Key-recovery Attacks on Various RO PUF Constructions via Helper Data Manipulation

Jeroen Delvaux and Ingrid Verbauwhede

ESAT/SCD-COSIC and iMinds, KU Leuven
Kasteelpark Arenberg 10, B-3001 Leuven-Heverlee, Belgium
Email: {firstname.lastname}@esat.kuleuven.be

Abstract. Physically Unclonable Functions (PUFs) are emerging as hardware security primitives. They are mainly used to generate secret keys which are inherently unique for every manufactured sample of a chip. Ring Oscillator (RO) PUFs are among the most widely researched PUFs. In this work, we claim various RO PUF constructions to be vulnerable against manipulation of their public helper data. Partial/full key-recovery is a threat for the following constructions, in chronological order. (1) Temperature-aware cooperative RO PUFs, proposed at HOST 2009. (2) The sequential pairing algorithm, proposed at HOST 2010. (3) Group-based RO PUFs, proposed at DATE 2013. (4) Or more general, all entropy distiller constructions proposed at DAC 2013.

Keywords: PUF, fuzzy extractor, helper data

1 Introduction

With the ubiquity of electronic computing devices (ICs) in our everyday lives, cryptographic algorithms have become an important building block. Hereby, one heavily relies on the ability to store secret information. However, an attacker can easily gain physical access to the IC. Hardware attacks, either invasive or noninvasive, are thus a significant threat for modern applications.

Traditionally, binary keys are stored in programmable on-chip Non-Volatile Memory (NVM). EEPROM and its successor Flash are the main technologies. However, this approach has proven to be vulnerable against hardware attacks [5]. The permanent nature of storage worsens the problem as no limits are posed on the time frame of the attacker. Circuits that detect hardware invasion offer additional protection. Unfortunately, they suffer from practical limitations. They might be expensive, bulky, battery powered, vulnerable to bypassing and/or not appropriate for lightweight environments.

Physically Unclonable Functions (PUFs) have been proposed as a more secure and more efficient alternative. Silicon PUFs quantify the unique manufacturing variability of nanoscale structures. The secret is stored in intrinsic physical features of a chip, resulting in some remarkable security advantages in comparison to on-chip NVM. First, PUFs are often assumed to be resistant against invasive attacks. One can argue that invasion damages the physical structure of the PUF. Second, keys are inherently unique for each manufactured sample of a chip and there is no need to explicitly program them. Third, the key is only generated and stored in on-chip Volatile Memory (VM) when key-dependent operations have to be performed, as such posing limits on the attacker's time frame.

Ring Oscillator (RO) PUFs are very popular, inter alia because they can be implemented on FPGA. We describe their architecture in section 2. In order to generate high-quality keys, post-processing logic is required, as described in section 3. Public helper data, stored in NVM, is employed hereby. Various helper data constructions are described in sections 4 and 5. We claim four of them to be vulnerable against manipulation of the helper NVM, as discussed in section 6. Partial or even full key recovery might be possible. An extensive reflection of our findings is given in section 7. Section 8 concludes the work.

2 Ring Oscillator PUFs

RO PUFs quantify the manufacturing variability of identically laid-out oscillators. Each RO, consisting of an odd number of inverters, will have a unique frequency f. Frequencies are typically measured

by counting rising or falling edges on a wire connecting two subsequent inverters. Figure 1 shows the PUF architecture as originally proposed in [6]. One can distinguish four components: a RO array, multiplexers to access individual ROs, counters providing a frequency measurement and a comparator.

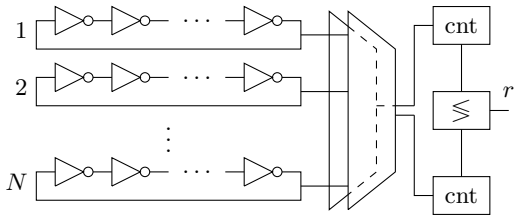

Fig. 1. RO PUF as originally proposed.

A pairwise frequency comparison ($\Delta f \lessgtr 0$) generates a single response bit r. There are $N(N-1)/2$ pairwise comparisons, although their response bits are interdependent. Consider the following minimal example given three ROs: $RO_A.f < RO_B.f$ and $RO_B.f < RO_C.f$ implicates $RO_A.f < RO_C.f$. The total PUF entropy is only $log_2(N!)$ bit as there are $N!$ ways to sort the frequency values. We hereby assume the ideal case, with all permutations equally likely.

For convenience, ring oscillators are typically implemented as a two-dimensional array. Without loss of generality, we still label each RO with a univariate index $i \in [1\ N]$. The multiplexer-counter-comparator architecture might greatly vary. The degree of parallelism for generating response bits, is directly affected. Consider the following two extreme cases for example: a dedicated counter per RO and a single counter accessing all ROs via a giant multiplexer.

3 Post-processing Logic

Cryptographic keys should be perfectly reproducible and uniformly distributed. Unfortunately, PUF response bits are not directly usable because of several imperfections. We distinguish two categories of problems: reliability and entropy. We list their root causes and provide examples for RO PUFs in particular. On-chip digital post-processing logic is required to address the former issues. Public helper NVM is employed hereby: data is generated during a one-time post-manufacturing enrollment phase.

3.1 Reliability

PUF response bits are not perfectly reproducible. The main responsible is CMOS device noise, which is a random time-dependent phenomenon. Instability of the environment, mostly defined by the IC supply voltage and the outside temperature, worsens the problem. RO frequencies increase with both increasing supply voltage and decreasing temperature. The environmental impact might be limited however, depending on the IC's application. The larger the nominal frequency discrepancy $|\Delta f|$ for a pairwise comparison, the more reliable the corresponding response bit.

3.2 Entropy

PUF response bits might possess undesired statistical properties, reducing the entropy of the secret key. Bias is a major concern for instance, whereby the probability of a response bit to be '1' (or '0') does not equal 50%. Correlations between response bits are another frequent issue. Asymmetries in the layout are one potential root cause. Systematic manufacturing variability, which is spatially correlated, is another root cause. As illustrated in figure 2, only random variation is desired. Furthermore, the occasional occurrence of $\Delta f = 0$ (counter values are discrete) introduces bias given that either '1' or '0' has to be returned.

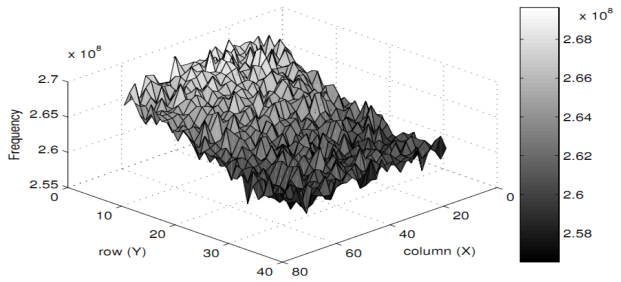

Fig. 2. Frequency topology of a RO array [4]. The slope corresponds with systematic variability. Only the random surface roughness is desired.

4 RO Pair Selection

There is a variety of methods to select pairs from a RO array. Their goal is to output many high-entropy bits, possibly with an incentive towards reliability to lighten the burden of post-processing logic. We now discuss four approaches, in order of increasing complexity. We do present attacks for the latter two methods in section 6.

4.1 Chain of Neighbors

Pairing neighboring ROs is perhaps the most intuitive approach. The reduced impact of spatial correlations is the main advantage [3]. By sharing ROs across pairs, up to $N - 1$ independent bits can be generated.

4.2 1-out-of-k Masking

A 1-out-of-k masking scheme [6] partitions ROs into groups of size k, each generating one response bit. During enrollment, the fastest and slowest RO are selected for every group, maximizing the frequency difference and favoring reliability as such. Their indices are saved in public helper NVM. Parameter k represents a trade-off between reliability and efficiency. A total of $\lfloor N/k \rfloor$ bits is generated.

4.3 Sequential Pairing Algorithm

The sequential pairing algorithm [8] selects up to $\lfloor N/2 \rfloor$ pairs, with each RO employed once at most. The frequency difference of every pair exceeds a given discrepancy threshold Δf_{th}. Pairing information is again stored in public helper NVM. Algorithm 1 provides simplified pseudocode. In the original proposal, one requires frequency measurements at two environmental extremes.

4.4 Temperature-aware Cooperative

Temperature-aware cooperative RO PUFs [7] operate within a user-defined temperature range: $T \in [T_{min},\ T_{max}]$. An on-chip temperature sensor is assumed to be available, posing strong limits on the applicability. Furthermore, RO frequencies are assumed to be linearly dependent on the temperature.

Neighboring ROs are paired, without overlap, leading to a total of $\lfloor N/2 \rfloor$ pairs. A frequency discrepancy threshold Δf_{th} is employed to assess their reliability. Pairs are classified in three groups, as illustrated on figure 3. Good pairs obey $|\Delta f(T)| > \Delta f_{th}$ within the whole operating range: they generate one reliable bit. Bad pairs obey $|\Delta f(T)| \leq \Delta f_{th}$ within the whole operating range: they are discarded. Some pairs are stable except for the region $[T_l,\ T_h]$ around their crossover point: they cooperate to generate reliable bits, assisted by public helper data.

For every cooperating pair, one does store the stability boundaries T_l and T_h in public helper NVM. Apart from the crossover region, no help is required. Although one has to compensate for the

Algorithm 1: Sequential Pairing (Simplified)

Input: Frequency measurements $RO_i.f$ with $i \in [1\ N]$
Frequency discrepancy threshold Δf_{th}
Output: List of pairs $\{RO_i, RO_j\}$
Sort frequencies in descending order and store indices as vector π:
$RO_{\pi(1)}.f > RO_{\pi(2)}.f > \ldots > RO_{\pi(N)}.f$
$i \leftarrow 1$
for $j \leftarrow \lceil \frac{N}{2} \rceil + 1$ ***to*** N **do**
  if $RO_{\pi(i)}.f - RO_{\pi(j)}.f > \Delta f_{th}$ **then**
    Create pair $\{RO_{\pi(i)}, RO_{\pi(j)}\}$
    $i \leftarrow i + 1$

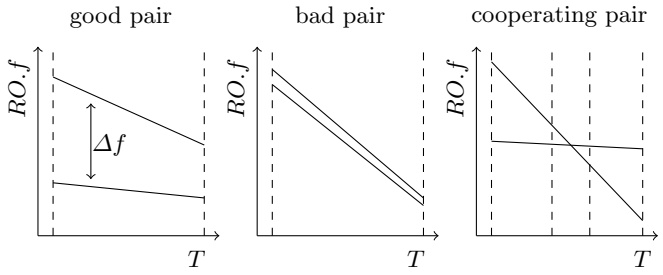

Fig. 3. Temperature-aware cooperative RO PUF: classification of RO pairs. The outer dashed lines represent the operating range: $[T_{min},\ T_{max}]$.

crossover: the response bit is inverted if $T > T_h$. Within the crossover region, one does rely on another cooperating pair with a nonintersecting crossover region. Its index is stored in public helper NVM as well.

Pairs cooperate in a masked manner, to prevent helper data leakage. Consider a first cooperating pair, requesting help and having response bit r_{c1}. A masking response bit r_{g1}, originating from a corresponding good pair, is assigned. A second cooperating pair, providing help and having response bit r_{ci}, should satisfy the following constraint: $r_{c1} \oplus r_{g1} = r_{ci}$. Note that both r_{g1} and r_{ci} are stable in the crossover region of the first cooperating pair, allowing for reconstruction.

However, we claim that the proposed masking scheme is not necessarily free from leakage. The second cooperating pair should be selected at random and hence not with a deterministic procedure that iterates over all candidates until the constraint is met. Otherwise, one exposes the following information for all non-selected candidates: $r_{cj} \neq r_{ci}$.

5 Group-based RO PUF

Group-based RO PUFs have first been introduced at HOST 2010 [8]. As the initial design had several shortcomings, the authors redefined their construction at DATE 2013 [9]. For ease of understanding, we make abstraction of the gradual development. In traditional designs, ROs are paired to generate a response bit. The group-based approach is very different in this regard. ROs are partitioned into groups, with their size not limited to two anymore.

The so-called entropy distiller, the main novelty, has been introduced in parallel at DAC 2013 too [10], although with more experimental evidence. Its use is not limited to the group-based approach. It can be employed with the pair selection methods of section 4 as well. For our attacks in section 6, we will also consider the latter use case.

The high-level architecture is represented by figure 4. We explicitly indicate the IC boundaries and interfaces to clarify an attacker's point of view. Like this, we also stress that all building blocks

do require an on-chip implementation. The resulting key is stored in on-chip VM, for as long as needed. An application with key-dependent operations communicates with the user, either directly or indirectly. We now discuss the building blocks separately.

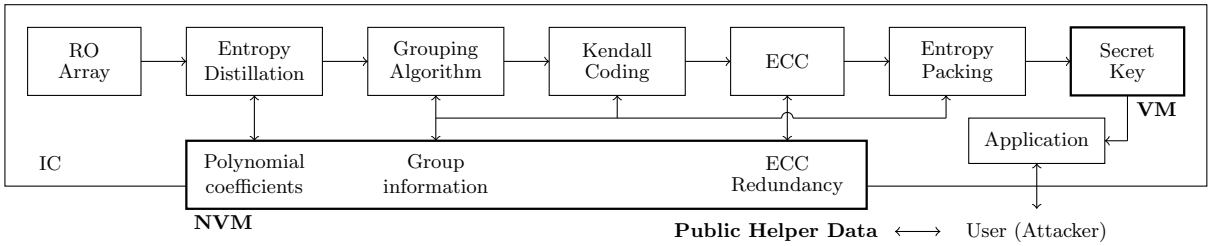

Fig. 4. Group-based RO PUF.

5.1 Entropy Distiller

The entropy distiller removes systematic variations, so that only random variations remain. Systematic variations are modeled via polynomial regression on the two-dimensional RO frequency map. The residuals represent the random variations hereby. An expression for the polynomial of degree p is given below. Experiments in [10] indicate $p = 2$ and $p = 3$ as good values, given an array of 16×32 ROs. Coefficients $\beta_{i,j}$ may be determined in a least mean squares manner. They are stored as public helper data. A subtraction procedure removes systematic variations for every regeneration of the key.

$$f(x, y) = \sum_{i=0}^{p} \sum_{j=0}^{i} \beta_{i,j} x^{i-j} y^{j}.$$

5.2 Grouping Algorithm

The grouping algorithm partitions the ROs into groups $G_1, G_2, \ldots$ The partitioning is strict: every RO is assigned to exactly one group. Within a group, every possible pair of ROs does exceed a frequency discrepancy threshold Δf_{th}, favoring reliability. Response bits will be extracted for each group independently. Algorithm 2 provides pseudocode for the grouping procedure. It optimizes the available entropy of $\sum_i log_2(|G_i|!)$ bits: having few large groups is more beneficial than having many small groups.

5.3 Kendall Coding

For every group G_i, there will be a particular order of the RO frequencies. A binary representation is required. Table 1 illustrates two schemes, assuming there are four ROs (A, B, C and D) in the group and hence $4! = 24$ possible orders. The most compact representations do require $\lceil log_2(|G_i|!) \rceil$ bit, as illustrated for the second column.

However, to facilitate the subsequent error-correction step, a non-minimum length coding scheme is proposed. One observes that errors mostly occur in form of a flip, e.g. BACD to BCAD. Using Kendall coding, one bit is generated for every possible pair of ROs, requiring $|G_i|(|G_i| - 1)/2$ bits in total. Error-correction requirements are relaxed in terms of error rate, as there is only one error per flip. Unfortunately, the workload increases quickly with the group size $|G_i|$.

5.4 Error-Correcting Code

Incoming bits are clustered in blocks, which are all error-corrected independently. An Error-Correcting Code (ECC), able to correct t errors, is employed hereby. The first generated instance of each block is selected as a golden reference. Public helper data allows regenerated instances to be error-corrected, so that they match the golden reference.

Algorithm 2: GROUPING

Input: Frequency measurements $RO_i.f$ with $i \in [1\ N]$
Frequency discrepancy threshold Δf_{th}
Output: Group assignments $RO_i.group$
Sort frequencies in descending order and store indices as vector π:
$RO_{\pi(1)}.f > RO_{\pi(2)}.f > \ldots > RO_{\pi(N)}.f$
$RO_0.f \leftarrow \infty$
for $i \leftarrow 1$ ***to*** N **do**
  $last(i) \leftarrow 0$
for $i \leftarrow 1$ ***to*** N **do**
  $done \leftarrow 0$
  $j \leftarrow 1$
  while $done = 0$ **do**
    if $RO_{\pi(last(j))}.f - RO_{\pi(i)}.f > \Delta f_{th}$ **then**
      $RO_{\pi(i)}.group \leftarrow j$
      $last(j) \leftarrow i$
      $done \leftarrow 1$
    $j \leftarrow j + 1$

Order	Compact	Kendall
ABCD	00000	000000
ABDC	00001	000001
ACBD	00010	000100
ACDB	00011	000110
ADBC	00100	000011
ADCB	00101	000111
BACD	00110	100000
BADC	00111	100001
BCAD	01000	110000
BCDA	01001	111000
BDAC	01010	101001
BDCA	01011	111001

Order	Compact	Kendall
CABD	01100	010100
CADB	01101	010110
CBAD	01110	110100
CBDA	01111	111100
CDAB	10000	011110
CDBA	10001	111110
DABC	10010	001011
DACB	10011	001111
DBAC	10100	101011
DBCA	10101	111011
DCAB	10110	011111
DCBA	10111	111111

Table 1. Coding of oscillator frequency order.

5.5 Entropy Packing

Kendall coding is noted to be non-uniform: many bit vectors are never used. To maintain entropy, conversion to a compact coding scheme (as in table 1) is proposed. However, one does not mention that the problem is only fixed partially, since $|G_i|!$ is not a power of two, given $|G_i| > 2$.

6 Attacks via Helper Data Manipulation

Before discussing the specifics for each RO PUF construction, we describe the general framework of our attacks. PUF response bits are retrieved one by one (or in small groups). For each iteration, two or more hypotheses H_i provide a statement about the bits of concern, of which exactly one is correct. Every hypothesis corresponds with a specific manipulation of the public helper data. We exploit differences in key regeneration failure rate to distinguish them. Failures are easy to observe, as key-dependent operations either execute with the wrong key or not execute at all.

Key regeneration fails if more than t errors occur within a certain ECC block. For ease of explanation, we assume all bits to fit within a single ECC block. However, extension to multiple blocks is fairly straightforward. For some constructions, the presence of an ECC might be optional: one can then consider the degenerate case $t = 0$. To quantify ECC failure behavior, the Probability Density

Function (PDF) of the numbers of errors is particularly useful. A binomial distribution, given an averaged error probability for a single bit, might be very accurate for large blocks. Although our attacks do not depend on the former assumption. Failures rarely occur in practice, assuming well-chosen ECC parameters. To accelerate the attack, we introduce additional errors accordingly for all sets of helper data.

Figure 5 provides an example, in case of two hypotheses H_0 and H_1. The nominal PDF, corresponding to unmodified helper data, serves as a reference. PDFs corresponding to modified helper data, are slightly shifted with respect to each other and hence distinguishable. The common offset originates from the additional errors.

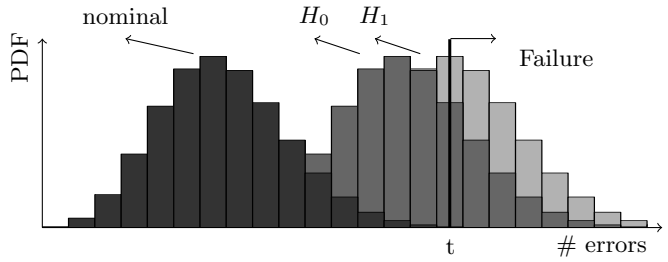

Fig. 5. Distinguishing hypotheses by observing key generation failure rates.

6.1 RO PUF with Sequential Pairing

Key recovery is fairly straightforward for the sequential pairing algorithm. Consider two RO pairs, resulting in response bits r_1 and r_2. We formulate two hypotheses as shown below. To distinguish them, we swap the order of the two pairs in public helper NVM. If H_0 is correct, the failure rate is not modified. However, if H_1 is correct, the failure rate does increase. Matching r_1 with all other response bits $r_2, r_3, \ldots r_{\lfloor N/2 \rfloor}$, only two possible values remain for the secret key. For the final decision, the performance of two corresponding sets of ECC helper data can be compared.

$$H_0 : r_1 = r_2. \quad H_1 : r_1 \neq r_2.$$

To accelerate the attack, more errors can be injected. For instance by swapping additional pairs, accordingly for the original and modified helper data. Initially, the additional pairs can be chosen at random. After revealing some response bit relations however, one can select those pairs which will introduce an error for sure.

6.2 Temperature-aware cooperative RO PUF

An attacker can retrieve the response bit relation for all cooperating pairs. Consider a first cooperating pair, having response bit r_{c1}. A second cooperating pair, having response bit r_{ci}, provides assistance. Consider another cooperating pair with response bit r_{cj} and having a nonintersecting crossover region too. We formulate the two hypotheses shown below. Helper data is modified so that r_{cj} provides assistance. If H_0 is correct, the failure rate is not modified. If H_1 is correct however, the failure rate does increase. To accelerate the attack, more errors can be injected. For instance, by manipulating the boundaries of crossover regions (T_l and T_h).

$$H_0 : r_{ci} = r_{cj}. \quad H_1 : r_{ci} \neq r_{cj}.$$

6.3 Group-based RO PUF

An attacker can retrieve the full key for group-based RO PUFs, due to the ability to directly reprogram the key. By injecting steep polynomials into the entropy distiller, one can completely overshadow random frequency variations. Hereby, the attacker's intended pattern is simply superimposed onto the original spatial correlation map. By repartitioning the groups, one can force response bits to be either '1' or '0'. As a final step, one does recompute the ECC helper data, which can be done as all response bits are known.

Consider the example of figure 6a, given an array of 4×10 ROs. A tilted plane provides a strong gradient in the horizontal direction, as represented by the grayscale. We repartition the groups so that they all contain two ROs. The responses of G_2 to G_{20} are fully determined by the attacker, as their ROs are put multiple columns apart. The response of G_1 however, is fully determined by random frequency variations. Figure 6b illustrates an identical scenario, but now with the ROs of G_1 in different columns. A quadratic polynomial then needs to be injected.

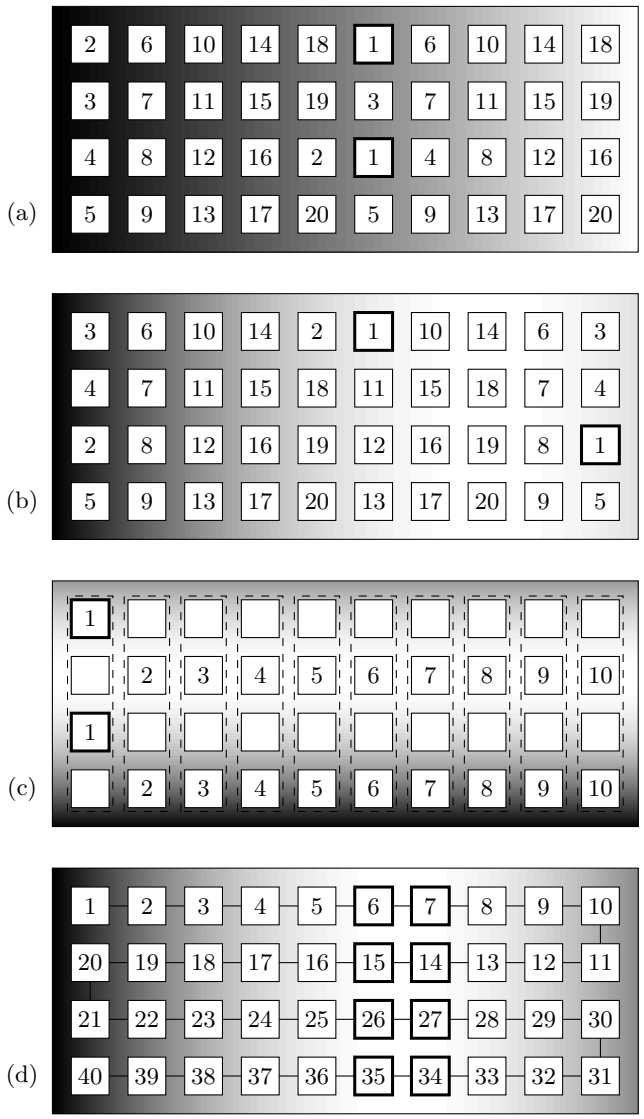

Fig. 6. Attacking entropy distiller constructions. (a) Group based RO PUF. (b) Group-based RO PUF. (c) 1-out-of-k masking. (d) neighbor pairing.

Suppose now that the ROs of G_1 already belonged to the same group for the original partitioning. Then their frequency order is of direct interest, as their response bit r_1 does influence a subkey. Consider the two hypotheses shown below. We compute a corresponding set of ECC helper data for both. The failure rate of the correct hypothesis is expected to be lower. Injecting additional errors is straightforward: we just compute the ECC redundancy given some inverted bit values.

$$H_0 : r_1 = 0. \quad H_1 : r_1 = 1.$$

6.4 Entropy Distiller with RO Pairing

Consider an entropy distiller being employed with a RO pairing scheme of section 4. We limit ourselves to chain of neighbors and 1-out-of-k masking here, as there is a 'stand-alone' attack for the latter two schemes. The attack methodology is similar as before, although there is limited and zero flexibility for picking groups (pairs) now respectively. Figure 6c provides an illustration for 1-out-of-k masking, given $k = 4$: again, only r_1 is determined by random variations. For neighbor pairing, it might be very difficult to isolate a single response bit. In figure 6d for instance, four response bits are fully determined by random variations. By increase the number of hypotheses, one can still perform the attack however.

7 Discussion

Former helper data constructions have all been proposed to solve reliability and entropy issues, as discussed in section 3. However, a well-established standard solution is available as well: the so-called fuzzy extractor [2]. We briefly discuss its architecture. Afterwards, we argue why helper data should be considered as public always, implicating that an attacker has both read and write access. Finally, we formulate best practices for both the users and developers of helper data schemes.

7.1 Fuzzy Extractor

Fuzzy extractors can be used with any PUF architecture: their use is not limited to RO PUFs. Their definition is very generic, but typical implementations always rely on an ECC and a cryptographic hash function, as shown in figure 7. Latter constructions deal with reliability and entropy respectively, in a sequential manner.

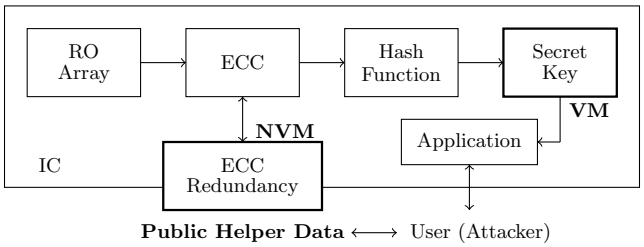

Fig. 7. RO PUF with fuzzy extractor.

7.2 Helper Data Considerations

In principle, programmable helper NVM can be implemented off-chip as well as on-chip. Therefore, we have drawn it on the IC boundary in figures 4 and 7. In the off-chip case, an attacker has full control: reading and modifying data is straightforward via the interface. However, also in the on-chip case, memory contents should be considered as public. Remember that PUFs have been proposed as a more

secure alternative for on-chip NVM. Labelling on-chip helper NVM as private would undermine the need for PUFs. A motivated attacker, able to afford expensive equipment, can still obtain read/write access.

Furthermore, off-chip helper NVM is highly preferable because of three major reasons. First, the overall efficiency might decrease in the opposite case: on-chip NVM remains while the PUF and its post-processing logic are extra. Second, on-chip NVM is expensive, as the standard CMOS manufacturing flow is insufficient. Third, typical FPGA platforms, for which RO PUFs are particularly interesting, do not contain on-chip NVM.

Secure and competitive PUF solutions do not pose read or write constraints on their helper data. For the fuzzy extractor, solid theory has been developed. Under certain assumptions, the ECC constructions of [2] do not leak information about the secret key. An extension of the architecture to prevent manipulation attacks is described in [1]. The RO PUF constructions under attack [7–10] do consider leakage as a threat. Manipulation is never mentioned however, although their prototypes are all developed on FPGA platforms without on-chip NVM (Xilinx Spartan-3 and Xilinx XC4010XL).

7.3 Best Practices

We encourage the use of fuzzy extractors, as solid helper data theory has been developed. New post-processing proposals should always be compared to this common reference. If efficiency (area, speed, power/energy, memory) and/or security (quality of the key, helper data leakage and manipulation, side-channels etc.) is not expected to improve, there is little argumentation to promote their use.

However, a thorough comparison is generally lacking: we strive for better practices in this regard. Sometimes the fuzzy extractor's existence is not even mentioned, as for the group-based RO PUF and the entropy distiller proposal for instance. Note that the former construction actually borrows its ECC notion. We strongly question the hardware efficiency of many proposals. Consider the collective overhead of group-based RO PUFs for instance (ECC excluded): it might surpass the requirements for a cryptographic hash function. Or consider temperature-aware cooperative RO PUFs, having severe applicability issues. Besides a temperature sensor, one does require an extension of the IC manufacturing flow (see [7]: measuring RO frequencies, disconnecting bad RO pairs from power supply, etc.).

Furthermore, many proposals are rather vague about their use of helper data. The precise storage format, parsing procedure and/or sanity checks are typically not specified. Although subtle differences might impact security tremendously. We provide a few examples and strive again for better practices. For both 1-out-of-k masking and the sequential pairing algorithm, pairs of RO indices are stored. However, there is no recommendation to store a pair's indices in an either randomized or sorted order. Otherwise there is direct leakage of the full key. The re-use of RO indices should also be prohibited somehow. For group-based PUFs, it is not clear whether grouping helper data is transferred three times or only once, with the former case offering more opportunities for an attacker.

8 Conclusion and Further Work

Like any other PUF, RO PUFs do require post-processing logic to generate reproducible and uniformly distributed secret keys. However, we showed various constructions to be vulnerable against manipulation of their public helper data. By observing system failure rates, an attacker can retrieve the key, or at least obtain some information about it. Actually, many more helper data constructions have been proposed in literature, not necessarily limited to the RO PUF. We do not claim to have studied them all and we advise to use them with great care. Instead, we encourage the use of fuzzy extractors, the well-established reference solution. We strive for better practices when proposing new helper data schemes. The following two items should be present: (1) an all-inclusive comparison with the reference solution and (2) a very precise specification of its helper data use.

Acknowledgment

This work was supported in part by the European Commission through the ICT programme under contract FP7-ICT-2011-317930 HINT. In addition this work is supported by the Research Council of KU Leuven: GOA TENSE (GOA/11/007), by the Flemish Government through FWO G.0550.12N and the Hercules Foundation AKUL/11/19. Jeroen Delvaux is funded by IWT-Flanders grant no. 121552.

References

1. X. Boyen, Y. Dodis, J. Katz, R. Ostrovsky and A. Smith, "Secure Remote Authentication Using Biometric Data," in *Eurocrypt*, pp. 147-163, May 2005.
2. Y. Dodis, R. Ostrovsky, L. Reyzin and A. Smith, "Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data," *SIAM J. Comput.*, vol. 38, no. 1, pp. 97-139, Mar. 2008.
3. A. Maiti and P. Schaumont, "Improving the Quality of a Physical Unclonable Function Using Configurable Ring Oscillators," in *Field Programmable Logic and Applications*, FPL 2009, pp. 703-707, Aug. 2009.
4. P. Sedcole and P.Y.K. Cheung, "Within-die delay variability in 90nm fpgas and beyond, in *Field Programmable Technology*, FPT 2006, pp. 97-104, Dec. 2006.
5. S. Skorobogatov, "Semi-invasive attacks - a new approach to hardware security analysis, Technical Report UCAM-CL-TR-630, University of Cambridge, Computer Laboratory, Apr. 2005.
6. G.E. Suh and S. Devadas, "Physical unclonable functions for device authentication and secret key generation," in *Design Automation Conference*, DAC 2007, pp. 9-14, Jun. 2007.
7. C.E. Yin and G. Qu, "Temperature-aware cooperative ring oscillator PUF," in *Hardware-Oriented Security and Trust*, HOST 2009, pp. 36-42, Jul. 2009.
8. C.E. Yin and G. Qu, "Lisa: Maximizing RO PUF's Secret Extraction," in *Hardware Oriented Security and Trust*, HOST 2010, pp. 100-105, Jun. 2010.
9. C.E. Yin, G. Qu and Q. Zhou, "Design and implementation of a group-based RO PUF," in *Design, Automation & Test in Europe*, DATE 2013, pp. 416-421, Mar. 2013.
10. C.E. Yin and G. Qu, "Improving PUF security with regression-based distiller," in *Design Automation Conference*, DAC 2013, pp. 1-6, May 2013.