arXiv:2303.17881v1 [cs.CR] 31 Mar 2023

Pentimento: Data Remanence in Cloud FPGAs

Colin Drewes
Stanford University
University of California San Diego
drewes@cs.stanford.edu

Olivia Weng
University of California San Diego
oweng@ucsd.edu

Andres Meza
University of California San Diego
anmeza@ucsd.edu

Alric Althoff
ARM
alric.althoff@arm.com

David Kohlbrenner
University of Washington
dkohlbre@cs.washington.edu

Ryan Kastner
University of California San Diego
kastner@ucsd.edu

Dustin Richmond
University of California, Santa Cruz
drichmond@ucsc.edu

***Abstract*—Cloud FPGAs strike an alluring balance between computational efficiency, energy efficiency, and cost. It is the flexibility of the FPGA architecture that enables these benefits, but that very same flexibility that exposes new security vulnerabilities. We show that a remote attacker can recover "FPGA pentimenti" – long-removed secret data belonging to a prior user of a cloud FPGA. The sensitive data constituting an FPGA pentimento is an analog imprint from bias temperature instability (BTI) effects on the underlying transistors. We demonstrate how this slight degradation can be measured using a time-to-digital (TDC) converter when an adversary programs one into the target cloud FPGA.**

This technique allows an attacker to ascertain previously safe information on cloud FPGAs, even after it is no longer explicitly present. Notably, it can allow an attacker who knows a non-secret "skeleton" (the physical structure, but not the contents) of the victim's design to (1) extract proprietary details from an encrypted FPGA design image available on the AWS marketplace and (2) recover data loaded at runtime by a previous user of a cloud FPGA using a known design. Our experiments show that BTI degradation (burn-in) and recovery are measurable and constitute a security threat to commercial cloud FPGAs.

1. Introduction

Amazon, Microsoft, Alibaba, Baidu, Huawei, TenCent, and Nimbix offer FPGAs as an on-demand cloud service. FPGAs efficiently accelerate common cloud applications including neural networks [18], video transcoding [1], genome sequencing [14], secure database transactions [5], networking [49], and homomorphic encryption [48].

Unfortunately, cloud FPGAs open the door to new security vulnerabilities related to confidentiality [19], [21], [55], [56], [69], integrity [11], [31], [37], [40], [50], and availability [24], [43]. Signal timing sensors [23], [70] have been used to extract cryptographic keys of active computation within the FPGA [55], identify the active computation running within the FPGA [25], implement covert channels across dies on a 2.5D integrated package [19], and perform attacks across chips on the same board [20], [56].

These attacks require that the attacker and victim are spatiotemporally co-located on the same system. For this reason, cloud FPGAs are often only temporally shared; they do not allow multiple users to co-exist spatially on the same FPGA. Conventional wisdom says information leakage will not occur if the FPGA is correctly erased after use. Thus, after a user relinquishes the cloud FPGA, the FPGA is wiped [7], [36], and at some point, rented to another user.

The attacks presented in this work exploit a side channel that allows an attacker to target previous users of the FPGA even after a wiping procedure is performed. The victim is no longer performing computation or renting the FPGA; the victim no longer resides on the device, *and the victim has left no logical information on the device.*

We show that data from previous users can be extracted via an analog side channel due to bias temperature instability (BTI) aka *burn-in*. We call these *"FPGA pentimenti"* – the analog residue of the previous digital data that remains on the FPGA due to BTI effects. FPGA pentimenti are recoverable by sensing BTI recovery using time-to-digital (TDC) sensors. Our experiments show that FPGA pentimenti are a real and extant threat to the security of cloud FPGAs. Much like infrared imaging can expose art work pentimenti (previous paint strokes since painted over by an artist) not visible to the naked eye, we demonstrate that attackers can exploit *FPGA pentimenti* (previous design and user data digitally wiped by the cloud provider).

BTI physically deteriorates transistors, thus negatively affecting their propagation delay. The BTI effect is caused by applying positive/negative (1/0) voltages to CMOS transistors. *BTI recovery* occurs when the transistors are no longer stressed; the transistors partially revert to their previ-

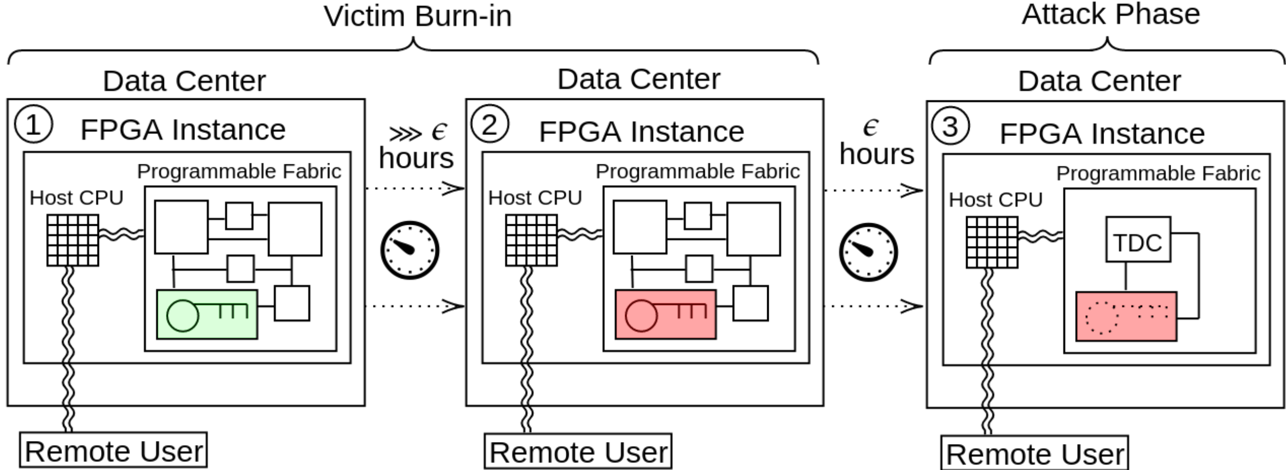

Figure 1: **Pentimenti FPGA Threat Models:** ① A design containing confidential information (green key) is loaded onto the FPGA. ② After this design runs for some number of hours, parts of the design are imprinted – a pentimento (red key) is left on the FPGA due to analog remanence from BTI effects. ③ The attacker loads their design with a BTI sensor to extract the pentimento based on the recovery timing effects.

ously faster state. Transistors undergo negative and positive BTI caused by applying logical 0 and 1 values, respectively. NBTI and PBTI degradation is not symmetric. NBTI effects are typically larger than PBTI. NBTI recovery is also faster. By measuring the speed and size of the recovery, an attacker can deduce if a previous value as a 1 or 0.

We show how to measure BTI effects on cloud FPGAs to exploit an analog temporal side channel that leaks data between successive users of a cloud FPGA. BTI is measured using a TDC sensor that records the time to propagate a pulse through FPGA resources. The change in propagation delay reflects the previous values held by those resources due to BTI effects. We show how this opens a side channel that can be exploited by an attacker to ascertain previous design and user data.

Our work describes, for the first time, how BTI effects are a security concern for commercial cloud FPGAs. We demonstrate the ability to recover pentimenti in remote FPGAs to expose two violations of the cloud FPGA security model when the "skeleton" (the physical structure, but not the contents) of the design is known to the attacker: An attacker can (1) extract proprietary details or keys from an encrypted bitstream accessible via the cloud platform (i.e., the AWS marketplace) and (2) recover non-transient runtime data from a previous user of a cloud FPGA device by observing the BTI recovery via circuit timing changes.

We experimentally validate the burn-in threat on a AWS F1 (Virtex UltraScale+) cloud FPGA and on a local ZCU102 (Zynq UltraScale+) FPGA. In both cases, we demonstrate a discernible difference of the burn-in behavior on an FPGA route before and after BTI degradation. The timing behavior of that route is dependent on what the previous value was, and thus an attacker is able to carry out attack (1) and (2) on the AWS F1 platform as detailed in Section 2.

We contextualize these findings on the OpenTitan hardware root of trust – an open-source hardware design with strict data security requirements [44]. Roots-of-trust carry out core, security-critical functionalities such as secure boot, configuration of operative modes (e.g., debug vs normal), and management of sensitive of data (e.g., cryptographic keys). OpenTitan security assets are vulnerable to the burn-in threat.

Section 2 presents the cloud FPGA threat model. Section 3 provides background on the effects of BTI transistor degradation, which are measured by the sensor presented in Section 4. We construct an experiment to verify the burn-in threat model generally on an Ultrascale+ FPGA and specifically on the AWS F1 platform in Section 5, which is carried out in Section 6. We relate this paper to prior efforts in Section 7, discuss mitigations in Section 8, and conclude in Section 9. **Disclosure:** The necessary steps have been taken to alert affected vendors. Amazon Web Services was originally notified July 2022. Xilinx was originally notified in August 2022.

2. Threat Models

The threat models extract side channel information about previous cloud FPGA user data via temporal analog residue, aka "pentimenti", that arise from BTI recovery effects. Our discussion is framed in the context of the AWS F1 platform though it applies to other cloud FPGA platforms. AWS enables customers to share/sell preexisting designs to other AWS users through the AWS marketplace. AWS provides these designs as an Amazon Machine Image (AMI) and Amazon FPGA Image (AFI). The AFI provides the FPGA bitstream. The AMI is a Linux machine that interfaces with the AFI.

Figure 1 shows the general approach of our threat models. ① A user rents and loads a design containing confidential information (denoted by the green key). ② The design remains programmed on the FPGA and computes for some number of hours, allowing the user data to experience BTI effects and burn-in (red key). The victim FPGA is

released back into the rental pool. It undergoes a design wipe performed by AWS to reset the system and clear out any data remanence [7], [36]. ③ The attacker gains access to the FPGA and loads the TDC sensor to extract the pentimenti – the analog residues of the previous digital data that remains on the FPGA due to BTI effects.

With this setup, we can extract two types of previously safe data using the techniques presented in this paper: **Type A** design data and **Type B** user data.

Type A (Design Data): FPGA designs often contain confidential information as netlist constants, e.g., cryptographic keys or machine learning weights. The AFI promises to keep such proprietary design information secret. A purchased AFI does not permit the user access to the FPGA source code or bitstream[1] to preserve intellectual property rights. But this sensitive information can be extracted via their pentimenti as we show in this paper. We call this secret information baked into the design **Type A** data, and the victim is the AFI publisher.

Type B (User Data): Type B data is from a previous user of the FPGA. The previous user loads confidential information onto an AFI at runtime. Since the attacker does not control the loading and unloading of the design, an attack cannot rely on gathering initial delay estimates (as can be done for Type A data). Thus, extracting Type B user data is more challenging but powerful attack that requires measuring BTI recovery.

*The difference between **Type A** and **Type B** is subtle but shifts the target of an attack from being the publisher of a design/AFI (**Type A**) to the user of design/AFI (**Type B**).* The threat models differ, but both follow the steps depicted in Figure 1.

Threat Model 1 - Proprietary Design Data Extraction: Threat Model 1 targets Type A Design Data encoded into the design itself, e.g., a netlist constant holding a cryptographic key or machine learning weight. The attacker is renting the design, satisfies Assumptions 1 and 2 (discussed later), and can control the loading and unloading of the design. AWS guarantees to keep design intellectual property secret [7], thus Threat Model 1 violates AWS F1 security guarantees.

The attacker extracts proprietary design information via the following:

1) A malicious AWS F1 user rents an FPGA instance with the intention extracting sensitive information from a third party design.
2) The attacker measures the routes that will hold the sensitive data and gather pre-burn-in route delay characteristics.
3) The attacker loads a target design in Stage ① of Figure 1 that contains sensitive information stored in the FPGA routes.
4) The attacker executes the target design until Stage ② of Figure 1 when the BTI effects burn-in the FPGA routes holding sensitive information.
5) The attacker initiates the attack phase (Stage ③ of Figure 1). They unload the victim design and load a measure design that contains the TDC sensor from Section 4 to measure the BTI degradation of the victim routes via their timing behavior.
6) The attacker analyzes sensor data to determine the sensitive information from the victim design with high probability.

Threat Model 2 - Confidential User Data Extraction: The attacker recovers confidential data from a previous victim tenant of the cloud FPGA (Type B). This model assumes that the attacker can requisition an FPGA after the victim has finished computing. The attacker extracts confidential user data via the following:

1) A non-malicious AWS F1 victim user loads a design in Stage ① of Figure 1. This design contains sensitive information either stored statically in FPGA bitstream (e.g., a netlist constant) or in data loaded at runtime.
2) The victim design executes, during which the sensitive data is statically held in the FPGA resources. After some time, the victim design has induced the burn-in effect (Stage ② of Figure 1).
3) The victim completes their computation and relinquishes the FPGA back into AWS's pool of available devices.
4) The attacker instantiates an AWS instance and is assigned the relinquished victim device.
5) The attacker loads in Stage ③ of Figure 1 an FPGA design that contains TDC sensors connected to victim resources that previously held sensitive information.
6) The attacker analyzes TDC sensor results to determine the sensitive victim data with high probability.

The difference between these two threat models shifts the attack target from the producers of the design IP (**Threat Model 1**) to a previous user the cloud FPGA (**Threat Model 2**). Both threat models are a fundamental violation of the AWS FPGA F1 security guarantees. AWS guarantees that "no FPGA internal design code is exposed" [7] through an AFI leased from the marketplace, meaning **Threat Model 1** should not occur. Furthermore, AWS states that they scrub "FPGA state on termination of an F1 instance," [7] meaning **Threat Model 2** leakage should not occur. Our results demonstrate the feasibility of these threat models, which show that burn-in is recoverable using a TDC sensor.

Our threat models rely on two assumptions. Assumption 1 covers Type A data and Threat Model 1. Type B data and Threat Model 2 make an additional assumption.

Assumption 1: The attacker knows the placement, or "skeleton", of the targeted design routes[2] that contains confidential design information (Type A) or sensitive user data (Type B).

The attacker's knowledge of the sensitive information's location could be derived from a publicly available design or bitstream. For example, the OpenTitan hardware root of trust distributes a prebuilt bitstream that a user loads with sensitive information like cryptographic keys [45]. Xilinx

1. The bitstream is the binary file used to program the FPGA.

2. The wire segments inside the FPGA holding the targeted data.

FINN provides prebuilt bitstreams for different neural network architectures [65]. In both cases, the complete source code and compilation scripts are available, which allows one to determine the locations of the sensitive data – the keys for OpenTitan and the neural network weights for FINN.

Other options to learn the target route places include: 1) the attacker is the original author of the AFI on the AWS marketplace and knows design route details. 2) Proprietary information about the design layout has been leaked to an attacker. Finally, when evaluating an implementation's security, it is common practice to assume the architecture is publicly visible [46].

Thus, we believe it is reasonable to assume that the attacker knows the placement information (Assumption 1). Loosen or removing this assumption would strengthen the threat model, and we are considering ways to expand the threat model without Assumption 1 in future work.

Assumption 2: The attacker has the ability to gain access to the same FPGA the victim relinquished. Gaining access to a relinquished cloud FPGA requires aspects of cloud cartography and co-location attacks [6], [30], [54], [66], [68] that check out devices en masse or leveraging cloud FPGA fingerprinting techniques [59]–[62]. Another potential option is a flash attack where the attacker locks up the available stock right before the victim releases their instance. If the attacker procures all the available resources, they are guaranteed to obtain the relinquished victim board. In our AWS experimentation we commonly received errors implying that we have reached the limit of F1 devices in the region, suggesting that this flash attack could be accomplished through acquiring only a few devices.

3. Bias Temperature Instability (BTI)

Bias temperature instability (BTI) is a degradation behavior of transistor transconductance, subthreshold slope, and linear and saturation drain current fundamental to modern field-effect transistors [38], [39]. *Negative BTI (NBTI)* occurs when the PMOS transistor gate voltage is negative relative to its other terminals (0/`False` logical value), which results in positive charge migration into the silicon dioxide insulation. *Positive BTI (PBTI)* affects NMOS transistors when its gate voltage is positive relative to the other terminals (1/`True` logical value), resulting in negative charge migration into the insulating dielectric. BTI effects accumulate under voltage stress, increasing the threshold voltage of their respective transistor types, and consequently increased transistor rise and fall transition delays [41].

CMOS logic gates (e.g., NAND, NOR, and INV) are built from PMOS and NMOS transistors. Static 0/1 inputs on logic gates cause NBTI/PBTI degradation on PMOS/NMOS transistors. The degradation affects the rising ($0 \rightarrow 1$) and falling ($1 \rightarrow 0$) propagation delays through the logic gates. Figure 2 details how a CMOS inverter undergoes data-dependent BTI effects which cause timing deviations captured by the difference in rising and falling propagation delays through the inverter.

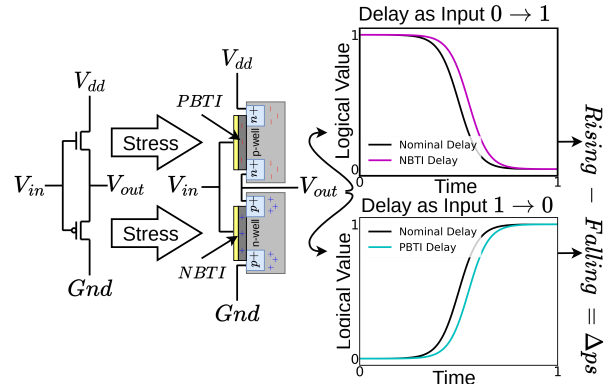

Figure 2: Bias temperature instability (BTI) effects induce circuit delays that differ depending on the data computed. An inverter is pictured on the left, composed of a PMOS (bottom) and NMOS (top) transistor. A V_{in} value of 0 (1) will allow current to flow through the PMOS (NMOS) transistor, causing it to degrade through NBTI (PBTI). As NBTI (PBTI) manifests, the speed in which the inverter stabilizes on an output, V_{out}, as V_{in} equals 0 (1) will begin to slow. Δps is the difference between the 0-input propagation speed (measuring PBTI) and the 1-input propagation speed (measuring NBTI). Δps will vary depending on whether the inverter was previously computing a 0 or 1 value; thus, it can be used to infer the values of a previous computation.

When BTI-causing values are removed, there is a partial threshold voltage recovery [8], [13], [51]–[53] that increases the transistor switching speed. NBTI and PBTI recovery differs in mechanism and recovery rate [27], [28], [35]. NBTI recovery is due to defect removal via the recuperation of broken bonds with positively-charged hydrogen atoms [33], [51]; PBTI recovery is due to the removal of trapped negatively-charged electrons in the transistor dielectric [39]. The defects from PBTI electron charge traps are energetically deeper than NBTI positive charge traps [67] and affect the recovery timescale of PBTI relative to NBTI. BTI recovery effects are measurable, the differences between positive and negative recovery characteristics are apparent, and we can exploit them as a temporal side channel.

BTI creates pentimenti – analog remenance of previous design state and data. BTI effects are differentiable; they encode data about the prior state, e.g., if the route under test was previous a 0 or 1 value. The rate and degree of BTI effects are driven by constant voltages and dynamic switching [58], and has a predictable effect. Remenance is an essential condition to recover any logical information and necessary for Threat Models 1 and 2. BTI effects are not-permanent. BTI degradation is not a permanent artifact and undergoes a recovery period where the transistors become faster [8], [13], [51]–[53]. This required for Threat Model 2 since it requires measuring remanance of previous user values that are no longer present.

FPGAs contain many resources that undergo BTI and can be targeted in pentimenti attacks: bitstream configura-

tion bits, programmable routing, configurable logic blocks (CLBs), digital signal processors (DSPs), and block RAMs (BRAMs). In order to perform a successful attack, the victim resource should meet the following conditions:

- **BTI effects must occur:** The target resource must be used in a manner that induces burn-in (BTI degradation) – a necessary condition to recover Type A and B data. For example, a route is statically held at constant value.
- **BTI effects must be differentiable:** The target resource should exhibit differences in circuit-level behavior due to BTI degradation and recovery. For example, the route delay profile differs based on whether it was previously held at logical 0 or 1.
- **BTI-effected resources must be observable:** Targeted resources must be in user-visible locations on AWS F1. Some cloud FPGA resources are inaccessible by the user, e.g., resources implementing the AWS shell. The attacker is limited by the interfaces exposed by the cloud provider. The attacker do not have physical access. They cannot use special sensing instrumentation. The BTI sensor must be implementable by any user, without elevated privilege, and pass design rule checks.

FPGA programmable routing meets all three conditions. Specifically, we target the route between an FPGA register and a CLB. The programmable route can be composed to be arbitrarily long sequences of transistors to increase the observable BTI effects, and it is trivial to use as a route. Additionally, programmable routes often carry sensitive data (e.g., encryption keys and machine learning weights). Thus, verifying that a route between a register and LUT is vulnerable to a pentimenti attack threatens the data integrity of most FPGA designs.

4. BTI Sensor

Our threat models depend on the ability to detect BTI effects in cloud FPGA resources. BTI degradation and recovery manifests as changes in the timing delays of the victim resource. BTI effects differ depending on the previous state that the victim held on that resource.

Time-to-digital converters (TDCs) measure ns-scale timing changes by sensing the propagation delay through an FPGA-instantiable delay line. There exists a large body of prior work on implementing TDC sensors within cloud FPGAs [15], [20], [22], [25], [55], [56]. Our experiments use the open-source Tunable Dual-Polarity TDC [15].

Figure 3 demonstrates how to use a Tunable Dual-Polarity TDC sensor [15] to measure BTI effects. The original sensor is designed for power measurement; we amend it to perform BTI timing measurement and exploit the threats presented in Section 2. The constituent structures of the TDC are presented below:

Programmable Clock Generator: This component generates the two clock domains: the `Launch Clock` and `Capture Clock`. These two clocks are identical in frequency with a runtime-programmable phase relationship defined by θ. Two clocks are necessary as the TDC is comparing how long it takes for the signal from the `Launch Clock` to reach a destination within the `Carry Chain` compared to when the `Capture Clock` causes the `Capture Registers` to sample. The `Launch Clock` first must be converted to a logic signal, which is performed by the `Transition Generator`.

Transition Generator: This component is responsible for sending positive ($0 \rightarrow 1$) and negative ($1 \rightarrow 0$) transitions through the the `Route Under Test`, `Carry Chain`, and into the `Capture Registers`. The same θ defines the relationship between the signal traversing the `Route Under Test`/`Carry Chain` is launched and when it measured in the `Capture Registers`. When the sensor is loaded onto an FPGA, θ is set to 0; an offset of θ is consistent between sensor design loadings. When θ is set correctly, a transition will be propagating through the delay line when the `Capture Registers` are clocked and record a metastable transition region. The distance that transition propagates is called the propagation distance, and is related to the propagation delay of the logic in the `Route Under Test` and `Carry Chain`.

Route Under Test: The primary intent of the TDC is to measure the timing delay through some FPGA programmable route that is affected by burn-in. When θ is properly configured the output of the `Capture Registers` reflects how far a rising or falling signal has propagated, and consequently the timing delay though the `Route Under Test`. BTI degradation causes the propagation delay to increase. The propagation delay decreases during BTI recovery.

Carry Chain: The primary structure of the TDC is a long linear array of delay elements. This is formed by a series of combinatorial logic elements that are able to propagate rising ($0 \rightarrow 1$) and falling transitions ($1 \rightarrow 0$). Ideally, each element is identical, with a timing delay of τ, so that the propagation of signals is uniform at every stage of the `Carry Chain`. To ensure consistency throughout the chain, the delay elements are uniformly placed and routed in consecutive physical locations. Our chosen sensor uses the fast look-ahead CARRY primitives of the Xilinx FPGA devices to construct this chain.

Capture Registers: Each element of the `Carry Chain` is output to a register, forming the `Capture Registers`. These registers are activated synchronously by the `Capture Clock`, that performs a capture of the state of the `Carry Chain`. If a rising ($0 \rightarrow 1$) or falling transition ($1 \rightarrow 0$) is propagating through the `Carry Chain`, and the `Capture Registers` are activated, the distance that signal has traveled will be captured.

Propagation Distance: Each falling and rising transition is captured at the output registers, as shown at the bottom of

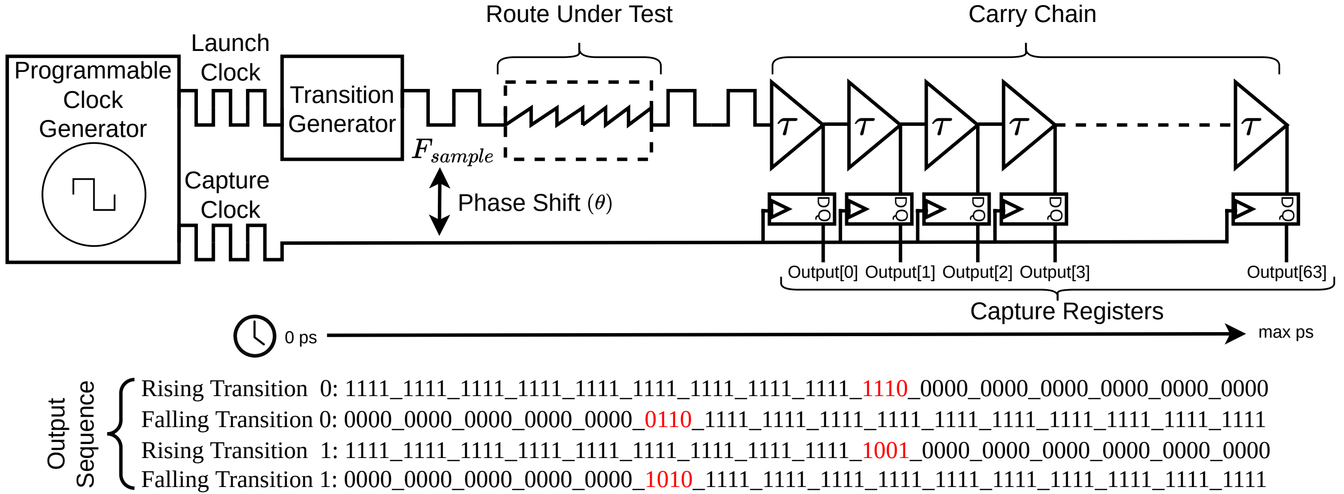

Figure 3: **Time-to-Digital Converter (TDC) sensor that measures BTI effects in FPGA routes.** A `Programmable Clock Generator` produces two clock domains: the `Launch Clock` and `Capture Clock`. The `Launch Clock` domain generates a signal that propagates through a test route and into a textttCarry Chain – an array of linear delay elements (64 in this example). The signal moves through the Route Under Test and into the TDC `Carry Chain`, when it is recorded as indicated by the `Capture Clock`. The TDC `Capture Registers` record the signal propagation distance. This provides is a measure of propagation delay of the `Route Under Test`. Taking measurements over time records changes of the propagation delay due BTI remnants stored on the transistors.

Figure 3. Rising Transition 0 shows that the $0 \rightarrow 1$ transition reached `Output[38]`. Falling Transition 0 shows that the $1 \rightarrow 0$ propagated to somewhere between `Output[21]` and `Output[23]`, with some metastability between the two points. Rising Transition 1 propagates differently; the $0 \rightarrow 1$ transition propagates to between `Output[36]` and `Output[39]`. Similarly, Falling Transition 1 propagates to between `Output[20]` and `Output[23]`. These changes can represent deviations in the timing through the `Route Under Test` (i.e. burn-in).

Post Processing: The output of the `Capture Registers` can be processed into a single value that represents the propagation time of a signal through a `Route Under Test`. This is done by computing the *Binary Hamming Distance* of the output registers. This is defined for the rising transitions as the binary Hamming distance from `64'h_0000_0000_0000_0000`, and for falling transitions, the binary Hamming distance from `64'h_ffff_ffff_ffff_ffff`. The *Binary Hamming Distance* of the example samples in Figure 4 will yield the sequence: 39, 22, 38, 22.

5. Experimental Setup

We perform a series of experiments to determine the extent of BTI effects on programmable routing and the ability to recover the digital data that was previously held on those routes. We perform the first experiments locally on a ZCU102 Ultrascale+ board. Subsequent experiments are performed remotely on AWS F1 instances. The experiments use a target design and a measurement design described in Section 5.1. Section 5.2 details the three experimental phases that calibrate the sensor, conditions (performs burn-in) on some target routes, and measures the BTI effects on the target routes. Section 5.3 determine the approximate target route lengths for assets from the OpenTitan hardware root of trust. The experiments in Section 6 demonstrate the ability to execute Threat Models 1 and 2 remotely on an AWS F1 instance.

5.1. Experimental Designs

Our experiments use two independent FPGA designs. The *Target* design holds the routes under test at a constant 1 or 0 value for a pre-determined duration to induce BTI effects. The *Measure* design records changes in route propagation delays caused by BTI effects. The data held on these routes represent the Type A or B logical data that an attacker wishes to recover.

Target Design: Figure 4 presents the *Target* design that biases a set of routes by holding them to a fixed 0 or 1 value. We represent a logical 1 bias with the color green and logical 0 bias with the color red.

The region of slices above the routes under test is explicitly left uninitialized (no logic may be placed there) during the compilation process. These slices will be used by the *Measure* design for the placement of `Carry Chains`. Using these slices could introduce noise, or worse, erroneous results, into our propagation delay measurements of the test route. In theory this noise would be minimal as the length of the route through these slices is significantly shorter than the route being tested for BTI. While this does not negate all possible burn-in effects in the slices, since the

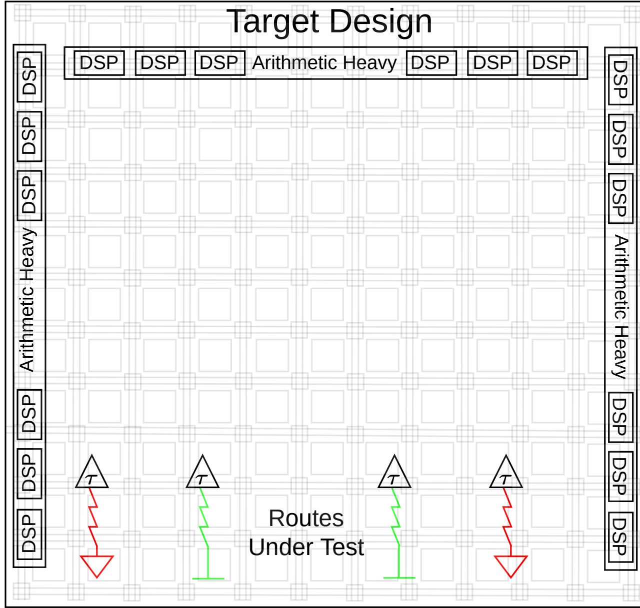

Figure 4: The **Target Design** conditions a set of pre-determined routes to 1/VCC (green) or 0/GND (red) aka the *burn value*. This induces a BTI effects on transistors of each route. The `Arithmetic Heavy` circuit is used emulate the surrounding logic of many FPGA computations, and also increases on-chip temperature to accelerate BTI. The center of the design is left empty; it will be used in the measure design (Figure 5).

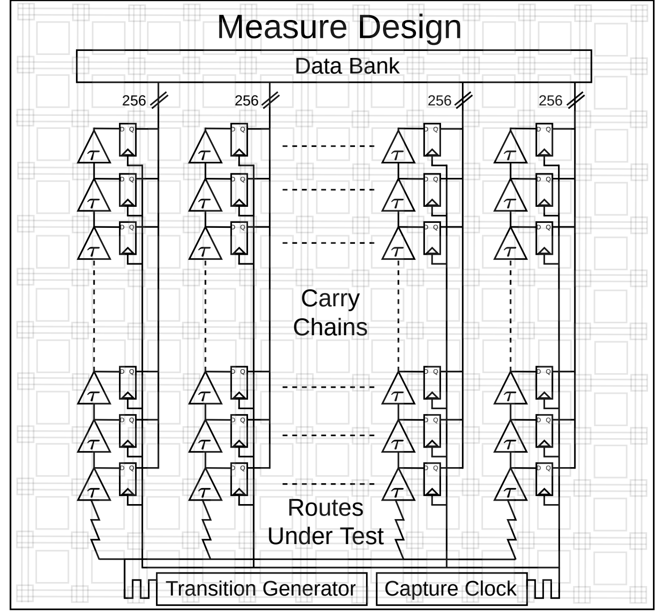

Figure 5: The **Measure Design** records the BTI degradation of multiple `Routes Under Test` using TDC sensors. As per Section 4, the `Transition Generator` is used to send rising ($0 \rightarrow 1$) and falling ($1 \rightarrow 0$) transitions through the tested routes. The changing propagation delay of these signals indicates the BTI effects on that route.

vendor-determined state of an uninitialized slice could also introduce BTI effects, it at least suggests that all slices will be affected equally. This is consistent with **Threat Model 1**, as the attacker may be the one publishing the maliciously constructed AFI, and thus has the control to leave these slices empty.

The routes containing **Type A** and **Type B** information is surrounded by other computation. We generalize these structures as `Arithmetic Heavy` circuits implemented as arrays of logic performing a pipelined fused multiply-add operation (similar to a machine learning or lattice cryptography accelerator). This has the added benefit of accelerating the BTI effect through increased heat generation.

Measure Design: Figure 5 presents a high level view of the architecture for measuring the propagation delay of the routes under test. These routes represent the **Type A** or **Type B** data an attacker intends to recover. By tracking the change in this propagation delay caused by BTI degradation and recovery, the side channel can be exposed to exploit **Threat Model 1** and **2**.

Each of these routes is a single `Route Under Test` from Figure 3 and we instantiate an array of TDCs as presented in Section 4. No extraneous routing outside of the the `Route Under Test` is used to connect the `Transition Generator` to the `Carry Chain`. Identical routing constraints from the *Target* design are used to generate the routes for the *Measure* design. The routes test a variety of lengths and placements in order to build a general understanding of how burn-in is affected by route characteristics.

5.2. Experimental Phases

These designs will then be used to form three experimental phases: *Calibration* to configure the TDCs of the *Measure* design, *Condition* to induce the BTI effect on a predefined set of routes, *Measurement* to measure BTI effects.

Calibration Phase: *Calibration* is the first phase, and determines a baseline θ value that captures rising and falling transitions traveling through the tested routes into the `Carry Chain` and `Capture Registers`. The TDC alone does not provide an absolute measure of the change in propagation delay through a tested route. We can use the TDC determine the change in the propagation delay by examining the increase or decrease over time of the *Binary Hamming Distance* output of the sensor tuned to this baseline θ value. To create a baseline, a short series of 2^4 samples (called a trace) is taken from each TDC as θ is iteratively reduced until the rising and falling transition appear in the output registers. We call this value θ_{init}, and an individual value is computed and saved for every route under test in the *Measure* design.

Condition Phase: The *Target* design is loaded onto the FPGA and a pre-defined set of *burn values* is applied to the

TABLE 1: **OpenTitan Earl Grey Distribution of Route Lengths (ps) on a Virtex UltraScale+.** This table reports the distribution of route lengths (in ps) for a selection of twenty security-critical assets in an OpenTitan Earl Grey implemented on a Virtex UltraScale+. Assets are sorted in ascending order by maximum route length. Route lengths of more than 1000 ps are common, and would increase when OpenTitan shares an FPGA with other logic.

#	Asset Paths	Type	Bus Width	MEAN	SD	MIN	25%	50%	75%	MAX
1	/otp_ctrl_otp_lc_data[state]	SV/T	320	169.5	98.1	39	95.5	157.5	228	509
2	/u_otp_ctrl/otp_ctrl_otp_lc_data[test_exit_token]	SV/T	128	197.5	115.4	37	114	170	242.2	534
3	/otp_ctrl_otp_lc_data[rma_token]	SV/T	101	239.8	122.8	38	148	222	325	583
4	/otp_ctrl_otp_lc_data[test_unlock_token]	SV/T	128	207.9	120.1	38	130.5	178.5	247.2	609
5	/keymgr_aes_key[key][1]_282	CK	32	538.3	106.4	380	433.5	551	614	738
6	/keymgr_otbn_key[key][0]_285	CK	384	219.8	150.9	41	99	167	327.2	919
7	/keymgr_kmac_key[key][0]_28	CK	256	317.6	141.7	49	213.8	291	408	1050
8	/otp_ctrl_otp_keymgr_key[key_share0]	CK	256	187.3	200.8	37	54	109	217	1064
9	/u_otp_ctrl/part_scrmbl_rsp_data	CK	64	353.4	146.1	116	267.2	348.5	411.2	1075
10	/keymgr_aes_key[key][0]_283	CK	256	360.3	154.2	86	270	333	412.2	1311
11	/u_otp_ctrl/u_otp_ctrl_scrmbl/gen_anchor_keys	CK	135	220.1	358.7	0	57	94	162.5	1333
12	/otp_ctrl_otp_keymgr_key[key_share1]	CK	256	262.5	273.4	37	51	158	335.5	1381
13	/csrng_tl_rsp[d_data]	S	32	1291.8	105.7	1031	1244.8	1323	1359.8	1432
14	/aes_tl_rsp[d_data]	S	32	1105.3	411.4	276	1135.8	1279	1369.5	1631
15	/keymgr_otbn_key[key][1]_284	CK	32	1062.7	281.2	480	854	1074.5	1270	1670
16	/u_otp_ctrl/part_otp_rdata	S	64	1298.9	213	933	1118.5	1311.5	1447.2	1784
17	/flash_ctrl_otp_rsp[key]	CK	128	1816.6	404.6	1215	1503	1717.5	2010.2	3245
18	/kmac_app_rsp	S	777	94.2	179.7	15	40	58	97	3398
19	/flash_ctrl_otp_rsp[rand_key]	CK	128	1908.1	670.7	553	1337	1882	2308.8	3706
20	/aes_tl_req[a_data]	S	32	2114.8	471.8	1455	1805	2079.5	2337.2	3946

routes under test. These *burn values* will either be **Type A** or **Type B** information that will induce variable BTI effect based on their value. The `Arithmetic Heavy` circuits are activated in this phase to emulate user computation and exacerbate BTI degradation.

Measurement Phase: The *Measurement* phase loads the *Measure* design and tunes all of the TDCs to their respective θ_{init}. Ten traces are taken from each TDC as θ is iteratively decreased from θ_{init}, to avoid relying on a single trace that could be affected by architectural irregularities [15], [17], [21]. For each route, the mean *Binary Hamming Distance* is computed on across all samples within each trace, and then the mean of all traces is computed to obtain a single value representing the propagation delay through the route under test. This value is converted to a measure of time based on a derived relationship of $\frac{2.8ps}{bit}$ for UltraScale+ parts [15], [64]. Any deviation in this value represents BTI-induced variation on a route.

5.3. OpenTitan Hardware Root of Trust

We study the OpenTitan hardware root of trust (RoT) to provide a context for a realistic target for our threat models. OpenTitan is a commercial-grade, open-source hardware root of trust [44]. The OpenTitan Earl Grey integrated into systems to carry out core, security-critical functionalities related to trusted platform module (TPM), platform integrity, and 2nd factor authentication. The OpenTitan consists of a security-enhanced RV32IMCB RISC-V Ibex core, cryptographic IP cores (e.g., AES, KMAC, HMAC), and memories (e.g., ROM, eFLASH, SRAM, OTP) protected by access control mechanisms.

OpenTitan is an open-source hardware design; all design files are available online. OpenTitan encourages a design flow where the user solely modifies the boot ROM (the data used to initialize the FPGA memory) of a pre-built bitstream. Thus, it adheres to threat model assumption 1 (Section 2) that the assets locations are known.

OpenTitan has many important security assets that govern its operation. Assets include cryptographic keys to encrypt data stored in off-chip memories (e.g., one-time programmable memory), keys to scramble data before transmission across on-chip busses to limit power side channels, and life-cycle related state values/tokens for attestation, identity management, and debug control.

We identify twenty security-critical assets in the following groups:

- **Cryptographic Keys (CK):** OpenTitan has a number of cryptographic keys that need to be protected that are spread across the design. This includes keys stored in the one-time programmable (OTP) memory and OpenTitan Key Manager. The OTP controller has an access control mechanism that arbitrates OTP data accesses and buffers key values;
- **State Values or Tokens (SV/T):** Assets stored in one-time-programmable (OTP) memory for use in the life-cycle controller, which influences the OpenTitan's DFT functionality, NVM backdoor access, debug, and CPU functionality;
- **Signals (S):** Variables carrying sensitive information from/to security peripherals.

Table 1 reports the route length distribution of twenty security-critical assets in OpenTitan implemented on a Virtex UltraScale+. **Bus Width** records the number of routes associated with each asset. **MEAN** and **SD** are the mean route length and standard deviation for each asset's routes,

respectively. **MIN** is the minimum route length for each asset. **25%**, **50%**, and **75%** are the route length for the 25th, 50th, and 75th percentiles of each asset's routes, respectively. **MAX** records the maximum route length for each asset. The assets are sorted in ascending order according to **MAX** route length.

Most routes are short – only a few hundred picoseconds. However, there are longer route lengths that approach 4 ns. When integrated with other cores or accelerators, it is logical that these route lengths will increase.

6. Experimental Results

The experiments interleave the calibration, condition, and measurement phases (Section 5.2) to extract pentimenti – BTI-induced side channels in FPGA programmable routing. The target route lengths are informed by a study of OpenTitan asset delay in Section 5.3. Experiment 1 uses a new ZCU102 Ultrascale+ FPGA development board. This allows us to characterize the burn-in effect while controlling temperature, eliminate previous FPGA usage, and minimize system computation during measurement. The remote cloud platform provides substantially less control over environmental conditions. Experiment 2 validates **Threat Model 1** on the AWS F1 platform. Experiment 3 validates **Threat Model 2** on the AWS F1 platform.

6.1. Experiment 1 (Lab Environment)

Experiment 1 studies BTI degradation and recovery effects on a local, new FPGA. The experiment validates that the burn-in degradation occurs and is differentiable, making **Threat Model 1** possible. Additionally, it shows that BTI is non-permanent. BTI recovery is observable and measurable, which is required for **Threat Model 2**.

A new ZCU102 Ultrascale+ is placed in a temperature-controlled forced convection oven (Lab Companion OF-01E) set to 60°C. The oven maintains a constant temperature, which ensures that temperature changes do not influence the delays. The ZCU102 is factory new; it will experience the largest BTI effects since no degradation has occurred.

64 routes are studied on the ZCU102 Ultrascale+ development board. The first group of 16 routes each has a delay of 1000 ps, the second 2000 ps, the third 5000 ps, and the fourth 10000 ps delay. The delay relates to the number of transistors used to form that programmable route subject to BTI effects.

Experiment 1 starts with a 200-hour burn-in period of the 64 routes. A randomly generated value X is applied to the routes. X is held constant over the entire 200 hours. The goal is to induce burn-in on the routes and understand the extent of BTI effects. Burn-in is followed by a 200-hour recovery period that applies a constant $\overline{X}$ to the routes under test to induce BTI recovery.

Experiment 1 is divided into three experimental periods consisting of phases from Section 5.2:

- **Hour 0:** The *Calibration* phase is executed to compute the θ_{init} for each of the 64 routes.
- **Hours [0,200):** The burn-in period alternates between *Condition* and *Measurement* phases. The *Condition* phase applies the *burn values* X to the 64 routes for one hour. Then, the *Measurement* phase is launched, which tunes the phase difference between the transition generator and capture clock (θ_{init}) to ensure the transition falls within the carry chain. The TDC sensors capture data for each of the 64 routes under test as described in Section 5.1. This condition/measurement sequence is repeated $200\times$ (approximately 200 hours). The *Measurement* phase runs once per hour to record the BTI degradation effects. Measurement is fast, taking less than a minute. Thus the vast majority of the time is spent in the condition phase.
- **Hours [200,400):** The recovery period inverts the values previously held on the routes under test. This is nearly identical to hours [0,200), except the *Condition* phase loads $\overline{X}$, the complement of X, into the 64 routes. This period focuses on understanding BTI recovery effects.

Figure 6 plots the 400-hour results of Experiment 1 for four different route delays. The four graphs each have 16 routes with various delays – (a) 1000ps, (b) 2000ps, (c) 5000ps, and (d) 10000ps. A switch from burn-in values $\overline{X}$ to recovery values $\overline{X}$ happens at the 200-hour mark, denoted by the transition between red and green backgrounds.

Every net is measured once per hour. A *Measure* phase is loaded, and the TDC sensor readings are recorded. Ten traces are taken from each TDC as the phase shift θ is decreased to avoid architectural irregularities [15], [17], [21]. We record the rising and falling transitions within each trace. For each transition type, the mean propagation distance is computed across all samples within a trace. Then, we compute the mean of the ten traces to obtain a single value representing the propagation delay through the route under test. Next, we subtract the rising transition distance from the falling transition distance to isolate the BTI effect to the *Route Under Test*. This bit value is converted to seconds based on a derived relationship of $\frac{2.8ps}{bit}$ for UltraScale+ parts [15], [64]. Finally, we center the data to the point at hour zero; any deviation from zero represents BTI degradation or recovery-induced variation on a route, which we call Δ ps. Measurement took about 52 seconds for all the routes in our cloud experiments; only 1.4% of the overall time is spent performing measurements.

The choice to plot the falling minus rising is based on the fact that architecturally the falling (rising) transition stresses the NMOS (PMOS) transistors in a route. This is due to NMOS (PMOS) devices being best suited to passing 0 (1) values. And so, the timing difference between the rising and falling signal reduces to a single value for each hour that considers both the NBTI and PBTI degradation. Figure 2 discusses this in more depth.

Datapoints are colored cyan if their burn value X is a logical 0 and purple if their burn value X is a logical 1. A red background indicates that the value applied to the

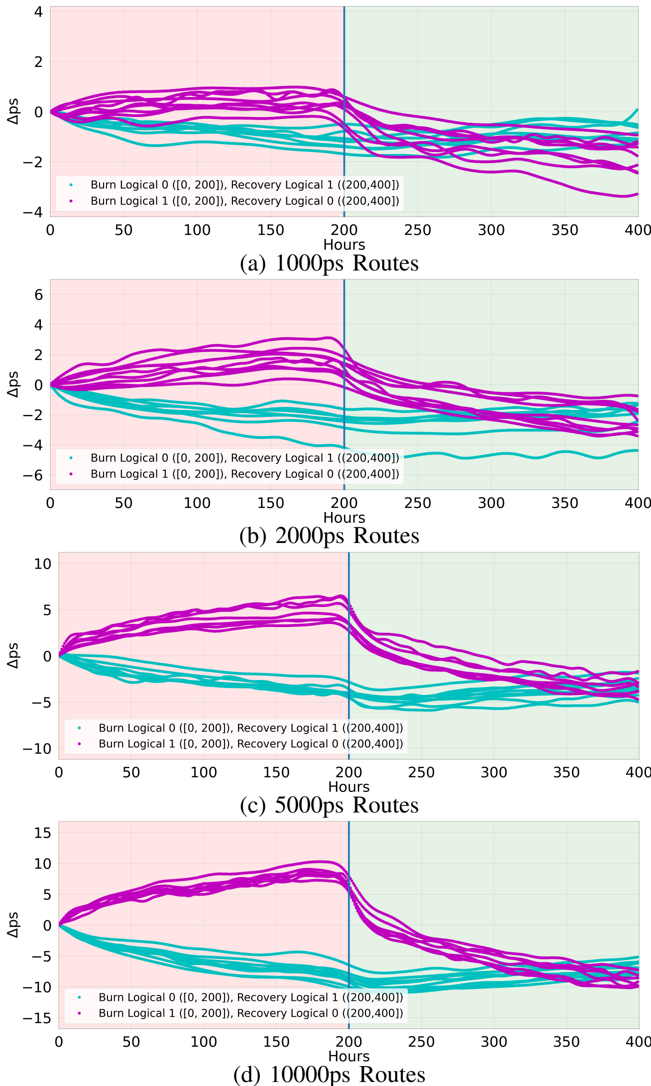

(a) 1000ps Routes

(b) 2000ps Routes

(c) 5000ps Routes

(d) 10000ps Routes

Figure 6: **Experiment 1 (Lab Environment):** Sets of 16 routes of varying lengths are initialized on a new ZCU102 FPGA. The experiment occurs at 60°C in a temperature-controlled environment. During the initial 200-hour burn-in period, a random burn value X is conditioned into the routes (red background). After which, 200-hour BTI recovery is induced by applying $\overline{X}$ into the routes under test (green background). The timing difference between each route's falling and rising transition delay is recorded every hour. Routes conditioned with logical 0 behave differently than routes conditioned with logical 1 in the burn-in and recovery periods. This reveals the unique effects of BTI degradation and recovery. It indicates that **Threat Model 1** and **Threat Model 2** are possible.

routes is the burn value X, and green background is the BTI recovery period where the values are complemented $\overline{X}$. The resulting time series are smoothed with a kernel regression, which finds a non-linear relationship between a pair of variables. Specifically, the Python statsmodels package's nonparametric kernel regression class is used in continuous mode with a local linear estimator.

A trend is immediately apparent in the first 200 hours (red half) of the charts in Figure 6. Burn value 0 (cyan) routes decrease from hour zero. Burn value 1 (magenta) routes increase from hour zero. This trend occurs regardless of the length of the route, but the magnitude differs. While a 10000 ps route is on the longer end for most designs, 1000 ps routes are commonplace, e.g., see the OpenTitan study in Section 5.3. Furthermore, this helps us understand the limits of our strategy. There appear to be no limitations in route length as to observable burn-in effects, with the 1000 ps tested routes showing a clear difference between GND and VCC burned routes.

*These results indicate that **Threat Model 1** is possible. If an attacker can observe BTI effects on a route before and after a design, they can easily deduce the burn value on a route and observe a side channel.*

Figure 6a shows that the 1000 ps routes have $\pm[1, 2]$ ps difference between the rising and falling transition at the 200-hour mark. The 2000 ps routes have a $\pm[2, 3]$ ps difference (Figure 6b), the 5000 ps routes have a $\pm[5, 6]$ ps difference (Figure 6c), and the 10000 ps routes have a $\pm[10, 11]$ ps difference. This data matches our expectation in burn-in behavior: the number of transistors in the route (i.e., the route length) directly increases the route's delay.

At the 200-hour mark, the experiment moves from the burn-in to recovery. The condition route values change from X to $\overline{X}$. The routes with logical 1 burn-in in X in the first 200 hours (and logical 0 in the recovery period) quickly return to their pre-burn state across all route lengths. This recovery takes approximately 30-50 hours before the propagation delay difference between the rising and falling transition has returned to the original state at hour 0. We do not see the same behavior in the routes that were logical 0 in the first 200 hours and logical 1 in the second 200 hours; they recover, but the process takes much longer (over 200 hours).

These results indicate that BTI is elastic and non-permanent. In addition, we can see that the BTI recovery in routes conditioned by burn 1 is substantially faster than in routes with burn 0. This pattern persists for all tested route lengths, suggesting a fundamental difference between the NBTI and PBTI effect on the 16nm FinFET transistors of the UltraScale+ device. The difference in BTI recovery enables **Threat Model 2**.

The quick recovery of the burn 1 routes indicates that they might be easier to detect when targeting Type B data (Threat Model 2). A Threat Model 2 attacker obtains the FPGA during the recovery period. They will not know the initial values and thus cannot complement them. The attacker must set the target values to logical 0 or 1. Since the Burn 1 degradation values see the greatest and fastest recovery, the attacker would set all recovery values to condition to logical 0 to observe the quick recovery. This motivates us to set the recovery values to logical 0 in Experiment 3 (Section 6.3).

6.2. Experiment 2 (Cloud Environment)

Experiment 2 tests the viability of **Threat Model 1** on an AWS F1 cloud FPGA, which aims to extract sensitive data from a rented third-party design. The attacker can load and unload the design and wants to extract design intellectual property, e.g., netlist constant holding cryptographic keys or neural network weights. The cloud environment provides no control over temperature, and it is likely the device is years old, making BTI effects less observable [3]. This experiment is performed in the eu-west-2 AWS region, which puts potentially fours years of wear on the device.[3]

We examine 16 1000 ps, 16 2000ps, 16 5000 ps routes and 16 10000 ps delay routes. 64 random bit values X are applied to these routes in the condition phase. The *Target* and *Measure* designs are built around these routes according to Section 5.1. The *Target* design is configured to utilize 3896 DSPs for this architecture. The *Target* design consumes 63 W of a maximum of 85 W imposed by AWS.

Experiment 2 is divided into two periods:

- **Hour 0:** The *Calibration* phase is executed to compute the θ_{init} for each of the 64 routes.
- **Hours [0,200]:** During each hour, we run a long *Condition* phase and a short *Measurement* phase. The *Condition* phase applies the *burn values* X to the 64 routes under test. The *Measurement* phase tunes TDCs to θ_{init} and captures delay estimates for each route as described in Section 5.1. After this brief measurement process (33 seconds), the data is saved. This is repeated 200 times (over 200 hours) to study **Threat Model 1**.

Figure 7 shows the results from Experiment 2, testing **Threat Model 1** on the AWS F1 platform. X is the **Type A** (Design Data) an attacker wishes to recover. The values of X are expressed in cyan for burn 0 values, and magenta for burn 1 values but are opaque to the attacker.

Figures 7a, 7b, 7c, 7d demonstrate the same trends as in Experiment 1. These results are expectedly noisier than from the ZCU102, which was a new part held at a constant temperature. The difference between the burn 0 rising and falling transitions decreases. In contrast, the burn 1 routes behave the exact opposite; the propagation delay difference between the rising and falling transition increases. This makes the routes easily distinguishable, with burn 0 (cyan) decreasing immediately from hour zero and burn 1 (magenta) increasing immediately from hour zero. This pattern persists irrespective of the route length (1000ps vs. 2000ps vs. 5000 ps vs. 10000 ps), but the magnitude differs. By examining the trends in data, X can be derived, demonstrating **Threat Model 1** recovery of **Type A** data is possible.

The 1000 ps routes of Figure 7a experience $\pm[0, .2]$ ps difference, the 2000 ps routes of Figure 7b experience $\pm[0, .4]$ ps difference, the 5000 ps routes of Figure 7c experience $\pm[0, 1]$ ps difference, and the 10000 ps routes

3. https://aws.amazon.com/about-aws/whats-new/2018/10/amazon-ec2-f1-instance-expands-to-more-regions-adds-new-features-and-improves-development-tools/

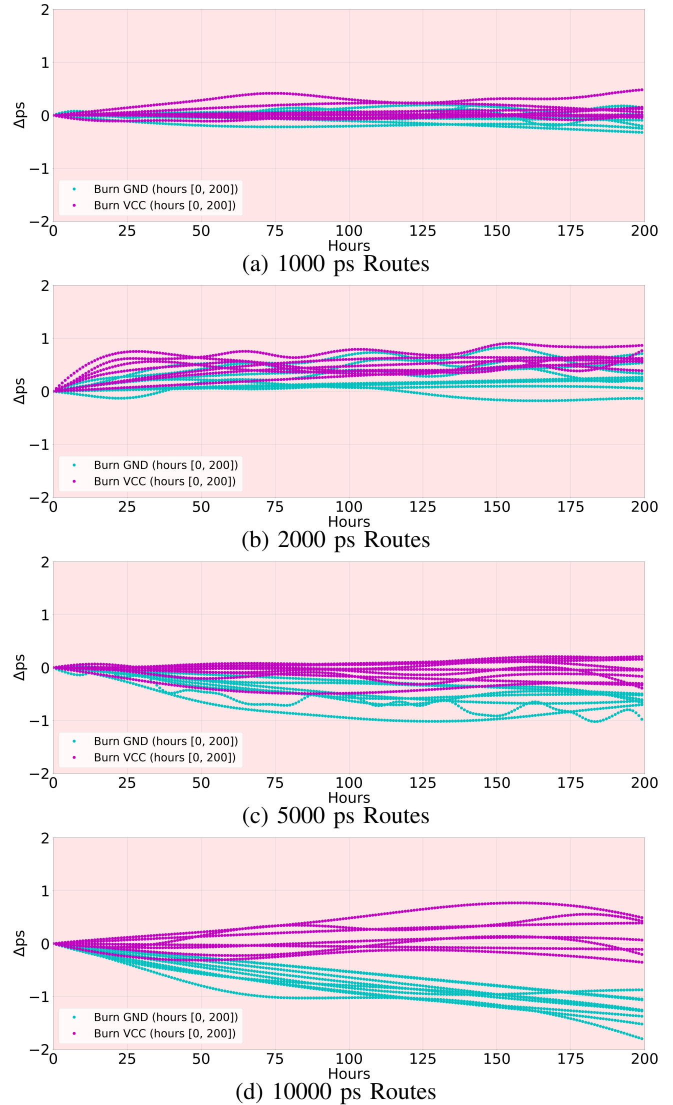

(a) 1000 ps Routes

(b) 2000 ps Routes

(c) 5000 ps Routes

(d) 10000 ps Routes

Figure 7: **Experiment 2 (Cloud Environment):** A random burn value X is conditioned into four sets of FPGA routes. 7a has 16 1000 ps routes, 7b has 16 2000 ps routes, 7c has 16 5000 ps routes and 7d has 16 10000 ps routes. The propagation delay difference (Δps) between the falling and rising transition is measured once per hour over 200 hours. Over time, the burn-in of the 0 (cyan) and 1 (purple) values produces observable changes in their route delays due to BTI degradation. This enables an attacker to recover **Type 1** data (Design Data) and execute **Threat Model 1.**

in Figure 7d experience $\pm[0, 2]$ ps difference. This further validates our previous conclusion that the amount of burn-in is proportional to the tested route length. It also indicates that the burn-in for the cloud FPGAs is lesser than that of the new ZCU102 from Experiment 1 (compare to Figure 6c and 6d). This is not surprising, given that cloud FPGAs are older and more used. Thus, it is more challenging to extract pentimenti from cloud FPGAs than a local FPGA.

It is now clear how **Threat Model 1** can be exploited by an attacker. In this model, an attacker instantiates a

design that contains **Type A** sensitive information on an AWS FPGA. The attacker knows the location of sensitive data routes, and they construct a *Measurement* design that maps TDCs to these routes. The attacker can interleave measurements every hour to eventually exposes the **Type A** data based on the relationship between the rising and falling transition over time. The attacker can continue the burn-in process until they are satisfied that the sensitive values are extracted.

6.3. Experiment 3 (Cloud Environment)

Experiment 3 studies the viability of **Threat Model 2**, which exploits BTI recovery as a side channel. Sixteen 1000 ps, 2000 ps, 5000 ps, and 10000 ps delay routes are instantiated on an AWS F1 FPGA. The routes undergo a burn-in period during which we do not measure the time delay. Then, they switch to the recovery phase. We aim to understand if it is possible to recover pentimenti of a previous user of the FPGA by only measuring during the recovery period. This meets the assumptions required for Threat Model 2. θ_{init} is consistent across all FPGAs of the same type, and so capturing θ_{init} once on any board is sufficient to assume that θ_{init} is known a priori for any attack.

Experiment 3 is divided into three periods:

- **Hours [0, 200]:** A burn-in period induced by the victim computation. The *Condition* phase executes with a constant, randomly generated X value loaded into the 64 routes under test. Calibration is **not** performed. The attacker does not have control of the FPGA, and thus measurement is not allowed. The condition phase is run uninterrupted for 200 hours.
- **Hour 200:** The victim relinquishes control of the FPGA, and the attacker gains control.
- **Hours (200, 225]:** The attacker launches the *Measurement* phase, which tunes TDCs to θ_{init}, and captures traces of the 64 routes under test as described in Section 5.1. After this quick process, the data is saved. Then, the attacker launches the *Condition* phase that runs for one hour. The Measurement/Condition sequence is repeated for 25 iterations (hours).

The target victim design holds a constant 64-bit X value on FPGA routes for 200 hours without interruption. X represents **Type B** (User Data) the attacker aims to recover. The burn-in X values are opaque to the attacker. But, X can be derived by the attacker based on the recovery behavior.

After 200 hours, the attacker gains control of the FPGA and thus can start measuring for BTI effects. The attacker sets all routes under test to logical 0 and measures the propagation delay once per hour over the next 25 hours. The attacker is looking for BTI recovery to extract pentimenti of the previously loaded design data. We do not assume the attacker has prior information about the FPGA before the victim computed upon it.

The choice of setting all routes to logical 0 is motivated by the results in Experiment 1; routes that were logical 1 in X and were switched to $\overline{X}$ quickly returned to the original value. Thus, it exhibits a more significant signal for detection. An attacker could also choose to set all lines to logical 1 or a mixture of 0 and 1.

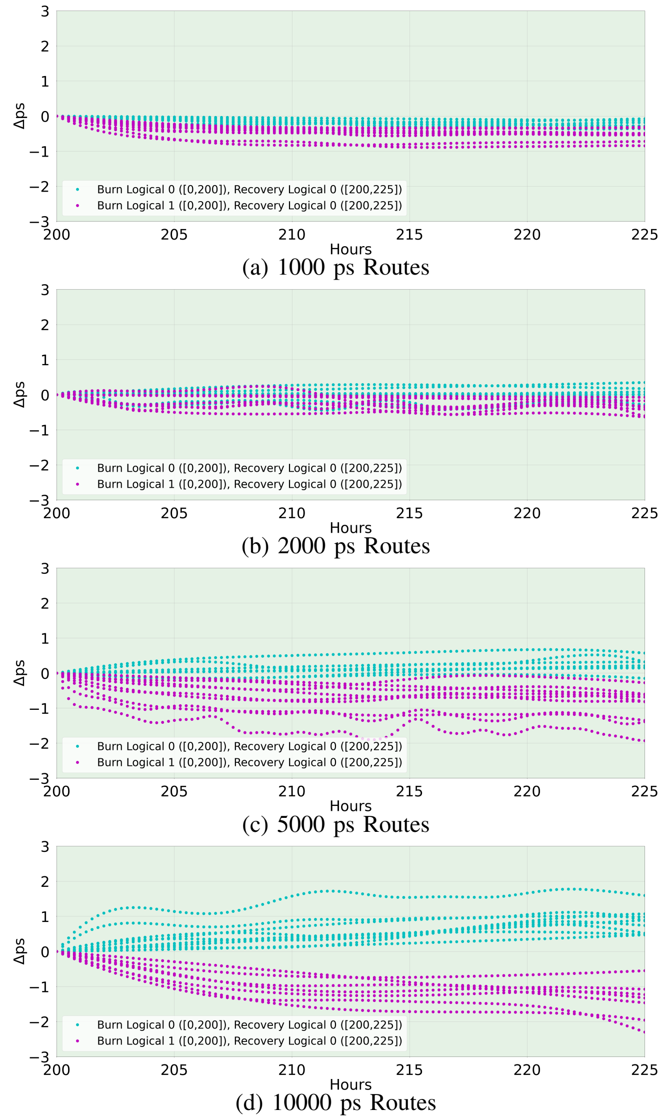

(a) 1000 ps Routes
(b) 2000 ps Routes
(c) 5000 ps Routes
(d) 10000 ps Routes

Figure 8: **Experiment 3 (Cloud Environment):** Four sets of 16 routes with different delays are initialized on an AWS F1 FPGA. A random constant burn value X is conditioned into the routes for 200 hours. The attacker gains control of the FPGA and instantiates TDC sensors to measure the timing delay of routes that previously held sensitive data. The timing difference between the falling and rising transition is plotted for the following 25-hour period—revealing the BTI recovery. This enables an attacker to execute on **Threat Model 2** and recover **Type 2** data (user data).

Figure 8 shows the 25-hour recovery period. Note that the graph starts at hour 200 after the burn-in. We have no data about the FPGA before that point. Burn 0 values are shown in cyan burn 1 values in magenta. The purple routes

that previously held logical 1 immediately begin to decrease in relation to the cyan routes that remain at logical 0 the entire time. The purple logical 1 routes are undergoing Burn 1 BTI recovery, which Experiment 1 showed was much more dramatic than Burn 0 recovery.

We do not observe the same magnitude and clarity of divergence of burn-in or recovery on AWS F1 as the ZCU102. This is likely due to more complex cloud operating factors, including non-constant temperature, FPGA age, and other computations simultaneously running on the AWS F1 system.

Despite these differences, the result is consistent with the elasticity we observed on the ZCU102.

We show an attacker can extract Type B Data demonstrating **Threat Model 2** on a cloud FPGA. An attacker generates a *Measure* design which maps TDCs to the routes which contained **Type B** (User Data) in a targeted design. After gaining access to an AWS F1 FPGA device running the target design that was previously loaded for tens of hours, the attacker recovers previous user data by measuring the timing behavior of the routes over time.

7. Related Work

Our attack is a *single-tenant temporal side channel* – state within the FPGA that is not wiped correctly, or it is impossible to remove between subsequent users [9]. It is common to "wipe" the FPGA device between successive users [36] as a security precaution. **Our approach subverts these efforts** as it measures analog remanence that remains even after wiping. We show that our data recovery techniques work even after performing the wiping done by AWS. It is impossible to mitigate burn-in risk via a logical erasure of the device because burn-in is a fundamental characteristic of the device transistors that reflects previous logical values.

Tian et al. [61] demonstrate a single-tenant temporal covert channel. They use ring oscillators to heat the FPGA (transmitter) and detect temperature (receiver). They can transmit hundreds of bits over a few minutes on cloud FPGA at Texas Advanced Computing Center using Microsoft Catapult hardware. To make their covert channel, the FPGA transmitter and receiver must alternatively obtain and release the same FPGA, which is possible but very difficult in other cloud infrastructures (e.g., AWS). Using temperature as a side channel requires the user to get on the FPGA quickly; temperature effects are short-term, e.g., the cloud FPGAs return to ambient temperatures within a few minutes [61]. Finally, BTI effects are a more pernicious temporal channel. Instead of measuring the tertiary effects of computation or a covert channel, it is a direct measurement of a previous user or proprietary design data. It can last hundreds of hours, as we have shown in our results.

Zick et al. [71] demonstrate a single-tenant temporal side channel on a local FPGA by recovering previous user data stored in LUT SRAMs. Their experiment has a burn-in period of 922 hours at high temperatures to induce burn-in. Then, the FPGA sat powered off for several weeks. Their experiments were performed on a local Xilinx Kintex-7 KC705 development board. Unfortunately, their experimental requirements are incompatible with the cloud FPGA attack model. They use a highly precise, off-chip oscillator to enhance the on-chip TDC sensor timing resolution. This results in femtosecond-level timing precision. Such precision is impossible on cloud FPGA TDC sensors since an attacker cannot use off-chip components. On-chip TDCs operate at approximately 10 ps precision on the UltraScale+, so it is an order of magnitude difference with their sensor. They perform recovery of data stored in FPGA LUTs (SRAM) and specifically target transistors in the output buffers of the SRAM bits. We ruled out the examination of this resource since their burn-in effects are too subtle to measure with cloud FPGA sensors, which is why they required femtosecond precision. We target FPGA programmable routing. We show that our attacks are deployable on cloud FPGAs (AWS F1 instances).

A significant body of prior work uses ring oscillators (RO)-based sensors to measure long-term FPGA BTI effects [4], [47], [58]. RO sensors build a combinatorial loop through a tested component and an inverter. The oscillation frequency through the loop reflects the time taken for the signal to propagate through that tested component, which changes due to BTI effects. While ROs measure BTI effects, they have two significant limitations. First, ROs have a single variable output—the frequency of oscillation – that integrates the propagation speed through the NMOS and PMOS transistors. This is an essential factor as BTI stresses PMOS vs. NMOS transistors differently. Our TDC sensor can separate the differences in BTI stress on PMOS and NMOS. We use this ability to differentiate between BTI degradation. Second, ROs are often not allowed on cloud FPGAs. ROs use combinatorial loops, which violate the design rule checks and can be detected [32], [34]. Cloud FPGA providers can disallow designs that contain self-oscillating circuits, e.g., as is done by AWS. Our TDC-based sensor is more challenging to detect since it uses computational structures that are common and many FPGA designs. It was implemented on an AWS F1 instance. Thus, it passes AWS design rule checks.

Previous works have recovered SRAM user data on recycled ICs [12], [29], [63]. Even though SRAMs are a form of volatile memory, where logical data is lost on power-off, an imprint is left behind and is recoverable. These techniques rely on measuring the statistical power-on state of SRAM bits. They assume a different threat model, e.g., requiring physical access to the chip.

8. Mitigations

This paper demonstrated that **Threat Model 1** and **Threat Model 2** are exploitable in cloud systems. A determined attacker could build more precise sensors to measure BTI on shorter routes with shorter burn-in periods. As a result, users should take precautions to manage sensitive data to mitigate burn-in effects, cloud FPGA providers should look to enforce stronger temporal boundaries between users,

and FPGA manufacturers should consider architectural solutions to mitigate BTI.

8.1. User Mitigations

The cloud FPGA user should not allow sensitive data to sit unchanged on the FPGA for long periods to avoid burn-in remnants that the following user could discover. When sensitive data must statically persist for long periods, the user should consider methods to mitigate its burn-in effects.

Techniques that periodically change sensitive data would reduce burn-in. For example, the data could be inverted at predetermined periods (e.g., every hour). Or it could be deterministically shuffled at the source and unshuffled at the receiver. Such data transformation approaches reduce the burn-in effects across the route at the expense of increasing the design complexity. Other ideas related to FPGA wear leveling [57] would likely reduce the burn-in effects but need to be verified.

If there are natural breaks in computation, the user could move between different FPGAs in the cloud. A new device should be checked out from the cloud provider, the application moved between FPGAs, and the burn-in would start fresh on the new FPGA. This would increase user design complexity, e.g., by requiring a robust process to stop, move application data, and restart an FPGA instance. This adds risks related to data corruption when moving between FPGAs.

The user should strive to make routes that hold sensitive data as short as possible. The longer the route, the more transistors affected by burn-in, and the larger the burn-in effects. Consequently, as we have shown in our results, shorter routes are a more secure FPGA design pattern.

FPGA physical design tools generally attempt to make routes as short as possible. Hand-placed routes for sensitive information could produce better results. Physical design tools often focus on routes on the critical path, often at the expense of other routes. The ability to specify that the physical design tools minimize sensitive routes would reduce vulnerability to pentimento-style attacks.

The user or design tools could place sensitive routes in a manner that makes them difficult to connect to a BTI sensor. The input to the route under test must be connected to the transition generator. The output of the route under test is connected to a TDC sensor. Placing the inputs or outputs in locations that make these connections challenging would make it more challenging to extract BTI information from that target route.

Verification tools could analyze the design or bitstream for sensitive data residing on long routes. The ability to provide reports about the route lengths of the sensitive information would allow hardware security verification engineers to better assess their data vulnerabilities w.r.t. to a pentimento attack. Providing a more precised measure of protection (e.g., vulnerability metric) enables even stronger hardware security verification.

Key rotation is common in cryptography [10], [16] and could be employed on cloud FPGAs. This is not always possible, especially if data needs to be embedded into the RTL directly, e.g., in random netlist constants as found in the OpenTitan.

Key masking [2], [26], [42] could help reduce the number and lengths of routes that hold a cryptographic key. Masking is specific to cryptographic algorithms and may not be feasible for other types of sensitive data.

The design could use partial reconfiguration to move the sensitive information – its storage and computation units – to different locations of the chip. This would act as a form of wear leveling. It lessens the burn-in effect at any one physical location. Yet, it also spreads the burn-in over more areas, which could potentially make it easier to exploit the information.

A cloud FPGA user could mitigate the BTI remnants by erasing their design and holding on to the instance for some time before relinquishing it back into the user pool. The tenant could invert the values of the sensitive routes to speed up the recovery and thus limit the remaining BTI signal. Or they could perform some other actions (perhaps toggling the routes). This costs the user money commensurate with the time they deemed sufficient to erase BTI effects.

8.2. Cloud Provider Mitigations

The primary issue cloud providers could hope to resolve is the rapid reallocation of FPGA devices once a user relinquishes them. The cloud provider could implement launch rate controls, by withholding devices after they are returned, for days, weeks, or longer to mitigate the ability to recover the burn-in. This would push the mitigation onto the cloud provider rather than the cloud user.

The cloud provider can attempt to combat the accelerators of the BTI effect: higher voltage and temperature. Some FPGAs that operate at different voltages and use a lower voltage would reduce the burn-in effects. Similarly, higher temperatures exacerbate burn-in. Temperature can be managed to some extent, but it would be very challenging to control the on-chip temperature to the point where an attacker can no longer observe BTI. Furthermore, cloud providers are already incentivized to control voltage and temperature to reduce FPGA power consumption and aging.

8.3. FPGA Manufacturer Mitigations

FPGA manufacturers can attempt to mitigate FPGA BTI effects. BTI mitigations are already commonly considered to increase reliability. It is unlikely that FPGA manufacturers will be able to eliminate BTI, especially at advanced design nodes. BTI effects are more negligible at less advanced process nodes; thus, falling back on older technology would be a potential mitigation. The performance and power benefits of advanced nodes are likely too much to sacrifice for cloud providers and users.

Manufacturers can help reduce BTI through voltage and temperature mitigations; however, this is already a primary directive due to their negative influence on power consumption. Thus, it is unlikely these mitigations will advance at

a faster pace. FPGA manufacturers could consider more advanced dynamic voltage scaling techniques to allow users to mitigate BTI selectively. This adds complexity to the design, which increases costs.

9. Conclusion

We find that a remote attacker can recover "FPGA pentimentos" – long-removed secret data belonging to a prior user or proprietary design image on a cloud FPGA. Just as a pentimento of a painting can be exposed via infrared imaging, FPGA pentimentos can be exposed via signal timing sensors instantiated on a remote cloud FPGA. The sensitive data constituting an FPGA pentimento is imprinted to the device through bias temperature instability effects on the underlying transistors. We demonstrate how this slight degradation can be measured using a time-to-digital converter when an adversary programs one into the target cloud FPGA. This technique allows an attacker to ascertain previously safe information, after it is no longer explicitly present, on cloud FPGAs. Notably, it can allow an attacker to (1) extract proprietary details or keys from an encrypted FPGA design image available on the AWS marketplace and (2) recover information from a previous user of a cloud-FPGA. Both threat models are experimentally validated on the AWS F1 platform.

References

[1] Video transcoding. https://www.xilinx.com/applications/data-center/video-imaging.html, 2022.

[2] Mehdi-Laurent Akkar and Christophe Giraud. An implementation of des and aes, secure against some attacks. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 309–318. Springer, 2001.

[3] Muhammad Ashraful Alam and Souvik Mahapatra. A comprehensive model of pmos nbti degradation. *Microelectronics Reliability*, 45(1):71–81, 2005.

[4] Abdulazim Amouri, Florent Bruguier, Saman Kiamehr, Pascal Benoit, Lionel Torres, and Mehdi Tahoori. Aging effects in fpgas: An experimental analysis. In *2014 24th international conference on Field Programmable Logic and Applications (FPL)*, pages 1–4. IEEE, 2014.

[5] Arvind Arasu, Ken Eguro, Manas Joglekar, Raghav Kaushik, Donald Kossmann, and Ravi Ramamurthy. Transaction processing on confidential data using cipherbase. In *2015 IEEE 31st International Conference on Data Engineering*, pages 435–446. IEEE, 2015.

[6] Amittai Aviram, Sen Hu, Bryan Ford, and Ramakrishna Gummadi. Determinating timing channels in compute clouds. In *Proceedings of the 2010 ACM workshop on Cloud computing security workshop*, pages 103–108, 2010.

[7] AWS. AWS FPGA - frequently asked questions. https://github.com/aws/aws-fpga/blob/master/FAQs.md, 2022.

[8] Sarvesh Bhardwaj, Wenping Wang, Rakesh Vattikonda, Yu Cao, and Sarma Vrudhula. Predictive modeling of the nbti effect for reliable design. In *IEEE Custom Integrated Circuits Conference 2006*, pages 189–192. IEEE, 2006.

[9] Christophe Bobda, Joel Mandebi Mbongue, Paul Chow, Mohammad Ewais, Naif Tarafdar, Juan Camilo Vega, Ken Eguro, Dirk Koch, Suranga Handagala, Miriam Leeser, et al. The future of fpga acceleration in datacenters and the cloud. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, 15(3):1–42, 2022.

[10] Dan Boneh, Kevin Lewi, Hart Montgomery, and Ananth Raghunathan. Key homomorphic prfs and their applications. In *Annual Cryptology Conference*, pages 410–428. Springer, 2013.

[11] Andrew Boutros, Mathew Hall, Nicolas Papernot, and Vaughn Betz. Neighbors from hell: Voltage attacks against deep learning accelerators on multi-tenant fpgas. In *2020 International Conference on Field-Programmable Technology (ICFPT)*, pages 103–111. IEEE, 2020.

[12] Cagla Cakir, Mudit Bhargava, and Ken Mai. 6t sram and 3t dram data retention and remanence characterization in 65nm bulk cmos. In *Proceedings of the IEEE 2012 Custom Integrated Circuits Conference*, pages 1–4. IEEE, 2012.

[13] GCKY Chen, KY Chuah, MF Li, Daniel SH Chan, CH Ang, JZ Zheng, Y Jin, and DL Kwong. Dynamic nbti of pmos transistors and its impact on device lifetime. In *2003 IEEE International Reliability Physics Symposium Proceedings, 2003. 41st Annual.*, pages 196–202. IEEE, 2003.

[14] Yu-Ting Chen, Jason Cong, Zhenman Fang, Jie Lei, and Peng Wei. When spark meets fpgas: A case study for next-generation DNA sequencing acceleration. In *8th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 16)*, 2016.

[15] Colin Drewes, Olivia Weng, Keegan Ryan, Bill Hunter, Christopher McCarty, Ryan Kastner, and Dustin Richmond. Turn on, tune in, listen up: Maximizing side-channel recovery in time-to-digital converters. In *Proceedings of the 2023 ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, FPGA '23, page 111–122, New York, NY, USA, 2023. Association for Computing Machinery.

[16] Adam Everspaugh, Kenneth Paterson, Thomas Ristenpart, and Sam Scott. Key rotation for authenticated encryption. In *Annual International Cryptology Conference*, pages 98–129. Springer, 2017.

[17] Claudio Favi and Edoardo Charbon. A 17ps time-to-digital converter implemented in 65nm fpga technology. In *Proceedings of the ACM/SIGDA international symposium on Field programmable gate arrays*, pages 113–120, 2009.

[18] Jeremy Fowers, Kalin Ovtcharov, Michael Papamichael, Todd Massengill, Ming Liu, Daniel Lo, Shlomi Alkalay, Michael Haselman, Logan Adams, Mahdi Ghandi, et al. A configurable cloud-scale dnn processor for real-time ai. In *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*, pages 1–14. IEEE, 2018.

[19] Ilias Giechaskiel, Kasper Rasmussen, and Jakub Szefer. Reading between the dies: Cross-slr covert channels on multi-tenant cloud fpgas. In *2019 IEEE 37th International Conference on Computer Design (ICCD)*, pages 1–10. IEEE, 2019.

[20] Ilias Giechaskiel, Kasper Rasmussen, and Jakub Szefer. C3APSULe: Cross-fpga covert-channel attacks through power supply unit leakage. In *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, 2020.

[21] Ognjen Glamočanin, Louis Coulon, Francesco Regazzoni, and Mirjana Stojilović. Are cloud fpgas really vulnerable to power analysis attacks? In *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1007–1010. IEEE, 2020.

[22] Dennis RE Gnad, Cong Dang Khoa Nguyen, Syed Hashim Gillani, and Mehdi B Tahoori. Voltage-based covert channels using fpgas. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 26(6):1–25, 2021.

[23] Dennis RE Gnad, Fabian Oboril, Saman Kiamehr, and Mehdi B Tahoori. Analysis of transient voltage fluctuations in fpgas. In *2016 International Conference on Field-Programmable Technology (FPT)*, pages 12–19. IEEE, 2016.

[24] Dennis RE Gnad, Fabian Oboril, and Mehdi B Tahoori. Voltage drop-based fault attacks on fpgas using valid bitstreams. In *2017 27th International Conference on Field Programmable Logic and Applications (FPL)*, pages 1–7. IEEE, 2017.

[25] Mustafa Gobulukoglu, Colin Drewes, Bill Hunter, Ryan Kastner, and Dustin Richmond. Classifying Computations on Multi-Tenant FPGAs. In *Design Automation Conference (DAC)*, DAC '21, 2021.

[26] Jovan D Golić and Christophe Tymen. Multiplicative masking and power analysis of aes. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 198–212. Springer, 2002.

[27] T Grasser, B Kaczer, Ph Hehenberger, W Gos, R O'connor, H Reisinger, W Gustin, and C Schlunder. Simultaneous extraction of recoverable and permanent components contributing to bias-temperature instability. In *2007 IEEE International Electron Devices Meeting*, pages 801–804. IEEE, 2007.

[28] Ph Hehenberger, H Reisinger, and Tibor Grasser. Recovery of negative and positive bias temperature stress in pmosfets. In *2010 IEEE International Integrated Reliability Workshop Final Report*, pages 8–11. IEEE, 2010.

[29] Joshua Hovanes, Yadi Zhong, and Ujjwal Guin. Beware of discarding used srams: Information is stored permanently, 2022.

[30] Mehmet Sinan Inci, Berk Gulmezoglu, Thomas Eisenbarth, and Berk Sunar. Co-location detection on the cloud. In *International Workshop on Constructive Side-Channel Analysis and Secure Design*, pages 19–34. Springer, 2016.

[31] Jonas Krautter, Dennis RE Gnad, and Mehdi B Tahoori. Fpgahammer: Remote voltage fault attacks on shared fpgas, suitable for dfa on aes. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 44–68, 2018.

[32] Jonas Krautter, Dennis RE Gnad, and Mehdi B Tahoori. Mitigating electrical-level attacks towards secure multi-tenant fpgas in the cloud. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, 12(3):1–26, 2019.

[33] AT Krishnan, C Chancellor, S Chakravarthi, PE Nicollian, V Reddy, A Varghese, RB Khamankar, and S Krishnan. Material dependence of hydrogen diffusion: Implications for nbti degradation. In *IEEE InternationalElectron Devices Meeting, 2005. IEDM Technical Digest.*, pages 4–pp. IEEE, 2005.

[34] Tuan Minh La, Kaspar Matas, Nikola Grunchevski, Khoa Dang Pham, and Dirk Koch. Fpgadefender: Malicious self-oscillator scanning for xilinx ultrascale+ fpgas. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, 13(3):1–31, 2020.

[35] Kyong Taek Lee, Wonchang Kang, Eun-Ae Chung, Gunrae Kim, Hyewon Shim, Hyunwoo Lee, Hyejin Kim, Minhyeok Choe, Nae-In Lee, Anuj Patel, Junekyun Park, and Jongwoo Park. Technology scaling on high-k & metal-gate finfet bti reliability. In *2013 IEEE International Reliability Physics Symposium (IRPS)*, pages 2D.1.1–2D.1.4, 2013.

[36] Miriam Leeser, Suranga Handagala, and Michael Zink. Fpgas in the cloud. *Computing in Science & Engineering*, 23(6):72–76, 2021.

[37] Yukui Luo, Cheng Gongye, Yunsi Fei, and Xiaolin Xu. Deepstrike: Remotely-guided fault injection attacks on dnn accelerator in cloud-fpga. In *2021 58th ACM/IEEE Design Automation Conference (DAC)*, pages 295–300. IEEE, 2021.

[38] S Mahapatra, N Goel, S Desai, S Gupta, B Jose, S Mukhopadhyay, K Joshi, A Jain, AE Islam, and MA Alam. A comparative study of different physics-based nbti models. *IEEE Transactions on Electron Devices*, 60(3):901–916, 2013.

[39] Souvik Mahapatra. *Fundamentals of Bias Temperature Instability in MOS Transistors*. Springer, 2016.

[40] Dina Mahmoud and Mirjana Stojilović. Timing violation induced faults in multi-tenant fpgas. In *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1745–1750. IEEE, 2019.

[41] Prasanth Mangalagiri, Sungmin Bae, Ramakrishnan Krishnan, Yuan Xie, and Vijaykrishnan Narayanan. Thermal-aware reliability analysis for platform fpgas. In *2008 IEEE/ACM International Conference on Computer-Aided Design*, pages 722–727. IEEE, 2008.

[42] Stefan Mangard, Thomas Popp, and Berndt M Gammel. Side-channel leakage of masked cmos gates. In *Cryptographers' Track at the RSA Conference*, pages 351–365. Springer, 2005.

[43] Kaspar Matas, Tuan Minh La, Khoa Dang Pham, and Dirk Koch. Power-hammering through glitch amplification–attacks and mitigation. In *2020 IEEE 28th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pages 65–69. IEEE, 2020.

[44] OpenTitan. *The OpenTitan Hardware Root of Trust offical website*. https://opentitan.org/.

[45] OpenTitan. OpenTitan FPGA setup. https://docs.opentitan.org/doc/getting_started/setup_fpga/, 2022.

[46] Fabien AP Petitcolas. Kerckhoffs' principle., 2011.

[47] Petr Pfeifer and Zdenek Pliva. On measurement of parameters of programmable microelectronic nanostructures under accelerating extreme conditions (xilinx 28nm xc7z020 zynq fpga). In *2013 23rd International Conference on Field programmable Logic and Applications*, pages 1–4, 2013.

[48] Thomas Pöppelmann, Michael Naehrig, Andrew Putnam, and Adrian Macias. Accelerating homomorphic evaluation on reconfigurable hardware. In *International Workshop on Cryptographic Hardware and Embedded Systems*, pages 143–163. Springer, 2015.

[49] Andrew Putnam, Adrian M Caulfield, Eric S Chung, Derek Chiou, Kypros Constantinides, John Demme, Hadi Esmaeilzadeh, Jeremy Fowers, Gopi Prashanth Gopal, Jan Gray, et al. A reconfigurable fabric for accelerating large-scale datacenter services. In *2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA)*, pages 13–24. IEEE, 2014.

[50] Adnan Siraj Rakin, Yukui Luo, Xiaolin Xu, and Deliang Fan. {Deep-Dup}: An adversarial weight duplication attack framework to crush deep neural network in {Multi-Tenant}{FPGA}. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 1919–1936, 2021.

[51] S Ramey, J Hicks, LS Liyanage, and S Novak. Bti recovery in 22nm tri-gate technology. In *2014 IEEE International Reliability Physics Symposium*, pages XT–2. IEEE, 2014.

[52] Sanjay Rangan, Neal Mielke, and Everett CC Yeh. Universal recovery behavior of negative bias temperature instability [pmosfets]. In *IEEE International Electron Devices Meeting 2003*, pages 14–3. IEEE, 2003.

[53] Hans Reisinger, O Blank, Wolfgang Heinrigs, A Muhlhoff, Wolfgang Gustin, and C Schlunder. Analysis of nbti degradation-and recovery-behavior based on ultra fast vt-measurements. In *2006 IEEE International Reliability Physics Symposium Proceedings*, pages 448–453. IEEE, 2006.

[54] Thomas Ristenpart, Eran Tromer, Hovav Shacham, and Stefan Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *Proceedings of the 16th ACM conference on Computer and communications security*, pages 199–212, 2009.

[55] Falk Schellenberg, Dennis RE Gnad, Amir Moradi, and Mehdi B Tahoori. An inside job: Remote power analysis attacks on fpgas. In *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1111–1116. IEEE, 2018.

[56] Falk Schellenberg, Dennis RE Gnad, Amir Moradi, and Mehdi B Tahoori. Remote inter-chip power analysis side-channel attacks at board-level. In *2018 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 1–7. IEEE, 2018.

[57] Edward Stott and Peter YK Cheung. Improving fpga reliability with wear-levelling. In *2011 21st International Conference on Field Programmable Logic and Applications*, pages 323–328. IEEE, 2011.

[58] Edward A Stott, Justin SJ Wong, Pete Sedcole, and Peter YK Cheung. Degradation in fpgas: measurement and modelling. In *Proceedings of the 18th annual ACM/SIGDA international symposium on Field programmable gate arrays*, pages 229–238, 2010.

[59] Shanquan Tian, Ilias Giechaskiel, Wenjie Xiong, and Jakub Szefer. Cloud fpga cartography using pcie contention. In *2021 IEEE 29th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pages 224–232. IEEE, 2021.

[60] Shanquan Tian, Andrew Krzywosz, Ilias Giechaskiel, and Jakub Szefer. Cloud fpga security with ro-based primitives. In *2020 International Conference on Field-Programmable Technology (ICFPT)*, pages 154–158. IEEE, 2020.

[61] Shanquan Tian and Jakub Szefer. Temporal thermal covert channels in cloud fpgas. In *Proceedings of the 2019 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 298–303, 2019.

[62] Shanquan Tian, Wenjie Xiong, Ilias Giechaskiel, Kasper Rasmussen, and Jakub Szefer. Fingerprinting cloud fpga infrastructures. In *Proceedings of the 2020 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 58–64, 2020.

[63] Stephen M Trimberger. System and method of detecting and reversing data imprinting in memory, January 22 2013. US Patent 8,359,447.

[64] Xilinx. Xilinx ultrascale clocking. https://docs.xilinx.com/v/u/en-US/ug572-ultrascale-clocking, 2018.

[65] Xilinx. FINN examples. https://github.com/Xilinx/finn-examples, 2022.

[66] Zhang Xu, Haining Wang, and Zhenyu Wu. A measurement study on co-residence threat inside the cloud. In *24th USENIX Security Symposium (USENIX Security 15)*, pages 929–944, 2015.

[67] Jian Fu Zhang, Rui Gao, Meng Duan, Zhigang Ji, Weidong Zhang, and John Marsland. Bias temperature instability of mosfets: Physical processes, models, and prediction. *Electronics*, 11(9):1420, 2022.

[68] Yinqian Zhang, Ari Juels, Alina Oprea, and Michael K Reiter. Homealone: Co-residency detection in the cloud via side-channel analysis. In *Security and Privacy (SP), 2011 IEEE Symposium on*, pages 313–328. IEEE, 2011.

[69] Mark Zhao and G Edward Suh. Fpga-based remote power side-channel attacks. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 229–244. IEEE, 2018.

[70] Kenneth M Zick and John P Hayes. Low-cost sensing with ring oscillator arrays for healthier reconfigurable systems. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, 5(1):1–26, 2012.

[71] Kenneth M Zick, Sen Li, and Matthew French. High-precision self-characterization for the lut burn-in information leakage threat. In *2014 24th International Conference on Field Programmable Logic and Applications (FPL)*, pages 1–6. IEEE, 2014.