



SCA secure and updatable crypto engines for FPGA SoC bitstream decryption: extended version

Florian Unterstein¹ · Nisha Jacob¹ · Neil Hanley² · Chongyan Gu² · Johann Heyszl¹

Received: 26 April 2020 / Accepted: 7 October 2020 / Published online: 24 December 2020
© The Author(s) 2020

Abstract

FPGA system on chips (SoCs) are ideal computing platforms for edge devices in applications which require high performance through hardware acceleration and updatability due to long operation in the field. A secure update of hardware functionality can in general be achieved by using built-in cryptographic engines and provided secret key storage. However, reported examples have shown that such cryptographic engines may become insecure against side-channel attacks at any later point in time. This leaves already deployed systems vulnerable without any clear mitigation options. To solve this, we propose a comprehensive concept that uses an alternative and side-channel protected cryptographic engine within the FPGA logic instead of the built-in one for the crucial task of bitstream decryption. Remarkably this concept even allows to update the cryptographic engine itself. As proof of concept, we describe an application to the Xilinx Zynq-7020 FPGA SoC in detail. We provide two options for a *leakage resilient* decryption engine which are based on the same primitive, a leakage resilient pseudorandom function (LR-PRF). Depending on a side-channel evaluation of this primitive on the target platform, either a version with additional side-channel countermeasures or a more efficient variant is deployed. The lack of accessible secret key storage poses a significant challenge and requires the use of a *physical unclonable function (PUF)* to generate a device intrinsic secret within the FPGA logic. At the same time this means that manufacturer-provided secret key storage or cryptography is no longer required; only a public key for signature verification of the first stage bootloader and initial static bitstream. We provide empirical results proving the side-channel security of the protected cryptographic engine as well as an evaluation of the PUF quality. The full design and source code is made available to encourage further research in this direction.

Keywords Secure boot · Leakage resilience · PUF · AES · Zynq

1 Introduction

High-performance edge computing is becoming increasingly important in automotive mobility, industrial control and other application domains. Vast amounts of sensor data need to be processed instantly within embedded systems for extraction of relevant information and fast reaction times. This is different to many current applications like speech recognition where data is sent to centralized cloud backends for processing. At the same time it is difficult to pre-determine all required edge functionality for devices which operate in the field for many years. For example, statistical analysis of sensor data through deep neural networks requires high computational capabilities best implemented in hardware, and will likely require fundamental updates during the field lifetime, e.g., to counteract so-called adversarial learning attacks on specific neural network instances or to improve recognition capabilities. FPGA SoCs are considered to be valuable

Florian Unterstein, Nisha Jacob and Neil Hanley have contributed equally to this research.

✉ Florian Unterstein
Florian.Unterstein@aisec.fraunhofer.de

Nisha Jacob
Nisha.Jacob@aisec.fraunhofer.de

Neil Hanley
n.hanley@qub.ac.uk

Chongyan Gu
cgu01@qub.ac.uk

Johann Heyszl
Johann.Heyszl@aisec.fraunhofer.de

¹ Fraunhofer Institute for Applied and Integrated Security AISEC, Garching near Munich, Germany

² Centre for Secure Information Technologies (CSIT), ECIT, Queen's University Belfast, Belfast, Northern Ireland

platforms in this regard. They offer powerful application CPUs integrated with hardware acceleration on an FPGA and update capabilities for improved functionality and/or security of the hardware as well as the software. The security of intellectual property (IP) which is implemented in such FPGA-based hardware against read-out and manipulation is the focus of this contribution.

FPGAs (i.e., SRAM-based ones) are configured with the hardware implementation at every power-up. The respective bitstream must be stored in a chip-external non-volatile memory (NVM). Since adversaries with physical access to the field devices must be expected, a read-out is often realistic and the contained IP prone to reverse-engineering. To counteract this, FPGA devices usually provide bitstream encryption and authentication features using dedicated built-in hardwired cryptographic engines. Unfortunately, it has been shown for several devices from different manufacturers that such cryptographic engines can be attacked using side-channel analysis (SCA) [30,31,38–40] which is widely known to attackers with physical access to devices. Another security issue can arise from hardware design flaws, as is evident by the work of Ender et al. [6]. They uncovered a hardware bug in the Xilinx 6 and 7 series FPGAs that allows them to read out user IP after it is decrypted by the built-in hardware cryptographic engine. As with the SCA vulnerabilities of hardwired engines, this bug cannot be fixed and requires a new revision of the hardware. Due to this, it is highly advisable to have a concept for retrofitting and updating the cryptographic engine which is used for bitstream decryption. This applies even if no successful attack is currently known against the cryptographic engine of a certain device. Hence, using an improved engine implemented within the FPGA logic instead of the built-in hardwired engine is the goal as soon as security issues arise. Unfortunately, with most currently available FPGAs and FPGA SoCs, it is difficult to use alternative cryptographic engines within the FPGA logic for this core functionality of bitstream decryption. While partial FPGA configuration is helpful and mostly available, on-chip key storage for symmetric decryption keys is often not accessible from the FPGA logic or not trusted in the long-term. Furthermore, it has been shown that readout of both eFuses and battery backed RAM (BBRAM), which are commonly used to implement key storage, can be possible with state of the art equipment [22,42]. These restrictions make it difficult to retrofit cryptographic engines with a clear security benefit without inadvertently creating new attack vectors.

In this contribution, we provide a sound concept for integrating an alternative side-channel hardened authenticated encryption with associated data (AEAD) engine into an FPGA SoC for decrypting hardware configuration bitstreams during startup and later operation. This concept allows to securely update the main FPGA functionality and

even allows to update the cryptographic engine itself (i.e., to use a different algorithm later). We describe a proof of concept implementation on a Xilinx Zynq-7020 FPGA SoC¹ and discuss the generalization to other devices. It is important to note that this is the first comprehensive solution allowing to *dismiss all manufacturer-provided secret key storage* options while enabling *side-channel secure field updates of user IP cores* and updates of the decryption engine.

As an alternative cryptographic engine for authenticated decryption of bitstreams, we use one that is side-channel protected by principles from *leakage resilient cryptography*. Protection through leakage resilience has gained significant attention in recent years. Unlike more mainstream protection mechanisms such as masking, these schemes do not require high-quality random numbers which might be hard to generate on embedded devices with limited entropy sources. The secure realization of masking schemes also poses implementation pitfalls [3,13] which we avoid. In this work, we provide two constructions for the leakage resilient AEAD (LR-AEAD) that are both based on the same LR-PRF primitive. The two options differ from each other in terms of their protection level against side-channel attacks and hardware overhead. We also provide side-channel evaluation results of the protected engine on the Xilinx Zynq device using high-precision laboratory equipment as evidence for its side-channel security. Based on the results of the side-channel evaluation of the target platform, either the LR-AEAD variant with additional side-channel protection or the variant that is more efficient in terms of hardware area can be chosen. The more efficient construction is newly added to the extended version of this paper and details of this construction are found in Sect. 8.

For complete independence from manufacturer-provided secret key storage, and since many devices do not allow access to a dedicated key storage facility from the FPGA logic, we use a PUF-based secret key storage in the FPGA logic. Note that in the case of the targeted Xilinx Zynq-7020 device, there is no secret key storage which is accessible from FPGA logic. Hence PUF-based key storage is the only option for alternative bitstream decryption engines. PUFs provide a mechanism for key storage by leveraging the manufacturing differences of an integrated circuit to derive a device intrinsic secret. This eliminates the need for a secure NVM. As evidence for the cryptographic quality of the PUF, we provide evaluation results from 20 Xilinx Zynq-7020 devices.

The concept describes the necessary integration into the FPGA SoC start-up and the precise use of partial reconfiguration (PR). Whenever a core functionality such as bitstream decryption employing secret keys is implemented in the FPGA logic instead of using the hardwired built-in options,

¹ Source code is available at https://github.com/Fraunhofer-AISEC/leakres_updatable_SoC.

this part of the FPGA logic configuration must be authenticated and integrity-protected. Otherwise, an attacker could modify or replace those initial hardware parts to leak secret keys. This makes using built-in bitstream authentication features unavoidable. However, when using public key signatures (e.g., the Xilinx Zynq-7000 series supports RSA signature authentication), this only requires storing and processing public keys on-chip. No secret private keys are stored on-chip, hence, leading to minimal additional attack surface. In our concept, the integrity of the bootloader and initial FPGA configuration containing the LR-AEAD, PUF, and PR handling are verified in this manner.

2 Related work

The general idea behind retrofitting an FPGA SoC with custom cryptographic cores to protect user IP has been previously outlined by Xilinx [33] in a white paper. Their concept includes an AES core together with a PUF and dynamic PR to secure user IP cores. They also suggest to use a manufacturer-provided authentication scheme like RSA to authenticate the static bitstream consisting of the custom cryptographic cores. However, little detail is given and no working implementation of this scheme was provided. Therefore, a large number of technical challenges, such as the PUF error correction, or buffering the authentication of bitstream segments, remained unaddressed. The white paper also lacks an analysis of the overall system security, and the various attack vectors which we later discuss are not considered. An implementation of that scheme was later published by Jacob et al. [16] where they use AES in Galois counter mode (GCM) for decryption and authentication with a twisted bistable ring PUF for the key storage. Their implementation is basic in the way that it provided no comprehensive evaluation of the used PUF and contains no countermeasures against side-channel attacks. Furthermore, the FPGA is updated after the decryption of the partial bitstreams but before its integrity is verified, allowing an adversary to physically damage the device via bitstream tampering. Owen et al. [18] follow a similar approach, but use the PUF to generate a key in a way that is sensitive to all changes within the bitstream to achieve a self-authenticating design. While they therefore do not require the manufacturer-provided authentication, they do however need physical access to the chip for the encryption of images since it can only be performed on-chip using a separate bitstream that contains the encryption core. This makes their proposal unsuitable to provide updates for devices in the field as it would entail that said devices need to be brought back to a secure environment for the update.

In contrast to the work of Owen et al., we use the PUF to embed an external secret key and only perform decryption on-chip thereby allowing off-chip encryption of updates.

Additionally, and extending the work of Jacob et al., we also perform an analysis of the stability and reliability of the PUF on a set of 20 Zynq devices, and implement side-channel countermeasures where required, in order to prevent leakage of secret information.

The work of Kashyap et al. [20] deals with security issues that arise if a partial bitstream is received during the runtime of a device and has to be stored in insecure external memory. They protect the partial bitstream by re-encrypting it with a fresh random key and keeping that key and certain values for integrity checks in internal memory. Their work assumes that the device has securely booted and is already running, thus their scope is a lot narrower than our approach which establishes security from boot until full configuration. Additionally, they also keep the keys in volatile internal memory, so after a reboot the configuration is lost.

3 System overview

In this section, we show how the user IP cores in the FPGA design can be protected using custom cryptographic cores on an FPGA SoC and provide a reference design for the Xilinx Zynq-7020 device. Custom cores are used for secure key storage and decryption instead of the manufacturer-provided built-in key storage and decryption core.

Figure 1 provides a system overview of the FPGA SoC consisting of the FPGA fabric and processing system on the same die. The cryptographic cores, namely the LR-AEAD,² PUF and PR controller (depicted in orange), are part of the initial static configuration of the FPGA. This part of the FPGA configuration is independent of the so-called *user design* which contains the actual design to fulfill the application's purpose. All respective user cores (depicted in green) are loaded using dynamic PR only once they have been successfully decrypted and authenticated using the initial static parts.

Prior to deployment of the device, a user supplied encryption key is linked to the PUF as a one-time operation. A dedicated software routine is loaded onto the device which retrieves the PUF helper data for that key. The PUF internally creates a device intrinsic secret which in combination with that helper data is later used to regenerate the supplied encryption key (this process is illustrated in Fig. 3). Neither the encryption key nor the PUF secret are permanently stored on the device which means that the secret cannot be recovered via offline attacks. After embedding the encryption key, an RSA public key is burned into the eFuses of the device and authenticated boot is enforced. This prevents the execution of

² In the original version of this paper the primitive is denoted as LR-AES; we changed the name to LR-AEAD to make it more generic and reflective of both solutions presented in this extended version.

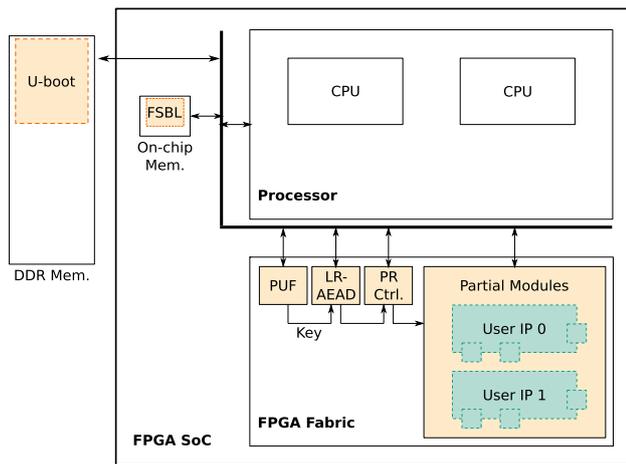


Fig. 1 System overview

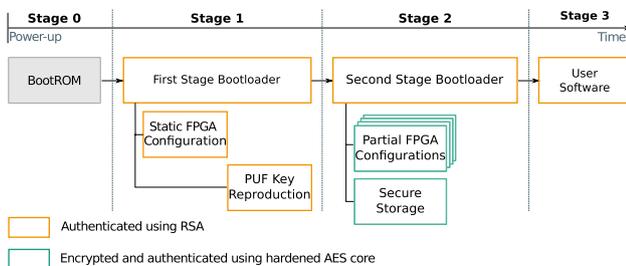


Fig. 2 Boot flow of the hardened boot process

an unauthentic boot loader and, importantly, also stops adversaries from embedding their own key into the PUF after the device is deployed. This does not prevent an authentic user from changing the key at a later point in time if required, once they have physical access to the device.

Figure 2 describes the boot flow of our Xilinx Zynq-7020 implementation once it is deployed. It shows the different stages of the secure boot process starting from power-up until the system has successfully booted and user software is running. After power-up the boot process begins with the hardwired BootROM code in the CPU that is provided by the manufacturer. To establish a chain of trust the first software and logic configuration that is loaded onto the system needs to be protected against manipulation by an adversary. As a result, the BootROM code has to be able to authenticate the first user provided boot code before control is handed over. The BootROM code itself cannot be altered. Hence, all implementations must rely on manufacturer-provided authentication of the first stage boot loader (FSBL), in our case RSA signature verification. However, since the signature check only requires the public part of the RSA key pair, *no secret keys* are stored or processed. Our concept is only based on the assumption that manufacturer-provided storage for the *public key* can be trusted. The FSBL, static bitstream and second stage boot loader (U-boot) are authenticated using

this RSA verification routine. After verification, the FSBL is loaded to the on-chip memory and control is handed to it.

The FSBL is a modifiable bootloader that initializes the system and authenticates and loads the static configuration to the FPGA fabric. We modify the FSBL to transfer the helper data to the PUF module and trigger the regeneration of the encryption key. The helper data is publicly accessible and is stored and authenticated together with the FSBL. Following the successful reproduction of the encryption key, the PUF is locked until the next reset. This prevents any entity from misusing the PUF at a later instance of time in order to attempt to reproduce the key. The encryption key generated by the PUF is directly transferred to the LR-AEAD engine and does not leave the FPGA fabric as can be seen in Fig. 1. The LR-AEAD core is now capable of authenticating and decrypting user hardware IP. The boot process continues with the FSBL authenticating and loading the second stage bootloader, U-boot. As the on-chip memory is too small for U-boot it is usually loaded to external memory.

U-boot is an open-source bootloader commonly used for embedded systems whose functions include system initialization and loading the kernel. In addition to this, we use U-boot to securely load the partial bitstreams. After completing the system initialization, U-boot begins sending encrypted partial bitstreams containing user IP to the LR-AEAD core for tag validation and decryption. Upon successful verification, the plain partial bitstreams are transferred from the LR-AEAD core directly to the PR controller, i.e., they never leave the FPGA fabric and are not transferred on any shared resources of the FPGA SoC. The PR controller dynamically reconfigures the relevant part of the FPGA without interrupting other regions and services. Depending on the application, those user cores may for example contain secret data for higher software layers. The last step in the boot process is for U-boot to load the software stack on the CPU. Currently, the software is only authenticated and not encrypted; this is explored in greater detail in Sect. 9.

This boot sequence comprises the same boot stages as the standard flow supported by Xilinx. The difference is that we leverage PR and how the partial bitstreams are decrypted and validated, namely by custom cores instead of vendor-provided ones. Looking back at Fig. 2, in case of the standard boot flow, all user cores would typically be loaded during Stage 1 following their validation using the built-in cryptographic cores. In contrast, we only load the static bitstream consisting of the custom cryptographic cores during this stage, all user cores are loaded during Stage 2 following their successful verification using the custom cores.

If over the lifetime of the device an update of a user core is necessary, the new version can be encrypted off-chip using the encryption key that was previously embedded during the enrollment process. It can then be sent to the device as remote update without requiring physical access. Remark-

ably, the LR-AEAD core for actual bitstream decryption can also be remotely updated by updating the static bitstream. When changes to the core are made, a new version of the entire static bitstream has to be generated, signed with the RSA private key and then transferred to the device to replace the old one. However, updating the PUF requires a secure connection or physical access. Due to its nature, changes to the PUF will inevitably change the intrinsic PUF secret and will prevent the recovery of the encryption key with existing helper data. In that case, the encryption key enrollment has to be repeated as it was done for a fresh device.

4 Building blocks

In this section, we describe the functionality of the hardware building blocks that make up the static FPGA configuration of our system: the PUF that generates a device intrinsic secret, the LR-AEAD which protects confidentiality and integrity of user cores and the PR controller which loads the user cores using dynamic PR.

4.1 Generating a device intrinsic secret

A PUF is a security primitive that utilizes manufacturing process variations to generate a unique digital fingerprint intrinsic to a physical piece of hardware [7]. As this natural variation between otherwise identical silicon devices is outside the control of even the manufacturer, PUFs are inherently difficult to clone. Constructions can be broadly split into two categories; challenge-response type PUFs which produce a device unique response for a given input challenge (often referred to as ‘strong’ PUFs), and identity-generator type PUFs which produce few or just a single response for the device (often referred to as ‘weak’ PUFs)³. While these constructions and how they are implemented provide different trade-offs for a designer to explore, for this work a method to generate a secret key is required. Hence we focus on identity-generator type PUFs where no input challenge is required and modeling type attacks, which have proven to be effective against many challenge-response constructions [36], are not in scope.

A number of publications suggest using PUFs for the generation of secret keys, for example the works in [1,24]. In order to be of use for security applications, a PUF must fulfill certain properties with respect to reliability, entropy and uniqueness. While the evaluation of these properties varies from work to work, they generally include Hamming distance measurements between responses from the same device (indicating the reliability) and between different devices

³ Note the terms strong and weak here refer to the size of the response space rather than the security of the PUF constructions.

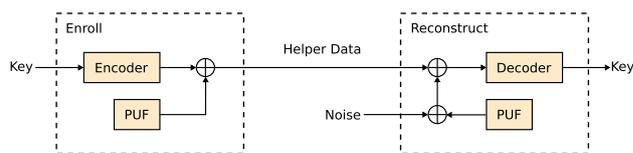


Fig. 3 Fuzzy commitment scheme

(indicating the uniqueness). These broad measurement metrics can hide subtle biases in the responses however, so care must be taken when performing the analysis. Recent work has looked to formalize this measurement analysis through the ISO standardization process [4].

Rather than using a PUF circuit to directly generate a device dependent key, a *fuzzy commitment* scheme is used where the PUF is used to mask the user-generated secret key [19]. A fuzzy commitment scheme is required as the PUF output is dependent on manufacturing variations hence there is an error probability when reproducing the PUF response, which empirically has been shown to vary across temperature and voltage variations. The fuzzy commitment scheme, as shown in Fig. 3, uses an error correcting code (encoder) to introduce redundancy to the key prior to masking with the PUF output. This expanded and masked key can be publicly stored as helper data to regenerate the secret key as required at a later stage. However, should the PUF response not have sufficient entropy, information leakage can occur through this helper data. Where a biased PUF is used to mask the helper data, key recovery could be possible through this publicly available helper data. Hence in practice, a de-biasing stage should be used in conjunction with an entropy extractor to securely generate a key. There are a number of approaches to implement such a stage, such as index-based syndrome (IBS) coding [46], von Neumann corrector [25] or wiretap coset codes [14] for example. In this work a de-biasing stage is not considered as the empirical bias of the PUF, as calculated over 20 Xilinx Zynq-7020 devices at room temperature, is close to ideal as shown in Sect. 6.1.

The PUF architecture used in this work is based on that in [10] which creates a cross coupled feedback loop contained within a single slice of an FPGA to generate a single PUF bit. For the error-correction, a (23, 12, 7) Golay linear block code is used, where every 12 bits of the key are encoded to a 23-bit codeword, with a Hamming distance of at least 7 between any codeword pair. This allows up to three errors per codeword to be corrected. Hence the 256-bit key is expanded to a 498-bit codeword and combined with the 498-bit PUF output bits to generate the helper data⁴. Once generated, this public helper data is then incorporated as part of the FSBL to enable key generation as outlined in Sect. 3. A Golay encoder

⁴ The key is zero padded to a multiple of 12, however these additional mask bits are not required in the final codeword.

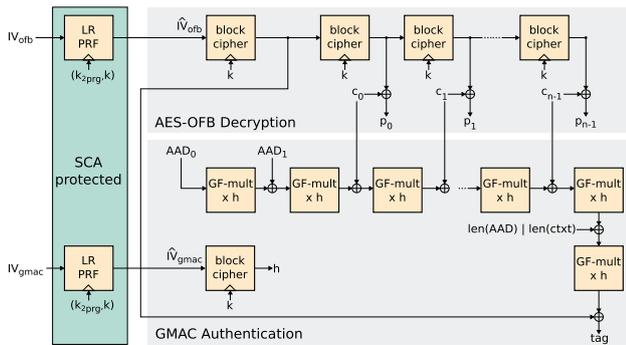


Fig. 4 Side-channel secure authenticated decryption

can be implemented very efficiently in hardware using a shift register, with the decoder block utilizing the same encoder block combined with an additional 12-bit look-up table of depth 2048. Both encoder and decoder architectures run in constant time to prevent unintended side channel leakage from the core.

4.2 Protecting confidentiality and integrity of HW IP

To protect integrity and confidentiality, hardware IP is usually encrypted and protected by a message authentication code (MAC). This mitigates offline attacks, but since edge nodes might be deployed in hostile environments we also need to protect the decryption against side-channel attacks.

Figure 4 shows a dataflow diagram of the authenticated decryption core that consists of three components: a side-channel secure stage with a LR-PRF, decryption using AES in output-feedback (OFB) mode and authentication using the Galois MAC (GMAC). Leakage resilient cryptography is an algorithmic countermeasure against SCA attacks. It aims to bound the leakage per execution such that an attacker cannot accumulate information about the processed secret. We use the LR-PRF proposed by Unterstein et al. [45] to derive a secret pseudorandom state from public initialization vectors (IVs). Specifically, that means processing two public IVs, IV_{ofb} and IV_{gmac} , with the LR-PRF to get two secret IVs, denoted by \hat{IV}_{ofb} and \hat{IV}_{gmac} , for the subsequent stream cipher and MAC. This LR-PRF is similar to the one presented by Medwed et al. [28], but it provides security against sophisticated attacks using high precision EM measurements at the cost of a higher latency. In our use case this is acceptable because the computational overhead occurs only once during the boot process. Internally the LR-PRF uses AES-128 encryption, so for the stream cipher we re-use the same AES core in OFB mode which makes the area overhead almost negligible.

For common SCA attacks the precondition is that the attacker can guess some internal value of the algorithm that depends on the secret and an input (e.g., the plaintext)

that is known to the attacker. The general idea behind our construction is that only the initial LR-PRF stage needs to be side-channel secure (through means of leakage resilient cryptography), because behind it, no public inputs are processed by the block cipher. This gives no surface to mount an SCA attack on the unprotected stages as there are no known inputs and thus guessing internal values is not possible. A detailed security analysis and a practical evaluation of our target platform is presented in Sect. 6.2.

For message authentication we chose GMAC [26] because it can be efficiently implemented in hardware since it only requires an additional Galois field multiplier. Differing from the specification, where the MAC key h is derived by encrypting a plaintext with all zeros, we derive the MAC key by processing another public IV_{gmac} with the LR-PRF and then encrypt the resulting \hat{IV}_{gmac} to get h . We divert from the specification here because, as previously explained, we need to prevent an unsecured encryption where the inputs are known to the attacker. This is achieved by processing the public input with the side-channel secured PRF first.

A general problem when decrypting and authenticating data in one pass is that the authentication tag can only be checked after processing all the data. But sending unauthenticated data to the PR controller, in our case the Xilinx proprietary internal configuration access port (ICAP), can be dangerous. Configuring the FPGA with unauthentic data could damage the FPGA due to short-circuits caused by false or malicious configurations as stated in [31]. Hence we need to buffer the decrypted data until it is verified before we pass it on. This buffering has to be done in on-chip memory (OCM), otherwise it would be prone to manipulation, for example by probing of the memory bus. Therefore, we implemented a first in, first out (FIFO) buffer immediately prior to the PR controller that releases data only after its tag has been verified. To keep the allocated block RAM (BRAM) for the FIFO small, we split the bitstream into segments which are decrypted and authenticated individually. While the FIFO is sending authenticated data to the PR controller, decryption of the next segment is concurrently being performed, thus reducing the latency. To prevent IV re-use, the IV must be unique for every segment. For this purpose, IV_{ofb} is split up into a 96-bit random value and a 32-bit counter value that is incremented with every segment. That also means that at the beginning of every segment, the LR-PRF is evaluated to generate a fresh \hat{IV}_{ofb} . Updating IV_{gmac} in between segments is not necessary because it is only used to generate the GMAC key h which we keep constant for the entire bitstream (note that in the original GMAC scheme h is always derived by encrypting an all zero plaintext).

U-boot only continues the boot process if the verification of all the segments of the partial bitstream succeeds. If the verification of any of the segment fails, U-boot aborts the

Table 1 Resource utilization

| Module | Slices | | BRAM RAM36 |
|---------------------------------|--------|-----------|---------------|
| | LUTs | Registers | |
| LR-AEAD (LR-PRF, AES-OFB, GMAC) | 4568 | 2810 | 0 |
| PicoPUF, Fuzzy Commitment | 3917 | 4179 | 1 |
| PR Controller | 50 | 53 | 0 |
| ICAP FIFO buffer | 99 | 124 | 1 |
| Overall, incl. interconnects | 9274 | 7911 | 2 |

boot process and system goes into a secure lockdown mode. A power cycle is necessary to remove the system from this state.

4.3 Partial reconfiguration of user IP cores

The PR controller receives the authenticated and decrypted partial bitstream from the corresponding LR-AEAD core and then passes them to the Xilinx ICAP interface. The ICAP is a proprietary interface from Xilinx typically used for configuration readback and reconfiguration of the FPGA. The ICAP has direct access to the configuration memory through the configuration registers, hence can be used to dynamically reconfigure the FPGA. In this work, the ICAP only operates in the device reconfiguration mode and no configuration data can be read back through the ICAP.

5 Implementation

Here, we describe the prototype implementation for the Xilinx Zynq-7020 FPGA SoC. The resource utilization of the various building blocks is provided in Table 1. Overall, we only use around 22 % of the available slices and 1.5 % of the available BRAM of this device, which is among the smaller, low-cost devices of the Xilinx Zynq-7000 product range. Hence, 78 % of the slices of this lightweight device are still available to implement for user IP cores.

All the hardware building blocks are configured and managed over the AXI interface by software drivers that are patched to a standard U-boot and FSBL. The helper data for the PUF is generated once and then hard coded into the FSBL. As the BRAM modules on the Xilinx Zynq-7000 series are each 4.5 KB large, we split the partial bitstreams into segments of 4 KB with an additional 16 bytes for the tag appended to each segment. This configuration leads to a storage overhead of less than 0.4 percent.

6 Security evaluation

This section presents the crucial elements of the security evaluation. We use an attacker model that allows physical access

to the device and full control over all inputs. The attacker is capable of mounting passive side channel attacks by taking measurements of the device. We will in our analysis consider the most powerful method, namely high-precision EM measurements directly over the decapsulated die. For our DUT we had to remove the packaging which technically makes it an invasive attack, but there are derivatives within this product family that are shipped as 'naked' flip chip. In that case the die is bonded upside down and the backside silicon is directly accessible for measurements, making it a non-invasive attack. In this work we do not consider active or invasive attacks like glitching, laser fault injection or focused ion beam (FIB) attacks.

6.1 Evaluation of the PUF

In this work, an identity type PUF design similar to that proposed in [10] is employed to regenerate the secret key. It comprises of 498 elementary 1-bit PUF cells, which are used to mask the output of the Golay linear block code used to encode the 256-bit user-generated key as shown in Fig. 3. Each PUF cell is designed to fit compactly in one FPGA slice as shown in Fig. 5. Figure 5a shows the architecture of each PUF cell comprising of four logic and three register components. The physical layout of a single PUF cell is shown in Fig. 5b with the registers and logic (implemented in the LUTs) depicted in blue, located in a single slice, and the timing-critical routing paths highlighted in black. The routing is fixed with the use of scripts as part of the Vivado design flow to ensure that all cells have identical routing paths, which have been selected to maximize the entropy while retaining sufficient reliability to allow minimal error correction of subsequent evaluations. As the register components all have the same clock source, they can be placed in the same slice. Therefore, this enables the timing-critical routing paths between PUF components, which largely determines the response value, to be placed within the local interconnect rather than the general purpose interconnect.

To assess the suitability of the PUF architecture for our side-channel protected secure boot design, a 256-bit output of the core PUF module was generated for testing prior to error

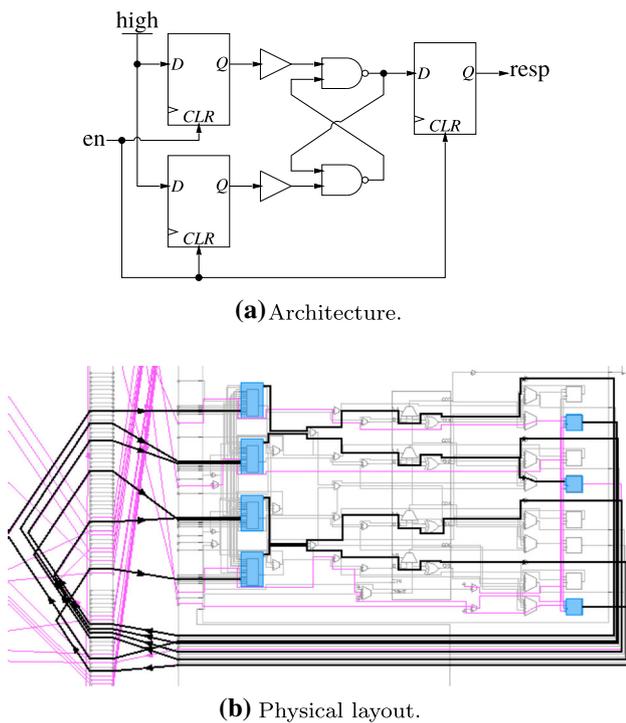


Fig. 5 PicoPUF design

correction being applied⁵. This output was regenerated 1001 times at room temperature for each of 20 Xilinx Zynq-7000 devices.

Analyzing the reliability of each bit, it was found that $\approx 80\%$ of bits on each device returned a stable value across *all* evaluations. While temperature and voltage variations will introduce additional noise to the PUF responses, empirical evidence indicates that constraining the design to a single slice helps minimize these effects [10]. The expected fractional Hamming distance (i.e., the bit error rate (BER)) between a sample evaluation and the reference response was empirically estimated to be 0.0217, or 5.56 bits of the 256 bit response. It was shown in [11], that the probability of a string of n bits having more than t errors is given by:

$$P_{total} = 1 - \sum_{i=0}^t \binom{n}{i} p_b^i (1 - p_b)^{n-i} \quad (1)$$

Given that the Golay encoder can correct up to $t = 3$ errors in an $n = 23$ bit codeword, this gives an expected $P_{total} = 1.4 \times 10^{-3}$. While not implemented here, using a simple (3, 1, 3) repetition code concatenated with the Golay encoder brings P_{total} down to 3.26×10^{-8} , trading off additional hardware requirements for the increased reliability. Alternatively, multiple PUF readouts could be taken with a

majority vote applied to determine the PUF response. While the bits are no longer independent and identically distributed (IID), this should still have a similar error rate to the equivalent sized repetition code. The hardware tradeoff here is an additional counter per PUF cell along with some minimal extra latency, removing the need for the additional PUF cells required when implementing a repetition code. Note that for smaller repetition code/majority voting implementations (e.g., 3 or 5), the area difference will be minimal.

The bias of the PUF response was close to ideal, with the average value of a bit expected to be 0.498, when averaged over all responses and devices. The expected fractional Hamming distance between the outputs of two different devices was empirically estimated to be 0.497. Given these results, it is expected that the helper data generated by different devices is sufficiently independent such that an adversary cannot learn anything useful about the key given access to the helper data from multiple devices.

SCA attacks on PUF designs can be categorized into attacks on the core PUF instance, and attacks on the post-processing. While attacks such as directly reading out PUF bits using a FIB are outside the scope of the attacker model, there has been work directly attacking the generation of the PUF output. For example, delay-based arbiter designs are attacked using power analysis in [2], while recovering frequencies from ring oscillator (RO) PUF designs is investigated in [29]. However, as the PUF design used in this work evaluates in a single clock cycle, it is not expected to be susceptible to attacks on its core bit generation.

Of greater threat are attacks on the post-processing stage [29,41]. Typically such attacks are differential attacks and require the attacker to manipulate the helper data. In our proposal, the helper data is authenticated before use and thus protected from tampering. This greatly reduces the number of *different* observable traces that an attacker can obtain, with differences only generated due to PUF response errors. The Golay error correction block consists of linear operations, and a single, constant time, table look-up indirectly based on the PUF response error. Linear operations have been shown to require a large number of side-channel traces to obtain information [34], hence are not considered practical in this scenario. While many attacks against block ciphers target look up tables, this is as their non-linearity makes them a suitable attack point [34]. In the implementation here, the look up table is not non-linear, and is used to both speed up and ensure constant time syndrome decoding by mapping the decoding error into memory. These errors are sparse vectors and, similar to the linear operation stage, the reduced number of different observable traces should prevent any power or EM-based SCA. Under these restrictions, simple power analysis (SPA) attacks using only a single or few traces, are likely infeasible.

⁵ The 256-bit output is for testing purposes only, the full design requires 498 PUF cells.

6.2 Leakage resilient authenticated decryption

We designed our leakage resilient authenticated decryption scheme following the principle described in [32]. As shown in Fig. 4, we rely on an initial “leak-free” stage using a LR-PRF that is protected against SCA. It establishes a pseudo-random intermediate state (\hat{V}_{ofb} and \hat{V}_{gmac}) that is unknown to the attacker and allows us to use unprotected decryption and authentication for the actual workload. Internally, the LR-PRF consists of a so-called GGM tree PRF [9] and optionally one or more length-doubling pseudorandom generators (2-PRGs) as described in [45]. Both the PRF and the (one or more) 2-PRG stages are implemented using the same hardware AES core that is also used in the stream cipher module. This means we only need to instantiate one AES core that is then time-shared between the different modules. The number of 2-PRG stages is a design choice that allows to compensate entropy loss after successful side-channel attacks on the very first AES execution which inevitably has to operate on public inputs. The required number of stages depends on the leakage behavior of the device and can be determined through laboratory analysis. Each stage adds a fresh 128-bit key to increase the remaining entropy within the internal secret state. We opted for one 2-PRG stage and thus require two 128-bit keys, k_{2prg} and k . k_{2prg} is used only in the initial 2-PRG stage, k in the rest of the construction.

The side-channel security of the LR-PRF is based on two principles: limited data complexity (i.e., the number of operations with the same key, but different inputs that an attacker can observe) and algorithmic noise from parallel S-boxes. As pointed out in [45], a laboratory analysis is mandatory to assess the security level. We implemented the LR-PRF using an AES core with parallel S-boxes on the Xilinx Zynq-7020 device and placed the AES as dense as possible as shown in the floorplan in Fig. 6. Following the method described in [45] we conducted a localized EM analysis using a near-field EM probe placed on the decapsulated die. We first identified the locations of the S-boxes through a grid scan and then ran template attacks on them using measurements from those locations. For each S-box we collected 400,000 traces for the profiling and 100,000 traces for the attack at its respective location. More traces typically lead to better results but the number of traces that can be used is limited by the acquisition time. However, during our experiments we found this number to be sufficient and that an increase in either the profiling or the attack set did not demonstrably improve the attack.

Usually, attacks with limited data complexity do not successfully recover all key bytes directly. Instead, some brute force effort is required to combine the most probable candidates for each key byte into full keys and then test each key candidate for its correctness. An evaluator that knows the correct key can use a key rank estimation algorithm [8] to determine the brute force effort even in cases where an

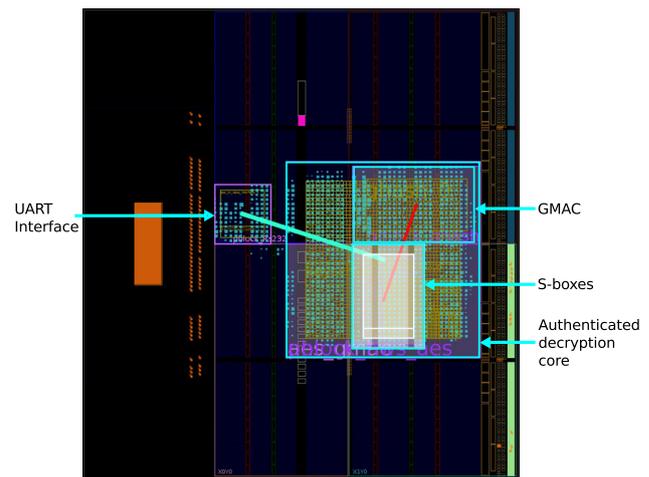


Fig. 6 Placement of the AES core for the SCA

actual enumeration of all candidates up to the correct key is not computationally feasible. The security level is then given as the brute force effort after running the template attack, i.e., the number of combined keys an attacker has to try until reaching the correct key. The natural attack vector in our case is the initial execution of the first 2-PRG stage, as this is the only stage where the AES is encrypting public plaintexts; the data complexity is 2. We found that the remaining security level after an attack on this 2-PRG stage is still 2^{120} , or 120 bits. This is significantly higher than the results reported on the Xilinx Spartan 6 platform in [45] where the security level was only 2^{48} and we attribute this difference to the smaller feature size (28nm compared to 45nm) and the different placement strategy. As shown in [44], the countermeasure is more effective if the placement of the parallel S-boxes is denser so that individual signals cannot easily be resolved by the EM probes. A smaller feature size allows for higher integration density (and consequentially shorter connections on the metal layers which are the main source of electromagnetic emanations) which should hinder localized EM attacks. Note that these results were achieved with a standalone design and not the entire system. We expect that the increased noise generated by the full design will only make attacks harder.

This means that the LR-PRF provides sufficient remaining guessing entropy after an attack even if we use the GGM tree directly without the additional 2-PRG stage and key upstream⁶. Nevertheless, we opted to keep the 2-PRG and second key to increase the security margin in case of more sophisticated, as yet unknown, attacks. The hardware overhead is minimal since we re-use the existing AES engine and

⁶ The LR-PRF is then identical to the construction in [27] with data complexity 2.

Table 2 Overview of system security

| Threat vector | Mitigation |
|--|---|
| Stage 0. BootROM | |
| Load malicious FSBL | Integrity check of FSBL using RSA |
| Stage 1. FSBL and static FPGA configuration | |
| Learning the PUF | Key reproduction triggered by FSBL, following which PUF is locked |
| Recovery of PUF secret | Key enrollment in the field is blocked |
| Readout encryption key | Key is directly transferred to LR-AEAD |
| Stage 2. U-boot | |
| Loading unauthenticated configuration | FIFO is used to buffer configuration |

thus only pay for this added security with latency and “key storage” in the form of a large PUF.

After the LR-PRF stage, we use AES in OFB mode for the actual decryption of the bitstream. We opted for the OFB mode, instead of e.g., GCM mode, because the plain- and ciphertext is not directly an input to the block cipher. As long as an attacker does not know both plain- and ciphertext, the in- and outputs to the block cipher remain secret and provide no surface for side channel attacks. Even partial knowledge, as is the case in counter modes where the initial counter value may be unknown but the increment of the counter is known, can be sufficient to mount an attack [17]. If on any platform portions of the plaintext are known, e.g., because they are all zeros, we propose not encrypting those parts and instead adding them to the additional authenticated data (AAD). An open question remains considering the effect of plaintexts that are unknown, but not uniformly random because they, e.g., consist of opcodes or addresses that do not utilize the entire value space. We are not aware of any published attacks that exploit such a scenario and imagine such an attack to be hard, nevertheless it remains as an interesting topic for future research.

6.3 Overall system security

The security evaluation of the building blocks, namely PUF and leakage resilient authenticated decryption cores, were presented in Sects. 6.1 and 6.2 respectively. However, the integration of the individual blocks can give rise to new attack vectors. Now we review the overall system security by recalling the boot process and listing the different threat vectors post-integration together with the implemented mitigation techniques. A summary of the findings for each boot stage is listed in Table 2.

Load malicious FSBL or U-boot: An adversary may attempt to bypass the implemented security mechanism by replacing the boot images with malicious ones. This is prevented by using the manufacturer-provided RSA signature verification. After power-up, the BootROM verifies the sig-

nature of the FSBL, only after successful verification of the signature is the control handed off to the FSBL. This check is enabled by burning an eFuse and cannot be disabled. Subsequently, the FSBL checks the signature of the U-boot. An adversary trying to forge any of those signatures would require the private key which is not stored on the device.

Learning the PUF: In our proof of concept implementation we use an identity-generator type PUF hence modeling attacks are out of scope. As well as ensuring that the PUF output is only ever combined with the helper data in order to generate the secret key, we also built in a mechanism to limit the access and interface to the PUF. Following the loading of the static bitstream by the FSBL, the FSBL is allowed to trigger a PUF key reproduction only once. After the successful reproduction of the key, the PUF is locked by the hardware and is not accessible until the next power-up.

Recovery of the PUF secret: As can be seen in Fig. 3, the helper data that is generated during key enrollment is an XOR combination of the encoded encryption key and the PUF secret. If an attacker could enroll his own keys, then it is trivial to calculate the PUF secret from the retrieved helper data. This is prevented by only allowing authorized FSBL code to be executed on the device. No authenticated code for key enrollment exists in the deployed device however (the respective code is executed before setting an authenticated code only configuration). The locking of the PUF after key reproduction also disables the ability to enroll keys during runtime.

Readout of the encryption key: An attacker may attempt to corrupt the bitstream so as to include additional lines to readout the key. This however is prevented by verifying the integrity (at the same time as its authenticity) of the bitstream before it is loaded. Thus, what remains is to attempt to read out the key at run-time, e.g., via existing shared interfaces. To protect the key from being read by an unauthorized entity, the reproduced key is directly transferred to the LR-AEAD core. The key never leaves the FPGA fabric and is not transferred over any shared resources of the FPGA SoC, this ensures that no other entity has access to the key.

Load unauthenticated data to FPGA: Authenticated encryption is used to protect confidentiality and integrity of the partial bitstreams where the decryption and authentication are performed in parallel. As the authentication tag can only be checked at the end, an adversary can send unauthenticated data and it will be decrypted before the invalid tag is noticed. If the switch boxes and LUTs of the FPGA are falsely configured, it could result in short-circuits in the fabric and thereby destroying parts of the device. To prevent decrypted data from being directly transferred to the PR controller before the authenticity can be verified, the bitstream is divided into segments. Each segment has its own verification tag and is buffered after decryption until it has been successfully authenticated. Only after the successful verification of a segment, are the buffered contents transferred to the PR controller.

7 Generalization

While we target a Xilinx Zynq-7020 FPGA SoC for our proof-of-concept, the method is agnostic to a specific FPGA manufacturer/family as long as it provides certain common functionalities. We now discuss the portability of the presented design to other platforms and the additional steps that need to be taken. The two key requirements to port this design are that the target provides:

1. Authentication and integrity checking of the static bitstream using public key cryptography
2. Partial reconfiguration

In the initial boot step, an authenticity and integrity check of the static bitstream containing the custom cryptographic cores is necessary. We used signature verification with public key cryptography on the Xilinx Zynq-7020 device to achieve this. Public key cryptography for this purpose is a widely adopted feature and is present in the majority of FPGA SoCs that are currently available. The benefit of using public key cryptography is that no additional *secret* key material is necessary and the operation does not need to be protected against key-recovery attacks.

PR is necessary to reconfigure parts of the FPGA fabric at run-time with user cores that were decrypted by our protected engine. This feature is supported by the leading manufacturers of FPGA SoCs such as Xilinx and Intel (formerly Altera). Thus this design can be ported to these devices with low effort.

As the Xilinx Zynq-7000 series devices do not provide any user key storage (only 32 bits of general purpose eFuses are available), we opted to use a PUF for key storage. Newer devices, however, often provide user accessible secure key storage. Some devices like the Xilinx Zynq Ultrascale+ and

Stratix 10 from Intel even include hard-core PUFs and allow secure boot using these PUFs. In such cases the custom PUF can be omitted if the key storage is trusted to be secure. If a custom PUF is used on a different platform, it needs to be re-evaluated to ensure that the key has sufficient entropy. Additionally, when porting the design to a different technology the leakage behavior of the decryption core can change and the side-channel security should also be re-evaluated as shown in [45]. Considering the large security margin that we found on the Xilinx Zynq-7020 device, we do not expect any issues on other comparable or newer technologies as a smaller feature size makes localized attacks harder.

8 Alternative leakage resilient decryption scheme with a smaller hardware footprint

In Sect. 6.2, we propose a LR-AEAD scheme that uses a LR-PRF and AES in OFB mode for decryption and GMAC for authentication. This choice was motivated by results of Unterstein et al. [45] that demonstrate that on FPGAs, it is possible to recover the key of an AES implementation in a side channel attack even if only two different operations can be observed (the case of data complexity 2). However, the remaining security level after such attacks is dependent on the layout and manufacturing process of the device, where smaller feature sizes usually make localized EM attacks harder. In the security evaluation on the Xilinx Zynq-7020 device we found, that our AES implementation resists localized EM attacks in case of data complexity 2 and retains a high security level of 120-bit. As was already mentioned in Sect. 6.2, this makes the extra 2-PRG stage obsolete on our target platform, but this also allows us to deploy an alternative construction that is easier to implement and requires less hardware resources.

8.1 Architecture

In this second implementation option, we drop the 2-PRG stage inside the LR-PRF, therefore we now require only a 128-bit key for the LR-PRF. We also replace the stream cipher part, i.e., the AES in OFB mode, and the GMAC authentication. Instead, we implement a LR-AEAD scheme based on the results of Krämer and Struck [21] who propose that a LR-AEAD scheme can be built using a LR-PRF, a leakage resilient PRG (LR-PRG) and a hash function. In reference to the introduction of this scheme in [5] we refer to it as $FGHF'$ -based LR-AEAD and to our original scheme as AES-OFB-based LR-AEAD. Figure 7 shows the hardware/software partitioning and the building blocks for the proposed $FGHF'$ -based LR-AEAD scheme. The left side describes the software implementation in pseudo code and makes references to the hardware on the right whenever

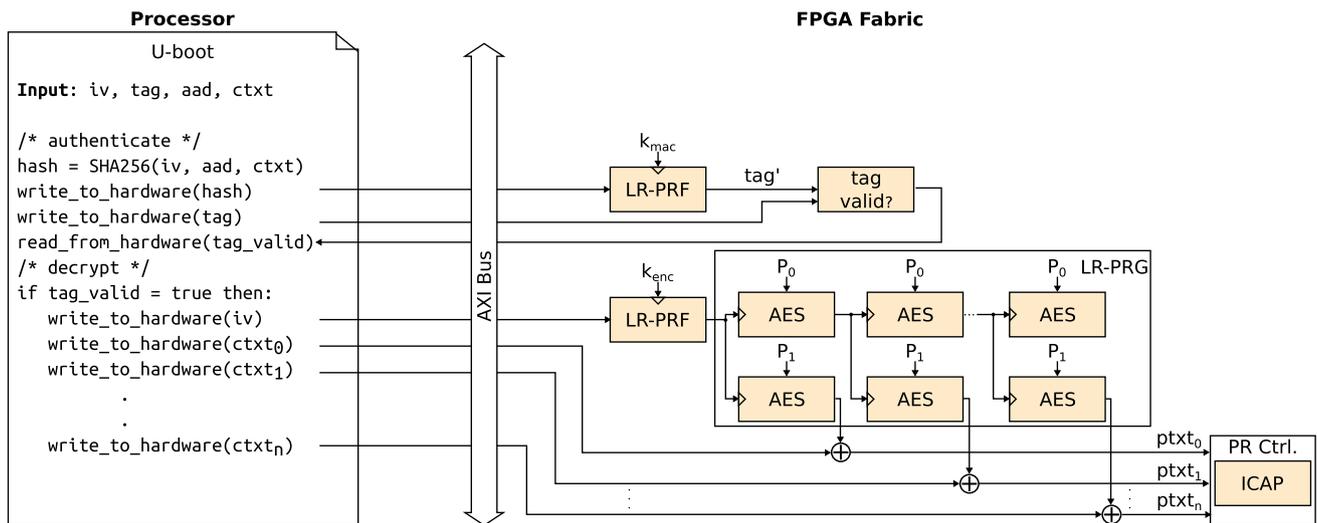


Fig. 7 Alternative side-channel hardened authenticated decryption

data is transferred between the two partitions. The right side describes the *dataflow* in the hardware implementation in a way similar to Fig. 4. Note that there is still only one hardware instance of the LR-PRF and that the (one) AES core used in the LR-PRG part is shared with the LR-PRF. The construction uses a single IV and two 128-bit keys, k_{mac} and k_{enc} .

Interestingly, the hash function in this construction does not need to be protected against side-channel attacks as it only processes public inputs and no secrets. Therefore, the authentication part of the scheme can be split between hardware and software. First, a hash function, here SHA-256, calculates the hash of the IV, AAD and ciphertext. This is implemented in software as part of U-boot. The hash is truncated to 128 bits and processed by the LR-PRF with key k_{mac} to calculate the tag. For this side-channel critical operation, the evaluation of the LR-PRF, we still rely on a protected hardware engine. The tag is now calculated over the entire bitstream at once and checked before decryption is started instead of calculating tags for multiple segments of the bitstream separately. This is different to our previous construction, where authentication and decryption are calculated in parallel hardware and require a hardware FIFO to buffer the plaintext until the tag is verified. This was necessary to avoid the performance penalty of a two-pass scheme that has to transfer the entire bitstream over the bus twice. In this new scheme, only the 128-bit hash is sent to the hardware during authentication. As the tag of the entire bitstream is now verified before decryption, the FIFO in front of the ICAP is not necessary and the decrypted bitstream can be directly configured.

Decryption of the bitstream is still implemented in hardware. If the tag is correct, the IV is processed by the LR-PRF using a second key, k_{enc} . The result is used as key for

a LR-PRG which produces the key stream to decrypt the ciphertext. The LR-PRG is based on AES and consists of a concatenation of 2-PRG stages. Within each stage, an all-zero plaintext p_0 and an all-one plaintext p_1 are encrypted with the same key, and the results are used as the key for the next iteration and as part of the key stream, respectively.

8.2 Side-channel security

In Sect. 6.2 we give results of a side-channel attack with data complexity 2 on a 2-PRG, specifically results of template attacks on an AES core which encrypts two distinct plaintexts. For the new construction we identify two attack vectors which we show to be equivalent to the case in Sect. 6.2:

Side-channel attack on the LR-PRF: Keys k_{mac} or k_{enc} can be targeted in an attack on the LR-PRF. Since we are not using the 2-PRG stage in front of the GGM tree PRF, the relevant attack vector is the first iteration of the GGM tree. In the context of an attack, this iteration behaves identically to the 2-PRG stage: it consists of an AES encryption that uses either all-zeros or all-ones as plaintext input.

Side-channel attack on the LR-PRG: The LR-PRG consists of multiple 2-PRG stages and changes the key between the stages. It is initialized with a key that is derived from the IV of the bitstream. Hence, an attack can only recover the ephemeral key to decrypt one specific bitstream, and not one of the long-term keys k_{mac} or k_{enc} . An attack on the LR-PRG can be mounted targeting any 2-PRG iteration, i.e., any intermediate key. However, in all cases the attack is limited to data complexity 2.

Both attack vectors are identical to the case investigated in Sect. 6.2. Therefore, we can apply the results of the side-

channel evaluation of the 2-PRG to this construction. As a result, the remaining entropy after a side-channel attack on k_{mac} , k_{enc} and the ephemeral LR-PRG keys is 120-bits per key. Since there is no meaningful way to combine the results of attacks on the different keys, the security level of the entire construction is thus 120-bits.

8.3 Resource utilization and performance

By using this scheme and moving part of the authentication to software, it is possible to significantly reduce the hardware overhead. The $FGHF'$ -based LR-AEAD hardware requires 3137 LUTs and 1779 slices, compared to a total of 4667 LUTs and 2934 registers for the AES-OFB-based LR-AEAD and FIFO combined. This is a reduction of 33% and 39%, respectively. The LR-AEAD hardware also does not now require any BRAM blocks as no buffering is required (previously the FIFO used one RAM36 block). In addition, the size of the encrypted bitstream is slightly reduced because only one tag for the entire bitstream is stored instead of one tag per segment. The requirement in order to deploy this variant is that the AES resists side-channel attacks with data complexity 2. On platforms where this is not the case, such as the example in [45], the AES-OFB based LR-AEAD must be used.

As the performance, i.e., the latency when decrypting one partial bitstream, is dependent on the concrete implementation parameters, we only provide an informal discussion. For the AES-OFB based LR-AEAD, decryption of one partial bitstream requires two initial LR-PRF executions and one initial AES encryption. The subsequent decryption of each 128-bit block of data requires one AES encryption. Since the decrypted bitstream has to be buffered in a FIFO, it is potentially split up into multiple segments if it exceeds the size of the buffer. For each new segment, an additional evaluation of the LR-PRF and an AES encryption is required. The duration of one LR-PRF evaluation is significant and depends on the configuration, as explained in [45]. Depending on the hardware platform, it can consist of up to 130 AES encryptions. How many LR-PRF evaluations the AES-OFB based LR-AEAD actually requires depends on the size of the FIFO and length of the partial bitstream.

For the $FGHF'$ -based LR-AEAD, initially the hash has to be calculated in software and the LR-PRF is evaluated twice. Then for each 128-bit block of data two AES encryptions are necessary. Hence, the $FGHF'$ -based LR-AEAD requires double the AES encryptions per 128-bit data block, but potentially less LR-PRF evaluations. Nevertheless, we expect that reconfiguration on such systems is a task that is most probably performed only once per boot. Therefore, we suspect that designers will aim for lower hardware footprint instead of lower latency.

9 Toward software encryption and runtime security

Our proposal is currently limited to bitstream decryption. User software is authenticated but not encrypted. For many applications this is sufficient, since rarely is the entire software stack confidential. If full confidentiality is desired for software, this usually goes hand in hand with runtime integrity protection and encryption of chip external RAM, which are hard problems on their own. However, if our AES-OFB based core is to be used for software decryption, the OFB mode of operation might not be the best choice and could lead to new attack vectors. The side-channel security of the decryption core relies on the fact that an attacker does not know the inputs to the underlying block cipher. In OFB mode this holds, as long as the attacker cannot observe both plain- and ciphertext. Otherwise, the XOR of both reveals the output of the block cipher. With that information a regular DPA on the last round of the cipher becomes possible. For the case of bitstream encryption, the decrypted bitstream is directly sent to the PR controller, i.e., the plaintext never leaves the hardware and is not exposed to the attacker. For software decryption, this is hard to guarantee since the plaintext is transferred back to the CPU and potentially ends up in external RAM.

A straightforward mitigation is to use key whitening to hide the in- and outputs to the cipher from the attacker. XEX [35] and XTS [15] are modes of operation that achieve this. However, there are several published side-channel attacks on those constructions. Luo et al. show an attack on the tweaking function of XTS that exploits the simple structure of the Galois field multiplication by 2 [23]. This attack can be prevented by using XEX which uses multiplications with arbitrary values. In contrast, Unterluggauer et al. attacked the AES in such a scenario directly by concatenating DPA attacks on the last two rounds [43]. This attack is feasible irrespective of the tweaking function but it requires the attacker to change the input while keeping the IV constant (e.g., writing the same sector multiple times in the case of disk encryption). Whether or not this is a relevant attack vector depends on the application.

In cases where the second proposed LR-AEAD core is applicable, this attack vector is prevented by the LR-PRG that is used instead of AES-OFB. Similar to AES-OFB, each block of ciphertext is decrypted by XORing with one block of the keystream. However, the main difference is that in case of the LR-PRG the key changes from block to block. This means that the attacker cannot mount DPA attacks even if multiple plain- and ciphertext pairs are known because for every key, only one encryption is observable. This corresponds to an attack with data complexity 1, which is effectively an SPA attack. Since we defined that the AES implementation has to resist attacks with data complexity 2 in a side-channel

evaluation, this attack is not feasible as it is even harder for the attacker.

We conclude, that only the LR-AEAD core based on a LR-PRG is capable of protecting firmware decryption where we expect that the attacker can guess large parts of the plaintext. Nevertheless, protecting the confidentiality of software demands much more than only secure decryption. For example, external memory is vulnerable to cold-boot attacks [12] and code can be reverse engineered by observing its behavior [37]. To provide a sound solution for software encryption, these issues need to be addressed and, depending on the specific use case, additional measures have to be taken which are out of the scope of this work.

10 Conclusion

Side-channel attack resistant cryptographic cores and PUF-based key storage are now available in the newest generation of devices of leading vendors Xilinx, Intel and Microsemi. However, all of them are closed-source and none are updatable; hence it is usually not feasible to upgrade already deployed devices to a new hardware platform. As a step toward a vendor agnostic solution, we present a SCA secure and fully updatable mechanism to securely configure the FPGA logic starting from power-up until the whole system is booted and running. To achieve this, we leverage the PR options of FPGAs, in conjunction with a PUF-generated device intrinsic key and a leakage resilient authenticated decryption core to securely load hardware IP cores. The presented work, to the best of our knowledge, is the first that allows side-channel secure field updates of user IP cores and the decryption engine without relying on any manufacturer-provided *secret key* storage. The concept requires very limited trust in the manufacturer and provides the necessary flexibility if demands change or new attacks arise. This approach is orthogonal to upgrading to newer and more expensive feature-rich devices, and is also suited to retro-fit older devices as it uses features that are already widespread in current hardware.

Funding Open Access funding enabled and organized by Projekt DEAL. This work is partly supported by the German Federal Ministry of Education and Research in the project ALESSIO through grant number 16KIS0629 and by the EPSRC (EP/N508664/CSIT2).

Compliance with ethical standards

Availability of data and material Measurement data can be made available on demand.

Code availability https://github.com/Fraunhofer-AISEC/leakres_updatable_SoC

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Aysu, A., Gulcan, E., Moriyama, D., Schaumont, P., Yung, M.: End-to-end design of a PUF-based privacy preserving authentication protocol. In: T. Güneysu, H. Handschuh (eds.) *Cryptographic Hardware and Embedded Systems - CHES 2015—17th International Workshop*, Saint-Malo, France, September 13–16, 2015, *Proceedings, Lecture Notes in Computer Science*, Vol. 9293, pp. 556–576. Springer (2015). https://doi.org/10.1007/978-3-662-48324-4_28
2. Becker, G.T., Kumar, R.: Active and passive side-channel attacks on delay based PUF designs. *IACR Cryptol. ePrint Arch.* **2014**, 287 (2014)
3. Cnudde, T.D., Bilgin, B., Gierlichs, B., Nikov, V., Nikova, S., Rijmen, V.: Does coupling affect the security of masked implementations? In: S. Guilley (ed.) *Constructive Side-Channel Analysis and Secure Design - 8th International Workshop, COSADE 2017*, Paris, France, April 13–14, 2017, *Revised Selected Papers, Lecture Notes in Computer Science*, vol. 10348, pp. 1–18. Springer (2017). https://doi.org/10.1007/978-3-319-64647-3_1
4. Danger, J.L., Guilley, S., Nguyen, P., Rioul, O.: PUFs: Standardization and evaluation. In: *2016 Mobile System Technologies Workshop (MST)*, pp. 12–18 (2016). <https://doi.org/10.1109/MST.2016.11>
5. Degabriele, J.P., Janson, C., Struck, P.: Sponges resist leakage: The case of authenticated encryption. In: Galbraith, S.D., Moriai, S. (eds.) *Advances in Cryptology—ASIACRYPT 2019—25th International Conference on the Theory and Application of Cryptology and Information Security*, Kobe, Japan, December 8–12, 2019, *Proceedings, Part II, Lecture Notes in Computer Science*, vol. 11922, pp. 209–240. Springer (2019). https://doi.org/10.1007/978-3-030-34621-8_8
6. Ender, M., Moradi, A., Paar, C.: The unpatchable silicon: A full break of the bitstream encryption of Xilinx 7-Series FPGAs. In: *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, Boston, MA (2020)
7. Gassend, B., Clarke, D.E., van Dijk, M., Devadas, S.: Silicon physical random functions. In: V. Atluri (ed.) *Proceedings of the 9th ACM Conference on Computer and Communications Security—CCS 2002*, Washington, DC, USA, November 18–22, 2002, pp. 148–160. ACM (2002). <https://doi.org/10.1145/586110.586132>
8. Glowacz, C., Grosso, V., Poussier, R., Schüth, J., Standaert, F.: Simpler and more efficient rank estimation for side-channel security assessment. In: Leander, G. (ed.) *Fast Software Encryption—22nd International Workshop, FSE 2015*, Istanbul, Turkey, March 8–11, 2015, *Revised Selected Papers, Lecture Notes in Computer Science*, vol. 9054, pp. 117–129. Springer (2015). https://doi.org/10.1007/978-3-662-48116-5_6
9. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions (extended abstract). In: *25th Annual Symposium on Foundations of Computer Science*, West Palm Beach, Florida,

- USA, 24–26 October 1984, pp. 464–479. IEEE Computer Society (1984). <https://doi.org/10.1109/SFCS.1984.715949>
10. Gu, C., O'Neill, M.: Ultra-compact and robust FPGA-based PUF identification generator. In: 2015 IEEE International Symposium on Circuits and Systems, ISCAS 2015, Lisbon, Portugal, May 24–27, 2015, pp. 934–937. IEEE (2015). <https://doi.org/10.1109/ISCAS.2015.7168788>
 11. Guajardo, J., Kumar, S.S., Schrijen, G.J., Tuyls, P.: Physical unclonable functions, FPGAs and public-key crypto for IP protection. In: Bertels, K., Najjar, W.A., van Genderen, A.J., Vassiliadis, S. (eds.) FPL 2007, International Conference on Field Programmable Logic and Applications, Amsterdam, The Netherlands, 27–29 August 2007, pp. 189–195. IEEE (2007). <https://doi.org/10.1109/FPL.2007.4380646>
 12. Halderman, J.A., Schoen, S.D., Heninger, N., Clarkson, W., Paul, W., Calandrino, J.A., Feldman, A.J., Appelbaum, J., Felten, E.W.: Lest we remember: Cold boot attacks on encryption keys. In: van Oorschot, P.C. (ed.) Proceedings of the 17th USENIX Security Symposium, July 28–August 1, 2008, San Jose, CA, USA, pp. 45–60. USENIX Association (2008)
 13. Heyszl, J., Miller, K., Unterstein, F., Schink, M., Wagner, A., Gieser, H.A., Freud, S., Damm, T., Klein, D., Kügler, D.: Investigating profiled side-channel attacks against the DES key schedule. *IACR Trans. Cryptogr. Hardw. Embed. Syst.* **2020**(3), 22–72 (2020). <https://doi.org/10.13154/tches.v2020.i3.22-72>
 14. Hiller, M., Önal, A.G.: Hiding secrecy leakage in leaky helper data. In: Fischer, W., Homma, N. (eds.) Cryptographic Hardware and Embedded Systems—CHES 2017—19th International Conference, Taipei, Taiwan, September 25–28, 2017, Proceedings, Lecture Notes in Computer Science, Vol. 10529, pp. 601–619. Springer (2017). https://doi.org/10.1007/978-3-319-66787-4_29
 15. IEEE standard for cryptographic protection of data on block-oriented storage devices. *IEEE Std 1619-2007* pp. c1–32 (2008). <https://doi.org/10.1109/IEEESTD.2008.4493450>
 16. Jacob, N., Wittmann, J., Heyszl, J., Hesselbarth, R., Wilde, F., Pehl, M., Sigl, G., Fischer, K.: Securing FPGA SoC configurations independent of their manufacturers. In: 30th IEEE International System-on-Chip Conference, SOCC 2017, Munich, Germany, September 5–8, 2017, pp. 114–119. IEEE (2017). <https://doi.org/10.1109/SOCC.2017.8226019>
 17. Jaffe, J.: A first-order DPA attack against AES in counter mode with unknown initial counter. In: Paillier, P., Verbauwhede, I. (eds.) Cryptographic Hardware and Embedded Systems—CHES 2007, 9th International Workshop, Vienna, Austria, September 10–13, 2007, Proceedings, Lecture Notes in Computer Science, vol. 4727, pp. 1–13. Springer (2007). https://doi.org/10.1007/978-3-540-74735-2_1
 18. Jor, D.O., Heeger, D., Chan, C., Che, W., Saqib, F., Arenó, M., Plusquellic, J.: An autonomous, self-authenticating, and self-contained secure boot process for field-programmable gate arrays. *Cryptography* **2**(3), 15 (2018)
 19. Juels, A., Wattenberg, M.: A fuzzy commitment scheme. In: Motiwalla, J., Tsudik, G. (eds.) CCS '99, Proceedings of the 6th ACM Conference on Computer and Communications Security, Singapore, November 1–4, 1999, pp. 28–36. ACM (1999). <https://doi.org/10.1145/319709.319714>
 20. Kashyap, H., Chaves, R.: Compact and on-the-fly secure dynamic reconfiguration for volatile fpgas. *ACM Trans. Reconfig. Technol. Syst.* **9**(2), 11:1–11:22 (2016). <https://doi.org/10.1145/2816822>
 21. Krämer, J., Struck, P.: Leakage-resilient authenticated encryption from leakage-resilient pseudorandom functions. *IACR Cryptol. ePrint Arch.* **2020**, 280 (2020)
 22. Lohrke, H., Tajik, S., Krachenfels, T., Boit, C., Seifert, J.: Key extraction using thermal laser stimulation: A case study on Xilinx Ultrascale FPGAs. *IACR Trans. Cryptogr. Hardware Embed. Syst.* **2018**(3), 573–595 (2018). <https://doi.org/10.13154/tches.v2018.i3.573-595>
 23. Luo, C., Fei, Y., Ding, A.A.: Side-channel power analysis of XTS-AES. In: Aienza, D., Natale, G.D. (eds.) Design, Automation and Test in Europe Conference and Exhibition, DATE 2017, Lausanne, Switzerland, March 27–31, 2017, pp. 1330–1335. IEEE (2017). <https://doi.org/10.23919/DATE.2017.7927199>
 24. Maes, R., Herrewewege, A.V., Verbauwhede, I.: PUFKY: A fully functional PUF-based cryptographic key generator. In: Prouff, E., Schaumont, P. (eds.) Cryptographic Hardware and Embedded Systems—CHES 2012—14th International Workshop, Leuven, Belgium, September 9–12, 2012. Proceedings, Lecture Notes in Computer Science, vol. 7428, pp. 302–319. Springer (2012). https://doi.org/10.1007/978-3-642-33027-8_18
 25. Maes, R., van der Leest, V., van der Sluis, E., Willems, F.M.J.: Secure key generation from biased PUFs: extended version. *J. Cryptogr. Eng.* **6**(2), 121–137 (2016). <https://doi.org/10.1007/s13389-016-0125-6>
 26. McGrew, D.A., Viega, J.: The Galois/counter mode of operation (GCM). Submission to NIST Modes of Operation Process (2004)
 27. Medwed, M., Standaert, F., Joux, A.: Towards super-exponential side-channel security with efficient leakage-resilient prfs. In: Prouff, E., Schaumont, P. (eds.) Cryptographic Hardware and Embedded Systems—CHES 2012—14th International Workshop, Leuven, Belgium, September 9–12, 2012. Proceedings, Lecture Notes in Computer Science, vol. 7428, pp. 193–212. Springer (2012). https://doi.org/10.1007/978-3-642-33027-8_12
 28. Medwed, M., Standaert, F., Nikov, V., Feldhofer, M.: Unknown-input attacks in the parallel setting: Improving the security of the CHES 2012 leakage-resilient PRF. In: Advances in Cryptology—ASIACRYPT 2016—22nd International Conference on the Theory and Application of Cryptology and Information Security, Hanoi, Vietnam, December 4–8, 2016, Proceedings, Part I, pp. 602–623 (2016). https://doi.org/10.1007/978-3-662-53887-6_22
 29. Merli, D., Schuster, D., Stumpf, F., Sigl, G.: Side-channel analysis of PUFs and fuzzy extractors. In: McCune, J.M., Balacheff, B., Perrig, A., Sadeghi, A., Sasse, M.A., Beres, Y. (eds.) Trust and Trustworthy Computing—4th International Conference, TRUST 2011, Pittsburgh, PA, USA, June 22–24, 2011. Proceedings, Lecture Notes in Computer Science, vol. 6740, pp. 33–47. Springer (2011). https://doi.org/10.1007/978-3-642-21599-5_3
 30. Moradi, A., Barenghi, A., Kasper, T., Paar, C.: On the vulnerability of FPGA bitstream encryption against power analysis attacks: extracting keys from Xilinx Virtex-II FPGAs. In: Chen, Y., Danezis, G., Shmatikov, V. (eds.) Proceedings of the 18th ACM Conference on Computer and Communications Security, CCS 2011, Chicago, Illinois, USA, October 17–21, 2011, pp. 111–124. ACM (2011). <https://doi.org/10.1145/2046707.2046722>
 31. Moradi, A., Schneider, T.: Improved side-channel analysis attacks on Xilinx bitstream encryption of 5, 6, and 7 series. In: Standaert, F., Oswald, E. (eds.) Constructive Side-Channel Analysis and Secure Design—7th International Workshop, COSADE 2016, Graz, Austria, April 14–15, 2016, Revised Selected Papers, Lecture Notes in Computer Science, vol. 9689, pp. 71–87. Springer (2016). https://doi.org/10.1007/978-3-319-43283-0_5
 32. Pereira, O., Standaert, F., Vivek, S.: Leakage-resilient authentication and encryption from symmetric cryptographic primitives. In: Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security, Denver, CO, USA, October 12–16, 2015, pp. 96–108 (2015). <https://doi.org/10.1145/2810103.2813626>
 33. Peterson, E.: Leveraging asymmetric authentication to enhance security-critical applications using Zynq-7000 all programmable SoCs. Xilinx (2015)
 34. Prouff, E.: DPA attacks and s-boxes. In: Gilbert, H., Handschuh, H. (eds.) Fast Software Encryption: 12th International Workshop,

- FSE 2005, Paris, France, February 21–23, 2005, Revised Selected Papers, Lecture Notes in Computer Science, vol. 3557, pp. 424–441. Springer (2005). https://doi.org/10.1007/11502760_29
35. Rogaway, P.: Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In: Lee, P.J. (ed.) *Advances in Cryptology—ASIACRYPT 2004*, 10th International Conference on the Theory and Application of Cryptology and Information Security, Jeju Island, Korea, December 5–9, 2004, Proceedings, Lecture Notes in Computer Science, vol. 3329, pp. 16–31. Springer (2004). https://doi.org/10.1007/978-3-540-30539-2_2
 36. Rührmair, U., Sehnke, F., Sölter, J., Dror, G., Devadas, S., Schmidhuber, J.: Modeling attacks on physical unclonable functions. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) *Proceedings of the 17th ACM Conference on Computer and Communications Security, CCS 2010*, Chicago, Illinois, USA, October 4–8, 2010, pp. 237–249. ACM (2010). <https://doi.org/10.1145/1866307.1866335>
 37. Schink, M., Obermaier, J.: Taking a look into execute-only memory. In: Gantman, A., Maurice, C. (eds.) *13th USENIX Workshop on Offensive Technologies, WOOT 2019*, Santa Clara, CA, USA, August 12–13, 2019. USENIX Association (2019)
 38. Skorobogatov, S., Woods, C.: Breakthrough silicon scanning discovers backdoor in military chip. In: Prouff, E., Schaumont, P. (eds.) *Cryptographic Hardware and Embedded Systems—CHES 2012—14th International Workshop*, Leuven, Belgium, September 9–12, 2012, Proceedings, Lecture Notes in Computer Science, vol. 7428, pp. 23–40. Springer (2012). https://doi.org/10.1007/978-3-642-33027-8_2
 39. Skorobogatov, S., Woods, C.: In the blink of an eye: there goes your AES key. *IACR Cryptol. ePrint Arch.* **2012**, 296 (2012)
 40. Swierczynski, P., Moradi, A., Oswald, D., Paar, C.: Physical security evaluation of the bitstream encryption mechanism of Altera Stratix II and Stratix III FPGAs. *TRETS* **7**(4), 34:1–34:23 (2015). <https://doi.org/10.1145/2629462>
 41. Tebelmann, L., Pehl, M., Sigl, G.: EM side-channel analysis of bch-based error correction for PUF-based key generation. In: *Proceedings of the 2017 Workshop on Attacks and Solutions in Hardware Security, ASHES '17*, pp. 43–52. ACM, New York, NY, USA (2017). <https://doi.org/10.1145/3139324.3139328>
 42. Trimberger, S.M., Moore, J.J.: FPGA security: motivations, features, and applications. *Proc. IEEE* **102**(8), 1248–1265 (2014). <https://doi.org/10.1109/JPROC.2014.2331672>
 43. Unterluggauer, T., Mangard, S.: Exploiting the physical disparity: Side-channel attacks on memory encryption. In: Standaert, F., Oswald, E. (eds.) *Constructive Side-Channel Analysis and Secure Design—7th International Workshop, COSADE 2016*, Graz, Austria, April 14–15, 2016, Revised Selected Papers, Lecture Notes in Computer Science, vol. 9689, pp. 3–18. Springer (2016). https://doi.org/10.1007/978-3-319-43283-0_1
 44. Unterstein, F., Heyszl, J., De Santis, F., Specht, R.: Dissecting leakage resilient PRFs with multivariate localized EM attacks—a practical security evaluation on FPGA. In: Guilley, S. (ed.) *Constructive Side-Channel Analysis and Secure Design—8th International Workshop, COSADE 2017*, Paris, France, April 13–14, 2017, Revised Selected Papers, Lecture Notes in Computer Science, vol. 10348, pp. 34–49. Springer (2017). https://doi.org/10.1007/978-3-319-64647-3_3
 45. Unterstein, F., Heyszl, J., De Santis, F., Specht, R., Sigl, G.: High-resolution EM attacks against leakage-resilient PRFs explained—and an improved construction. In: *Topics in Cryptology—CT-RSA 2018—The Cryptographers' Track at the RSA Conference 2018*, San Francisco, CA, USA, April 16–20, 2018, Proceedings, pp. 413–434 (2018). https://doi.org/10.1007/978-3-319-76953-0_22
 46. Yu, M.M., Devadas, S.: Secure and robust error correction for physical unclonable functions. *IEEE Des. Test Comput.* **27**(1), 48–65 (2010). <https://doi.org/10.1109/MDT.2010.25>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.