Glitter PUF: A Passive Anti-Tamper PUF Based On Images Of Glitter Reflections

Moeskops, N.; Aljuffri, A.; Hamdioui, S.; Taouil, M.

DOI
10.1109/ITC58126.2025.00057

Publication date
2025

Document Version
Final published version

Published in
Proceedings of the 2025 IEEE International Test Conference (ITC)

Glitter PUF: A Passive Anti-Tamper PUF Based On Images Of Glitter Reflections

Noeël Moeskops, Abdullah Aljuffri, Said Hamdioui and Mottaqiallah Taouil
Computer Engineering Lab, Delft University of Technology, Delft, the Netherlands

2025 IEEE International Test Conference (ITC) | 979-8-3315-7041-5/25/$31.00 ©2025 IEEE | DOI: 10.1109/ITC58126.2025.00057

***Abstract*—In this paper, we introduce a novel passive physical anti-tampering Physical Unclonable Function (PUF) based on glitters that can protect an entire Integrated Circuit (IC) and/or Printable Circuit Board (PCB). A prototype of the proposed glitter based PUF has been developed. The glitters are dropped randomly in a resin layer during its formation and their positioning is used as the basis of a PUF. The PUF response is created by taking a picture inside the coating layer. To get a stable response resilient against noise and different temperature cycles, the picture is processed using filtering, image processing, and error correction. Using actual drill measurements, our findings indicate that even drilling with a 0.1mm diameter drill can be detected and lead to a wrong PUF response.**

***Index Terms*—Glitter PUF, anti-tampering, passive PUF.**

I. Introduction

The field of electronic device security goes beyond protecting digital assets; it also critically includes the physical security of the devices themselves. Physical security is a cornerstone of ensuring the overall integrity of electronic systems, as unauthorized physical access can inflict damages comparable to, or even exceeding, the impact of cyberattacks. To mitigate this, an anti-tamper technique can be applied, like those described in [1], [2] and [3].

Several anti-tamper technologies exist, which can be divided into two categories: active and passive anti-tampering schemes. Active anti-tampering schemes require a power source to operate, whereas passive schemes do not. A common active approach is to use a protective foil over and under the most critical parts of the design [4], [5], [6]. By continuously monitoring either the resistance or capacitance of the foil, intruders can be detected. Active anti-tamper schemes have a major drawback: power might be disabled or not available during tampering (e.g., during transportation), leaving the system vulnerable. A passive system, on the other hand, can still protect the system even when powered off. Once the system is started or used, a tampering check can be performed. Unfortunately, there are limited passive anti-tampering methods. One example is the B-TREPID envelope [7], which provides a wrapper around the device that has a unique capacitive property that changes when modified. The capacitance is measured during startup and is used to decrypt the content of the device if it remains intact. In order for an attacker to make alterations to the design, the foil and (thus the key) has to be destroyed. The downside is that the key needs to be reliably extracted in the presence of noise and during temperature changes / cycles. Furthermore, this foil only protects a single IC (or limited area) on the PCB. Hence, better schemes are needed.

This paper presents a novel, fully operational, passive (i.e., battery-less) anti-tamper method. The PUF consists of glitters placed randomly in a resin layer around a chip or PCB. The PUF responses consist of images taken from a camera inside the resin layer. The proposed technique has been validated by means of experimentation.

The remainder of this paper is organized as follows. Section II presents the proposed Glitter PUF concept and design. Section III characterizes the Glitter PUF. Section IV provides details on the Glitter PUF implementation and results. Finally, Section V concludes this paper.

II. Glitter PUF Concept and Design

A. Attack Model

We assume that the target of the adversary is (i) to alter the functionality of the device once deployed in the field (by e.g., adding hardware Trojans) and (ii) to extract secret data such as cryptographic keys or information which is only available at runtime. Additionally, we assume that the adversary has physical access to the device to perform such attacks.

B. Glitter PUF Concept

The Glitter PUF response is generated by taking an image that contains information of the complete chip / PCB which is independent from the technology node of the chips. The core principle of this approach is that any tampering would undoubtedly destroy the visual aspects of the system. To prevent relying on a static image of the product, which could easily be reproduced by an attacker, the image must capture something inherently unique per device. This unique feature must be intrinsically tied to the device and irreversibly destroyed in the event of tampering.

Fig. 1 illustrates a side-view diagram of the PUF. The design works as follows: a wide-angle camera is placed in the middle of the secure PCB, surrounded by Light Emitting Diode (LED)s. Transparent resin mixed with tiny glitter particles are poured over the PCB, filling the entire volume. The system is then closed using a light-tight box. This build procedure is the same for every device. However, each time with a different glitter composition, and thus with a different image. Furthermore, the physical glitter particles composition cannot be replicated once settled. This intrinsic randomness makes it a good PUF candidate. Next, the camera is used to take pictures.

This work is funded by "Resilient Trust" project of the EU's Horizon Europe research and innovation programme, grant agreement No. 101112282.

Poster

2378-2250/25/$31.00 ©2025 IEEE
DOI 10.1109/ITC58126.2025.00057

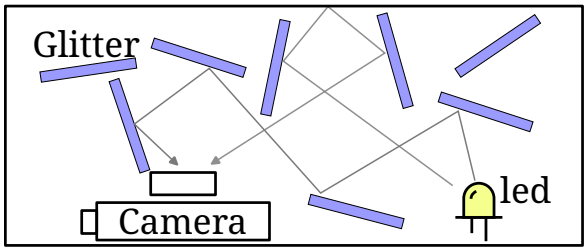

Fig. 1: Sideview of the Glitter PUF

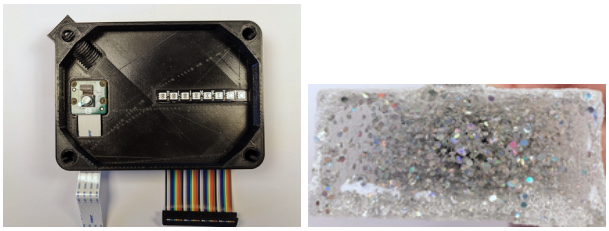

(a) Without Resin. (b) Glitters Inside Resin Layer.

Fig. 2: Front View of the PUF Enclosure When Opened.

The picture response is used to derive a key; the key can be used to decrypt sensitive encrypted data from the memory at start-up. When an adversary attempts to tamper with the device, the resin and glitter composition will change and will lead to a different picture response. By using glitters, the goal is to maximize light reflections inside the resin and make sure that each pixel response contains as much information as possible. In an ideal scenario, all the pixels of the image should be influenced by all glitter positions. This will detect any drilling attacks from adversaries, even when the drilling happens outside the Field-of-View (FoV) of the camera.

In addition to ensuring data confidentiality, the Glitter PUF can be used to verify the integrity and authenticity of the device by hashing the PUF response and use that to compare to the baseline response. The passive PUF can also work in an active manner by regularly checking the picture response. Fig. 2 shows the prototype implementation used for experimentation in this paper.

III. Glitter PUF characterization

A. Image Capturing Settings

Optimal camera settings are determined using Shannon Entropy to maximize captured information. Key parameters explored include shutter time (ST) and exposure time. To reduce noise, we oversample images by averaging 20 responses.

Using 5 LEDS in the PUF, we tested 31 LED configurations (excluding all-off) and 21 shutter speeds, totaling 651 configurations. Each is evaluated using the entropy (based on the average of 20 samples) and Structural Similarity Index Measure (SSIM) [8]. Typically, 40–55s exposures yield the best results, except when all LEDs are on, where 20–24s works better. Long exposures are impractical in production but acceptable for this prototype. The high exposure demand may result from glitter particles obstructing weak LED light, which could be improved with smaller, more transparent particles or brighter LEDs.

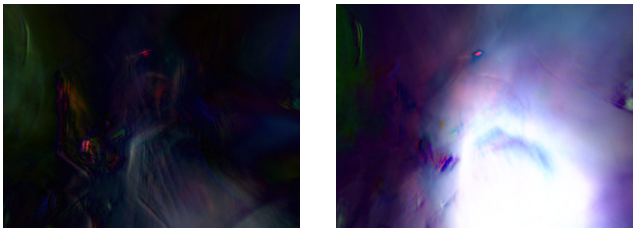

(a) With a 0.8mm Drill. (b) With a 2mm Drill.

Fig. 3: Color Differences per Pixel (Scaled 25x).

A picture can hide information in low-contrast parts that may be important for tamper detection, especially when the drilling is localized. This is not desired, as these low-contrast parts might contain vital information of the image. To spread the information equally, a contrast equalizer can be used. The Shannon Entropy of pictures where adaptive equalization (AE) is applied to is approximately 3 times bigger than without this enhancement technique. Meaning that the pixel values differ more and that thus possibly more information is in the picture.

B. Aging Test

The aging test was designed to simulate aging and to figure out how picture responses change over time. To accelerate aging, we performed 23 temperature cycle tests. In each cycle we controlled the environment temperature using an oven/freezer from 22°C to -10°to 60°C. After experimentation the average error compared to the initial response is flatting out and stabilizing, which indicates that a certain amount of ECC is sufficient to guarantee functional operation.

As discussed in [9], a stable or ”golden” sample of the hardware must be taken in order to get a stable consistent response from the PUF. For the Glitter PUF these golden pixels can be determined by going through several temperature cycles (20+) and checking which pixels change the least. A new enrollment picture should be based after the temperature cycling test is completed, as this stabilizes the PUF and reduces the Error Correcting Code (ECC) requirement. We used an enrollment temperature of 22 °C.

C. Drill Test

Penetration tests have been performed using drills with diameters 0.1mm, 0.5mm, 0.8mm, 1mm, 1.5mm, 2mm and 3mm in ascending order. Pictures are taken after each drill twice, one after drilling halfway and one when the resin has been completely drilled. The pictures after drilling are subsequently compared to the reference.

Fig. 3 shows the absolute pixel differences between the reference picture and the picture taken after a hole was drilled with a 0.8mm and 2mm drill. The figure shows a clear difference in the bottom right part as the drill size grows. The left part of the image remains almost unchanged and from an optical point of view, has barely been impacted.

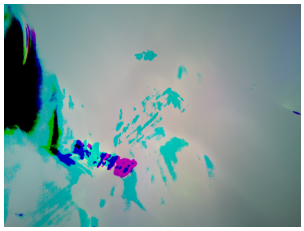

(a) 90% Most Stable Pixels.

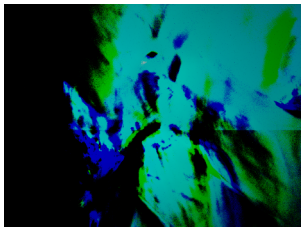

(b) 30% Most Stable Pixels.

Fig. 4: PUF Images Showing Only the Most Stable Channels.

The more white the pixels are, the larger the difference, as all Red Green Blue (RGB) channels changed. This behavior differs from the aging experiment, where mostly only a single channel is impacted. The pixels that are affected by drilling are typically stable during aging and with an appropriate pixel noise scheme, tampering detection can be improved.

IV. Glitter PUF Implementation & Results

The PUF has four stages: increase Signal-to-noise ratio (SNR), compression, error correction, and key acquisition.

A. Increasing SNR

Twenty samples are taken using the camera and averaged. This averaged image is further processed to increase the SNR by using AE. Moreover, data filtering is needed to only allow the most stable parts of the image.

1) Data Filtering Techniques: Some parts of the image are more susceptible to noise than others, see Fig. 4. Filtering unstable channels can greatly enhance the usability of the data. The total error of a picture can be brought down by selecting the best pixels and channels for authentication. One method could be to let the device go through numerous temperature cycles and then check the ones which change the least. Dividing the image at sub pixel level (i.e. Red, Green and Blue channels) is important as color channels may have different stability. An example of this can be seen in Fig. 4.

2) Data Filtering Amount: Excluding less stable pixel regions significantly improves the PUF's reliability but results in missing information. Ideally, the color distance per pixel between drilling and reference is as big as possible, as drilling should be detected. Truncating the Least Significant Bit (LSB) bits (different colored lines) reduces the distance and hence the ECC requirement. Finding a balance between these two is therefore necessary. From experimentation, we observe that taking between 40% and 50% of the most stable channels pixels lead to the best results. Taking more cells into account increases the noise (and hence more ECC is needed), while reducing this amount leads to information loss and hence a lower tampering detectability.

3) Data Filtering Enhancement: Using different post-processing steps on the image could possibly give rise to smaller detail changes. This is, however, a two edge sword, as the impact of noise and temperature changes is also amplified. However, the difference from drilling might be more significant. AE worked best for getting the most detail out of the

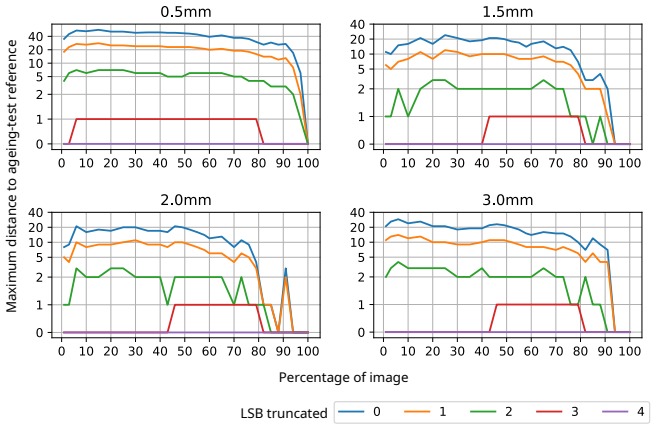

Fig. 5: Color Difference When Drilling Using AE

(a) 20% Most Stable Cells with AE and Greyscaling.

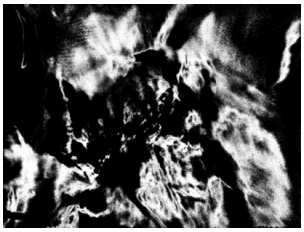

(b) Parts of Figure 6a that Failed After ECC.

Fig. 6: Image After 1mm Drilling (Single Light, 45s ST).

picture and gives stable results. However, the overall error with AE is also large. Still using this equalization method could be beneficial since the stable percentage of the image used can be reduced to 20%, which will also reduce the ECC requirement.

B. Compression

As discussed previously, only 20% of the stable sub-pixels are considered of the taken images. Moreover, the 2 LSB bits of each color channel of each pixel are truncated as well. After truncating the 2 LSB bits a color difference of (63_{10}) only 5 bits needs to be corrected for errors. When using AE (20% of the image with the two LSBs truncated) only 15% of the images are used. This value can be further reduced by converting the image to grey scale. In that case only 5% of the original data is required. It is important that the weights of the RGB color channels are properly tweaked when an image is converted to greyscale. Having weights of $1/3$ for each color channel is not the best thing to go for as it does not account for the stability of each channel. In order to properly convert the color image to greyscale the accuracy of the channels is measured by weighting the total color distance of each channel. This results in a final pixel value of: $B * 0.37 + G * 0.36 + R * 0.27$.

C. Error Correction

As discussed earlier, an error is described as the distance to the correct value. Therefore, in order to have a correct error correcting scheme, only the distance to the source

Poster

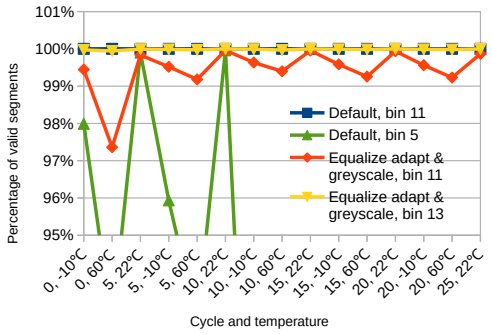

(a) Various Temperature Cycles

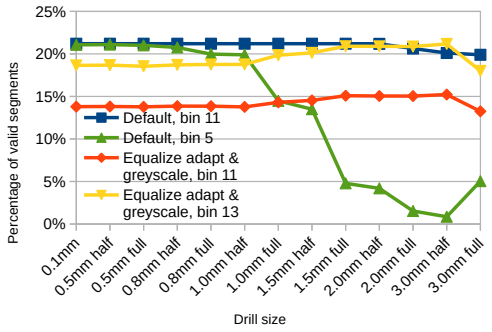

(b) Drilling with Various Sizes

Fig. 7: Prototype Results Using Average Glitter PUF Settings.

should matter. In order to accomplish this, a binning algorithm is used as shown by Equations 1 and 2. These equations are applied on the reference picture. In these equations, b is a byte (range 0-255), v is the binned value, s size of the bin, and o is the offset (between $-(\frac{s}{2})$ and $\frac{s}{2}$).

$$v = \frac{b}{s} \quad (1) \qquad o = \begin{cases} -\frac{b \bmod s}{2}, & \text{if } \frac{b \bmod s}{2} < \frac{s}{2} \\ \frac{b \bmod s}{2}, & \text{otherwise} \end{cases} \quad (2)$$

D. Key Acquisition

Faulty pixel responses are not uniformly divided after attacks. As can be seen in Fig. 3, drilling causes a big local error, while random noise through time and temperature deviations occurs spread out over the image. To detect changes locally, the image can be subdivided into a grid where each part of the grid holds a part of the secret key. As long as one part fails due to drilling, the key cannot be reconstructed anymore. As long as the relation between the pixels and the key is obstructed. One of the most secure methods is to XOR the grids together to get a new value. In the picture, a total of approximately 538795 segments ($(3280 * 2464 * 0.2)/3$) are present; each segment contains 3 grayscaled pixels that are XORed with a 24 bit random key value. These random values can be further used to generate a proper key length (e.g., 256 bit for Advanced Encryption Standard (AES)). We chose to concatenate all 24-bit segments (consisting of a random public key XORed with image responses) and subsequently hash this value using a secure hash algorithm (i.e., SHA-256).

E. Final Results

Figure 7 shows the percentage of segments that were successfully recovered for different temperatures, drill sizes, and configurations. The equalization AE with a bin size of 11 also (ever so slightly) suffers from some segments filling during temperate cycle tests. Using AE with a bin size of 13 performs similarly to a bin size of 11. However, it only requires 13.33% of the data needed as compared to the default setting. Hence, it improves the memory usage and key derivation time.

Fig. 7b shows that even a small hole triggers an invalid response, as only a small fraction of segments leads to correct responses. Comparing this predrilling to the reference picture already shows a similar result in succeeding segments as with a 0.1mm drill. Looking at the default configuration with a bin size of 5 in Fig. 7b, a sharp decrease in correct segments starts to occur around the 1mm drill size.

V. Conclusion & Discussion

This research presented a novel, low-cost passive anti-tampering technique based on a unique glitter-based PUF design. The study demonstrated that with minimal hardware components such as a camera, a resin layer, and glitter particles, an effective anti-tampering system can be developed. This approach is capable of detecting drill holes as small as 0.1mm, effectively preventing tampering attempts and securing the device's integrity. Our goal was to demonstrate that creating such a PUF is feasible and not necessarily fully optimizing it. Using different light sources, camera's ST, and experimenting with glitter positions could significantly improve the PUF quality, reduce ST and the required ECC. In conclusion, the Glitter PUF has shown considerable promise as a passive anti-tampering mechanism. However, more testing is needed to ensure a practical and secure implementation.

References

[1] T. Mosavirik *et al.*, "Silicon Echoes: Non-Invasive Trojan and Tamper Detection using Frequency-Selective Impedance Analysis," *Transactions on Cryptographic Hardware and Embedded Systems*, 2023.

[2] F. Zhang *et al.*, "Robust counterfeit pcb detection exploiting intrinsic trace impedance variations," in *IEEE 33rd VLSI Test Symposium (VTS)*, 2015.

[3] V. Immler *et al.*, "Secure Physical Enclosures from Covers with Tamper-Resistance," *Transactions on Cryptographic Hardware and Embedded Systems*, 2018.

[4] C. Gaine *et al.*, "Active shielding against physical attacks by observation and fault injection: Chaxa," *Journal of Hardware and Systems Security*, 2023.

[5] D.-C. Vasile *et al.*, "Protecting the secrets: Advanced technique for active tamper detection systems," 2019.

[6] H. Eren *et al.*, "Fringe-effect capacitive proximity sensors for tamper proof enclosures," in *Sensors for Industry Conference*, 2005.

[7] V. Immler *et al.*, "B-trepid: Batteryless tamper-resistant envelope with a puf and integrity detection," in *IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*, 2018.

[8] Z. Wang *et al.*, "Image quality assessment: From error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.

[9] H. Ning *et al.*, "Physical unclonable function: Architectures, applications and challenges for dependable security," *IET Circuits, Devices & Systems*, 2020.

Poster