ROBoost: A Study of FPGA Logic-Based Power-Wasting Primitives

Dina G. Mahmoud[1][0000−0003−0720−1342], Simone Andreani[2], Vincent Lenders[3][0000−0002−2289−3722], and Mirjana Stojilović[2][0000−0001−5649−5020]

[1] The American University in Cairo, Egypt dina-mahmoud@aucegypt.edu
[2] EPFL, Lausanne, Switzerland simone.andreani@epfl.ch, mirjana.stojilovic@epfl.ch
[3] Cyber-Defence Campus, armasuisse, Thun, Switzerland vincent.lenders@armasuisse.ch

Abstract. Heterogeneous computing systems increasingly leverage FPGAs in the cloud and embedded use cases. With cloud FPGAs being remotely accessible, security is a critical concern. Recent studies show adversaries can exploit FPGA logic to create and remotely deploy malicious power-wasting circuits that consume excessive dynamic power, potentially injecting faults or causing denial of service. This work analyzes the most common reconfigurable power-wasting primitives to assess their power consumption, detection challenges, and attack effectiveness. We further propose new, logic-based, and resource-efficient variations of these circuits and experimentally evaluate them on two families of AMD FPGAs. Finally, we discuss factors influencing attack effectiveness and compare the studied designs' trade-offs.

Keywords: FPGAs · remote attacks · voltage drop · hardware security

1 Introduction

The rising demand for accelerated computing has outpaced general-purpose CPUs, pushing embedded and cloud systems to adopt specialized processing units. Field-programmable gate arrays (FPGAs) are favored for their fine-grained parallelism and reconfigurability. Cloud service providers (CSPs) like Amazon and Alibaba now offer FPGA-accelerated instances [27], while AMD and Intel integrate FPGAs with CPUs in systems-on-chip (SoCs).

The growing adoption and remote accessibility of FPGAs in the cloud have made their security critical [27]. Users with low-level control on the FPGA fabric can deploy malicious bitstreams, creating *power wasters*—FPGA circuits that draw excessive power. These circuits can overwhelm power supplies, causing voltage drops and, in turn, timing violations or even FPGA resets [12]. Other remote exploits include passive circuits monitoring activity for side-channel attacks [10].

This paper focuses on FPGA logic-based power-wasting primitives. While prior research explores their use in attacks and defenses [11,12,15,17,19], implementation details and voltage drop capabilities remain underexplored. To effectively counter current and future threats, a deeper understanding of power waster

characteristics and the extent of possible improvements is needed. Additionally, evaluating the ease of implementation, portability, and associated constraints is vital to assess their risk to cloud and remote FPGA applications.

In this work, we compare known logic-based power-wasting primitives on two AMD FPGAs [8, 24] and examine the factors influencing their success. Building on these findings, we propose new variations that validate these factors, including the ability to confine primitives to specific logic regions*. Some of these variations rival the best-known designs while bypassing design rule check (DRC) warnings used by CSPs like Amazon (offering AMD FPGA instances) to block malicious circuits [14]. This makes them deployable in current cloud environments.

The paper is structured as follows: Section 2 provides background on power-wasting attacks and voltage measurement. Section 3 reviews power waster types. Section 4 details the experimental setup. Section 5 presents the results. We discuss findings in Section 6 and conclude in Section 7.

2 Background

2.1 Power-Wasting Attacks

Power consumption in electronic circuits depends on static leakage and dynamic signal changes. Dynamic power varies with the circuit implementation and operations, in the function of voltage, switching frequency, and load capacitance:

$$P_{\text{dyn}} \propto C_{\text{L}} \times V_{\text{cc}}^2 \times f. \tag{1}$$

Here, P_{dyn} is the dynamic power consumption, C_{L} the load capacitance, V_{cc} the supply voltage, and f the switching frequency [7]. The clock frequency and the frequencies of combinational signals toggling determine the switching frequency.

Remote power-wasting exploits on FPGAs attracted attention after Gnad et al. demonstrated the first DoS attack using ring oscillators (ROs) [12]. These attacks leverage short combinational feedback paths, creating high-frequency oscillations that increase power consumption. Power-wasting primitives can also target higher load capacitance to amplify power consumption further. Current variations caused by signal switching lead to voltage drops in the power distribution network (PDN), affecting signal propagation delays and potentially causing faults in memory elements [5, 13, 22, 25, 32]. Voltage drops also increase flip-flop (FF) setup and hold times [6], potentially causing unsafe operating conditions.

To combat malicious combinational loops, Amazon prevents their use in cloud FPGAs by using AMD Vivado's DRCs [3], but exploits still exist. Later research employs FFs and latches to break the combinational loop and bypass these checks [15, 23]. Continuing research efforts are devoted to detecting power wasters on one side and developing stealthier malicious designs on the other [27].

* For the reproducibility of the experiments and the results of this work, we make the associated artifacts openly available [16].

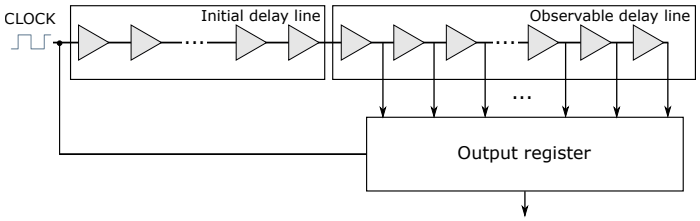

Fig. 1: High-level view of a TDC sensor for FPGA on-chip voltage variations [26].

2.2 On-Chip Voltage Measurement

Similarly to power-wasting primitives, FPGA logic can turn into sensors capturing delay changes that correlate with PDN voltage variations caused by on-chip activity [10, 26]. These sensors are typically used in remote power side-channel attacks, where an adversary steals secrets or detects voltage drops [10, 32]. The most common sensor is the time-to-digital converter (TDC, Fig. 1). It measures the propagation depth of a clock signal edge while it travels down a fine-grained delay line. This delay line, sensitive to voltage variations, is typically implemented using the carry chain logic. The delay line is tapped, meaning the sensor has an output register that periodically captures the clock's propagation depth. The output register value is converted to the Hamming weight (HW) to obtain one sensor *sample*, which directly correlates with the on-chip voltage. Finally, the on-chip voltage profile can be reconstructed by collecting the sequence of sensor samples for a given time. We use on-chip TDC sensors to measure voltage drops caused by the power wasters, mimicking the real remote undervolting attack scenario [10].

3 Types of Power Wasters

Early power wasters, known as *combinational power wasters*, used only combinational elements in FPGAs to create self-oscillating circuits that rely on feedback loops around combinational logic (e.g., lookup tables (LUTs)). However, FPGAs also offer flip-flops. Accordingly, a second category of power wasters, noncombinational or *FF-based power wasters*, becomes a possibility. FF-based power-wasting primitives introduce sequential elements into the feedback path and may require a clock signal. Next, *improved power wasters* refine previous designs to enhance resource utilization, stealth, and power consumption by leveraging FPGA fabric properties. Finally, there are *hidden power wasters* that embed malicious designs within circuits that appear benign; hiding techniques are beyond the scope of this work.

Effective power wasters maximize dynamic power consumption by increasing switching frequency, load capacitance, or both, as shown by equation (1). High switching frequencies can be achieved through fast clocks, feedback loops, or glitches, while high fanout boosts load capacitance. We discuss these categories of power-wasting primitives in detail in the following subsections.

(a) RO (b) RO2-*cf*

Fig. 2: A ring oscillator implemented with (a) one six-input LUT and (b) two five-input LUTs (I5 needs to be a logical `1` for the two outputs to be independent). The combinational feedback connections are highlighted in red.

3.1 Combinational Power Wasters

Combinational power wasters use only combinational logic elements within the FPGA fabric. Targeting the goal of high switching frequency, a combinational RO (RO-*cf*) is the most straightforward design of a power waster; here, *cf* stands for the combinational feedback. An RO-*cf* consists of an odd number of inverters in a loop. In the simplest FPGA-based implementation, one LUT implementing an inverter is sufficient and would result in the lowest combinational delay. Consequently, the oscillation frequency of one such RO-*cf* would be extremely high, leading to a high power draw with enough instances of RO-*cf*s. In practical attack scenarios, the attacker needs control over when the RO-*cf*s should start (and stop) oscillating [19]. Therefore, typical RO-*cf* designs resemble the circuit in Fig. 2a, where a NAND gate replaces the inverter and is controlled by an enable input. As many FPGAs now support fracturable LUTs, it is also possible to implement two ROs in one LUT for more efficient use of logic resources [23]. Fig. 2b illustrates the design of the corresponding *dual* RO (RO2-*cf*).

Modern FPGAs contain combinational logic elements other than LUTs that can also be programmed to create a self-oscillating circuit. Depending on the target FPGA, users can control MUXes, carry chain elements (CARRY), or digital signal processing (DSP) blocks, creating MUX-based ROs, CARRY-based ROs, and DSP-based ROs [15]. These designs bypass the design rule checks on commercial tools such as AMD Vivado, which can detect the feedback loop only through a LUT [15].

3.2 FF-Based Power Wasters

FF- and latch-based power wasters are designed to break the combinational loop using a flip-flop or a latch to avoid detection. Inserting these elements increases the feedback path length, effectively decreasing the oscillation frequency. It also results in higher resource usage. However, the additional connections increase the load capacitance and FF-based power wasters still prove effective for DoS attacks on cloud FPGA instances [14].

Fig. 3a shows an example design of the FF-based power waster proposed by Giechaskiel et al., who used it not as an attack primitive but as a power side-channel attack sensor [9]. The LUT part is identical to the combinational RO.

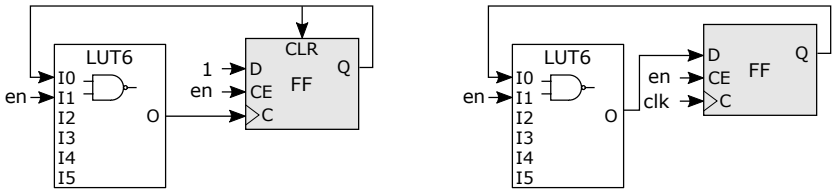

(a) RO-FF (b) RO-FF (PLL)

Fig. 3: FF-based RO with (a) self-oscillations and (b) clock signal from a PLL.

However, the output of the RO LUT acts as the clock for the FF. The `D` input of the FF is fixed to `1`, and according to the states of the clock and the clear signals, the output oscillates. The oscillating FF output acts as an input to the LUT and controls the clear signal of the FF. The continuous change of the FF output due to the back and forth between it being cleared and it becoming a `1` at the rising edge of the clock results in the oscillating behavior of this design [23].

Another variation of FF-based ROs, shown in Fig. 3b, uses a high-frequency clock generated from a phase-locked loop (PLL) on the FPGA. In this case, the `D` input of the FF is the output of the LUT, and the enable signal is used to control both the FF and the LUT [23]. Suppose a latch is available in the programmable logic. In that case, the design may be simplified by controlling the latch's enable with the enable signal, connecting the `D` input to the output of an inverter LUT, and connecting a constant `1` to the clock of the latch. The input of the LUT would be driven by the latch, creating the RO-L design [9].

Additionally, modern FPGAs can implement shift registers in their programmable fabric. Suppose the clock that controls the FFs of the shift register is fast, and the register is initialized with a sequence of values to ensure the outputs change every clock cycle (an alternating sequence of `1`s and `0`s). In that case, the outputs of each FF will oscillate at the high clock frequency, resulting in considerable power consumption [23].

FF-based power wasters are not limited to the patterns of FFs and LUTs. For instance, deliberately created long routing paths with different delays, when connected to a logic gate, can result in inputs taking a long time to stabilize and, consequently, multiple transitions (i.e., glitches) at the output. Matas et al. [20] proposed and evaluated an example of such a design with XOR logic.

3.3 Improved Power Wasters

While effective, the typical primitives with ROs and FFs can be improved to increase the dynamic power consumption further. High switching frequency is the main feature of RO-*cf*s and RO2-*cf*s, making them suitable for power attacks. Therefore, ROs can be enhanced to take advantage of the other characteristic of good power wasters: the load capacitance. La et al. proposed a design for enhanced ring oscillators (EROs), increasing the fanout of each RO and the routing used to consume more power [15]. One instance of an ERO-*cf* is shown in Fig. 4;

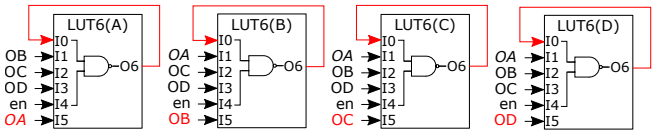

Fig. 4: Enhanced ring oscillator (ERO). The feedback connections are in red. The output of the first LUT is in italic to highlight its connections to the other LUTs.

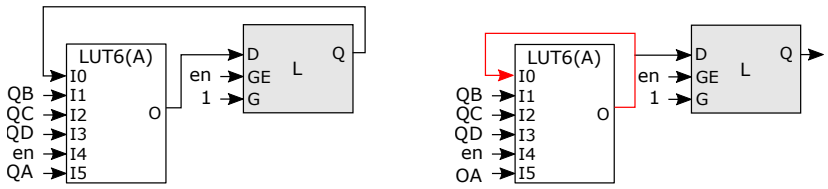

(a) ERO-L (b) ERO-L-*cf*

Fig. 5: Latch-based ERO with (a) latch feedback and (b) combinational feedback.

it comprises four LUTs implementing a NAND of the enable and the feedback signals. The enable signal of the ERO-*cf* drives one input of each LUT. Two LUT inputs are connected to the output of the same LUT, forming a combinational loop. The remaining three inputs are driven by the outputs of the other three LUTs in the ERO-*cf*. Compared to the traditional ROs, LUTs in the ERO-*cf* drive a higher capacitive load, resulting in increased power consumption [15].

Following the same reasoning behind the design of the EROs and the use of FFs and latches in the RO-FF and RO-L primitives, we design new variations that aim to combine the best features of existing designs:

- We combine EROs with latches, once breaking the combinational loop using the latch and preserving the increased routing of the EROs (ERO-L, Fig. 5a), and once maintaining the combinational loop of the EROs (ERO-L-*cf*, Fig. 5b).
- We combine EROs with FFs, instead of latches. Two variations are designed and tested, one breaking the combinational feedback with an FF (ERO-FF, Fig. 6a) and one preserving it (ERO-FF-*cf*, Fig. 6b).

For the versions preserving the feedback loop, we configure the additional inputs of the LUTs to be driven by the FF outputs due to the constraints on routing the LUT outputs when all the FFs are used (i.e., if both FFs are used, only one LUT output, `O6` in Fig. 6, can be routed out of the slice back to the LUT inputs). Typically, a PLL drives the FF clock input. We also propose a variation clocked from an RO-L, avoiding the requirement of a PLL.

4 System Design

We implement several power waster designs to compare them. The first designs we consider are RO-*cf* and RO2-*cf* (Figs. 2a and 2b). These two designs are

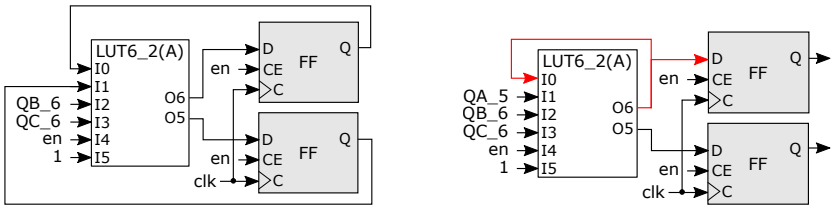

(a) ERO-FF (b) ERO-FF-*cf*

Fig. 6: FF-based ERO with (a) FF feedback and (b) combinational feedback.

the standard baseline against which to compare because many of the demonstrated DoS and fault-injection exploits leveraged combinational ROs [27]. Our analysis also considers ERO-*cf* circuits, representing the most effective power wasters in literature [11,15]. Additionally, we consider two designs without combinational loops: latch-based ROs (RO-L) [15] and FF-based ROs clocked from a phase-locked loop (RO-FF (PLL)) [23]. Neither of these two designs generates DRC warnings. We do not consider self-oscillating FF-based wasters for two reasons: first, they generate a gated clock warning [15] and, second, they cannot be densely packed because the FFs within a slice on an AMD FPGA must share the clock signal if they are to be used simultaneously [29]. We also evaluate our new variations of power wasters: ERO-L (Fig. 5a), ERO-L-*cf* (Fig. 5b), ERO-FF (Fig. 6a), and ERO-FF-*cf* (Fig. 6b). For the two FF-based designs, we evaluate them when (a) a PLL or (b) an ERO-L generates the clock.

The key result of power wasters activity is a voltage drop. An effective power waster causes a more significant voltage drop than other designs. To compare malicious designs, we monitor the on-chip FPGA voltage variations during their activity. The design causing the most variation and the largest voltage drop is deemed the most effective. While external measurements are possible, internal FPGA-based measurements avoid external equipment and directly capture the power wasters' impact on the collocated and concurrently executing FPGA applications. We use TDC on-chip voltage-variations sensors, described in Section 2, whose readings directly correlate with voltage changes [21].

We perform experiments on Pynq Z1 and Genesys-ZU boards, covering two AMD FPGA families. The variety helps assess design portability and how FPGA fabric features affect power waster implementation.

4.1 Pynq-Z1 Setup

The FPGA PL of the PYNQ-Z1 SoC is of the AMD 7-series. It contains 13,300 logic slices, each with four six-input LUTs and eight flip-flops [24]. Four of the eight FFs per slice can act as latches. The SoC also includes a dual-core Cortex-A9 processing system (PS). We use the PS to control the PL, and the two parts communicate through advanced extensible interface (AXI) general-purpose inputs/outputs (GPIOs). The PL clock frequency is 100 MHz. This frequency ensured the correct operation of all circuits implemented within the PL. An

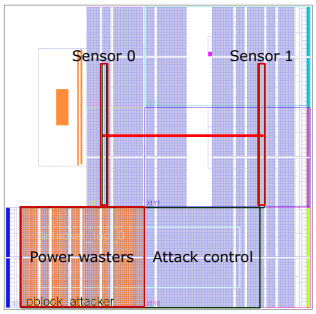

Fig. 7: Design floorplan on the Pynq-Z1.

experimental attack run lasts for 256 clock cycles (2.56 µs).Vivado 2020.1 was used to generate all of the tested hardware circuits. The enable signal of the power wasters has a period of 150 clock cycles and a duty cycle of 50%.

To mimic the limitations on the region available to an adversary and to avoid accidental reset, we limit the testing of all of our power wasters to the clock region X0Y0, which contains 50 columns, where each column has 50 slices, as shown in Fig. 7. The power wasters' control is limited to the region adjacent to it. Two TDC sensors are used, one placed on the same side as the power wasters on top of X0Y0 (sensor 0) and the second (sensor 1) placed on the other side of the PL farther away from the wasters. Each sensor uses 64 CARRY4 elements, resulting in 256 bits in the sensor output (sensor readings—i.e., the Hamming weight of the output register—lie in the 0–256 range).

We organize the power wasters in *block*s, each covering one column of the chosen region. Depending on the number of primitives we can use in an FPGA slice, we either group two blocks in one *node* or restrict the node size to one block. The PS sends the signals to the control circuit to determine how many power waster nodes are active in an experimental run. The nodes are activated one by one in a staggered manner with 25 steps at most (for designs with 50 nodes, two nodes are activated at each step). The staggered activation makes the effect of the attack gradual. Additionally, it has been shown to be more effective because it reduces the magnitude of the sudden voltage changes occurring when the power wasters are activated and deactivated [18].

4.2 Genesys-ZU Setup

The second testing platform is the Genesys-ZU board, which includes a Zynq UltraScale+ MPSoC [8]. The FPGA PL contains 71,000 LUTs, organized in slices of eight, where each slice also contains 16 FFs, for a total of 8,875 logic slices. Unlike the Pynq-Z1, all the FFs within a slice can act as latches. The MPSoC includes a quad-core ARM Cortex-A53 application processing unit (APU). We leverage this APU to control the PL through the AXI GPIOs. The PL clock frequency is 150 MHz, which guarantees correct design operation and takes ad-

vantage of the faster logic within the Zynq UltraScale fabric. One experimental attack run corresponds to 384 clock cycles (2.56 µs).

The enable signal of the power wasters has a period of 800 clock cycles and a duty cycle of 50% (i.e., circuits are active for 400 clock cycles). This period is longer than the attack duration, thus we do not get the full-length voltage trace captured. However, this choice improves the likelihood of getting reasonably accurate voltage readings from the on-chip system monitor so that we can compare them to the TDC sensor readings. The analog-to-digital converter (ADC) used in the system monitor operates at 0.2 MHz [28], a much lower rate compared to the TDC sensor (150 MHz). Hence, we take the readings from the system monitor only as an indication of the voltage values rather than an accurate measure of the minimum voltage. This is because the minimum value may be skipped due to the reduced sampling speed of the ADC with respect to the TDC. Having the power wasters active for the entire attack duration means that the lowest voltage is expected to last longer, improving the likelihood of the system monitor capturing it.

The power wasters and their control are constrained to the clock regions X0Y0 and X1Y0. The only constraint on the power wasters within the two clock regions is that each group of logic elements forming power wasters that can fit into one slice is packed and not spread across slices. We implement 25 nodes of 400 LUTs each (and 800 FFs when FFs are used). The power wasters are also activated in a staggered fashion. We use one TDC sensor constrained to the right side of the FPGA. The TDC comprises 64 CARRY8 elements, resulting in a sensor reading in the 0–512 range.

5 Comparison of Power Wasters

The experiments commence with recording the sensors' calibration parameters (for reproducibility), followed by baseline readings from the sensor when no power wasters are active. Then, we repeat the power waster's activity ten times to ensure the repeatability and consistency of the results. Each run records the sensor's average, maximum, and minimum readings. We then compute the sample-wise averages across all runs, noting the combined results' maximum, minimum, and average sensor readings. All values are reported relative to the baseline sensor average (e.g., negative value means the activity resulted in lower voltage compared to the baseline in the absence of activity).

Voltage and power variations depend on the circuit design, which determines its frequency of oscillation and load capacitance. Measuring or estimating the oscillation frequency using the synthesis tools are two ways to assess that comparison metric. Regarding capacitance, even though it is not immediately available, the designs can be compared based on their fanout and use of routing resources (e.g., higher fanout increases routing demand and load capacitance).

Table 1: FPGA resource use and DRC warnings on Pynq-Z1. Critical warnings are italicized. Parentheses show the clock source. In gray, wasters introduced in this work.

Power waster	**LUTs**	**Flip-Flops (FFs)**	**Warnings**
RO [12]	10k	0	*LUTLP-1*
RO2 [15]	10k	0	*LUTLP-1*
ERO [15]	10k	0	*LUTLP-1*
RO-L [15]	10k	10k	N/A
RO-FF (PLL) [23]	10k	20k	N/A
ERO-L-*cf*	10k	10k	*LUTLP-1*
ERO-L	10k	10k	PDCN-1569
ERO-FF-*cf* (PLL)	10k	20k	*LUTLP-1*
ERO-FF (PLL)	10k	20k	PDCN-1569
ERO-FF-*cf* (L)	10k	19.8k	*LUTLP-1*
ERO-FF (L)	10k	19.8k	PDCN-1569

5.1 Experimental Results on Pynq-Z1

Packing of Power Wasters Within clock region X0Y0, we implement the power wasters listed in Table 1. Since they use local routing, Vivado reports no routing congestion. Therefore, Table 1 reports only logic resources used. We also report DRC warnings, as they correlate with the ease of detection and deployment within a cloud environment. All power wasters use all LUTs within the slice. RO2-*cf* uses each LUT as two LUT5s, but while each LUT generates two outputs, the number of LUTs remains unchanged. Aside from purely combinational power wasters, the designs use the FFs within the slice. Given the limitation of the 7-series FPGAs that only four of each eight FFs in a slice can act as latches, all latch-based designs use half of the available FFs. The FF-based power wasters use all available FFs. Versions of ERO-FF-*cf* and ERO-FF clocked using a RO-L (last two rows) can use only half the registers within in slices where a clock signal is generated. With one clock signal per column and 50 columns, the design uses only 19,800 FFs (50×4 FFs unused).

The designs without a latch or an FF in the feedback path generate a critical DRC warning (`LUTLP-1`), pointing to the combinational loop. All designs with increased routing following the ERO design generate a warning (`PDCN-1569`) related to unused inputs of the LUT. Such a warning is not critical and can occur for many designs, including AMD IPs, and can be safely ignored [4]. All other designs have no DRC warnings, making them suitable for generating oscillations when the target platform forbids combinational loops. One of the designs we excluded is the self-oscillating FF, where the clock for the FF is generated from its inverted output. This circuit generated a gated clock warning, removing the stealth advantage [15]. Moreover, in 7-series FPGAs, the FFs within a slice must use the same clock signal to be all used simultaneously, making the self-oscillating FF unable to fully utilize the resources [29]. Similar constraints exist for other architectures, including the UltraScale [30].

Sensor Readings Starting with combinational power wasters, we compare them to understand how their features affect power consumption and to establish a baseline against which to measure other wasters' effectiveness. As expected, our

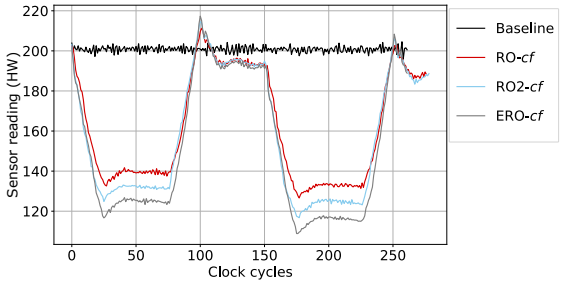

Fig. 8: Comparison of the sensor readings for the combinational power wasters.

Table 2: Max, min, average, and peak-to-peak Sensor 0 readings relative to the baseline (Pynq-Z1). Wasters are sorted from the most to the least effective. The best values are highlighted in italic if the waster generates DRC warnings and in bold if it does not. The parentheses specify the clock source. In gray, wasters introduced in this work.

Power waster	**Max. reading***	**Min. reading***	**Avg. reading***	**Peak-to-peak**
ERO-L-cf	*21.83*	*-118.07*	*-62.47*	*139.9*
ERO	16.79	-95.61	-50.07	112.40
ERO-FF-*cf* (L)	6.78	-94.82	-49.63	101.6
ERO-L	**16.00**	**-92.90**	**-47.83**	**108.9**
ERO-FF-*cf* (PLL)	6.88	-92.72	-48.02	99.6
RO2	5.21	-82.09	-41.61	87.3
RO	11.84	-76.16	-38.55	88.00
ERO-FF (L)	4.66	-71.44	-36.87	76.1
RO-L	12.52	-69.88	-34.64	82.40
RO-FF (PLL)	3.78	-60.02	-27.02	63.80
ERO-FF (PLL)	4.00	-59.60	-27.37	63.60

*Relative to the baseline average obtained before each experiment

results show that Sensor 0, being closer, is more sensitive to the changes induced by the power wasters. Therefore, the results we report are those from Sensor 0. Fig. 8 shows the drop in the TDC sensor readings for the averaged ten runs when using all attacker nodes. All three designs have comparably high oscillation frequencies, as the combinational loop in all cases includes one LUT. However, the additional routing within ERO-*cf*s and the resulting increase in capacitance make the ERO-*cf*s more effective at wasting power. The denser packing of the RO2-*cf*s results in a more significant voltage drop than the RO-*cf*s. The results in Fig. 8 are consistent with what we expected and the previous work [15].

Table 2 shows the maximum, minimum, and average sensor readings relative to the baseline when the power wasters are using all the available resources within the clock region. Each value in Table 2 is the average of the corresponding quantity over the ten experimental runs. We find that the variations between the runs are minimal. The peak-to-peak is the difference between the averaged maximum and minimum readings, indicating the effectiveness of the power waster. A power waster that causes a substantial drop with respect to the baseline will likely cause a noticeable peak in the voltage when the activity is stopped due to

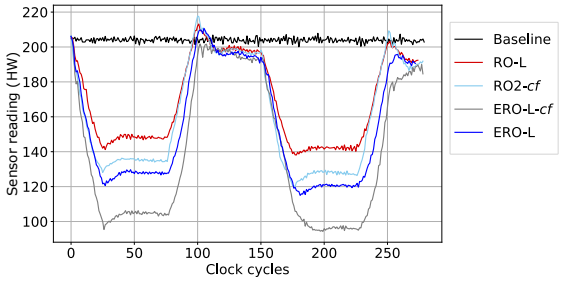

Fig. 9: Comparison of the sensor readings for the latch-based power wasters.

the reaction of the voltage regulator. The peak-to-peak captures the extent of that voltage swing. However, we base our analysis on the drop with respect to the baseline, because the drop is important for the success of a fault-injection or denial-of-service exploit.

For the FF-based power wasters, we first examine the self-oscillating latch (RO-L), where the latch is constantly enabled when the enable signal of the waster is high. The latch's existence on the path decreases the oscillation frequency with respect to the RO designs. While the reduced frequency affects the voltage drop, the effect is not detrimental to the power-wasting capability, as can be seen in Fig. 9 (RO2-*cf* and RO-L lines). Also, given the additional routing and logic elements involved, some additional capacitance is added, resulting in a design with power-wasting capabilities comparable to the ROs, but that is also implemented without generating any warnings (Table 1). The RO-FF power waster fares worse. The decrease in oscillation frequency is only in part caused by the increased path length, due to passing through the register. In this case, the frequency is governed by the maximum frequency that the PLL can generate, which in our case is 465 MHz. The frequency of self-oscillating designs is much higher than that, usually more than double that frequency [15]. Therefore, the FF-based designs clocked from the PLL always perform the worst.

While the RO-L power wasters are a good alternative for combinational power wasters, their power consumption is slightly lower. If an exploit requires a specific voltage drop and the adversary is limited to a particular region, latch-based variants may not generate the needed voltage drop when their combinational counterparts would. Therefore, we explore new design variations, where we combine latches and FFs with EROs to increase the load capacitance. The frequency, minimally affected, is not a primary factor to consider since it is either limited by the PLL or the delay of the feedback path. We also implement versions of these improved power wasters with combinational feedback to be able to analyze the effect of the combinational loops.

Fig. 9 shows the sensor readings obtained for the latch-based designs, compared against those for the RO2-*cf*. Combining the ERO with the latch while maintaining the combinational feedback means that the latch only adds extra routing and load capacitance while not breaking the feedback loop. Therefore,

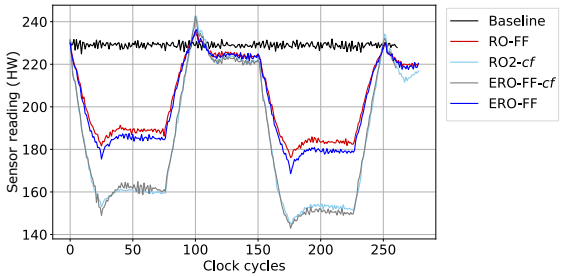

Fig. 10: Comparison of the sensor readings for the PLL-clocked FF-based power wasters.

this achieves the most significant voltage drop. Looking at Table 2, we can see that it even outperforms the ERO-cf. The version without the combinational feedback breaks the loop using the latch. While that decreases the oscillation frequency and makes it consume less power than the ERO-cf, the additional routing and capacitance more than compensate for it, and we observe that it performs better than the ROs.

Including an ERO with sequential elements while maintaining the feedback loop preserves the effectiveness of the EROs for the FF-based designs as well, as shown in Fig. 10. Using the PLL clock or the latch clock induces minimal differences. This is not unexpected, as the delays of passing through the latch are much more significant than those of only passing through a LUT, as reported in the device datasheet [31]. Therefore, the frequency achievable by an RO-L is not significantly higher than that of the PLL. The limited frequency of the change of the FF output compared to the transparent latch makes the improved FF-based designs induce a less significant voltage drop. They also perform slightly worse than standard ERO-cfs, potentially because the increased capacitance at the LUT output decreases the ERO oscillation frequency. Removing the feedback loop makes the FF-based EROs comparable to and slightly better than the RO-FF power waster. The version using a latch clock benefits from the slightly increased frequency of the clock and the increased routing from the ERO, and thus performs somewhat better.

5.2 Experimental Results on Genesys-ZU

Packing of Power Wasters Within the clock regions for the attacker, we instantiate the same number of power wasters as for the Pynq-Z1. As a result, the utilization numbers are similar to those in Table 1. However, there are a few differences. First, all of the FFs within a slice can act as latches, so the number of utilized FFs in all designs that are not purely combinational is the same. Second, a slice contains eight LUTs and 16 FFs, so fewer slices are needed to have the same resource utilization. Our designs use the same hardware description as for the other board, with the only variation being the constraint of two power waster instances (four LUTs and eight FFs each) per slice to ensure that the wasters are packed effectively.

Table 3: Max, min, average, and peak-to-peak Sensor 0 readings relative to the baseline, along with the min and max voltage reported by Vivado (Genesys-ZU). Wasters are ranked by effectiveness. The best values are in italics if triggering critical DRC warnings and in bold if not. In gray, wasters introduced in this work.

Power waster	Max. reading*	Min. reading*	Avg. reading*	Peak-to-peak	Min. Voltage	Max. Voltage
ERO-L-cf	*1.32*	*-371.28*	*-311.03*	*372.60*	*0.762 V*	*0.864 V*
ERO-*cf*	-1.29	-294.29	-244.73	293.00	0.797 V	0.864 V
ERO-L	**1.51**	**-261.69**	**-214.77**	**263.20**	**0.785 V**	**0.864 V**
ERO-FF-*cf* (L)	5.03	-216.37	-177.04	221.40	0.803 V	0.861 V
RO2	5.52	-196.09	-161.19	201.61	0.803 V	0.864 V
ERO-FF-*cf* (PLL)	2.82	-173.38	-132.31	176.20	0.812 V	0.861 V
RO-L	-5.00	-186.51	-154.25	181.51	0.812 V	0.864 V
RO	2.26	-139.44	-112.75	141.70	0.844 V	0.864 V
ERO-FF (L)	11.23	-95.77	-74.72	107.00	0.820 V	0.861 V
ERO-FF (PLL)	-1.49	-48.79	-36.38	47.30	0.844 V	0.864 V
RO-FF (PLL)	3.06	-32.44	-22.61	35.50	0.844 V	0.861 V

*Relative to the baseline average obtained before each experiment

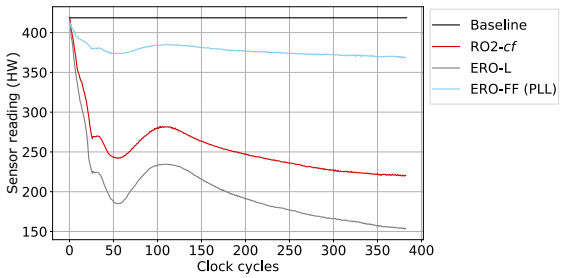

Fig. 11: Comparison of the sensor readings on the Genesys-ZU.

Sensor Readings The results for the Genesys-ZU board are summarized in Table 3. The power wasters' performance follows the same trends as for the Pynq-Z1. The results correspond to the enable signal period of 800 clock cycles, but an attack duration of 384 clock cycles (we have tested the power wasters with various enable signal periods and validated that they remain consistent). We note that the sensor calibration is board-specific and, hence, different (the range for the readings is also different); accordingly, the readings also differ with respect to those on the Pynq-Z1. Fig. 11 shows a sample of the sensor readings for three power wasters on the Genesys-ZU. The main difference between the Pynq-Z1 and the Genesys-ZU is due to the number of latches used in a slice. Therefore, we see latch-based power wasters ranking better in Table 3 than in Table 2.

We also report the voltages collected from the system monitor on the chip. The reported voltage values make use of the fact that the supply voltage and the corresponding voltage regulator for the programmable logic are shared with the processing system and the BRAM. Therefore, Table 3 reports the minimum voltage values across these three components as recorded by Vivado. We choose to do this as the limited sampling frequency of the ADC sometimes leads to the minimum reading on one of the monitored voltages not changing after the ten

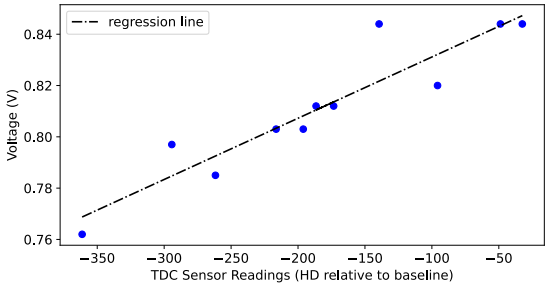

Fig. 12: Minimum TDC sensor readings relative to the baseline vs. minimum voltage readings reported by the system monitor (SYSMON). The figure also shows the regression line to underscore the correlation between the two values.

experimental runs despite the other voltage values changing. The effect of the low sampling frequency of the ADC is apparent in Table 3. While the most effective power waster (ERO-L-*cf*) also results in the lowest voltage, corresponding to a drop of about 10% of the nominal voltage, not all voltage readings are consistent. For example, the ERO-L design has a minimum recorded voltage lower than the ERO-*cf*, despite the TDC sensor readings, which are directly correlated with the voltage [21], showing the opposite trend. However, we validate the correlation between the TDC readings and the voltage and find that the Pearson correlation coefficient is 0.97385. Fig. 12 shows how the minimum voltage and the minimum TDC readings (as Hamming distance (HD) from the baseline) are correlated. In agreement with that, the voltage readings still show that the effect of the top half of the table, i.e., the more effective wasters, is more significant and affects the voltage in a way that the on-chip sensors can measure.

Overall, the results highlight how the expected effect of a power waster can be predicted based on its implementation details. This understanding allowed us to build new power wasters (e.g., ERO-L-*cf* and ERO-L) that outperform existing ones and work with and without combinational loops. The results also show that the designs are portable across FPGA families and that combinational loops are unnecessary for an effective FPGA-based voltage-drop attack.

6 Discussion

On an FPGA, the fine granularity and variety of hardware primitives translate to possible extreme variations in circuit power consumption, as demonstrated by the power-wasting primitives we examined in this work. Consequently, an adversary has a plethora of options to choose from in the function on the desired exploit, the target FPGA platform, and the available resources.

Attack circuits' effectiveness: As expected from Equation (1), the most effective power wasters are those that combine signals switching at a high frequency with a large fanout and a lot of routing. Therefore, all variants that

use combinational feedback loops cause more significant variations in the sensor readings than their counterparts that break the combinational loop. Increasing the routing and the fanout also make a substantial difference, with ERO-*cf* and similar designs performing better than those using only one LUT with the corresponding FFs. The only case where the capacitance effect is not apparent is when the designs are significantly limited by the oscillation frequency, such as those in the last two rows of Table 2. Additionally, due to the limits on the clocks generated by the clock managers within an FPGA, self-oscillating designs tend to perform better than those that use a PLL clock.

Portability: Modern FPGAs share similar organization and resources, facilitating the portability of power-wasting primitives across platforms. LUT-based designs are simple to port, as most FPGAs support similar LUTs with two outputs. On the other hand, ERO-*cf*s may need adjustments for LUT input counts and grouping. Designs using latches and FFs require more effort to be adapted to available FFs and control signal origins. Routing LUT and FF outputs might offer different options for synchronous and combinational outputs. For all designs, the control mechanism will also need to be adapted to the platform. As shown in Section 5, porting designs across FPGA families is feasible, requiring only adjustments to account for the potentially different organization of resources.

Stealth: Oscillating circuits can enable side-channel, fault-injection, and DoS attacks, leading to efforts to detect them. Amazon leverages Vivado to detect the loops to prevent the deployment of ROs [2]. FF-based power wasters, which break the combinational loop, can be deployed on commercial cloud FPGA instances and, as a result, they offer more stealth than their combinational counterparts. Despite the potential for detection [1, 15], the threat remains critical, and studying these exploits is essential for developing future robust protections.

7 Conclusion

Due to their widespread deployment, the security of remotely accessible FPGAs is increasingly critical. This paper explores the potential for implementing power wasters using logic primitives within FPGA programmable fabrics. We classify FPGA-based power-wasting circuits as combinational, FF-based, improved, and hidden. Key comparison features include power consumption, voltage drop, resource efficiency, and the severity of design rule check warnings. We implemented known logic-based power wasters, analyzed them, and proposed new variations. Our analysis validated the factors influencing power consumption, with some proposed designs outperforming the standard ring oscillator. Evaluation on two hardware platforms demonstrated the portability of these designs and validated findings across FPGA families. Our results reveal the potential for new waster designs that evade critical warnings while causing significant voltage drops, posing a threat to remotely accessible systems. Future work can explore hiding wasters within benign designs and developing countermeasures to detect and disable these malicious circuits while preserving legitimate functionality.

Acknowledgments. This research is supported by armasuisse Science and Technology.

References

1. Alrahis, L., Nassar, H., Krautter, J., Gnad, D., Bauer, L., Henkel, J., Tahoori, M.: MaliGNNoma: GNN-based malicious circuit classifier for secure cloud FPGAs. arXiv (Mar 2024), arXiv:2403.01860 [cs]
2. AWS EC2 FPGA HDK+SDK errata (2019), https://github.com/aws/aws-fpga/blob/master/ERRATA.md.
3. FPGA-based Amazon EC2 F1 computing instances (2022), https://aws.amazon.com/ec2/instance-types/f1/
4. 66906 - UltraScale soft error mitigation (SEM) IP - [DRC 23-20] rule violation (PDCN-1569) LUT equation term check (Sep 2021), https://support.xilinx.com/s/article/66906?language=en_US
5. Amer, H.H.: Behavior of memory elements in the presence of power supply disturbances. In: 34th Annual Spring Reliability Symposium, "Reliability - Investing in the Future". pp. 45–51. Boxborough, MA, USA (Apr 1996)
6. Chen, C.H., Bowman, K., Augustine, C., Zhang, Z., Tschanz, J.: Minimum supply voltage for sequential logic circuits in a 22nm technology. In: International Symposium on Low Power Electronics and Design. pp. 181–186. Beijing, China (Sep 2013)
7. García, A.D.G., Pérez, L.F.G., Acuña, R.F.: Power consumption management on FPGAs. In: 15th International Conference on Electronics, Communications and Computers. pp. 240—245 (Feb 2005)
8. Genesys ZU: Zynq UltraScale+ MPSoC development board (2022), https://digilent.com/reference/programmable-logic/genesys-zu/reference-manual
9. Giechaskiel, I., Rasmussen, K.B., Szefer, J.: Measuring long wire leakage with ring oscillators in cloud FPGAs. In: 29th International Conference on Field-Programmable Logic and Applications. pp. 45–50. Barcelona, Spain (Sep 2019)
10. Glamočanin, O., Coulon, L., Regazzoni, F., Stojilović, M.: Are cloud FPGAs really vulnerable to power analysis attacks? In: Design, Automation & Test in Europe Conference & Exhibition (DATE). pp. 1–4. Grenoble, France (Mar 2020)
11. Glamočanin, O., Kostić, A., Kostić, S., Stojilović, M.: Active wire fences for multitenant FPGAs. In: DDECS. pp. 13–20 (May 2023)
12. Gnad, D.R.E., Oboril, F., Tahoori, M.B.: Voltage drop-based fault attacks on FPGAs using valid bitstreams. In: FPL. pp. 1–7. Ghent, Belgium (Sep 2017)
13. Gupta, M.S., Oatley, J.L., Joseph, R., Wei, G.Y., Brooks, D.M.: Understanding voltage variations in chip multiprocessors using a distributed power-delivery network. In: Design, Automation & Test in Europe Conference & Exhibition (DATE). pp. 1–6. Nice, France (Apr 2007)
14. La, T., Pham, K., Powell, J., Koch, D.: Denial-of-Service on FPGA-based cloud infrastructures — attack and defense. IACR Transactions on Cryptographic Hardware and Embedded Systems **2021**(3), 441–464 (Jul 2021)
15. La, T.M., Matas, K., Grunchevski, N., Pham, K.D., Koch, D.: FPGADefender: Malicious self-oscillator scanning for Xilinx UltraScale + FPGAs. ACM Transactions on Reconfigurable Technology and Systems **13**(3), 15:1–15:31 (Sep 2020)
16. Mahmoud, D.G., Andreani, S., Lenders, V., Stojilović, M.: ROBoost: A study of FPGA logic-based power-wasting primitives. Artifacts (Feb 2025). https://doi.org/10.5281/zenodo.14840696

17. Mahmoud, D.G., Dervishi, D., Hussein, S., Lenders, V., Stojilović, M.: DFAulted: Analyzing and exploiting CPU software faults caused by FPGA-driven undervolting attacks. IEEE Access **10**, 134199–216 (Dec 2022)
18. Mahmoud, D.G., Hussein, S., Lenders, V., Stojilović, M.: FPGA-to-CPU undervolting attacks. In: Design, Automation & Test in Europe Conference & Exhibition (DATE). pp. 999–1004 (Mar 2022)
19. Mahmoud, D.G., Shokry, B., Lenders, V., Hu, W., Stojilović, M.: X-Attack 2.0: The risk of power wasters and satisfiability don't-care hardware Trojans to shared cloud FPGAs. IEEE Access **12**, 8983–9011 (Jan 2024)
20. Matas, K., La, T.M., Pham, K.D., Koch, D.: Power-hammering through glitch amplification – attacks and mitigation. In: 28th Symposium on Field-Programmable Custom Computing Machines. pp. 65–69. Fayetteville, AR, USA (May 2020)
21. Moini, S., Deric, A., Li, X., Provelengios, G., Burleson, W., Tessier, R., Holcomb, D.: Voltage sensor implementations for remote power attacks on FPGAs. ACM Trans. Reconfigurable Technol. Syst. **16**(1) (Dec 2022)
22. Pant, S.: Design and Analysis of Power Distribution Networks in VLSI Circuits. Ph.D. thesis, The University of Michigan (2008), https://deepblue.lib.umich.edu/bitstream/handle/2027.42/58508/spant_1.pdf%3Fsequence%3D1
23. Provelengios, G., Holcomb, D., Tessier, R.: Power wasting circuits for cloud FPGA attacks. In: 30th International Conference on Field-Programmable Logic and Applications. pp. 231–35. Gothenburg, Sweden (Aug 2020)
24. Digilent reference for PYNQ-Z1, https://digilent.com/reference/programmable-logic/pynq-z1/start
25. Salman, E., Dasdan, A., Taraporevala, F., Kucukcakar, K., Friedman, E.G.: Exploiting setup–hold-time interdependence in static timing analysis. IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems **26**(6), 1114–25 (Jun 2007)
26. Spielmann, D., Glamočanin, O., Stojilović, M.: RDS: FPGA routing delay sensors for effective remote power analysis attacks. IACR Transactions on Cryptographic Hardware and Embedded Systems **2023**(2), 543–567 (Mar 2023)
27. Stojilović, M., Rasmussen, K., Regazzoni, F., Tahoori, M.B., Tessier, R.: A visionary look at the security of reconfigurable cloud computing. Proceedings of the IEEE **111**(12), 1548–71 (Dec 2023)
28. UltraScale architecture system monitor user guide (Sep 2021)
29. Xilinx Inc.: 7 series FPGAs configurable logic block user guide (UG474) (Sep 2016)
30. Xilinx Inc.: UltraScale architecture configurable logic block user guide (UG574) (Feb 2017)
31. Xilinx Inc.: Zynq-7000 SoC: DC and AC switching characteristics (DS187) (Dec 2020)
32. Zhao, M., Suh, G.E.: FPGA-based remote power side-channel attacks. In: IEEE Symposium on Security and Privacy (SP). pp. 229–244. San Francisco, CA, USA (May 2018)