Lukasz MAKOWSKI

Warsaw University of Technology, Institute of Theory of Electrical Engineering, Measurement and Information Systems
ORCID: 0000-0002-4586-6890

doi:10.15199/48.2025.03.53

Evaluation of computational performance of PUF implementation for IoT device authentication

***Abstract**. One of the major challenges in the Internet of Things (IoT) is effectively authenticating large numbers of active node devices without relying on information stored on them and without requiring technical knowledge from users. Physical Unclonable Functions (PUFs) offer a viable alternative by exploiting the unique physical properties of electrical circuits to securely and reliably generate device-specific signatures. This paper examines the computational efficiency of processing data streams from a nonlinear oscillator circuit used to generate identifiers on a microcontroller.*

***Streszczenie.** Jednym z głównych wyzwań w Internecie Rzeczy (IoT) jest skuteczne uwierzytelnianie dużej liczby aktywnych urządzeń bez polegania na przechowywanych na nich informacjach i bez wymogu posiadania wiedzy technicznej przez użytkowników. Fizyczne funkcje nieklonowalne (PUF) oferują realną alternatywę, wykorzystując unikalne właściwości fizyczne obwodów elektrycznych do bezpiecznego i niezawodnego generowania sygnatur specyficznych dla urządzeń. Niniejsza praca bada efektywność obliczeniową przetwarzania strumieni danych z obwodu oscylatora nieliniowego do generowania identyfikatorów na mikrokontrolerze. (**Ocena wydajności obliczeniowej implementacji PUF do uwierzytelniania urządzeń IoT**)*

Keywords: Internet of Things, authentication, performance measurement, nonlinear systems
Słowa kluczowe: Internet Rzeczy, autentykacja, pomiary wydajności, systemy nieliniowe

Introduction

Since its inception, the Internet of Things (IoT) has become increasingly prevalent in diverse sectors such as industry, home automation, agriculture, and medicine. These IoT systems handle sensitive and private data, including business activities and personal biometrics. Consequently, one of the foremost challenges encountered by designers, administrators, and developers in the IoT domain is ensuring the robust security of these devices [1].

An attacker deploying a device that mimics a genuine IoT node poses a serious security threat to the IoT system. The operator of this counterfeit node could potentially gain unauthorized access to sensitive information within the system or corrupt it with maliciously crafted data. Therefore, reliable authentication of these devices before they can communicate with the network is crucial [2].

In conventional networks, certificates and public-private key pairs are well-established solutions. However, implementing these methods in IoT systems presents significant challenges. IoT devices often operate under resource constraints, limiting their computational power and memory, which makes traditional security measures difficult to deploy [3]. Additionally, these solutions often demand technical expertise from end users, which may exceed general expectations.

Storing read-only numeric identifiers during the production stage of devices provides a straightforward initial layer of security for end users but is accompanied by significant drawbacks. Inexpensive identifiers, simply randomized during production and deeply embedded in the device's hardware, are prone to collisions, with the likelihood increasing as more devices are deployed. Ensuring true uniqueness, however, substantially increases production costs and requires oversight by a trusted authority to manage the issuance process. Despite these measures, the identifiers still remain susceptible to spoofing and malicious exploitation, ultimately undermining the effectiveness of traditional security mechanisms.

An alternative solution is Physically Unclonable Functions (PUFs), which are investigated in this research. PUFs offer a method to enhance the security of Internet of Things (IoT) devices by generating unique identifiers on demand [4, 5]. This approach eliminates the need for long-term storage of identifiers and makes devices resistant to duplication. Even when using electronic circuits with identical nominal values, PUFs produce unique identifiers because minute variations in the actual values of components arise from their inherent sensitivity to environmental conditions and manufacturing processes [6].

The remainder of this paper is organized as follows: Section 2 briefly reviews related work on IoT device authentication and PUF technology. Section 3 details the methodology, focusing on the implementation of the PUF on an ARM Cortex-M microcontroller, and presents the experimental setup. Section 4 analyzes the results and discusses their implications for IoT security. Finally, Section 5 concludes the paper and outlines potential directions for future research.

Physically Unclonable Functions

Implementing PUFs as electronic circuits is a favorable choice because it enables the production of devices with truly unique identifiers at a low cost and on a large scale. This approach is well-established and proven in practice of IoT devices [7].

PUFs are typically implemented as nonlinear electrical circuits, which benefit from being both cost-effective and scalable. These circuits operate in the meta-stable region, generating signals that, while seemingly complex, are not completely random. The generator discussed in this article could stably produce a tone at a specific frequency. However, modifications to the circuit, under specific circumstances, can displace its operating point from this simple equilibrium, resulting in a signal of greater intricacy. The shape of the resulting signal is not random but is deterministically influenced by the values of the circuit components and the supply voltage. Despite this complexity, the generator can reliably produce the same signal during consecutive runs for a brief period after power-up, after which the output transitions into true chaos.

One might expect that re-implementing the circuit with components of identical nominal values would yield a circuit generating signal identical to the original. However, even the highest quality electronic components have variations in their actual values, and each component introduces a degree of freedom in the circuit around its operating point. This variability, combined with the circuit's operation in the meta-stable region, makes it impossible to perfectly

replicate or copy such a circuit as even the smallest change has significant effect [8].

PUFs can also be implemented as integrated circuits, which significantly increases the difficulty of obtaining detailed information about the component values used in the circuit, thereby minimizing the chances of an attack. This approach also allows for the utilization of more subtle component characteristics, such as the individual delay times of logic gates [4]. Due to cost constraints research presented here will utilize a circuit made of discrete elements.

Methodology and implementation

The circuit used in this research is based on a reference design, which was extensively analyzed in a previously published article by its authors [9]. The adaptation employed here incorporates a more modern transistor and makes corresponding adjustments to the RC network elements. The schematic is shown in Figure 1.

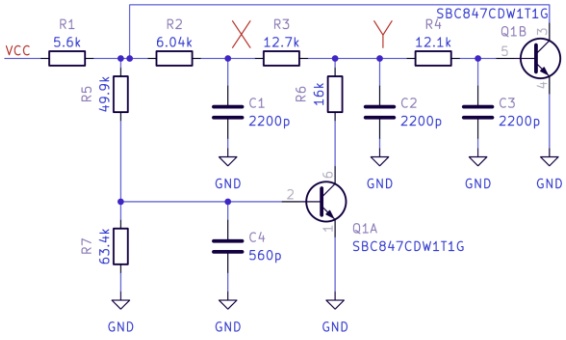

Fig.1. Schematic of the chaotic generator (adapted from [9])

The circuit is based on a phase-lag 3-stage RC network oscillator, with a transistor serving as the amplifying element. This core configuration is augmented by the addition of another RC network and a second transistor of the same type. The auxiliary network's role is to draw a small amount of charge from the main RC oscillator, thereby perturbing its operation and pushing it into a meta-stable region. While multiple points in the circuit can be considered as outputs, the most significant signals are observed at two specific capacitors (C1 and C2) in the main oscillator section. The test points are labeled X and Y in Figure 1.

The signals generated by this circuit can have relatively low voltage levels, necessitating conditioning to ensure they are suitable for sampling by modern microcontrollers. To achieve this, the circuit incorporates several operational amplifiers that serve multiple functions: buffering the signal to prevent loading of the original source, shifting the signal level, and amplifying it. Precision operational amplifiers (ADA4622) are employed as unity-gain buffers for the raw signals X and Y. Each signal is then conditioned using low-noise amplifiers (ADA4807), with one serving as a reference voltage and two others configured in a differential amplifier setup.

To ensure the circuit operates as intended, it was extensively simulated using LTSpice with Monte Carlo and worst-case analysis methods. Figure 2 shows five representative simulated signals obtained at the channel X.

Four distinct stages can be observed in the signal plot. Initially, all capacitors are discharged. When the power supply is activated, the output signal enters its transient state, characterized by fluctuations between extreme values until the operational amplifiers stabilize after about 0.18 ms.

The second stage is the oscillation growth phase, during which the circuit begins oscillating at its base frequency. Before reaching a steady state, the amplitude continues to increase, as expected from this type of generator.

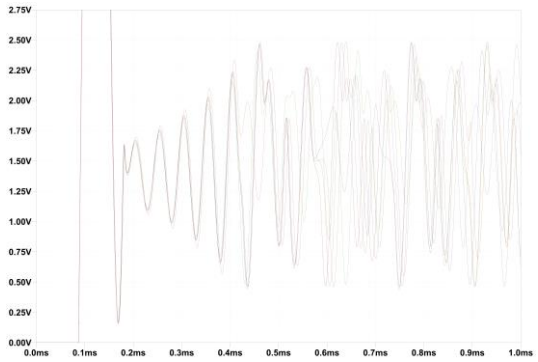

Fig.2. Simulated signal plots from channel X

However, once the voltage is sufficient to activate the auxiliary transistor, the third stage begins, leading to deterministic chaos. This stage is termed 'deterministic' because it consistently produces the same signal pattern each time the specific circuit is started. The shape of the signal is influenced by the precise values of the circuit components, and even minor alterations can significantly affect this part of the signal.

For simulation purposes, it was assumed that the components would be of high quality: capacitors with a tolerance of 1 % and resistors with a tolerance of 0.1 %, or better. Despite these narrow tolerance ranges, the signals generated by circuits built according to the same schematic will diverge shortly after startup.

After a relatively short period, fully developed chaos emerges. At this stage, the signal becomes unpredictable and highly sensitive to even the smallest changes in resistance and capacitance values. Factors such as the temperature coefficients of the components can have a significant impact at this point. The presence of a human body near the circuit can act as a large capacitor and may visibly alter the signal.

Plots similar to those obtained from channel X, as shown in Figure 2, can also be acquired from channel Y. Together, these signals enable the visualization of attractors, which are graphical representations of the system's state over time and reveal the unique dynamic behavior of each device. To analyze the system and determine the appropriate RC network values, this process was repeated numerous times throughout the research using Monte Carlo and worst-case methods considering quality of R and C components. For clarity, only a representative example of five X-Y plots simulated in LTSpice is presented in Figure 3. The shape closely resembles the original attractor [9].

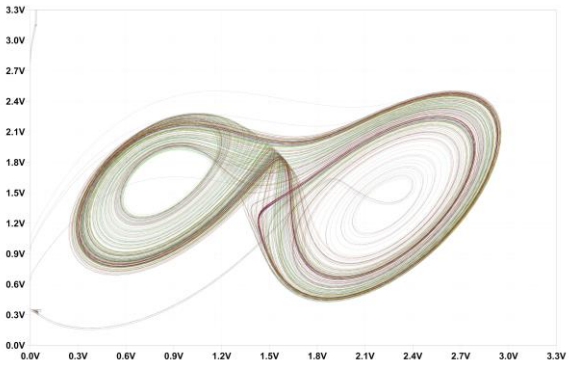

Fig.3. Simulated attractor in XY plane

Fig.4. Rendering of PCB design and the assembled board

Following the successful verification of the circuit through simulations, a PCB was designed and the device was constructed. A 3D model of the board alongside its physical implementation is shown in Figure 4. The board features a popular Arduino Uno -compatible pinout, which is also utilized by many STM32 development boards, including the STM32WB55-NUCLEO, which was later employed in this research.

Analysis of results

After assembling the device, its operation was tested using laboratory signal generators and an oscilloscope, triggered by the signal's initial transient state. By utilizing oscilloscope features such as high-resolution sampling and digital persistence, a plot of X and Y channels was obtained, as shown in Figure 5. The plot closely resembles the simulated one in key aspects, including the transient state, the increasing amplitude of oscillations, and the deterministic chaos phase. As true chaos emerges, the values observed in previous runs are displayed as a shaded envelope behind the current plot. This effect is made possible by the oscilloscope's digital persistence feature, which retains previous plots.

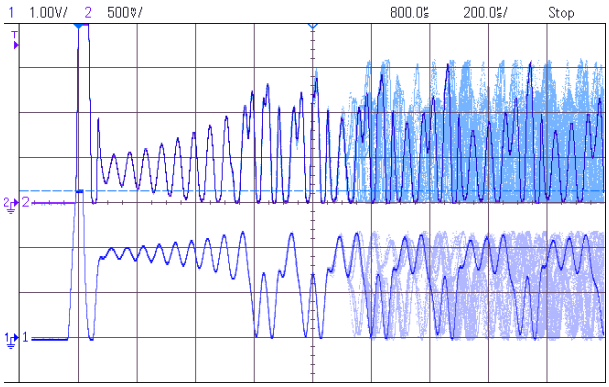

Fig.5. Oscillogram of captured X and Y signals

It can be assumed that the truly chaotic portion of the signal behaves like a random process. If this is the case, the signal should diminish in that section when averaged. This hypothesis was tested and confirmed using an oscilloscope. Subsequently, an application for the STM32WB55 was developed to repeatedly capture the initial few milliseconds of the signal after startup and average the consecutive runs. The parameters for this application were as follows:

- ADC clock: 24 MHz,
- Channels: 2,
- Sampling of single channel: 2.5 cycles,
- Conversion of single channel: 12.5 cycles,
- Effective sampling rate per channel: 800 kHz,
- Samples per series: 2048,
- Series sampling time: 2.56 ms.

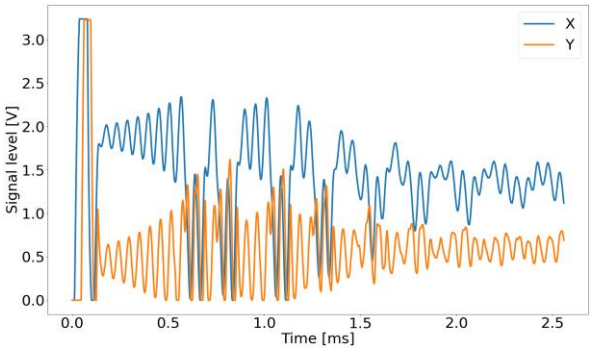

Fig.6. Rendering of PCB design and the assembled board

As the signal enters the truly chaotic phase, the averaged values begin to diminish, as shown in Figure 6, where 256 signal captures were averaged.

To further investigate the reliability of the deterministic component, entropy computation was employed. The original sequences, acquired using an STM32 microcontroller, had a resolution of 12 bits, which was also maintained for all averaged data. These 12-bit signals were then requantized to resolutions ranging from 1 to 11 bits. For each resolution, entropy computations were conducted across sequences that were consistent in terms of the number of averages. Samples corresponding to the same time points were used in the analysis. The time range remains consistent with previous analysis, limited to 2.56 ms. Representative results are shown in Figure 7.

The subplots in Figure 7 show the signal from the X channel of the chaotic generator after four averages, displayed at different resolutions following requantization. Due to space limitations in the paper, only four quantization levels were selected for publication. The signal is represented by solid lines, using the left ordinate axis. These plots serve as a visual reference for the entropy calculations.

Entropy values at specific time points are shown using a scatter plot, where each dot represents one of the 2048 entropy values across 10 acquired sequences. Some dots may appear darker if multiple points overlap at the same location. The entropy plot uses the right ordinate axis. As observed, the entropy value increases and reaches a plateau just above 3 bits after approximately 1.2 ms.

This analysis aligns with earlier observations, confirming that between the termination of oscillation and the emergence of true chaos, there is a brief phase of deterministic chaos. Although difficult to predict, this phase exhibits relatively lower entropy than the fully chaotic stage that follows. Additionally, we observe that entropy varies with the resolution of the analyzed signal. In this specific circuit and implementation, reducing the signal resolution to 4 bits could be considered, as it clearly distinguishes the two chaotic phases.

Above observations support the concept of deriving a reliable digital basis from the analog signal, which can then be used as a unique identifier for authenticating IoT devices.

The remaining concern is the computational feasibility of this approach. Given that modern hardware typically supports hashing functions, the microcontroller in this method is primarily burdened with the averaging process.

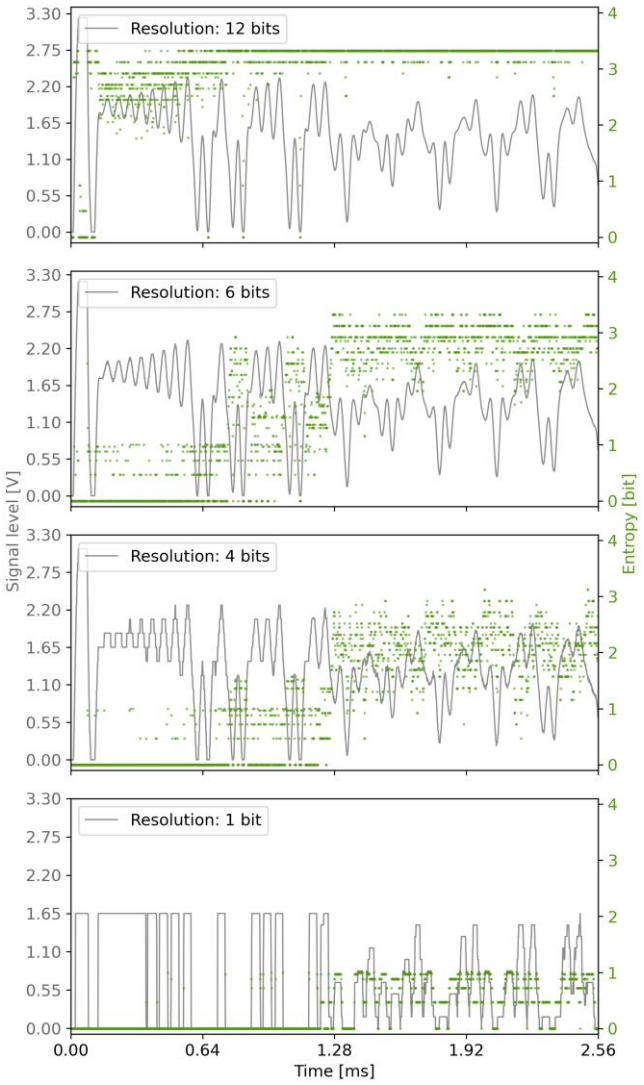

Fig.7. Entropy calculations for 10 sequences of 4-times averaged signals from the chaotic generator

Averaging times were tested with 2^n averages, where n ranged from 0 to 8. Both integer and floating-point averaging were evaluated. The duration of each averaging process was measured using the internal STM32 clock ticks, with the microcontroller operating at a 48 MHz clock frequency.

The linear relationship between the number of averages and the computation time, as shown in Figure 8, is expected. However, the absolute time required to acquire the signal from the chaotic generator and compute the corresponding PUF function holds significant practical value.

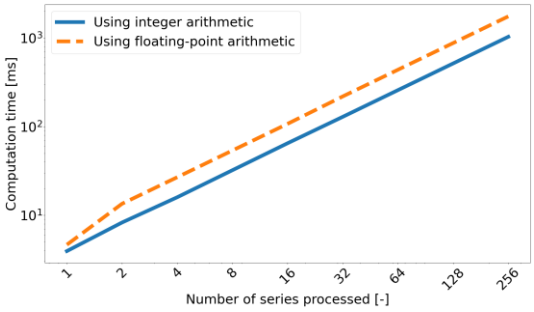

Fig.8. Times of series averaging in reference to number of averages

Conclusions

A chaotic generator based on discrete electronic components is examined both analytically and practically in this article. It highlights the necessity of effective signal conditioning in the analog domain using operational amplifiers. This conditioning is aimed at preparing the signal for sampling by a microcontroller and optimizing the use of its ADC range.

Analysis using laboratory equipment and microcontroller experiments has shown that averaging techniques can stabilize the deterministic chaos component of signals, making it a reliable source of unique information. This stabilization allows the PUF data to serve as a robust basis for authentication. The observed reliability of the deterministic chaos signal after averaging indicates significant potential for real-world applications, particularly for devices deployed in large quantities, such as those used in the IoT.

In conclusion, the findings indicate that PUF-based solutions represent a promising approach to IoT device authentication in the future. With increasingly interconnected environments, this approach could provide a reliable method for identifying and securing devices.

***Authors**: dr inż. Lukasz Makowski, Institute of Theory of Electrical Engineering, Measurement and Information Systems, Faculty of Electrical Engineering, Warsaw University of Technology, ul. Koszykowa 75, 00-662 Warszawa, Poland, e-mail: lukasz.makowski.ee@pw.edu.pl*

REFERENCES

[1] Barros T.G.F., Teixeira E.S., Common Requirements for IoT Platforms, *2023 IEEE 9th World Forum on Internet of Things (WF-IoT)*, Aveiro, Portugal, 2023, pp. 1-8, doi, 10.1109/WF-IoT58464.2023.10539561

[2] Neshenko N., Bou-Harb E., Crichigno J., Kaddoum G., Ghani N., Demystifying IoT Security, An Exhaustive Survey on IoT Vulnerabilities and a First Empirical Look on Internet-Scale IoT Exploitations, *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, pp. 2702-2733, 2019

[3] El-Hajj M., Beune P., Lightweight public key infrastructure for the Internet of Things, A systematic literature review, 2024, *Journal of Industrial Information Integration*, Volume 41, https,//doi.org/10.1016/j.jii.2024.100670.

[4] Bauer T., Hamlet J., Physical Unclonable Functions, A Primer, *IEEE Security & Privacy*, vol. 12, no. 06, pp. 97-101, 2014, doi, 10.1109/MSP.2014.123

[5] Garcia-Bosque M., Díez-Señorans G., Sánchez-Azqueta C., Celma S., Introduction to Physically Unclonable Fuctions, *Properties and Applications, European Conference on Circuit Theory and Design (ECCTD)*, Sofia, Bulgaria, 2020, pp. 1-4, doi, 10.1109/ECCTD49232.2020.9218404

[6] Mustafaev A.G., Buchaev A.Y., A Reliable Authentication Method for the Internet of Things Devices, *2020 International Conference on Information Technologies (InfoTech)*, Varna, Bulgaria, 2020, pp. 1-3, doi, 10.1109/InfoTech49733.2020.9211069.

[7] Al-Meer A., Al-Kuwari S., Physical Unclonable Functions (PUF) for IoT Devices, *ACM Comput. Surv*. 55, 14s, Article 314 (December 2023), https,//doi.org/10.1145/3591464

[8] Basiński R., Siwek K., Efektywny dwupołożeniowy algorytm stabilizacji orbit okresowych w układach chaotycznych, *Przegląd Elektrotechniczny*, R. 97 NR 12/2020, doi,10.15199/48.2020.12.28

[9] Keuninckx L., Van der Sande G., Danckaert J., Simple Two-Transistor Single-Supply Resistor–Capacitor Chaotic Oscillator, *IEEE Transactions on Circuits and Systems II*, Express Briefs, vol. 62, no. 9, pp. 891-895, Sept. 2015, doi, 10.1109/TCSII.2015.2435211