Optical PUFs Reloaded

Ulrich Rührmair
Computer Science Dept.
TU München
80333 München, Germany
ruehrmair@in.tum.de

Christian Hilgers
Computer Science Dept.
TU München
80333 München, Germany
hilgers@muc.zae-bayern.de

Sebastian Urban
Computer Science Dept.
TU München
80333 München, Germany
surban84@gmail.com

Agnes Weiershäuser
Computer Science Dept.
TU München
80333 München
agnes.weiershaeuser@gmail.com

Elias Dinter
Computer Science Dept.
TU München
80333 München
elias@dinternetz.de

Brigitte Forster
FIM
Universität Passau
94032 Passau, Germany
brigitte.forster@uni-passau.de

Christian Jirauschek
Electrical Engineering Dept.
TU München
80333 München, Germany
jirauschek@mytum.de

ABSTRACT

We revisit optical physical unclonable functions (PUFs), which were proposed by Pappu et al. in their seminal first publication on PUFs [40, 41]. The first part of the paper treats *non-integrated* optical PUFs. Their security against modeling attacks is analyzed, and we discuss new image transformations that maximize the PUF's output entropy while possessing similar error correction capacities as previous approaches [40, 41]. Furthermore, the influence of using more than one laser beam, varying laser diameters, and smaller scatterer sizes is systematically studied. Our findings enable the simple enhancement of an optical PUF's security without additional hardware costs. Next, we discuss the novel application of non-integrated optical PUFs as so-called *"Certifiable PUFs"*. The latter are useful to achieve practical security in advanced PUF-protocols, as recently observed by Rührmair and van Dijk at Oakland 2013 [48]. Our technique is the first mechanism for Certifiable PUFs in the literature, answering an open problem posed in [48].

In the second part of the paper, we turn to *integrated* optical PUFs. We build the first prototype of an integrated optical PUF that functions without moving components and investigate its security. We show that these PUFs can surprisingly be attacked by machine learning techniques if the employed scattering structure is linear, and if the raw interference images of the PUF are available to the adversary. Our result enforces the use of non-linear scattering structures within integrated PUFs. The quest for suitable materials is identified as a central, but currently open research problem. Our work makes intensive use of two prototypes of optical PUFs. The presented integratable optical PUF prototype is, to our knowledge, the first of its kind in the literature.

Categories and Subject Descriptors

B.0 [**Hardware**]: General; C.3 [**Special Purpose and Application-Based Systems**]: Smartcards

Keywords

Optical Physical Unclonable Functions (PUFs), Machine Learning, Implementation

1. INTRODUCTION

Mobile security devices have become ubiquitous in our life. Their widespread use makes them an attractive and accessible target for adversaries. The majority of known attacks are thereby not directed against the employed cryptographic primitives themselves. Rather, they often attempt to obtain the used secret keys by physical techniques or malware. Such key-extracting strategies are not just a theoretical concern, but have been demonstrated several times in widespread, commercial systems [30, 14, 1]. The fact that the security devices shall be inexpensive aggravates the problem, leaving little room for elaborate key protection measures.

The described situation was one motivation that led to the development of *Physical Unclonable Functions (PUFs)*. A PUF is a (partly) disordered physical system that can be challenged with so-called external stimuli or challenges C_i, upon which it reacts with corresponding responses R_i. The tuples (C_i, R_i) are thereby often called the *challenge-response pairs (CRPs)* of the PUF. The responses R_i shall be a function of the applied challenge and the micro- or nanoscale structural disorder present in the PUF. It is assumed that this disorder cannot be cloned or reproduced exactly, not even by the PUF's original manufacturer. Due to its complex internal structure, it is usually harder to read out, predict, or derive the responses of a PUF than to obtain digital keys stored in standard non-volatile memory. This can make PUF-based systems more resilient against hardware and malware attacks than classical approaches.

PUFs can be used in various cryptographic and security applications. Examples include their employment as tamper sensitive secret key storage [59, 23, 19, 58], or their use as a complex identifier for a hardware system that embeds the PUF [40, 41, 18, 58]. Recent research has also discovered their application in more complex protocols such as oblivious transfer, bit commitment, key exchange, or multi-party computation [42, 3, 38]. The practical security of the latter protocols has been discussed intensively in an attempt to keep PUF theory and PUF applications closely together [46, 47, 48].

On the implementation side, most recent research has focused on electrical PUFs. The most common types include SRAM PUFs and variants thereof [23, 27, 31, 55], the Arbiter PUF and its various modifications [18, 58, 51, 34, 19, 33], and Ring Oscillator PUFs [58]. Other examples include Crossbar PUFs [45, 49], analog PUFs based on cellular non-linear networks [9], or the Bistable Ring PUF [5, 6]. Several of these electrical PUFs have been attacked by machine learning based modeling attacks [51, 53], however, and also successful cloning attacks on certain types of electrical PUFs have been reported recently [25].

Optical PUFs.

Under their original name *"Physical One-Way Functions (POWFs)"*, the optical systems proposed by Pappu et al. [40, 41] were one of the first suggested PUFs. In their original, non-integrated form, they possess several advantages:

- Low costs per piece: A non-integrated optical PUF merely consists of an inexpensive plastic platelet with randomly distributed light scatterers inside. No microelectronic or silicon circuitry on the PUF-carrying object is required.
- High output complexity: Each PUF-response consists of thousands of bits, and results from a very complex optical interference process inside the token.
- High security against modeling attacks (compare [51, 53] for such attacks on certain types of electrical PUFs). No similar attacks on non-integrated optical PUFs are known to date.
- High security against physical cloning attacks (compare [25] for such attacks on certain types of electrical PUFs). No such attacks on non-integrated optical PUFs have been reported.
- Non-integrated optical PUFs are suited as "Certifiable PUFs". It can be proven within certain limits that they have not been modified or exchanged by malicious parties (see Section 6).

On the downside, non-integrated optical PUFs à la Pappu et al. require an optical precision mechanism for read-out [40, 41]. This mechanism must establish exactly the same relative positioning of the light scattering token, the laser beam, and the CCD camera upon every single read-out. Its implementation is expensive and potentially error prone. This downside is hard to overcome, as Pappu et al.'s initial construction cannot be integrated or miniaturized easily.

Related Work and Our Contributions.

Despite their advantages, there has been surprisingly little activity on optical PUFs in recent years. Two examples are Tuyls et al. [63], who investigate the information content of optical PUFs, and Tuyls and Skoric [61], who briefly discuss integrated optical PUFs in theory, but present no experimental prototypes or security analyses [61]. A further, general source is [62]. Perhaps more activity on the security use of scattering phenomena has recently occured in a related, but not identical area. In this strand, the complex patterns emerging from laser-illuminated paper surfaces are used to authenticate objects, products, packages, documents, or digital content. Recent works on this topic have appeared in the Nature magazine [4], the IEEE Security and Privacy Symposium [8], ACM CCS [54], and other venues [57]. Furthermore, authors have suggested at CHES 2009 [24] and FC 2009 [65] how the digital content stored on compact discs can be authenticated by exploiting the individual scattering behavior of each disc.

Within this research landscape, we make the following contributions. Firstly, we investigate the optimized implementation of optical PUFs. We describe how their output complexity and security can be enhanced by simple measures such as varying the laser diameter or choosing the right size of the light scattering elements. Application of this technique allows us to increase the number of effectively independent CRPs of our non-integrated optical PUF prototype by a factor of around 24, without causing substantial additional costs. Secondly, we observe that the Gabor image transform applied by Pappu et al. [40, 41] leads to strong regularities in the derived cryptographic keys. We present a number of alternative transformations, showing that they can enhance the estimated response entropy by a factor of almost four. Our new transformations can be used with great benefits in *any* security application that exploits optical scattering phenomena, including the abovementioned approaches [4, 8, 54, 57].

Thirdly, we suggest that non-integrated optical PUFs could be used as so-called "Certifiable PUFs". As described recently at Oakland 2013 [48], Certifiable PUF are useful, and in a certain sense even necessary, to implement secure and practical PUF protocols for advanced cryptographic tasks like oblivious transfer. We describe a method for the offline certification of non-integrated optical PUFs at low costs. Our suggestion is the first construction for Certifiable PUFs in the literature.

In the second part of the paper, we investigate *integrated* optical PUFs. We build a first prototype and evalute its security. One surprising result is that integrated optical PUFs can be machine learning successfully if linear optical scattering structures are used, and if the adversary can access the raw scattering images before they are postprocessed. The quest for suited non-linear optical materials is identified as a central open research problem.

In all our analyses, real experimental data from two prototypes is used: A non-integrated optical PUF à la Pappu et al. [41, 40], and an integratable optical PUF constructed from inexpensive consumer components. Our integrated PUF prototype is the first of its kind in the literature. It functions without moving components and possesses an exponential number of challenges.

Organization of this Paper.

The first part of the paper deals with non-integrated optical PUFs: Section 2 describes our experimental set-up for this PUF-type. Section 3 discusses their security against modeling attacks. Sections 4 and 5 optimize their output complexity by tuning the laser diameter, scatterer size, and employed image transformations, respectively. Section 6 describes their use as Certifiable PUFs. Subsequently we turn to integrated optical PUFs: Section 7 details our experimental prototype of an integrated optical PUF. Section 8 reports modeling attacks on this PUF-type. We conclude the paper in Section 9.

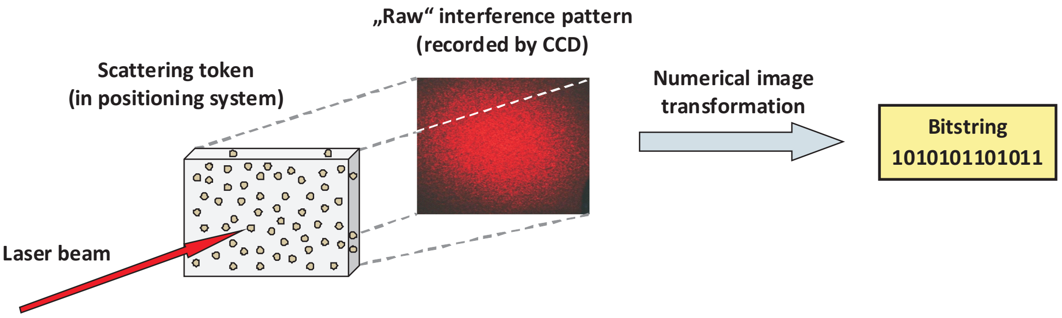

Figure 1: Schematic illustration of our implementation of a non-integrated PUF à la Pappu et al. [40, 41] (*“Pappu's PUF”*), and the necessary postprocessing for the derivation of the PUF response. As shown, we measure the intrference pattern (often also called “speckle pattern”) in transmission, whereas Pappu et al. [40, 41] measure it in reflection. A numerical image transformation is usually applied to obtain the eventual PUF response.

2. IMPLEMENTATION OF NON-INTEGRATED OPTICAL PUFS

In our implementation of non-integrated optical PUFs, we employed a slightly modified version of Pappu et al.'s original set-up [40, 41]. Instead of measuring the reflected light as in [40], we detect the speckle pattern in transmission of the light through the optical PUF. The new set-up enabled simpler and more cost-effective realizations and thus seemed preferrable. Note that said modification does not affect the PUF's input-output complexity in any way. The schematics of our set-up are given in Figure 1. A challenge C_i to this PUF consists of a specific point and angle of incidence of the applied laser beam. The corresponding response R_i of the PUF is the result of an image transformation applied to the raw speckle pattern captured by the CCD camera. Pappu et al. apply the Gabor transformation to this end [40, 41], but several alternatives are possible (see Section 5).

As light source we used a red Lasiris SNF Laser with a wavelength of 635 nm and a power of less than 5 mW. The distance between the laser and the probe is about 980 mm. The integrated optic of the laser is adjustable, which enables the focussing of the laser beam. The optical interference pattern or “speckle pattern” of the transmitted laser light is captured with a MV-D1024E camera from Photonfocus. The distance between the PUF and the camera lens is 26 mm. The integrated CMOS sensor of the camera prevents the common CCD blooming effects. For our experiments we took eight bit gray scale images with a resolution of 1024x1024 pixels. To apply different challenges to the PUF, we used a positioning system for moving and rotating the scattering token. The positioning system is based on three stepper motors. We used two linear tables of type LM 45 for horizontal and vertical movements and also one MOGO 40 to adjust the angle of the scattering token. With this system we can move the optical PUF about 12 mm in vertical and horizontal direction. We can rotate the PUF, and thus change the angle of incidence of the laser beam, by about ± 15°.

The scattering tokens were prepared by distributing glass spheres of a certain size range randomly in a transparent matrix material. Various size ranges have been examined in our experiments: We used glass spheres from Mühlmeier in the ranges 400-600μm and 300-400μm, and from Worf Glaskugeln in the ranges 250-420μm, 105-210μm, 90-106μm, and 40-80μm. From the tested matrix materials, the consumer glue UHU Plus Schnellfest was simpler to handle and led to better results than the more expensive optical adhesive NOA 61 from Norland. For the manual preparation of of the tokens, we used the following method: We applied the glue on a glass slide and scattered the glass spheres in a dense layer. After a short drying time, we apply the next layer of glue and glass spheres. Depending on the size of the spheres, we varied the number of layers between four and five, resulting in equally sized tokens of ca. 1cm$\times$1cm$\times$2.5mm.

Read-Out Stability of Our Set-Up.

To evaluate the stability of our set-up, we used our above positioning system to apply the same challenge C_0 twice, but positioned the system to some intermediate challenge C_I in between the two measurements for C_0. Due to inaccuracies in the step motors, the re-positioning to challenge C_0 is prone to small errors. After application of the Gabor transformation the observed noise level (measured in the hamming distance of the two obtained Gabor-transformed images for the challenge C_0) was about 6%. This noise level poses no problem to practical applications (compare Pappu et al. [40, 41]).

3. SECURITY OF NON-INTEGRATED OPTICAL PUFS

While successful modeling attacks have been published on several electrical PUFs [51, 53], no such attacks have ever been reported on non-integrated optical PUFs. We re-investigated the hardness of this problem, and collected 100,000 raw speckle images from each of three different scattering tokens, using our set-up of the last section. The images were generated by dividing the x- and y-axis into in 317 equal segments and by moving the laser beam to the resulting grid points, using a fixed laser angle. We then considered in theory and/or practice the applicability of machine learning (ML) algorithms to this data, including those techniques used in earlier attacks on electrical PUFs [51, 53]. All of our efforts remained unsuccessful, however. For example, the application of Support Vector Machines with linear kernels to predict single bits of the raw optical PUF output led to error rates around 50%, i.e., the prediction quality was essentially equal to random guessing [13].

There are three main theoretical reasons for the practical failure of any investigated ML method on non-integrated optical PUFs. The first is their large information content. As argued in [41], every volume unit of size around the laser wavelength of around 600 nm in principle can have an influence on the scattering process. Given the size of our tokens, in theory up to 10^{11} volume units would have to be considered. Even though this figures represent the extreme, theoretical case, also in practice the ML feature vectors are too large to be handled well. In opposition, the ML problems for

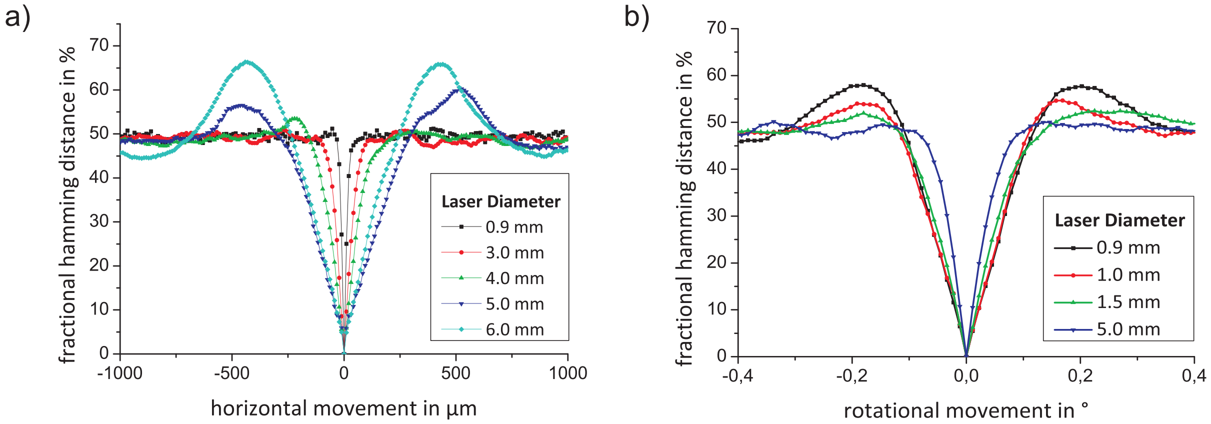

Figure 2: Decorrelation speed of the PUF output for horizontal and rotational shifts of the laser beam, as a function of different laser diameters. A fractional Hamming distance of 0.5 indicates full decorrelation between the two outputs. Smaller laser diameters cause faster decorrelation for horizontal shifts, but slower decorrelation for changes in the laser angle. A token with scatterers in the size range of 300-400μm was used.

known electrical PUFs lead to much smaller feature vectors: a 128-bit Arbiter PUF, for example, has an associated feature vector with only 129 entries [51, 53].

A second reason is the complex optical interference process inside optical PUFs. For all methods known to us, exact simulation is too laborious to be carried out in practice (compare Section 4.3). This prevents the application of ML techniques that require such simulation in the evaluation of the so-called *"fitness"* of a given ML feature vector. It rules out so-called *"Evolution Strategies"* and related methods, which had been applied successfully to electrical PUFs in the past [51, 53]. In comparison, the internal mechanisms of current electrical PUFs can often be described by simplified models, such as the well-known linear additive delay model in the case of Arbiter PUFs [51, 53]. The decisive feature of optical systems here seems that the massless photons facilitate a very complex scattering interaction, which is yet stable against varying ambient conditions, and does not alter or wear off the PUF internally. The same effect seems hard to obtain for electrical structures.

Finally, in theory each PUF-challenge should illuminate the entire scattering token. We observed that this is not the case in practice, however. The incident laser causes a light cone inside the token, meaning that for different PUF-challenges different and *independent* regions of the token are illuminated predominantly and cause the respective responses. This independence complicates or even directly prevents straightforward forms of modeling attacks. In addition, the vast number of different laser positions prohibits that the PUF output is modeled by some simple form of superposition of known signals, as it is the case for integrated optical PUFs (see Section 8). No comparable simplifications could be derived by us for non-integrated optical PUFs. Taken together, these factors result in an unprecedented level of security of non-integrated PUFs against modeling attacks.

4. ENHANCING THE CHALLENGE SPACE AND SCATTERING COMPLEXITY

Pappu et al. report that their non-integrated optical PUF possesses around 2.37×10^{10} challenges for which the corresponding Gabor-tranformed responses are virtually independent and decorrelated. Given its high resilience against modeling (see Section 3), this makes a complete read-out the currently most viable attack strategy on this PUF type. The relatively small size of the challenge space is also exploited in a number of recent quadratic attacks on PUF protocols reported at CHES 2012 and in the Journal of Cryptographic Engineering 2013 [46, 47]. In this section, we therefore systematically investigate methods to increase the number of decorrelated CRPs. The measures we discuss in Sections 4.2 and 4.3 are particularly inexpensive and simple to realize.

4.1 Influence of Multiple Laser Beams

One seemingly straightforward step to raise the size of the challenge space is to use several lasers beams with different frequencies, for example one red and one green laser. Due to the differing wavelengths, the interference pattern resulting from a green laser incident at point $\vec{p}$ and angle Θ *differs* from the pattern resulting from a red laser incident at exactly the same point and angle. It hence seems suggestive to use two such lasers, and to define one PUF challenge C_i to consist of the incidence points $\vec{p}$ and angles Θ of both lasers, i.e., $C_i := (\vec{p}_{\text{red}}, \Theta_{\text{red}}; \vec{p}_{\text{green}}, \Theta_{\text{green}})$. This promises to quadratically enhance the size of the challenge space. Similar suggestions have been existent as folklore in the community for some time.

There is a problem with this approach, however, as long as linear scattering structures are used. In this case, the interference pattern resulting from the challenge $C_i = (\vec{p_{\text{red}}}, \Theta_{\text{red}}; \vec{p}_{\text{green}}, \Theta_{\text{green}})$ is nothing else than the sum of the two patterns resulting from the challenges $C_i^{\text{red}} = (\vec{p}_{\text{red}}, \Theta_{\text{red}})$ and $C_i^{\text{green}} = (\vec{p}_{\text{green}}, \Theta_{\text{green}})$. More precisely, the intensity in each CCD pixel for the challenge C_i is exactly equal to the intensity resulting from challenge C_i^{red} plus the intensity resulting from challenge C_i^{green}. A second reason for this simple behavior is that the red and green light are not coherent as long as two separate standard lasers are used. No complex interference process can take place, and the resulting intensities simply add up in the CCD image.

This allows a simplified full read-out of the above type of PUF that possesses a red and green laser as follows: The adversary first reads out the interference patterns for all challenges from the red laser alone, then for all challenges from the green laser alone. Subsequently he can derive the patterns for all combined challenges $C_i = (\vec{p_{\text{red}}}, \Theta_{\text{red}}; \vec{p}_{\text{green}}, \Theta_{\text{green}})$ by simply adding the known pattern from the challenge $C_i^{\text{red}} = (\vec{p_{\text{red}}}, \Theta_{\text{red}})$ to the known pattern from challenge $C_i^{\text{green}} = (\vec{p}_{\text{green}}, \Theta_{\text{green}})$. Using two lasers hence effectively increases the challenge space only by a factor of about two. On the other hand, however, it results in a significantly increased implementation effort, as the two lasers need to be positionable independently from each other. A further practical aspect is that in Pappu et al.'s and our set-up, the scattering token is made

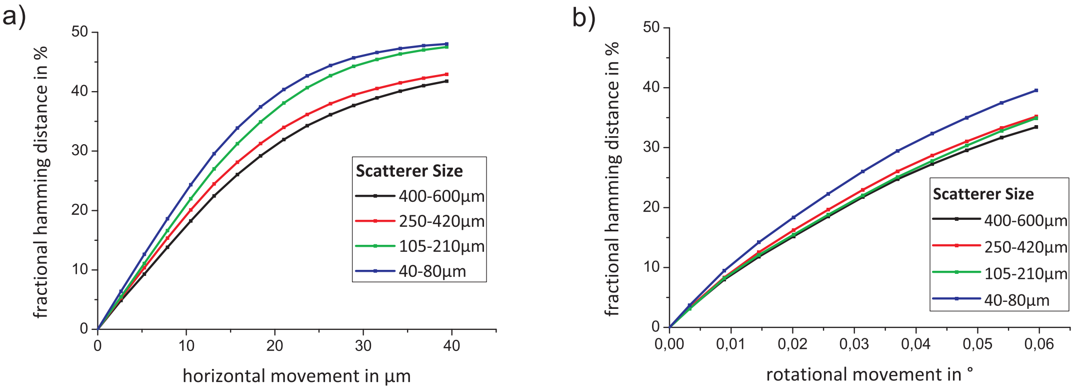

Figure 3: Decorrelation speed of the PUF output for variation of the incidence point and angle of the laser, and for different scatterer sizes. Scatterers in an overall range from 400μm to 40μm were examined. Smaller scatterers lead to faster decorrelation.

positionable and the laser is immobile. The simple reason is that the laser is much larger and heavier. With two lasers that shall generate independent challenges, this method is no longer applicable. Similar considerations hold for the case of k lasers.

To summarize, using k laser beams with a linear scattering medium only increases the effective challenge space by a factor of k, but is very expensive. Better methods are discussed over the next subsections.

4.2 Influence of the Laser Diameter

Next, we investigated the sensitivity of the PUF-output in dependence of the laser diameter to (i) variation of the point of incidence of the laser in the x-y-directions, and to (ii) alteration of its incidence angle. This sensitivity is a good measure for an optical PUF's security: It determines the number of virtually independent challenge-response pairs of the PUF, and thus its resilience against full read-out attacks.

In the set-up of Figure 1, the laser diameter can be adjusted by simply focusing the laser. The focal point lies beyond the entrance point of the PUF, and we measured the effective laser diameter at the entrance point. During our experiments we used a PUF with five layers and 300-400μm glass spheres. During the experiment, we move the PUF by the positioning system in equally spaced shifts and take pictures of the speckle patterns. The center position of the token was chosen as reference position. For all measured raw images, we then computed the Gabor transformed images as in [40, 41]. We determined their fractional Hamming distances [1] to the Gabor image of the reference position. A fractional Hamming distance of 0.5 signals a virtual decorrelation between the Gabor images.

Figures 2a and b depict our findings. For each shown curve, the critical parameter is how quickly the curve not reaches, but *settles at* the 0.5-level of the fractional Hamming distance. The faster this occurs, the more virtually independent and decorrelated challenge-response pairs of the PUF exist. The presented data illustrates that smaller diameters lead to faster decorrelation for horizontal (and also vertical) movements, but slower decorrelation in the rotational movement. This turns the choice of the optimal laser diameter into an optimization problem. In practice, its solution depends on the experimental set-up and the used materials in the scattering token. It must thus be solved in each application of optical PUFs empirically by the above method to achieve optimal security. In our case, laser diameters between 1.5 mm and 2.5 mm were optimal. Estimating from the above diagrams, we conclude that the optimal laser diameter (in comparison to a non-optimal diameter of 5.0mm or larger) can increase the effective challenge space by a factor of around 3 to 5.

4.3 Influence of the Scatterer Size

We also systematically investigated the influence of the scatterer sizes. From theoretical considerations, an optimal complexity of the scattering process can be expected for a size similar to the wavelength of the used laser light. In this case, which is also referred to as Mie-regime [36, 11, 28], the interference pattern critically depends on the particle sizes and shapes. The resulting electromagnetic field distribution within the PUF can *only* be adequately described by *exactly* solving the Maxwell equations, which is numerically very demanding. For particle sizes *significantly larger* than the optical wavelength, simplified methods such as ray tracing approaches become valid [20], which avoid the numerical complexity involved in solving the full Maxwell equations. On the other hand, for particles or refractive index inhomogeneities on a length scale *significantly smaller* than the wavelength, the scattered light field can be described by an approximation to the Mie solution, the much simpler Rayleigh scattering [28]. In the extreme case of a large ensemble of far sub-wavelength objects, the single scatterers cannot be optically resolved at all. The medium is then comprehensively described by a volume-averaged effective refractive index, again exhibiting low optical complexity.

This suggests that for our red laser of wavelength 635nm, particles of approximately the same sizes should be used to facilitate maximally complex scattering phenomena. Note that this is about a factor of 1,000 smaller than the sizes used by Pappu et al. in their original experiments [40, 41]. The experimental exploration of this regime requires nanofabrication techniques, which we did not have available at the time of writing. In order to lead at least a qualitive proof of concept that optical PUFs become increasingly complex for smaller scatterers, we carried out experiments with size ranges 400-600μm, 250-420μm, 105-210μm, and 40-80μm (compare Section 2). These could be handled by our manual fabrication methods. The particles we used are still up to a factor of ten smaller than the scatterers employed by Pappu et al. [40, 41].

As formal measure for the resulting internal complexity of the scattering process, we used again the sensitivity of the speckle pattern against variations in the laser coordinates and angles. We determined how quickly the transformed images decorrelate in prac-

[1] The *fractional Hamming distance* of two bitstrings of the same length l is the number of all bits on which the two strings differ divided by the length l of the string(s).

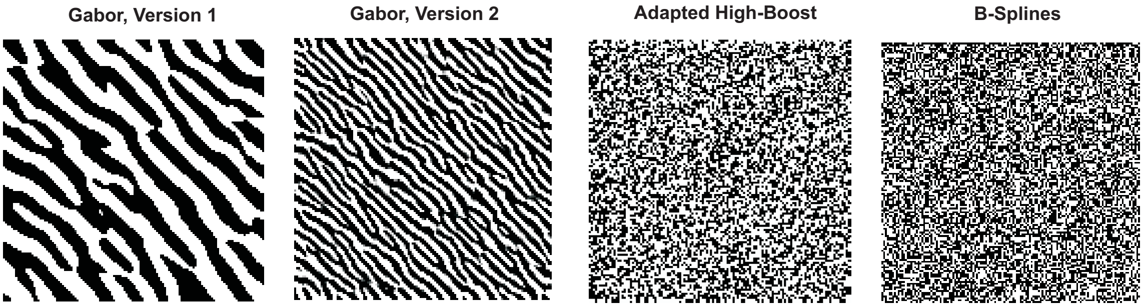

Figure 4: Images obtained from different image transformations. The two leftmost images are obtained from the Gabor transformation for different parameters, illustrating the typical, zebra stripe like regularities. The two right images are obtained from our adapted high-boost transformation and from B-spline wavelets. They cause considerably less regularities.

tice, depicting our findings in Figure 3. They confirm that in the length regimes examined by us, the internal complexity steadily increases for smaller scatterers. They also show that the effective challenge space can be increased by choosing smaller scatterers. Extrapolating from our above diagrams, using scatterers of size 40-80μm instead of 500-800μm as in Pappu et al. [41, 40] improves the size of the challenge space by about a factor of two in the horizontal directions, and the same in the vertical direction and for varying angles. This will lead to an overall improvement of a factor around eight. Probably yet better improvements could be achieved in the nano-regime.

Our conclusion is that the use of small scatterers is a very simple and inexpensive means of increasing the security of optical PUFs. In practical applications, they should hence be chosen as small as possible under the given fabrication, cost, and stability constraints.

5. ENHANCING RESPONSE ENTROPY

As already noted by Pappu et al. [40, 41], the "raw" optical interference images (as recorded by a CCD camera) should not be used directly as response of an optical PUF, as they are too large and unstable. A numeric transformation must be applied to distill shorter, more stable bitstrings. Pappu et al. utilize the well-established Gabor transform to this end. This transformation is applied to the raw CCD images, and the resulting the two-dimensional Gabor-transformed images (after a threshold step for conversion into binary data) are subsequently be converted into a cryptographic key by simply reading out the transformed images line by line.

There are, however, two downsides of this approach. The Gabor transformed images have very strong, zebra-stripe like regularities (see Figure 4). These regularities cause strong patterns in the cryptographic keys, i.e., they lead to regularly alternating, medium-length sequences of consecutive ones and zeros in the keys. Secondly, since the Gabor images do not contain maximal entropy (as they exhibit said patterns), they do not reflect the small-scale physical randomness of the optical PUF in an optimal way. This makes it in principle easier to build PUFs with the same challenge-response behavior. In order to achieve a maximal security level against cloning, an optimal bitwise entropy in the PUF responses should be achieved.

The problem of regularities in the Gabor images has also been observed in [60]. As a countermeasure, the authors propose to choose a random subset of the bits of the Gabor image. However, their positions of these bits need to be stored together with the bit values themselves, increasing storage requirements. Furthermore, it is non-trivial how to select bits that are stable and carry a large amount of information at the same time. Finally, the method does not address the problem that much information about the unclonable random structure is wasted by the Gabor transform. Instead of applying some postprocessing to the Gabor transform, we believe that a better approach is to use alternative image transformations from the start.

From theoretical considerations [10], so-called wavelet transformations seemed good candidates, since they induce less structure in the transformed image. We empirically tested a considerable number of transformations from this family for their practical performance on optical PUFs, including Daubechies wavelets [10], symlets [10], polyharmonic isotropic B-spline wavelets [64] and quincunx wavelets based on on the McLellan transformation [15]. We empirically optimized the exact parameters of the transformations during our experiments. As expected, all tested transformations could significantly increase the bitwise entropy of the PUF responses, with the B-spline wavelets being best in class. All of them exhibited slightly worse robustness than the Gabor-transformation, however. Even with our relatively inexpensive set-up and stepper motors, this posed no practical problem. But in order to be equipped for any type of application scenario, we devised a new transformation in-house [26]. It realizes essentially the same stability level as the Gabor transformation, but still causes much less regularities in the transformed images.

In our method, which we call *adapted high-boost transform (AHB)*, we used a convolution kernel, which compares the intensity of a pixel against its neighborhood. As in the Gabor transformation, the conversion into a binary key is accomplished by a threshold filter. The method calculates the arithmetic mean of the intensity of pixels in the neighborhood of a center pixel. If the intensity of the center pixel plus an adjustable offset is smaller than the arithmetic mean, the pixel is set to "1", otherwise to "0".

In practice, this can be implemented with a simple convolution matrix A and a threshold t. For the case of a 3×3 matrix the convolution kernel is a 2D Laplace filter which also response to 45° edges. The definition for a $n \times n$ convolution kernel A is given as:

$$A = \left.\underbrace{\begin{bmatrix} 1 & 1 & \cdots & 1 & \cdots & 1 & 1 \\ \vdots & \vdots & \ddots & 1 & \ddots & \vdots & \vdots \\ 1 & 1 & \cdots & -n^2+1 & \cdots & 1 & 1 \\ \vdots & \vdots & \ddots & 1 & \ddots & \vdots & \vdots \\ 1 & 1 & \cdots & 1 & \cdots & 1 & 1 \end{bmatrix}}_{n}\right\} n$$

The binary key of speckle image $I(u, v)$ can be extracted by reading the binary image $B(x, y)$ line by line. We obtained the best results by a 7×7 convolution matrix. Filters comparable to ours have been termed "high-boost" in [22], hence the name AHB.

$$B(x,y) = \begin{cases} 1, & \text{if } \int_v \int_u I(u,v) \cdot A((x-u),(y-v))\, du\, dv > t \\ 0, & \text{if } \int_v \int_u I(u,v) \cdot A((x-u),(y-v))\, du\, dv \leq t \end{cases}$$

Before the convolution matrix is applied, the picture can be scaled to a desired size. The number of scales gives, as in the Gabor transformation, the level of the transformation. The scaling algorithm has a significant effect on the robustness and information density. The implementation was realized with the OpenCV with a pixel area relation, which is the preferred method for image decimation that gives Moiré-free results [37].

Figure 4 qualitatively illustrates the difference in entropy and randomness for the applied image transformations, which is already visible well by the sheer eye. A more quantitative analysis can be obtained by applying a procedure already suggested by Pappu et al. (compare [41] and section 8.1 of [40]): One collects a large sequence of transformed images (for m randomly chosen challenges, say), and compares the statistical variance to the variance that would have occured if the sequence would have been generated by an ideal binomial distribution of length k. As argued by Pappu et al. [40, 41], this can be used as an estimator for the number of independent bits. We applied this procedure to our case, collecting 400 transformed images. The used challenges were distributed equidistantly over the token, and the angle of incidence was perpendicular to the token surface. We obtained the results of Table 1. They quantitatively underpin the significant improvement of our new transformations.

Image Transform	Gabor	AHB	B-Splines
Estimated Independent Bits	25%	90%	94%

Table 1: Quantitative comparison of image transformations. The number of independent bits is estimated by the method of Pappu et al. [40, 41] (see above).

We stress again that the AHB transform has essentially the same robustness as the Gabor transform. The B-splines exhibited better stability than most other tested wavelet transformations, but slighty worse robustness than the Gabor transform and our AHB transform. Details are given in in Appendix A and Figure 7.

6. USE OF NON-INTEGRATED PUFS AS CERTIFIABLE PUFS

In a recent paper at Oakland 2013 [48], two new, practically relevant attack models on Strong PUF protocols have been discussed. In one of the new attack models, the so-called *"bad PUF model"*, malicious parties may replace PUFs by other, malicious hardware which looks like a PUF from the outside, but possesses hidden extra properties that allow cheating. This approach represents a practically viable attack strategy as long as the PUF cannot be "certified" or "attested" for being unmanipulated and benign. The attack method is most relevant for integrated electrical Strong PUFs, since they communicate with external parties merely via a digital challenge-response interface. What is behind the interface hence remains hard to detect or verify for users.

One specific example of bad PUFs are so-called *"Simulatable PUFs"*. These are hardware systems which look like a proper PUF from the outside, but possess a simulation code by which the manufacturer (or other malicious parties) can simulate the PUF-responses to arbitrary challenges. They can hence obtain challenges without being in physical possession of the PUF. One way of a realizing a simulatable bad PUF is to construct a hardware system that looks like an integrated electrical PUF from the outside, possessing a standard CRP interface. But internally, the system generates the PUF-responses by use of a numerical pseudo-random number generator or a pseudo-random function, whose secret seed is known to the manufacturer/the malicious party. The tacit use of such simulatable bad PUFs can spoil the security of several PUF protocols, for example schemes for oblivious transfer, as fully detailed in [48].

As a countermeasure, the authors of [48] propose the design and use of *"Certifiable PUFs"*: These are PUFs for which it can be verified that they do possess (at least some of) the expected properties, for example, that they are unpredictable for all parties and that they have not been manipulated or altered after their production. Up to this date, however, no strategies to "certify" PUFs in the above sense have been proposed in the literature. Electrical integrated PUFs seem very hard to certify in said manner, since they are accessed via a digital challenge-response interface (see above). Non-integrated optical PUFs show better potential, since their complex analog responses are hard to imitate for malicious PUFs, and are measured directly, i.e. not through any digital interface. We follow this line of thought and present below the first scheme for the offline certification of non-integrated optical PUFs. It uses an oblivious transfer protocol of Rührmair and van Dijk as basis, which employs interactive hashing as a substep (see [47] for details).

Protocol 1: SECURE OBLIVIOUS TRANSFER BASED ON CERTIFIABLE OPTICAL PUFS

Set-Up Assumptions:

- The used optical PUF is fabricated by a manufacturer who is trusted by both the OT-sender and the OT-receiver.
- The manufacturer uses a digital signature scheme $\mathsf{DS}_{\mathsf{Man}}$ with signing key SK and corresponding verification key VK.
- VK is known to both the OT-sender and the OT-receiver.

Pre-Protocol Steps:

- The manufacturer fabricates the optical PUF. He applies k randomly chosen challenges $C_1, \ldots, C_k$ to the PUF, for k being a small, one-digit security parameter. He obtains the responses $R_1, \ldots, R_k$.
- He generates the signature $\mathsf{Sig} := \mathsf{DS}_{\mathsf{Man}}(C_1, \ldots, C_k, R_1, \ldots, R_k)$, and defines the certificate as
$$\mathsf{Cert} := (C_1, \ldots, C_k, R_1, \ldots, R_k, \mathsf{Sig})\,.$$
- The PUF is distributed together with its certificate to the OT-receiver. In practice, the certificate can be stored inexpensively via a barcode, for example, which is printed on the item in which the optical PUF is embedded.

Protocol:

Let the sender's input be two strings $s_0, s_1 \in \{0,1\}^\lambda$ and the receiver's input be a bit $b \in \{0,1\}$. The protocol then proceeds as follows:

1. The receiver verifies the certificate of the PUF. To that end, he applies the challenges $C_1, \ldots, C_k$ to the PUF, and verifies that the obtained responses are equal to $R_1, \ldots, R_k$.

2. The receiver chooses a challenge $c \in \{0,1\}^\lambda$ uniformly at random. He applies c to the PUF, obtaining the response r. He transfers the PUF together with the certificate to the sender.

3. The sender verifies the certificate of the PUF in the same manner as above.

4. The sender and receiver execute an IH protocol, where the receiver has input c. Both get outputs c_0, c_1. Let i be the value where $c_i = c$.

5. The receiver sends $b' := b \oplus i$ to the sender.

6. The sender applies the challenges c_0 and c_1 to the PUF. Denote the corresponding responses as r_0 and r_1.

7. The sender sends $S_0 := s_0 \oplus r_{b'}$ and $S_1 := s_1 \oplus r_{1-b'}$ to receiver.

8. The receiver recovers the string s_b that depends on his choice bit b as $S_b \oplus r = s_b \oplus r_{b \oplus b'} \oplus r = s_b \oplus r_i \oplus r = s_b$.

The above certification step only works since the analog responses of the PUF are measured by the involved parties themselves. The raw, two-dimensional speckle patterns are too complex to be imitated by a malicious, bad PUF. Therefore a verification of a very small number of sample CRPs suffices to exclude that the PUF has been altered or exchanged against another PUF. The verification can be executed offline, i.e., without additional communication with the manufacturer. As already noted by Rührmair and van Dijk, this feature is essential: If an online communication with a trusted authority would be a regular step in the protocol, then the OT could be executed much simpler via this trusted authority itself.

It is interesting to consider the above protocol under the aspect of the involved computational or other assumptions. The protocol requires two assumptions: (i) an unpredictable PUF (see [43, 3] for formal definitions of the latter); (ii) a secure digital signature scheme $\mathsf{DS}_{\mathsf{Man}}$. It uses these two assumptions to implement OT. It is long known that secure digital signature schemes exist if and only if one-way functions exist [21], but currently no construction is known that implements OT merely from one-way functions. The use of unpredictable, certifiable PUFs hence is necessary and creates additional value in the protocol.

The technique of digitally signing a unique, reflective speckle pattern seems also promising in combination with future generations of electrical Erasable PUFs [48, 44]. The speckle pattern could be recorded directly from the surface of the electrical PUF, or an optical encapsulation could be used that enables certification of the Erasable PUF inside. This eventually seems a promising technique to finally realize PUFs that are *both* erasable and certifiable, as required in [48].

Example Implementation.

We carried out an example implementation of certifiable optical PUFs by use of 2D barcodes. We chose the libdmtx library [32] for the implementation of the widely used Data Matrix Code. In order to save area requirements for the barcode, our implementation is based on the bilinear pairing based scheme by Zhang, Safavi-Naini und Susilo (ZSS) [66]. For the implementation we chose the PBC library [39] with the elliptic curve type F and a signature of 200 bits. We assumed that the following information must be stored on the product: Manufacturer ID (16 bit), PUF ID (48 bit), Signature (200 bits), and image transformed speckle pattern. With a barcode module width of 0.25 mm this leads to a barcode of size on the order of 1 cm^2. Barcodes of these sizes can easily and inexpensively accompany optical PUFs in practice, for example on bank cards.

7. IMPLEMENTATION OF INTEGRATED OPTICAL PUFS

One central, practically motivated research goal is to embed optical PUFs into microelectronic systems, i.e., to design secure integratable optical PUFs. Pappu's PUF is not very well suited to this end, since it requires movable components that must be positioned with high accuracy. A first miniaturized version of Pappu's optical PUF, which was briefly reported in [56], thus had to use expensive piezo positioners. Such constructions can merely achieve comparably slow read-out speeds, and still exhibit a polynomial number of challenges only, exactly like Pappu et al.'s original optical PUF. Are there other possibilities?

General Architectures of Integrated Optical PUFs.

Figures 5a and 5b describe two possible approaches. The goal of the constructions is to design optical PUFs without moving components which still allow an exponential number of different challenges. Figure 5a shows an immobile laser diode array with k phase-locked diodes $D_1, \ldots, D_k$ [67], which is used to excite a disordered, random scattering medium. The diodes can be switched on and off independently, leading to 2^k challenges C_i. These can be written as $C_i = (b_1, \ldots, b_k)$, where each $b_i \in \{0,1\}$ indicates whether diode D_i is switched on or off. At the right hand side of the system, an array of l light sensors $S_1, \ldots, S_l$, e.g. photodiodes, measures the resulting light intensities locally. A response R_i consist of the intensities $I_1, \ldots, I_l$ in the l sensors.

As depicted in Figure 5b, instead of phase-locked diode arrays, also a single laser source with a subsequently placed, inexpensive light modulator (LCD array) can be employed. Comparable suggestions have been made earlier by Gassend [17] and by Tuyls and Skoric [61] in theory, i.e., without experimental realization. The k pixels of the LCD can be switched on and off independently, again leading to 2^k possible challenges. The whole system can be encapsulated by a reflective layer to facilitate the internal interference process. Both systems easily lend themselves to miniaturization.

In order to allow optical interference and to generate complex behavior also with linear scattering media, *all* laser light inside the scattering structure must be coherent. This necessitates the use of a phase-locked diode array (as in Figure 5a) or the employment of only one single laser source plus a subsequent light modulator/LCD (as in Figure 5b).

Our Prototype.

We built the first prototype of the above class of integratable optical PUFs from commercial components, including an LCD array from a customary beamer. The aim was not yet miniaturization, but a first proof of concept and a subsequent security analysis. The schematics are given in Figure 5c.

As light source we used HRP050 HeNe laser from Thorlabs with a wavelength of 632 nm and a power of 5.0 mW. The laser beam was widened to approximately 3/4 the size of the LCD array using an assembly of a bi-concave lense ($f = -25.0$mm) and a plano-convex lense ($f = 150.0$mm) with a space of 125.00mm inbetween. In a miniaturized system it is unnecessary to widen the laser beam, as done here using lenses, since the mirrored encapsulation as shown in Figure 5b would ensure that the scattered beam

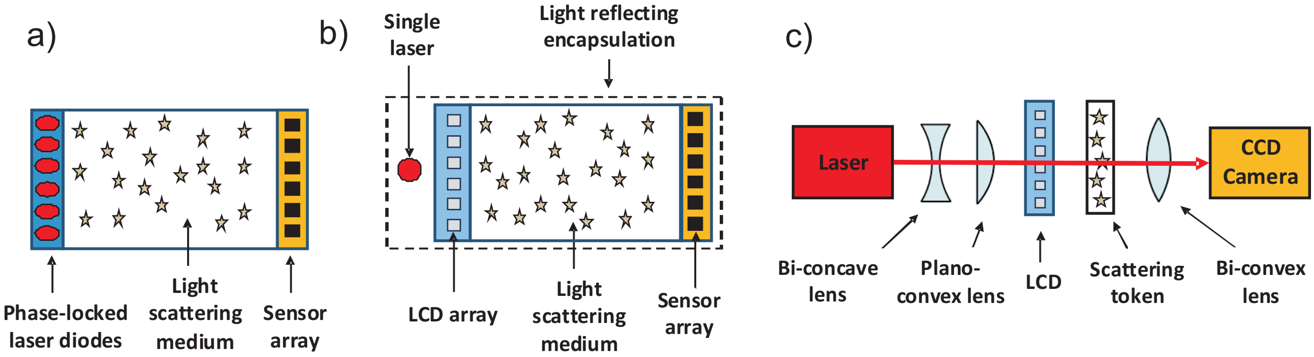

Figure 5: a) and b): Two possible theoretical types of integrated optical PUFs (compare [61, 17]). c): Schematic illustration of our prototype (not true to scale).

eventually passes through the whole LCD array and scattering token.

The widened beam was then modulated as it passed an LCD array, which had a resolution of 1024x768 pixels and was extracted, together with the associated control electronics, from a commercial Geha compact 640 LCD projector. Since the laser illuminated only 3/4 of the LCD array, approximately only 750x550 pixels were useable for the modulation of the beam. During initial testing we determined that flipping *single* pixels on the LCD did not result in a significant change of the optical response as recorded by the camera. Thus in order to get a detectable influence of each single challenge bit on the optical response we grouped adjacent pixels into blocks of 35 $\times$ 35 pixels. Hence the useable area of the LCD array was divided into 15 $\times$ 15 $=$ 225 rectangular blocks. The k-th block was associated with the k-th bit b_k of the PUF-challenge. This bit determined whether *all* pixels of the whole k-th block were switched on or off. This methods leads to PUF-challenges C_i of length 225 bits and a challenges space of 2^{225}. As scattering objects we used the same structures as in Section 2, specifically glass spheres by Mühlmeier with a size range of 300-400μm.

The modulated beam passes through the scattering structure and is then fed using a bi-convex lens (f = 25.4 mm) onto an Edmund Optics EO-0413BL Monochrome CMOS sensor with a resolution of 752$\times$480 pixels, where the response of the optical system is recorded. Of those pixels an area of about 40,000 pixels were illuminated by the beam. As bleeding and overexposure lead to nonlinearities, only 38,663 were used effectively.

The whole setup was controlled by a standard PC, and the LCD array was driven using the VGA output port of the PC. We measured that after a change of the VGA output signal the used LCD array requires about 13ms until the new output is fully displayed and the picture is stable. In order to ensure that we do not record the output of the optical system while the LCD is still transitioning between the current and the previous challenge we always waited 30ms before recording the optical response of a new challenge.

Read-Out Stability of our Set-Up.

Similar as in Section 2, we investigated the read-out stability of our prototype by applying a challenge C_0 followed by an intermediate challenge C_I followed by switching back to C_0. We found that even without subsequent error correction, the euclidian distance $d_{\text{Euclidian}}(p,q) := \sqrt{\sum_{i=1}^{n}(q_i - p_i)^2}$ in the raw speckle images was on average only 0.80% per pixel (highest value here 0.84% per pixel). This illustrates one first significant advantage of integrated optical PUFs: As they have no moving parts, they can achieve unprecedented stability levels.

8. SECURITY OF INTEGRATED OPTICAL PUFS

After our prototype was functional, we investigated the security of integrated optical PUFs. We found that under the provision that a *linear* scattering medium is used in the integrated optical PUFs of Figure 5, the following analysis holds.

Consider that all blocks on the LCD array are turned off, so that no light can pass through them, except the block corresponding to challenge bit b_j, which is turned on. Then the electric field at each detector cell i is the result of the laser beam passing through LCD block b_i and being scattered by the optical token. Thus the amplitude E_i of the electric field at the CMOS cell i is given by

$$E_i = T_{ij} e^{i\omega t},$$

where $T_{ij} \in \mathbb{C}$. Since the optical medium is linear, the electric fields at the CMOS cells combine linearly if we turn on more than one LCD block, thus we have

$$E_i = \sum_{j=1}^{N} T_{ij} b_j e^{i\omega t},$$

where b_j is 0 if block j is turned off and 1 if it is turned on. The corresponding intensities are given by

$$I_i = |E_i|^2 = |\sum_{j=1}^{N} T_{ij} b_j|^2.$$

To faciliate linear learning, this can be written as

$$I_i = \sum_{j=1}^{N} \sum_{k=1}^{N} T_{ij} T_{ik}^* b_j b_k,$$

since $b_k \in \{0,1\}$. By defining the $M \times N^2$ matrix $R_{i,j\cdot N+k} = T_{ij} T_{ik}^*$ and the vector $\beta_{j\cdot N+k} = b_j b_k$ with N^2 elements, we can rewrite the resulting intensities as

$$I_i = \sum_{l=1}^{N^2} R_{il} \beta_l \,, \tag{1}$$

and thus we see that there is a linear relationship between the measured pixel intensities on the CMOS sensor and the state of the

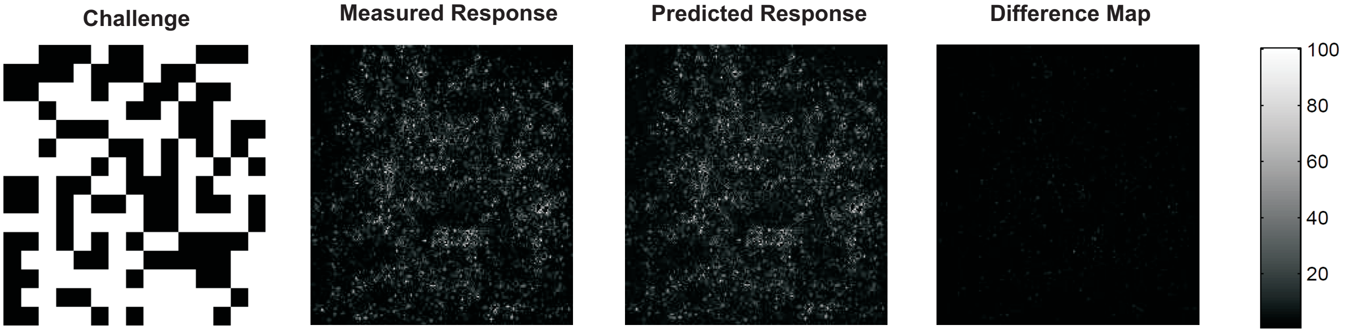

Figure 6: A randomly chosen 15 × 15 excitation pattern or challenge to the PUF; a CCD image of the response of the optical integrated PUF; the numerically predicted response; and the difference map between the latter two.

LCD blocks. To determine R we present $l = 53701$ randomly chosen challenges $\vec{b}^{(n)}, n = \{1, \ldots, l\}$ on the LCD and record the corresponding optical responses $\vec{I}^{(n)}$. During the recording of the reponses of some challenges some pixels of the CMOS sensors were overexposed or underexposed and thus the intensities I_i returned by the sensor for these pixels were clipped. We cannot expect a linear relation of these clipped intensities to the challenge bits and therefore we ignore all intensities that are either over- or underexposed in any response. Calculation of R from the acquired challenge response pairs (CRPs) is then done using standard linear regression with the squared error function

$$E(R) = \frac{1}{2}\sum_{n=1}^{l} |\vec{I}^{(n)} - R\vec{\beta}^{(n)}|_2^2,$$

where $\vec{\beta}^{(n)}$ is defined as above. Its minimum is given by [2]

$$R^* = B^\dagger \psi$$

where $B^\dagger = (B^T B)^{-1} B^T$ is the Moore-Penrose pseudo-inverse of B, $B_{ni} = \beta_i^{(n)}$ and $\psi_{jn} = I_j^{(n)}$ is the intensity measured at the jth CCD cell when the nth challenge is applied. Once the matrix T is known, the simulation of a response $R_{C_m} = (I_1, \ldots, I_l)$ to a given challenge $C_m = (b_1, \ldots, b_k)$ can be executed by simple calculation following (1).

We applied this strategy described to data that was collected from our prototype of Figure 5c. For evaluation 1% of the recorded CRPs were withheld from training and used as a test set. The success for two different excitation patterns is shown in Figure 6. The difference map between the actually acquired optical image and the prediction shows that the deviations are extremely small. The average error for a test set of 300 CRPs was $d_{\text{Euclidian}}(p, q) = 1.23\%$. The values for the example shown in Figure 6 are

$$d_{\text{Euclidian}}(p, q) = 1.21\%$$

per pixel. These differences are exactly in the range of recording a response for the same challenge twice (compare last section).

We stress again that our above attack assumes that linear scattering structures are used, and that the attacker has access to the raw speckle images that are produced by the setup. The latter occurs if the attacked can invasively probe the PUF, if the postprocessing is carried out outside the optical PUF, or if the raw speckle images (or other unprocessed sensor data) are directly used as PUF-output. The latter two cases are a realistic scenario for integrated optical PUFs due to their stability, and since one would like to save computational resources for postprocessing inside the PUF.

One direct consequence of our findings is that in general, nonlinear optical materials must be used in integrated optical PUFs of the above type to achieve maximal security. The identification of suited substances constitutes an important open problem, which we pose to the community in this work. *Linear* integrated optical PUFs seem only secure as long as they are used within a secure perimeter and with additional postprocessing to the PUF responses, e.g. within Controlled PUF architectures (compare [17]).

9. SUMMARY

We revisited integrated and non-integrated optical PUFs, their optimal implementation, and their security in this paper, drawing on a large basis of experimental data from two dedicated prototypes. We began our journey with *non-integrated* optical PUFs à la Pappu et al. [40, 41]. Using data from our prototype, we analyzed the security of these PUFs against machine-learning based modeling attacks, finding no vulnerabilities at all. The most relevant attack point on these PUFs hence remains their comparatively low number of decorrelated challenge-response pairs. For this reason, we next investigated simple and inexpensive measures to enlarge the effective challenge space. They included the use of multiple laser beams, optimizing the laser diameter, and choosing well-suited scatterer sizes. It turned out that the latter two steps can achieve a better enhancement than multiple laser beams. Conservatively estimated, they could enlarge the challenge space for our examined set-ups and systems by an overall factor of around 3×8 = 24, as detailed in Sections 4.2 and 4.3. At the same time, they are much cheaper and simpler to realize than multiple lasers. We then investigated new image transformations that can improve the bit entropy of PUF-responses. Our motivation was that the Gabor transformation leads to strong regularities in the transformed images and the derived cryptographic keys. We showed that new transformations can get rid of such obvious patterns and can increase the bitwise response entropy by a factor of almost four. The methodology we introduced in the respective sections can be applied to any practical or commercial optical PUF systems in order to optimize their security with inexpensive means. The new image transformations will also be useful in other optical security applications, for example in the use of scattering images of randomly structured paper surfaces [4, 8, 54].

Subsequently, we revealed an entirely new application of Pappu et al.'s non-integrated optical PUFs as so-called *"Certifiable PUFs"*. A few digitally signed responses of these PUFs can serve as a fingerprint that certifies their input-output complexity and non-simulatability. Assuming trust in the manufacturer who issues the signature, and assuming the possession of an associated public verification key, this allows the *offline* certification of non-integrated optical PUFs. It enables their secure one-time use in advanced protocols such as oblivious transfer [38, 48]. Non-integrated optical PUFs are uniquely qualified for this approach: The resulting speckle patterns are too complicated to physically reproduce for a fraudster with a malicious system like a bad PUF, even if the patterns are known to him. This is in strong contrast to the simple, single-bit digital outputs of integrated electrical PUFs. Furthermore, the optical responses are measured directly from the non-integrated optical PUF, and are not communicated via a (potentially malicious) digital interface. Our construction is the first Certifiable PUF, addressing an open question posed at Oakland 2013 [48].

In the final part of the paper, we turned to *integrated* optical PUFs without moving components and exponential challenge spaces, and presented the first prototype of this kind. It was not yet embedded into a microelectronic system, since this was not the goal of this paper, but it easily lends itself to miniaturization. We used our set-up to examine the security of this PUF type, and surprisingly found that it can be successfully machine learned under two premises: (i) A linear scattering structure is used. (ii) The adversary has direct access to the resulting raw speckle images. We argued why this case is realistic in practice, and gave a theoretical security analysis. We proved the validity of the analysis in practice by predicting entire raw speckle images with extremely high accuracy. Our findings enforce the use of non-linear scattering materials in this PUF type. The search for suitable non-linear optical materials which must be stable, non-toxic, inexpensive, and should exhibit their non-linearities already at low light intensities, is posed as an essential open problem in this work.

Overall, our investigations show that there are some very good reasons to study and optimize optical PUFs further. The input/output complexity of non-integrated optical PUFs is simply unmatched, thus overcoming the security problems of electrical Strong PUFs against modeling attacks [51, 53]; their isotropically disordered 3D structure and their complex responses gives them an extreme security against cloning, surpassing the recent cloning attacks on electrical PUF-types with one challenge and 1-bit responses [25]; they are the very first PUFs for which “certification” or “attestation” is possible, a feature that has not been realized for any other class of PUFs yet; the use of non-linear materials promises inexpensive integrated optical PUFs, overcoming the practical downsides of Pappu at al.’s non-integrated approach [41, 40]; and, last but not least, they allow fascinating research at the cross-section of security, embedded systems, machine learning, and nanotechnology.

10. REFERENCES

[1] R. J. Anderson: *Security Engineering: A guide to building dependable distributed systems*. Wiley, 2010.

[2] C. M. Bishop: *Pattern Recognition and Machine Learning (Information Science and Statistics*. Springer-Verlag New York, Inc., 2006.

[3] C. Brzuska, M. Fischlin, H. Schröder, S. Katzenbeisser: *Physical Unclonable Functions in the Universal Composition Framework.* CRYPTO 2011.

[4] J. Buchanan, R. Cowburn, A. Jausovec, D. Petit, P. Seem, G. Xiong, D. Atkinson, K. Fenton, D. Allwood, and M. Bryan: *Forgery: Fingerprinting documents and packaging.* Nature, vol. 436, 2005.

[5] Q. Chen, G. Csaba, P. Lugli, U. Schlichtmann, U. Rührmair: *The Bistable Ring PUF: A New Architecture for Strong Physical Unclonable Functions*. HOST 2011.

[6] Q. Chen, G. Csaba, P. Lugli, U. Schlichtmann, U. Rührmair: *Characterization of the Bistable Ring PUF*. DATE 2012.

[7] C. Christensen: *Frames and Bases*. Birkäuser, Boston, 2008.

[8] W. Clarkson, T. Weyrich, A. Finkelstein, N. Heninger, J. Halderman, E. Felten: *Fingerprinting blank paper using commodity scanners.* IEEE S&P, pp. 301-314, 2009.

[9] G. Csaba, X. Ju, Z. Ma, Q. Chen, W. Porod, J. Schmidhuber, U. Schlichtmann, P. Lugli, U. Rührmair: *Application of Mismatched Cellular Nonlinear Networks for Physical Cryptography*. IEEE CNNA, 2010.

[10] I. Daubechies: *Ten lectures on wavelets*. Society for industrial and applied mathematics (SIAM), 1992.

[11] W. Demtröder: *Experimentalphysik 2: Elektrizität und Optik*. Springer 2004. ISBN- 10: 3540202102.

[12] M. van Dijk, U. Rührmair: *Physical Unclonable Functions in Cryptographic Protocols: Security Proofs and Impossibility Results.* Cryptology ePrint Archive, Report 2012/228, 2012.

[13] E. Dinter: *Physikalische Einwegfunktionen.* Diplomarbeit, Technische Universität München, 2009.

[14] T. Eisenbarth, T. Kasper, A. Moradi, C. Paar, M. Salmasizadeh, M. T. Manzuri Shalmani: *On the Power of Power Analysis in the Real World: A Complete Break of the KeeLoqCode Hopping Scheme*. CRYPTO 2008.

[15] M. Feilner, D. Van De Ville, and M. Unser: *An orthogonal family of quincunx wavelets with continuously adjustable order*. IEEE Trans. on Image Processing, 2005.

[16] B. Forster, P. Massopust (Eds.): *Four Short Courses in Harmonic Analysis. Wavelets, Frames, Time Frequency Methods, and Applications to Signal and Image Analysis*. Birkhäuser, 2009.

[17] B. Gassend, *Physical Random Functions*. MSc Thesis, MIT, 2003.

[18] B. Gassend, D. E. Clarke, M. van Dijk, S. Devadas: *Silicon physical random functions*. ACM Conference on Computer and Communications Security 2002, pp. 148-160, 2002

[19] B. Gassend, D. Lim, D. Clarke, M. van Dijk, S. Devadas: *Identification and authentication of integrated circuits.* Concurrency & Computation: Practice & Experience, 2004.

[20] A. S. Glassner: *An introduction to ray tracing*. Morgan Kaufmann Pub., 1989.

[21] O. Goldreich: *The Foundations of Cryptography – Volume 2*. ISBN 0-521-83084-2, Cambridge University Press, 2004.

[22] R.C. Gonzales, R.E. Woods: *Digital Image Processing (2nd Edition)*. Prentice Hall, 2002.

[23] J. Guajardo, S. S. Kumar, G. J. Schrijen, P. Tuyls: *FPGA Intrinsic PUFs and Their Use for IP Protection*. CHES 2007.

[24] G. Hammouri, A. Dana, B. Sunar: *CDs have fingerprints too.* CHES 2009.

[25] C. Helfmeier, D. Nedospasov, C. Boit and J.-P. Seifert: *Cloning Physically Unclonable Functions*. HOST 2013, to appear.

[26] C. Hilgers: *Praktische Realisierung von Verfahren aus der Physikalischen Kryptographie*. MSc Thesis, Technische Universität München, 2009.

[27] D. E. Holcomb, W. P. Burleson, K. Fu: *Power-Up SRAM State as an Identifying Fingerprint and Source of True Random Numbers*. IEEE Trans. Computers, 2009.

[28] A. Ishimaru: *Wave propagation and scattering in random media*. Vol. 1 & 2. New York: Academic press, 1978.

[29] C. Jaeger, M. Algasiner, U. Rührmair, G. Csaba, M. Stutzmann: *Random pn-junctions for physical cryptography*. Applied Physics Letter 96, 172103, 2010.

[30] T. Kasper, M. Silbermann, C. Paar: *All You Can Eat or*

Breaking a Real-World Contactless Payment System. Financial Cryptography and Data Security (FC), 2010.

[31] S. Kumar, J. Guajardo, R. Maes, G. J. Schrijen, P. Tuyls: *The Butterfly PUF: Protecting IP on every FPGA.* HOST 2008.

[32] The libdmtx library. See http://www.libdmtx.org/

[33] J.-W. Lee, D. Lim, B. Gassend, G. E. Suh, M. van Dijk, and S. Devadas. *A technique to build a secret key in integrated circuits with identification and authentication applications.* IEEE VLSI Circuits Symposium, 2004.

[34] M. Majzoobi, F. Koushanfar, M. Potkonjak: *Lightweight Secure PUFs.* IC-CAD 2008: 607-673.

[35] S. Mallat: *A wavelet tour of signal processing*. Academic Press, San Diego, 1997.

[36] G. Mie: *Beiträge zur Optik trüber Medien, speziell kolloidaler Metallösungen.*, Annalen der Physik 330(3), 1908.

[37] OpenCV documentation, see http://docs.opencv.org/ modules/imgproc/doc/geometric_transformations.html

[38] R. Ostrovsky, A. Scafuro, I. Visconti, A. Wadia: *Universally Composable Secure Computation with (Malicious) Physically Uncloneable Functions*. Eurocrypt 2012.

[39] The pairing based cryptography library (PBC). See http://crypto.stanford.edu/pbc/

[40] R. Pappu: *Physical One-Way Functions*. PhD Thesis, Massachusetts Institute of Technology, 2001.

[41] R. Pappu, B. Recht, J. Taylor, N. Gershenfeld: *Physical One-Way Functions*, Science, vol. 297, 2002.

[42] U. Rührmair: *Oblivious Transfer based on Physical Unclonable Functions*. TRUST 2010.

[43] U. Rührmair, H. Busch, S. Katzenbeisser: *Strong PUFs: Models, Constructions and Security Proofs.* In A.-R. Sadeghi, P. Tuyls (Editors): Towards Hardware Intrinsic Security: Foundation and Practice. Springer, 2010.

[44] U. Rührmair, C. Jaeger, M. Algasinger: *An Attack on PUF-based Session Key Exchange, and a Hardware-based Countermeasure: Erasable PUFs*. Financial Cryptography and Data Security (FC), 2011.

[45] U. Rührmair, C. Jaeger, M. Bator, M. Stutzmann, P. Lugli, and G. Csaba: *Applications of high-capacity crossbar memories in cryptography*. IEEE Trans. on Nanotechn., 2011.

[46] U. Rührmair, M. van Dijk: *Practical Security Analysis of PUF-based Two-Player Protocols*. CHES 2012.

[47] U. Rührmair, M. van Dijk: *On the Practical Use of Physical Unclonable Functions in Oblivious Transfer and Bit Commitment Protocols*. Journal of Cryptogr. Engin., 2013.

[48] U. Rührmair, M. van Dijk: *PUFs in Security Protocols: Attack Models and Security Evaluations*. IEEE S&P, 2013.

[49] U. Rührmair, C. Jaeger, C. Hilgers, M. Algasinger, G. Csaba, and M. Stutzmann: *Security applications of diodes with unique current-voltage characteristics.* Financial Cryptography and Data Security (FC), 2010.

[50] U. Rührmair, S. Devadas, F. Koushanfar: *Security based on Physical Unclonability and Disorder.* In M. Tehranipoor and C. Wang (Editors): Introduction to Hardware Security and Trust. Springer, 2011

[51] U. Rührmair, F. Sehnke, J. Sölter, G. Dror, S. Devadas, J. Schmidhuber: *Modeling Attacks on Physical Unclonable Functions*. ACM CCS, 2010.

[52] U. Rührmair, J. Sölter, F. Sehnke: *On the Foundations of Physical Unclonable Functions.* IACR Cryptology ePrint Archive, Report 2009/277, 2009.

[53] U. Rührmair, J. Sölter, F. Sehnke, X. Xu, A. Mahmoud, V. Stoyanova, G. Dror, J. Schmidhuber, W. Burleson, S. Devadas: *PUF Modeling Attacks on Simulated and Silicon Data*. IACR Cryptology ePrint Archive, Report 2013/112, 2013.

[54] A. Sharma, L. Subramanian, E. A. Brewer: *PaperSpeckle: microscopic fingerprinting of paper.* ACM CCS, 2011.

[55] P. Simons, E. v.d. Sluis, V. v.d. Leest: *Buskeeper PUFs, a promising alternative to D Flip-Flop PUFs.* HOST 2012.

[56] B. Skoric, G.J. Schrijen, W. Ophey, R. Wolters, N. Verhaegh, J. van Geloven: *Experimental hardware for coating PUFs and optical PUFs*. In: P. Tuyls, B. Skoric, T. Kevenaar (Ed.): Security with Noisy Data (pp. 255-268). Springer London, 2007.

[57] J. R. Smith, A. V. Sutherland: *Microstructure-Based Indicia.* Second Workshop on Automatic Identification Advanced Technologies, Morristown, 1999.

[58] G. E. Suh, S. Devadas: *Physical Unclonable Functions for Device Authentication and Secret Key Generation.* DAC 2007.

[59] P. Tuyls, G. J. Schrijen, B. Skoric, J. van Geloven, N. Verhaegh, R. Wolters *Read-Proof Hardware from Protective Coatings.* CHES 2006.

[60] P. Tuyls, G.J. Schrijen, F. Willems, T. Ignatenko, B. Skoric: *Secure key storage with PUFs*. In: P. Tuyls, B. Skoric, T. Kevenaar (Ed.): Security with Noisy Data (pp. 255-268). Springer London, 2007.

[61] P. Tuyls, B. Skoric: *Strong Authentication with Physical Unclonable Functions.* In: Security, Privacy and Trust in Modern Data Management, M. Petkovic, W. Jonker (Eds.), Springer, 2007.

[62] P. Tuyls, B. Skoric, T. Kevenaar: *Security with Noisy Data*. Springer London, 2007.

[63] P. Tuyls, B. Skoric, S. Stallinga, and A. Akkermans, and W. Ophey: *Information-theoretic security analysis of physical uncloneable functions*. Financial Cryptography and Data Security (FC2005), 2005.

[64] D. Van De Ville, T. Blu, and M. Unser: *Isotropic polyharmonic B-Splines: Scaling functions and wavelets*. IEEE Trans. on Image Processing, 2005.

[65] D. Vijaywargi, D. Lewis, D. Kirovski: *Optical DNA*. Financial Cryptography and Data Security (FC), 2009.

[66] F. Zhang, R. Safavi-Naini, W. Susilo: *An efficient signature scheme from bilinear pairings and its applications.* Public Key Cryptography (PKC), 2004.

[67] D. Zhou, L.J. Mawst: *Two-dimensional phase-locked antiguided vertical-cavity surfaceemitting laser arrays*. Applied Physics Letters, 77(15), pp. 2307-2309, 2000.

APPENDIX

A. ERROR CORRECTION CAPACITY OF TESTED IMAGE TRANSFORMATIONS

Figure 7 depicts the error correction capacity of the six image transformations we evaluated in our experiments. As described already in Section 5, these were Daubechies wavelets [10], symlets [10], polyharmonic isotropic B-spline wavelets [64] and quincunx wavelets based on on the McLellan transformation [15], together with our own, adapted high-boost transform (AHB) and the Gabor transform as base value. The error correction capacity is measured by examining how quickly the transformed images decorrelate for small horizontal misplacements of the token.

The figure shows that our AHB transform has virtually the same error correction behavior as the Gabor transform, while it leads to much more randomness and entropy in the images (see Section 5 and Table 1). From the other transforms, B-splines led to the largest entropy (compare again Section 5 and Table 1), but interestingly still had better robustness than other transforms. This illustrates well that instability and induced response entropy are not the same.

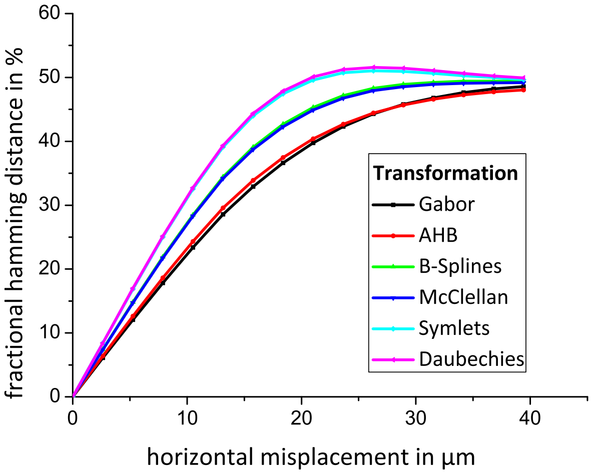

Figure 7: Robustness of the various image transformations tested by us against horizontal misplacement of the probe.