Characterizing Power Distribution Attacks in Multi-User FPGA Environments

George Provelengios, Daniel Holcomb, and Russell Tessier
Department of Electrical and Computer Engineering
University of Massachusetts, Amherst, MA, USA

***Abstract*—Multi-tenant FPGAs that contain circuits from multiple users are emerging as a new usage model in cloud and embedded computing environments. Interactions between untrusting tenant applications in an FPGA can enable new security exposures and the risk of side channel attacks or fault injection. In this work, we investigate the ability for aggressive power consumption of one application to disturb the power network to an extent that causes delay faults in a second application on the same FPGA. In particular, we identify the mechanisms by which the supply voltage is disturbed by the attack, and we characterize the magnitude of the disturbance as a function of time, power consumed by attacker, and position of the victim relative to the attacker. We highlight strategies that can be used to mitigate attacks, including low-cost monitoring circuits that can identify the source of an attack so that the attacker's use of the FPGA can be revoked.**

I. Introduction

FPGAs are quickly growing in importance in a variety of computing spaces, especially in cloud computing. As FPGAs grow in size and complexity, cloud FPGA deployments aim to leverage economies of scale to share FPGAs among different, untrusting cloud users who wish to accelerate their machine learning, data search, or other applications with FPGAs. In some cases, numerous independent applications may simultaneously reside in a single FPGA. Such uses of multi-tenant FPGAs open the door to numerous potential attack vectors on unsuspecting circuits. At least one integrated approach [1] has been proposed that leverages an operating system to dynamically allocate and simultaneously execute multiple untrusting circuits in a cloud FPGA.

It is well-known that a drop in FPGA supply voltage can cause circuit timing faults [2], [3], [4]. These fluctuations typically occur in the event of a power supply failure or the use of over-aggressive power reduction techniques. However, if multiple tenants share the same FPGA device, one tenant may deliberately and maliciously cause the chip supply voltage to drop in an effort to impact the behavior of a neighbor's circuit. If on-chip logic is used to induce the drop, the attack can be performed remotely, without physical access to the target device. Relative to multi-core chips or GPUs, the flexibility of programmable FPGA logic allows attackers to create arbitrary malicious structures, and thus gives rise to a large and diverse attack surface.

In this work, we attempt to address several questions related to on-chip FPGA voltage attacks that are directly relevant to multi-tenant FPGAs.

- We evaluate the ability of power wasting circuits to exploit properties of the FPGA power distribution network in order to induce delay faults in a neighboring circuit.
- We explore how the effect on a victim circuit depends on disruption time, distance to the voltage disruption circuitry, and the power consumption of the attacker.
- We examine the use of a network of small on-chip voltage sensors to quickly identify impending attacks and mitigate their effectiveness.

DE1-SoC boards, each containing an Intel Cyclone V FPGA, were used to evaluate these effects. In a series of experiments, we characterize voltage drops in the on-chip power distribution network (PDN) caused by an activation of a fraction of the available FPGA logic as power wasters. Our experiments show that voltage drops caused by inductance ($L\frac{di}{dt}$) can be used to create fault attacks that can even target tenants located far from the power wasting area. To address the possibility of power wasting attacks by adversaries, we introduce a monitoring approach using FPGA logic to quickly identify attackers attempting to deploy power wasting circuits.

The remainder of this paper is organized as follows. Background on FPGA voltage attacks and sensors and the threat model for FPGA cloud attacks are described in Section II. Section III describes and analyzes our approach to causing power fluctuations. Section IV examines techniques to cause FPGA faults using voltage fluctuations. Our monitor-based remediation approach is described in Section V, and Section VI concludes the paper.

II. Background and Related Work

A. Multi-Tenant FPGA Threat Model

We consider the following threat model for attacks on the FPGA PDN. Multiple independent users can implement and execute circuits in an FPGA at the same time. Their logic and interconnect resources may be isolated, and each user only has access to the logic design (e.g. bitstream) of their own circuit. There are no physical connections (e.g. wires) shared by the circuits. The software accessed by the designers which interacts with the FPGA is secure as is the interface logic provided in the FPGA. Each user has the flexibility to implement any circuit in their assigned portion of the FPGA.

An example cloud-based system that follows this threat model was outlined in Khawaja *et al.* [1]. This research demonstrated an operating system for shared use of a cloud-based FPGA. The system allows for multiple users to execute

circuits at the same time that may or may not be physically isolated on an FPGA. Device I/O and memory interfaces are fairly shared across users. The PDN in the Xilinx or Intel FPGA is also shared in this model.

B. FPGA Voltage Sensing

One approach to identifying FPGA voltage attacks is to implement distributed voltage sensors fashioned from FPGA logic throughout the logic fabric. The ability to identify voltage levels on an FPGA has many uses ranging for verifying safe FPGA operation [5], [6], [7] to the extraction of secret information [8]. Contemporary FPGAs often contain at least one hardened power supply voltage sensor [9] per chip for power supply voltage measurement. Additional on-chip FPGA voltage measurement circuits typically are based on either ring oscillators or time-to-digital converters (TDCs). A ring oscillator (RO) consists of an asynchronous loop containing an odd number of inverters. The frequency of the oscillation can be measured by connecting the RO to a counter. Although RO frequency is affected by temperature [10], voltage fluctuations have a much stronger effect [11]. TDC-based sensors are based on a combinational chain of buffers that are triggered by a clock edge [7]. The output of each buffer is sampled by a clock-triggered flip flop and voltage values can be determined by how far the initial pulse travels in the chain. Although more difficult to implement effectively, TDCs can be used to measure instantaneous voltage changes on the order of a clock cycle [12]. Given our interest in voltage changes due to attacks, we select a network of simpler but highly-effective ROs for our monitoring system.

Voltage sensors have been used previously to assess voltage changes in the PDN under typical FPGA workloads. Gnad *et al.* [13] examined spatial and temporal voltage effects across an FPGA for a variety of workload characteristics. This work did not consider malicious attempts to waste power by an attacker.

C. FPGA Voltage Attacks

A diverse collection of FPGA voltage attacks that can be exploited in multi-tenant scenarios have been reported. One tenant may try to maliciously induce localized instability in the supply voltage through LUT-based shift registers [14] or deliberate short circuits [15], possibly also exploiting resonances of the power grid. This attack could cause errors in neighboring tenants. Prior work has also shown that a shared FPGA PDN creates coupling between applications. The coupling has been exploited for side channel attacks [8], [16] in which an encryption key is extracted from an unsuspecting victim crypto circuit. Both RO [8] and TDC-based voltage sensors [16] were used successfully for key extraction. Several recent works [2], [3], [4] show that in some cases, RO-based power wasters can be used in an FPGA to cause voltage instability. However, these works do not characterize the nature of the instability, or consider remediation approaches.

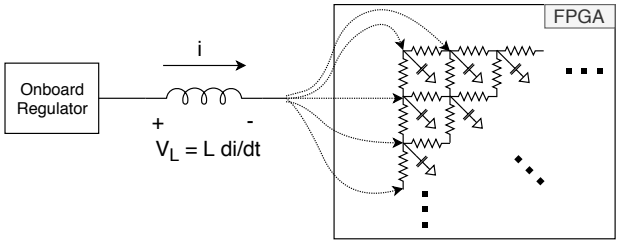

Fig. 1: Schematic of on-chip FPGA power system. A voltage drop occurs across the inductor due to di/dt. A steady-state voltage drop occurs in the PDN due to its resistance.

III. On-chip Attack on an FPGA PDN

The Intel Cyclone V FPGAs (5CSEMA5F31C6) used for this work are located on Terasic DE1-SoC boards [17]. The Cyclone V device does not include on-chip voltage sensors or hardened monitors. Power to the board is from a 12V DC source. The 1.1V internal FPGA core voltage (*VCCINT*) is created by a Linear Technology LTC3608 step-down switching regulator. This switching regulator power supply approach is standard for SRAM-based FPGAs. It supplies power to the FPGA core at 617 kHz through a 1 μH inductor. A schematic of a typical on-chip FPGA PDN is shown in Fig. 1. Although publicly-available information about on-FPGA PDNs is limited, the PDN performance of several SRAM-based Xilinx FPGAs are characterized in Klokotov *et al.* [18]. The basic components of our characterization for instantaneous current changes are similar. After passing through the inductor, the supply voltage is distributed to core voltage inputs of the FPGA. The resistance and capacitance of the PCB traces and on-die PDN network allow localized voltage fluctuations to occur within the chip, such that different parts of the fabric may have different supply voltages at the same time instant [18].

A. Methodology and Calibration

1) On-die Voltage Sensors: A voltage monitoring system is needed to observe the response of the PDN to adversarial power consumption during an attack. Because the Cyclone V device lacks on-chip voltage sensors, we measure the voltage at selected positions of the PDN using ring oscillator-based voltage sensors. The frequency of each oscillator decreases in a consistent way to voltage drops, and a calibration procedure is used to learn the correspondence between voltage and frequency. After calibration, frequency measurements made at each sensor are translated into the voltages that cause them.

Fig. 2 illustrates the architecture of the monitoring system. The sensors are placed on the die forming a regular rectangular grid which in recent related work [8] proved to be sufficient for performing power analysis attacks. Each sensor consists of a 19-stage RO triggering a 20-bit frequency counter. With 19 inverting stages, the design meets the timing constraints, minimizes the effect of local variations [19], and allows for stacking the ROs in a single Cyclone V LAB. Although shorter ROs are also possible by inserting open latches in the ring to

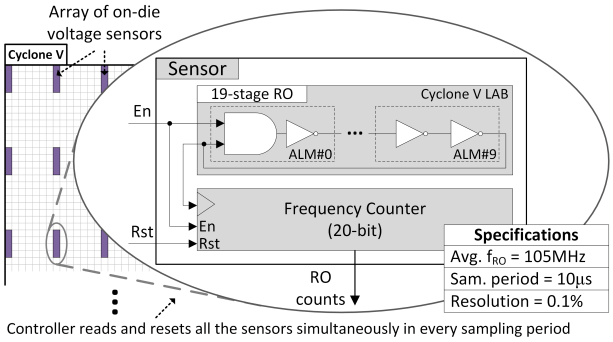

Fig. 2: Schematic of the RO-based voltage sensor.

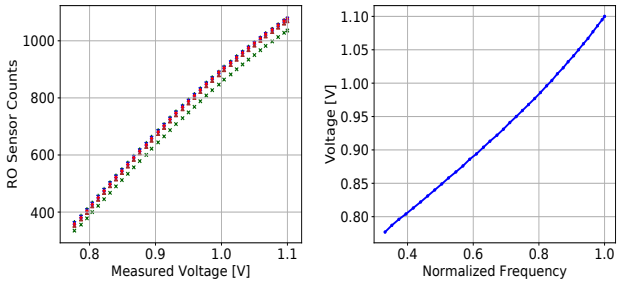

Fig. 3: Figure at left shows oscillation counts from seven different 19-stage ring oscillators when the FPGA voltage is varied. The measurements are used to generate the sensor calibration curve at right, which relates frequency changes to the supply voltage values that account for them. The frequency of a sensor is inversely proportional to the propagation delay of the oscillating signal.

increase the path delay [6], the lack of built-in latch elements in a Cyclone V device makes this technique unsuitable. In our selected Cyclone V device, the 19 inverting stages of the RO design shown in Fig. 2 achieve an average frequency of 105 MHz. We use a 10μs measurement period, which gives us the capability of detecting 0.1% frequency changes, corresponding to a sub-millivolt resolution in supply voltage measurement. We found that the chosen period provides sufficient resolution for performing voltage characterization experiments without complicating the design of the sensor.

To control the voltage when calibrating the sensors, we desoldered the switching regulator and its output inductor from one DE1-SoC board, and supplied the FPGA core voltage to that board directly from a Keysight E36312A benchtop power supply. We varied the supplied voltage, and at each step measured the FPGA input voltage with a Keysight MSOX4154A oscilloscope, and also recorded the frequency of the sensors using test logic on the FPGA. To prevent any localized voltage drops and ensure that the measured voltage matches the voltage at the sensors, only the test logic and sensors are active during calibration, which minimizes the power drawn by the FPGA. Fig. 3 shows the measured correspondence between voltage and frequency of the sensors. The measurements from the RO sensors exhibit a consistent trend across voltages,

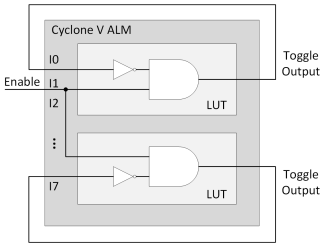

Fig. 4: Power waster circuit.

TABLE I: Power consumed by each power waster instance.

Number Instances	Power / Instance
160	1.13 mW
1600	1.02 mW
3200	0.91 mW
4800	0.84 mW
6400	0.75 mW

and the same trend is observed on all sensors, allowing us to calibrate the relationship between voltage and normalized frequency which is shown at right in Fig. 3.

Although the Cyclone V device and DE1-SoC board do not include a temperature sensor it is expected that voltage gradients will have a much stronger impact on the measured RO delay than temperature. To further minimize heating effects the experiments were conducted using sampling periods at the sub-millisecond range (e.g. 10μs) with no more than a hundred samples being taken each time and between iterations an idle period of a few seconds was introduced. In addition, the ambient temperature during the calibration and characterization experiments was kept at 24°C. Therefore, we anticipate that thermal effects are negligible in our characterization.

2) Adversarial Power Consumption Circuit: We assume that an application on one part of the FPGA is adversarial, and implements a design capable of high power consumption to disturb the PDN. An area of 1,408 LABs (44 rows by 32 columns) was arbitrarily chosen as one representative example of the FPGA real estate an adversary might occupy, which is 32.8% of the total LABs on the chip. Section V considers a second attacker area with a different size. Dynamic power is maximized by circuits with a high amount of switching (see Eq. 1), so we allow the adversary to instantiate various quantity of single-stage ring oscillators as power waster circuits. Fig. 4 shows an ALM implementing two power wasters. Up to 20 power wasters can be implemented in each LAB. When instantiating a desired number of power wasters, a script places them uniformly at random locations throughout the allocated region.

$$p_{dyn} = C * V_{DD}^2 * f_{SW} \quad (1)$$

The power consumed by each instance is given in Tab. I, measured using the modified board and benchtop supply. Note that the power consumed per instance is diminished as the number of instances grows. This occurs because the power wasters cause a local drop in supply voltage which slows down their oscillation (reducing f_{SW} in Eq. 1) and causes the switching to occur at lower voltage (reducing V_{DD}^2). Although our later experiments use up to 12,000 power waster instances, Tab. I ends at 6,400 because the 5A current limit is reached on the benchtop supply used to power the modified board.

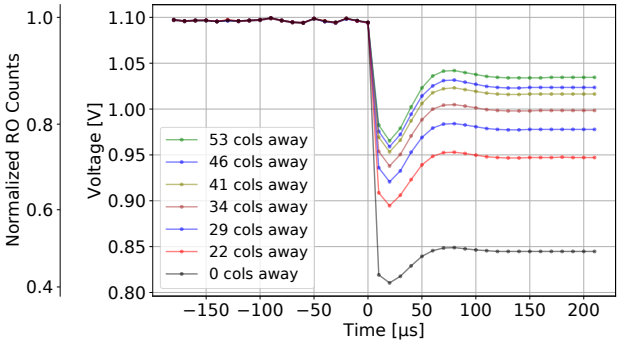

Fig. 5: Normalized RO sensor counts (left axis) and their corresponding voltages (right axis) measured by sensors before and during a power wasting attack that begins at time 0. The legend shows the distance between each sensor and the center of the power wasting region.

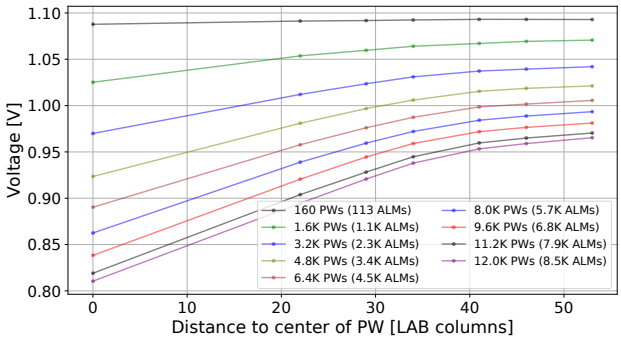

Fig. 6: Voltage change across distance for various number of power wasters.

Note that the modified board is used only to calibrate the sensors and measure consumption of the power wasters (Sections III-A1 and III-A2). All experiments in the remainder of the paper are performed on unmodified DE1-SoC boards with the original switching regulators.

B. Physical Characterization of Voltage Drop

To evaluate the PDN response to high power consumption, an experiment is performed with sensors placed at various distances away from a region with 12,000 power wasters. At time 0 the power wasters turn on and the frequency of the sensors, or equivalently their supply voltages (Fig. 5), drop in response to the attacker's power consumption. The supply voltage measured by each sensor initially drops, undershoots, and then settles back to a steady-state voltage that is lower than the nominal 1.1 V for as long as the power wasters remain active. At the center of the power consumption area, the supply voltage drops to a minimum of 811mV and reaches a steady state of 846mV. Sensors farther away observe a similar behavior but a smaller magnitude of voltage drop.

1) Varying the Amount of Power Consumed: As one might expect, attacks consuming more power will cause larger voltage drops. The voltage drops will be observed at the site of the attack and also in the surrounding area of the chip. Fig. 6 shows voltage plotted against distance from the center of the attack; each line in the figure corresponds to a different number of power wasters being instantiated and used in the attack. We can observe in each attack that the supply voltage change can have a far-reaching impact on other circuitry. Even 53 columns away from the center of attack, the supply voltage is reduced from 1.1V to 967mV in the strongest attack.

2) Role of the Inductor in Undershoot: The voltage undershoot observed in Fig. 5 is caused by the large and sudden change in the current drawn from the FPGA core supply when the power wasters all turn on simultaneously. The sudden change in current creates a voltage drop across the inline inductor of the switching regulator, which thereby reduces the voltage supplied to the chip (Eq. 2). Fig. 7b shows the core voltage dropping when 12,000 power wasters turn on, as captured by a Keysight MSOX4154A oscilloscope. The waveform shows that the peak voltage drop of 85mV occurs roughly 15μs after the power wasters turn on. Integrating the measured inductor voltage with respect to time shows that the current draw is increasing by more than 2.5A within just 60μs.

$$V_{core} = V_{reg} - V_L = V_{reg} - L\frac{di}{dt} \tag{2}$$

The 85mV voltage drop measured across the inductor impacts every part of the FPGA that shares the same supply, which can allow an attacker to affect victim circuits regardless of their position on the chip. Unlike the L$\frac{di}{dt}$ drop, the iR voltage drop due to resistances in the PDN depends only on the current, and not on the change in current. Therefore, L$\frac{di}{dt}$ drop is maximal when the current is changing, and iR drop is maximal after the current has changed, so they do not both contribute their peak values at the same time. The largest total voltage drop is observed to be a combination of L$\frac{di}{dt}$ drop from the inductor combined with a iR drop of the power grid. At the same time that the core voltage is being measured on the board using the oscilloscope (Fig. 7b), sensors are measuring the internal voltage at different locations on the chip. At each sensor location, we extract the minimum voltage reached and the voltage reached in steady-state when the power wasters are active and the current is constant, which is purely an iR drop. Fig. 7a plots these two voltages against the distance between sensor and center of power consumption. In the sensor farthest from the power consumption, the PDN voltage has a relatively small steady state iR drop. Due to the inductor, the minimum voltage reached is 70mV below steady state, which is almost the full 85mV drop observed on the oscilloscope measurement.

IV. Causing Faults via PDN

A decrease in supply voltage causes an increase in propagation delay of combinational logic. Path delay faults will be caused by a reduced supply voltage if the completion time of the combinational results do not satisfy the setup time requirement of the capturing flops. Having shown that aggressive power consumption can cause a far-reaching drop in supply voltage, we now turn to examining whether that voltage

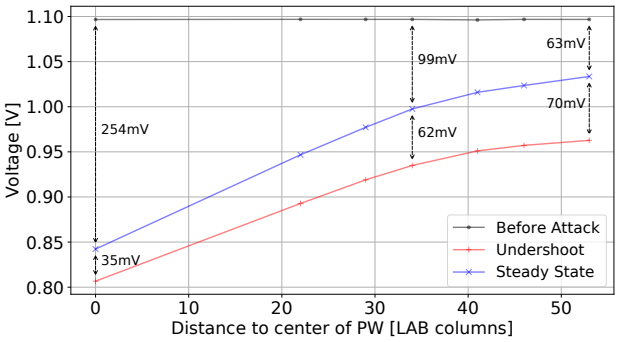

(a) Voltage versus position on chip.

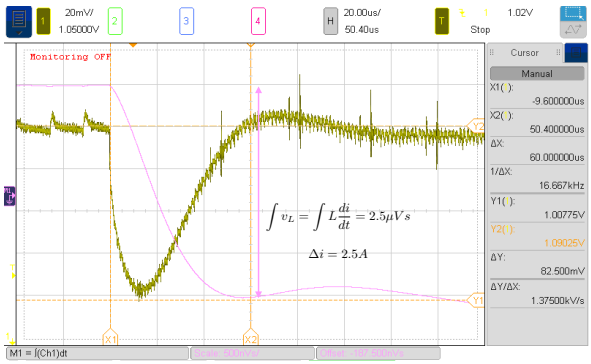

(b) Voltage drop from inductor, measured at FPGA power pin.

Fig. 7: Turning on power waster circuit causes a large instantaneous change in current, and a higher steady state current. The instantaneous change causes a voltage drop on the off-chip inductor, which can be observed at the FPGA power pin and effects all parts of the chip. The high steady-state current additionally causes an IR drop on-chip in the immediate vicinity of the power consumption, with an effect that diminishes moving away from the location of the power consumption.

drop can induce path delay faults in a victim circuit. For simplicity, we use ripple carry adders as test circuits because their carry chain presents different path lengths that can be sensitized by applying appropriate vector pairs.

A. Demonstration of Path Delay Faults

Our first path delay fault experiment has an attacker using 12,000 power wasters within a block of 1,408 LABs, and ripple carry adders instantiated at distances of 22, 26, 30, 35, 38, 42, 47, 50, and 54 columns away from the center of the attack. Vectors are generated to sensitize paths with lengths of 49, 54, 59, 64, 69, and 74 carry stages during the attack. The timing slack of each path in each adder instance is reported using the TimeQuest Timing Analyzer [20] with the slow 1100mV 85°C timing model, and vectors are discarded for any path with a negative slack exceeding 3 ns according to this model. The vectors are repeatedly applied during attack iterations and a log is kept with the times during the attack at which faults occur. Fig. 8 shows the faults that occur from the attack. The x and y coordinates of each point denote the time

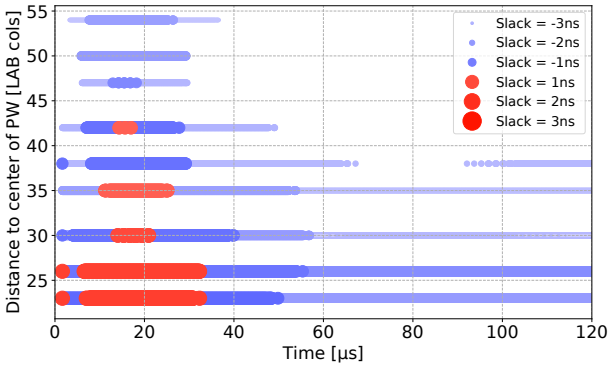

Fig. 8: Delay faults on adder circuits at nine different locations when attacker turns on 12,000 power wasters at time 0. The supply voltage undershoot in the early part of the attack causes the most severe faults. Faults are observed on legal paths with positive slack that are 42 columns away from the center of the attacker.

and location of a fault, and the size of the point corresponds to the reported slack of the path on which the fault occurred. Larger points denote paths with more slack, which are less susceptible to delay faults. Every point on the plot depicts an incorrect result being captured. Red points are from paths with positive slack, which are timing faults in legal paths that meet timing constraints even under the conservative timing model used. Blue are from paths that have negative slack according to the conservative timing model but are always error free in the absence of an attack. These paths help to show the overall trend of faults.

Faults are induced on legal paths only during the first 40μs, when the undershoot causes the lowest voltages to occur (see Fig. 5). Given that the attack causes faults on legal paths with positive timing slack that are 42 columns away from the center of the attack, it is reasonable to conclude that spatial isolation between tenants may be insufficient for protecting against PDN attacks in multi-tenant FPGA applications.

B. Relating Voltage and Timing Slack to Fault Sensitivity

Having demonstrated the capability to cause delay faults, and characterizing PDN voltage in response to power consumption, we now connect the two and show experimentally the combinations of slack and voltage that lead to faults. In this experiment, 1,024 random attack scenarios are created and implemented. In each attack scenario the following parameters are chosen at random from the stated ranges:

- The position of the victim adder circuit (between 23 and 53 columns from center of attacker).
- The sensitized path of the victim adder (uses between 53 and 64 stages of carry logic implemented on the hardened carry circuitry of the FPGA).
- The number of power wasters used by the attacker (between 3,200 and 12,000 instances).

The minimum voltage at the victim circuit during each attack is inferred by interpolation on the data shown in Fig. 6

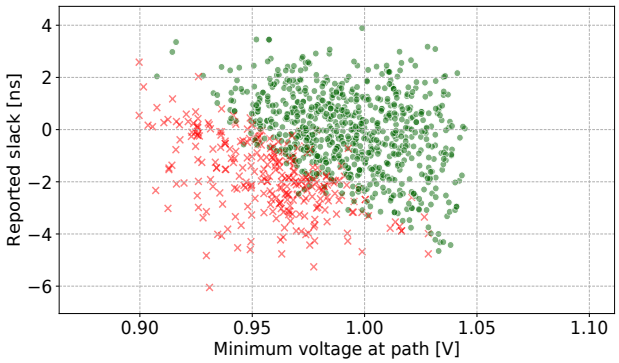

Fig. 9: Scatter plot shows which randomly generated attack scenarios caused faults and which did not. X-coordinate denotes voltage in victim circuit during attack. Y-coordinate is the reported timing slack of path exercised during attack.

according to the victim location and number of power wasters in the attack. As in the prior subsection, the path is repeatedly sensitized during the attack and the result is checked for faults. Red and green marks in Fig. 9 denote attack scenarios in which faults did or did not occur, respectively. The x-coordinate of each point is the minimum voltage at the victim during the attack. The y-coordinate of each point is the timing slack of the victim path as reported by TimeQuest Timing Analyzer using the slow 1100mV 85°C timing model.

Timing models are conservative with respect to operating conditions and process variation, and the effects of the conservative timing model can be seen in Fig. 9. Paths reported as having 0 slack are typically fault free even when their voltage drops by 140mV, although Fig. 3 shows that a 140mV drop should cause a significant increase in propagation delay.

The pattern of faulty and fault-free points in Fig. 9 shows that the combination of voltage and timing slack are largely sufficient to explain which adder paths will experience faults during an attack. This finding supports the supply voltage drop being the cause of the fault, and not some other artifact of power consumption. The results also show that conservative timing models provide some inherent margin against attack.

V. Monitoring System for PDN Attacks

PDN attacks require power consumption, transiently or in steady-state, beyond what the power distribution network can handle. Our results have shown that the power consumption of one adversarial block can cause a measurable and significant difference in the voltage of other blocks. Circuits closest to the power consumption experience the largest voltage drop, and the voltage drop becomes smaller moving farther away (Fig. 6). The voltage gradients effectively provide a map pointing toward the center of the attack, which will have the lowest voltage. A spatially distributed network of voltage sensors can enable a resource manager to monitor voltage gradients and identify the source of any attacks that occur. The resource manager can then prevent further instances of the same attack by taking the offending application offline, or banning it from co-tenant settings.

TABLE II: Resources used in voltage monitoring network for various numbers of sensors.

Num. RO sensors	ALMs (Avail.: 32,070)	Flip flops (Avail.: 128,280)
10	390 (1.2%)	200 (<1%)
20	780 (2.4%)	400 (<1%)
30	1,170 (3.6%)	600 (<1%)
40	1,560 (4.9%)	800 (<1%)
46	1,794 (5.6%)	920 (<1%)
Controller	430 (1.3%)	111 (<1%)

* Device: Intel Cyclone V 5CSEMA5F31C6 FPGA

A. Monitor Network

A network of 46 sensors monitor voltage fluctuations and log the data for processing. The area cost of the monitor network is given in Tab. II. Each sensor uses 39 ALMs and 20 flip-flops. The 46 sensors collectively consume 5.6% of the ALMs and less than 1% of the flip flops on the chip. The controller logic that logs the sensor data to memory is only synthesized for the full 46 sensor network as shown in the table. The controller for a network with fewer than 46 sensors would consume less resources.

Fig. 10 shows the voltage contours of the chip based on sensor data during two different power attacks. The specific data used to generate the plot is the minimum value observed by each sensor in the 500μs time period that contained the attack. A cubic interpolation algorithm reconstructs the smoothed voltage contours from the samples collected at the discrete sensor locations.

In the first attack (Fig. 10a), the attacker turns on 12,000 power wasters within an area spanning 44 columns and 32 rows shown in purple in the figure. As denoted on the voltage contour lines, the voltage at the center of the attack drops below 825mV, and the voltage at the farthest corner of the FPGA drops to 975mV.

In the second attack (Fig. 10b), the attacker turns on 3,200 power wasters within an area spanning 20 columns and 20 rows at the bottom right corner of the chip. The voltage close to the attack drops below 975mV, and the voltage at the farthest corner of the FPGA remains above 1.050V.

B. Attack Attribution

A goal for the monitoring network is to determine the source of any attacks that occur. In the case of PDN attacks, the attacker cannot easily mask their identity, because of the spatial extent of the voltage drops that they cause. Here we evaluate the number of sensors required to find an attacker based on voltage contours. For each attack scenario from the previous subsection, we consider how precisely the attacker can be located using 10, 20, 30, or 40 of the 46 sensors (Tab. II), which would reduce the cost of the monitoring network. For each number of sensors, we randomly choose 100 different subsets containing that number of sensors, and from each subset try to predict the location of the attacker.

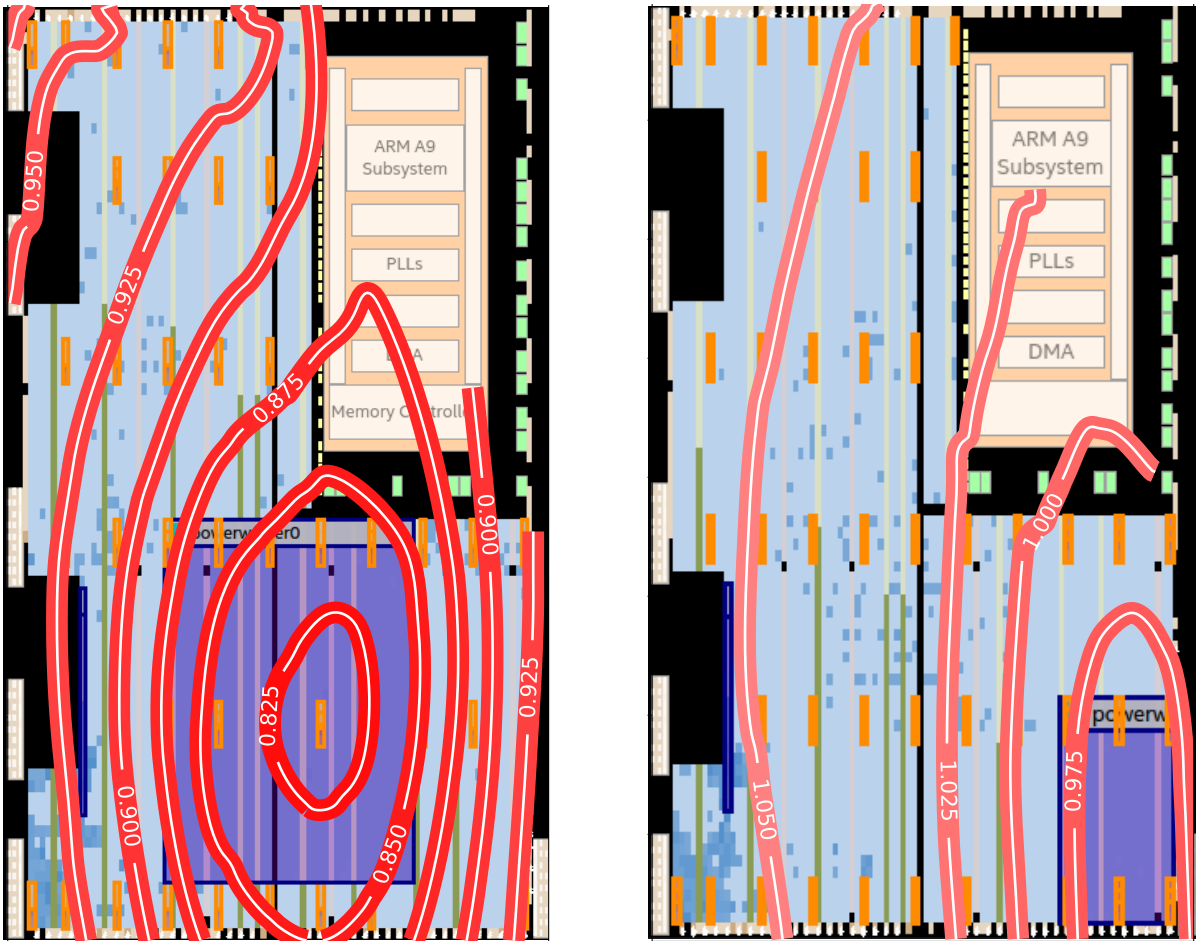

(a) 12,000 power waster attack.

(b) 3,200 power waster attack.

Fig. 10: Map of voltage contours on chip during power attacks, reconstructed from sensor data. Purple rectangle denotes location of the attacker's power waster circuits. Orange rectangles are the sensors.

Fig. 11 shows the results of this analysis. The dots on each plot are the 100 different predictions of the attacker location. As one might expect, the chance of successfully locating the attack increases with the number of sensors. The predictions based on 10 sensors are imprecise, but when 20 or more sensors are used, the predictions converge to a location within the attacking circuit. These results show that a network of monitors can be used to find the attacker with less than 46 sensors, which can reduce the cost of the monitor network. The overall low hardware overhead of the monitoring system should not interfere with the design of other circuits.

VI. Conclusion

The recent emergence of FPGAs in the cloud has made the possibility of multi-tenant FPGAs used by multiple independent circuit designers a reality, as evidenced by a recent paper from industry [1]. Xilinx and Intel FPGAs have been already deployed to Amazon EC2 F1 [21] and Microsoft Catapult [22] systems, respectively. In this paper, we show that multiple power-waster circuits implemented in one portion of an FPGA can induce faults in other independent parts. These effects are carefully characterized for an Intel Cyclone V FPGA located on a Terasic DE1-SoC board. Specifically, we characterize the magnitude of the disturbance as a function of time, power consumed by attacker, and position of the victim relative to the attacker. For mitigation, we propose the use of a series of small voltage sensors that collect voltage information in real time and pass it to a central controller. We demonstrate that the source of a voltage-altering attack can be easily identified by a small number of sensors consuming less than 5% of FPGA logic. Upon identification, the attacker's FPGA privileges can be revoked. Future work will include experiments with voltage monitoring networks on an existing cloud FPGA platform to identify how quickly an attacker can be identified and an attack prevented. Other forms of power wasters that do not depend on asynchronous ROs and TDCs will also be explored.

Acknowledgement

This research was funded by NSF/SRC grant CNS-1619558 and a grant from Intel's Corporate Research Council.

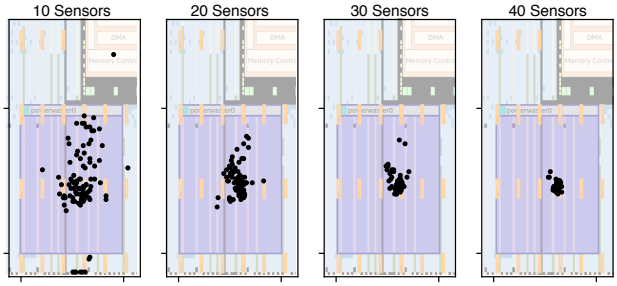

(a) Locating attack that uses 12,000 power wasters.

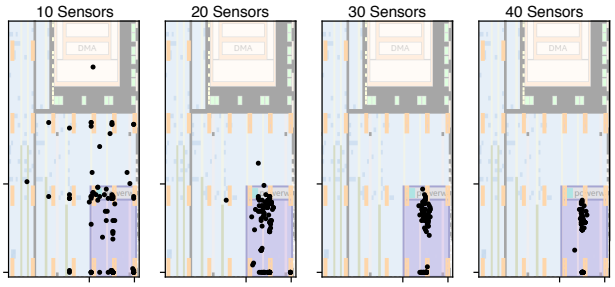

(b) Locating attack that uses 3,200 power wasters.

Fig. 11: Marks represent predicted center of attack based on a randomly selected subset of sensors. Each subplot contains 100 points.

REFERENCES

[1] A. Khawaja, J. Landgraf, R. Prakash, M. Wei, E. Schkufza, and C. J. Rossbach, "Sharing, protection, and compatibility for reconfigurable fabric with AmorphOS," in *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, Carlsbad, CA, 2018, pp. 107–127.

[2] J. Krautter, D. Gnad, and M. Tahoori, "FPGAhammer: Remote voltage fault attacks on shared FPGAs, suitable for DFA on AES," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, vol. 2018, no. 3, pp. 44–68, Aug. 2018.

[3] D. R. Gnad, F. Oboril, and M. B. Tahoori, "Voltage drop-based fault attacks on FPGAs using valid bitstreams," in *2017 27th International Conference on Field Programmable Logic and Applications (FPL)*. IEEE, 2017, pp. 1–7.

[4] D. Mahmoud and M. Stojilović, "Timing violation induced faults in multi-tenant FPGAs," in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2019, pp. 1745–1750.

[5] K. M. Zick and J. P. Hayes, "On-line sensing for healthier FPGA systems," in *Proceedings of the 18th annual ACM/SIGDA International Symposium on Field programmable gate arrays (FPGA)*. ACM, 2010, pp. 239–248.

[6] ——, "Low-cost sensing with ring oscillator arrays for healthier reconfigurable systems," *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, vol. 5, no. 1, p. 1, 2012.

[7] K. M. Zick, M. Srivastav, W. Zhang, and M. French, "Sensing nanosecond-scale voltage attacks and natural transients in FPGAs," in *Proceedings of the ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA)*. ACM, Feb. 2013, pp. 101–104.

[8] M. Zhao and G. E. Suh, "FPGA-based remote power side-channel attacks," in *IEEE Symp. Security and Privacy*, May 2018, pp. 805–820.

[9] *Intel FPGA Voltage Sensor IP Core User Guide*, Intel Corporation, Feb. 2018.

[10] M. Barbareschi, G. Di Natale, and L. Torres, "Implementation and analysis of ring oscillator circuits on Xilinx FPGAs," in *Hardware Security and Trust*. Springer, 2017, pp. 237–251.

[11] D. R. E. Gnad, F. Oboril, S. Kiamehr, and M. B. Tahoori, "Analysis of transient voltage fluctuations in FPGAs," in *IEEE International Conference on Field Programmable Technology*, Dec. 2016, pp. 12–19.

[12] M. Ueno, M. Hashimoto, and T. Onoye, "Real-time supply voltage sensor for detecting/debugging electrical timing failures," in *2013 IEEE 27th International Symposium on Parallel and Distributed Processing Workshops and PhD Forum*, 2013, pp. 301–305.

[13] D. R. E. Gnad, F. Oboril, S. Kiamehr, and M. B. Tahoori, "An experimental evaluation and analysis of transient voltage fluctuations in FPGAs," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 26, no. 10, pp. 1817–1830, Oct 2018.

[14] D. Ziener, F. Baueregger, and J. Teich, "Using the power side channel of FPGAs for communication," in *IEEE International Symposium on Field-Programmable Custom Computing Machines*, May 2011, pp. 1–8.

[15] D. C. Savory, "Power side-channel DAC implementations for Xilinx FPGAs," Master's thesis, Dept. of Electrical and Computer Engineering, Brigham Young University, Apr. 2012.

[16] F. Schellenberg, D. R. Gnad, A. Moradi, and M. B. Tahoori, "An inside job: Remote power analysis attacks on FPGAs," in *Design, Automation & Test in Europe Conference & Exhibition, DATE 2018*, 2018.

[17] *DE1-SOC User's Manual*, Terasic Corporation, Feb. 2014.

[18] D. Klokotov, J. Shi, and Y. Wang, "Distributed modeling and characterization of on-chip/system level PDN and jitter impact," in *DesignCON*, 2014, pp. 1–22.

[19] T. Takahashi, T. Uezono, M. Shintani, K. Masu, and T. Sato, "On-die parameter extraction from path-delay measurements," in *2009 IEEE Asian Solid-State Circuits Conference*. IEEE, 2009, pp. 101–104.

[20] *TimeQuest Timing Analyzer Quick Start Tutorial*, Intel Corporation, Dec. 2009.

[21] "Amazon F1 web site," https://aws.amazon.com/ec2/instance-types/f1/.

[22] A. Putnam et al., "A reconfigurable fabric for accelerating large-scale datacenter services," *ACM SIGARCH Computer Architecture News*, vol. 42, no. 3, pp. 13–24, 2014.