Secure Lightweight Entity Authentication with Strong PUFs: Mission Impossible?

Jeroen Delvaux^{1,2}, Dawu Gu², Dries Schellekens¹, and Ingrid Verbauwhede¹

¹ ESAT/COSIC and iMinds, KU Leuven,
Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium
{jeroen.delvaux,dries.schellekens,ingrid.verbauwhede}@esat.kuleuven.be
² CSE/LoCCS, Shanghai Jiao Tong University,
800 Dongchuan Road, Shanghai 200240, China
dwgu@sjtu.edu.cn

Abstract. Physically unclonable functions (PUFs) exploit the unavoidable manufacturing variations of an integrated circuit (IC). Their input-output behavior serves as a unique IC ‘fingerprint’. Therefore, they have been envisioned as an IC authentication mechanism, in particular for the subclass of so-called strong PUFs. The protocol proposals are typically accompanied with two PUF promises: lightweight and an increased resistance against physical attacks. In this work, we review eight prominent proposals in chronological order: from the original strong PUF proposal to the more complicated converse and slender PUF proposals. The novelty of our work is threefold. First, we employ a unified notation and framework for ease of understanding. Second, we initiate direct comparison between protocols, which has been neglected in each of the proposals. Third, we reveal numerous security and practicality issues. To such an extent, that we cannot support the use of any proposal in its current form. All proposals aim to compensate the lack of cryptographic properties of the strong PUF. However, proper compensation seems to oppose the lightweight objective.

Keywords: physically unclonable function, entity authentication, lightweight.

1 Introduction

In this work, we consider a common authentication scenario with two parties: a **low-cost resource-constrained token** and a **resource-rich server**. Practical instantiations could be the following: RFID, smart cards and a wireless sensor network. One-way or possibly mutual entity authentication is the objective. The server has secure computing and storage at its disposal. Providing security is a major challenge however, given the requirements at the token side. Tokens typically store a secret key in non-volatile memory (NVM), using a mature technology such as EEPROM and its successor Flash, battery-backed SRAM or fuses. Cryptographic primitives import the key and perform an authentication protocol. Today’s problems are as follows. First, implementing cryptographic

primitives in a resource-constrained manner is rather challenging. Second, an attacker can gain physical access to the integrated circuit (IC) of a token. NVM has proven to be vulnerable to physical attacks [22]: the secret is stored permanently in a robust electrical manner. Third, most NVM technologies oppose the low-cost objective. EEPROM/Flash requires floating gate transistors, resulting in additional manufacturing steps with respect to a regular CMOS design flow. Battery-backed SRAM requires a battery. Circuitry to protect the NVM contents (e.g. a mesh of sensing wires) tends to be expensive.

Physically Unclonable Functions (PUFs) offer a promising alternative. Essentially, they are binary functions, with their input-output behavior determined by IC manufacturing variations. Therefore, they can be understood as a unique IC ‘fingerprint’, analogous to human biometrics. They might alleviate the aforementioned problems. Many PUFs allow for an implementation which is both resource-constrained and CMOS compatible. Furthermore, the secret is hidden in the physical structure of an IC, which is a much less readable format. Invasive attacks might easily destroy this secret, as an additional advantage. Several PUF-based protocols have been proposed, in particular for the subclass of so-called strong PUFs. We review the most prominent proposals: controlled PUFs [4–6], Öztürk et al. [17], Hammouri et al. [8], logically reconfigurable PUFs [11], reverse fuzzy extractors (FEs) [24], slender PUFs [16, 20] and the converse protocol [12]. The novelty of our work is threefold. First, we employ a unified notation and framework for ease of understanding. Second, we initiate direct comparison between protocols, which has been neglected in each of the proposals. Third, we reveal numerous security and practicality issues. To such an extent, that we cannot support the use of any proposal in its current form.

The remainder of this paper is organized as follows. Section 2 introduces notation and preliminaries. Section 3 describes and analyzes the strong PUF protocols. Section 4 provides an overview of the protocol issues. Section 5 concludes the work. We operate at protocol level, considering PUFs as a black box. The low-level protocol of Hammouri et al. [8] is therefore largely discussed in Appendix B, preceded by a discussion of PUF internals in Appendix A.

2 Preliminaries

2.1 Notation

Binary vectors are denoted with a bold lowercase character, e.g. $\mathbf{c} \in \{0, 1\}^{1 \times m}$. All vectors are row vectors. Their elements are selected with an index $i \geq 1$ between round brackets, e.g. $\mathbf{c}(1)$. The null vector is denoted as $\mathbf{0}$. Binary matrices are denoted with a single bold uppercase character, e.g. $\mathbf{H}$. Operations are the following: addition modulo 2 (XOR), e.g. $\mathbf{x} \oplus \mathbf{c}$, multiplication modulo 2, e.g. $\mathbf{e} \cdot \mathbf{H}^T$, concatenation, e.g. $\mathbf{x} || \mathbf{c}$, and bit inversion, e.g. $\bar{\mathbf{r}}$. Variable assignment is denoted with an arrow, e.g. $d \leftarrow d - 1$. Functions are printed in *italic*, with their input arguments between round brackets, e.g. Hamming weight $HW(\mathbf{r})$ and Hamming distance $HD(\mathbf{r}_1, \mathbf{r}_2)$.

2.2 Physically Unclonable Functions: Black Box Description

The m -bit input and n -bit output of a PUF are referred to as challenge $\mathbf{c}$ and response $\mathbf{r}$ respectively. Unfortunately for cryptographic purposes, the behavior of the challenge-response pairs (CRPs) does not correspond with a random oracle. First, the **response bits are not perfectly reproducible**: noise and various environmental perturbations (supply voltage, temperature, etc.) result in non-determinism. The reproducibility (error rate) differs per response bit. Second, the response bits are **non-uniformly distributed**: bias and correlations are present. The latter might enable so-called modeling attacks. One tries to construct a predictive model of the PUF, given a limited set of training CRPs. Machine learning algorithms have proven to be successful [21].

PUFs are often subdivided in two classes, according to their number of CRPs. **Weak PUFs** offer few CRPs: their total content (2^{m+n} bits) is of main interest. Architectures typically consist of an array of identically laid-out cells (or units), each producing one response bit. E.g. the SRAM PUF [9] and the ring oscillator PUF¹ [23] are both very popular. The total bit-content scales roughly linear with the required IC area. Although there might be some spatial correlation or a general trend among cells, a predictive model is typically of no concern. The response bits are mostly employed to generate a secret key, to be stored in volatile memory, in contrast to NVM. Post-processing logic, typically a fuzzy extractor (FE) [3], is required to ensure a reproducible and uniformly distributed key.

Strong PUFs offer an enormous number of CRPs, often scaling exponentially with the required IC area. They might greatly exceed the need for secret key generation and have been promoted primarily as lightweight authentication primitives. Architectures are typically able to provide a large challenge (e.g. $m = 128$), but only a very small response, mostly $n = 1$. CRPs are highly correlated, making modeling attacks a major threat. The most famous example is the arbiter PUF [13], described in appendix A. The definition of strong PUFs has shifted over the years. The original more specific notion in [7] assumes a large response space in addition to strong cryptographic properties: resistance against modeling and tamper evidence. Although more relevant than ever, we stick to the more practical recent notion.

2.3 Secure Sketch

The non-determinism of a PUF causes the regenerated instance of a response $\mathbf{r}$ to be slightly different: $\tilde{\mathbf{r}} = \mathbf{r} \oplus \mathbf{e}$, with $HW(\mathbf{e})$ small. Secure sketches [3] are a useful reconstruction tool, as defined by a two-step procedure. First, public helper data is generated: $\mathbf{p} = Gen(\mathbf{r})$. Second, reproduction is performed: $\mathbf{r} = Rep(\tilde{\mathbf{r}}, \mathbf{p})$. Helper data $\mathbf{p}$ unavoidably leaks some information about $\mathbf{r}$, although this entropy loss is supposed to be limited. Despite the rather generic definition, two

¹ We consider the most usable read-out modes which avoid correlations, e.g. pairing neighboring oscillators.

constructions dominate the implementation landscape, as specified below. Both the code-offset and syndrome construction employ a binary $[n, k, t]$ block code $\mathcal{C}$, with t the error-correcting capability. The latter construction requires a linear block code, as it employs the parity check matrix $\mathbf{H} \in \{0, 1\}^{(n-k) \times n}$. Successful reconstruction is guaranteed for both constructions, given $HW(\mathbf{e}) \leq t$. Information leakage is limited to $n - k$ bits. The hardware footprint is asymmetric: *Gen* is better suited for resource-constrained devices than *Rep* [24].

code-offset	<i>Gen</i>	<i>Rep</i>	syndrome	<i>Gen</i>	<i>Rep</i>
	Random $\mathbf{w} \in \mathcal{C}$ $\mathbf{p} \leftarrow \mathbf{r} \oplus \mathbf{w}$	$\tilde{\mathbf{w}} \leftarrow \tilde{\mathbf{r}} \oplus \mathbf{p} = \mathbf{w} \oplus \mathbf{e}$ Error-correct $\tilde{\mathbf{w}}$ to $\mathbf{w}$ $\mathbf{r} \leftarrow \mathbf{p} \oplus \mathbf{w}$		$\mathbf{p} \leftarrow \mathbf{r} \cdot \mathbf{H}^T$	$\mathbf{s} \leftarrow \tilde{\mathbf{r}} \cdot \mathbf{H}^T \oplus \mathbf{p} = \mathbf{e} \cdot \mathbf{H}^T$ Determine $\mathbf{e}$ $\mathbf{r} \leftarrow \tilde{\mathbf{r}} \oplus \mathbf{e}$

3 Lightweight Authentication with Strong PUFs

We analyze all strong PUF authentication schemes in chronological order. One can read the protocol discussions in arbitrary order, although we highly recommend to read Sections 3.1 and 3.2 first. All schemes employ two phases. The first phase is a one-time enrollment in a secure environment, following IC manufacturing. The server then obtains some information about the PUF, CRPs or even a predictive model via machine learning, to establish a shared secret. The destruction of one-time interfaces might permanently restrict the PUF access afterwards. The second phase is in-the-field deployment, where tokens are vulnerable to physical attacks. Token and server then authenticate over an insecure communication channel. In general: challenge $\mathbf{c}$ and response $\mathbf{r}$ are required to be of sufficient length, e.g. $m = n = 128$, to counteract brute-force attacks and random guessing.

3.1 Reference

For proper assessment, we define two reference authentication methods. Reference I employs a token with a secret key $\mathbf{k}$ programmed in NVM, as represented by Figure 1(a). Additional cryptographic logic performs the authentication. For ease of comparison, we opt for a hash function, hereby limiting ourselves to token authenticity only. The server checks whether a token can compute $\mathbf{a} \leftarrow \text{Hash}(\mathbf{k}, \mathbf{n})$, with $\mathbf{n}$ a random nonce. Reference II employs PUF technology, potentially providing more physical security at a lower manufacturing cost. We employ a weak PUF² to generate a secret key, as represented by Figure 1(b). The reproducibility and non-uniformity issue are resolved in a sequential manner, using a FE. A secure sketch first ensures reproducibility. *Gen* is executed only once during enrollment. Public helper data $\mathbf{p}$ is stored by the server, or alternatively at the token side in insecure (off-chip) NVM. A hash function performs entropy compression, hereby compensating the non-uniformity of

² Logic for generating challenges is implicitly present and might be as simple as reading out the full cell array.

Fig. 1. Token representation for all protocols and the two references. The following IC logic is not drawn: expansion of the strong PUF responses, intermediary registers (volatile) and control. A dashed line represent the interface with the server. One-time interfaces destructed after enrollment are marked by the symbol $\times$.

$\mathbf{r}$, in addition to the entropy loss caused by the helper data. One could generate a key as $\mathbf{k} \leftarrow \text{Hash}(\mathbf{r})$. We perform an optimization by merging this step with the authentication hash: $\mathbf{a} \leftarrow \text{Hash}(\mathbf{r}, \mathbf{n})$.

3.2 Naive Authentication

The most simple authentication method employs an unprotected strong PUF only [18], as shown in Figure 1(c). Figure 2 represents the corresponding protocol. The server collects a database of d arbitrary CRPs during enrollment. We assume the use of a true random number generator (TRNG). A genuine token should be able to reproduce the response for each challenge in the server database. Only an approximate match is required, taking PUF non-determinism into account: Hamming distance threshold ϵ implements this. To avoid token impersonation via replay, CRPs are discarded after use, limiting the number of authentications to d . Choosing e.g. $m = 128$, an attacker cannot gather and tabulate all CRPs and clone a token as such. Choosing e.g. $n = 128$, randomly guessing $\mathbf{r}$ is extremely unlikely to succeed.

Fig. 2. Naive authentication protocol. The thick arrow points from verifier to prover.

Modeling Attacks. Strong PUFs are too fragile for unprotected exposure, as demonstrated by a history of machine learning attacks [21]. A predictive PUF model would enable token impersonation. So far, no architecture can claim to be practical, well-validated and robust against modeling. Stated otherwise: no architecture does satisfy the original strong PUF definition given in [7], as has been observed by others (e.g. [14]). Two fundamental problems undermine the optimism for a breakthrough. First, strong PUFs extract their enormous amount of bits from a limited IC area only, hereby using a limited amount of circuit elements. A highly correlated structure is the unavoidable consequence. The arbiter PUF model in appendix A.2 provides some insights in this matter. Second, the more entangled and diffusing the structure of the PUF, the more robust against modeling, but the less reproducible the response as it accumulates more contributions from local noise sources. Appendix A.3 provides some insights in this matter.

Limited Response Space. In practice, strong PUFs provide a small response only, often $n = 1$. Replicating the PUF circuit is a simple but unfortunately very expensive solution. The lightweight approach is to evaluate a list of n challenges, hereby concatenating the response bits. Various methods can be employed to generate such a list. The server could generate the list, requiring no additional IC logic, but resulting in a large communication overhead [8]. A small pseudorandom number generator (PRNG), such as a linear feedback shift register (LFSR), is often employed [8, 16, 19, 20, 24]. Challenge $\mathbf{c}$ is then used as a seed value: $\tilde{\mathbf{r}} \leftarrow \text{PUF}(\text{PRNG}(\mathbf{c}))$. A variety of counter-based solutions can be applied as well [11]. Most protocol proposals in the remainder of this work suggest a particular response expansion method. We make abstraction of this, except when there is a related security problem.

3.3 Controlled PUFs

Controlled PUFs [4–6] provide reinforcement against modeling via a cryptographic hash function (one-way). Two instances, preceding and succeeding the PUF respectively, are shown in Figure 1(d). Figure 3 represents the corresponding protocol³. The preceding instance eliminates the chosen-challenge advantage of an attacker. The succeeding instance hides exploitable correlations due to the PUF structure. The latter hash seems to provide full protection by itself, but requires the use of a secure sketch: its avalanche effect would trigger on a single bit flip. CRPs stored by the server are accompanied by helper data.

Fig. 3. Authentication with controlled PUFs

Inferior to Reference II. The proposal seems to be inferior to reference II. First, the PUF is required to have an enormous instead of modest input-output space. This is inherently more demanding, even if one would extend reference II with a few challenge bits to enable the use of multiple keys. Second, server storage requirements scale linearly with the number of authentications, in contrast to constant-size.

³ Controlled PUFs were proposed in a wider context than CRP-based token authentication only.

3.4 Öztürk et al.

Öztürk et al. [17] employ two easy-to-model PUFs, as shown in Figure 1(e). Figure 4 represents the corresponding protocol. The outer PUF is assumed to possess a large response space, without defining challenge expansion logic however: it equips the token with CRP behavior. To prevent its modeling by an attacker, an internal secret $\tilde{\mathbf{x}}$ is XORed within the challenge path. A feedback loop, containing a (repeated) permutation and an inner PUF with a single response bit, is employed to update $\tilde{\mathbf{x}}$ continuously. During enrollment, the server has both read and write access to $\tilde{\mathbf{x}}$ via a one-time interface. This allows the server to construct models for either PUF, followed by an initialization of $\tilde{\mathbf{x}}$. The server has to keep track of $\tilde{\mathbf{x}}$, which is referred to as synchronization. The non-determinism of the inner PUF makes this non-trivial. One assumes an excellent match between the responses of the inner PUF and its model. At most one bit of $\tilde{\mathbf{x}}$ is assumed to be affected, in the seldom case of occurrence. An authentication failure (violation of ϵ) provides an indication thereof. One proposes a simple recovery procedure at the server side: bits of $\mathbf{x}'$ are successively flipped until the authentication succeeds.

Fig. 4. Authentication protocol of Öztürk et al

Issues Regarding $\tilde{\mathbf{x}}$. There are several issues related to the use of $\tilde{\mathbf{x}}$. First, it implicates the need for secure reprogrammable NVM, hereby undermining the advantages of PUFs. Either read or write access would enable an attacker to model the system, as during enrollment. Second, the synchronization effort is too optimistic. PUFs and their models typically have a 1 – 10% error rate. The server faces a continuous synchronization effort, not necessarily limited to a single error. Third, it enables denial-of-service attacks. An attacker can collect an (unknown) large number of CRPs, desynchronizing $\tilde{\mathbf{x}}$ and $\mathbf{x}'$. Propagation of errors across authentications makes the recovery effort rapidly infeasible.

Feedback Loop Comments. The permutation has to be chosen carefully. Repetition is bound to occur, as determined by the order k : $\pi^k = \pi$. This would implicate $\tilde{\mathbf{x}}$ to have identical bits, opposing its presumed non-uniformity.

A simple simulation confirms the significance for 64-bit vectors. The estimated probability of a randomly chosen permutation to have $k \leq 63$ equals $\approx 9\%$. Furthermore, the need for the inner PUF is questionable. First, its non-determinism poses a limit on the complexity of the feedback loop, and hence the modeling resistance of the overall system. Second, the outer PUF and the initialization of $\tilde{\mathbf{x}}$ already induce IC-specific behavior. Using a cryptographic hash function as feedback would resolve all foregoing comments. The resulting system would then be remarkably similar to a later proposal: logically reconfigurable PUFs [12].

3.5 Hammouri et al.

Hammouri et al. [8] employ again two strong PUFs, as shown in Figure 1(f). As before, both PUFs are modeled during enrollment. The outer PUF is an arbiter PUF. The inner PUF is a custom architecture based on the arbiter PUF. *Lin* largely compensates the non-linearity of the outer arbiter PUF. Figure 5 represents the authentication protocol. The proposal is non-generic, in comparison with all other protocols: correct functioning strongly depends on internal PUF details. We consider this a bad practice. Only a brief summary of the issues is provided here: we refer to Appendix B for the low-level argumentation.

Fig. 5. Authentication protocol of Hammouri et al

Unusual Inner PUF, Prone to Modeling Deficiencies. The inner PUF is a rather unusual extension of the arbiter PUF. The ability to construct a model, under the given procedure, is strongly layout-dependent and hence prone to deficiencies.

Unusual Modeling Procedure: Contradictive and Overcomplicated. Reading out $\tilde{\mathbf{x}}$, the response of the inner PUF, via a one-time interface would have made the enrollment easy. This would allow to model both PUFs separately. However, one designed a rather complicated procedure to model the inner PUF

via the response of the outer PUF. For this purpose, one did introduce a function *Lin* to linearize the outer arbiter PUF. This is rather contradictive as it degrades the overall modeling resistance. Furthermore, the enrollment might be problematic for a significant fraction of the fabricated tokens, depending on the IC-specific behavior of the outer PUF.

Non-functional: Error Propagation. We believe the proposal to be non-functional: non-determinism of the inner PUF is strongly amplified, leading to a persistent authentication failure. A minor modification could resolve this issue.

3.6 Logically Reconfigurable PUF

Logically reconfigurable PUFs [11] were proposed in order to make tokens recyclable, hereby reducing electronic waste. An internal state $\mathbf{x}$ is therefore mixed into the challenge path, as shown in Figure 1(e). Read access to $\mathbf{x}$ is allowed, although not required. There is no direct write access, although one can perform an update: $\mathbf{x}^{(t+1)} \leftarrow \text{Hash}(\mathbf{x}^{(t)})$. Figure 6 represents the authentication protocol.

Fig. 6. Authentication protocol for logically reconfigurable PUFs

Exacting PUF Requirements. The proposal does not aim to prevent PUF modeling attacks, despite providing forward/backward security proofs with respect to $\mathbf{x}$. Therefore, the practical value of the proposal is rather limited: the protocol cannot be instantiated due to lack of an appropriate strong PUF.

Issues Related to NVM. The proposal requires reprogrammable write-secure NVM, which is not free of issues. First, it undermines a main advantages of PUFs: low-cost manufacturing. Second, it enables denial-of-service attacks. An attacker can update $\mathbf{x}$ one or more times, invalidating the CRP database of the server. The proposal does not describe an authentication mechanism for the reconfiguration.

code size limitations, $\tilde{\mathbf{r}}$ is subdivided in non-overlapping sections: a set of helper data vectors, of length $n - k$ each, is transferred rather than a single $\tilde{\mathbf{p}}$. One employs 7 sections with $k = 21$, aiming at a security level of $128 < 7 \cdot 21$ bit. The proposal does not provide an explicit warning to refuse fixed points of the LFSR, e.g. $\mathbf{c} = \mathbf{0}$. This would have been appropriate in order to avoid a trivial server impersonation attack. Fixed points will result in either $\tilde{\mathbf{r}} = \mathbf{0}$ or $\tilde{\mathbf{r}} = \bar{\mathbf{0}}$, assuming stability for one particular response bit. An attacker could hence then guess $\mathbf{b}$ with success probability $1/2$. A token impersonation threat, which is far less obvious, is described next.

First consider the unrealistic but desired case of a perfectly deterministic PUF. Consider an arbitrary response section, denoted as $\mathbf{r}$. Helper data leakage can be understood via an underdetermined system of linear equations $\mathbf{H} \cdot \mathbf{r}^T = \mathbf{p}^T$, having $n - k$ equations for n unknowns. Consider a challenge $\mathbf{c}_1$ leading to an expanded response $\mathbf{r}_1 = (r_1 \ r_2 \ \dots \ r_n)$. One could easily construct a challenge $\mathbf{c}_2$ leading to a response $\mathbf{r}_2 = (r_2 \ r_3 \ \dots \ r_{n+1})$, given the use of an LFSR. Repeating the former mechanism, we can construct challenges $\mathbf{c}_3, \mathbf{c}_4, \dots, \mathbf{c}_q$, with $\mathbf{r}_q = (r_q \ r_{q+1} \ \dots \ r_{q+n-1})$. An attacker collects all data in a single system of equations:

$$\mathbf{A} \cdot (r_1 \ r_2 \ \dots \ r_{q+n-1})^T = \begin{pmatrix} \mathbf{p}_1^T \\ \mathbf{p}_2^T \\ \vdots \\ \mathbf{p}_q^T \end{pmatrix} \quad \text{with } \mathbf{A} = \begin{pmatrix} \mathbf{H} & \mathbf{0}^T & \dots & \mathbf{0}^T \\ \mathbf{0}^T & \mathbf{H} & \dots & \mathbf{0}^T \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{0}^T & \mathbf{0}^T & \dots & \mathbf{H} \end{pmatrix}.$$

Equation dependencies have to be considered. We performed experiments where $\mathbf{A}$ is transformed to reduced row echelon form (*rref*). A distinction between cyclic and non-cyclic codes seems to be crucial, as illustrated in Figure 8. In the latter case, q can be small and an attacker is quasi able to solve the system. The persistence of few unknowns is only a minor inconvenience. In the former case, the number of unknowns remains k . However, a repeated machine learning attack can be performed instead, for large q , exploiting the introduction of sections. Arbiter PUFs can be modeled with only a few thousand CRPs, so consider e.g. $q = 10000$, including both training and verification data. The correct combination of unknowns (only $2^k = 2^{21}$ possibilities) would result in the observable event of a high modeling accuracy.

We now consider a PUF which is not perfectly deterministic. To minimize this inconvenience, an attacker could ensure the IC's environment to be very stable. Code parameters, which are normally chosen in order to maintain robustness against a variety of environmental perturbations (supply voltage, outside temperature, etc.), become relatively relaxed then. Subsequently, we consider a transformation $\mathbf{H}^\star = \mathbf{T} \cdot \mathbf{H}$, which selects linear combinations of the rows of $\mathbf{H}$. The transformation is chosen so that the rows of $\mathbf{H}^\star$ do have a low Hamming weight, making it feasible to incorporate a limited set of stable bits only. One could develop a variety of algorithms in order to find a suitable transformation.

Suboptimal. Even with $d = 1$, the protocol still allows for an unlimited number of authentications. Token impersonation via replay has already been prevented

Fig. 8. PRNG exploitation for the reverse fuzzy extractor: illustration for non-cyclic Hamming and cyclic BCH codes, with $[n = 7, k = 4, t = 1]$ and $q = 5$. Black and white represent 0 and 1 respectively.

by the use of nonce $\mathbf{n}$. This observation could result in numerous protocol simplifications, as has been acknowledged in [24]. However, one does not state clearly that there would a simplification for the PUF as well: the enormous input-output space is no more required, even a weak PUF could be employed. Despite all the former, one still promotes the use of $d > 1$, arguing that it offers an increased side channel resistance. We argue that this countermeasure might not outweigh the missed advantages and that there might be numerous more rewarding countermeasures. The modified protocol proposal [14] does not discuss the foregoing efficiency matter, although it uses $d = 1$, hereby providing a weak PUF (with implicit challenge) as an example.

3.8 Slender PUF

The slender PUF proposal [16] includes three countermeasures against modeling, while avoiding the need for cryptographic primitives, as clear from Figure 1(i). First, one requires a strong PUF with a high resistance. A model is constructed during enrollment via auxiliary one-time interfaces. One employs a variant of the XOR arbiter PUF for this purpose (see appendix A.3). Second, the exposure of $\tilde{\mathbf{r}}$ is limited to random substrings $\tilde{\mathbf{s}}$, hereby obfuscating the CRP link. The corresponding procedure *SubCirc* treats the bits in a circular manner. Third, server and token both contribute to the challenge via their respective nonces $\mathbf{c_S}$ and $\mathbf{c_T}$, counteracting chosen-challenge attacks. Figure 9 represents the protocol.

A protocol extension has been proposed in [20]. One presents a fourth countermeasure against modeling. Substring $\tilde{\mathbf{s}}$ is padded with random bits before transmission, again in a circular manner: $\text{SubCirc}(\tilde{\mathbf{s}}_P, j) \leftarrow \tilde{\mathbf{s}}$ with $\tilde{\mathbf{s}}_P \in \{0, 1\}^{1 \times n_P} \xleftarrow{U} \text{TRNG}()$ and $j \in [1 \ n_P] \xleftarrow{U} \text{TRNG}()$. Furthermore, the optional establishment of a session key is introduced, via concatenation of secret indices i and j . A repeated execution of the protocol is required to obtain a key of sufficient length.

PRNG Exploitation. The protocol employs a PRNG to expand the PUF response space: $\tilde{\mathbf{r}} \leftarrow \text{PUF}(\text{PRNG}(\mathbf{c_S}, \mathbf{c_T}))$. However, the PRNG construction of the proof-of-concept implementation might allow for token impersonation:

Fig. 9. Slender PUF protocol

$\text{PRNG}(\mathbf{c}_S, \mathbf{c}_T) = \text{LFSR}(\mathbf{c}_S) \oplus \text{LFSR}(\mathbf{c}_T)$. We assume an identical feedback polynomial for both LFSRs, which is the most intuitive assumption⁵. A malicious token might then return $\mathbf{c}_T \leftarrow \mathbf{c}_S$, resulting in an expanded list of challenges all equal to $\mathbf{0}$. The server's PUF model outputs either $\mathbf{r}' = \mathbf{0}$ or $\mathbf{r}' = \bar{\mathbf{0}}$. So provided a substring $\tilde{\mathbf{s}} = \mathbf{0}$ (or $\tilde{\mathbf{s}}_P = \mathbf{0}$), authentication does succeed with a probability $1/2$. Via replay, one could increase the probability to 1. Eavesdropping on a single genuine protocol execution is required: $\mathbf{c}_S^{(1)}, \mathbf{c}_T^{(1)}$ and $\tilde{\mathbf{s}}^{(1)}$. The malicious prover gets authenticated with the following: $\mathbf{c}_T^{(2)} \leftarrow \mathbf{c}_S^{(2)} \oplus \mathbf{c}_S^{(1)} \oplus \mathbf{c}_P^{(1)}$ and $\tilde{\mathbf{s}}^{(2)} \leftarrow \tilde{\mathbf{s}}^{(1)}$. One can easily replay old sessions keys as well. We stress that careful PRNG redesign can resolve all of the former.

Exacting PUF Requirements. The PUF requirements are rather exacting and partly opposing. On one hand, the PUF should be easy-to model, requiring a highly correlated structure. On the other hand, CRP correlations enable statistical attacks, due to the lack of cryptographic primitives. Such attacks exploit the knowledge of a function $P(r_u = r_v) = f(\mathbf{c}_u, \mathbf{c}_v)$ for a certain strong PUF architecture. CRPs with $|P(r_u = r_v) - 1/2| > 0$ are correlated and hence exploitable. One might be able to retrieve indices i (and j) as such⁶, circumventing the *SubCirc* countermeasure and hence restoring the CRP link. Although the proposed PUF seems to offer this delicate balance (see Appendix A.4), it

⁵ The proof-of-concept implementation employs 128-bit LFSRs, with $\mathbf{c}_S$ and $\mathbf{c}_T$ as the initial states, without specifying feedback polynomials. Furthermore, FPGA implementation results (Table III in [16] and Table 8 in [20]) strongly suggest the use of identical feedback polynomials.

⁶ The slender PUF protocol has a direct predecessor: the pattern matching key generator [19], which serves as an alternative for a conventional fuzzy extractor [3]. CRP correlation attacks are equally dangerous for this construction and might lead to key-recovery. However, the risk has not been analyzed as thoroughly as in [16]. Furthermore, another statistical attack of fundamentally different nature has recently been published [2].

comes at a price: several one-time interfaces⁷, high response non-determinism and a limited modeling accuracy. Furthermore, the user's control regarding the challenge list should be highly restricted: statistical attacks are very powerful if an attacker has a (partial) chosen-challenge advantage. The PRNG-TRNG construction should accomplish this goal.

3.9 Converse Authentication

Figures 1(j) and 10 represent the converse authentication protocol [12]. The authentication is one-way, in a less conventional setting where tokens verify the server. The difference vector (XOR) between two responses is denoted as $\mathbf{d}$. The restriction $\mathbf{d} \neq \mathbf{0}$ prevents an attacker from impersonating the server, choosing $\{\mathbf{c}_i = \mathbf{c}_j, \mathbf{p}_i = \mathbf{p}_j\}$. Optionally, one can extend the protocol with the establishment of a session key $\mathbf{k} \leftarrow \text{Hash}(\mathbf{r}_i || \mathbf{r}_j)$. The attacker capabilities are restricted: invasive analysis of the prover IC is assumed to be impossible. Furthermore, one assumes an eavesdropping attacker, trying to impersonate the server given a log of genuine authentication transcripts.

Fig. 10. Converse authentication protocol

Attacker Overly Restricted. The attacker capabilities are unclear and overly restricted to be practical. First, restricting invasion would automatically extend to physical attacks in general. Furthermore, this greatly reduces the need for PUFs, with respect to the traditional approach of storing a key in NVM. It is also not clear whether an attacker is allowed to freely query the server, as exploited hereafter. We argue that this should be possible, as the protocol is initiated by the token.

Scalability Issues. The probability of success for a randomly guessing attacker would be $1/2^n$, when trying to impersonate the server. In certain sense,

⁷ Chains are modeled separately via machine learning, before XORing takes place.

the server faces a similar probability, posing an upper limit on n for practicality reasons. Successful authentication of the server relies on the availability of a given $\mathbf{d}$ within its database. A scalable method to (construct and) search within a database was not discussed and seems far from obvious. Assume a cumbersome trial and error procedure as a search method: a pairwise selection has a probability of $1/2^n$ to be usable. Memory requirements are mild compared to the search workload, as there are $\frac{d(d-1)}{2}$ pairwise selections, although still enormous in comparison to all other proposals. A database of size $d = 2^{25}$ is mentioned to be realistic. A major threat for server impersonation is the following. An attacker might query the server and construct a personal database: $\langle \mathbf{d}_i, \mathbf{c}_{i1}, \mathbf{p}_{i1}, \mathbf{c}_{i2}, \mathbf{p}_{i2} \rangle$. Authentication will succeed, although a session key cannot be retrieved if present.

Predecessor Issues. Although the protocol is in essence identical to a direct predecessor [1], it is not described in terms of modifications. Nevertheless, a few issues have (quietly) been resolved. We observe five modifications. First, *Rep* is acknowledged to require helper data. Second, one introduces the restriction $\mathbf{d} \neq \mathbf{0}$. However, this event occurs so seldom that an explicit check could be regarded as overhead. Third, $\mathbf{d}$ is generated by a TRNG instead of a PRNG. To avoid replay, the latter construction would require secure NVM, hereby undermining the advantages of PUFs. Fourth, the protocol is initiated by a token instead of the server. Unfortunately, this enables exhaustion of the server database as described before. Fifth, PUF responses are reinforced by a cryptographic hash function. In its absence, we see a direct opportunity for prover impersonation in the occasional case that $HW(\mathbf{d}) < t$. Consider an arbitrary response $\mathbf{r} = Rep(PUF(\mathbf{c}), \mathbf{p})$. For both secure sketch constructions, an attacker might be able to produce a given $\mathbf{d}$:

Code-offset construction: $\mathbf{r} \oplus \mathbf{d} = Rep(PUF(\mathbf{c}), \mathbf{p} \oplus \mathbf{d})$.

Syndrome construction: $\mathbf{r} \oplus \mathbf{d} = Rep(PUF(\mathbf{c}), \mathbf{p} \oplus \mathbf{d} \cdot \mathbf{H}^T)$.

4 Overview and Discussion

Tables 1 and 2 provide an overview of Section 3. We do not support the use of any strong PUF proposal in its current form, given the numerous amount of issues. However, we do not object the use of both weak PUF protocols: reference II and the modified reverse FE proposal [14]. We now discuss the strong PUF issues.

Two proposals rely on **reprogrammable NVM**: Öztürk et al. [17] and logically reconfigurable PUFs. Their respective assumptions of R/W- and W-security undermine a major benefit of PUF technology: an increased resistance against physical attacks. Furthermore, the need for reprogramming undermines a second potential benefit: low-cost manufacturing, as CMOS-compatible fuses cannot be used. Finally, updating the NVM state was identified as a denial-of-service attack for both proposals.

Table 1. For all protocols: token hardware (left), the authenticity provided (middle) and the secret stored by the server (right)

	Weak PUF	Strong PUF ¹	NVM	TRNG	Gen	Rep	Hash	1× interface	Other	Token auth.	Server auth.	# Auth.	CRPs	Model	Key
Reference I	×	×	✓	×	×	×	✓	✓	×	✓	×	∞	×	×	✓
Reference II	✓	×	×	×	×	✓	✓	✓	×	✓	×	∞	×	×	✓
Naive	×	✓	×	×	×	×	×	×	×	✓	×	<i>d</i>	✓	×	×
Controlled	×	✓	×	×	×	✓	✓	✓	×	✓	×	<i>d</i>	✓	×	×
Öztürk et al.	×	✓ ²	✓ ⁶	×	×	×	×	✓	✓	✓	×	∞	×	✓	×
Hammouri et al.	×	✓ ²	×	✓	×	×	×	✓	✓	✓	×	∞	×	✓	×
Reconfiguration	×	✓ ³	✓ ⁶	×	×	×	✓	×	✓	✓	×	<i>d</i>	✓	×	×
Reverse FE	×	✓ ⁴	×	×	✓	×	✓	✓	✓	✓	✓	∞	✓	×	×
Slender	×	✓ ⁵	×	✓	×	×	×	✓	✓	✓	×	∞	×	✓	×
Converse	×	✓	×	✓	×	✓	✓	✓	✓	✓	✓	∞	✓	✓	×

¹ Including response expansion.⁴ Non-determinism lower bound.² Easy-to-model.⁵ Both easy- and hard-to-model.³ Robust against modeling.⁶ Reprogrammable.**Table 2.** Issues revealed in this work. Implementation-dependent issues are printed in *italic*.

Protocol	Issues
Controlled	- Inferior to reference II.
Öztürk et al. [17]	- NVM undermines the advantages of PUFs. - Synchronization effort is presented too optimistic. - Denial-of-service attack. - <i>Choice of permutation requires care: avoid low orders.</i>
Hammouri et al. [8]	- Unusual inner PUF, prone to modeling deficiencies. - Non-functional: internal error propagation. - Unusual modeling procedure: contradictory and overcomplicated.
Reconfiguration	- Unrealistic PUF requirement: robust against modeling. - NVM undermines the advantages of PUFs. - Denial-of-service attack.
Reverse FE	- Exacting PUF requirements to counteract replay attacks. - <i>PRNG exploitation, leading to token/server impersonation.</i> - No need for strong PUF.
Slender PUF	- <i>PRNG exploitation, leading to token impersonation.</i> - Exacting PUF requirements to counteract statistical attacks.
Converse	- Attacker model too restricted. - Scalability issues, leading to server impersonation. - Predecessor issues.

PUF responses are not perfectly reproducible. Protocols employ two approaches to overcome this issue: error correction and error tolerance. Unfortunately, two proposals struggle with the latter approach. PUF non-determinism is greatly underestimated in Öztürk et al. [17]: the protocol synchronization effort is presented too optimistic. We believe the proposal of Hammouri et al [8]. to be non-functional because of internal error propagation.

In practice, **strong PUFs do have a small output space.** There are various methods to resolve this issue, all imposing a certain efficiency burden. However, the protocol proposals have very little attention for this topic. Even though system security is not necessarily unaffected. We specified attacks for the proof-of-concept implementations of the reverse FE and slender PUF proposal, both exploiting the challenge expansion PRNG.

In practice, **strong PUFs are insecure against modeling.** Two proposals are therefore too demanding: logically reconfigurable PUFs and the slender PUF protocol. Although the latter offers some countermeasures, there is an opposing requirement for the PUF to be easy-to-model, leading to a delicate balancing exercise. In general, proposals not reinforced by a cryptographic hash function are much more likely to be vulnerable.

Several proposals rely on a **secure TRNG**. Tampering with its randomness opens new perspectives for a physical attack. Its use seems unavoidable however if server authentication is a must, as for the converse protocol, to avoid replay attacks. The reverse FE proposal extracts its non-determinism from the PUF instead, which has been identified as a delicate balancing exercise. Two protocols without server authentication employ their TRNG as a modeling countermeasure: Hammouri et al. [8] and the slender PUF protocol.

All proposals aim to provide lightweight entity authentication, which addresses a highly relevant need. However, in many use cases, there will be accompanying security requests: message confidentiality and integrity, privacy, etc. We did not consider **protocol extensibility** in this work, although it might be of interest when designing a new protocol. References I and II, which employ a secret key, might benefit from a huge amount of scientific literature. Like-minded, the establishment of a session key has been proposed as an extension for the converse and slender PUF proposals.

5 Conclusion

Various protocols utilize a strong PUF to provide lightweight entity authentication. We described the most prominent proposals using a unified notation, hereby creating a first overview and initializing direct comparison as well. We defined two reference authentication methods, to identify the misuse of PUFs. Our analysis revealed numerous security and practicality issues. Therefore, we do not recommend the use of any strong PUF proposal in its current form. Most proposals aim to compensate the lack of cryptographic properties of the strong PUF. However, proper compensation seems to be in conflict with the lightweight objective. More fundamental physical research is required, aiming

to create a truly strong PUF with great cryptographic properties. If not, we are inclined to recommend conventional PUF-based key generation as a more promising alternative. The observations and lessons learned in this work can facilitate future protocol design.

Acknowledgment. The authors greatly appreciate the support received. The European Commission through the ICT programme under contract FP7-ICT-2011-317930 HINT. The Research Council of KU Leuven: GOA TENSE (GOA/11/007), the Flemish Government through FWO G.0550.12N and the Hercules Foundation AKUL/11/19. The national major development program for fundamental research of China (973 Plan) under grant no. 2013CB338004. Jeroen Delvaux is funded by IWT-Flanders grant no. 121552. The authors would like to thank Anthony Van Herrewege (KU Leuven, ESAT/COSIC) for his valuable comments.

References

1. Das, A., Kocabaş, Ü., Sadeghi, A.-R., Verbauwhede, I.: PUF-based secure test wrapper design for cryptographic SoC testing. In: Design, Automation & Test in Europe, DATE 2012, pp. 866–869 (March 2012)
2. Delvaux, J., Verbauwhede, I.: Attacking PUF-Based Pattern Matching Key Generators via Helper Data Manipulation. In: Benaloh, J. (ed.) CT-RSA 2014. LNCS, vol. 8366, pp. 106–131. Springer, Heidelberg (2014)
3. Dodis, Y., Ostrovsky, R., Reyzin, L., Smith, A.: Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. *SIAM J. Comput.* 38(1), 97–139 (2008)
4. Gassend, B., Clarke, D.E., van Dijk, M., Devadas, S.: Silicon physical random functions. In: ACM Conference on Computer and Communications Security, CCS 2002, pp. 148–160 (November 2002)
5. Gassend, B., Clarke, D.E., van Dijk, M., Devadas, S.: Controlled Physical Random Functions. In: Annual Computer Security Applications Conference, ACSAC 2002, pp. 149–160 (December 2002)
6. Gassend, B., van Dijk, M., Clarke, D.E., Torlak, E., Devadas, S., Tuyls, P.: Controlled physical random functions and applications. *ACM Trans. Inf. Syst. Secur.* 10(4) (2008)
7. Guajardo, J., Kumar, S.S., Schrijen, G.-J., Tuyls, P.: FPGA Intrinsic PUFs and Their Use for IP Protection. In: Paillier, P., Verbauwhede, I. (eds.) CHES 2007. LNCS, vol. 4727, pp. 63–80. Springer, Heidelberg (2007)
8. Hammouri, G., Öztürk, E., Sunar, B.: A tamper-proof and lightweight authentication scheme. *Journal Pervasive and Mobile Computing* 6(4) (2008)
9. Holcomb, D.E., Burleson, W.P., Fu, K.: Power-Up SRAM State as an Identifying Fingerprint and Source of True Random Numbers. *IEEE Trans. Computers* 58(9) (2009)
10. Hospodar, G., Maes, R., Verbauwhede, I.: Machine Learning Attacks on 65nm Arbiter PUFs: Accurate Modeling poses strict Bounds on Usability. In: IEEE Workshop on Information Forensics and Security (WIFS) 2012, pp. 37–42 (December 2012)

11. Katzenbeisser, S., Kocabaş, Ü., van der Leest, V., Sadeghi, A.-R., Schrijen, G.-J., Schröder, H., Wachsmann, C.: Recyclable PUFs: Logically Reconfigurable PUFs. In: Preneel, B., Takagi, T. (eds.) CHES 2011. LNCS, vol. 6917, pp. 374–389. Springer, Heidelberg (2011)
12. Kocabaş, Ü., Peter, A., Katzenbeisser, S., Sadeghi, A.-R.: Converse PUF-Based Authentication. In: Katzenbeisser, S., Weippl, E., Camp, L.J., Volkamer, M., Reiter, M., Zhang, X. (eds.) Trust 2012. LNCS, vol. 7344, pp. 142–158. Springer, Heidelberg (2012)
13. Lee, J.W., Lim, D., Gassend, B., Suh, G.E., van Dijk, M., Devadas, S.: A technique to build a secret key in integrated circuits for identification and authentication applications. In: 2004 Symposium on VLSI Circuits, pp. 176–179 (June 2004)
14. Maes, R.: Physically Unclonable Functions: Constructions, Properties and Applications. PhD Thesis, KU Leuven (August 2012)
15. Majzoobi, M., Koushanfar, F., Potkonjak, M.: Testing techniques for hardware security. In: IEEE International Test Conference (ITC), pp. 1–10 (October 2008)
16. Majzoobi, M., Rostami, M., Koushanfar, F., Wallach, D.S., Devadas, S.: Slender PUF Protocol: A Lightweight, Robust, and Secure Authentication by Substring Matching. In: IEEE Symposium on Security and Privacy (SP), pp. 33–44 (May 2012)
17. Öztürk, E., Hammouri, G., Sunar, B.: Towards Robust Low Cost Authentication for Pervasive Devices. In: IEEE Conference on Pervasive Computing and Communications, PerCom (March 2008)
18. Pappu, R.: Physical One-Way Functions. PhD Thesis, MIT, ch. 9 (2001)
19. Paral, Z., Devadas, S.: Reliable and Efficient PUF-Based Key Generation Using Pattern Matching. In: Hardware-Oriented Security and Trust, HOST 2011, pp. 128–133 (June 2011)
20. Rostami, M., Majzoobi, M., Koushanfar, F., Wallach, D.S., Devadas, S.: Robust and Reverse-Engineering Resilient PUF Authentication and Key-Exchange by Substring Matching. *IEEE Transactions on Emerging Topics in Computing*, 13 (2014)
21. Rührmair, U., Sehnke, F., Sölter, J., Dror, G., Devadas, S., Schmidhuber, J.: Modeling attacks on physical unclonable functions. In: ACM Conference on Computer and Communications Security, CCS 2010, pp. 237–249 (October 2010)
22. Skorobogatov, S.: Semi-invasive attacks - a new approach to hardware security analysis, University of Cambridge, Computer Laboratory (April 2005)
23. Suh, G.E., Devadas, S.: Physical unclonable functions for device authentication and secret key generation. In: Design Automation Conference, DAC 2007, pp. 9–14 (June 2007)
24. Van Herrewege, A., Katzenbeisser, S., Maes, R., Peeters, R., Sadeghi, A.-R., Verbauwhede, I., Wachsmann, C.: Reverse Fuzzy Extractors: Enabling Lightweight Mutual Authentication for PUF-Enabled RFIDs. In: Keromytis, A.D. (ed.) FC 2012. LNCS, vol. 7397, pp. 374–389. Springer, Heidelberg (2012)

A Arbiter PUF

A.1 Architecture

Arbiter PUFs [13] quantify manufacturing variability via the propagation delay of logic gates and interconnect. The high-level functionality is represented by Figure

11(a). A rising edge propagates through two paths with identically designed delays, as imposed by layout symmetry. Because of nanoscale manufacturing variations however, there is a delay difference Δt between both paths. An arbiter decides which path ‘wins’ the race ($\Delta t \leq 0$) and generates a response bit r .

Fig. 11. Arbiter PUF (a) and its XOR variant (b)

The two paths are constructed from a series of m switching elements. The latter are typically implemented with a pair of 2-to-1 multiplexers. Challenge bits determine for each stage whether path segments are crossed or uncrossed. Each stage has a unique contribution to Δt , depending on its challenge bit. Challenge vector $\mathbf{c}$ determines the time difference Δt and hence the response bit r . The number of CRPs equals 2^m . The response reproducibility differs per CRP: the smaller $|\Delta t|$, the easier to flip side because of various physical perturbations.

A.2 Vulnerability to Modeling Attacks

Arbiter PUFs show additive linear behavior, which makes them vulnerable to modeling attacks. A decomposition in individual delay elements is given in Figure 12. Both intra- and inter-switch delays contribute to Δt , as represented by white and gray/black squares respectively. We incorporate the latter in their preceding switches, without loss of generality, to facilitate the derivation of a delay model. In the end, delay elements are important as far as they generate delay differences between both paths. Therefore, each stage can be described by two delay parameters only: one for each challenge bit state, as illustrated in Figure 13. The delay difference at the input of stage i flips in sign for the crossed configuration and is incremented with δt_i^1 or δt_i^0 for crossed and uncrossed configurations respectively.

Fig. 12. Arbiter PUF decomposed in individual delay elements, represented by small squares which are prone to manufacturing variability. The interconnecting lines have zero delay.

Fig. 13. Delay behavior of an arbiter stage

The impact of a δt on Δt is incremental or decremental for an even and odd number of subsequent crossed stages respectively. By lumping together the δt 's of neighboring stages, one can model the whole arbiter PUF with $m + 1$ independent parameters only (and not $2m$). A formal expression for $\Delta t = \gamma \tau^T$ is shown below. Vector $\gamma \in \{\pm 1\}^{1 \times (m+1)}$ is a transformation of challenge vector $\mathbf{c}$. Vector $\tau \in \mathbb{R}^{1 \times (m+1)}$ contains the lumped stage delays. The more linear a system, the easier to learn its behavior. By using γ instead of $\mathbf{c}$ as ML input, a great deal of non-linearity is avoided. The non-linear threshold operation $\Delta t \leq 0$ remains however. Only 5000 CRPs were demonstrated to be sufficient to model non-simulated 64-stage arbiter PUFs with an accuracy of about 97% [10].

$$\tau = \frac{1}{2} \begin{pmatrix} \delta t_0 & \delta t_1^0 - \delta t_1^1 \\ \delta t_1^0 + \delta t_1^1 & + \delta t_2^0 - \delta t_2^1 \\ & \vdots \\ \delta t_{m-1}^0 + \delta t_{m-1}^1 & + \delta t_m^0 - \delta t_m^1 \\ \delta t_m^0 + \delta t_m^1 & \end{pmatrix}^T \text{ and } \gamma = \begin{pmatrix} 1 - 2(\mathbf{c}(1) \oplus \dots \oplus \mathbf{c}(m)) \\ 1 - 2(\mathbf{c}(2) \oplus \dots \oplus \mathbf{c}(m)) \\ \vdots \\ 1 - 2\mathbf{c}(m) \\ 1 \end{pmatrix}^T.$$

A.3 XOR Variant

Several variants of the arbiter PUF increase the resistance against ML. They introduce various forms of non-linearity for this purpose. We only consider the XOR variant. The response bits of multiple arbiter chains are XORED to produce a single response bit, as shown in Figure 11(b). All chains have the same challenge as input. The more chains, the more resistance against ML: the required number

of CRPs and the computation time both increase rapidly [21]. However, the reproducibility of r decreases with the number of chains as well: each additional chain injects non-determinism into the overall system. A practical limit on the ML resistance is hence imposed. Modifications to the architecture can improve former trade-off [15]: e.g. a permutation of the challenge bits across chains, as employed for the slender PUF protocol [16, 20].

A.4 CRP Correlations: Enabling Statistical Attacks

Machine learning attacks exploit CRP correlations in an implicit manner. Statistical attacks benefit from their explicit exploitation, hereby assuming the knowledge of a function $P(r_1 = r_2) = f(\mathbf{c}_1, \mathbf{c}_2)$. Such functions have already been determined via simulations [15]. We are the first to derive an analytical model, as demonstrated for the the arbiter PUF and its XOR variant in Figure 14. Modified architectures can converge to the ideal curve $f = 1/2$ more rapidly. Let $\mathcal{N}(\mu, \sigma)$ denote a normal distribution with mean μ and standard deviation σ . Let $\phi(x, \sigma)$ and $\Phi(x, \sigma)$ denote the probability density function and cumulative distribution function respectively, assuming $\mu = 0$. We assume all δt 's to have a distribution $\mathcal{N}(0, \sigma_1)$, which is a common and well-validated practice in previous work. Although not fully correct⁸, we then assume the elements of $\boldsymbol{\tau}$ to have a distribution $\mathcal{N}(0, 2\sigma_1)$. We introduce the variable $h = HD(\gamma_1, \gamma_2)$.

Fig. 14. Correlations for the arbiter PUF and its XOR variant. Dots represent simulations results. The mathematical model, drawn continuously although of discrete nature, matches reasonably well. Vertical dashed lines enclose 99% of the data for randomly chosen challenges. The more chains being XORed, the better one approximates the ideal behavior $f = 1/2$, but the larger the response non-determinism.

⁸ We neglect dependencies within $\boldsymbol{\tau}$ and we also ignore the different form of $\boldsymbol{\tau}(1)$ and $\boldsymbol{\tau}(m+1)$

B Hammouri et al. [8]

B.1 Unusual Inner PUF, Prone to Modeling Deficiencies

The inner PUF is a custom architecture based on the arbiter PUF. One proposes a rather unusual extension of the challenge space. Out of two individual chains, one aims to construct a single reconfigurable chain. For each stage, one out of two switching elements is selected, hereby introducing a second challenge $\mathbf{s}$. The proposal ignores the need to describe the reconfiguration logic: Figure 15 (right) provides a generic schematic, including both intra- and inter-switch delays. Via reconfiguration, one aims to provide a large response space. One does evaluate $\mathbf{c}_I$ for a fixed list of configurations vectors $\{\mathbf{s}_1, \mathbf{s}_2, \dots\}$, hereby concatenating the response bits. The configuration vectors are generated by a PRNG, having $\mathbf{s}_1$ as initial state. Note that one could have provided a large response space with a regular arbiter PUF as well, given the use of a PRNG.

Fig. 15. Reconfigurable stage of the inner PUF

The architecture is required to be easy-to-model: the server has to construct a model during enrollment. The overall structure is additive, as for a regular arbiter PUF, and therefore we expect modeling to be feasible. Although one could have derived a generic (but complicated) delay model, possibly leading to an efficient modeling method, the authors propose a shortcut. The overall delay model is supposed to be separable in terms of the individual chains. One does construct a model τ_U for the upper chain, applying $\mathbf{s} = \mathbf{0}$. Similar, one obtains a model τ_L for the lower chain, applying $\mathbf{s} = \mathbf{\bar{0}}$. The overall model $\tau(\mathbf{s})$ selects elements from both vectors: $\tau(i) = \mathbf{s}(i)\tau_L(i) + \mathbf{s}(i)\tau_U(i)$. The variability of intra-stage delays is being neglected, justified by placing the stages far apart in the circuit lay-out. An approximating delay model for upper/lower chain is derived as well, as shown in Figure 15 (left).

Although it might be possible to make all the former workable (there is no proof-of-concept implementation), the approach is strongly layout-dependent and prone to modeling deficiencies. Apart from being area consuming: positioning stages far

apart might not be sufficient to justify separability. Intra- and inter-stage variations originate from CMOS transistors and metal interconnect respectively. We highly doubt that the former would be per se negligible with respect to the latter. It is possible though to enhance separability: upsizing transistors of the switches, inserting minimum-sized inverter chains in between switches, etc. Furthermore, one should distinguish the metal interconnect before and after the stage selection logic, as shown in Figure 15 (right). Variability of the latter would undermine the separability. We stress that all these complications could have been avoided easily, e.g. by using a regular arbiter PUF.

B.2 Unusual Modeling Procedure: Contradictive and Overcomplicated

Reading out $\tilde{\mathbf{x}}$, the response of the inner PUF, via a one-time interface would have made the enrollment easy. This would allow to model both PUFs separately. However, one designed a rather complicated procedure to model the inner PUF via the response of the outer arbiter PUF. The latter has a single bit response: during authentication, a list of challenges $\mathbf{c}_O/\mathbf{c}_I$ is transferred between token and server to expand its response space. One does introduce a function Lin , compensating the non-linearity of the arbiter PUF, apart from the final thresholding step: Δt is a linear function of $Lin(\tilde{\mathbf{x}})$. This transformation is relatively simple, as shown below. We note that this is rather contradictive: it degrades the overall modeling resistance. Fixing $\mathbf{c}_O = \mathbf{0}$, one obtains a system $\tilde{r} \leftarrow PUF_O(Lin(\tilde{\mathbf{x}}))$. Error propagation from $\tilde{\mathbf{x}}$ to $\tilde{r}$ is very limited then. An error in bit $\tilde{\mathbf{x}}(i)$ would flip the sign of $\gamma(i)$ only, corresponding to $h = 1$ in Figure 14. During enrollment, one can force the PRNG to maintain either $\mathbf{0}$ or $\overline{\mathbf{0}}$ as its state, allowing to model the upper and lower chain separately. This requires some sort of destructive interface, similar to our proposal to read out $\tilde{\mathbf{x}}$ directly. Depending on $\mathbf{c}_I$, $\tilde{\mathbf{x}}$ will be either $\mathbf{0}$ or $\overline{\mathbf{0}}$, apart from potential noisiness, hereby maximizing h . As a consequence, one can distinguish either case, as $\tilde{r}$ is very likely to flip. This enables modeling of the inner PUF. Although some IC samples might be problematic: the pair $\tilde{\mathbf{x}} = \mathbf{0}/\overline{\mathbf{0}}$ occasionally result in $\Delta t \approx 0$, maintaining $\tilde{r}$ in a noisy state.

$$Lin(\tilde{\mathbf{x}}) = \begin{pmatrix} \tilde{\mathbf{x}}(1) \oplus \tilde{\mathbf{x}}(2) \\ \tilde{\mathbf{x}}(2) \oplus \tilde{\mathbf{x}}(3) \\ \vdots \\ \tilde{\mathbf{x}}(m-1) \oplus \tilde{\mathbf{x}}(m) \\ \tilde{\mathbf{x}}(m) \end{pmatrix}^T.$$

B.3 Non-functional: Error Propagation

We believe the proposal to be non-functional: non-determinism of the inner PUF is strongly amplified, leading to a persistent authentication failure. The minimization of h does not hold for the system $\tilde{r} \leftarrow \widetilde{PUF_O}(Lin(\tilde{\mathbf{x}}) \oplus \mathbf{c}_O)$, in the general case that $\mathbf{c}_O \neq \mathbf{0}$. A single error in $\tilde{\mathbf{x}}$ will flip the sign of $\tilde{r}$ with a probability close to $1/2$. This could have been avoided by implementing the system $\tilde{r} \leftarrow PUF_O(Lin(\tilde{\mathbf{x}} \oplus \mathbf{c}_O))$ instead.