



1 Article

2 **Correlation-Based Robust Authentication (Cobra)**
3 **using Helper Data Only**4 **Jim Plusquellic^{1*} and Matt Areno²**5 ¹ University of New Mexico and IC-Safety; jimp@ece.unm.edu6 ² Trusted and Secure Systems; matt@trusecsys.com7 * Correspondence: jimp@ece.unm.edu; Tel.: +240-475-1882

8 Received: date; Accepted: date; Published: date

9 **Abstract:** PUF-based authentication protocols have been proposed as a strong challenge-response
10 form of authentication for IoT and embedded applications. A special class of so called strong PUFs
11 are best suited for authentication because they are able to generate an exponential number of
12 challenge-response-pairs (CRPs). However, strong PUFs must also be resilient to model-building
13 attacks. Model-building utilizes machine learning algorithms and a small set of CRPs to build a
14 model that is able to predict the responses of a fielded chip, thereby compromising the security of
15 chip-server interactions. In this paper, response bitstrings are eliminated in the message exchanges
16 between chips and the server during authentication and therefore, it is no longer possible to carry
17 out model-building attacks in the traditional manner. Instead, the chip transmits a Helper Data
18 bitstring to the server and this information is used for authentication instead. The server constructs
19 Helper Data bitstrings using enrollment data that it stores for all valid chips in a secure database
20 and computes correlation coefficients (CCs) between the chip's Helper Data bitstring and each of
21 the server-generated Helper Data bitstrings. The server authenticates (and identifies) the chip if a
22 CC is found that exceeds a threshold, which is determined during characterization. The technique
23 is demonstrated using data from a set of 500 Xilinx Zynq 7020 FPGAs, subjected to industrial-level
24 temperature and voltage variations.

25 **Keywords:** PUF-based authentication; Helper data correlation; hardware security

26

27 **1. Introduction**

28 Robust authentication and key generation are critically important to defining a root of trust and
29 to providing data integrity and confidentiality in communications between internet-of-things (IoT)
30 devices. Physical unclonable functions (PUFs) are proposed as replacements to traditional non-
31 volatile memory (NVM) for storing keys and to using cryptographic primitives in authentication
32 protocols [1-7]. PUFs are able to reproduce keys and bitstrings on-the-fly by measuring small changes
33 in the signal behavior of an integrated circuit that occur because of the finite, non-zero tolerances that
34 exist in manufacturing processes. A special class of so-called strong PUFs are able to generate an
35 uncountable numbers of reproducible bits making it possible to construct unique response bitstrings
36 for authentication protocols without the need to employ entropy-enhancing cryptographic primitives
37 such as secure hash, thereby reducing energy and area overheads in IoT devices.

38 PUF-based authentication protocols that allow direct access to the embedded PUF through an
39 unprotected interface [7], where challenges and responses are not obfuscated using cryptographic
40 primitives, represent the most attractive usage scenario because such protocols are typically very
41 simple and compact. The drawback of protocols with unprotected interfaces is that the embedded
42 PUF is susceptible to model-building attacks. Model-building is typically carried out using machine
43 learning (ML) algorithms where the goal of the adversary is build a model of the challenge-response

44 behavior of the PUF and then later use the model to predict the PUF's response to any arbitrary
45 challenge. If this can be accomplished, then the model can be used to impersonate the actual device.
46 Although ML attacks require a large amount of computing effort, they represent a serious threat to
47 PUF-based authentication protocols with unprotected interfaces.

48 In this paper, we propose a PUF-based, privacy-preserving, mutual authentication protocol with
49 unprotected interfaces that is resilient to model-building attacks. The technique is demonstrated
50 using the hardware-embedded delay PUF (HELP) [7-8] but is applicable to any PUF architecture that
51 is able to produce **soft data**. Soft data refers to digital values that represent the magnitude of the
52 signal being measured, e.g., path delays and frequencies. Soft data captures the inherent, random
53 variations that occur in these signals from one chip to another, and represents the source of entropy
54 for the PUF. The proposed methodology can therefore be applied to an enhanced version of the
55 arbiter PUF [11], and to traditional weak PUFs, such as the RO PUF [12], with additional but simple
56 enhancements to expand the challenge-response space from n to 2^n as required for authentication
57 applications.

58 The PUF architecture defines a set of functions that convert soft data into bitstrings and keys to
59 be used for encryption and authentication. The authentication protocol that we propose uses soft
60 data, and the corresponding Helper Data bitstrings that are produced by the PUF architecture, as
61 input to a correlation technique. We refer to the protocol as **Cobra** for **Correlation-based robust**
62 **authentication**. The results presented in this paper show that by correlating Helper Data bitstrings,
63 a server can correctly and securely authenticate a fielded chip. The significance of this claim is that
64 there is no need to reveal the response bitstrings in the message exchanges between the chip and
65 server and therefore, the traditional approach of applying machine-learning algorithms to the
66 challenge-response-pairs (CRPs) is no longer possible.

67 In the Cobra protocol, the chip constructs and transmits a Helper Data bitstring to the server.
68 The Helper Data bitstring transmitted is traditionally used by the server to identify weak bits in the
69 response bitstring, but otherwise reveals no information about the response bitstring itself [7]. The
70 server constructs a set of Helper Data bitstrings using enrollment data that it stores for all valid chips
71 in a secure database. The chip's Helper Data bitstring is correlated to each of the server generated
72 Helper Data bitstrings by bitwise AND'ing the two bitstrings and then counting the number of '1's
73 in the AND'ed bitstring. (Alternatives to bitwise AND correlation include bitwise XNOR (discussed
74 later) and traditional time and frequency domain digital signal processing forms of correlation.) The
75 server authenticates the chip if **exactly one** of the Helper Data bitstrings constructed using enrollment
76 data correlates, i.e., has a large number of '1's, to the chip's Helper Data bitstring. A similar process
77 is carried out in reverse from server to chip to enable mutual (two-way) authentication but without
78 the exhaustive search component. Therefore, no information regarding the PUF secrets is revealed to
79 the adversary despite the fact that the Helper Data bitstrings are derived from random variations that
80 occur within the PUF's circuit components.

81 The remainder of this paper is organized as follows. Section 2 provides background on PUF
82 architectures and defines key concepts such as soft data, error avoidance methods, the HELP PUF
83 architecture, Helper Data generation, methods of correlation and an analysis illustrating proof-of-
84 concept. Section 3 describes the Cobra privacy-preserving, mutual authentication protocol.
85 Experimental results are presented in Section 4 on data collected from a set of 500 Xilinx Zynq FPGAs,
86 showing both the effectiveness and the limitations of Cobra on data collected across the range of
87 industrial-level temperature and voltage specifications. A security analysis is presented in Section 5
88 and conclusions in Section 6.

89 2. Background

90 2.1. PUF Architectures and Soft Data

91 A wide range of PUF architectures have been proposed since the initial papers on PUFs were
92 published [12][13]. The source of entropy (randomness) for the PUF is chip-to-chip and within-die
93 process variations that occur between and within chips during production. The PUF architecture

94 defines the mechanism that is used to measure small signal variations introduced by process
95 variations effects. In some cases, the measurement process leverages the existing architectural
96 features of the chip, e.g., the SRAM PUF measures its entropy source, i.e., the state of the individual
97 SRAM cells, by simply applying power to the SRAM array [14]. For most PUFs, however, circuit
98 components need to be added to the chip, e.g., the RO PUF requires a set of MUXs and counters to
99 select and measure the frequencies associated with elements in the array of ROs [12]. In another
100 example, the HELP PUF adds a launch-capture clocking mechanism to precisely time the delays of
101 combinational logic paths [7-8].

102 For PUF architectures that measure and digitize the signal behavior associated with the entropy
103 source, the digitized values provide additional information that can be leveraged in strong challenge-
104 response-pair (CRP) forms of authentication. The digitized values represent the magnitude of the
105 signal behavior, e.g., RO frequency or path delay, and are often used as input to mathematical
106 processes defined by the PUF architecture. The goal of the mathematical operations is to isolate and
107 amplify the random differences that occur among multiple copies of the individual circuit
108 components. The digitized values are eventually converted to a bit and used in the response of the
109 CRP. We refer to the digitized values as **soft data**.

110 Of the PUF architectures that create soft data (the RO and HELP PUFs are just two examples),
111 the conversion to bits can be accomplished in a variety of ways. For example, the RO PUF typically
112 selects a pair of ROs and then computes a difference by subtracting the soft data associated with the
113 two ROs. The sign of the difference can then be used to generate a bit, with, e.g., negative differences
114 producing a '0' and positive differences producing a '1'. The HELP PUF also computes differences
115 among pairings of path delays and uses a modulus operation to assign a '0' or '1' to the differences.

116 2.2. Error Correction and Avoidance Methods

117 Nearly all PUF architecture need to deal with bit-flip errors, which are differences in the
118 response bitstring that occur when the response is regenerated. Bit-flip errors are most probable when
119 the magnitude of the difference between a pair of soft data values is close to zero. In these cases,
120 regeneration of the bitstring, which takes place later in the field and under potentially adverse
121 environmental conditions, can result in bits flipping from '0' to '1' or vice versa. Although
122 authentication protocols can be designed to be tolerant to a small number of bit-flip errors, the
123 number of bit-flip errors that can occur is too large in most PUF architectures to guarantee that
124 authentication works correctly when regeneration is carried out in harsh environments.

125 To deal with this issue, nearly all PUF architectures define an **error correction** or **error avoidance**
126 method to improve reliability during regeneration. Error correction is the more popular of the two
127 reliability-enhancing methods. Error correction typically processes the PUF response bits into a final
128 response bitstring using algorithms based on linear block codes [15] or Bose-Chaudhuri-
129 Hochquenghen (BCH) codes [16]. However, nearly all of the error correction schemes ignore the
130 magnitude of the difference in the soft data and use all of the PUF response bits to construct a smaller
131 but reproducible bitstring response, including bits that have a high probability of changing value.

132 Error avoidance schemes, on the other hand, integrate a thresholding method that skips bits that
133 are deemed unreliable. The reliability of a bit is often directly related to the soft data associated with
134 the bit, and in particular, the distance of the soft data value to the bit-flip line. The bit-flip line is
135 defined by the PUF architecture as a soft threshold between '0' or a '1'. Unlike error correction
136 methods, error avoidance methods can only be used with PUF architectures that produce soft data
137 values, e.g., the RO [12], metal resistance [17], NVM [18] and HELP [8-10] PUFs are some examples.
138 Notable exceptions here are memory-based PUFs, including SRAM, DRAM, FF and latch-based
139 PUFs, which as originally proposed, are not capable to producing soft data.

140 In typical usage scenarios, an **enrollment** phase is carried out in which challenges are applied to
141 the PUF in a secure facility and response bitstrings are generated for the first time. In addition to the
142 response bitstring, PUF architectures that use either error correction or error avoidance methods also
143 produce Helper Data. Helper Data is typically maintained in the secure facility and transmitted to
144 the fielded device later during **regeneration** to enable the PUF to precisely reproduce the response

145 bitstring. The form of the generated Helper Data varies dramatically depending on the error
 146 correction or avoidance method employed. A key contribution of the method proposed here relates
 147 to Helper Data, and in particular to Helper Data that is generated by error avoidance methods. The
 148 next two subsections discuss a simple error avoidance scheme used by the HELP PUF as well as




Figure 1 HELP Architecture illustrating Clock Strobing concept used to create path delay (Timing) soft data.

149 features of the generated Helper Data that are leveraged in the proposed Cobra protocol.

150 2.3. HELP PUF

151 To better exemplify the principles of the proposed technology, we use the HELP PUF
 152 architecture and data collected from a set of Zynq FPGAs in the illustrative examples that follow [9-
 153 10]. HELP uses a *launch-capture* technique to obtain accurate digital timing values of path delays
 154 through a combinational logic block. The combinational logic block for a full adder is shown in Fig.
 155 1 but any functional unit can be used (the combinational logic from one column of the Advanced
 156 Encryption Standard is used in [7-10] and for the experimental results presented in Section 4). Logic
 157 signal transitions are launched from the *Launch FFs* shown on the left using Clk₁ and captured in the
 158 *Capture FFs* shown on the right using Clk₂.

159 The digital clock manager (DCM) on the FPGA is used to create the clocks, with *dynamic fine*
 160 *phase shift* enabled for Clk₂. Dynamic fine phase shift allows a state machine running in the
 161 programmable logic (PL) of the FPGA to increment the phase shift by increments of 18 ps (see right
 162 side of Fig. 1). A path through the full adder is timed by repeatedly applying a 2-vector sequence to
 163 the *Launch FFs* until the signal propagating along the highlighted path is successfully captured in the
 164 *Capture FF*. A successful capture occurs in this example when the '0' produced from the first vector
 165 V₁ of the 2-vector sequence is overwritten by the '1' produced by V₂ (see center portion of Fig. 1).
 166 When this occurs, the current fine phase shift value, which is an integer typically between 100 and
 167 500, is recorded by HELP as the digitized timing value for this path. These timing values represents
 168 the soft data associated with HELP.

	PN ₀	PN ₁	PN ₂	PN _x
C ₁	380.1	294.8	366.9	276.2
C ₂	366.6	282.8	352.7	286.1
C ₃	366.3	288.4	355.7	282.3
⋮				
C _n	387.5	301.2	373.5	272.1

Secure Server Database

Figure 2 HELP PN database created during enrollment.

169 During enrollment, a set of timing values, called **PUF Numbers** or PN, for each chip are stored
 170 in the rows of a secure database as shown in Fig. 2. This data is consulted by the secure server to
 171 authenticate these chips after they are deployed in fielded systems. The storage of soft data on the
 172 server, in contrast to response bitstrings, is the first of several significant differences that exist
 173 between Cobra and the PUF-based protocols proposed by others [1-6].

174 As indicated above, the proposed Cobra protocol leverages Helper Data to carry out
 175 authentication. The Helper Data is derived from the PNs stored in the secure database on the server,
 176 and from an instance of HELP that is programmed into the programmable logic of an FPGA, which
 177 represents the fielded chip. The entropy that is associated with each chip is captured by the PN stored
 178 in the server database. An adversary carrying out machine learning attacks against the protocol
 179 would attempt to learn the timing information stored in the database by reverse-engineering the
 180 bitstring responses that are exchanged openly over the network. Once learned, the adversary can
 181 then impersonate the chip. Therefore, it is vital that the relationship between the PN and the response
 182 bitstring be obscured and remain hidden to make this task difficult or impossible for the adversary.

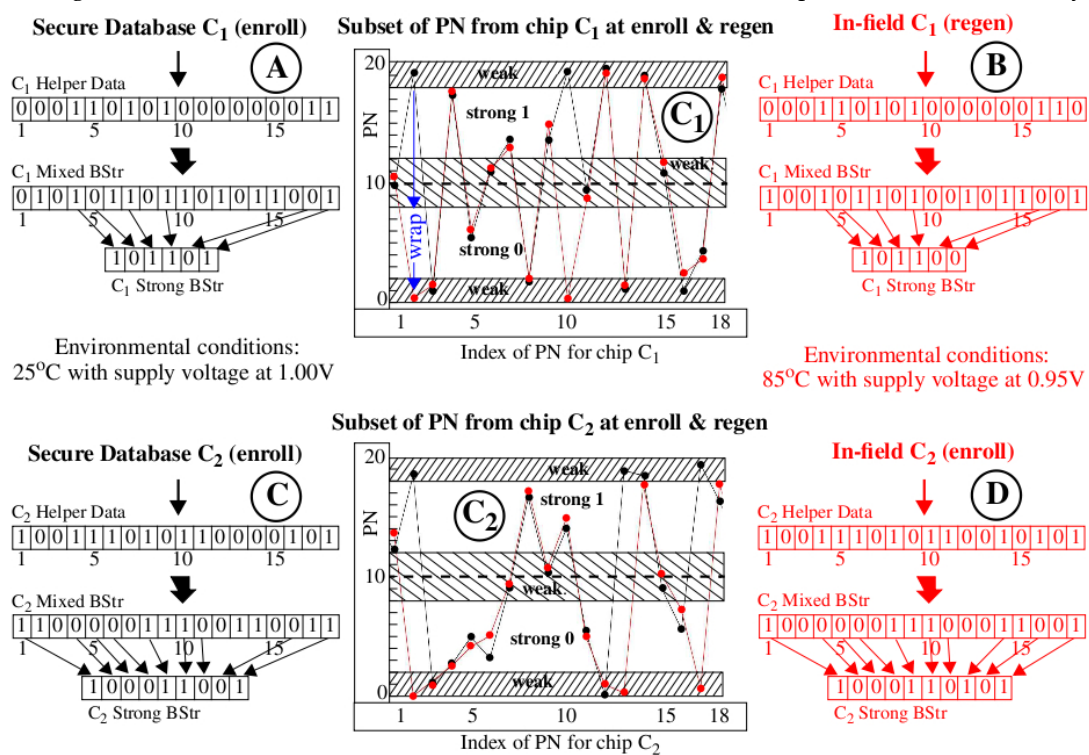


Figure 3 (Center) PN for chips C1 and C2 under enrollment (black points), and regeneration (red points), conditions, (top left A) Helper Data and response bitstrings for C1 under enrollment conditions, (top right B) Helper Data and response bitstrings for C1 under regeneration conditions, (bottom left C and right D) same for chip C2.

183 2.4. Helper Data Generation using an Error Avoidance Technique

184 The illustration presented in Fig. 3 shows how bitstrings and Helper Data are generated using
 185 an error avoidance scheme called Margining, as a precursor to our discussion on the proposed Helper
 186 Data correlation method. The graphs labeled with “C1” and “C2” (in large circles) in the center of the
 187 figure plot a set of 18 PN (timing values) along the x-axis for two chips C1 (top) and C2 (bottom), with
 188 the black curves depicting PN obtained from the secure server database and the red curves depicting
 189 PN generated on-the-fly by these chips during regeneration in the field. The environmental
 190 conditions for data collected during enrollment and stored in the secure database are specified as
 191 25°C under nominal supply voltage conditions (1.00V) while to fielded chips are subjected to high

192 temperature (85°C) and low supply voltage conditions (-5% or 0.95V). The data shown are actual
193 measurements obtained from two Zynq FPGAs used in our experiments.

194 The HELP PUF converts the PN into a bitstring by applying the following algorithm. First, a
195 pseudo-random number generator selects pairs of PN and creates differences. A temperature-voltage
196 compensation method called **TVComp** is then applied to the differences to compensate the measured
197 timing values for changes introduced by environmental conditions (we omit the details of these
198 operations [19] here to focus the discussion on the proposed correlation technique). Finally, a
199 modulus operation is applied to the compensated differences to remove path length bias effects. The
200 modulus operation is defined in the standard way as returning the positive remainder after dividing
201 by the modulus. The value of the modulus is one of several parameters to the HELP algorithm. The
202 graphs shown in Fig. 3 use a modulus of 20, which is reflected as the maximum value given on the
203 y-axis.

204 The TVComp process implemented within HELP is very effective at compensating the chip
205 regenerated PN (red values) but is not ideal. The red data points are vertically offset above and below
206 the black (enrollment) data points because of **uncompensated temperature-voltage noise** (TVNoise).
207 An interesting example labeled 'wrap' is shown for the 2nd data point in the upper "C₁" graph where
208 TVNoise has caused the point to 'wrap' from the enrollment value near 20 back around to a value
209 near 0 during regeneration. Despite these anomalies, the black and red curves in each graph track
210 very closely, i.e., they are correlated. In contrast, the black curves from both graphs (for C₁ and C₂)
211 are not correlated. This key observation serves as the basis for the innovation proposed within the
212 Cobra protocol.

213 As mentioned earlier, the error avoidance scheme implemented within HELP is called
214 Margining. The Margining scheme skips soft data values (PN in our example) for cases in which the
215 probability of a bit-flip error is large. These highly probable bit-flip regions are labeled "weak" in the
216 center graphs of Fig. 3 and are located adjacent to the bit-flip lines at 0, 10 and 20. In other words, PN
217 that are within a distance of 2.0 of these bit-flip lines have the highest probability of changing value.
218 For example, the PN labeled "wrap" represents a bit-flip error which is introduced by TVNoise. PN
219 that are located in the "weak" regions are assigned '0' in the Helper Data bitstring. For example, the
220 "C₁ Helper Data" bitstrings in the region labeled with the circled "A" in the figure begins with "000",
221 which reflects that the status of the first 3 PN in the black curve of graph "C₁". In contrast, the 4th bit
222 is '1' because the PN at position 4 in graph "C₁" falls within the "strong 1" region.

223 The "C₁ Mixed BStr" in region "A" of Fig. 3 records the bit value associated with each of the 18
224 PN, with PN < 10.0 assigned '0' and those >= 10.0 assigned '1'. This response bitstring corresponds
225 one-to-one to the "C₁ Helper Data" bitstring and contains both strong and weak bits. The "C₁ Strong
226 BStr" is constructed from the "C₁ Mixed BStr" by selecting only those bits identified as strong in the
227 "C₁ Helper Data" bitstring. The region labeled "B" shows the corresponding bitstrings generated
228 using the red (regeneration) data points from graph "C₁". The graphs and annotations labeled "C₂",
229 "C" and "D" are completely analogous to "C₁", "A" and "B" except the PN and bitstrings are derived
230 from second chip C₂.

231 The HELP authentication protocol from [7] proposes a DualHelperData scheme where both the
232 Helper Data and Strong BStr are exchanged between the server modeled on the left side of Fig. 3 and
233 the fielded chips modeled on the right side. As discussed earlier, exposing the Strong BStr to the
234 adversary enables model-building attacks where the adversary attempts to reverse engineer the PN
235 stored in the secure database using machine learning (ML) algorithms. Although attempts to model-
236 build HELP have not been successful (see [20]), the exposure of the response bitstrings (Strong BStr)
237 still represents a vulnerability that enables ML attacks. If it becomes possible to construct an ML
238 attack that is able to deduce the relationships among the PN, then the response bitstrings to other
239 challenges can be predicted, and the chip impersonated.

240 2.5. Poof-of-Concept: Helper Data Correlation

241 As a mitigation strategy, we propose an alternative authentication protocol that exchanges only
242 the Helper Data bitstrings. Authentication is carried out in Cobra by correlating the Helper Data

243 bitstrings generated using the enrollment data on the server with the Helper Data bitstring generated
 244 on-the-fly by the chip. The simplest correlation strategy is to compute a new bitstring by bitwise
 245 AND'ing the Helper Data bitstrings from the server and chip and then counting the number of '1's
 246 in the AND'ed version. We refer to the number of '1's as the **correlation coefficient (CC)** because it

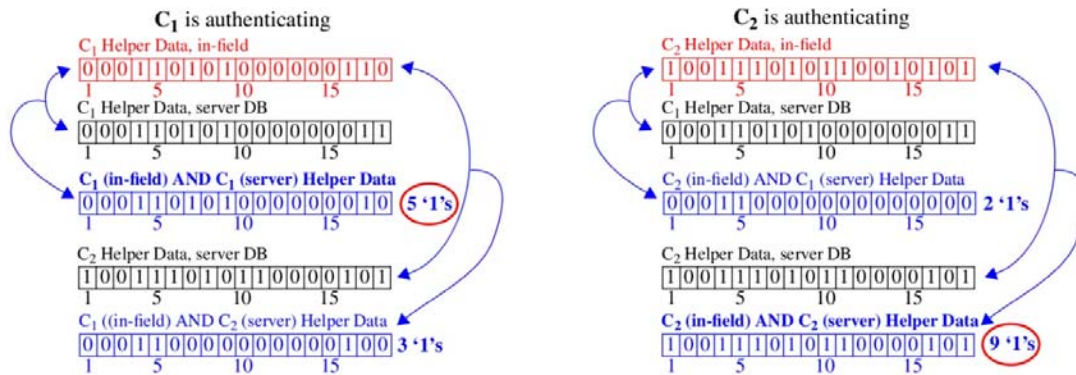


Figure 4 Illustration showing Helper Data bitstring correlation using “AND” operator with C₁ authenticating to the server on the left and C₂ on the right. Helper Data bitstrings are obtained from data in Fig. 3.

247 reflects the level of similarity that exists (among the '1's) between the two Helper Data bitstrings.

248 As an example, the left column of Fig. 4 shows an authentication attempt by chip C₁ while the
 249 right column shows an attempt by chip C₂. The top-most red-colored bitstrings in each column are
 250 the Helper Data bitstrings transmitted by the chips to the server. For each Helper Data bitstring
 251 received during authentication, the server carries out an exhaustive search operation using the PN
 252 stored its secure database. It constructs the black-colored Helper Data bitstrings in each column
 253 using the technique described in Fig. 3 (in fact, the bitstrings shown here are identical to those shown
 254 in Fig. 3).

255 For each Helper Data bitstring it constructs, the server bitwise AND's it with the received chip's
 256 Helper Data bitstring. For example, the AND'ed versions are labeled “C₁ (in-field) AND C₁ (server)
 257 Helper Data” and “C₁ (in-field) AND C₂ (server) Helper Data” in the left column of Fig. 4. As
 258 discussed, the bitwise AND operation is a form of correlation that acts to preserve more '1's in cases
 259 where the bitstrings are similar. The CCs (number of '1's) are reported to the right of AND'ed Helper
 260 Data bitstrings. The results for both authentication attempts show higher correlation for cases where
 261 the server-generated Helper Data bitstring is derived using PN collected earlier during enrollment
 262 from the same chip. In other words, the authenticating chip is correctly identified to the server using
 263 only the information provided by the CC.

264 This example uses only a small number of 18 PN from the larger set of 2048 that are produced
 265 by each iteration of the HELP algorithm [7]. The results shown in Fig. 5 expand the example
 266 illustration to the full length bitstrings and to PN collected across 9 TV corners using 500 Xilinx Zynq
 267 FPGAs. Each of the 9 curves plots the CCs for the 500 chips along the x-axis. The authenticating chip
 268 is labeled C₂₅₀ and is highlighted in the center region of the figure.

269 The graphic illustration given in Fig. 6 shows how the curves in Fig. 5 are constructed. Here, the
 270 chip's Helper Data bitstring on the left, regenerated under 25°C, 1.00V, is transmitted to the server
 271 on the right. A second authentication request is also shown in red where the chip in this case is
 272 regenerating Helper Data with environmental conditions set to -40°C, 1.05V. The server carries out
 273 an exhaustive search using data stored in the Enroll DB separately for each of these two
 274 authentication attempts and computes a set of 500 CCs. The CCs are plotted along the x-axis in Fig.
 275 5 as the top-most and bottom-most curves. A similar process is carried out to construct the CCs for
 276 the remaining 7 curves in Fig. 5 but using PN from the other TV corners.

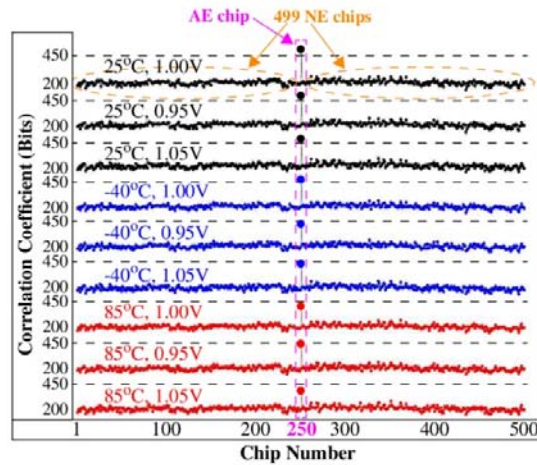


Figure 5 CCs (y-axis) for 500 chips (x-axis) across 9 TV corners obtained by correlating Helper Data bitstrings from the HELP PUF with a Margin of 3 and Modulus of 18. Peak at 250 occurs when Helper Data is derived using PN generated by the chip and matched with PN collected during enrollment for this same chip (called the authentic-enrolled or AE chip). All other CCs are derived by correlating

277 The CCs in Fig. 5 for C_{250} (referred to as the authentic-enrolled (AE) chip) vary from more than
 278 450 bits in the top curve to approx. 380 bits in the bottom curve. Although the correlation is weakened
 279 when the chip is exposed to harsh environmental conditions, it still remains high with respect to the
 280 CCs produced by the remaining 499 chips (referred to as non-authentic-enrolled (NE) chips) from the
 281 DB. The largest value associated with a NE chip is approx. 260. The large margin between the AE and
 282 NE CCs suggests that it should be possible for the server to define a threshold to distinguish
 283 successful authentications from unsuccessful authentications with very high probability, e.g., any
 284 value between 260 and 380 works in this example.

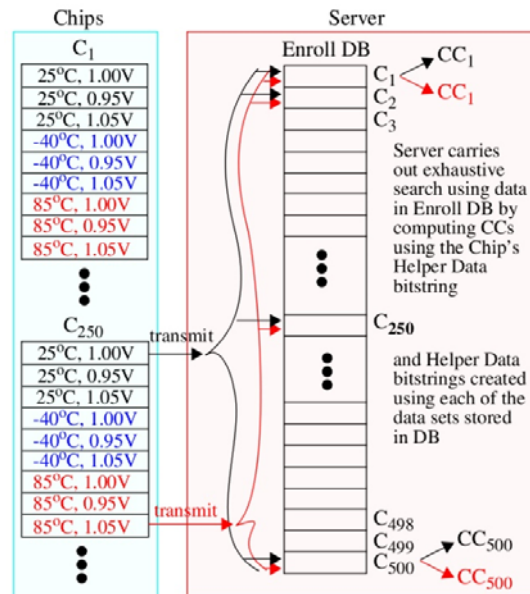


Figure 6 Graphic depicting process used to construct curves shown in Fig. 5. Each same-colored set of CCs displayed vertically are concatenated and shown along the x-axis in Fig. 5. This process represents exactly what the sever would do to identify and authenticate each Chip's Helper Data bitstring request transmitted to the server. Red-colored CCs correspond to the same C_{250} authenticating when the environmental conditions are 85°C, 1.05V as shown by the bottom-most curve in Fig. 5.

285 Note that our analysis considers only AE and NE authentication attempts. Two other
 286 possibilities include non-authentic-not-enrolled (NN) and non-authentic-counterfeit (NC)
 287 authentication attempts. Modeling NN authentication attempts is trivially accomplished by
 288 removing their data from the Enroll DB. It follows that attempts to authenticate by chips with no data
 289 in the Enroll DB would produce CCs similar to those produced for the 499 NE CCs shown in Fig. 5.
 290 Modeling NC authentication attempts is difficult without employing some type of machine learning
 291 algorithm if in fact an attack model can be devised. We leave this non-trivial task for future work.

292 We emphasize here that under the AND correlation scheme, it is possible for adversaries to
 293 construct Helper Data bitstrings with all ‘1’s, which guarantees a large number of matches. However,
 294 large CCs would occur for ALL data sets in the secure DB, which, in turn, would be flagged by the
 295 server and result in a failed authentication attempt. This is true because the server allows only one
 296 CC to be above the threshold in order for an authentication attempt to be classified as successful.




Figure 7 Cobra Protocol: Enrollment (top) and Chip Authentication (bottom) operations and message exchanges of proposed Helper Data bitstring correlation technique for implementing a privacy-preserving, mutual authentication protocol between chip (left) and server (right).

297 **3. The Cobra Protocol**

298 In this section, we describe the general structure of the proposed Cobra protocol. A graphical
 299 illustration of the Enrollment and Authentication operations, including the message exchanges
 300 between the chip and server, are presented in Fig. 7 along the top and bottom, respectively.
 301 Enrollment is performed in a secure facility using a confidential FPGA programming bitstream that
 302 allows access to the PUF’s soft data. The server on the right generates challenges $\{c_k\}$ and transmits
 303 them to the chip on the left. The chip then applies the challenges to its PUF to generate the set $\{PN_j\}$,
 304 which is returned to the server. The server generates a chip identifier ID_i and stores the soft data set
 305 $\{PN_j\}$ under ID_i in its secure database DB .

306 The first phase of authentication is called Chip Authentication. Here, a fielded chip i requests
 307 authentication to the server (note, no chip ID is transmitted to the server in order to preserve privacy).
 308 The server generates a nonce n_2 , selects a set of challenges $\{c_m\}$ and transmits them to the chip. The
 309 nonce n_2 is used to select a *Margin* parameter, and the challenges $\{c_m\}$ are a subset of the challenges
 310 $\{c_k\}$ used during enrollment. The chip applies the challenges $\{c_m\}$ to its PUF, along with a *Margin*,

311 which is selected using the function $ParamSel(n_1 \text{ XOR } n_2)$, and generates a Helper Data bitstring h'
312 using the Margining scheme described earlier. Both h' and n_1 are transmitted to the server.

313 The server then carries out a search by processing soft data sets $\{PN_j\}_i$ it stores in its database for
314 each chip i . The routine *Correlate* produces a Helper Data bitstring h internally (not shown) using each
315 of the stored data sets $\{PN_j\}_i$ and the same *Margin* that was used by the chip. The Helper Data bitstring
316 h' is then correlated with h . Correlation based on a bitwise AND operation was shown in the previous
317 section, but other possibilities exist including bitwise XNOR and/or standard waveform correlation
318 methods. The output of *Correlate* is a correlation coefficient CC_i that is then compared to a threshold
319 t . The authentication is considered successful if **exactly one** CC is larger than the threshold. The
320 'exactly one CC ' constraint implements a countermeasure against simple adversarial attacks which
321 use Helper Data bitstrings constructed with all 1's. This issue is discussed in detail in Section 5. The
322 threshold is determined in advance using a characterization process in a secure facility. The goal of
323 characterization is to select a threshold that unambiguously distinguishes authentic-enrolled chips
324 (AE) from non-authentic-enrolled (NE), non-authentic-not-enrolled (NN) and non-authentic-
325 counterfeit (NC) chips.

326 The last phase of the Cobra mutual authentication protocol is for the chip to authenticate the
327 server. This phase is not shown in Fig. 7 but is similar except the message exchanges are reversed and
328 the search process is omitted. Moreover, the AND-based correlation scheme used by the server to
329 authenticate the chip cannot be used because the chip does not have access to the secure database.
330 Instead, the chip uses XNOR correlation, which, as we discuss further below, requires matches to
331 both 0's and 1's in the Helper Data bitstring received from the server and the bitstring produced on-
332 the-fly by the fielded chip.

333 Server authentication is not performed unless chip authentication succeeds, in which case, the
334 server has identified the chip's soft data set $\{PN_j\}_i$. The server uses this $\{PN_j\}_i$ to generate another
335 Helper Data bitstring, which is transmitted to the chip. Note that the Helper Data bitstrings generated
336 during server authentication are distinct from those generated during chip authentication because
337 the challenge subset $\{C_m\}$ and nonces n_1 and n_2 are selected differently in this phase. Although the
338 nonces select only the *Margin* parameter in this example, other parameters can be introduced to
339 further expand the CRP space, as described below (and in reference [7]).

340 4. Experimental Results

341 This section carries out a worst case analysis using a larger set of the HELP CRP space and
342 discusses the security properties of the Cobra protocol in more detail. Similar to the preliminary
343 analysis presented in Section 2.5, the data analyzed here is collected from a set of 500 chip-instances
344 under enrollment and 9 temperature-voltage (TV) corners.

345 4.1. HELP Challenge Space

346 The full CRP space of the HELP algorithm is defined by 1) sets of challenges (2-vector sequences)
347 and corresponding *Path-Select-Masks* where each challenge set produces 4096 PN and 2) a set of
348 parameters, consisting of two *LFSR seeds*, two floating point parameters called the *reference mean* and
349 *range*, and a *Modulus* and *Margin* as discussed earlier in Section 2.4. The challenges and *Path-Select-*
350 *Masks* create a CRP space with size exponentially related to the size of the functional unit used as the
351 entropy source [21]. The parameters increase the CRP space by approx. 2^{20} , i.e., there are 2^{20} 2048-bit
352 bitstrings that can be generated by varying these parameters for each set of challenges and *Path-Select-*
353 *Masks*. Only one set of challenges and *Path-Select-Masks* are used for the analysis carried out in this
354 paper and instead we focus on analyzing Helper Data bitstrings produced by varying only the
355 parameters. Although this represents only a small subset of the entire CRP space, our results show
356 that the Cobra technique works well across a statistically significant sample.

357 The two *LFSR seed* parameters allow up to 2048 distinct bitstrings to be generated, each of length
358 2048 bits. The *reference mean* and *range* increase the number of distinct bitstrings by a factor of approx.

359 128. The analysis performed here analyzes the Helper Data bitstrings generated using 10 distinct
 360 *LFSR seeds*, one combination of the *reference mean* and *range* and a set of 11 *Moduli* and 3 *Margins*.

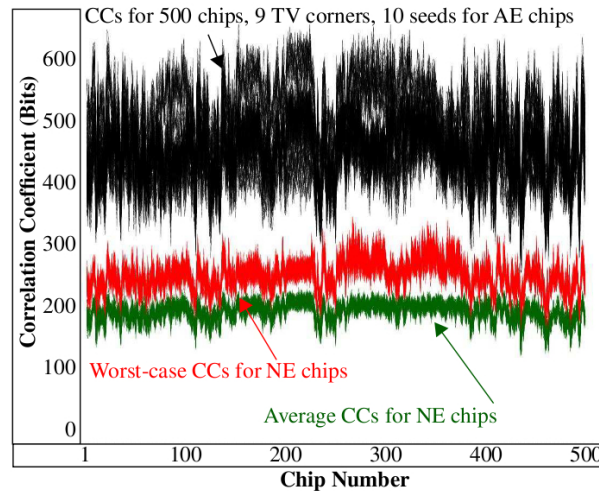


Figure 8 AE chip (black), worst-case NE chip (red), average-case NE chip (green) CCs with a Margin of 3 and Modulus of 18.

361 4.2. Illustration of Worst Case Scenario

362 The approach taken for the analysis of the worst case is illustrated using Figs. 8 and 9. The data
 363 presented is derived using only one *Modulus* and *Margin* in this section, and is expanded to the larger
 364 set in Section 4.3. The black curves in Fig. 8 plot the CCs computed by the server when each of the
 365 500 AE chips identified along the x-axis authenticates. These curves are constructed as we illustrated
 366 for C_{250} in Fig. 6 but now include data for all 500 chips and 9 TV corners shown on the left in that
 367 figure, and for the 10 *LFSR seeds*. Therefore, there are 90 curves, each with 500 AE CCs. Similarly, the
 368 red curves in Fig. 8 plot the worst-case (largest) NE CC among the remaining 499 CCs in each
 369 authentication attempt while the green curves plot the average NE CC.

370 Note that the curves shown in Fig. 8 are in a different format than the curves shown in Fig. 5. In
 371 particular, the CCs in the black curves of Fig. 8 correspond to the AE chip only and the red curves
 372 plot only one CC (the largest, worst case value) from the 499 NE CCs in each of the curves of Fig. 5.
 373 Therefore, only two points from each of the curves in Fig. 5 are plotted in Fig. 8, and appear as column
 374 pairs in the black and red curves. The two points in each pair represent the worst case (smallest)
 375 separation between the AE and NE CCs and visually portray how close a NE chip gets to being falsely
 376 authenticated as the AE chip on the server. In summary, the number of black and red CC pairs is
 377 given by $500 * 9 \text{ TV corners} * 10 \text{ LFSR seeds} = 45,000$. Note that each pair corresponds to 500
 378 authentication attempts so in total, there are $45,000 * 500 = 22,500,000$ (22.5M) authentication attempts.

379 As indicated, the key feature of this graph is the distance (separation) between pairs of points in
 380 the black and red curves. This separation is key to the server's ability to distinguish between AE and
 381 NE, NN and NC chip authentications. Unfortunately, a hard threshold (horizontal line) between the
 382 black and red points cannot be drawn without some AE authentications failing (false negatives) and
 383 some NE authentications succeeding (false positives).

384 To deal with this issue, we propose a new metric that computes the *percentage change CC*, called
 385 *PCC*, as follows. First, a set of CCs are computed using the chip's Helper Data bitstring against each
 386 of the server computed Helper Data bitstrings. Then the two largest CCs are identified and plugged
 387 in Eq. 1 to obtain the *PCC* value. The server successfully authenticates the chip if the *PCC* is larger
 388 than a hard threshold value, and fails otherwise. In cases where the authentication is successful, the

389 soft data in the Enroll DB associated the largest CC, $CC_{largest}$, identifies the authenticating chip.

$$PCC = \frac{(CC_{largest} - CC_{2nd_largest})}{CC_{largest}} \quad \text{Eq. 1.}\]$$

390

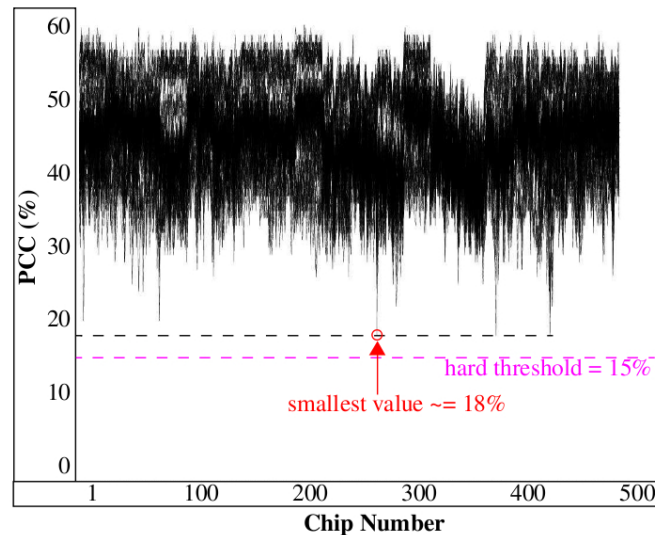


Figure 9 CCs as percentage change using AE chip and worst-case NE chip from Fig. 8.

391 The curves in Fig. 9 plot a closely related metric, defined as PCC_{AE_WC} in Eq. 2, using the CCs
 392 from Fig. 8. Here, CC_{AE} is associated with an AE chip obtained from the black curves in Fig. 8 and
 393 $CC_{worst_case_NE}$ is the other point in the pair obtained from the red curves. Note that in practice, we would
 394 not know the authenticate chip and therefore this analysis is somewhat artificial. However, it turns
 395 out for every CC pair in Fig. 8, $CC_{largest} = CC_{AE}$ and $CC_{2nd_largest} = CC_{worst_case_NE}$. Therefore, Eq. 1 and Eq. 2
 396 produce the same results for this set of CCs. In fact, the PCC_{AE_WC} would be negative if any of the CC_{AE}
 397 is not the largest CC among the 500 generated by the server for the authentication attempt. The
 398 smallest PCC_{AE_WC} present is circled in Fig. 9 and is approx. 18%. This indicates that all of the CC_{AE} are
 399 significantly larger than the NE CCs across the 22.5M authentications. Expressed in terms of bits, the
 400 smallest CC from Fig. 8 is approx. 300 bits. Therefore, the smallest separation between AE and NE
 401 CCs is approx. $300 \cdot 0.18 = 54$ bits. A **hard threshold** can now be defined, e.g., at 15% as shown in Fig.
 402 9, that enables the server to properly identify and authenticate chips with high probability.

$$PCC_{AE_WC} = \frac{(CC_{AE} - CC_{worst_case_NE})}{CC_{AE}} \quad \text{Eq. 2.}\]$$

403

404 4.3. Validation using the Larger CRP Space

405 An analysis reporting PCC_{AE_WC} is expanded in this section to a set of 11 *Moduli* and 3 *Margins*.
 406 The analysis is carried out using bitwise AND correlation as described in the previous sections and
 407 is repeated using bitwise XNOR correlation. XNOR correlation further restricts the matching criteria
 408 over AND correlation to count matches to both '1's and '0's in the two Helper Data bitstrings. Bitwise
 409 XNOR correlation relaxes the criteria used for a successful authentication in the Cobra protocol from
 410 'exactly one' to any CC that exceeds the threshold. This is possible because bitwise XNOR correlation
 411 measures the degree of matching between all bits in the Helper Data bitstrings, in contrast to bitwise
 412 AND which counts only the number of matching '1's. Interestingly, the two forms of correlation
 413 behave differently and are somewhat complementary as we discuss below and further in Section 5.

414 The PCC_{AE_WC} values are shown in Fig. 10 for Helper Data bitstrings derived using bitwise AND
 415 correlation along the top row and bitwise XNOR correlation along the bottom row. The bar heights

416 represent the worst case PCC_{AE_WC} for a set of *Moduli* given along the x-axis and *Margins* given along the y-axis. Note that only one PCC_{AE_WC} is reported for each *Margin-Moduli* combination, and in
 417 the y-axis. Note that only one PCC_{AE_WC} is reported for each *Margin-Moduli* combination, and in
 418 particular, the value that corresponds to the circled point in Fig. 9 labeled ‘smallest value’. Negative
 419 height bars represent that at least one authentication failure has occurred, i.e., at least one CC_{AE} is not
 420 the largest CC among the 500 CCs computed as we discussed in the previous section.

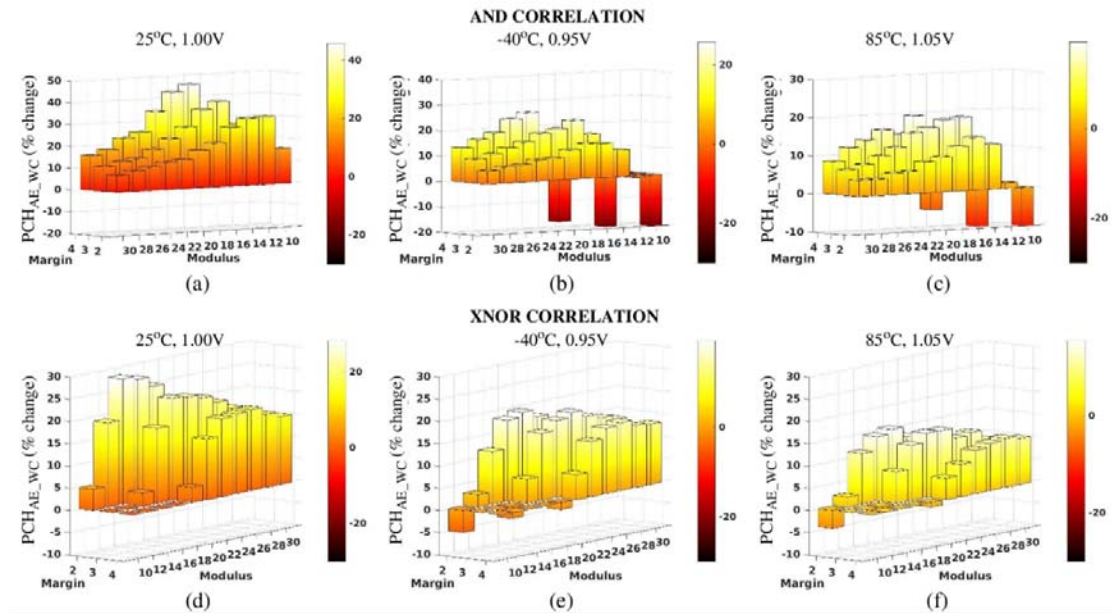


Figure 10 Worst case Correlation Coefficient (CC) differences between AE and NE chips expressed as percentage change using Eq. 2 under nominal environmental conditions, 25°C, 1.00V in (a) and (d), and worst case environmental conditions, -40°C, 0.95V in (b) and (e), and 85°C, 1.05V in (c) and (f) for 22.5M authentication attempts for Margins 2 through 4 and Moduli 10 through 30. Top row gives results using bitwise AND correlation and bottom row gives results using bitwise XNOR correlation. Positive bars indicate server identifies AE chip correctly in all 22.5M authentications while negative bars indicate at least one false authentication with an NE chip occurred.

421 The bar graphs in each of the three columns in Fig. 10 gives the results for three TV corners,
 422 namely 25°C, 1.00V in (a) and (d), -40°C, 0.95V in (b) and (e) and 85°C, 1.05V in (c) and (f). The latter
 423 two columns represent worst case TV corners, i.e., the results for the remaining 6 TV corners (not
 424 shown) produce bars that are larger. From these results, it is clear that the *Margin-Moduli*
 425 combinations with negative bar heights cannot be used in the Cobra authentication protocol.
 426 However, *Moduli* between 18 to 22 for *Margin* 3 and for 22 and 24 for *Margin* 4 produce bar heights
 427 greater than 15% when using AND correlation. For XNOR correlation, the behavior is reversed
 428 regarding the *Margin* where *Margins* of 2 and 3 produce better results (note that the bar graph
 429 orientation in the top row is rotated 180° in the bottom row). Here, the $PCCs$ exceed 15% for *Moduli*
 430 of 16 and 18 for a *Margin* of 2 and for *Moduli* 20 and 22 for a *Margin* of 3.

431 The bar graphs in the left column for TV corner 25°C, 1.00V show the $PCCs$ generated under
 432 nominal conditions. Unlike the results for the other TV corners, the bar heights for all *Moduli* and
 433 *Margins* are positive, indicating that the CC_{AES} are the largest among the sets of 22.5M authentications
 434 carried out in each analysis.

435 5. Security Analysis

436 The effectiveness of the proposed Helper Data bitstring correlation method is directly related to
 437 two components of the soft data processing algorithm carried out within the PUF architecture. The
 438 first is temperature-voltage compensation, referred to as TVComp earlier. Under the HELP
 439 algorithm, TVComp scales and shifts the path delays (PN) measured on-the-fly by the chip to make

440 them as similar as possible to the values measured during enrollment under nominal conditions. The
441 proposed correlation techniques depend heavily on the effectiveness of TVComp. For other PUF
442 architectures, e.g., the ARB, RO, NVM and metal PUFs, a similar form of TV compensation can be
443 applied as a means of enabling correlation-based authentication as described here for HELP.

444 A second critically important component of the correlation methods is related to the Margining
445 scheme portrayed within the center graphs of Fig. 3. Each of the two response bit regions, with '0'
446 assigned to the region between 0-10 and '1' assigned to the region between 10-20, also contain two
447 strong-weak region boundaries. Bit assignments within the Helper Data bitstrings depend only on
448 these four strong-weak region boundaries and are independent of response bit boundaries at 0 and
449 10. In other words, a Helper Data bit can be assigned '0' or '1' in either of the response bit regions
450 with equal probability. The symmetrical placement of the strong-weak region boundaries within each
451 response bit region eliminates leakage in the Helper Data bitstrings and is the basis for the claimed
452 improvements to model-building resistance of the correlation methods.

453 As mentioned earlier, when the AND correlation method is used by Cobra to generate *PCCs*, an
454 authentication is deemed successful only in cases where **exactly one** server-computed *PCC* is above
455 the threshold. The requirement of 'exactly one *PCC*' represents a countermeasure to adversarial
456 attacks in which Helper Data bitstrings are artificially constructed with all '1's or a large fraction of
457 '1's. The server is effectively screening for an outlier, i.e., one *PCC* that is significantly larger than all
458 the others it computes during the exhaustive search process. A successful impersonation attack then
459 requires the adversary to construct a Helper Data bitstring that is consistent with a matching server-
460 generated version **at all bit positions**, i.e., both the '0's and '1's must be correlated. Otherwise
461 authentication fails because more than *PCC* is above the threshold.

462 From the results shown in Fig. 10 (a), the AND correlation method performs best, i.e., produces
463 the largest *PCCs*, when the fraction of '1's in the Helper Data bitstrings is small. The fraction of '1's
464 (and '0's) is determined by the ratio of the *Margin* and *Modulus*. The Margining scheme requires the
465 $Modulus \geq 4 * Margin + 2$. As an example, when the *Margin* is 4, the *Modulus* must be at least 18.
466 Assuming the PN are evenly distributed across the range defined by the *Modulus* (which is not valid
467 for individual PN but is valid across a large collection of PN), the fraction of '1's is approximately
468 equal to the sum of the two strong bit regions divided by the *Modulus*. For AND correlation, the best
469 results are obtained when the strong response bit regions are size 1 or 2. In particular, the largest CC_{AE}
470 is nearly 45% larger than $CC_{worst_case_NE}$ in Fig. 10 (a) for a *Margin* of 4 and a *Modulus* of 18. The fraction
471 of '1's in this case is $2/18 \approx 11\%$.

472 The strong-weak bit regions acts as selection functions for the correlation methods. The AND
473 correlation method is a one-sided selection function because only the 'strong bit' side of the boundary
474 impacts the *PCC*. From the results, AND correlation performs best when the selection regions are
475 asymmetrically skewed toward narrow strong bit regions.

476 In contrast, both sides of the strong-weak boundaries affect the XNOR *PCC*. This fundamental
477 difference in the two correlation functions is reflected in the *PCCs* computed using different *Margins*.
478 For AND correlation in Fig. 10 (a), the largest *PCCs* occur using a *Margin* of 4 for the majority (not all)
479 of the *Margin-Moduli* combinations, while best case for the XNOR correlation occurs for a *Margin* of
480 2. The bar graphs in the second row are rotated 180° to more clearly illustrate this characteristic. For
481 example, the best results for XNOR occur for a *Margin* of 2 and for *Moduli* of 14 and 16. For these
482 *Margin-Moduli* combinations, the size of weak and strong bit regions are nearly equal (they are equal
483 for a *Modulus* of 16). Therefore, the two-sided XNOR selection function appears to work better for
484 *Margin-Moduli* combinations that divide the entire region more evenly into weak and strong regions.
485 Given the distinctive behavior of the AND and XNOR correlation functions, the Cobra protocol can
486 in fact leverage both *PCCs* as a means of reducing the probability of false negative and false positive
487 authentication decisions.

488 The overall decrease in the *PCC* magnitudes under adverse environmental conditions (Figs. 10
489 (b-c) and (e-f)) is caused by TVNoise. From the results, it is clear that the level of sensitivity of the
490 *PCCs* to TVNoise is higher for smaller *Margins* and *Moduli*.

491 5.1. Overhead of the Cobra Protocol

492 The largest component of the overhead associated with the Cobra protocol is related to the PUF
493 architecture. Reference [10], Table II gives the resource requirements for the HELP algorithm and
494 original HELP protocol as 2,350 LUTs and 1,454 FFs with an additional 706 and 3,494 LUTs needed
495 to implement the functional unit (entropy source) using an instance of AES SBOX and AES
496 *sbox_mixcol*, resp. An instance of *sbox_mixcol* is used to generate the data presented in this paper. It
497 consists of 4 AES SBOXs and 1 32-bit copy of AES mixed column implemented in a hazard-free
498 combinational logic style. Cobra eliminates the need to generate the response bitstring and therefore,
499 is slightly smaller by approx. 100 LUTs. For example, the size of the implementation used for
500 experiments in this paper is 5,750 LUTs and 1,454 FFs, excluding the LUTs and FFs used by the Xilinx
501 IP. Total size with Xilinx IP is 6,770 LUTs and 3,271 FFs, plus 1 MMCM and 1 25-bit multiplier.

502 5.2. Analysis of Cobra's Challenge-Response Space

503 As outlined earlier in Section 4.1, the full CRP space of HELP is defined by two components; 1)
504 the values associated with a set of six parameters including 2 11-bit *LFSR seeds*, two floating point
505 parameters called the *reference mean* and *range*, and a *Modulus* and *Margin*, and 2) the challenges and
506 *Path-Select-Masks*. The challenges and *Path-Select-Masks* are used to select two sets of 2,048 PN from
507 the larger set $\{PN_j\}$ stored in the secure database (see Fig. 7). The number of distinct 2048-bit Helper
508 Data bitstrings that can be generated by varying the parameters is given by the product: $2048 * 128 * 2 * 2 = 2^{20}$ or 1 million. HELP first creates 2048 PN differences (PND) by subtracting unique pairs of
509 elements from the two sets of 2,048 PNs. The factor 2,048 in the above product represents the number
510 of ways unique sets of 2,048 PND can be created using the *LFSR seeds*. HELP applies a procedure
511 called TVComp that shifts and scales the PND using two floating point parameters called the *reference*
512 *mean* and *range*. The factor 128 in the product represents a conservative estimate on the number of
513 ways the PND can be modified to produce unique response and Helper Data bitstrings by varying
514 these parameters. Finally, the factors of 2 represent a conservative estimate on the number of *Margins*
515 and *Moduli* that can be applied, as we discussed in Section 4.2. Note that 2^{20} represents the number of
516 unique Helper Data bitstrings that can be produced using the 4,096 PN selected by one set of
517 challenges and *Path-Select-Masks*.
518

519 The challenges and *Path-Select-Masks* represent the larger component of the CRP space. In [21],
520 we used two sets of 7,500 PN as the enrollment data (15,000 PN), where each PN can be stored as a
521 16-bit fixed-point value, resulting in less than 32 KB of storage in the secure database per fielded
522 device. Although the number of distinct PND that can be created from the two sets of 7,500 PN is
523 $7,500^2 \approx 56$ million, the number of distinct response and Helper Data bitstrings that can be produced
524 is much larger because of the Distribution Effect (which is the main topic of [21]). The challenges and
525 *Path-Select-Masks* allow any two subsets of 2,048 PN from the 7,500 sets to be selected. The
526 Distribution Effect is an artifact of the TVComp process which transforms any given PND into one of
527 approx. 100 different compensated PND (PND_c). More importantly, it is not possible to predict the
528 value of a PND_c unless the entire set of 2,048 PND are known. Although the Distribution Effect only
529 increases the number of distinct PND_c to approx. 5.6 billion, the number of different distributions of
530 2,048 PND is characterized as $(7,500 \text{ select } 2,048)$ which is a very large exponential.

531 The entire CRP space with 15,000 PN per fielded device would then be lower bounded by the
532 product of $2^{20} * 2^{32} = 2^{52}$, which accounts for both the parameters and challenge components. The large
533 diversity of HELP's CRP space prevents replay attacks, and adds significantly to the difficulty of
534 model-building attacks if in fact such attacks are possible using only Helper Data bitstrings.

535 6. Conclusions

536 A privacy-preserving, mutual PUF-based authentication protocol called Cobra is described in
537 this paper. Cobra exchanges and correlates Helper Data bitstrings instead of PUF response bitstrings
538 as a means of authenticating the chip and server. Helper Data is derived from the PUF response
539 bitstrings and therefore the Helper Data bitstrings inherit the randomness and uniqueness

540 characteristics associated with the PUF's source of entropy. By eliminating PUF response bitstrings
541 in the message exchanges between the chip and server, attacks such as model-building are much
542 more difficult to carry out. Cobra is demonstrated on a statistically significant set of FPGAs using the
543 HELP algorithm, and a simple thresholding method is proposed that the server and chip can use for
544 authentication. Although the HELP algorithm is used in this paper, the method is applicable to any
545 PUF that produces soft data, i.e., digitized values that capture the magnitude of signal behavior such
546 as delay or metal resistance. Future work will investigate the application of the Helper Data
547 correlation method to other PUF architectures and will evaluate more traditional forms of correlation
548 that are used in digital signal processing applications.

549 **Author Contributions:** "Conceptualization, Jim Plusquellic; Methodology, Jim Plusquellic; Validation, Matt
550 Areno; ; Writing-Original Draft Preparation, Jim Plusquellic and Matt Areno;

551 **Funding:** "This research received no external funding"

552 **Conflicts of Interest:** "The authors declare no conflict of interest."

553 References

- 554 1. L. Bolotny and G. Robins, "Physically Unclonable Function-based Security and Privacy in RFID Systems",
555 *PerCom*, 2007, pp. 211-220.
- 556 2. G. Hammouri, E. Ozturk, and B. Sunar, "A Tamper-Proof and Lightweight Authentication Scheme",
557 *Pervasive and Mobile Computing*, Vol. 4, Issue 6, 2008, 807-818.
- 558 3. A.-R. Sadeghi, I. Visconti, and C. Wachsmann, "Enhancing RFID Security and Privacy by Physically
559 Unclonable Functions", *Information Security and Cryptography*, 2010, pp. 281-305.
- 560 4. U. Kocabas, A. Peter, S. Katzenbeisser, and A. Sadeghi, "Converse PUF-Based Authentication", *TRUST*,
561 2012, pp. 142-158.
- 562 5. M. Majzoobi, M. Rostami, F. Koushanfar, D. S. Wallach, and S. Devadas, "Slender PUF Protocol: A
563 Lightweight, Robust, and Secure Authentication by Substring Matching", *Symposium on Security and Privacy
564 Workshop*, 2012, pp. 33-44.
- 565 6. A. Aysu, E. Gulcan, D. Moryama and P. Schaumont, "Compact and Low-power ASIP Design for
566 Lightweight PUF-based Authentication Protocols", *IET Information Security*, Vol. 10, Issue 5, 2016, pp. 232-
567 241.
- 568 7. W. Che, M. Martin, G. Pocklassery, V. K. Kajuluri, F. Saqib, and J. Plusquellic, "A Privacy-Preserving,
569 Mutual PUF-Based Authentication Protocol", *Cryptography* 2017.
- 570 8. J. Aarestad, P. Ortiz, D. Acharyya and J. Plusquellic, "HELP: A Hardware-Embedded Delay-Based PUF",
571 *Design and Test of Computers*, Mar., 2013, pp. 17-25.
- 572 9. W. Che, V. K. Kajuluri, M. Martin, F. Saqib and J. Plusquellic, "Analysis of Entropy in a Hardware-
573 Embedded Delay PUF", *Cryptography*, 2017.
- 574 10. W. Che, F. Saqib and J. Plusquellic, "Novel Offset Techniques for Improving Bitstring Quality of a
575 Hardware-Embedded Delay PUF", *Trans. on VLSI*, 2018.
- 576 11. http://ece-research.unm.edu/jimp/pubs/ARB_PUF.pdf, *UNM publication*, Mar. 2013.
- 577 12. B. Gassend, D. E. Clarke and M. van Dijk, S. Devadas, "Silicon Physical Unknown Functions", *Conference
578 on Computer and Communications Security*, 2002, 148-160.
- 579 13. K. Lofstrom, W. R. Daasch, D. Taylor, "IC Identification Circuits using Device Mismatch", *International
580 Solid State Circuits Conference*, 2000, pp. 372-373.
- 581 14. J. Guajardo, S. S. Kumar, G.-J. Schrijen, and P. Tuyls, "FPGA Intrinsic PUFs and their use for IP Protection",
582 *Cryptographic Hardware and Embedded Systems (CHES)*, 2007, pp. 63-80.
- 583 15. Y. Dodis, L. Reyzin, A. Smith, "Fuzzy Extractors: How to Generate Strong Keys from Biometrics and Other
584 Noisy Data", *Advances in cryptology (EUROCRYPT)*, 2004, pp. 523-540.
- 585 16. C. Bosch, J. Guajardo, A.-R. Sadeghi, J. Shokrollahi and P. Tuyles, "Efficient Helper Data Key Extractor on
586 FPGAs", *Workshop Cryptographic Hardware and Embedded Systems (CHES)*, pp. 181-197.
- 587 17. R. Helinski, D. Acharyya and J. Plusquellic, "A Physical Unclonable Function Defined Using Power
588 Distribution System Equivalent Resistance Variations", *Design Automation Conference*, 2009, pp. 676-681.
- 589 18. W. Che, S. Bhunia and J. Plusquellic, "A Non-Volatile Memory based Physically Unclonable Function
590 without Helper Data", *International Conference on Computer-Aided Design*, 2014.

- 591 19. <http://ece-research.unm.edu/jimp/HOST/index.html>, see video tutorials labeled 'HELP' and 'HELP
592 Protocol'
- 593 20. W. Che, M. Martinez-Ramon, F. Saqib, J. Plusquellic, "Delay Model and Machine Learning Exploration of
594 a Hardware-Embedded Delay PUF", *Symposium on Hardware-Oriented Security and Trust (HOST)*, 2018.
- 595 21. W. Che, V. K. Kajuluri, F. Saqib and J. Plusquellic, "Leveraging Distributions in Physical Unclonable
596 Functions", *Cryptography*, MDPI, 2017.



599

1. © 2018 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).