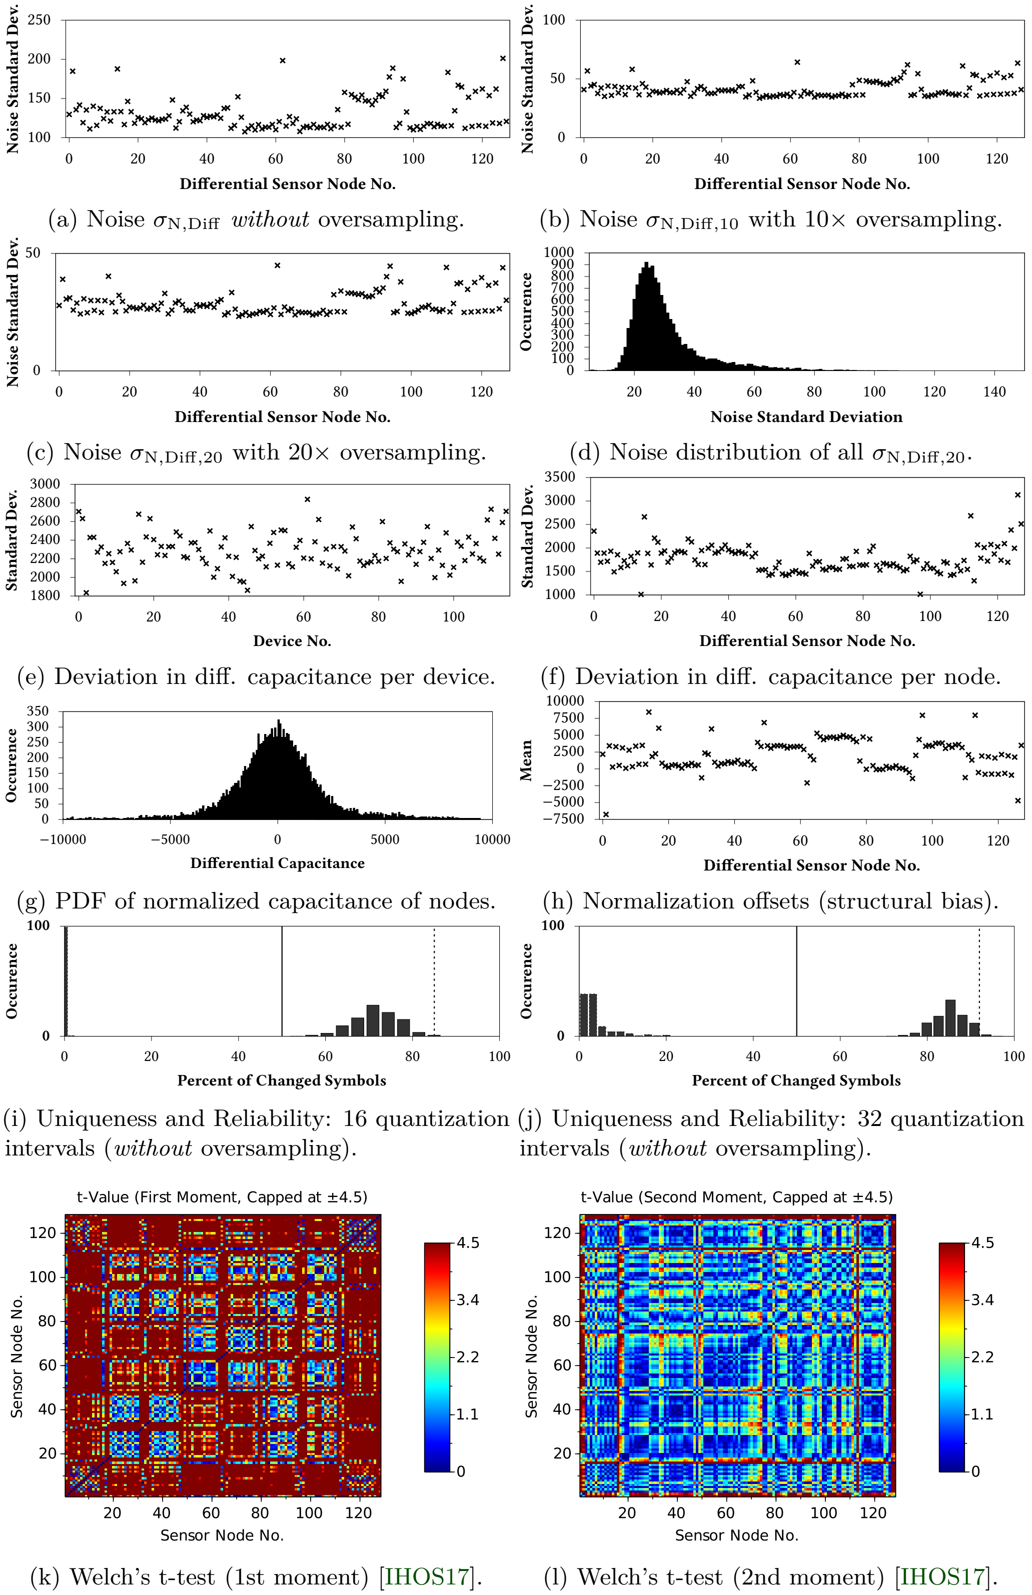

(a) Noise $\sigma_{\mathrm{N,Diff}}$ *without* oversampling.

(b) Noise $\sigma_{\mathrm{N,Diff,10}}$ with 10× oversampling.

(c) Noise $\sigma_{\mathrm{N,Diff,20}}$ with 20× oversampling.

(d) Noise distribution of all $\sigma_{\mathrm{N,Diff,20}}$.

(e) Deviation in diff. capacitance per device.

(f) Deviation in diff. capacitance per node.

(g) PDF of normalized capacitance of nodes.

(h) Normalization offsets (structural bias).

(i) Uniqueness and Reliability: 16 quantization intervals (*without* oversampling).

(j) Uniqueness and Reliability: 32 quantization intervals (*without* oversampling).

(k) Welch's t-test (1st moment) [IHOS17].

(l) Welch's t-test (2nd moment) [IHOS17].

Figure 14: Statistical evaluation of 115 flexPCB covers (differential capacitance).

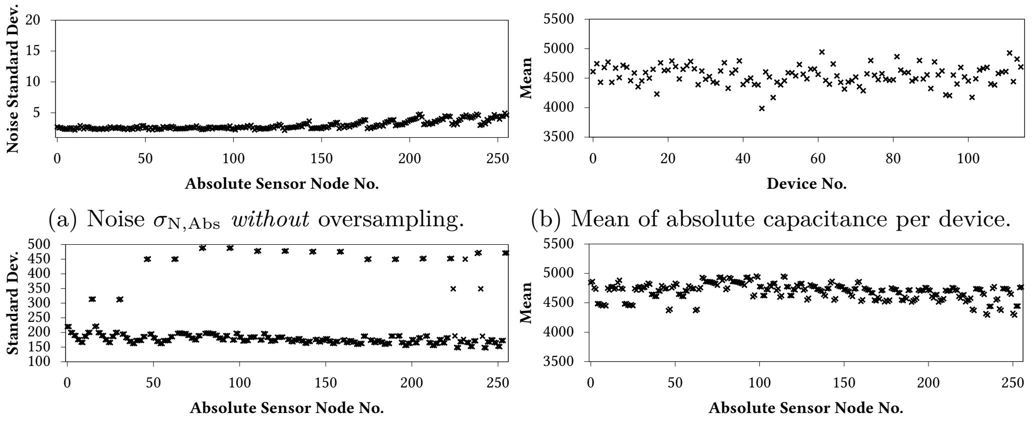

(a) Noise $\sigma_{\mathrm{N,Abs}}$ *without* oversampling.

(b) Mean of absolute capacitance per device.

(c) Deviation of absolute capacitance per node.

(d) Mean of absolute capacitance per node.

Figure 15: Statistical evaluation of 115 flexPCB covers (absolute capacitance).

interface to the subsequent ECC. Considering the classical definition of Uniqueness with k being the number of PUF devices and n the number of bits of each PUF

$$\text{Uniqueness}_{\text{HD,non-weighted}} = \frac{2}{k(k-1)} \sum_{i=1}^{k-1} \sum_{j=i+1}^{j=k} \frac{\text{HD}(R_i, R_j)}{n} \cdot 100\% \tag{7}$$

it is evident that it is based on the Hamming Distance (HD) as metric to rate how many *substitutions* are necessary to change one fixed-length bit string into the other. Please note, the definition of Hamming Distance not only holds true for binary-strings *but also* for strings from a higher-order alphabet, i.e., it is possible to substitute the bits in Equation 7 with the symbols from the higher-order alphabet. In general, this equation shows how many percent of the positions differ between PUF responses. For the binary case, the optimum for the Uniqueness is 50 %, i.e., half of the bits change on average. This optimum is based on a uniform distribution of a binary-PUF.

However, as the alphabet size increases from binary to more symbols, the expected output of this metric changes. For an alphabet size of 4 which is equal to having 4 quantization intervals, it is expected that 75 % of the positions differ, again assuming a uniform distribution. This already increases to 87.5 % for 8 symbols. For non-uniform distributions, e.g., Gaussian, the expected number of symbols to change decreases in comparison to the uniform distribution. This can be calculated for i.i.d symbols by

$$\text{ExpectedChange} = \sum_{i=1}^{i=q} \Pr(A_i)(1 - \Pr(A_i)) \cdot 100\% \tag{8}$$

for an alphabet $A = \{A_1, A_2, \dots A_q\}$ of size q. Each individual probability of a symbol is dependent on the PUF distribution and its specific quantization interval from which it is generated from. To reflect the properties of the Uniqueness for the binary case, we choose a lower bound of 50% and the upper bound as [50%, ExpectedChange], i.e., the resulting histogram must be in this range to consider the PUF as unique. Alternatively, the lower bound could be chosen based on a stochastic model to provide a stronger rationale. Since ExpectedChange is the best value a PUF can achieve given a Gaussian distribution without noise, we expect that most empirical data will fail to actually reach that bound. Unlike the Uniqueness for binary-PUFs, we now have a metric that better complements the entropy contained in the PUF, as a range of values is acceptable to consider a PUF as unique. Please note, it is still possible to adapt the given metric to mimic the behavior of the binary uniqueness scenario by weighting the output. Moreover, to better reflect the properties of certain q-ary codes in the future, it is of interest using either the *Lee* or *Manhattan* distance and defining the Uniqueness correspondingly.

For the Reliability, we adhere to the previous definition of [MCMS10], i.e., a change of a symbol to any other (no matter its distance) is counted as one. A suitable ECC could then either be based on Reed-Solomon (RS) codes [RS60] employed in a fuzzy commitment scheme or insertion/deletion codes with a variable-length bit mapping of symbols [IHL+17].

The resulting behavior for both metrics based on our data set is illustrated in Figure 14i and Figure 14j (*without* using oversampling). The minimum boundary of 50% is illustrated as a solid vertical line, while ExpectedChange as a dotted line. For 16 quantization intervals, the reliability is very high while the Uniqueness is centered between the two defined boundaries. Now, when increasing the number of intervals this increases the entropy we can extract from the PDF and the histogram of the Uniqueness moves closer to the dotted line which by itself also moves towards 100%. At the same time, since the width of the quantization interval reduces, the effect of the noise becomes more dominant, thereby clearly affecting the Reliability. Overall, the results show that interpreting the PUF output as a higher-order alphabet nicely complements previous works in this domain, while opening up a new path for ECCs, i.e., working on higher-order alphabets instead of a binary PUF output.

8.4 Practical Security Analysis

In the following, we provide practical evidence for the difficulty of tampering with the cover without causing detection. Clearly, it is not possible to exhaustively cover all possible attacks in a single paper. Hence we do not claim a complete protection against all attacks. Instead, it should be considered as a study on the presented enclosure concepts to demonstrate that practically carrying out a successful attack would be challenging, in particular when considered as a black-box with limited prior knowledge, i.e., a real-world design would include additional obfuscation techniques to increase uncertainty for the attacker. The choice of parameters for the quantization is based on the results of Section 8.5 and accounts for possible changes in the environment, too. Hence, realistic parameters are chosen to assess the intrusion detection. These parameters are $\Delta_{\mathrm{Q}} \approx 500$ as quantization width Q_{w} which leads to 40 quantization intervals and a Shannon entropy of 4.18 bit per differential node. This corresponds to a *min*-entropy of 3.46 bit per differential node. We note that this is based on a 10× oversampling for system startup.

8.4.1 Invasive Attacks: Drilling

To investigate the tamper-evident properties of our enclosure with respect to the assumed attacker model, we attacked several covers by drilling various types of holes and carrying out attempted repairs. Thus, our focus is on open-circuits and corresponding repairs, as short-circuits, especially on the Tx side are prone to cause damage to the circuit. There is no plausible benefit for the attacker to deliberately cause such short-circuits. For drilling, we used a multifunction rotary tool (a “dremel”) with corresponding workstation as shown in Figure 16a. High revolutions per minute (RPM) are required to not break the fragile drill bits illustrated in Figure 16c. Since the structure size is chosen with respect to the assumed minimum drill diameter of 0.3 mm, it is guaranteed that at least one Tx *and* Rx electrode are cut-off. Therefore, larger drill sizes will cause even more damage.

For smaller drill sizes than 0.3 mm that are outside of the assumed attacker model, there is still a reasonable chance of sufficient damage to cause detection, e.g., a diameter of 0.2 mm is still guaranteed to break at least one Tx *or* Rx electrode. Even for 0.1 mm, there is still a chance left to break electrodes based on the position of the drill hole. Please note, for all drilling attempts that severed electrodes, we had to *disable* the integrity check first, i.e., without attempted repairs and *independent* from the PUF-properties this would already allow the system to determine that an attack was carried out. In the following, we study several attack profiles that we chose based on our understanding of the system[5].

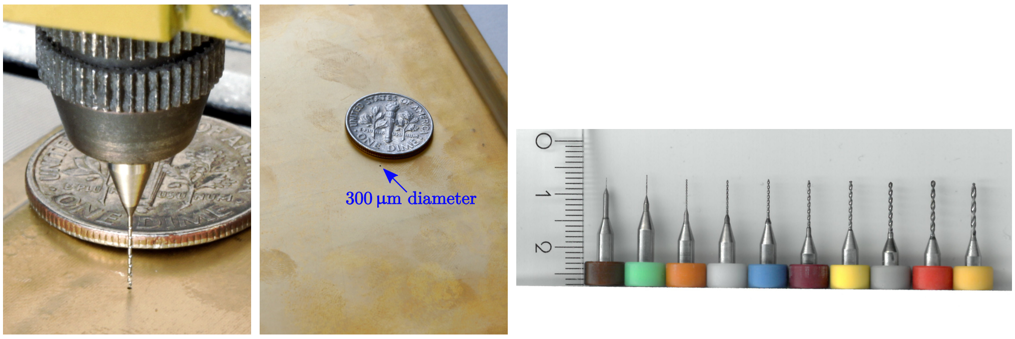

(a) Attack close-up. (b) Hole from distance. (c) Drill bits from 0.1 mm to 1.0 mm.

Figure 16: Exemplary attack on cover with 300 µm drill and a US dime as reference showcasing the disproportion of attack size to overall size of cover.

Attack Profile 1 (P1): Single 0.3 mm Hole. Beginning of Tx. As a start, we created a single hole of 0.3 mm relatively close to the beginning of a Tx electrode. In this case, the affected electrodes were Tx8 and Rx2. The resulting plot in Figure 17a shows the noise-free *difference* of the differential capacitances from before and after the attack, i.e., the nodes were measured 200 times and averaged to remove the noise.

As the Tx pair consisting of electrodes Tx7 and Tx8 is no longer balanced, a dramatic change for the whole group of differential nodes is observed. Since Rx2 is destroyed also, it shows up as significant change in all the other Tx groups. Rx1 also appears to have taken damage but is not flagged by the integrity check as broken. Moreover, cut-off electrode parts lead to improper grounding, creating a changed coupling behavior which in turn results in additional shifts for a majority of the other nodes at the stage of the discretized PUF data. For the specific attack considered, all but one of the nodes have significantly moved away from their enrollment such that they would have had a different value during reconstruction. Hence, recovery of the key either by direct measurement of the cover or extracting the circuit's data would have been infeasible.

To complement the differential measurements, we show the result of the difference in absolute capacitance in Figure 17b. The significant change in values is easily detectable by Tamper Detection C, i.e., the change is larger than 15% of the absolute capacitances' mean. By computing the difference to the mean, drift effects such as temperature would be accounted for even under different environmental conditions (see Section 8.5).

Attack Profile 2 (P2): Single 0.3 mm Hole. Center of Tx. As next step, we started over with a new cover. This time we created a single hole of 0.3 mm in the center of a Tx electrode to balance the cut-off parts of both Tx and Rx electrodes. The affected electrodes were Tx9 and Rx10. Figure 18a shows the resulting plot of the change in differential capacitance.

Since Tx9 does no longer create a balanced Tx pair with Tx10, again a severe change for the whole corresponding group of differential values is observed. As Rx10 is destroyed also, it shows up as significant change in all the other Tx groups. Due to a more centered destruction of Tx and Rx is the global change in the coupling behavior not as significant when compared to P1. Still, some additional shifts in several other nodes occur.

In general, experimental results support the argument that for a hole of 0.3 mm, the whole Tx group (one column) and the affected Rx group (one row) are always sufficiently

[5] Here, we want to emphasize that for some of these attacks, it took us several attempts in carrying out the attack strategy as intended, even though the text neglects this fact.

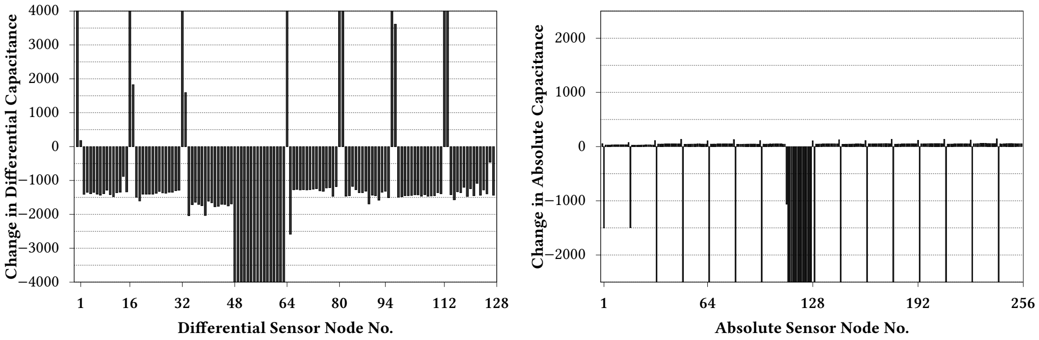

(a) Change in differential capacitance (P1).

(b) Change in absolute capacitance (P1).

Figure 17: Attack Profile 1 (P1): result of a single hole of 0.3 mm in diameter, severing electrode Tx8 and Rx2. Clearly visible is the significant change in values.

altered, resulting in at least $8 + 16 - 1 = 23$ destroyed nodes, i.e., nodes that shift by ≥ 500 points. Hence, we expect that at least $23 \cdot 3.46$ bit $= 80$ bits of *min*-entropy are destroyed by a single hole without attempted repairs. Taking into account that only a fraction of differential nodes happen to reside on the center of a quantization interval, it is likely that for most practical experiments more nodes differ from the quantized value of their enrollment even for smaller shifts. For the specific cover of Figure 18, we observed a total of 47 differential nodes that would have moved away from their enrollment. This is still idealized in the sense that we are considering noise-free values, i.e., an attacker would need to deal with noisy values which increases the difficulty of an attack. Hence, the actual loss in entropy would have been even higher under real-world conditions. Moreover, results for the difference in absolute capacitance in Figure 18b again provide strong evidence that in addition to the loss in entropy, the attack would have been detected upon power-on prior to generating the key.

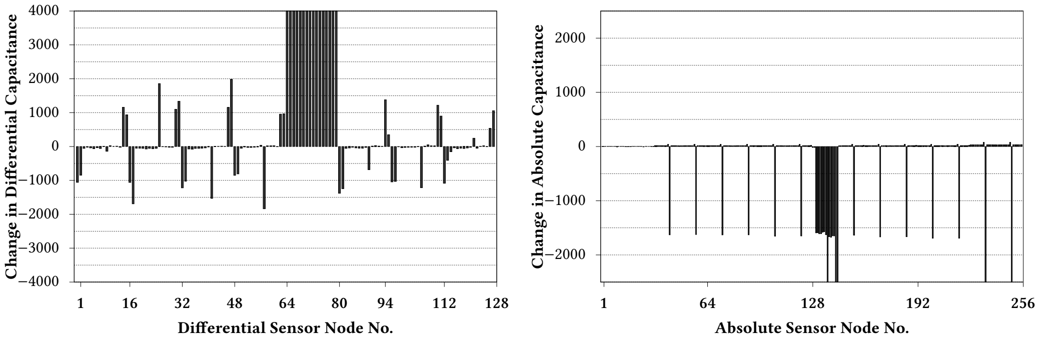

(a) Change in differential capacitance (P2).

(b) Change in absolute capacitance (P2).

Figure 18: Attack Profile 2 (P2): result of a single hole of 0.3 mm in diameter, severing electrode Tx9 and Rx10. Clearly visible is the significant change in values.

Attack Profile 3 (P3): Two-Holes of 0.3 mm. Additional Tx Damage. For the next step of the analysis, again a new cover was used. This time, two holes of 0.3 mm in diameter were created *while aiming at shortest cut-offs of the electrodes.* The first hole severed Tx5 and Rx10. To minimize the damage of the overall attack, we created the second hole such that only Tx10 was additionally cut-off. This is possible by penetrating the cover at a spot where Rx10 is cut-off once more. The resulting damage of the differential capacitance measurement is shown in Figure 19a. As expected, we see two devastating shifts in two Tx groups. Moreover, we see a result that is consistent with P1, i.e., a global shift occurs which indicates a severely degraded behavior within the cover due to improper grounding of unused signals. This would again render almost all capacitive nodes destroyed. From this result, we deduce that the more damage to Tx electrodes is done, the worse is the global shift. We confirmed this behavior for other attacks causing more damage. Hence, even when aiming at shortest cut-offs, it is improbable for an attacker to succeed without attempted repairs.

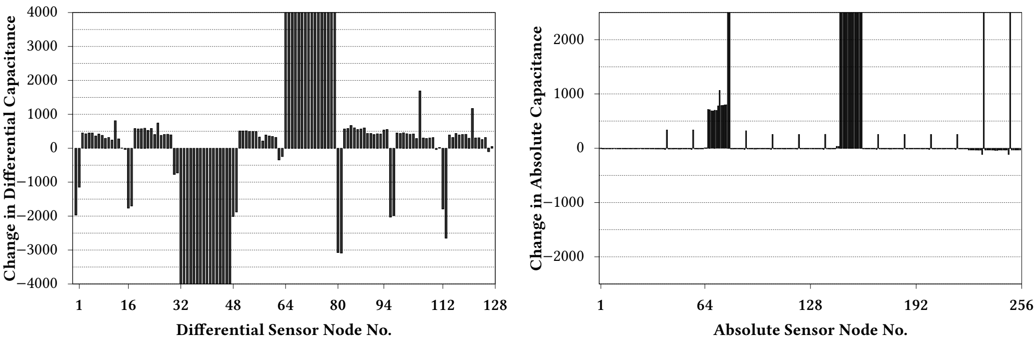

(a) Change in differential capacitance (P3).

(b) Change in absolute capacitance (P3).

Figure 19: Attack Profile 3 (P3): attack with two holes of diameter 0.3 mm severing Tx5, Tx10, and Rx10.

In the plot of Figure 19b showing the difference in absolute capacitance we again see severe changes in the capacitive behavior, too. Clearly visible are the two groups as result of the two broken Tx electrodes. Moreover, when comparing the data between the first and second hole, we see a difference in the change of the Rx10 electrode which is owed to the two different points where it was damaged (plot omitted). Hence, by using the information drawn from the absolute capacitance measurement it is possible to provide a spatial estimate of where the attack took place.

Attack Profile 4 (P4): Single Hole of 0.33 mm, Symmetric Rx Cut-Off. For the analysis of the next attack profile, again a new cover was used. This time, an uncommon drill bit of 0.33 mm in diameter was used to create a hole of approximately the same diameter. The affected Tx electrodes were Tx2 and Rx1 and Rx2. Based on geometrical considerations, we consider this as a perfect symmetric cut-off of the electrodes Rx1 and Rx2. This attack leads to the change in differential capacitance as shown in Figure 20a. Again, we hit the Tx electrode more towards its beginning, resulting in a severe shift in all values due to a much larger portion of the electrode that has been cut-off. Hence, if an attacker would not be able to repair any damage, the best strategy for the current circuit implementation (e.g., when not measuring from both sides) would be to attack electrodes such that the cut-off parts are the shortest and farthest away from the excited input.

Clearly visible is the overall severe damage that does not justify a more detailed analysis. Furthermore, the change in absolute capacitance as shown in Figure 20b also indicates an attack. Hence, there is no advantage in attempting a symmetric Rx cut-off.

Attack Profile 5 (P5): Single Hole of 0.33 mm, Symmetric Tx Cut-Off. For this attack, we continued with the cover used in P1. To do so, we hit the previous 0.3 mm hole with our 0.33 mm drill bit. This caused the additional destruction of Tx7, creating a symmetric cut-off with Tx8. The resulting change in capacitance of Figure 21 should be compared to Figure 17 of P1. It is interesting to see that the previously assumed damage of Rx1 is now mostly gone in addition to the observed global shift in the values. However, the damage in Rx2 remains, as expected from the result of the failed integrity check. Moreover, while the damage in the Tx group was significantly lowered from more than 10 000 points to slightly less than $\sim$ 4 000, it is still present, clearly indicating an

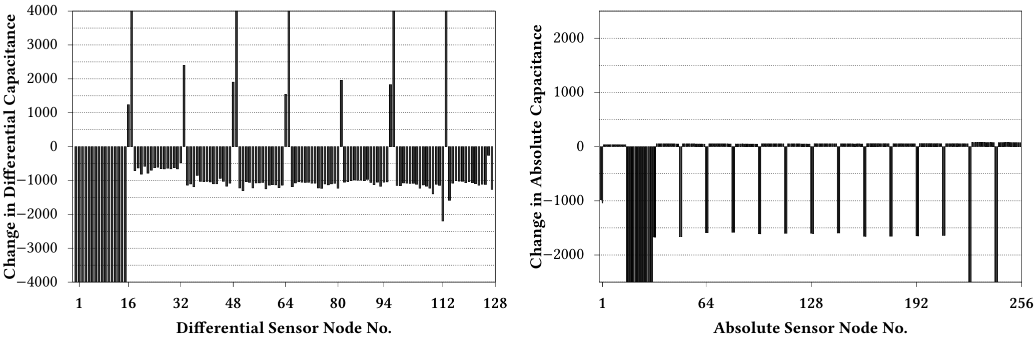

(a) Change in differential capacitance (P4).

(b) Change in absolute capacitance (P4).

Figure 20: Attack Profile 4: attack with a single hole of diameter 0.33 mm and symmetric Rx cut-off. Here, severing electrodes Tx2, Rx1, and Rx2.

attack. Taking the results of this attack and previous attack profiles into account, it is highly improbable to succeed in attacking the device without doing attempted repairs.

An attacker may still want to aim for symmetric Tx cut-offs to minimize the effects due to imbalanced Tx pairs. However, when aligning these results with the absolute capacitance measurement of Figure 21b, it is evident that the attack would have been detected both by the differential and absolute capacitance measurement. Hence, the absolute capacitance measurement provides additional assurance to detect attacks that aim at tricking the behavior of the differential measurement. Careful consideration of the absolute capacitance measurement data as a kind of associated data for the privacy amplification stage of the PUF key derivation will be an option investigated in the future.

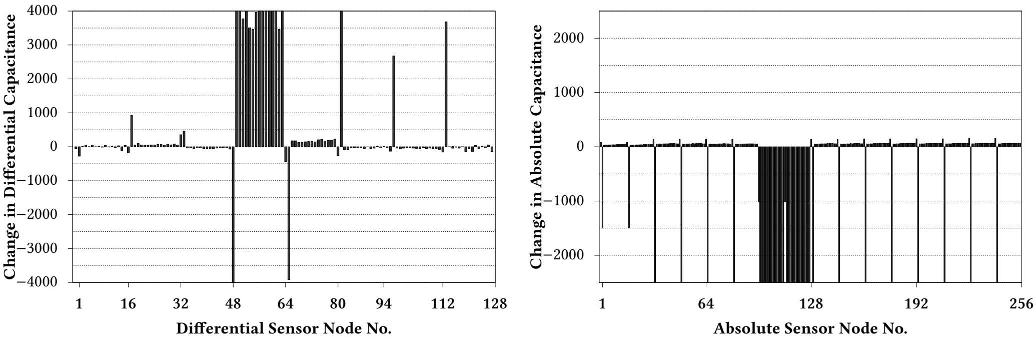

(a) Change in differential capacitance (P5).

(b) Change in absolute capacitance (P5).

Figure 21: Attack Profile 5 (P5): result of a single hole of 0.33 mm in diameter, severing electrode Tx7, Tx8, and Rx2. Due to having a single hole is the cut-off of Tx7 and Tx8 considered symmetric.

Attack Profile 6 (P6): Advanced Attack with Attempted Repair. As a next step, we push the concept to its limits by first drilling a hole with 5 mm and then simulating a real attack by means of analyzing the localized electromagnetic emanation (EM) of an IC as shown in Figure 22a. We chose the position for the hole such that the attacker would minimize the cut-off parts of the electrodes and at the same time, allow for the largest hole possible without exceeding a 2 × 2 node square as illustrated in Figure 6a. Moreover, we repaired the damage caused by the attack by reconnecting the severed electrodes, namely Tx11, Tx12, Rx11, and Rx12 using ultra-thin copper wire. A larger hole would have affected more electrodes and make this attack more complex in terms of repair.

To account for attackers exceeding our own capabilities and to simulate tasks we consider practically extremely challenging, we simplified the following steps as part of the attack. Prior to mounting the cover and carrying out the attack, the IC was decapsulated. No heatsink was mounted such that between the drilled hole and the IC no material had to be removed. While the repair of the affected Tx electrodes was done from the outside, we reconnected the broken Rx electrodes *on the inside* prior to mounting the cover. Since the finalized assembly prevents a non-destructive cover removal this is a noticeable simplification to not consider the effort required of reaching the Rx layer through the Tx layer and performing a miniature repair. Alternatively, a hole would need to be made to pull the bottom layer of the cover outwards and do the same (without breaking the remainder of the electrodes).

The resulting plot of the change in differential capacitance is shown in Figure 22b. While the damage is quite significant, it can be seen also that it is not as devastating thanks to the repairs. Still, a total of 18 nodes would have been destroyed, i.e., still exceeding the threshold of the subsequent ECC scheme and causing a total of $18 \cdot 3.46\,\text{bit} = 62\,\text{bits}$ of *min*-entropy to be destroyed. While the loss in entropy drops to a level that is no longer considered computationally infeasible, we need to emphasize that the practical complexity of carrying out the attack in addition to the computational effort is still high, especially when considering the corresponding amount of Shannon entropy. Moreover, there is no doubt that based on the results of the absolute capacitance measurement as shown in Figure 22c would raise an alarm, too.

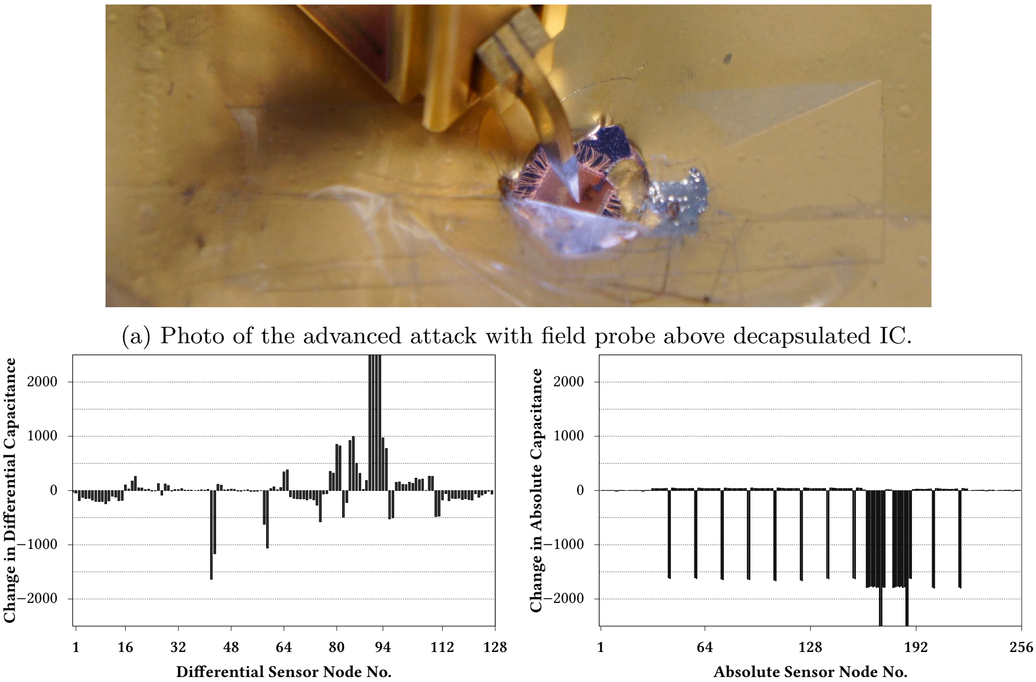

(a) Photo of the advanced attack with field probe above decapsulated IC.

(b) Change in differential capacitance (P6).

(c) Change in absolute capacitance (P6).

Figure 22: Attack Profile 6 (P6): Using drill of 5 mm with subsequent Tx *and* Rx repair.

Attack Profile 7 (P7): Advanced Attack with Attempted Repair. We performed another advanced attack by testing the limits of this concept with holes of 300 µm in diameter and attempted repairs. The corresponding attack is shown in Figure 23a. As stated beforehand, we are of the opinion that compromising the enclosed system by making one hole only is not practically feasible due to the complex IC-level checks made. Instead, multiple such holes would need to be made at several strategic positions, necessitating more rework which in turn increases the likelihood for an attacker to make mistakes.

The drilled hole of 300 µm in diameter destroyed the integrity of Tx3 and Rx4 as result of the attack (before the repair). Figure 23b presents the change in differential capacitance from before the attack to after the attack including the attempted repair. Clearly visible is that the imbalance in the Tx pair due to the repair is insufficient to cause a shift in the values across the group. What remains is the Rx damage in all Tx excitation groups. To take advantage of the specific behavior of such attacks which we derived from previous analyses, we chose the specific location for the attack based on our knowledge of the actual values. Still, a total of 8 nodes would have moved away from their designated values, allowing for the attack to be detected but no longer representing an effort considered computationally infeasible, assuming the attacker would be able to obtain the measurement data just by using this hole alone. While we are unaware of how such a small hole with attempted repair could be used to compromise the underlying system, we fairly show the limits of our concept when using commercially available manufacturing technology only, i.e., a customized technology limiting the repairability of holes will help mitigate the risk of such attacks, e.g., by doping the carrier substrate with randomized dielectric particles and/or customized material for the electrode tracks.

When considering the results of the absolute capacitance measurement in Figure 23c, we again see a striking difference in the capacitive behavior, allowing the detection of the attack. This emphasizes the importance of combining different measurement principles to make physical attacks more difficult to perform. As pointed out later in Section 8.4.3, minimizing the impact on the absolute capacitance could possibly be minimized by performing a "Frankenstein" attack for which we lacked the necessary tools.

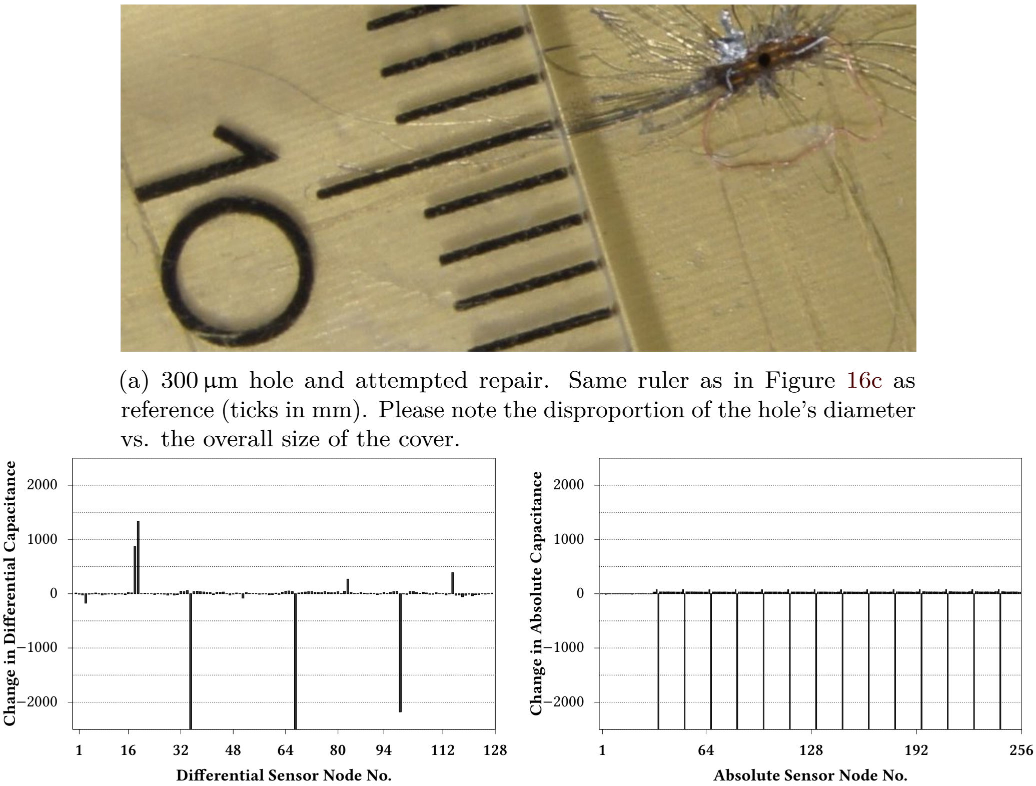

(a) 300 µm hole and attempted repair. Same ruler as in Figure 16c as reference (ticks in mm). Please note the disproportion of the hole's diameter vs. the overall size of the cover.

(b) Change in differential capacitance (P7).

(c) Change in absolute capacitance (P7).

Figure 23: Attack Profile 7 (P7): Using drill of 300 µm with subsequent repair.

Conclusions on Attack Profiles. We have practically and fairly evaluated the security of the cover based on the assumed attacker model under various drilling attacks including attempted repairs. The overall result is that attacks without attempted repairs are detected with very high probability. By carrying out more advanced attacks with attempted repairs while allowing some simplifications to be made, we have also openly shown the limits of the concept that cannot be fully overcome without more advanced manufacturing technology for the enclosure. Still, the combined use of differential and absolute capacitance measurement is a promising approach to detect a majority of physical intruders even when repairs are attempted. Moreover, during our white-box testing, we could disable countermeasures at will and focus on the effects seen in the measurement data. In other situations this was also helpful, e.g., when reconnecting electrodes, as this is a laborious task and alignment errors are easily made such that the wrong electrodes would be mistakenly connected. Since the PUF data acquisition and tamper detection is done within a complex IC, disabling the detection logic while not destroying more entropy of the cover appears challenging.

8.4.2 Non-Invasive Attacks: Optical Inspection and Probing

One of the possible threats of PUF-based enclosures is that an attacker may learn the PUF by means of optical inspection, i.e., contactless techniques that are non-invasive and therefore impossible to detect after attempted use when the device is powered on. As part of a more detailed analysis, we studied drill holes with the help of a Shimadzu SMX 6000 scanning system which is intended for PCB failure analysis and allows 2D and 3D X-ray imaging. The resulting 2D X-ray image of a drill hole with 200 µm and its surrounding mesh is shown in Figure 24.

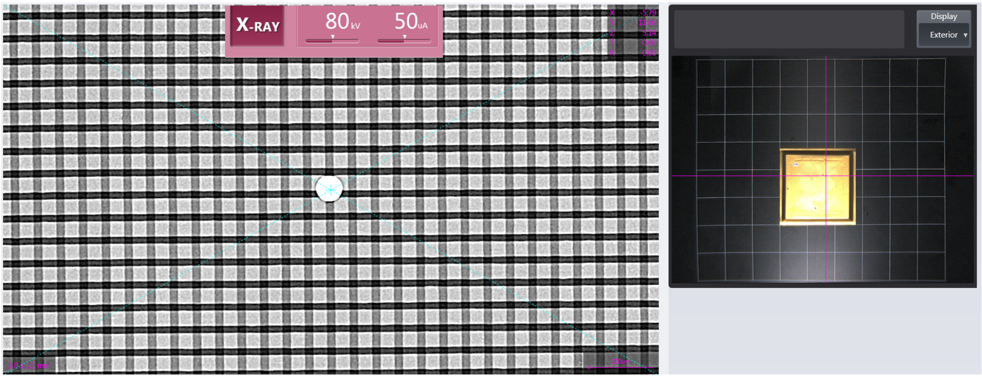

Figure 24: X-ray based two-dimensional (2D) optical inspection of cover with 200 µm hole.

It was *not* necessary to remove the cover's shield, i.e., it is possible to see through the solid copper plane. The same applies when considering the resulting 3D X-ray image as shown in Figure 25. Neither in 2D, nor in 3D, it is possible to identify locations from which specific information on the PUF could be derived, i.e., other than a highly regular structure of the mesh there is no revealing information visible. This is within the scope of our expectation due to the following reasons:

- For the 2D case, the obtained image is from a bird's eye view, i.e., the 3D structure of the fuzzy edges of the PCB tracks cannot be resolved. Likewise it is not possible to analyze the surface roughness in between the tracks from the outside, even for the 3D case, at least with the imaging technology we had at hand.

- While the 3D structure of the mesh becomes visible under 3D imaging technology, we still could not derive useful information from these images about the PUF values.

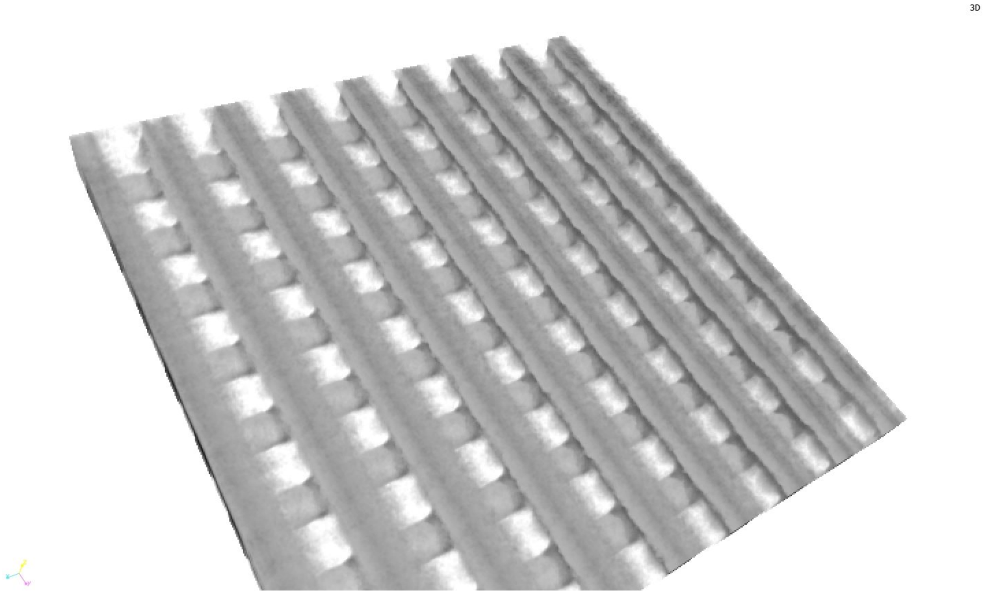

Figure 25: X-ray based three-dimensional (3D) optical inspection of mesh.

- Assuming the PUF deviation could be observed to a certain degree, it is still mandatory to look at the accumulated deviation over all sensor cells per node, i.e., an automated tool would need to extract the deviation per sensor cell which entails a certain error due to limited resolution, etc. This error then accumulates over the sum of all cells per node and would severely falsify the obtained value when compared to the actual value.
- Upon manual inspection of the images, there are no obvious patterns or marks visible (aside from manufacturing defects) that would justify further analysis with regard to optical inspection.

Other optical attacks include Laser Voltage Probing (LVP), as for example used in [YPER99]. To the best of our knowledge is this technique designed for IC analysis only, as it requires a p-n junction to work correctly. Moreover, it is beyond our own expertise if a *current* (as opposed to voltage) signal in the lower nanoampere range could be optically probed. We are currently unaware of other analysis techniques in this domain that could help to optically probe signals on bare tracks inside the flexPCB.

8.4.3 Discussion of Additional Attacks

Since it is not possible to exhaustively cover all possible attacks in a single paper, let us briefly consider a selection of other attacks and how they have been considered in the design. Note that some attacks require additional countermeasures which are outside the scope of the cover itself, e.g., preventing data remanence or having a sufficiently internally buffered supply to enable zeroization even if an attacker pulls the power during runtime.

Bending/Prying Open the Cover. In general, there are two types of flexPCB offered. One type is for static flexing, i.e., a one-time bending to fit the flexPCB in the packaging design. When targeting this application, it is common to choose an adhesiveless carrier, i.e., the same we use. In contrast, for dynamic flexing where the flexPCB must be bent multiple times as part of the functionality, it is common to choose carriers with flexible adhesives to minimize strain when bending the flexPCB. Since our flexPCB is intended for one-time bending and has been manufactured correspondingly, it is difficult to not create cracks when bending it in reverse direction of the previous assembly process. As prying open the cover causes severe mechanical stress, either breaking it or creating cracks in the copper tracks. Moreover, without X-rays, such cracks cannot be located through the solid copper plane of the shield which makes it difficult to repair them, too.

Careful Cover Disassembly and Measurement with Attacker's Circuit. The goal of this attack would be to extract the cover's PUF key without the actual device, i.e., to carefully disassemble it without destroying the PUF behavior. Since the packaging concept including its potting have been specifically designed to thwart such attempts of an easy cover removal, it is not possible to remove the cover without severely damaging it. The whole unit has *not* been designed to allow servicing of its components, even by its legitimate device owner.

Assuming the cover could be removed, the attacker would still need to replicate the measurement circuit with utmost care. Due to the specifics of the electrode setup, e.g., its massively parallel structure, disproportion of different capacitances contributing to the measurement, and the small-scale differential capacitance, it is highly unlikely to use a standard LCR-meter to carry out the measurement of C^{V} in a useful way. In a certification process, this would add to the complexity of the attack even despite the fact that this is not theoretically impossible.

Imposter Attack. The goal of this attack would be an undetected disassembly, successful tampering with an IC on the inside, and re-assembly. Due to the same reasons stated above, we consider cover disassembly not as a well-founded choice for the attacker. In addition to these difficulties related to that would an imposter attack imply that an attacker has not only been able to secretly circumvent all countermeasures that are checked by the device itself but also tamper-evident properties that are visible to the human eye. For example, optical inspection of the unit prior to putting it in the field would notice differences in the particle-mix of the potting, possible damage, etc.

"Frankenstein" Attack. Since our laser intended for IC failure analysis could not be used to cut or drill flexPCB material, we could not carry out attacks where pieces of one flexPCB cover are used to repair damage done to another flexPCB cover. This requires a precision setup to not violate the underlying design rules of the electrodes, i.e., a matched cell-overlap of differential electrode pairs. We point out that cutting and putting back pieces of flexPCB material entails a significant amount of work for reconnecting each of the cut lines, adding to the complexity of the attack the larger the piece is. A better approach would be knowing the size of the piece targeted for removal and to manufacture a corresponding piece where the wiring is done internally, such that only the outside connectors would need to be reconnected appropriately, thereby reducing the work of reconnecting lines. However, our findings for batches from two different manufacturers and even batch-to-batch differences indicate that there will still be noticeable differences in the PUF behavior, making this attack still reasonably difficult to perform.

Physically Probing Electrodes. An attacker might try to probe electrodes directly to measure their capacitance or eavesdrop signals. This requires access to all electrodes, as properly connecting unused ones is mandatory for the measurement. This claim is not only supported by our practical experience but also the plots presented as part of the attack profiles as unconnected parts of an electrode degrade the measurement. At the same time, the shield would need to be partially removed at multiple spots, causing the surrounding field to change, thereby falsifying the results. Repeatedly carrying out these steps without making errors along the way is considered challenging. Moreover, even state-of-the-art micro probes [GGB04a] add a capacitive load of $> 20\,\mathrm{fF}$ which exceeds the observed standard deviation in differential capacitance. We are currently developing customized circuitry to investigate the feasibility of such an attack.

Side-Channel Attacks. Emanations of the system are prevented by the heatsink, shielding layers, and the supply lines are additionally protected with filters. Moreover, the

Tx layer carries only insensitive excitation signals, i.e., the attacker would only see the 33.3 kHz of the excitation signal without the possibility to derive useful information from it. In contrast, the Rx layer carries sensitive signals in the lower nano-ampere range, making it difficult to eavesdrop on them. The measurement itself is otherwise time-constant.

8.5 Environmental Tests

To analyze the robustness of our approach, we carried out tests in the temperature range of −20 °C to +60 °C using a VT 4011 temperature chamber by Vötsch as illustrated in Figure 26a. We tested this with a single board and three top covers, i.e., the assembly was not finalized and no potting was used to enable the measurement of different covers using the same circuit. Overall, we observed a highly similar behavior for the covers.

When both cover and measurement circuit are subject to these environmental influences this causes a certain temperature drift in the values as shown in Figure 26c for the absolute capacitance measurement. The plateau regions illustrate the differences in temperature with steps of 10 °C. Clearly visible is the direct relation of temperature to change in value and that the spread of values relative to the overall mean per sample point in time is relatively constant. In fact, the absolute capacitance measurement could be exploited as a coarse-grained temperature sensor for Environmental Failure Protection (EFP), too.

This behavior is incomparable to the raw differential capacitance prior to compensation, as shown in Figure 26d. Here, we see a much weaker pattern from the temperature cycle which is only barely visible. Moreover, as the differential nodes have different values, they behave slightly different. For a constant temperature level, the lines would be going straight from left to right. Here, we do see that larger differential capacitances tend to have a larger drift when compared to smaller differential capacitances that are apparently less affected by temperature. For a representative group of values with larger and smaller capacitances which is based on one Tx pair, the maximum drift d_E after compensation is less than 130 points as illustrated in Figure 26b.

To counteract this remaining drift effect, we need to lower the number of quantization intervals, i.e., increase their width to $Q_\mathrm{w} = 2 \cdot y \cdot \sigma_{\mathrm{N,Diff,10}} + 2 \cdot d_\mathrm{E}$. Accordingly, for $y = 3.29$, the number of quantization intervals is reduced to 40 while the Shannon entropy drops to 4.17 bit per node. Even for drifts that are much higher, e.g., up to $d_\mathrm{E} = 400$ points, we conveniently stay above 3 bit of Shannon entropy per node. Without implementing more advanced compensating techniques, the temperature drift is fully accounted for by the increased width of the quantization interval. Hence, there is no need to improve the error-correcting capability of the subsequent ECC scheme. As the quantization interval width is only ~ 500 points (based on 40 intervals), it is still possible to reliably detect the damage of the physical attacks as presented beforehand. Erroneous differential nodes as result from an erratic behavior under temperature effects is typically less than three to five nodes such that a sufficient gap is ensured to destroyed nodes from physical attacks. Hence, an attacker would try to attack the system at the temperature of enrollment to exploit the ECC to possibly correct damage made by the attack.

Aging. We also performed tests for accelerated aging of the foils, i.e., heating up to +110 °C for drying at a relative humidity of < 10%, then exposing the covers to +90 °C at a relative humidity of 85% with another drying cycle afterwards. This procedure was repeated several times. In between each step, we measured the values to determine their behavior, i.e., the measurement circuit was *not* subject to this accelerated aging to assess the properties of the covers independently of a possible aging in circuit components. After this test, the majority of values returned to their designated values of the enrollment with very small error margin (typically much less than 30 points). This is not unexpected, since flexPCB is typically rated for much worse conditions. The only nodes with critical behavior were located in the flaps, as they were not mechanically secured by a conformal coating or potting for our tests, resulting in mechanical stress due to expansion of the

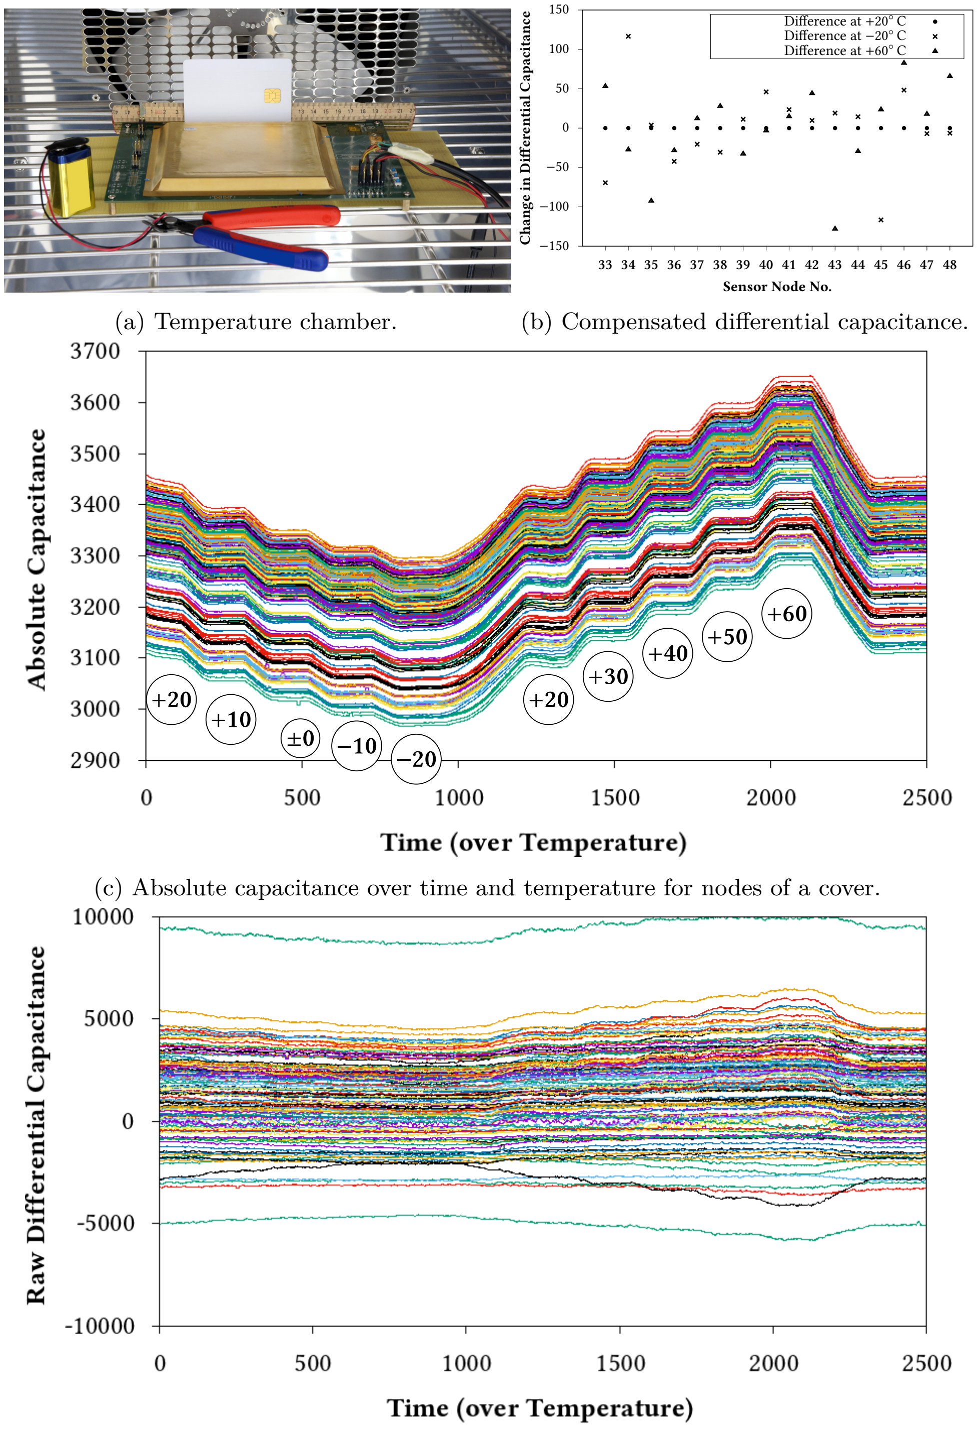

(a) Temperature chamber.

(b) Compensated differential capacitance.

(c) Absolute capacitance over time and temperature for nodes of a cover.

(d) *Raw* (unprocessed) differential capacitance over time and temperature.

Figure 26: Environmental tests and results. Plots in Figure 26c and Figure 26d have the identical time axis, i.e., they both cover the temperature range from +20 °C to −20 °C, then to +60 °C, and back to room temperature during the same test cycle.

material. This is owed to the fact that for the purpose of measuring the covers, we needed to mount and unmount the covers which would not have been possible when finalizing their assembly, i.e., applying the potting and securing the seams would have prevented this. In the future, a measurement IC is developed which is why the aging behavior of the chosen COTS components was not of relevance. We additionally point out that for aging, it is always an option to re-enroll the device in the field if necessary.

9 Conclusion and Outlook

In this work we analyzed how to enclose a device with a cover that is evaluated using a batteryless security concept while still detecting a majority of physical intruders. We implemented our proposed full-stack approach and experimentally verified the PUF-behavior based on the statistical measurements of 115 covers. Overall, the results are promising and a step towards tamper-resistant enclosures without battery. The advantage of such enclosures is the protection of arbitrary components, thereby providing an additional layer of security as required for high-performance cryptographic modules and their compliance with certification standards. As part of this work, we also designed and tested the analog circuit in addition to the digital data processing, resulting in a fully usable system based on FreeRTOS that could directly be employed in such systems.

The basic idea of our approach is to extend tamper-evident PUFs by the concept of integrity, i.e., an intertwined mechanism that allows a binary decision w.r.t. to the structural integrity of the device. Moreover, we complement this concept by several runtime tamper detection mechanisms that ensure continuous protection during run-time without the necessity of repetitive PUF key generations. A “full scope” measurement by means of a differential *and* absolute capacitance measurement makes it practically impossible to tailor an attack that is able to trick both measurements at the same time.

Our comprehensive tests provide initial evidence that the concept fulfills the targeted requirements, i.e., statistical results in addition to attacks and environmental tests confirm the chosen design rationale. When comparing our academic study with previous industrial solutions, it is however evident that our material properties should be further improved to provide an even higher level of security by making attempted repairs more difficult. Moreover, a layout randomization is currently not implemented, due to the limitation of using COTS components for the measurement circuit. We developed a corresponding concept for layout randomization that allows to dynamically change the layout based on a challenge for the PUF measurement, i.e., based on this challenge, different pieces of the cover are then combined for the measurement similar to a puzzle. Further improvements could be the measurement from both sides of the electrodes. Hence, our results clearly facilitate future research as the presented concepts are generic and do not depend on the chosen manufacturing technology or circuit implementation. Hence, it should be considered as a hint of what could be achieved with different manufacturing technologies such as panel level integration [OBSL15] or more advanced manufacturing technologies with custom tailored materials for either the tracks [PMT08] or the carrier materials.

Several ongoing and future developments have been pointed out in the paper, e.g., transferring our knowledge from a discrete circuit into an analog sensor IC thereby reducing the circuit’s area requirement, speeding up the data acquisition, and providing a higher level of integration for improved security. Furthermore, we plan to update the physical design such that the outcome results in a bimodal distribution, i.e., a double-peaked PDF with a local minimum in the center which is aligned with the value 0 of the differential measurement. This has the benefit of increased value shift upon attacks and consequently makes them more difficult to perform.

Acknowledements

The authors of Fraunhofer AISEC would like to express their deepest gratitude for having the opportunity to collaborate with DSO National Laboratories. Moreover, we would like to thank all former and present team members including but not limited to: (in alphabetical order): Lukas Auer, Maxim Hennig, Matthias Hiller, Ludwig Kürzinger, Aysun Önalan, Karthik Uppund. Results of the Spatial Context Tree Weighting (SCTW) are joint work with Michael Pehl of TUM (*to be published*).

References

[ADDS91] D. G. Abraham, G. M. Dolan, G. P. Double, and J. V. Stevens. Transaction security system. *IBM Systems Journal*, 30(2):206–229, 1991.

[BDHV07] Ileana Buhan, Jeroen Doumen, Pieter Hartel, and Raymond Veldhuis. Fuzzy extractors for continuous distributions. In *ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, page 353, 2007.

[BGS+08] Christoph Bösch, Jorge Guajardo, Ahmad-Reza Sadeghi, Jamshid Shokrollahi, and Pim Tuyls. Efficient helper data key extractor on FPGAs. In Elisabeth Oswald and Pankaj Rohatgi, editors, *Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, volume 5154 of *LNCS*, pages 181–197. Springer Berlin / Heidelberg, 2008.

[BOU07] BOURNS INC. Application note – security housing, 2007. http://application-notes.digchip.com/176/176-48205.pdf.

[Bri04] Gary A. Brist. Design Optimization of Single-Ended and Differential Impedance PCB Transmission Lines. In *PCB West Conference Proceedings*, 2004.

[CT06] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. John Wiley & Sons, New York, second edition, 2006.

[DGSV15] Jeroen Delvaux, Dawu Gu, Dries Schellekens, and Ingrid Verbauwhede. Helper data algorithms for PUF-based key generation: Overview and analysis. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 34(6):889–902, 2015.

[dGVL12] Joep de Groot, Boris Škorić, Niels de Vreede, and Jean-Paul Linnartz. Quantization in continuous-source zero secrecy leakage helper data schemes. Cryptology ePrint Archive, Report 2012/566, 2012. https://eprint.iacr.org/2012/566.

[DRS04] Yevgeniy Dodis, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In Christian Cachin and Jan L. Camenisch, editors, *Advances in Cryptology (EUROCRYPT)*, volume 3027 of *LNCS*, pages 523–540. Springer Berlin / Heidelberg, 2004.

[EFK+12] Thomas Esbach, Walter Fumy, Olga Kulikovska, Dominik Merli, Dieter Schuster, and Frederic Stumpf. A New Security Architecture for Smartcards Utilizing PUFs. In *ISSE Conference*, 2012.

[ES05] H. Eren and L.D. Sandor. Fringe-Effect Capacitive Proximity Sensors for Tamper Proof Enclosures. In *Sensors for Industry Conference*, 2005.

[GCDD02] Blaise Gassend, Dwaine Clarke, Marten van Dijk, and Srinivas Devadas. Silicon Physical Random Functions. In *ACM CCS*, 2002.

[GGB04a] GGB Industries Inc. Picoprobe Model 19C, 2004. Available online: www.ggb.com/PdfIndex_files/mod18c.pdf, as of October 10, 2016.

[GGB04b] GGB Industries Inc. T-4 Series Tungsten Probe Tips, 2004. Available online: http://www.ggb.com/t-4.html, as of October 10, 2016.

[GLM$^+$04] Rosario Gennaro, Anna Lysyanskaya, Tal Malkin, Silvio Micali, and Tal Rabin. Algorithmic tamper-proof (ATP) security: Theoretical foundations for security against hardware tampering. In Moni Naor, editor, *Theory of Cryptography Conference (TCC)*. 2004.

[Goe58] Gerald Goertzel. An algorithm for the evaluation of finite trigonometric series. *The American Mathematical Monthly*, 65(1):34–35, 1958.

[Hew09] Hewlett-Packard Company. Atalla Cryptographic Subsystem (ACS) Security Policy (compliant to FIPS 140-2 level 4), July 2009. https://csrc.nist.gov/CSRC/media/projects/cryptographic-module-validation-program/documents/security-policies/140sp1174.pdf.

[Hew10] Hewlett-Packard Company. Atalla Cryptographic Subsystem (ACS) Security Policy (compliant to FIPS 140-2 level 3), October 2010. https://csrc.nist.gov/csrc/media/projects/cryptographic-module-validation-program/documents/security-policies/140sp1441.pdf.

[HMSS12] Matthias Hiller, Dominik Merli, Frederic Stumpf, and Georg Sigl. Complementary IBS: Application specific error correction for PUFs. In *IEEE International Symposium on Hardware-Oriented Security and Trust (HOST)*, pages 1–6, 2012.

[HNT$^+$13] Clemens Helfmeier, Dmitry Nedospasov, Christopher Tarnovsky, Jan Starbug Krissler, Christian Boit, and Jean-Pierre Seifert. Breaking and entering through the silicon. In *ACM Conference on Computer and Communications Security (CCS)*, 2013.

[HYKD14] Charles Herder, Mandel Yu, Farinaz Koushanfar, and Srinivas Devadas. Physical Unclonable Functions and Applications. *Proceedings of the IEEE*, 102, 2014.

[HYS16] Matthias Hiller, Mandel Yu, and Georg Sigl. Cherry-picking reliable PUF bits with differential sequence coding. *IEEE Transactions on Information Forensics and Security*, 11(9):2065–2076, 2016.

[IBM12] IBM. IBM 4765 Cryptographic Coprocessor Security Module Security Policy (compliant to FIPS 140-2 level 4), December 2012. https://csrc.nist.gov/CSRC/media/projects/cryptographic-module-validation-program/documents/security-policies/140sp1505.pdf.

[IHKS16] Vincent Immler, Maxim Hennig, Ludwig Kürzinger, and Georg Sigl. Practical aspects of quantization and tamper-sensitivity for physically obfuscated keys. In *Workshop on Cryptography and Security in Computing Systems (CS2)*, pages 13–18. ACM, 2016.

[IHL$^+$17] Vincent Immler, Matthias Hiller, Qinzhi Liu, Andreas Lenz, and Antonia Wachter-Zeh. Variable-Length Bit Mapping and Error-Correcting Codes for Higher-Order Alphabet PUFs. In *Security, Privacy, and Applied Cryptography Engineering (SPACE)*, 2017.

[IHOS17] Vincent Immler, Matthias Hiller, Johannes Obermaier, and Georg 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.

[IMJFC13] Phil Isaacs, Thomas Morris Jr, Michael J. Fisher, and Keith Cuthbert. Tamper Proof, Tamper Evident Encryption Technology. In *Pan Pacific Symposium*. SMTA, 2013.

[IOK+18] Vincent Immler, Johannes Obermaier, Martin König, Matthias Hiller, and Georg Sig. B-TREPID: Batteryless tamper-resistant envelope with a PUF and integrity detection. In *IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*, pages 49–56, 2018.

[ISS+06] T. Ignatenko, G. Schrijen, B. Skoric, P. Tuyls, and F. Willems. Estimating the secrecy-rate of physical unclonable functions with the context-tree weighting method. In *2006 IEEE International Symposium on Information Theory*, pages 499–503, July 2006.

[Joi15] Joint Interpretation Library. *Application of Attack Potential to Hardware Devices with Security Boxes*. SOGIS, December 2015.

[JW99] Ari Juels and Martin Wattenberg. A fuzzy commitment scheme. In *ACM Conference on Computer and Communications Security (CCS)*, pages 28–36, 1999.

[KLR08] Wolfgang Killmann and Kerstin Lemke-Rust. Common Criteria Protection Profile - Cryptographic Modules, Security Level "Enhanced", July 2008.

[MCMS10] A. Maiti, J. Casarona, L. McHale, and P. Schaumont. A Large Scale Characterization of RO-PUF. In *IEEE International Symposium on Hardware-Oriented Security and Trust (HOST)*, pages 94–99, 2010.

[MGS11] Abhranil Maiti, Vikash Gunreddy, and Patrick Schaumont. A systematic method to evaluate and compare the performance of physical unclonable functions. Cryptology ePrint Archive, Report 2011/657, 2011. https://eprint.iacr.org/2011/657.

[MSSS11] Dominik Merli, Dieter Schuster, Frederic Stumpf, and Georg Sigl. Semi-invasive em attack on fpga ro pufs and countermeasures. In *Proceedings of the Workshop on Embedded Systems Security*, page 2. ACM, 2011.

[MV10] Roel Maes and Ingrid Verbauwhede. A discussion on the Properties of Physically Unclonable Functions. In *TRUST*, 2010.

[MVHV12] Roel Maes, Anthony Van Herrewege, and Ingrid Verbauwhede. *PUFKY: A Fully Functional PUF-Based Cryptographic Key Generator*, volume 7428 of *LNCS*, pages 302–319. Springer Berlin / Heidelberg, 2012.

[Nat02] National Institute of Standards and Technology (NIST). *FIPS PUB 140-2: Security Requirements for Cryptographic Modules*. NIST, Gaithersburg, MD, USA, May 2002.

[OBSL15] A. Ostmann, C. Boehme, K. Schrank, and K. Lang. Development of a micro-camera with embedded image processor using panel level packaging. In *2015 European Microelectronics Packaging Conference (EMPC)*, pages 1–4, Sept 2015.

[OHHS18] Johannes Obermaier, Florian Hauschild, Matthias Hiller, and Georg Sigl. An embedded key management system for puf-based security enclosures. In *2018 7th Mediterranean Conference on Embedded Computing (MECO)*, June 2018.

[OI18] Johannes Obermaier and Vincent Immler. The past, present, and future of physical security enclosures: From battery-backed monitoring to PUF-based inherent security and beyond. *Journal of Hardware and Systems Security*, Aug 2018.

[OIHS18] Johannes Obermaier, Vincent Immler, Matthias Hiller, and Georg Sigl. A measurement system for capacitive puf-based security enclosures. In *Proceedings of the 55th Annual Design Automation Conference, DAC 2018, San Francisco, CA, USA, June 24-29, 2018*, pages 64:1–64:6, 2018.

[Pay12] Payment Card Industry Security Standards Council. *Payment Card Industry PTS HSM Security Requirements, v2.0.* PCI, Wakefield, MA, USA, May 2012.

[Pay13] Payment Card Industry Security Standards Council. *Payment Card Industry PTS POI Modular Derived Test Requirements, v4.0.* PCI, Wakefield, MA, USA, May 2013.

[PMT08] P. Paul, S. Moore, and S. Tam. Tamper protection for security devices. In *2008 Bio-inspired, Learning and Intelligent Systems for Security*, pages 92–96, Aug 2008.

[RS60] Irving Reed and Golomb Solomon. Polynomial codes over certain finite fields. *Journal of the Society of Industrial and Applied Mathematics*, 8(2):300–304, 06/1960 1960.

[SAS17] Taras Stanko, Fitria Nur Andini, and Boris Skoric. Optimized quantization in zero leakage helper data systems. *IEEE Transactions on Information Forensics and Security*, 2017.

[SFIC14] Merrielle Spain, Benjamin Fuller, Kyle Ingols, and Robert Cunningham. Robust keys from physical unclonable functions. In *IEEE International Symposium on Hardware-Oriented Security and Trust (HOST)*, pages 88–92, 2014.

[Sko05] Sergei P. Skorobogatov. Semi-invasive attacks – A new approach to hardware security analysis. Technical Report UCAM-CL-TR-630, University of Cambridge, Computer Laboratory, April 2005.

[Sko17] Sergei Skorobogatov. How microprobing can attack encrypted memory. In *Euromicro Conference on Digital System Design, DSD 2017, Vienna, Austria, August 30 - Sept. 1, 2017*, pages 244–251, 2017.

[SSAQ02] 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. Proceedings.*, 2002.

[TFL$^+$17] Shahin Tajik, Julian Fietkau, Heiko Lohrke, Jean-Pierre Seifert, and Christian Boit. Pufmon: Security monitoring of fpgas using physically unclonable functions. In *23rd IEEE International Symposium on On-Line Testing and Robust System Design, IOLTS 2017, Thessaloniki, Greece, July 3-5, 2017*, pages 186–191, 2017.

[TSS$^+$06a] Pim Tuyls, Geert-Jan Schrijen, Boris Skoric, Jan van Geloven, Nynke Verhaegh, and Rob Wolters. Read-Proof Hardware from Protective Coatings. In *CHES*, volume 4249 of *LNCS*. 2006.

[TSŠ+06b] Pim Tuyls, Geert-Jan Schrijen, Boris Škorić, Jan Van Geloven, Nynke Verhaegh, and Rob Wolters. Read-proof hardware from protective coatings. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 369–383. Springer, 2006.

[TSS+06c] Pim Tuyls, Geert-Jan Schrijen, Boris Skoric, Jan van Geloven, Nynke Verhaegh, and Rob Wolters. Read-proof hardware from protective coatings. In Louis Goubin and Mitsuru Matsui, editors, *Workshop on Cryptographic Hardware and Embedded Systems (CHES)*, volume 4249 of *LNCS*, pages 369–383. Springer Berlin Heidelberg, 2006.

[Ult17] Ultra Electronics AEP. Advanced Configurable Cryptographic Environment (ACCE) v3 HSM Crypto Module (compliant to FIPS 140-2 level 4), December 2017. https://csrc.nist.gov/CSRC/media/projects/cryptographic-module-validation-program/documents/security-policies/140sp2793.pdf.

[VNK+15] Michael Vai, Ben Nahill, Josh Kramer, Michael Geis, Dan Utin, David Whelihan, and Roger Khazan. Secure architecture for embedded systems. In *IEEE High Performance Extreme Computing Conference (HPEC)*, 2015.

[Wei00] Steve H. Weingart. Physical Security Devices for Computer Subsystems: A Survey of Attacks and Defenses. In *CHES*. 2000.

[W.L07a] W.L. GORE & Associates Inc. GORE Secure Encapsulated Module (Commercial Brochure), 2007.

[W.L07b] W.L. GORE & Associates Inc. GORE Tamper Respondent Surface Enclosure (Commercial Brochure), 2007.

[WSL+15] Lingxiao Wei, Chaosheng Song, Yannan Liu, Jie Zhang, Feng Yuan, and Qiang Xu. BoardPUF: Physical Unclonable Functions for Printed Circuit Board Authentication. In *IEEE/ACM ICCAD*, 2015.

[WST95] F. M. J. Willems, Y. M. Shtarkov, and T. J. Tjalkens. The context-tree weighting method: basic properties. *IEEE Transactions on Information Theory*, 41(3):653–664, May 1995.

[WTM08] Hidehito Watanabe, Hideo Tsuzaka, and Masami Masuda. Microdrilling for Printed Circuit Boards – Influence of radial run-out of microdrills on hole quality. *Precision Engineering Journal of The International Societies for Precision Engineering and Nanotechnology – PRECIS ENG*, 2008.

[YD10] Mandel Yu and Srinivas Devadas. Secure and robust error correction for physical unclonable functions. *IEEE Design & Test of Computers*, 27(1):48–65, 2010.

[YPER99] Wai Mun Yee, M. Paniccia, T. Eiles, and V. Rao. Laser voltage probe (lvp): a novel optical probing technology for flip-chip packaged microprocessors. In *Proceedings of the 1999 7th International Symposium on the Physical and Failure Analysis of Integrated Circuits*, pages 15–20, 1999.

A More on Related Work

Selected details of battery-backed tamper-responding approaches are presented in [OI18]. In general, public information on these approaches is scarce. Therefore, we provide additional information regarding their differences to our implementation in the following.

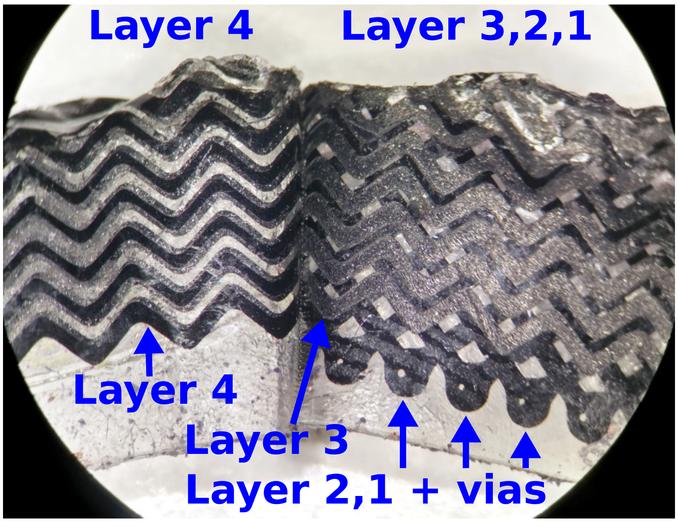

Figure 27: Cut-off piece of the GORE envelope, showing all four layers. The sheets comprising layer 3 and 4 have been partially pulled appart for illustration purposes.

Figure 27 shows all four layers of the GORE envelope which has been partially pulled apart. This process revealed that layer 1 and 2 are based on diagonally running traces whose width approximately equals their trace spacing. Vias are present near the edge of the envelope that connect layer 1 to layer 2 and vice versa. This forms a regular mesh of traces that runs diagonally across the surface. However, these traces still leave gaps between them that would be unprotected.

Atop of these traces, two additional layers were placed that close all existing gaps or holes, i.e., layer 3 and 4. They are running in a zig-zag pattern from left to right. Their trace width is larger than their spacing, thereby increasing the sensitive portion of each sensor layer. As depicted in the figure, the traces of layer 3 run at exactly along the gaps that are present in the tracks on layer 4. Thus, this layout creates an entirely protected surface without any visible gaps.

As seen in the figure, pulling apart the layers has a destructive effect on the traces, i.e., the layer stack-up has been strategically designed to provide low tensile strength only and therefore separates easily. The resistance of layer 3 significantly changed since the trace still partially sticks to the backside of the substrate of layer 4 while its remainder is left at the opposite side. Apart from this are the sensor tracks rather fragile, i.e., upon mechanical contact they tend to damage easily. Altogether, the GORE envelope has shown to be based on custom-tailored material properties that significantly contribute to system security.

Compared to the GORE envelope which has a rather coarse structure is our solution based on traces with a width and space of only 100 µm. Both are shown in a side-by-side comparison in Figure 28, whereas the GORE envelope is at the bottom and our approach is at the top. According to the image, the trace width and spacing differs by a factor of five.

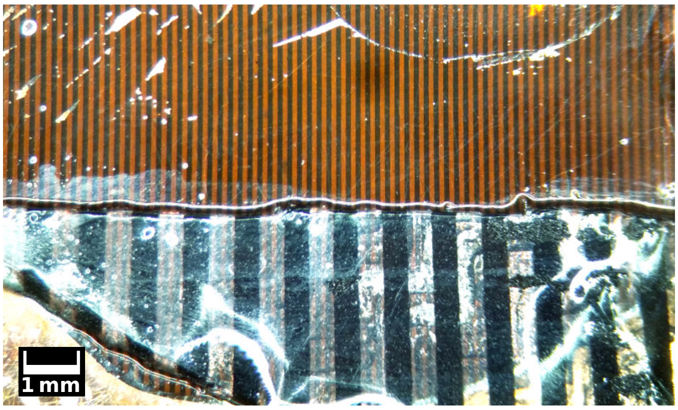

Figure 28: Comparison between the GORE structure (bottom) and our solution (top).