



Article

Analysis of Entropy in a Hardware-Embedded Delay PUF

Wenjie Che ^{1,*}, Venkata K. Kajuluri ¹, Mitchell Martin ¹, Fareena Saqib ² and Jim Plusquellic ^{1,*}¹ Department of Electrical and Computer Engineering, University of New Mexico, Albuquerque, NM 87131, USA; vkishorekajuluri@unm.edu (V.K.K.); mmartin3@unm.edu (M.M.)² Department of Electrical and Computer Engineering, Florida Institute of Technology, Melbourne, FL 32901, USA; fsaqib@fit.edu

* Correspondence: wjche@unm.edu (W.C.); jimp@ece.unm.edu (J.P.); Tel.: +1-505-277-0785 (J.P.)

Academic Editor: Kwangjo Kim

Received: 27 February 2017; Accepted: 2 June 2017; Published: 7 June 2017

Abstract: The magnitude of the information content associated with a particular implementation of a Physical Unclonable Function (PUF) is critically important for security and trust in emerging Internet of Things (IoT) applications. Authentication, in particular, requires the PUF to produce a very large number of challenge-response-pairs (CRPs) and, of even greater importance, requires the PUF to be resistant to adversarial attacks that attempt to model and clone the PUF (model-building attacks). Entropy is critically important to the model-building resistance of the PUF. A variety of metrics have been proposed for reporting Entropy, each measuring the randomness of information embedded within PUF-generated bitstrings. In this paper, we report the Entropy, MinEntropy, conditional MinEntropy, Interchip hamming distance and National Institute of Standards and Technology (NIST) statistical test results using bitstrings generated by a Hardware-Embedded Delay PUF called HELP. The bitstrings are generated from data collected in hardware experiments on 500 copies of HELP implemented on a set of Xilinx Zynq 7020 SoC Field Programmable Gate Arrays (FPGAs) subjected to industrial-level temperature and voltage conditions. Special test cases are constructed which purposely create worst case correlations for bitstring generation. Our results show that the processes proposed within HELP to generate bitstrings add significantly to their Entropy, and show that classical re-use of PUF components, e.g., path delays, does not result in large Entropy losses commonly reported for other PUF architectures.

Keywords: entropy analysis; physical unclonable function; authentication

1. Introduction

The number of independent sources of information used to distinguish a system is a measure of its complexity, and relates to the amount of effort required to copy or clone it. The relationship between complexity and effort can be exponential, particularly for systems designed to conceal or mask the information and only provide *controlled* access to it. A physical unclonable function (PUF) is an information system that can meet these criteria under certain conditions. The information embedded in a PUF is random, enabling it to serve hardware security and trust roles related to key generation, key management, tamper detection and authentication [1]. PUFs represent an alternative to storing keys in non-volatile-memory (NVM), thereby reducing cost and hardening the embedding system against key-extraction-based attacks. PUFs are widely recognized as next-generation security and trust primitives that are ideally suited for authentication in industrial, automotive, consumer and military IoT-based systems, and for dealing with many of the challenges related to counterfeits in the supply chain.

PUFs enable access to their stored random information using a challenge-response-pair (CRP) mechanism, whereby a server or adversary ‘asks a question’ usually in the form of a digital bitstring and the PUF produces a digital response after measuring a set of circuit parameters within the chip. The nanometer size of the integrated circuit (IC) features and the analog nature of stored information makes it extremely difficult to read out the information using alternative access mechanisms. The circuit parameters that are measured vary from one copy of the chip to another, and can only be controlled to a small, but non-zero, level of tolerance by the chip manufacturer. This feature of the PUF makes it unclonable and provides each copy of the chip with a distinct ‘personality’, in the spirit of fingerprints or DNA for biological systems.

Strong PUFs are a special class of PUFs that are distinguished from *weak PUFs* by the amount of information content they possess. The traditional definition for distinguishing between weak and strong PUFs is to consider only the number of CRPs that can be applied. For weak PUFs, the number of CRPs is polynomial while strong PUFs have an exponential number, e.g., the number of challenges for an n -binary-input weak PUF can be n^2 while a strong PUF typically has 2^n . Unfortunately, this traditional definition leads to a misnomer as to the true strength of the PUF to adversary attacks. For example, the original Arbiter PUF [2,3] is classified as strong even though machine-learning-based model-building attacks have shown that only a small, polynomial, number of CRPs are needed to predict its complete behavior.

Therefore, a truly strong PUF must have both an exponential number of CRPs and an exponential number of unique, uncorrelated responses, i.e., a large input challenge space is necessary but is not a sufficient condition. This requires the PUF to have access to a large source of entropy, either in the form of IC features from which random information is extracted, or in an artificial form using a cryptographic primitive, such as a secure hash function. Either mechanism makes the PUF resilient to machine learning attacks. However, using a secure hash for expanding the CRP space of the PUF and for obfuscating its responses consumes additional area and increases the required reliability of the PUF. Therefore, the former scenario, i.e., a large source of entropy, is more attractive but more difficult to achieve.

In this paper, we present results that support this more attractive alternative using a hardware-embedded delay PUF called HELP. HELP generates bitstrings from delay variations that occur along paths in an on-chip macro, i.e., the source of entropy for HELP is within-die manufacturing process variations that cause path delays to be slightly different in each copy of the chip. Macros or functional units that implement cryptographic algorithms and common data path operators such as multipliers typically possess at least 32 inputs and therefore, HELP meets the large input space requirement of a strong PUF.

Moreover, the wire interconnectivity within the macro used by HELP provides a large number of testable paths, on order of 2^n for n inputs, satisfying the large output space requirement of a strong PUF. Unlike other PUFs that meet these conditions, the task of generating input test sequences (challenges) that test all of the testable paths is an NP-complete problem. Although this may appear to be a drawback, it, in fact, makes the task of model-building HELP much more difficult. For example, the adversary not only must devise a machine learning strategy that is able to predict output responses, but he/she must also expend a large effort on generating the challenges, which is typically accomplished using automatic test pattern generation (ATPG) algorithms. Note that these characteristics of HELP, namely, the use of a functional unit as a source of Entropy, paths of arbitrary length and the ATPG requirement, distinguish HELP from other delay-based PUFs such as the Arbiter and Ring Oscillator (RO) PUFs.

This paper investigates entropy of HELP using 500 instances of a functional unit (the entropy source) embedded on a set of 20 Xilinx Zynq 7020 Field Programmable Gate Arrays (FPGAs). The specific contributions of this paper include the following:

- Strong experimental evidence that HELP leverages within-die variations (WDV) almost exclusively as its source of entropy.
- A statistical evaluation of Entropy, MinEntropy, conditional MinEntropy, Interchip hamming distance and NIST statistical test results on hardware generated bitstrings.

- A special worst-case analysis that maximizes correlations and dependencies introduced by (1) **full path reuse** and (2) **partial path reuse** where the *same* paths in different combinations or paths with many common segments are used to generate distinct bits.

The rest of this paper is organized as follows. Related work is presented in Section 2 and an overview of HELP is given in Section 3. Statistical results are described in Section 4 using FPGA-based path delay data and bitstrings. A worst-case correlation analysis is presented in Section 5 and conclusions in Section 6.

2. Related Work

The source of random information varies widely among proposed PUF architectures, and includes transistor threshold voltages [4], delay chains and ring oscillators (RO) [2–6], FPGAs [7,8], SRAMs [9], leakage current [10], metal resistance [11], transistor transconductance [12], the path delays of core logic macros [13–15], memristors [16], scan chains [17], phase change memory [18], plus many others.

One of the earliest delay-based PUFs, called the Arbiter PUF, uses n -bit differential delay lines and a latch to generate a 1-bit PUF response [19,20]. Because of the limited amount of entropy, model-building attacks are effective against the Arbiter PUFs [21]. Ring Oscillator (RO) PUF [22] measure the frequency difference between two identical ring oscillators by counting the transitions on the output of each RO and then comparing counter values to generate a PUF bit. The number of challenges is limited to the number of pairings (n^2) and therefore the RO PUF is a weak PUF. The authors of [23] analyze RO frequency differences, selecting those pairings where the frequency difference is large enough to avoid any bit flip errors caused by environmental variations. The authors of [24] propose a scheme to produce $(n - 1)$ reliable bits, and Ref. [25] proposes a longest increasing subsequence-based grouping algorithm (LISA) for FPGAs that sequentially pairs RO-PUF bits and can generate $n/2$ reliable bits out of n ring oscillators. In [26], the authors propose a regression based distiller to remove systematic variations.

PUF responses are affected by the environmental variations such as temperature and voltage variations, thus processing is required to extract the entropy from the noise. Several schemes including helper data and fuzzy extractor schemes are proposed to improve the reliability of bitstring regeneration and improve randomness [27]. Helper data is generated during the enrollment phase, which is carried out in a secure environment and is later used with the noisy responses during regeneration to reconstruct the key. Bosch et al. [28] demonstrated a hardware implementation of concatenated codes based fuzzy extractors that have been used to produce bitstrings with high reliability. Reference [29] discusses a fuzzy extractor scheme based on repetition codes that can limit the usable entropy and show that such a scheme is not applicable to the PUFs with small entropy. Dodis et al. [30] provided a formal definition and analysis of entropy loss in fuzzy extractors. The authors of [31] evaluated the reliability and unpredictability properties of five different types of PUFs (Arbiter, RO, SRAM, flip-flop and latch PUFs) from an Application-specific integrated circuit (ASIC) implementation.

3. HELP Overview

HELP attaches to an on-chip functional unit, such as a portion of the Advanced Encryption Standard (AES) labeled *sbox-mixedcol* on the left side of Figure 1. The logic gate structure of the functional unit defines a complex interconnected network of wires and transistors. This combinational data path component includes 64 primary inputs (PIs) and 64 primary outputs (POs) and is implemented in Wave Dynamic Differential Logic (WDDL) logic-style [32] on a Xilinx Zynq FPGA using approx. 2900 LUTs and 30 K wire segments.

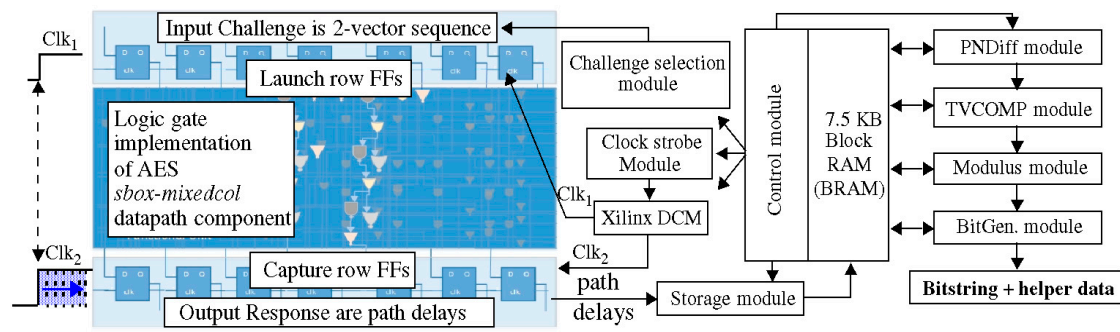


Figure 1. Instantiation of the HELP entropy source (left) and HELP processing engine (right).

Path delay is defined as the amount of time (Δt) it takes for a set of 0-to-1 and 1-to-0 bit transitions introduced on the PIs of the functional unit (input challenge) to propagate through the logic gate network and emerge on a PO. HELP uses a clock-strobing technique to obtain high resolution measurements of path delays as shown on the left side of Figure 1. A series of launch-capture operations are applied in which the vector sequence that defines the input challenge is applied repeatedly to the PIs using the Launch row flip-flops (FFs) and the output responses are measured on the POs using the Capture row FFs. On each application, the phase of the capture clock, Clk_2 , is incremented forward with respect to Clk_1 , by small Δt s (approx. 18 ps), until the emerging signal transition on a PO is successfully captured in the Capture row FFs. A set of XOR gates connected to the Capture row FF inputs and outputs (not shown) provide a simple means of determining when this occurs. When an XOR gate value becomes 0, then the input and output of the FF are the same (indicating a successful capture). The first occurrence in which this occurs during the clock strobe sweep causes the current phase shift value to be recorded as the digitized delay value for this path. The current phase shift value is referred to as the launch-capture-interval (LCI). The *Clock strobe* module is shown in the center portion of Figure 1, which utilizes features on Xilinx Digital Clock Manager (DCM).

The digitized path delays are collected by a *storage* module and stored in an on-chip block RAM (BRAM) as shown in the center of Figure 1. Each digitized timing value is stored as a 16-bit value, with 12 binary digits serving to cover a signed range between ± 2048 and 4 binary digits of fixed point precision to enable up to 16 samples of each path delay to be measured and averaged. The digitized path delays are stored in the upper half of the 16 KByte BRAM. We configure the applied challenges to test 2048 paths with rising transitions and 2048 paths with falling transitions. The digitized path delays are referred to as PUFNums, or PN, with PNR used to refer to rising path delays and PNF for falling. Once a set of 4096 PN are collected, a sequence of operations implemented in VHDL are started to produce the bitstring and helper data, as shown on the far right of Figure 1. These operations are described below.

3.1. Implementation Details

We created 25 instances of *sbox-mixedcol* on each of 20 chips, for a total of 500 implementations (25 separate programming bitstreams are generated). Figure 2 shows a screen snapshot of Xilinx Vivado Implementation view, which depicts a completed instance of the functional unit in the lower right corner (labeled as *instance₁*). The VHDL code for *sbox-mixedcol* is synthesized and implemented into a *pblock*, which is shown as a magenta rectangle surrounding *instance₁*. Once completed, *tcl* commands are issued that save a set of constraints for the wire and LUT components of the functional unit to a file called a *check-point*. The base *y* coordinate of the *pblock* is then incremented by 3 to create a sequence of *pblock* implementations, each of which is synthesized into a separate bitstream. In this fashion, a sequence of identical and overlapping *pblock* instances of the functional unit are created and tested, one at a time. The rationale for doing this is two-fold. First, it increases the statistical significance of the analysis without requiring a corresponding increase in the number of chips. Second,

data from overlapping instances on the same chip implicitly eliminate chip-to-chip process variations, and provide a basis on which we can prove experimentally that HELP leverages within-die variations almost exclusively.




Figure 2. sbx-mixedcol functional unit instance placement in Xilinx Zynq 7020 using Vivado implementation view.

3.2. PN, PND and PND_c Processing Steps

The PN processing operations shown on the far right in Figure 1 are designed to eliminate both chip-to-chip performance differences and environmental variations, while leaving only within-die variations as a source of entropy for HELP. In order to accomplish this, the following modules and operations are defined. The *PNDiff* module creates unique, pseudo-random pairings between elements of the PNR and PNF groups using two seeded linear feedback shift registers (LFSR). The LFSRs are used to generate 11-bit addresses to access any of the 2048 PNR and PNF values. The two 11-bit *LFSR seeds* are configuration parameters. The PN differences are referred to as **PND**. The primary reason for creating PND is to increase the magnitude of within-die variations, i.e., path delay variations are doubled (in the best case) over those available in the PNR and PNF.

Figure 3a shows an example of this process using a pairing of paths from the PNR and PNF sets. The graph contains curves for 500 PNR and 500 PNF, one for each of the 500 chip-instances. Although it is difficult to distinguish between the two groups in the figure, the PNF have a larger delay and are displayed above the PNR. The 13 line-connected points in each curve represent the PN measured under a range of environmental conditions, called temperature-voltage (TV) corners. The PN at the *x*-axis position given by 0 are those measured under nominal conditions (referred to as **enrollment** values below), i.e., at 25 °C, 1.00 V. The PN at positions 1, 2 and 3 are also measured at 25 °C but at supply voltages of 0.95, 1.00 and 1.05 V. Similarly, the other groups of three consecutive points along the *x*-axis are measured at these supply voltages but at temperatures 0 °C, −40 °C and 85 °C. The PN measured under TV corners numbered 1 to 12 are referred to as **regeneration** PN. Figure 3b plots the PND defined by subtracting pointwise, each PNF from a PNR for each chip-instance.

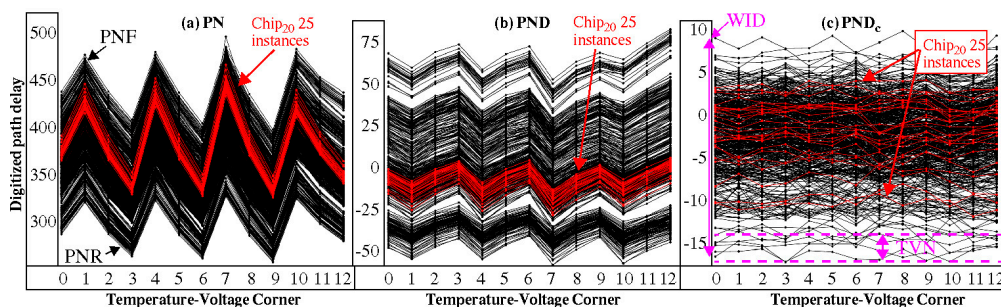


Figure 3. (a) best-case set of rising and falling path delays (PNs); (b) (rise delay - fall delay) (PND) and (c) TV Compensated PND (PND_c).

TV-related effects on delay negatively impact bitstring reproducibility. It is clear that subtraction alone, which is used to create the PND, is not effective at removing all of the variations introduced by different environmental conditions (if it was, the curves would be horizontal lines).

$$zval_i = \frac{(PND_i - \mu_{TV_x})}{Rng_{TV_x}}. \quad (1)$$

We propose a *TV compensation* (TVCOMP) process that is applied to the PND as a mechanism to eliminate most of the remaining temperature-voltage variations (called TV-noise).

TVCOMP is applied to the entire set of 2048 PND measured for each chip-instance at each of the 13 TV corners separately (note, Figure 3b shows only one of the PND from the larger set of 2048 that exist for each chip-instance and TV corner). The TVCOMP procedure first converts the PND to ‘standardized’ values. Equation (1) represents the first transformation, which makes use of two constants, i.e., μ_{chip} (mean) and Rng_{chip} (range), obtained by measuring the mean and range of the distribution defined by the PND. The second transformation is represented by Equation (2), which translates the standardized $zvals$ to a new distribution with mean μ_{ref} and range Rng_{ref} . The **reference** mean and range values are also configuration parameters. In our experiments, we fix μ_{ref} and Rng_{ref} in the TVCOMP operation for all chip-instances as a means of eliminating chip-to-chip performance differences.

$$PND_c = zval_i Rng_{ref} + \mu_{ref}. \quad (2)$$

Figure 3c illustrates the effect of TVCOMP under these conditions. The PND_c (‘c’ for compensated) plotted in the graph are obtained by applying the TVCOMP procedure to the 2048 PND measured under each of the 13 TV corners for each chip, i.e., 13 TV corners \times 500 chip-instances = 6500 separate applications. Several features of TVCOMP are evident. First, the transformation significantly reduces TV-noise which is evident by the flatter curves (note that the scale used on the y -axis is amplified over that shown in Figure 3b). Second, global (chip-wide) performance differences are also nearly eliminated between the chip-instances, leaving only within-die variations. This is illustrated nicely by the highlighted red curves (25 instances) for chip₂₀. The curves shown in Figure 3a,b for the 25 instances on chip₂₀ are grouped together, illustrating that these instances have similar performance characteristics as expected, since they are obtained from the same chip. However, the corresponding curves in Figure 3c are distributed across most of the y -range, and are indistinguishable from the 450 curves from the other 19 chip-instances. The dispersion of the chip₂₀ curves across the entire range illustrates that the random information leveraged by HELP is based on within-die variations (WDV), and not on global performance differences that occur from chip-to-chip.

The differences that remain in the PND_c are those introduced by WDV and *uncompensated* TV noise (TVN). The range of TVN for the bottom-most curve in Figure 3c is labeled and is approx. 3, which translates to approx. 90 ps. In general, PND_c with larger amounts of TVN are more likely to introduce bit flip errors. Therefore, it is desirable to make TVN as small as possible, and is the main driver for using the TVCOMP process.

The last operation applied to the PNs is represented by the *Modulus* operation shown on the right side of Figure 1. Modulus is a standard mathematical operation that computes the positive remainder after dividing by the modulus. The Modulus operation is required by HELP to eliminate the path length bias that exists in the PND_c , which acts to reduce randomness and uniqueness in the generated bitstrings. The value of the Modulus is also a configuration parameter, similar to the LFSR seed, μ_{ref} and Rng_{ref} parameters, and is discussed further in the following. The term **mod** PND_c is used to refer to the values used in the bitstring generation process.

3.3. Offset Method

An optional *offset* can also be applied to PND_c values prior to the application of the Modulus to further improve the statistical quality of the bitstrings. An offset is computed for each PND_c separately

in a characterization process. The offset is simply the *median value* of the PND_c , derived using PNs from a sample of chips or from a nominal simulation. The offsets are transmitted to the token and are therefore a second component of the challenges. The token adds the individual offsets to each of the PND_c as they are generated. The offset shifts the PND_c upwards and *centers the population* over the 0–1 line associated with the Modulus. We use the term PND_{co} to refer to the PND_c with offsets applied. Since the offset is a population-based value, it leaks no information regarding the bit values generated from the $modPND_{co}$ (to be discussed).

As an example, three randomly selected PND_c are shown in Figure 4. The PND_c from the 500 chip-instances are given on the left in the same format as that used in Figure 3c, while the corresponding ‘shifted’ PND_{co} are shown to their immediate right. The 0–1 lines associated with a Modulus of 24 are superimposed as dashed horizontal lines. The Modulus creates vertical partitions of size 24, with 0–1 lines at Modulus/2 and Modulus. The corresponding bit values for each region are shown on the far right.

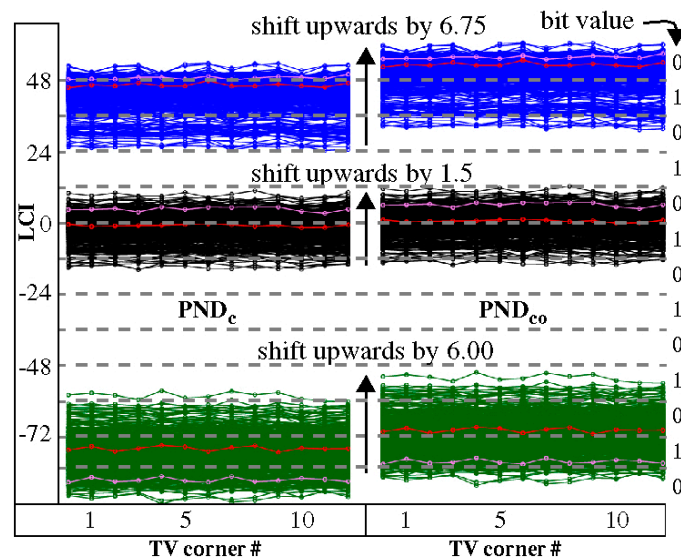


Figure 4. Three example PND_c from 500 chip-instances (y -axis) at each of the 13 TV corners. PNs before 4-bit offset is added (**left**) and afterwards (**right**) using a Modulus of 24. Dashed lines identify 0–1 lines, with corresponding bit values associated with each region shown on the far right. Two chip-instances are highlighted as red and magenta to illustrate their random occurrence among different sets of PND_c , which is caused by within-die variation effects.

The shift amounts are shown between the two sets of waveforms. The centering of the population over the 0–1 lines ensures that nearly equal numbers of chips produce 0 s and 1 s for each of the corresponding PND_{co} . We restrict the offset encoding to 4 bits, making it possible to shift the population by $Modulus/(2 \times 16)$. The additional factor of 2 in the denominator accounts for the fact that the maximum shift required to reach one of the 0–1 lines is half the Modulus.

3.4. Margining

A *Margin* technique is used to improve reliability by identifying and excluding bits that have the highest probability of ‘flipping’ from 0 to 1 or 1 to 0. As an illustration, Figure 5 plots 18 of the 2048 $modPND_{co}$ from $Chip_1$ along the x -axis. The red curve line-connects the data points obtained under enrollment conditions while the black curves line-connects data points under the 12 regeneration TV corners. A set of margins are shown of size 2 surrounding two strong bit regions of size 8. Designators along the top given as ‘s0’, ‘s1’, ‘w0’ and ‘w1’ classify each of the enrollment data points as either a strong 0 or 1, or a weak 0 or 1, resp. Data points that fall on or within the hatched areas are classified

as weak as a mechanism to avoid bit flip errors introduced by uncompensated TV noise (TVN) that occurs during regeneration.

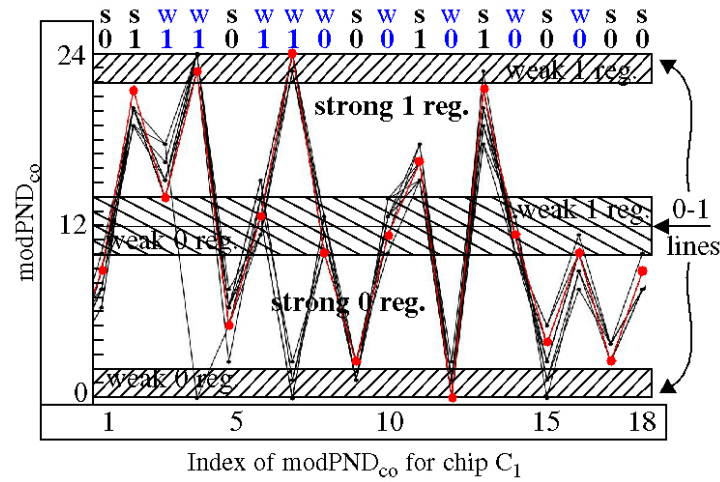


Figure 5. Strong/Weak TVCOMP modPNDco classification using margining.

The Margin method improves bitstring reproducibility by eliminating data points classified as ‘weak’ in the bitstring generation process. For example, the data points at indexes 4, 6, 7, 8, 10 and 14 would introduce bit flip errors at one or more of the TV corners during regeneration because at least one of the regeneration data points is in the opposite bit value region, i.e., they cross one of the annotated 0–1 lines, from the corresponding enrollment value. A *helper data* string is constructed during enrollment that records the strong/weak status of each modPNDco, which is used during regeneration to identify which modPNDco generate bits (strong) and which are skipped (weak).

4. Statistical Results

4.1. Entropy Analysis

The statistical analysis is carried out using the bitstrings generated from the 500 chip-instances. Entropy is defined by Equation (3) and MinEntropy by Equation (4). The frequency p_{ij} of ‘0’ s and ‘1’ s is computed at each bit position i across the 500 chip-instance bitstrings of size 2048 bits, i.e., no Margin is used in this analysis.

$$H(X) = \sum_{i=1}^{2048} \left(- \sum_{j=0}^1 p_{ij} \cdot \log_2 p_{ij} \right), \quad (3)$$

$$H_{\infty}(X) = \sum_{i=1}^{2048} (-\log_2(\max(p_{ij}))). \quad (4)$$

Figure 6 plots incremental Entropy and MinEntropy for both the original modPNDco and the 4-bit offset technique using black and blue curves, resp., as chip-instances are added, one at a time, to the analysis (a similar analysis is presented in [33]). The x -axis gives the index of the chip-instance starting with two chip-instances on the left and ending with 500 chip-instances on the right. The 4-bit offset technique shifts and centers the population of chip-instances associated with each modPNDc over a 0–1 line as discussed in Section 3.3. The centering has a significant impact on Entropy and MinEntropy, which is reflected in the larger values and the gradual approach of the curves to the ideal value of 2048 as chip-instances are added.

Figure 7a,b depict bar graphs of Entropy and MinEntropy for Moduli 10 through 30 (x -axis). The height of the bars represents the average values computed using the 2048-bit bitstrings from 500

chip-instances, averaged across 10 separate LFSR seeds. Entropy varies from 2037 to 2043, and is close to the ideal value of 2048 independent of the Moduli. MinEntropy varies between 1862 at Moduli 12 up to 1919, which indicates that, in the worst case, each bit contributes between 91% and 93.7% bits of Entropy.

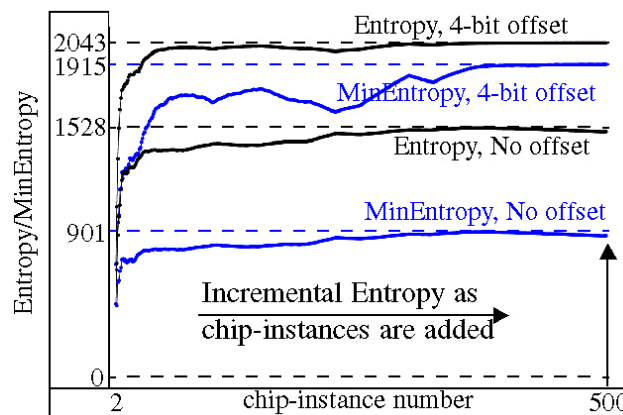


Figure 6. Entropy (black) and MinEntropy (blue) change as chips are added to the analysis along the x-axis. Maximum value is 2048 bits. Top curves show results using 4-bit offset while lower curves show analysis with no offset using a Modulus of 24 and Mean scaling.




Figure 7. (a) Entropy; (b) MinEntropy and (c) InterChip hamming distance (HD) computed as average values across 10 seeds of 2048 bits each using 500 chip-instances with TVCOMP reference set to ‘Mean’ scaling. Bars of zero height for InterChip HD are invalid combinations of Margin and Modulus. Entropy varies over the range 2037 to 2043, and MinEntropy from 1862 to 1919, with 2048 as the ideal value. InterChip HD varies from 49.4% to 51.2% with ideal at 50%.

4.2. Uniqueness

The InterChip hamming distance (InterChipHD) results are shown in Figure 7c, again computed using the bitstrings from 500 chip-instances, averaged across 10 separate LFSR seed pairs. Hamming distance is computed between all possible pairings of bitstrings, i.e., $500 \times 499/2 = 124,750$ pairings for each seed and then averaged.

The values for a set of Margins of size 2 through 4 (y-axis) are shown for each of the Moduli. Figure 8 provides an illustration of the process used for dealing with weak and strong bits under the Margin scheme in the InterchipHD calculation. The helper data bitstrings HelpD and raw bitstrings BitStr for two chips Cx and Cy are shown along the top and bottom of the figure, resp. The HelpD bitstrings classify the corresponding raw bit as weak using a ‘0’ and as strong using a ‘1’. The InterchipHD is computed by XOR’ing only those BitStr bits from the Cx and Cy that have both HelpD bits set to ‘1’, i.e., both raw bits are classified as strong. This process maintains alignment in the two bitstrings and ensures the same modPNDc from Cx and Cy are being used in the InterchipHD calculation.

InterChip HD, HD_{Inter} , is computed using Equation (5). The symbols NC , NB_a and NCC represent ‘number of chips’, ‘number of bits’ and ‘number of chip combinations’, resp. (NCC is 124,750 as

indicated above) This equation simply sums all the bitwise differences between each of the possible pairing of chip-instance bitstrings BS as described above and then converts the sum into a percentage by dividing by the total number of bits that were examined. *Bit cnter* from the center of Figure 8 counts the number of bits that are used for NB_a in Equation (5), which varies for each pairing of chip-instances a . The HD_{Inter} is computed separately for each of the 10 seeds and the average value is given in Figure 7c. The HD_{inter} vary from 49.4% to 51.2% and therefore are close to the ideal value of 50%.

$$HD_{inter} = \left(\frac{1}{NCC} \cdot \sum_{i=1}^{NC} \sum_{j=i+1}^{NC} \frac{\left(\sum_{k=1}^{NB_a} (BS_{i,k} \oplus BS_{j,k}) \right)}{NB_a} \right) \times 100 \quad (5)$$




Figure 8. Hamming distance illustration for results shown in Figure 7.

4.3. NIST Test Evaluation

The NIST statistical test suite is used to evaluate randomness of the bitstrings [34]. The bitstrings are constructed as described above for Interchip HD. All tests are passed with at least 488 bitstrings passing of the 500 bitstrings as required by NIST except for CumulativeSums (NIST test #4) under two Moduli. The two failing cases failed with 487 and 482 bitstrings passing, resp., so the failures were only by at most six chips in the worst case.

5. Correlation Analysis

Correlation analysis measures whether a relationship exists between modPND_{co} in which the bit response from one allows the response from a second to be predicted with probability greater than 50%. All strong PUF architectures to date have the potential to exhibit correlation because the 2^n response bits are generated from a much smaller set of m components, with the m components representing the underlying random variables. For the case of a 64-stage Arbiter PUF, the 256 path segments are all reused in every challenge, and therefore, the potential for correlation introduced by *path segment reuse* is very high. HELP also reuses path segments, but the probability of two paths sharing a large number of path segments is very small. The following analysis focuses on the reuse of path segments within HELP despite the fact that, in practice, it is statistically rare.

Our correlation analysis of *path segment reuse* (called **Partial Reuse**) is carried out using a set of ‘unique’ paths, and therefore, it ensures that at least one path segment is different in any pairing of PN used to create PND , PND_c , PND_{co} and modPND_c (note: we refer to PND_c in the following because the analysis focuses on how the Offset and Modulus operations affect the results). An example of partial reuse is shown in Figure 9. The highlighted red wire on the left indicates that the two paths, labeled ‘path #1’ and ‘path #2’, share all of the initial path segments, and are only different at the fanout point where they diverge into LUT_a and LUT_b . The two paths then reconverge at the next gate and form a ‘bubble’ structure.

It is also possible to pair the **same** PNs in different combinations to produce a much larger set of PND_c (on order of n^2 with n PNs). We refer to this as **Full Reuse**. Full path reuse can result in *dependent bits*, i.e. bits that are completely determined by other bits. Reference [25] investigates these dependencies for ROs and proposes schemes designed to eliminate and/or reduce the number of dependent bits.




Figure 9. Hamming distance illustration for results shown in Figure 7. Reuse worst-case example of two paths forming a ‘bubble’. The path segments which define the bubble are unique to each path while the remaining components are common to both paths.

We show in the following that the Offset and Modulus operations break the correlations found in classic dependency analysis typically exemplified using RO frequencies as $f(\text{RO}_A) > f(\text{RO}_B)$ and $f(\text{RO}_B) > f(\text{RO}_C)$ implies $f(\text{RO}_A) > f(\text{RO}_C)$. Therefore, partial reuse and full reuse of paths have a smaller penalty in terms of Entropy and MinEntropy when they occur within HELP.

5.1. Preliminaries

As indicated earlier, the HELP algorithm creates differences (PND) between PNR and PNF using a pair of LFSR seeds, which are then compensated using TVComp to produce PND_c . A key objective of our analysis is to purposely create worst case conditions for correlations by crafting the PND such that *partial reuse* and *full reuse* test cases are created. The analysis of correlations requires the set of PND that are constructed to be adjacent to each other in the arrays on which the analysis is performed. Therefore, the LFSRs used in the HELP algorithm are not used to create the PND and instead a linear, sequential pairing strategy is used.

The Offset and Modulus operations in the HELP algorithm are the key components to improving Entropy. As an aid to help with the discussion that follows, Figure 10 illustrates how these two operators modify the PND_c . The figure shows four groups of 10 vertical line graphs, with each line graph containing 500 PND_c data points corresponding to the 500 chip-instances. The line graph on the left and bottom illustrates that the vertical spread in the line-connected points is caused by within-die delay variations.

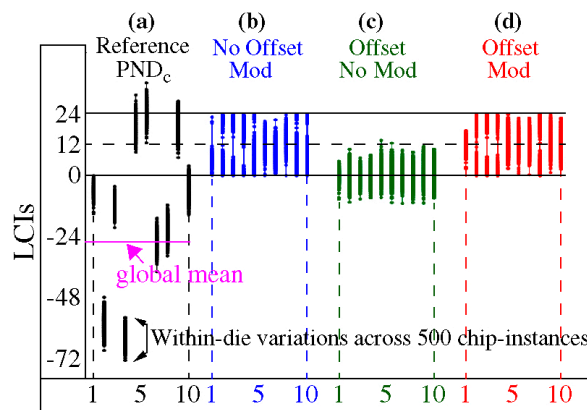


Figure 10. A sample of 10 PND_c from 500 chip-instances illustrating four experimental scenarios. (a) Reference represents PND_c that have no Offset or Modulus applied; Scenarios (b) No Offset; Mod (c) Offset; No Mod and (d) Offset, Mod show how the PND_c change under different combinations of these parameters.

The *Reference* PND_c shown on the left are the compensated differences before the Offset and Modulus operations are applied. The DC bias introduced by differences in the lengths of the paths changes the vertical positions of the line graphs, which spans a range from -72 to $+40$ launch-capture intervals (LCIs) (Recall that 1 LCI = 18 ps, and represents the phase adjustment resolution of the Xilinx

DCM.). The Offset and Modulus operations are designed to increase the Entropy in the PND_c by eliminating this bias. For example, the *No Offset, Mod* group show the PND_c from the *Reference PND_c* group after a Modulus of 24 is applied. Similarly, the *Offset, No Mod* group show the *Reference PND_c* after subtracting the median value from each line graph, which effectively centers the populations of 500 PND_{co} over the 0 horizontal line. Finally, the *Offset, Mod* group shows the PND_c with both operations applied, and represents the values used in the HELP algorithm. Here, an Offset is first applied to center the populations over the closest multiple of 12 and then a modulus of 24 is applied (the boundaries used to separate the '0' and '1' bit values are 12 and 24 for a Modulus of 24, see Figure 5). We analyze the change in Entropy and MinEntropy as each of the operations are applied. Note that HELP processes 2048 PND_c at a time during bitstring generation, of which only 10 are shown in Figure 10.

5.2. Partial Reuse

Although we defined *path segment reuse* above as a pair of PNs with at least one path segment that is different in a given PND_c , we do not want to restrict our analysis to these types of specific physical characteristics but instead want to analyze the actual worst case. The Xilinx Vivado implementation view does not provide information that directly reflects the chip layout, and therefore, a broader approach to correlation analysis is required to ensure the worst case correlations are found.

We use Pearson's correlation coefficient (PCC) [35] to measure the degree of correlation that exists among PND_c and then select a subset of the most highly correlated for Entropy and MinEntropy analyses. Figure 11 depicts the construction process used to create an exhaustive set of PND_c , from which the most highly correlated are identified. In order to simplify the construction process, the TVComp operation is applied to a set of 2048 PNR and 2048 PNF separately for each of the 500 chip-instances (HELP normally applies TVComp only once, and to the PND as discussed in Section 3.2, for processing efficiency reasons, but the results using either method are nearly identical.). Note the 'c' subscript is not used in the PNR/PNF designation for clarity. TVComp eliminates chip-to-chip delay variations and makes it possible to compare data from all chips directly in the following analysis.




Figure 11. PND pairing creation process for partial reuse analysis using Pearson's correlation coefficient.

Note: all PN are TVCOMP'ed but subscript 'c' is removed for clarity.

Only one of the PNR, PNR_0 , is used to create a set of 2048 PND_c by pairing it as shown with each of the PNF. Correlations that occur in the generated bitstring are rooted in correlations among the PND_c . Therefore, the 2048 PND_c are themselves paired, this time with each other under all combinations for $2048 \times 2047 / 2 = 2,096,128$ pairing combinations. The same process is carried out using the first PNF, PNF_0 , with all of the PNR (not shown) to create a second set of PND_c , which are again paired under all combinations. We use only one rising reference PN, PNR_0 , and one falling reference PN, PNF_0 , because the value of the PCC is identical for other choices of these references.

For each of the 2 million+ PND_c pairings, the Pearson correlation coefficient (PCC) given by Equation (6) is computed using enrollment data from the 500 chip-instances. PCC can vary from highly correlated (-1.0 and 1.0) to no correlation (0.0). The absolute value of the PCC in each group of

2 million+ rising and falling PND_c are then sorted from high to low. Scatterplots of the most highly and least correlated PND_c pairings are shown in Figure 12 from the larger set of more than 4 million pairings. The most highly correlated 1024 PND_c pairings (for a total of 2048 PND_c since each pairing contains two PND_c) are used in the bitstring generation process for the Entropy and Conditional MinEntropy (CmE) evaluation below. Highly correlated PND_c are stored as adjacent values to facilitate analysis of the corresponding 2-bit sequences:

$$PCC = \frac{\sum (X_i - \bar{X}) \cdot (Y_i - \bar{Y})}{\left[\sum (X_i - \bar{X})^2 \cdot \sum (Y_i - \bar{Y})^2 \right]^{1/2}}, \text{ where } -1 \leq PCC \leq 1. \quad (6)$$

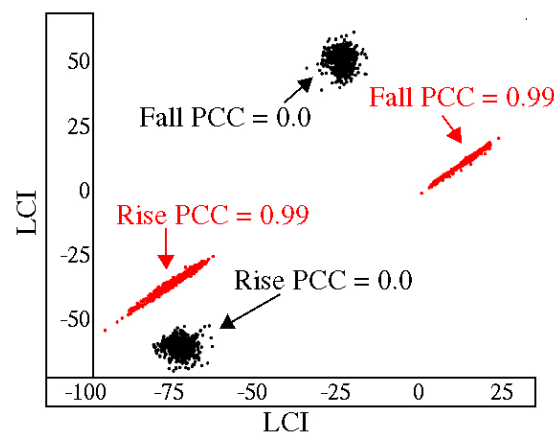


Figure 12. Scatterplot showing most correlated rising and falling PN and least correlated.

The 2048 PND_c are processed into bitstrings under four different scenarios as shown in Figure 10. For example, the PND_c are compared to a *global mean* under the *Reference* scenario (see annotation in figure). The global mean is the average PND_c across all chip-instances and all 2048 PND_c (500×2048). A '0' is assigned to the bitstring for cases in which the PND_c for a chip-instance falls below the *global mean* and a '1' otherwise. Given the large DC bias associated with the PND_c under the *Reference* scenario, the Entropy and CmE statistics are expected to be very poor.

The *No Offset, Mod* and *Offset, Mod* bitstring generation scenarios use the value 12 as the boundary between '0' and '1' (for Modulus 24 as shown in the figure), i.e., $PND_c \geq 0$ and < 12 produce a '0' and those ≥ 12 and < 24 produce a '1'. The '0'–'1' boundary for the *Offset, No Mod* scenario is 0 and the sign bit is used to assign '0' (for negative PND_c) and '1' (for positive PND_c). The *Offset, Mod* scenario represents the operations performed by the HELP algorithm. The analysis is extended for this scenario by evaluating Entropy and CmE over Moduli between 14 and 30 to fully illustrate the impact of the Modulus operation.

The PND_c from a *normal* use case are also analyzed using these four bitstring generation scenarios to determine how much Entropy/CmE is lost when compared to the *highly correlated* case analysis. For the normal use case, no attempt is made to correlate PND_c and instead random pairings of PNR and PNF are used to construct the PND_c. Table 1 provides a summary of the eight scenarios investigated.

Table 1. Summary of scenarios for partial reuse analysis

Cases		Scenarios		
Highly Correlated	No Offset No Modulus (Reference)	No Offset Modulus	Offset No Modulus	Offset Modulus (HELP)
Normal use				

Figure 13 provides a graphic that depicts the process used to compute Entropy and Conditional MinEntropy (CmE), (modeled after the technique proposed in [31]). As indicated earlier, highly correlated PND_c and the corresponding bits that they generate are kept in adjacent positions in the array. The bitstrings are of length 2048. Therefore, each chip-instance provides 1024 sets of 2-bit sequences.

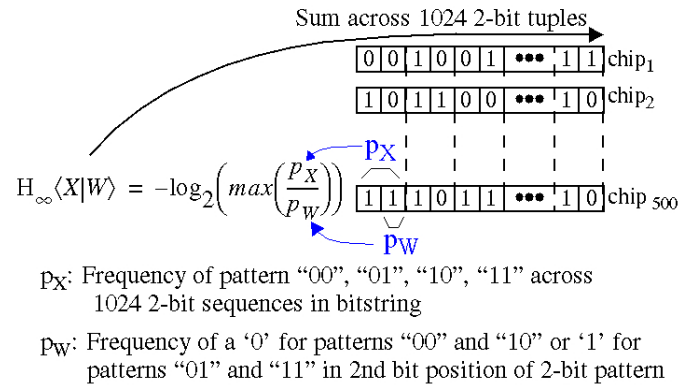


Figure 13. Conditional MinEntropy (CmE) expression and illustration of its application.

Equation (7) is used to compute the Entropy of the 1024 2-bit sequences for each chip-instance, which is then divided by 1024 to convert into Entropy/bit. The p_i represents the frequencies of the four 2-bit patterns as given in Figure 13. The Entropy/bit value reported below is the average of 500 chip- instance values. CmE is computed using Equation (8) (also from [31]). The expression $\max(p_X/p_W)$ represents the maximum conditional probability among the four values computed for each 2-bit sequence. Again, the sum over the 1024 2-bit sequences is converted to CmE/bit for each chip-instance and the average across all 500 chip-instances is reported:

$$H(X) = -\sum_{i=0}^3 p_i \cdot \log_2 p_i, \quad (7)$$

$$H_{\infty}(X|W) = -\log_2\left(\max\left(\frac{p_X}{p_W}\right)\right). \quad (8)$$

The Entropy and CmE results are plotted in Figure 14 for both the highly correlated and normal use scenarios. The x -axis represents the experiment, with 0 plotting the results using the *Reference* bitstring generation method (from Figure 10), 1 representing the *No Offset, Mod*, 2 representing *Offset, No Mod* and 3 through 11 representing the *Offset, Mod* method for Moduli between 30 and 14, respectively. The maximum Entropy/bit is 2 while the maximum CmE is 1. From the trends, it is clear that both Offset and Modulus improve the statistical quality of the bitstrings over the Reference. However, Modulus appears to provide the biggest benefit, which is captured by the drops in Entropy and CmE for experiment 2 in which the Modulus is not applied. Moreover, the loss in Entropy is almost zero between the normal use and highly correlated scenarios and CmE drops on average by only 0.2 bits for experiments 3 through 11 for the *Offset, Mod* method. Therefore, partial reuse under worst case conditions introduces only a small penalty on the quality of the bitstrings generated by the HELP algorithm.

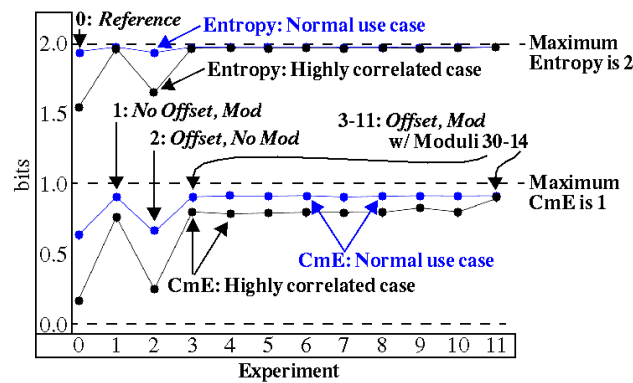


Figure 14. Entropy and Conditional MinEntropy results under processing schemes from Figure 10 and cases as given in Table 1.

5.3. Full Reuse

Full reuse refers to the **repeated use of the PN** in multiple PND_c as shown for the 2-PN reuse example in Figure 15. Here, two rise PN, PNR_0 and PNR_1 are paired in all combinations with two fall PN, PNF_0 and PNF_1 . The traditional analysis predicts that, because of correlation, only a subset of the 16 possible bit patterns can be generated when using PND_A through PND_D to generate a 4-bit response. In particular, patterns “0110” and “1001” are not possible. However, as indicated earlier, the Modulus and Offset operations break the classical dependencies and allow all patterns to be generated, as we show below.

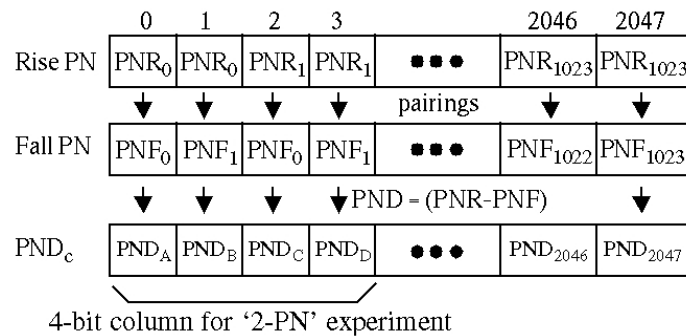


Figure 15. PND_c construction process for full reuse analysis, called 2-PN. Every column of 4-bits, with the first one labeled PND_A through PND_D , are correlated because two rise PNs and two fall PNs are subtracted under all combinations to create PND_c .

The frequency of the 16 patterns for the 2-PN experiment are shown in Figure 16. Here, PND_c are created for each of the 500 chip-instances according to the illustration in Figure 15. With 2048 bits/chip-instance, there are 512 4-bit columns each with 500 instances. The graph simply plots the percentage of each pattern across this set of $500 \times 512 = 256,000$ samples under each of the PND_c scenarios described earlier with reference to Figure 10. The ideal distribution is uniform with the percentage $1/16 \times 100 = 6.25\%$ for each ‘Pattern Bin’ along the x -axis as annotated in the figure.

The distributions associated with the *Reference* (black) and *Offset, No Mod* (red) experiments are clearly not uniform. Pattern bins 6 and 9 are zero for *Reference*, as predicted by the classical dependency analysis. Although the differences are small, the *Offset, No Mod* distribution is slightly better with non-zero values in pattern bins 6 and 9 and most of the other pattern bins closer to the ideal value of 6.25%. The Modulus operation, particularly in combination with the Offset operation, produce much better results. The percentages for the *Offset, Mod* experiment (yellow curve) vary by at most 1.2% from the ideal value of 6.25%.

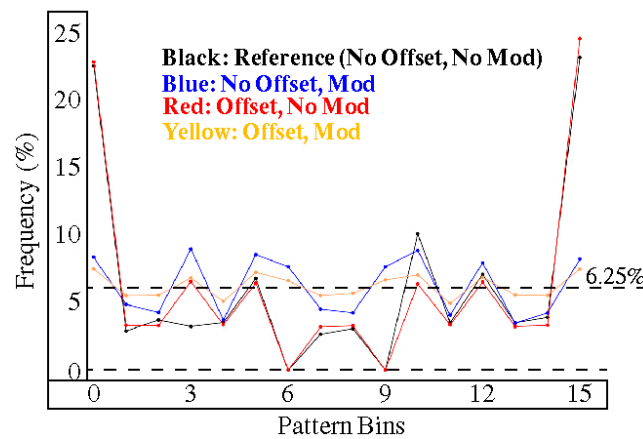


Figure 16. Frequency of 4-bit patterns for bin 0 with pattern “0000” through bin 15 with pattern “1111” using 2-PN reuse data under four scenarios from Figure 10. Ideal frequency value is $1/16 = 6.25\%$. Reference PND_c exhibits the worst case behavior with frequencies of 0% for patterns “0110” and “1001”, while Offset, Mod exhibits the best behavior.

The positive impact of the Offset and Modulus operations on Entropy is further supported by an analysis carried out in a 3-PN experiment, where 3 rise and 3 fall PN are combined under all combinations to produce a 9-bit column (analogous to 2-PN illustration in Figure 15). With 9-bit columns, there are 512 possible pattern bins. Using the 2048 bitstrings from 500 chip-instances, we were able to construct 227 full 9-bit columns (left over columns were discarded), for a total sample size of 113,500. A scatterplot showing the results for the 3-PN experiment is given in Figure 17 using *Offset, Mod* PND_c bitstring data (black dots). The ideal percentage is $1/512 \times 100 = 0.195\%$. As a reference, the results using PND_c constructed without reusing any rising or falling PN (referred to as the normal use scenario above) are superimposed in blue. The smaller variation of the frequencies under the normal use scenario, when compared with the 3-PN full reuse scenario, clearly shows that there is a penalty associated with reuse, but none of the pattern bins are empty in either case and most of the frequency values are within 0.1% of the ideal value at 0.195%.




Figure 17. Frequency of 9-bit patterns for bin 0 with pattern “000,000,000” through bin 512 with pattern “111,111,111” using 3-PN reuse data (black) and normal data (blue). The distribution should be uniform with each bin percentage at $1/512 = 0.195\%$ as shown by the dotted line.

Table 2 presents the MinEntropy computed using Equation (4) for each of the PND_c scenarios (rows) for the 2-PN and 3-PN experiments described above, and an additional 4-PN experiment. For the 4-PN experiments, all combinations of 4 PNs are used and the frequency of the 65,536 possible

patterns in the set of 128 16-bit columns are analyzed. The corresponding MinEntropy values under the normal use scenario (with column labeled ‘Normal’) are also given for reference.

In all cases, except for row 3, column 2, the MinEntropy values in the last row are larger than those in the first three rows. Moreover, the drop in MinEntropy over the normal use case scenario in the last row is 0.19, 1.45 and 1.9 bits, resp., illustrating that the penalty associated with reuse is very modest.

Table 2. MinEntropy for compensated PND (PND_c) and χ -PN experiments.

Scenarios	2-PN	Normal	3-PN	Normal	4-PN	Normal
No Offset, No Mod	2.11 of 4	3.05 of 4	3.17 of 9	6.15 of 9	4.3 of 16	7.0 of 16
No Offset, Mod	3.92 of 4	3.81 of 4	6.10 of 9	7.89 of 9	8.3 of 16	10.2 of 16
Offset, No Mod	2.02 of 4	2.82 of 4	3.05 of 9	5.34 of 9	4.1 of 16	8.5 of 16
Offset, Mod	3.73 of 4	3.92 of 4	6.95 of 9	8.40 of 9	9.2 of 16	11.1 of 16

6. Conclusions

An analysis of the statistical characteristics of a Hardware-Embedded Delay PUF (HELP) are presented in this paper, with emphasis on Interchip Hamming Distance, Entropy, MinEntropy, conditional MinEntropy and NIST statistical test results. The bitstrings generated by the HELP algorithm are shown to exhibit excellent statistical quality. An experiment focused on purposely constructing worst case correlations among path delays is also described as a means of demonstrating the Entropy-enhancing benefit of the Offset and Modulus operations carried out by the HELP algorithm. Special data sets are constructed which maximize physical correlations and dependencies introduced by reusing components of the underlying Entropy. Although statistical quality is reduced under these worst case conditions, the reduction is modest. Therefore, the Modulus and Offset operations harden the HELP algorithm against model-building attacks.

A quantitative analysis of the relationship between Entropy as presented in this paper and the level of effort required to carry out model-building attacks on HELP is the subject of a future work. Developing a formal quantitative framework that expresses the relationship between Entropy and model-building effort is inherently difficult because of the vastly different mathematical domains on which each is based. Best practices relating Entropy to security properties that predict attack resilience is focused on correlating results from separate analyses of Entropy and model-building resistance. A thorough treatment of model-building resistance requires a wide range of machine-learning experiments. Work on this topic is on-going and will be reported in a separate paper in the near future.

Author Contributions: Wenjie Che and Jim Plusquellic conceived and designed the experiments; Venkata K. Kajuluri performed the experiments; Mitchell Martin and Fareena Saqib collected the data; Jim Plusquellic wrote the paper.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Che, W.; Martin, M.; Pocklassery, G.; Kajuluri, V.K.; Saqib, F.; Plusquellic, J. A Privacy-Preserving, Mutual PUF-Based Authentication Protocol. *Cryptography* **2017**, *1*, 3. [[CrossRef](#)]
2. Gassend, B.; Clarke, D.; van Dijk, M.; Devadas, S. Controlled Physical Random Functions. In Proceedings of the Conference on Computer Security Applications, Washington, DC, USA, 9–13 December 2002.
3. Gassend, B.; Clarke, D.E.; van Dijk, M.; Devadas, S. Silicon Physical Unknown Functions. In Proceedings of the Conference on Computer and Communications Security, Washington, DC, USA, 18–22 November 2002; pp. 148–160.
4. Lofstrom, K.; Daasch, W.R.; Taylor, D. Identification Circuits using Device Mismatch. In Proceedings of the International Solid State Circuits Conference, Piscataway, NJ, USA, 31 May 2000; pp. 372–373.

5. Maiti, A.; Schaumont, P. Improving the quality of a Physical Unclonable Function using Configurable Ring Oscillators. In Proceedings of the 2009 International Conference on Field Programmable Logic and Applications, Prague, Czech Republic, 31 August–2 September 2009.
6. Meng-Day, Y.; Sowell, R.; Singh, A.; M'Raihi, D.; Devadas, S. Performance Metrics and Empirical Results of a PUF Cryptographic Key Generation ASIC. In Proceedings of the 2012 IEEE International Symposium on Hardware-Oriented Security and Trust (HOST), San Francisco, CA, USA, 3–4 June 2012.
7. Simpson, E.; Schaumont, P. Offline Hardware/Software Authentication for Reconfigurable Platforms. *Cryptogr. Hardw. Embed. Syst.* **2006**, *4249*, 10–13.
8. Habib, B.; Gaj, K.; Kaps, J.-P. FPGA PUF Based on Programmable LUT Delays. In Proceedings of the Euromicro Conference on Digital System Design, Santander, Spain, 4–6 September 2013; pp. 697–704.
9. Guajardo, J.; Kumar, S.S.; Schrijen, G.; Tuyls, P. Brand and IP Protection with Physical Unclonable Functions. In Proceedings of the Symposium on Circuits and Systems, Seattle, WA, USA, 18–21 May 2008; pp. 3186–3189.
10. Alkabani, Y.; Koushanfar, F.; Kiyavash, N.; Potkonjak, M. Trusted Integrated Circuits: A Nondestructive Hidden Characteristics Extraction Approach. In Proceedings of the 10th International Workshop Information Hiding, Santa Barbara, CA, USA, 19–21 May 2008.
11. Helinski, R.; Acharyya, D.; Plusquellic, J. Physical Unclonable Function Defined Using Power Distribution System Equivalent Resistance Variations. In Proceedings of the Design Automation Conference, San Francisco, CA, USA, 26–31 July 2009; pp. 676–681.
12. Chakraborty, R.; Lamech, C.; Acharyya, D.; Plusquellic, J. A Transmission Gate Physical Unclonable Function and On-Chip Voltage-to-Digital Conversion Technique. In Proceedings of the Design Automation Conference, Austin, TX, USA, 29 May–7 June 2013; pp. 1–10.
13. Aarestad, J.; Plusquellic, J.; Acharyya, D. Error-Tolerant Bit Generation Techniques for Use with a Hardware-Embedded Path Delay PUF. In Proceedings of the Symposium on Hardware-Oriented Security and Trust (HOST), Austin, TX, USA, 2–3 June 2013; pp. 151–158.
14. Saqib, F.; Arenó, M.; Aarestad, J.; Plusquellic, J. An ASIC Implementation of a Hardware-Embedded Physical Unclonable Function. *IET Comput. Digit. Tech.* **2014**, *8*, 288–299. [[CrossRef](#)]
15. Che, W.; Saqib, F.; Plusquellic, J. PUF-Based Authentication. In Proceedings of the 2015 IEEE/ACM International Conference on ICCAD, Austin, TX, USA, 2–6 November 2015.
16. Rose, G.S.; McDonald, N.; Lok-Kwong, Y.; Wysocki, B.; Xu, K. Foundations of Memristor Based PUF Architectures. In Proceedings of the International Symposium on Nanoscale Architectures, Brooklyn, NY, USA, 15–17 July 2013; pp. 52–57.
17. Yu, Z.; Krishna, A.R.; Bhunia, S. ScanPUF: Robust Ultralow-Overhead PUF using Scan Chain. In Proceedings of the Asia and South Pacific Design Automation Conference, Yokohama, Japan, 22–25 January 2013; pp. 626–631.
18. Konigsmark, S.T.C.; Hwang, L.K.; Deming, C.; Wong, M.D.F. CNPUF: A Carbon Nanotubebased Physically Unclonable Function for Secure Low-Energy Hardware Design. In Proceedings of the Asia and South Pacific Design Automation Conference, Singapore, 20–23 January 2014; pp. 73–78.
19. Majzoobi, M.; Koushanfar, F.; Devadas, S. FPGA PUF using Programmable Delay Lines. In Proceedings of the Workshop on Information Forensics and Security, Seattle, WA, USA, 12–15 December 2010; pp. 1–6.
20. Hori, Y.; Yoshida, T.; Katashita, T. Satoh Quantitative and Statistical Performance Evaluation of Arbiter Physical Unclonable Functions on FPGAs. In Proceedings of the Conference on Reconfigurable Computing and FPGAs, Cancun, Mexico, 13–15 December 2010; pp. 298–303.
21. Gassend, B.; Lim, D.; Clarke, D.; van Dijk, M.; Devadas, S. Identification and Authentication of Integrated Circuits. *Concurr. Comput. Pract. Exp.* **2014**, *16*, 1077–1098. [[CrossRef](#)]
22. Xin, X.; Kaps, J.; Gaj, K. A Configurable Ring-Oscillator-Based PUF for Xilinx FPGAs. In Proceedings of the Conference on Digital System Design, Oulu, Finland, 31 August–2 September 2011; pp. 651–657.
23. Suh, E.; Devadas, S. Physical Unclonable Functions for Device Authentication and Secret Key Generation. In Proceedings of the Design Automation Conference, San Diego, CA, USA, 4–8 June 2007; pp. 9–14.
24. Maiti, A.; Inyoung, K.; Schaumont, P. A Robust Physical Unclonable Function with Enhanced Challenge-Response Set. *Trans. Inf. Forensics Secur.* **2012**, *7*, 333–345. [[CrossRef](#)]
25. Chi, E.; Yin, D.; Qu, G. LISA: Maximizing RO PUF's Secret Extraction. In Proceedings of the 2010 IEEE International Symposium on Hardware-Oriented Security and Trust (HOST), Anaheim, CA, USA, 13–14 June 2010.

26. Chi, E.; Yin, D.; Qu, G. Improving PUF Security with Regression-based Distiller. In Proceedings of the Design Automation Conference, Austin, TX, USA, 29 May–7 June 2013.
27. Delvaux, J.; Gu, D.; Schellekens, D.; Verbauwhede, I. Helper Data Algorithms for PUF-based key generation: Overview and analysis. *Trans. Comput. Aided Des. Integr. Circuits Syst.* **2015**, *34*, 889–902. [CrossRef]
28. Bosch, C.; Guajardo, J.; Sadeghi, A.-R.; Shokrollahi, J.; Tuyls, P. Efficient Helper Data Key Extractor on FPGAs. *Workshop Cryptogr. Hardw. Embed. Syst.* **2008**, *5154*, 181–197.
29. Koeberl, P.; Li, J.; Rajan, A.; Wu, W. Entropy Loss in PUF-based Key Generation Schemes: The Repetition Code Pitfall. In Proceedings of the Symposium on Hardware-Oriented Security and Trust, Arlington, VA, USA, 6–7 May 2014; pp. 44–49.
30. Dodis, Y.; Ostrovsky, R.; Reyzin, L.; Smith, A. Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other Noisy Data. *SIAM J. Comput.* **2008**, *38*, 97–139. [CrossRef]
31. Katzenbeisser, S.; Kocabas, Ü.; Rozic, V.; Sadeghi, A.-R.; Verbauwhede, I.; Wachsmann, C. PUFs: Myth, Fact or Busted? A Security Evaluation of Physically Unclonable Functions (PUFs) Cast in Silicon. In Proceedings of the 14th international conference on Cryptographic Hardware and Embedded System, Leuven, Belgium, 9–12 September 2012.
32. Tiri, K.; Verbauwhede, I. A Logic Level Design Methodology for a Secure DPA Resistant ASIC or FPGA Implementation. In Proceedings of the Design, Automation and Test in Europe Conference and Exhibition, Paris, France, 16–20 February 2004; pp. 246–251.
33. Claes, M.; van der Leest, V.; Braeken, A. Comparison of SRAM and FF PUF in 65 nm Technology. In Proceedings of the Nordic Conference on Secure IT Systems, Karlskrona, Sweden, 31 October–2 November 2011; pp. 47–64.
34. National Institute of Standards and Technology. Available online: http://csrc.nist.gov/groups/ST/toolkit/rng/documentation_software.html (accessed on 1 January 2017).
35. Pearson Correlation Coefficient, Wikipedia. Available online: https://en.wikipedia.org/wiki/Pearson_correlation_coefficient (accessed on 1 January 2017).



© 2017 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).