Refutation and Redesign of a Physical Model of TERO-based TRNGs and PUFs

Jeroen Delvaux[0000−0003−0684−8427]

jeroen.delvaux@osr-tech.com
Open Security Research (OSR), Room 29–31, Floor 8, Building 12B,
Shenzhen Bay Tech-Eco Park, 518000 Shenzhen, China

Abstract. In an article from CHES 2015, which appears in extended form in the Journal of Cryptology in 2019, Bernard, Haddad, Fischer, and Nicolai modeled the physical behavior of a *transient effect ring oscillator* (TERO), thereby providing a means to certify its operation as a *true random number generator* (TRNG). In this work, we disprove the physical assumption on which the whole model is based. Moreover, we show that the convenient use of tractable, closed-form equations stems from a mathematical error. On a more constructive note, we are the first to point out that TEROs and *Bistable Ring physically unclonable functions* (PUFs) are closely related, thereby not only laying the foundations of a more accurate physical model but also revealing a new design trade-off between throughput, entropy, and reliability. Furthermore, we demonstrate that most TERO implementations in the literature are prone to counter value corruptions, and propose a solution to this problem. Measurements performed on a *field-programmable gate array* (FPGA) substantiate our claims.

Keywords: Transient effect ring oscillator · True random number generator · Physically unclonable function · Stochastic model.

1 Introduction

Devices that take part in a cryptographic protocol usually require both irreproducible random numbers, such as nonces for providing freshness and masks for countering *side-channel attacks*, and reproducible random numbers, such as *symmetric keys*. To fulfill both needs using a single circuit, i.e., a joint *true random number generator* (TRNG) and *physically unclonable function* (PUF), *ring oscillators* (ROs) are an obvious choice. Conventionally, a TRNG output bit is produced by sampling the output node of a high-frequency RO at the rate of a low-frequency RO [12], whereas a PUF output bit is produced by comparing the frequencies of two identically laid-out ROs. As an alternative, Varchola and Drutarovský [28] proposed the use of a *transient effect ring oscillator* (TERO). Contrary to a regular RO, the number of inverting gates of a TERO is even rather than odd, it has two propagating events rather than only one, and its oscillation is temporary rather than permanent, i.e., the oscillation stops when

one event racing through the ring catches the other one. The duration of the oscillation is measured by counting the number of rising edges that pass through to the output node of the TERO. The *least significant bit* (LSB) of this counter is used as a TRNG output, whereas the *most significant bits* (MSBs) can be used to craft a complementary PUF, as pointed out by Varchola et al. [29] and Bossuet et al. [5].

Procedures for the *Common Criteria* evaluation of PUFs are still premature, but for TRNG designs targeting cryptographic applications, certification bodies such as the German BSI through AIS 31 [13] dictate that a stochastic model must be derived. Such a model is a form of *white-box* testing and provides a formal guarantee on the entropy provided, whereas *black-box* statistical tests, such as those specified by the United States' *National Institute of Standards and Technology* (NIST) [27], are not guaranteed to detect all exploitable non-uniformities. Bernard et al. [4] try to fulfill this need for a stochastic model for long TEROs, i.e., TEROs cascading tens of gates, in particular. Their model relies on the assumption that the oscillatory behavior of a TERO is dominated by an interplay between incomplete charging and discharging cycles. Hars [11] and Li et al. [14], however, provide two radically different explanations of the oscillatory behavior, although not as extensively elaborated as in the mathsy article of Bernard et al. [4]. As neither one out of three articles provides conclusive experimental evidence, which theory should an outsider believe?

1.1 Contributions

Through a series of experiments performed on a *field-programmable gate array* (FPGA), we confirm the claim of Li et al. [14] that process variations do not affect the two propagating events of a TERO equally and, therefore, dominate its oscillatory behavior. In this regard, we point out that TEROs and Bistable Ring PUFs [7] operate in an almost identical fashion and, consequentially, reveal a new design trade-off between throughput, entropy, and reliability. This revelation on the physical behavior of a TERO invalidates the mathematical model proposed by Bernard et al. [4] in its entirety, but the given formulas are equally disputable in their own right. Their derivation includes a trivial, easy-to-evaluate equation that is supposed to describe the *first-passage time* distribution of a Gaussian *random walk* with a complicated *drift function*, which opposes the academic consensus that a tractable, closed-form equation might not even exist in a non-trivial form. Not to give up on the quest of certifying TERO-based TRNGs, we develop a new model that is aimed to be mathematically simple yet physically accurate. As a final contribution, we demonstrate that previously published TERO implementations are prone to counter value corruptions. The corrupted values are hard to detect as they fall within the same range as the uncorrupted values, but can easily be avoided through a circuit modification.

1.2 Structure

The remainder of this work is structured as follows. Section 2 introduces the preliminaries. Section 3 scrutinizes the stochastic models of Hars [11] and Bernard et al. [4]. Section 4 proposes a new model that builds on the physical insight of Li et al. [14]. Section 5 describes experiments performed on a FPGA. Section 6 concludes this work.

2 Preliminaries

2.1 Notation

Constants are denoted by characters from the Greek alphabet, whereas variables are denoted by characters from the Latin alphabet. In the latter case, random variables and their realizations are distinguished by using uppercase and lowercase characters respectively. The *probability mass function* (PMF) and *cumulative distribution function* (CDF) of a discrete random variable X are denoted by $\mathsf{f}_X(x)$ and $\mathsf{F}_X(x)$ respectively. The same notation is used for the *probability density function* (PDF) and CDF of a continuous random variable X. A multivariate normal distribution is denoted by $X \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where $\boldsymbol{\mu}$ is the mean vector and where $\boldsymbol{\Sigma}$ is the symmetric covariance matrix.

2.2 Shannon Entropy

The Shannon entropy of a discrete random variable X is defined as $\mathrm{H}(X) \triangleq -\sum_x \mathsf{f}_X(x) \log\big(\mathsf{f}_X(x)\big)$, where summands $0 \log(0) \triangleq 0$. For binary vectors $X \in \{0,1\}^\lambda$, a convenient choice for the base of the logarithm is 2. It then holds that $\mathrm{H}(X) \in [0, \lambda]$, where $\mathrm{H}(X) = \lambda$ if and only if X is uniformly distributed on $\{0,1\}^\lambda$. For a vector $X = (X_1\, X_2 \cdots X_\lambda)$ that consists of *independent and identically distributed* (i.i.d.) bits, where $\forall j \in [1, \lambda], \Pr(X_j = 1) \triangleq \rho_{\text{bias}}$, it holds that $\mathrm{H}(X) = -\lambda\big(\rho_{\text{bias}} \log_2(\rho_{\text{bias}}) + (1 - \rho_{\text{bias}}) \log_2(1 - \rho_{\text{bias}})\big)$.

2.3 TERO-based TRNGs and PUFs

TEROs used in the literature have greatly different lengths. The most compact designs [28,10,29,5] consist of four gates, as illustrated in Fig. 1a, whereas the most bulky designs consist of tens of gates [8,4,19,14], as illustrated in Fig. 1b. The depicted circuits can be understood as a set-reset flip-flop where the feedback paths are slightly and considerably lengthened respectively, and where the set and reset signals are shorted to form a single control signal V_{en}. Each output bit is the result of a two-step process:

- **Reset phase** ($V_{\text{en}} = 0$). Following a universal guideline for designing TRNGs by Bucci and Luzzi [6], the depicted circuits are reset to a stable, predefined state prior to generating a new bit. The benefit of this design methodology is that the entropy source and the post-processing logic are both memoryless,

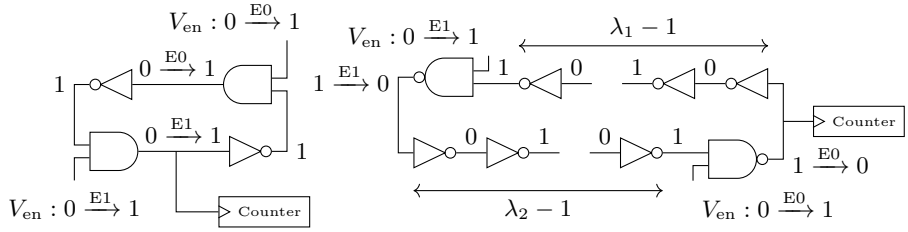

(a) Short. (b) Long.

Fig. 1: Example architectures of (a) a short TERO [28,10,29,5] and (b) a long TERO [8,14,4,19], where $\lambda_1, \lambda_2 \in \{3, 5, 7, \cdots\}$. In the latter case, Bernard et al. [4] recommend choosing $\lambda_1 \neq \lambda_2$. In either case, the signal V_{en} can be reused as an active-low reset for the counter.

and the generated bits can thus be assumed to be independent. Observe that upon assigning $V_{\text{en}} = 0$, the outputs of the AND and NAND gates unconditionally become 0 and 1 respectively, and once these changes are propagated by the inverters, either circuit reaches a stable state where the second inputs of the AND and NAND gates are 1.

– **Oscillation phase** ($V_{\text{en}} = 1$). Upon assigning $V_{\text{en}} = 1$, events E0 and E1 start to propagate through the ring. The ring is bistable now owing to an even number of inverting gates and is, therefore, left in either one of two states when the oscillation stops. Until that happens, a counter increments on each rising edge passing through the output node of the ring. Averagely speaking, the longer the TERO, the higher the counter values [8].

The obtained counter values serve two purposes:

– **TRNG**. Transistor-level noise sources in the TERO circuit cause minor, time-variant changes of the eventual counter value. These noise sources are represented by a parasitic current $I(t)$ and are called *white* if their *power spectral density* (PSD) $\mathsf{S}_I(f) \propto 1$, and *pink* if their PSD $\mathsf{S}_I(f) \propto 1/f$, where f denotes the frequency. For example, Johnson–Nyquist noise, which is caused by the thermal agitation of charge carriers, is white, whereas fluctuating occupancies of *traps* generate pink noise. As only the LSB of the counter is used as a TRNG output, an utterly compact but poorly testable implementation could constrict the counter to a single *toggle* (T) flip-flop.
– **PUF**. Process variations, such as *random dopant fluctuations* in transistors, render each manufactured TERO unique. Bossuet et al. [5] demonstrated that the final state of the TERO output node exhibits PUF behavior, but the reported *bit error rate* (BER) is too high to be practical. The authors suggest using the major, time-invariant changes of the counter value instead. Although it would be possible to extract a PUF response bit by comparing the counter value of a single TERO to a hardwired threshold,

a more technology-independent and environmentally robust approach omnipresent in PUF design is to compare the output of two identically laid-out structures. Stated otherwise, the counter values obtained from two TEROs are compared to generate a single response bit. Initially, electromagnetic side-channel attacks were claimed to provide no information on the number of oscillations [5], but Tebelmann et al. [26] and Mureddu et al. [21] later demonstrated the opposite.

2.4 Physicals Models of TEROs

A prerequisite for deriving a stochastic model of a TERO-based TRNG is knowledge on precisely what physical phenomenon dominates the oscillatory behavior. For short TEROs, Varchola and Drutarovský [28] and Bossuet et al. [5] agree that previous studies on *oscillator metastability* in flip-flops [23] provide this knowledge. As detailed in Fig. 2 for *complementary metal–oxide–semiconductor* (CMOS) technology, the process of charging and discharging circuit nodes is not instantaneous. Therefore, the narrowest of two pulses in a TERO, which is either the positive or the negative one, gradually disappears. Stated otherwise, the overtaking event accelerates given that the charging cycle caused by the event falling behind is not yet entirely completed and thus needs to be partially reversed only. This phenomenon is further referred to as *drafting* [30], in analogy to a cyclist's reduction in wind resistance when riding closely behind another cyclist. For long TEROs, which are the focus in the remainder of this article, the given explanations of the oscillatory behavior contradict one another.

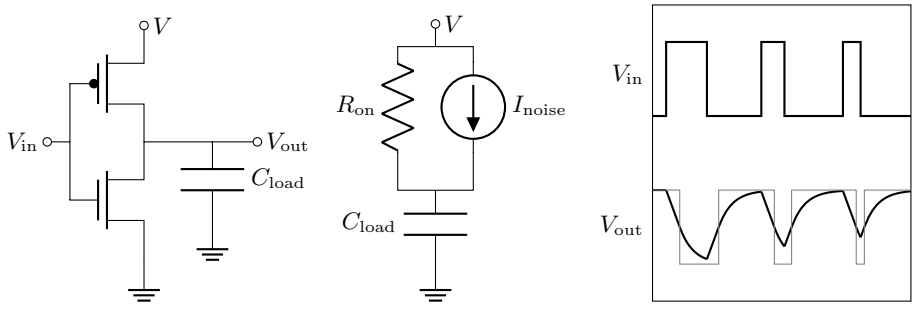

(a) Circuit. (b) Charging. (c) Transient response.

Fig. 2: A static CMOS inverter. (a) The circuit consists of an nMOS and a pMOS transistor. The load capacitance comprises an aggregate of all gates driven by the inverter. (b) In a first-order model of the charging process, only the pMOS transistor contributes [22, Chapter 5]. The discharging process is not drawn, but is similar in the sense that only the nMOS transistor contributes. Optionally, noise can be modelled by a current source [1]. (c) As the charging and discharging processes are non-instantaneous, narrow pulses tend to disappear [23].

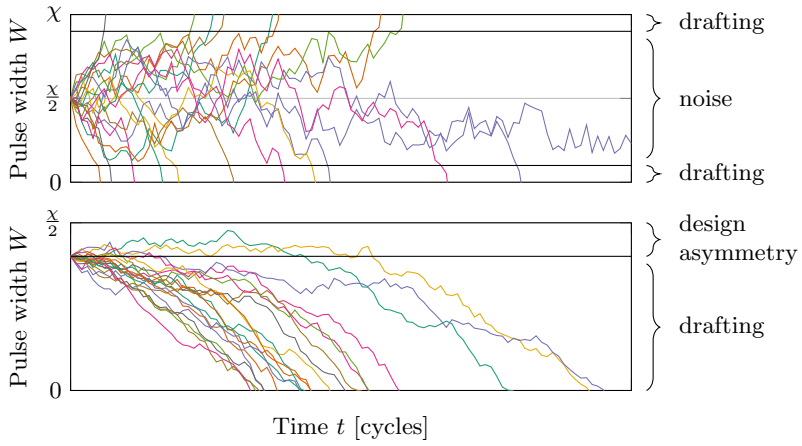

Fig. 3: The physics behind the stochastic models of (a) Hars [11] and (b) Bernard et al. [4]. Either model is developed for long TEROs.

According to Hars [11], transistor-level noise sources dominate the oscillatory behavior. The positive pulse width W as a function of the time T is modelled as a one-dimensional, discrete-time, driftless, Gaussian *random walk*, as illustrated in Fig. 3a. Let the *stochastic process* $\{W_t\}_{t\in\mathbb{N}}$ capture the pulse width $W_t \in [0, \chi]$ at the output node after traversing t cycles, where constant χ is the propagation delay of the whole loop. Starting from a certain W_0, the process is described as $W_t = W_0 + \sum_{j=1}^{t} U_j$, where increments U_j are i.i.d. zero-mean random variables $U_i \sim N\big(0, \sigma^2_{\text{noise}}\big)$. However, once noise has driven the two events sufficiently close to one another, the oscillation ceases quickly owing to the drafting effect. The author assumes that this effect emerges once $W_t \notin [\omega_{\min}, \chi - \omega_{\min}]$, where constant $\omega_{\min}$ is the minimum pulse width required for gates to switch properly.

According to Bernard et al. [4], the drafting effect dominates the oscillatory behavior entirely, i.e., long TEROs are operationally equivalent to short TEROs. Even more, the authors aim to kick-start this positive-feedback mechanism by instantiating the two inverter chains in Fig. 1b in a slightly asymmetric manner, i.e., $|\lambda_1 - \lambda_2|/(\lambda_1 + \lambda_2) \in [0.05, 0.35]$, roughly speaking. Similarly to the model of Hars [11], $W(T)$ is described as a one-dimensional, discrete-time, Gaussian random walk, but now there is drift. Increments U_i are modeled as random variables $U_j \sim N\big(\mu_j, \sigma_j^2\big)$, where mean μ_j is nonzero and dependent on the previous pulse width W_{t-1}. Analytical expression for μ_j and σ_j^2 are derived from a parameterized inverter model. Cherkaoui et al. [8, Section III] endorse the same physical foundations as Bernard et al. [4], but do not introduce a design asymmetry and let initial conditions be determined by process variations.

Marchand et al. [19] and Mureddu et al. [20,21] mention that process variations are the dominant phenomenon, but do not substantiate their claim by a logical explanation. Li et al. [14], however, point out that the periods of the two

propagating events, X_0 and X_1, are determined by the process variations from two disjoint sets of transistors, thereby implying $X_0 \neq X_1$. For any given circuit node in Fig. 1b, one event induces a rising edge in every cycle, whereas the other event always induces a falling edge. Note that for static CMOS inverters, as shown in Fig. 2, rise and fall times are determined by the pMOS and nMOS transistors respectively.

3 Scrutiny of Existing Physical Models

We now show that the stochastic models of Hars [11] and Bernard et al. [4] suffer from severe shortcomings, both from a physical and a mathematical perspective. Li et al. [14] provide a viable physical insight, which is not to be confused with a comprehensive, well-validated model.

3.1 The Model of Hars

Physics The assumption of Hars [11] that noise dominates the oscillatory behavior of a TERO can be readily dismissed. Experimental results from Bernard et al. [4, Fig. 1], Mureddu et al. [20], and also ourselves in Section 5 show that the pulse width W increases or decreases monotonically with the time T, thereby opposing the capricious dynamics shown in Fig. 3a.

Mathematics Hars [11] crudely simplifies theory and simulations by dropping Gaussianity and resorting to walks on the integer number line $\mathbb{Z}$, where moves are chosen randomly, uniformly, and independently from $\{-1, 1\}$. In addition to a loss of physical relevance, the problem with this simplification is that the given formulas for characterizing the distribution of the counter value Q, which is defined as the two-sided first-passage-time $Q \triangleq \min\{t \in \mathbb{N} \mid W_t \notin [\omega_{\min}, \chi - \omega_{\min}]\}$, still do not go beyond the expected value $\mathrm{E}[Q]$, thereby impeding a subsequent derivation of the entropy rate.

3.2 The Model of Bernard, Haddad, Fischer, and Nicolai

Physics We question the assumption of Bernard et al. [4, Fig. 1] that the drafting effect can bridge a gap of tens of inverters. When monitoring a node of an RO with an oscilloscope [11,4], the waveforms are usually observed to have relatively sharp rising and falling edges. As even 3-inverter ROs are known to exhibit rail-to-rail swings, we presume that the drafting effect must be contained in a span of at most two gate propagation delays. To be clear: we do not contradict the predictions of gate-level models that node voltages V converge to the supply voltages in an asymptotic sense, i.e., the voltage gap approaches zero when $T \to \infty$. Instead, we point out that the voltage gap drops quickly below noise levels and, therefore, becomes insignificant. Note also that the oscillatory model borrowed from Reyneri et al. [23] was developed for flip-flops consisting of a few

gates only, i.e., long chains were not anticipated for. Similarly, the feedback loops in the *self-timed rings* of Winstanley and Greenstreet [30], to which Cherkaoui et al. [8] refer, are short.

Mathematics For the model of Bernard et al. [4], a one-sided barrier suffices, and the first passage time becomes $Q \triangleq \min\{t \in \mathbb{N} \mid W_t \geq 0\}$. Bias, and thus the entropy rate, is computed by accumulating the probability masses for all odd values of Q, i.e., $\rho_{\text{bias}} = \sum_{j \in \mathbb{N}} \mathsf{f}_Q(2j+1)$. The PMF of Q is trivially derived from the CDF of Q as follows: $\mathsf{f}_Q(q) = \mathsf{F}_Q(q) - \mathsf{F}_Q(q-1)$. Unfortunately, the CDF of Q is wrongly computed as $\mathsf{F}_Q(q) = \Pr(W(q) \leq 0) = \mathsf{F}_N(-\mu(q)/\sigma(q))$. The complex situation of dealing with walks that have multiple barrier crossings, as illustrated in Fig. 4, is mishandled this way. As CDF $\mathsf{F}_Q(q)$, the authors inadvertently computed the probability that a walk has an odd number of crossings up until the given time q. Consequentially, PMF $\mathsf{f}_Q(q)$ is thus inadvertently the probability that a walk has an outbound crossing in the time interval $[q-1, q]$ minus the probability that a walk has an inbound crossing in the time interval $[q-1, q]$. As can be seen in Fig. 4(b), the correct and incorrect distribution are quite similar and both right-skewed. As moves are neither independent nor identically distributed, we conjecture that a tractable, closed-form equation for the correct result does not exist. In fact, first-passage-time distributions are easier to handle in the continuous-time domain, where random walks are referred to as Wiener processes or Brownian motions, often aided by the *inverse Gaussian* distribution, but also here there are limitations on the analytical complexity [16].

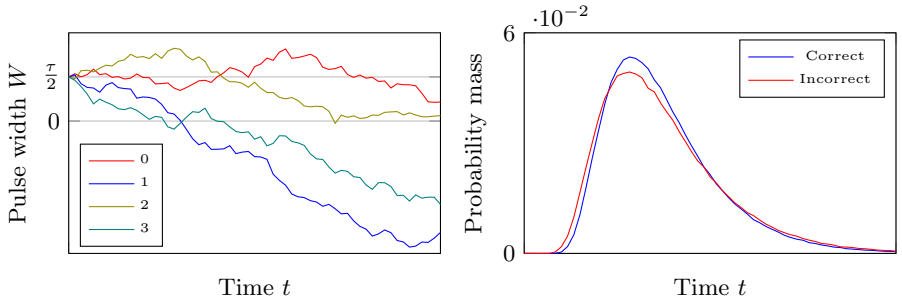

Fig. 4: (a) Random walks having 0, 1, 2, and 3 barrier crossings. (b) Monte Carlo experiment of 10^6 walks estimating the PMF of C.

3.3 The Model of Li, Wu, Zhang, Wu, Zhou, and Wang

Physics Li et al. [14] provide a solid argument for how process variations dominate the oscillatory behavior of a TERO, but strong experimental evidence is missing. Also, physical phenomena that influence but therefore not dominate the oscillatory behavior are neither analyzed nor put into perspective.

Mathematics Li et al. [14] do not propose a mathematical model.

4 Proposal of a New Physical Model

Using the physical insight of Li et al. [14] as a starting point, we now develop a comprehensive model for the oscillatory behavior of a long TERO.

4.1 Physics

Guided by the overview in Fig. 5, we now discuss several physical phenomena that affect the oscillatory behavior of a TERO:

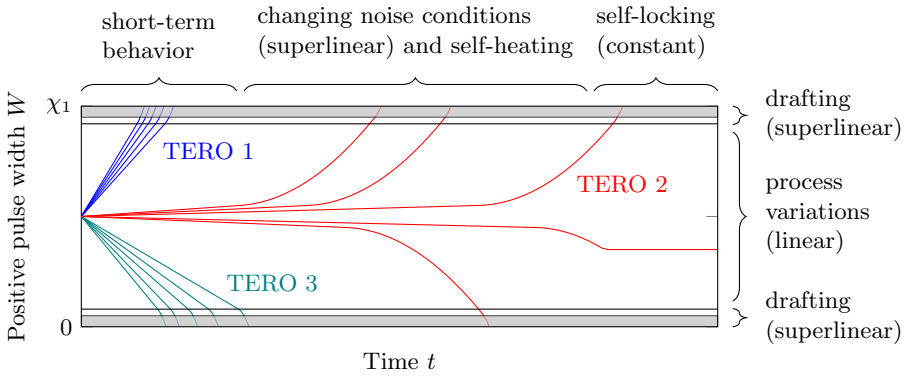

Fig. 5: The physics behind our newly proposed stochastic model. Narrow, low-amplitude pulses do not increment the counter, as reflected by the gray regions. The drawn curves are not intended to be as accurate as real-world measurements, i.e., proportions are purely indicative.

Process Variations Similarly to models of Arbiter [15] and other delay-based PUFs [24], we abstract gate propagation delays as real numbers. This way, 2λ parameters arise for a TERO consisting of $\lambda \in \{6, 8, 10, \cdots\}$ inverting gates: the low-to-high propagation delay $D_{j,\mathrm{LH}}$ and the high-to-low propagation delay $D_{j,\mathrm{HL}}$ of any given circuit node $j \in [1, \lambda]$. On FPGAs, where inverters are implemented as a *lookup table* [31], D_{LH} and D_{HL} differ not only because the pull-up and pull-down networks of a gate are realized by different transistors but also because two different paths in a *multiplexer* (MUX) are followed. If by convention, we label λ as the output node of the TERO, and E0 and E1 as the events that respectively induce a falling and a rising edge on this node, the respective periods X_0 and X_1 are given in Eq. (1). The eventual state of the

output node is 0 if $X_0 < X_1$ and 1 otherwise. Bistable Ring PUFs [7], which can be understood as TEROs having not 2 but λ propagating events, were modeled in a mathematically identical fashion by Xu et al. [32]. They, however, use the abstract quantity "gate strength" instead of the concrete quantity "gate propagation delay".

$$X_0 = \sum_{j=1}^{\lambda/2} D_{2j,\mathrm{HL}} + D_{2j-1,\mathrm{LH}}, \quad X_1 = \sum_{j=1}^{\lambda/2} D_{2j,\mathrm{LH}} + D_{2j-1,\mathrm{HL}}. \tag{1}$$

If the values of X_0 and X_1 remain stable until the end of the oscillation, then the positive pulse width W is a linear function of the time T, as formalized in Eq. (2). The given expression for W_0 assumes that events E0 and E1 are initiated at nodes ι_0 and ι_1 respectively, where $\iota_0, \iota_1 \in \{2, 4, \cdots, \lambda\}$ and $\iota_0 < \iota_1$.

$$W(T) = W_0 + \frac{X_0 - X_1}{X_1} T, \quad \text{where } W_0 = \sum_{j=\iota_0/2}^{\iota_1/2} D_{2j,\mathrm{HL}} + D_{2j-1,\mathrm{LH}}. \tag{2}$$

Internal Noise In reality, periods X_0 and X_1 are weakly time-dependent as the transistors that constitute a TERO contain noise sources. Existing stochastic models of both RO-based [17,3] and TERO-based [11,4] TRNGs are plagued by the misassumption that stochastic processes describing consecutive periods, i.e., $\{X_j\}$, and/or consecutive pulse widths, i.e., $\{W_j\}$, exhibit high-frequency variations owing to white noise sources in these transistors. Formally, it is a common practice to assume i.i.d. random variables $W_j \sim N(\mu, \sigma^2)$ and/or i.i.d. random variables $X_j \sim N(\mu, \sigma^2)$, which implies that the process $\{X_1 + X_2 + \cdots + X_i\}$ is assumed to describe a Gaussian random walk. However, for ROs and TEROs implemented in static CMOS technology, as shown in Fig. 2, the parasitic noise current $I(t)$ generated by each transistor is integrated by the load capacitance of the next inverter, thereby realizing a low-pass filter. Abidi [1] derived that the white and pink components of noise current $I(t)$ transfer to the oscillatory voltage signal $V(t)$ as PSDs $\mathsf{S}_V(f) \propto 1/(f - f_0)^2$ and $\mathsf{S}_V(f) \propto 1/|f - f_0|^3$ respectively, where f_0 is the frequency in absence of noise. As high frequencies are attenuated, the curves drawn in Fig. 5 are smooth and do not mimic the shaky trajectories described by the random walks in Fig. 3.

Nevertheless, the integrated noise contributions imply that TERO periods X_0 and X_1 change from the start of one oscillation to another, as reflected by the bundles of lines having slightly differing slopes in Fig. 5. For slowly-converging TEROs, i.e., TEROs for which process variations only generate a small difference between X_0 and X_1, there is enough time for noise conditions to change significantly during the process of generating a single bit. Evidently, a condition where $X_0 \approx X_1$ can only be escaped from by an increase in the absolute value $|X_0 - X_1|$ and, therefore, curves drawn for the slowly-converging TERO in Fig. 5 are superlinear.

External Perturbations Not only noise sources contained in the TERO itself, but also perturbations external to this circuit might change the values of periods X_0 and X_1. For example, surrounding on-chip components such as a microcontroller might generate heat and/or supply voltage spikes and, consequentially, influence the TERO operation. Also changes external to the chip, which includes the outside temperature and the power source, are a contributor. A crucial observation is that the externally induced changes to X_0 and X_1 are correlated. For example, an increase in the temperature or a decrease in the supply voltage will increase the value of both X_0 and X_1 [22]. Similar to internal noise sources, external conditions change from the start of one oscillation to another, or for slowly-converging TERO in particular, even during the process of generating a single bit.

Partially Detectable Drafting Effect Towards the end of the oscillation, when the drafting effect kicks in, the output node of the TERO does not exhibit rail-to-rail voltage swings anymore. Initially, the counter might still detect such deteriorated pulses, but once their amplitudes drop below the *threshold voltage* of certain transistors within the counter circuit, pulses pass by unnoticed.

Self-heating ROs generate heat, as demonstrated by Agne et al. [2]. In a process known as *Joule heating*, charge carriers transfer part of their kinetic energy to the (semi)conductor lattice upon colliding and scattering. TEROs, which have twice as many switching events per unit of time as ROs, are equally if not more adept at raising their own local temperature. Consequentially, periods X_0 and X_1 are expected to increase over time [22]. This process is self-regulated given that the effect counteracts its cause, i.e., the increased periods reduce the generation of heat. Furthermore, we expect this process to be partially cyclic: the reset phase lacks switching events, thereby allowing for cooling.

Self-locking The two propagating events of a TERO might self-induce injection locking, as demonstrated by Mureddu et al. [20, Fig. 13]. If process variations and initial noise conditions create an exceptionally small difference between periods X_0 and X_1 at time $T = 0$, circumstances are ideal for the locking phenomenon to occur and further reduce this difference until it is practically 0. Eventually, disturbances will cause the two self-locked events to unlock. The prevalence of self-locking is layout-dependent as opposite nodes of the TERO are required to interact, e.g., through a capacitive coupling in the silicon *substrate*.

4.2 Mathematics

We propose a two-parameter stochastic model for quickly-converging TEROs. Due to the swift convergence, periods X_0 and X_1 are assumed to be constant during any given evaluation, i.e., $W(T)$ in Eq. (2) is a linear function. By solving $W(T) = 0$ for $X_0 < X_1$, we obtain that the oscillation stops when

$T = W_0\,X_1/(X_1 - X_0)$. If we further assume a trivariate normal distribution $(X_0, X_1, W_0) \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, where mean vector $\boldsymbol{\mu}$ is determined by process variations, and where the covariance matrix $\boldsymbol{\Sigma}$ aggregates internal and external noise contributions, we are saddled with nine parameters, i.e., three means μ, three variances σ^2, and three covariances ρ. However, because the normally distributed denominator $(X_1 - X_0) \sim N(\mu_{X_1} - \mu_{X_0}, \sigma^2_{X_1} + \sigma^2_{X_0} - 2\rho_{X_0X_1})$ has a near-zero mean, the numerator can be approximated by a constant. Therefore, the distribution of the eventual counter value Q can be modeled by two parameters, as shown in Eq. (3). As for Bernard et al. [4], the bias and, subsequently, the entropy rate are computed by accumulating the probability masses for all odd values of Q, i.e., $\rho_{\text{bias}} = \sum_{j\in\mathbb{N}} \mathsf{f}_Q(2j+1)$.

$$Q = \left\lfloor \frac{1}{Y} \right\rfloor, \text{where } Y \sim N(\mu, \sigma^2) \implies \mathsf{f}_Q(q) = \mathsf{F}_{N(\mu,\sigma^2)}\left(\frac{1}{q}\right) - \mathsf{F}_{N(\mu,\sigma^2)}\left(\frac{1}{q+1}\right). \tag{3}$$

The modeling of slowly-converging TEROs is a daunting task: changing noise conditions and effects such as self-heating and self-locking, of which the precise characteristics are not even 100% understood in isolation, then need to be combined with accurate proportions. However, as TRNGs deployed in cryptographic protocols usually have stringent latency constraints, the long-term behavior might not even be relevant in practice. Unfortunately, the inherent latency of a TERO-based TRNG cannot accurately be controlled prior to device manufacturing, unlike traditional TRNG designs consisting of a low-f RO and a high-f RO [12]. A workaround, such as reconfigurable gate propagation delays, or a selection procedure among multiple TEROs laid-out in parallel, is hence instrumental in obtaining a quick convergence. Evidently, the lack of a priori control is a problem for the entropy rate too, given that the entropy increases with the convergence time.

5 Experiments

5.1 TERO Architecture and Implementation

Our TERO architecture, which is optimized for testability, is shown together with its immediate surroundings in Fig. 6. The oscillating core comprises a circular chain of 2-to-1 MUXs, as previously used by Malik [18]. MUX data inputs $(c_{2j-1}\,c_{2j})$ are reconfigurable such that each stage $j \in [1, \lambda]$ can realize either a buffer (input 01) or an inverter (input 10) or a constant 0 (input 00) or a constant 1 (input 11). This way, thousands or even millions of TERO instances and, therefore, a statistically significant amount of data, can be generated without having to defy area constraints by laying out as many loops in parallel. A notable feature of our design is that the number of propagating events, ψ, is arbitrary. In addition to TEROs, we can realize ROs and a generalized version of Bistable Ring PUFs where $\psi \in \{4, 6, 8, \cdots, \lambda\}$, as illustrated in Table 1.

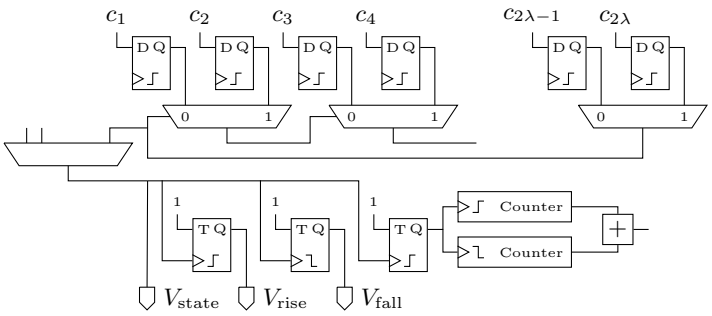

Fig. 6: Our TERO architecture. For each length $\lambda \in \{8, 16, 24, 32, 40, 48, 56, 64\}$, we lay out 8 TEROs in parallel, i.e., a total of 64 loops. To avoid injection locking, at most one TERO is oscillating at any given time. Pentagons represent test pins that are used for performing oscilloscope measurements.

Table 1: Configuration pair examples for a 16-stage module.

Mode		Configuration $\boldsymbol{c} = (c_1\, c_2\, c_3\, \cdots\, c_{32})$	BUF ↑	BUF ↓	INV ↑	INV ↓
RO ($\psi = 1$)	$\boldsymbol{c}_{\text{rst}}$	01|10|01|10|10|**00**|10|10|01|01|01|10|10|10|01|10				
	$\boldsymbol{c}_{\text{osc}}$	|01|	7	7	9	9
Unbalanced TERO ($\psi = 2$)	$\boldsymbol{c}_{\text{rst}}$	10|10|**11**|01|01|10|01|01|10|10|**11**|10|10|01|10|10				
	$\boldsymbol{c}_{\text{osc}}$	| | |01| |	2	4	5	5
		|10| | |	4	2	5	5
Balanced TERO ($\psi = 2$)	$\boldsymbol{c}_{\text{rst}}$	01|10|**00**|01|10|01|01|10|01|01|**00**|10|10|01|01|10				
	$\boldsymbol{c}_{\text{osc}}$	| | |10| |	4	4	4	4
		|10| | |	4	4	4	4
Unbalanced Bistable Ring ($\psi = 4$)	$\boldsymbol{c}_{\text{rst}}$	01|01|**00**|10|10|10|**11**|10|01|01|**00**|01|01|10|**11**|01				
	$\boldsymbol{c}_{\text{osc}}$	| | |10| | | |10|	2	6	4	4
		|01| | | |10| | |	6	2	4	4
Balanced Bistable Ring ($\psi = 8$)	$\boldsymbol{c}_{\text{rst}}$	**00**|10|**00**|01|**00**|01|**00**|10|**00**|01|**00**|01|**00**|10|**00**|10				
	$\boldsymbol{c}_{\text{osc}}$	| |01| | | |10| | | |10| | | |01|	4	4	4	4
		01| | | |10| | | |01| | | |10| | |	4	4	4	4

Similarly to Arbiter [15] and other delay-based PUFs [24], the input–output behavior of a reconfigurable TERO in state-based PUF mode can be described by a quasi-linear model. The difference between both periods, $Z = X_0 - X_1$, is rewritten as a dot product $Z = M^T S$ further defined in Eq. (4), where M depends on gate propagation delays D and where S depends on the configuration C_{osc}. If $Z > 0$, the response $R = 1$, otherwise, $R = 0$. This model also points out the danger of two types of biases. First, there is architectural bias [25,9]. Even in the ideal case $M \sim N(\mathbf{0}_{2\lambda}, \sigma^2 \boldsymbol{I}_{2\lambda})$, the probability $\Pr(R = 1)$ might differ from the ideal value 50% for any given TERO instance. This is evidenced by the observation $\mathrm{E}_{C_{\text{osc}}}[Z|M] = (M_{2\lambda-1} + M_{2\lambda})/2 \neq 0$. In fact, by having both inverters and buffers, we largely mitigated the problem. An alternative reconfigurable TERO design where each out of λ stages selects one out of two identically laid-out inverters [7, Fig. 2] has $\mathrm{E}_{C_{\text{osc}}}[Z|M] = \sum_{j=1}^{\lambda}(-1)^j(M_{2j-1} + M_{2j})/2$ and is, therefore, more prone to architectural bias. Second, bias can be caused by asymmetries in the implementation, which includes the circuit, the layout, and the manufacturing process. The distributions of $D_{\text{BUF,LH}}$, $D_{\text{BUF,HL}}$, $D_{\text{INV,LH}}$, $D_{\text{INV,HL}}$ are not necessarily identical. A partial solution to this problem is to choose configurations in a balanced way [24, Fig. 1] such that the four types of delays occur in equal quantities for both propagating events.

$$M = \begin{pmatrix} D_{1,\text{BUF,LH}} - D_{1,\text{BUF,HL}} \\ D_{1,\text{INV,LH}} - D_{1,\text{INV,HL}} \\ D_{2,\text{BUF,LH}} - D_{2,\text{BUF,HL}} \\ D_{2,\text{INV,LH}} - D_{2,\text{INV,HL}} \\ \vdots \\ D_{\lambda,\text{BUF,LH}} - D_{\lambda,\text{BUF,HL}} \\ D_{\lambda,\text{INV,LH}} - D_{\lambda,\text{INV,HL}} \end{pmatrix}, S = \begin{pmatrix} (1 - C_{1,\text{osc}}) \prod_{j=2}^{\lambda}(1 - 2\,C_{2j-1,\text{osc}}) \\ C_{1,\text{osc}} \prod_{j=2}^{\lambda}(1 - 2\,C_{2j-1,\text{osc}}) \\ (1 - C_{3,\text{osc}}) \prod_{j=3}^{\lambda}(1 - 2\,C_{2j-1,\text{osc}}) \\ C_{3,\text{osc}} \prod_{j=3}^{\lambda}(1 - 2\,C_{2j-1,\text{osc}}) \\ \vdots \\ 1 - C_{2\lambda-1,\text{osc}} \\ C_{2\lambda-1,\text{osc}} \end{pmatrix}. \tag{4}$$

Using Vivado, a Verilog implementation of our design in Fig. 6 is mapped to the Xilinx 7-series programmable logic [31] of a Zynq-7000 *system on chip* (SoC), which is part of a Digilent *Zynq Evaluation and Development* (Zed) board. As asymmetries in the implementation are of limited importance in this study, we avoid the ordeal of creating *hard macros* and use automated placement and routing. The SoC also features an ARM Cortex-A9 processor, which is used to control the TERO modules and transfer their outputs via a *universal asynchronous receiver-transmitter* (UART) to a *personal computer* (PC).

Avoiding Counter Value Corruptions To the best of our knowledge, all previously published TERO implementations featuring multi-bit counters are flawed. As shown in Fig. 1, it is a common practice [5,8,4,19,21] to wire the output node of a TERO to the clock input of a counter. Towards the end of the oscillation, however, the duty cycle of this clock becomes extreme, whereas the flip-flops embedded in the counter are only guaranteed to function properly for clocks having 50% duty cycle. Even worse, these flip-flops have slightly

different critical timing characteristics owing to process variations and routing asymmetries. Conceivably, for the last detectable clock edge, corruptions can occur where some flip-flops successfully update their states in accordance with the latest counter value, while the remaining flip-flops stick to the second-to-last counter value. This intuition is confirmed through FPGA experiments; a few of the observed corruptions are listed in Table 2.

Table 2: Examples of experimentally observed counter value corruptions. To detect such corruptions reliably, we had an improperly clocked counter, as is shown in Fig. 1, and a pair of properly clocked counters, as is shown in Fig. 6, operate in parallel.

Counter value						**Duty cycle**
Second-to-last		**Last**		**Corrupted**		
31	(011111)	32	(100000)	16	(010000)	$\to 0\%$
31	(011111)	32	(100000)	63	(111111)	$\to 100\%$
63	(111111)	64	(1000000)	127	(1111111)	$\to 100\%$
79	(1001111)	80	(1010000)	64	(1000000)	$\to 0\%$
95	(1011111)	96	(1100000)	127	(1111111)	$\to 100\%$
127	(01111111)	128	(10000000)	255	(11111111)	$\to 100\%$
255	(011111111)	256	(100000000)	15	(000001111)	$\to 100\%$
415	(110011111)	416	(110100000)	447	(110111111)	$\to 100\%$
511	(0111111111)	512	(1000000000)	783	(1100001111)	$\to 100\%$
1023	(01111111111)	1024	(10000000000)	1007	(01111101111)	$\to 100\%$

On the bright side, outliers are only produced when the two last counter values have a significant Hamming distance, i.e., when a carry ripples through several consecutive adder stages. More devastating corruptions were observed for a less conventional counter implementation in the form of a maximum-length *linear-feedback shift register* (LFSR). In the latter case, state updates are characterized by larger Hamming distances, and involve flipping even the MSBs. For TRNG mode, the observed counter value corruptions are not overly dramatic, given that the primary output consists of the LSB only. Nevertheless, the problem cannot be ignored: full counter values facilitate stochastic modeling and/or online health tests, and are instrumental for PUF mode as well. To solve the problem, we restore the duty cycle of the clock to approximately 50% using a T flip-flop. This solution has the adverse side effect of dividing the clock frequency by two. Not to sacrifice temporal resolution, the positive-edge-triggered flip-flops embedded in the counter could be replaced by dual-edge-triggered flip-flops. As the latter type of flip-flop is not readily available on our FPGA platform, we sum the outputs of a positive-edge-triggered counter and a negative-edge-triggered counter instead.

Oscilloscope Measurements The analog signal of the TERO output node, i.e., $V_{\text{state}}(t)$ in Fig. 6, is crucial for gaining new insights but is ill-suited for direct

oscilloscope measurements. Towards the end of the oscillation, when duty cycles become extreme, only the highest frequencies contain useful information, whereas the capacitors contained in test pins, connectors, and probes realize a low-pass filter. When connecting probes to the Digilent Pmod connectors, i.e., the only readily available interface on a Zed board, the last few pulses are not even visible. To improve the signal quality, we resorted to the *FPGA Mezzanine Card* (FMC) connector, which necessitated soldering wires and jumpers to a *breakout board* supplied by *Instrumentation And Measurement* (IAM) Electronic, as shown in Fig. 7a. As this only partially solved the problem, we added two extra test pins where the two propagating events can be monitored separately, i.e., $V_{\text{rise}}(t)$ and $V_{\text{fall}}(t)$ in Fig. 6. The fourth channel, which is not shown in Fig. 6, is a trigger pulse generated at the start of each oscillation, thereby allowing for the synchronization of measurements. The signals in Fig. 7b are measured using a Tektronix 5104B *mixed signal oscilloscope* (MSO). Its *sampling frequency* is set to 5 GHz, which is the maximum value. Each TPP1000 *Tektronix passive probe* (TPP) has a 3 dB bandwidth of 1 GHz. Counters are disabled during oscilloscope measurements to avoid potential interference.

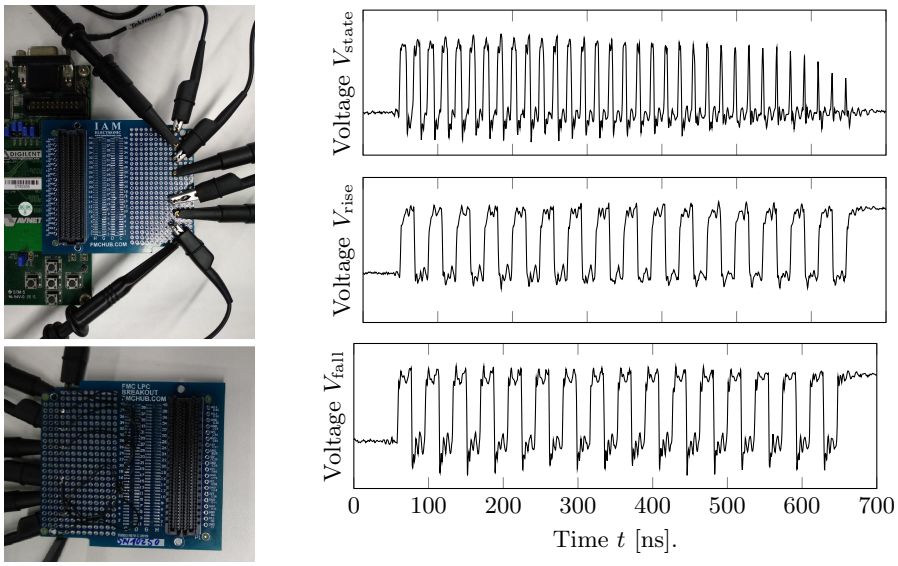

(a) FMC breakout board.

(b) Oscilloscope signals.

Fig. 7: Our experimental setup for performing oscilloscope measurements. (a) A breakout board facilitates connecting probes to four *differential pairs* of the FMC connector. (b) The measured signals for a 64-stage TERO. Its configuration pair $(\boldsymbol{c}_{\text{rst}}, \boldsymbol{c}_{\text{osc}})$ is not randomly chosen but is selected to generate few pulses only, for easy of display.

To compute periods X and pulse widths W as a function of the time T, we adopt a two-step approach. First, an edge detection filter is applied to determine rough estimates of (i) the sample where the oscillation starts, (ii) the sample where the oscillation stops, and (iii) the period X. For RO mode, we apply this filter to $V_{\text{state}}(t)$, whereas for TERO mode, we use $V_{\text{rise}}(t)$ and $V_{\text{fall}}(t)$ instead. Second, we select a small subsequence from either $V_{\text{state}}(t)$ or $V_{\text{rise}}(t)$ or $V_{\text{fall}}(t)$ and slide it along the original signal at a given time offset, while computing the Pearson correlation coefficient and finding its peak location using quadratic interpolation such that the time resolution is not limited by the sampling period of 200 ps. To compute $X(T)$, we continuously update the window contents, whereas for $W(T)$, we use the same window for the whole signal.

5.2 Refutation of the Physical model of Bernard et al.

Bernard et al. [4] assume that the drafting effect causes event E0/E1 to lose the race if its initial lead over event E1/E0 is smaller than, approximately, half the number of stages. We test the validity of this assumption for eight 64-stage TEROs by measuring the probability that event E0/E1 wins the race as a function of its initial lead. Fig. 8 shows that even under extreme initial conditions where event E0/E1's lead is only 1, 2, or 3 stages, it is still capable of winning the race. Therefore, the drafting effect cannot be the dominant factor in determining the winner of the race. Beware of using Fig. 8 for estimating the exact significance of the drafting effect near the end of the race. Another contributor to the measured probability falling below 50% is the occasional occurrence of configuration pairs where process variations generate a small difference between periods X_0 and X_1. If the initial lead of event E0/E1 is small, a relatively quick start of event E1/E0 then instantly ends the race, whereas in a fair game, event E0/E1 would have more time to overturn its relatively slow start and eventually win the race. Stated otherwise, there is no time for initially unfavorable noise conditions to change in favor of E0/E1.

5.3 Validation of the New Physical Model

Pulse Width Measurements The curves $W(T)$ manually drawn in Fig. 5 can be measured using an oscilloscope. Mureddu et al. [20, Fig. 11, Fig. 13] previously presented similar measurements for two configurations, which is insufficient to properly validate our model. Moreover, their results are open to interpretation as a precise specification of the digital signal processing methods has not been given. Figure 9a shows that for quickly converging TEROs, curves are approximately straight, which is consistent with two events propagating at different speeds owing to process variations from two disjoint sets of transistors. Moreover, the smoothness of the curves indicates that there is an integration of noise.

The Revival of State-Based PUFs Figure 10a confirms the existing knowledge that the expected counter value of a TERO, $\mathrm{E}[Q]$, increases with the number of stages, λ. Figure 10b confirms the existing knowledge that the Shannon

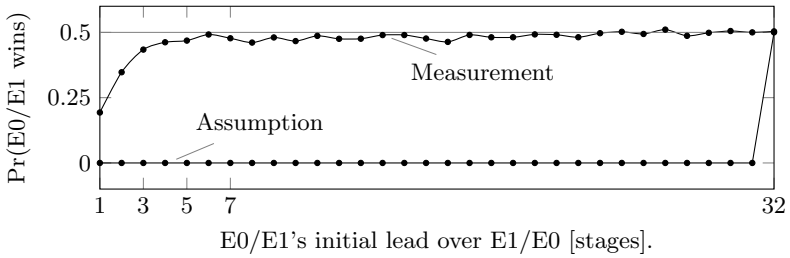

Fig. 8: Refutation of the physics behind the stochastic model of Bernard et al. [4]. For several 64-stage TEROs, we vary event E0/E1's initial lead from 1 to 32 stages, and measure its probability of winning the race. Each measured dot averages the outcome of 32000 evaluations, or more precisely, 10 evaluations of 400 different configurations pairs for 8 different TEROs.

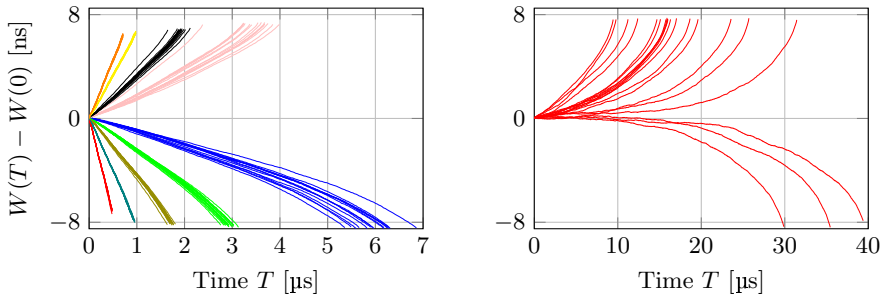

(a) Quickly converging TEROs.

(b) Slowly converging TERO.

Fig. 9: The pulse width W as a function of the time T, as measured using an oscilloscope. We plot increments $W(T) - W(0)$, where time $T = 0$ at the start of any given oscillation, for different 64-stage TERO instances and different configurations.

entropy of the counter-based TRNG output, H(N), increases with E[Q]. Figure 10c, however, shows a novel observation that validates our model: the BER of the state-based PUF response R is zero when the initial periods X_1 and X_2 significantly differ, and becomes nonzero when $X_1 \approx X_2$. Moreover, the BER averaged over all configuration pairs is less than 1%, which opposes the observation of Bossuet et al. [5, Section III-A] that responses bits R cannot reliably be extracted as the state keeps oscillating in 29% of the cases. Their numbers, however, cannot be trusted due to the self-contradictory statement that counter values Q, which should increment as long as the state keeps oscillating, are observed to become stable in all their experiments. To conclude, the literature's exclusive focus on counter-based PUFs might not be justified. State-based PUFs are, at the very least, competitive, also in part because they only require half the number of loops for a given number of uncorrelated response bits[1]. The omission of counters and comparators provides even more benefits: (i) less area is consumed, (ii) counter value corruptions become irrelevant, and (iii) their electromagnetic emissions, which might facilitate side-channel attacks [26,21] and induce self-locking [20], become irrelevant.

Machine Learning If the delay-based model in Section 4 and Eq. (4) is correct, then we should be able to determine its parameters using machine-learning techniques such that the unseen response r to any given configuration pair $(\boldsymbol{c}_{\text{rst}}, \boldsymbol{c}_{\text{osc}})$ is correctly predicted with high probability. We adopt an *artificial neural network* (ANN) that consists of a single neuron only, given that it realizes a dot product of which the result is compared to a threshold. *Resilient backpropagation* is used as a training algorithm. To overcome the inconvenience that ANNs serve a regression purpose rather than a classification purpose, we train two single-neuron networks: a model $\hat{\boldsymbol{m}}_1$ obtained using original responses r and a model $\hat{\boldsymbol{m}}_0$ obtained using inverted responses $1 - r$. The predicted response $\hat{r}$ is 1 if $s \cdot \hat{\boldsymbol{m}}_1 > s \cdot \hat{\boldsymbol{m}}_0$ and 0 otherwise. Figure 11 shows that the obtained accuracy approaches the theoretical maximum.

TEROs vs. Bistable Rings To study the similarities and differences between TEROs and Bistable Rings, we vary the number of propagating events, ψ, for a given configuration C_{osc}. Under the assumption that periods X_1 and X_2 of the propagating events remain roughly the same, Fig. 12a shows that the throughput of a Bistable Ring PUF/TRNG is roughly proportional to ψ. This competitive edge of a high ψ is counterbalanced by a slightly decreased entropy in TRNG mode and a slightly increased BER in state-based PUF mode, as shown in Fig. 12b and Fig. 12c respectively. Fig. 12d shows that the state-based PUF response R does not vary with ψ, except when $X_1 \approx X_2$.

[1] From 64 TEROs, Bossuet et al. [5] extract 63, 126, 189, or 252 counter-based PUF bits, even though only 32 uncorrelated bits can be extracted [9, Section 3.2.4]. Note that 64 TEROs allow for 64 uncorrelated state-based PUF bits.

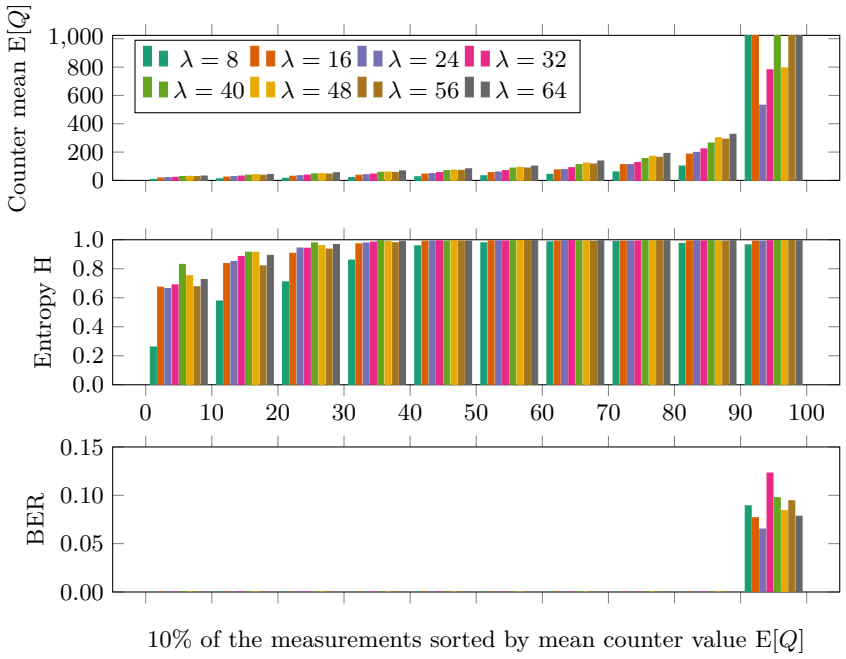

Fig. 10: The output statistics of TEROs having $\lambda \in \{8, 16, 24, 32, 40, 48, 56, 64\}$ stages. For each value of λ, we randomly generate 100 configuration pairs $(C_{\text{rst}}, C_{\text{osc}})$ for each out of 8 parallel circuit instances, and evaluate their outputs 100 times. Subsequently, configuration pairs $(C_{\text{rst}}, C_{\text{osc}})$ are sorted in ascending order of the mean counter value, $\text{E}[Q]$, and subdivided into 10 bins. For each bin, we plot averages of (a), the mean counter value, $\text{E}[Q]$, (b) the Shannon entropy of the counter-based TRNG output, $\text{H}(N)$, and (c) the BER of the state-based PUF response, $\Pr(R \neq R_{\text{vote}})$, where R_{vote} is the most frequently occurring bit among 100 evaluations.

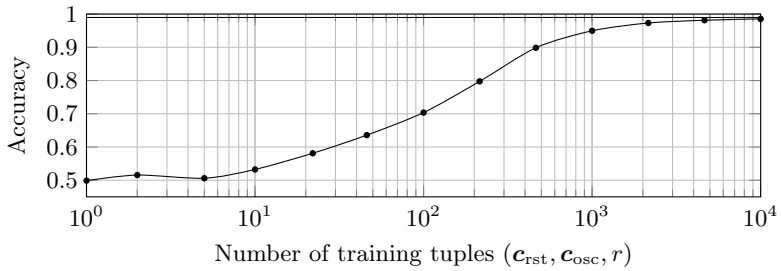

Fig. 11: Machine learning of our reconfigurable TERO in PUF mode. The test set, which is used to estimate the accuracy, consist of 1000 tuples $(\boldsymbol{c}_{\text{rst}}, \boldsymbol{c}_{\text{osc}}, r)$. The *bit error rate* (BER) is circa 1% and determines an upper bound for the accuracy.

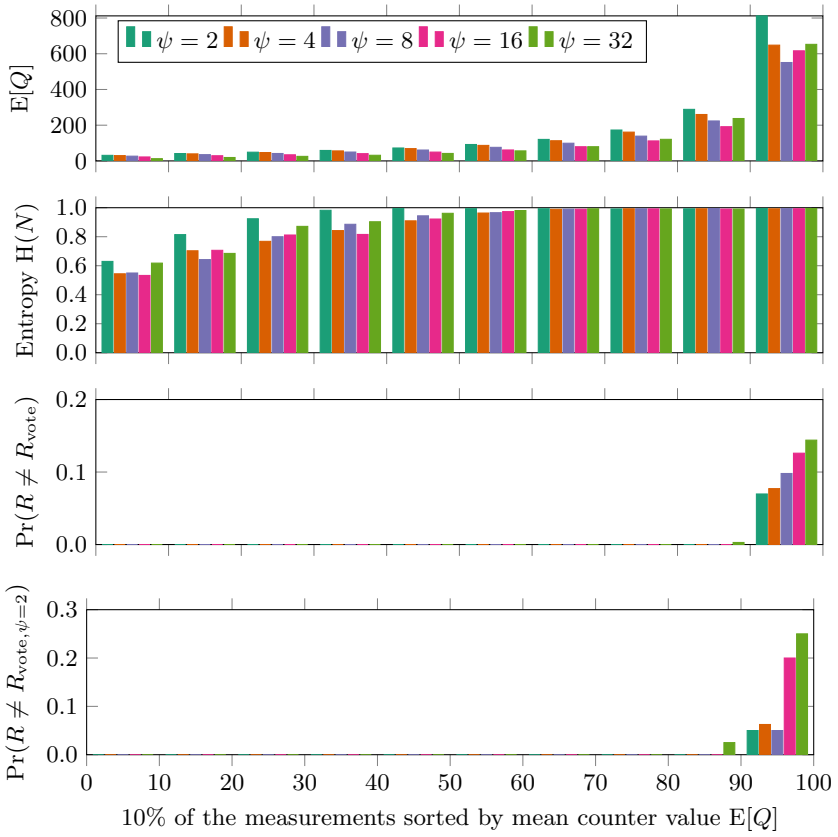

Fig. 12: The output statistics of Bistable Rings having $\lambda = 64$ stages and $\psi \in \{2, 4, 8, 16, 32\}$ propagating events. The configuration C_{osc} is independent of ψ, i.e., we randomly generate 100 tuples ($C_{\text{rst},\psi=2}$, $C_{\text{rst},\psi=4}$, $C_{\text{rst},\psi=8}$, $C_{\text{rst},\psi=16}$, $C_{\text{rst},\psi=32}$, C_{osc}) for each out of 8 parallel circuit instances, and evaluate their outputs 100 times. Subsequently, configuration pairs $(C_{\text{rst}}, C_{\text{osc}})$ are sorted in ascending order of the mean counter value, E[Q], and subdivided into 10 bins. For each bin, we plot averages of (a), the mean counter value, E[Q], (b) the Shannon entropy of the counter-based TRNG output, H(N), (c) the BER of the state-based PUF response, $\Pr(R \neq R_{\text{vote}})$, and (d) the BER of the state-based PUF response using TERO mode as a reference, $\Pr(R \neq R_{\text{vote},\psi=2})$.

Model Fitting Whenever mathematical models are fit to experimental data, beware that the strength of the validation decreases with the number of degrees of freedom. Bernard et al. [4, Fig. 9] model four measured curves of PMF $\mathsf{f}_Q(q)$ using a total of twelve parameters, so it is not surprising that four tight fits can be produced despite the physical flaws. To avoid this pitfall, we model 15 curves using 16 parameters in Fig. 13. Given that all 15 curves are measured using a single reconfigurable TERO instance, there is no reason to assume that noise parameter σ in Eq. (3) would greatly vary. We thus fit a vector $(\mu_1\,\mu_2\,\cdots\,\mu_{15}\,\sigma)$ using non-linear optimization. We only selected curves where the duty cycles approaches 100%, as σ might change for the 0% direction.

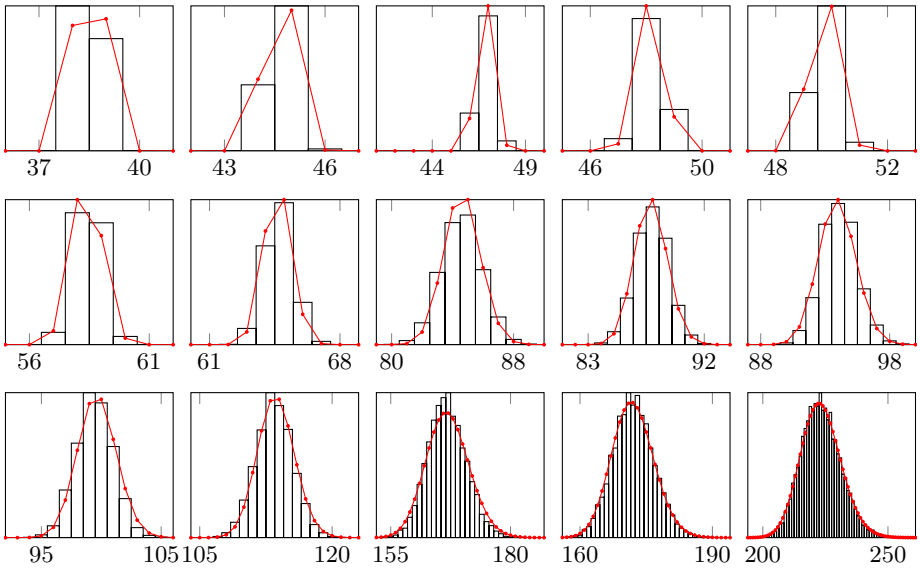

Fig. 13: Model fitting of 15 PMFs $\mathsf{f}_Q(q)$. The bars represent experimental data using 20000 evaluations each. The red curves are the model in Eq. (3).

Indirect temperature analysis Unfortunately, the relevance of self-heating is hard to assess, given that one cannot simply measure the local, on-chip temperature of a TERO circuit. Equally problematic, a global, off-chip temperature measurement would be technically feasible but is likely to misrepresent such a small circuit, in part because other nearby elements generate heat too. These setbacks, however, do not prevent us from measuring the minuscule changes of the period X as a function of the time T using an oscilloscope, as shown in Fig. 14. To enable long measurements and to exclude injection locking as a contributing factor, we opt for RO mode rather than TERO mode. Most notably, periods X increase up until time $T \approx 20\,\mu\text{s}$, then slightly decrease, and finally

level off. This behavior is not inconsistent with an initial heating followed by self-regulation. Nevertheless, due to the indirect approach, we cannot exclude that another physical phenomenon or a measurement artifact is responsible.

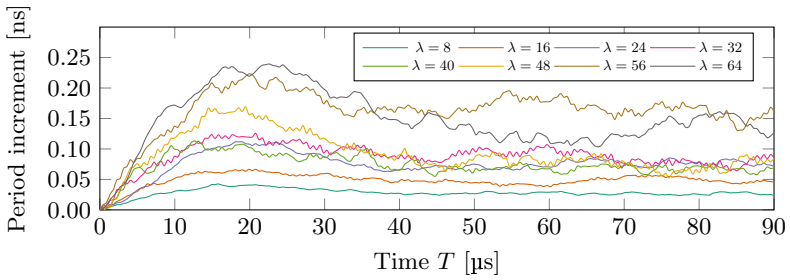

Fig. 14: An indirect temperature analysis for ROs having $\lambda \in \{8, 16, 24, 32, 40, 48, 56, 64\}$ stages. We plot the period increment $X(T) - X(0)$, where time $T = 0$ at the start of any given oscillation.

6 Conclusion

Despite the large number of publications analyzing TEROs, we conclude that this building block is not yet properly understood on a basic architectural level. For long TEROs, we have demonstrated that the oscillatory behavior is dominated not by a drafting effect but by process variations from two mutually exclusive sets of transistors. Furthermore, by establishing a strong link between TEROs and Bistable Rings, we discovered a previously unexplored design trade-off that can benefit both PUF and TRNG applications. For PUF applications in particular, we have also shown that the literature's exclusive focus on counter-based designs is not necessarily justified, given that state-based designs appear to be competitive. The latter type of designs also renders counter value corruptions, which have first been identified as a problem in this article, irrelevant.

Acknowledgement

The author thanks Yiming Lu and Yizhong Hu, who are both coworkers at OSR, for troubleshooting with the use of Xilinx tools and electronic equipment. For Yiming, this includes his soldering work displayed in Fig. 7a.

References

1. Abidi, A.A.: Phase noise and jitter in CMOS ring oscillators. IEEE Journal of Solid-State Circuits **41**(8), 1803–1816 (Aug 2006). https://doi.org/10.1109/JSSC.2006.876206

2. Agne, A., Hangmann, H., Happe, M., Platzner, M., Plessl, C.: Seven recipes for setting your FPGA on fire – a cookbook on heat generators. Microprocess. Microsyst. **38**(8), 911–919 (Nov 2014). https://doi.org/10.1016/j.micpro.2013.12.001
3. Amaki, T., Hashimoto, M., Mitsuyama, Y., Onoye, T.: A design procedure for oscillator-based hardware random number generator with stochastic behavior modeling. In: Chung, Y., Yung, M. (eds.) Information Security Applications(WISA 2010). Lecture Notes in Computer Science, vol. 6513, pp. 107–121. Springer (2010). https://doi.org/10.1007/978-3-642-17955-6_8
4. Bernard, F., Haddad, P., Fischer, V., Nicolai, J.: From physical to stochastic modeling of a TERO-based TRNG. Journal of Cryptology **32**(2), 435–458 (Apr 2019). https://doi.org/10.1007/s00145-018-9291-2
5. Bossuet, L., Ngo, X.T., Cherif, Z., Fischer, V.: A PUF based on a transient effect ring oscillator and insensitive to locking phenomenon. IEEE Transactions on Emerging Topics in Computing **2**(1), 30–36 (2014). https://doi.org/10.1109/TETC.2013.2287182
6. Bucci, M., Luzzi, R.: Design of testable random bit generators. In: Rao, J.R., Sunar, B. (eds.) 7th Workshop on Cryptographic Hardware and Embedded Systems (CHES 2005). Lecture Notes in Computer Science, vol. 3659, pp. 147–156. Springer (Sep 2005). https://doi.org/10.1007/11545262_11
7. Chen, Q., Csaba, G., Lugli, P., Schlichtmann, U., Rührmair, U.: The bistable ring PUF: A new architecture for strong physical unclonable functions. In: International Symposium on Hardware-Oriented Security and Trust (HOST 2011). pp. 134–141. IEEE (Jun 2011). https://doi.org/10.1109/HST.2011.5955011
8. Cherkaoui, A., Bossuet, L., Marchand, C.: Design, evaluation, and optimization of physical unclonable functions based on transient effect ring oscillators. IEEE Trans. Information Forensics and Security **11**(6), 1291–1305 (2016). https://doi.org/10.1109/TIFS.2016.2524666
9. Delvaux, J.: Security Analysis of PUF-Based Key Generation and Entity Authentication. Ph.D. thesis, KU Leuven (2017), dawu Gu and Ingrid Verbauwhede (promotors)
10. Drutarovský, M., Varchola, M.: Analysis of randomness sources in transition effect ring oscillator based TRNG. In: 8th Workshop on Cryptographic Architectures Embedded in Reconfigurable Devices (CryptArchi 2010)
11. Hars, L.: Random number generation based on oscillatory metastability in ring circuits. Cryptology ePrint Archive, Report 2011/637 (Nov 2011), https://eprint.iacr.org/2011/637
12. Jun, B., Kocher, P.: The Intel random number generator. Tech. rep., Cryptography Research, Inc. (Apr 1999)
13. Killmann, W., Schindler, W.: A proposal for: Functionality classes for random number generators. Tech. rep. (Sep 2011)
14. Li, T., Wu, L., Zhang, X., Wu, X., Zhou, J., Wang, X.: A novel transition effect ring oscillator based true random number generator for a security SoC. In: Conference on Electron Devices and Solid-State Circuits (EDSSC 2017). pp. 1–2. IEEE (Oct 2017). https://doi.org/10.1109/EDSSC.2017.8126539
15. Lim, D.: Extracting Secret Keys from Integrated Circuits. Master's thesis, Massachusetts Institute of Technology (May 2004)
16. Lin, X.S.: Double barrier hitting time distributions with applications to exotic options. Insurance: Mathematics & Economics **23**(1), 45–58 (Oct 1998). https://doi.org/10.1016/S0167-6687(98)00021-3

17. Ma, Y., Lin, J., Chen, T., Xu, C., Liu, Z., Jing, J.: Entropy evaluation for oscillator-based true random number generators. In: Batina, L., Robshaw, M. (eds.) Cryptographic Hardware and Embedded Systems (CHES 2014). Lecture Notes in Computer Science, vol. 8731, pp. 544–561. Springer (Sep 2014). https://doi.org/10.1007/978-3-662-44709-3_30
18. Malik, K.: Ring oscillator for determing select-to-output delay of a multiplexer (Jul 2007), US patent 2007/0126515 A1
19. Marchand, C., Bossuet, L., Mureddu, U., Bochard, N., Cherkaoui, A., Fischer, V.: Implementation and characterization of a physical unclonable function for IoT: A case study with the TERO-PUF. IEEE Trans. on CAD of Integrated Circuits and Systems **37**(1), 97–109 (2018). https://doi.org/10.1109/TCAD.2017.2702607
20. Mureddu, U., Bochard, N., Bossuet, L., Fischer, V.: Experimental study of locking phenomena on oscillating rings implemented in logic devices. IEEE Transactions on Circuits and Systems I: Regular Papers (2019)
21. Mureddu, U., Colombier, B., Bochard, N., Bossuet, L., Fischer, V.: Transient effect ring oscillators leak too. Cryptology ePrint Archive, Report 2019/300 (Mar 2019), https://eprint.iacr.org/2019/300
22. Rabaey, J.M., Chandrakasan, A., Nikolic, B.: Digital Integrated Circuits, p. 800. Pearson, 2 edn. (Jan 2003)
23. Reyneri, L.M., Del Corso, D., Sacco, B.: Oscillatory metastability in homogeneous and inhomogeneous flip-flops. IEEE Journal of Solid-State Circuits **25**(1), 254–264 (Feb 1990). https://doi.org/10.1109/4.50312
24. Rioul, O., Solé, P., Guilley, S., Danger, J.L.: On the entropy of physically unclonable functions. In: International Symposium on Information Theory (ISIT 2016). pp. 2928–2932. IEEE (Jul 2016). https://doi.org/10.1109/ISIT.2016.7541835
25. Sahoo, D.P., Nguyen, P.H., Chakraborty, R.S., Mukhopadhyay, D.: On the architectural analysis of arbiter delay PUF variants. Cryptology ePrint Archive, Report 2016/057 (Jan 2016), https://eprint.iacr.org/2016/057
26. Tebelmann, L., Pehl, M., Immler, V.: Side-channel analysis of the TERO PUF. In: Polian, I., Stöttinger, M. (eds.) Constructive Side-Channel Analysis and Secure Design (COSADE 2019). Lecture Notes in Computer Science, vol. 11421, pp. 43–60. Springer (Apr 2019). https://doi.org/10.1007/978-3-030-16350-1_4
27. Turan, M.S., Barker, E., Kelsey, J., McKay, K.A., Baish, M.L., Boyle, M.: NIST special publication 800-90B – recommendation for the entropy sources used for random bit generation. Tech. rep., NIST (Jan 2018). https://doi.org/10.6028/NIST.SP.800-90B
28. Varchola, M., Drutarovský, M.: New high entropy element for FPGA based true random number generators. In: Mangard, S., Standaert, F.X. (eds.) 12th Workshop on Cryptographic Hardware and Embedded Systems (CHES 2010). Lecture Notes in Computer Science, vol. 6225, pp. 351–365. Springer (Aug 2010). https://doi.org/10.1007/978-3-642-15031-9_24
29. Varchola, M., Drutarovský, M., Fischer, V.: New universal element with integrated PUF and TRNG capability. In: Conference on Reconfigurable Computing and FPGAs (ReConFig 2013). pp. 1–6. IEEE (Dec 2013). https://doi.org/10.1109/ReConFig.2013.6732311
30. Winstanley, A., Greenstreet, M.: Temporal properties of self-timed rings. In: Margaria, T., Melham, T. (eds.) Correct Hardware Design and Verification Methods (CHARME 2001). pp. 140–154. Springer (Sep 2001). https://doi.org/10.1007/3-540-44798-9_12
31. Xilinx: 7 series FPGAs configurable logic block user guide. https://www.xilinx.com/support/documentation/user_guides/ug474_7Series_CLB.pdf (Sep 2016)

32. Xu, X., Rührmair, U., Holcomb, D.E., Burleson, W.P.: Security evaluation and enhancement of bistable ring PUFs. In: Mangard, S., Schaumont, P. (eds.) 11th Workshop on Radio Frequency Identification: Security and Privacy Issues (RFIDsec 2015). Lecture Notes in Computer Science, vol. 9440, pp. 3–16. Springer (Jun 2015). https://doi.org/10.1007/978-3-319-24837-0_1