An Investigation on Data Center Cooling Systems Using FPGA-based Temperature Side Channels

Yuan Liang
Virginia Tech
yl6194@vt.edu

Xing Gao
University of Delaware
xgao@udel.edu

Kun Sun
George Mason University
ksun3@gmu.edu

Wenjie Xiong
Virginia Tech
wenjiex@vt.edu

Haining Wang
Virginia Tech
hnw@vt.edu

***Abstract*—As power and cooling cost has become a major factor in the total cost of ownership (TCO) of large-scale data centers, it is important to investigate how data centers run their cooling systems in practice. The data centers of Amazon Web Services (AWS) have been continuously expanding worldwide, and their restrictive security policies keep many management aspects of data centers private. In this paper, we make an attempt to explore the cooling systems of AWS data centers without privileged accesses. We first demonstrate PVT (process, voltage, and temperature) variations in AWS FPGAs (Field Programmable Gate Arrays) using time-digital converters (TDC). We further leverage the DRAM temperature side channel and improve the usage of the TDC to measure the temperature change accurately. We conduct a measurement on the daily temperatures of AWS data centers worldwide and find that temperature changes of some data centers are closely related to local weathers. Thus, we deduce they adopt free cooling techniques. This measurement study motivates us to re-think the vulnerability of data centers to power/thermal attacks.**

***Index Terms*—FPGA, side channel, Data Center, Cooling System.**

I. Introduction

As data centers have expanded their scales with more powerful servers to meet the increasing service demands, the amount of heat emitted by those servers also significantly increases and requires the cooling system to prevent overheating[28]. Improving the efficiency of data centers and using most of energy for computing are important[22][32]. The power and cooling systems have played a major role in the total cost of ownership (TCO) of large-scale data centers, which motivates data centers to adopt the power over-subscription and aggressive cooling strategies for cost reduction. Nowadays, there are various technical solutions for data center cooling systems[1][10]. However, although there are significant improvements, they are still far from the ideal due to many factors like the non-linearity of air dynamics. More importantly, warehouse-scale data centers keep their inside information private for security reasons[5][6][22]. Thus, it is difficult to comprehensively investigate their cooling systems.

In this paper, we take advantage of the temperature information leakage through side channels to learn the cooling system inside a data center without privileged accesses. We choose physical side channels of the FPGA (Field Programmable Gate Array) to investigate AWS (Amazon Web Services) data centers. As the cloud FPGA becomes popular, covert/side channels research on FPGAs has been intensive. However, when more and more FPGA-based covert/side channels are identified, some of them are not practical for remote attacks on current data centers[35]. In addition, most of the existing works focus on information leakage among different tenants. Based on previous research, we choose the DRAM and the time-digital converter (TDC) as tools for the temperature estimation. They are used to measure the temperatures of data centers that power the public cloud.

To implement a precise and high-resolution FPGA-based side channel in order to sense a temperature change, we introduce the spatial average technique for TDCs. It recovers power ripple precisely so that the power ripple can be filtered out, and we can observe the temperature effect on the signal propagation in the FPGA. We also use this technique to measure the switching frequency and transient response of the power system of AWS FPGAs, and we demonstrate the process variation of AWS FPGAs. Leveraging the FPGA-based temperature side channel in TDC and DRAM, we measure the daily temperature changes of AWS data centers worldwide (in region-level and zone-level) that provide the FPGA service. We analyze the collected temperature data, and we observe that temperature changes of some AWS data centers are highly related to local weathers, which evidences that these data centers have adopted the free air cooling technique for cost saving and environmentally friendly purposes. Moreover, the behaviors of the cooling system and temperature dynamics caused by computing equipment inside data centers are studied. Finally, we discuss the threats of power/thermal attacks and cloud cartography on the WSC (Warehouse-scale computing) data center.

Note that in this paper, precision, high resolution, and accuracy are not equivalent terms. The precision is the consistency of repeated measurements, and the accuracy represents the closeness to the true value. High resolution means the uncertainty of a measurement value is small[34]. Each availability zone in AWS corresponds to a subnet, so the availability zone and subnet are used interchangeably. An instance in AWS is a virtual machine configured with designed resources. AWS is continuously growing, and its security policies may change over time.

The rest of this paper is organized as follows. Section II presents the background and related work. Section III

describes the experiments on TDCs to show PVT variations of FPGAs. Section IV focuses on investigation of data centers with temperature side channels. Section V discusses the limited application of free cooling in practice and security vulnerability posed by temperature information leakage. Finally, Section VI concludes this work.

II. Background and Related Work

A. Data Center Cooling

The design consideration of a data center cooling system mainly includes effectiveness, cost, and reliability. The critical part of a cooling system is building loops to collect the heat and propel it to the outside. CRACs (computer room air conditioners) are typically used to dissipate the collected heat to the outside[10]. The CRAC is effective, but its operational cost is high.

Since a cooling system once costed a large portion of energy consumption in data centers[26], free cooling is an innovative design to minimize the power consumption of the cooling system by taking advantages of the cool air outside. However, free cooling still requires artificial cooling, and the temperature management should consider the overall condition of the data center[21]. Temperature is an important factor in the power consumption of the cooling system and various computing equipment[13]. Understanding the temperature pattern can help us to deduce the conditions of computing equipment. Even if physical accesses to data centers are limited, FPGA-based temperature side channels allow us to remotely investigate data centers.

Cool weather is applicable for cooling data centers in many places on the Earth. The cold nights and winter seasons provide excess cold air, and so turning on the air conditioning or chiller under such scenarios is wasteful. But using cool outside air is not entirely free[1]. Note that water is used as a heat medium for some advanced designs. First of all, outside air should be processed to remove dust. Secondly, cooling fans are required to transfer the air, or pumps are used to move the water. Finally, when the outside temperature is high, the chiller operates to cool down the air. In the case of freezing weather, the cold outside air should be mixed with hot outlet air to maintain a warm environment for the computing equipment, and water should be prevented from being frozen. Since water cannot work as the heat medium in extremely low temperatures, radiators can be used as coolers. Radiators rely on coolants that are not frozen below water freezing points to transfer heat.

Figure 1 shows the elements considered in this study. The temperature condition of a data center is affected by the local weather, the cooling system, and the heat from the computing equipment. FPGAs in the computing equipment leaks the temperature information to a remote adversary, and the further information can be deduced based on the leakage.

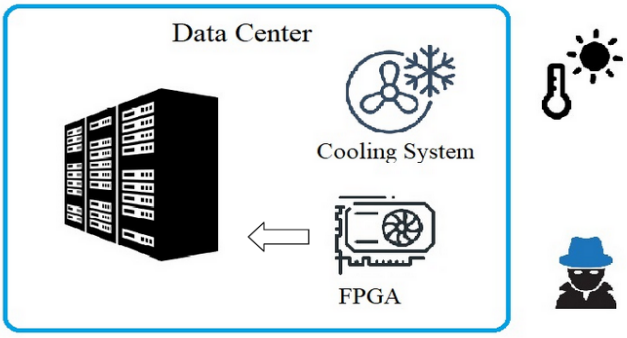

Fig. 1: A temperature change in the data center can be observed by FPGA side channels.

B. Cloud FPGA

FPGA services have been increasingly provided by data centers recently. Compared to CPUs and GPUs, FPGAs are reconfigurable hardware that can implements digital logic with higher performance and lower power consumption for accelerating certain workloads. We focus on Amazon EC2's F1 instance as an example. Current public available regions that support F1 instances are US East (Northern Virginia), US West (Oregon), Asia(Sydney) and Europe (Ireland, Frankfurt, and London)[2]. The signal propagation in FPGA is affected by process variation, voltage, and temperature (PVT). Based on the effects of voltage and temperature, many side/covert channels are developed. The research on FPGA-based side/covert channels has progressed rapidly in the past decade. It advances from the lab setting to the cloud setting[35]. Ziener et al.[57] used the power supply pin to communicate with the FPGA. The time-digital converter (TDC) and the ring oscillator implemented with look-up tables (LUT-RO) are tools available to measure the signal propagation in the FPGA. Zick et al.[56] showed the measurement of transient on FPGA with TDC. Because the signal of the power supply correlates with the delay of FPGA, the oscilloscope is not needed. The LUT-RO can be used to develop side/covert channels through a power delivery system[15][55], but it requires restrictive settings and is not generally applicable for the current cloud setting. Because temperature affects the signal propagation of FPGA, the LUT-RO can be used as a temperature sensor to develop covert channels. Iakymchuk et al.[25] built a temperature covert channel between two electrically isolated parts of FPGA. Tian et al.[44] used stressors to heat up multiple cloud FPGAs and measure temperatures of FPGAs to build a covert channel. Generally, the TDC and LUT-RO are similar and interchangeable in many cases, except the TDC outputs higher resolution data and requires more memory.

In this paper, we use the AWS F1 instance to implement the temperature sensor for the measurement study. Previous studies show that temperature information can be used to estimate power consumption[30], evaluate the vulnerability of cooling system[14], and receive information in the covert channel[24].

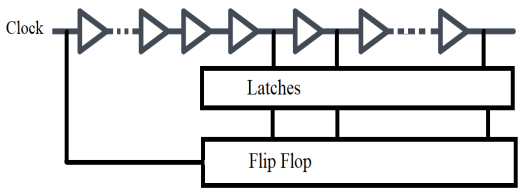

Fig. 2: Time Digital Converter (Gate-level Illustration)

C. Co-Residence and Remote Side Channels

Co-residence is essential for launching many malicious attacks in cloud environments. For example, exploiting co-residence, an adversary can locate victims with fingerprints and then build covert channels. Xu et al.[53] studied co-residence threats inside the AWS data centers and revealed policies and patterns of VM placement. However, AWS often change their policies. Now the `traceroute` command no longer reports any meaningful information for private IPs in AWS. Thus, we have to explore the physical side/covert channels of cloud FPGAs.

Besides the signal propagation of FPGA we discussed before, DRAM and PCIe accessed by the FPGA also can be used to capture the temperature change, fingerprint, and develop a cross-FPGA covert channel[16][17][42][45][51]. Furthermore, all computing equipment in the data center generates heat, and so the airflow design is critical for the cooling system. Since Guri et al. [24] used air as the medium to create a covert channel between two PCs, it is possible that the airflow in the data center is used for the construction of a covert channel.

D. Power/Thermal Attacks on the Data Center

The threats of power/thermal attacks have been recognized in the past decade[14][40][54], in which adversaries exploit over-subscriptions of power supplies and reduced redundancy of cooling systems to cause power outages or/and heat up machines to disrupt or degrade the reliability and performance of data centers. Xu et al.[54] demonstrated that attackers can force victim servers to reach their power peaks at the same time and then trip the circuit breaker, causing power outages. Gao et al.[14] introduced rack-level and data center-level thermal attacks where attackers run thermal-intensive workloads to rapidly generate a large amount of heat, forcing the victim servers into a high temperature, which can potentially cause hardware damage or even server shutdown. Wang et al.[47] analyzed data center hardware failure reports over four years, and found that failure rates are higher in some rack positions. This is because hotspots are unavoidable and difficult to be cooled down, due to the non-linearity of airflow.

In particular, attackers can exploit side/covert channels to further assist mounting their power/thermal attacks. Adversaries can place different types of sensors (e.g., cooling fan sound and some frequency components of the power distribution unit) in the servers to estimate the power consumption[29][30][31], and leveraging benign workloads in the background to amplify their attacks. Adversaries can also utilize covert channels to verify co-residence, which is helpful to improve the effectiveness of both thermal and power attacks[14][53].

Warehouse-scale computing (WSC) refers to a data center in a single organization that owns all equipment. It allows clients to access resources through cloud computing, but its physical access is restrictive. Thus, placing sensors in the WSC data center is not realistic. We focus on the AWS data centers, which are WSC data centers. AWS has strict rules for physical access to its data centers[5][6]. However, leaking information from operation engineers and other sources is possible. For example, the unreliable source Wikileaks published the internal information about AWS data centers[11][48]. Because FPGAs are close to the physical layer, they can be applied for power/thermal attacks. The power hammer shutting down the individual FPGA is a type of power attack[33]. For the large scale, FPGAs are potential tools to capture information leakage in the physical layer for evaluating the physical condition of a data center.

III. Time-Digital Converter

The time-digital converter (TDC) reflects the delay change in FPGA. The TDC layout is shown in Figure 2. The input signal does not reach the last output within a clock cycle so that registers can capture how far the signal propagates. The components of TDC include look-up tables (LUT), CARRY8, latches, and flip-flops. It requires manual placement to fix these components to the designed locations. The input source usually is a clock signal. It starts from LUTs, which are buffers, then the signal is passed to CARRY8s and latches, and then its propagation distance is stored in flip-flops (registers)[23].

The TDC can capture the voltage changes and transient responses[19][20][56]. It is useful to test delay changes in various settings and detect voltage attacks from the power system. The previous works demonstrate that TDC has high accuracy for power analysis of the AES algorithm, the BNN accelerator, and versatile tensor accelerator[18][36][39][43].

The TDC is affected by PVT variations, and we use the temperature effect to deduce the temperature change. The process variation and voltage effect in AWS FPGAs are detailed in Section III-D and Section III-F, respectively. The temperature effect is demonstrated in Section III-G and Section IV.

A. Temporal Average and Spatial Average

The resolution and precision of the value in a single run of a TDC is not high enough because it is affected by noise from parasitic elements, process variation, and thermal difference. The previous research uses a temporal averaging technique to improve the resolution and precision[18][36] [39][43]. We introduce the spatial average technique to tackle the noise. The layout of TDCs is shown in Figure 3. We define temporal average and spatial average as follows:

- Temporal average is the average of data that are collected in different time instances.

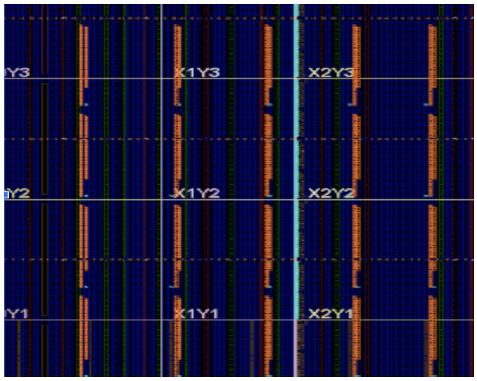

Fig. 3: An example layout of 20 TDCs (Other FPGA designs are not shown)

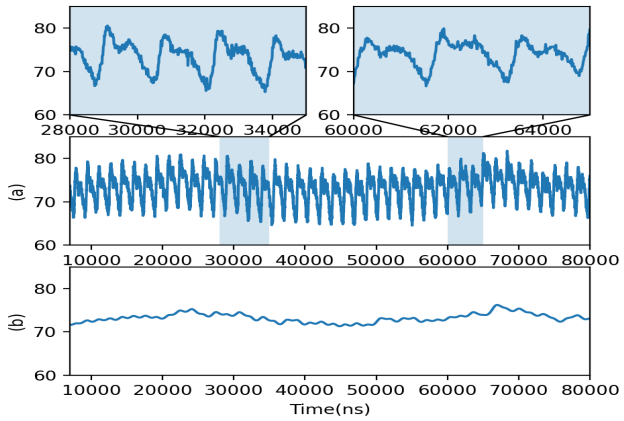

Fig. 4: (a) An averaged TDC signal (b) The low-pass filter output

- Spatial average is the average of data that are collected in different positions in the FPGA.

The temporal average technique is useful to observe the FPGA internal activities or power consumption of PDN (power distribution network). Although it improves the precision and resolution, this technique filters out a power ripple with probability. It may be related to the theory of probability sampling, but sampling timing cannot be controlled. Its uncertainty is difficult to be understood and explained theoretically. We prefer techniques from the field of signal processing to filter out the power ripple.

The power ripple and the oscillation from the feedback loop of the voltage regulator generate noises in certain frequency components. The spatial average technique detects switching frequencies from the power delivery system. It has high-resolution and precise outputs, so that the power ripple can be recovered and removed with a low pass filter.

B. Amazon EC2 F1 Instance Implementation

AWS does not allow users to upload and run their FPGA images directly. Users have to use AWS FPGA toolkits that are available in GitHub[4]. AWS provides three types of F1 instances. They are F1 2xlarge, F1 4xlarge and F1 16xlarge. The FPGA chip is Xilinx Virtex UltraScale+ VU9P[3]. The CPU of the F1 instance is Intel Xeon E5-2686. We cannot find further information about the FPGA board in the AWS, voltage regulators, and power supply. It seems that AWS designs a custom board for that FPGA chip. Bittware designs the XUP-P3R board for the VU9P FPGA[9], and we believe the F1 instance uses a similar board.

AWS allows synthesizing the FPGA design with constraint files (XDC files)[49]. We use up to 20 TDCs and set up the layout as shown in Figure 3. We make TDCs manually using Xilinx Vivado[50], and it generates a XDC file so that we can use it for synthesis with AWS toolkits[4].

We add the TDCs implementation to the `cl_hello_world` example in the AWS toolkit to reuse the interface. The data from TDCs is stored in the BRAM of FPGA then transferred to the software. We debug and develop our design based on the `cl_hello_world` example. We adjust an appropriate number of LUTs and CARRY8s for the TDC to make the signal propagation passing the last LUT but not reaching the last bit output of the last CARRY8. Otherwise, it is not usable. The place and route algorithm may be run multiple times for an acceptable FPGA image.

It is costly to make all TDCs in our design usable, due to the inconsistent clock routing and process variation. The details about the process variation of AWS FPGA are discussed in Section III-D. Suppose we attempt to make all TDCs usable, we need to adjust individual TDCs to synthesize multiple times for an acceptable routing solution, and the solution has to be dedicated to a physical FPGA. After we stop and restart the F1 instance, the physical FPGA can be changed. If some TDCs become unusable, the FPGA design has to be re-synthesized.

Making all TDCs usable is unnecessary because we find 15 usable TDCs are sufficient for this work. Thus, we make sure that at least 15 out of 20 TDCs are usable before starting the data collection. Each TDC output is stored in a separate file, and the average of them is computed with a software program. Because the average of TDC outputs can recover the switching frequency precisely, the low-pass filter can remove the noise caused by the switching frequency better. However, if anomalous data is found, we can check individual TDC output for reasoning.

C. Noise from Power Delivery System

The large scale FPGA in the cloud connects to the power supply directly, and these power supplies and the voltage regulators in the board are likely switch-mode. They generate and regulate power with the switch. Thus, we detect these switching frequencies with the spatial average of TDCs. Figure 4 shows the averaged signal of TDCs. It is more precise than signals from individual TDCs shown in Figure 5.

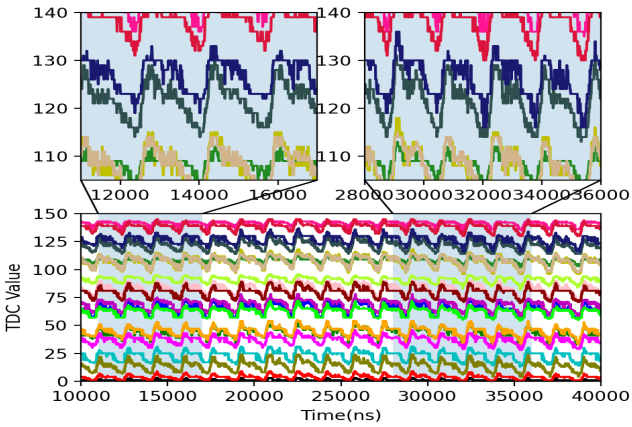

Fig. 5: Individual TDC signals correspond to the averaged signal in Figure 4

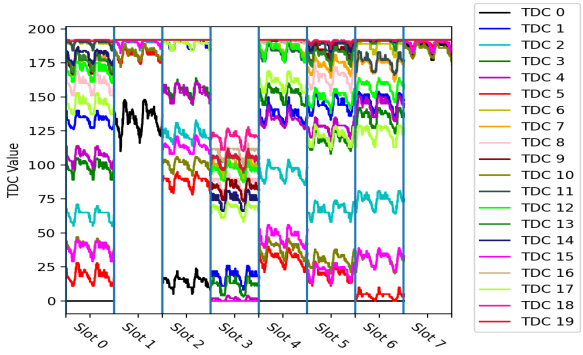

Fig. 6: Individual TDC signals of a image in different FPGAs of the f1.16xlarge instance

D. Process Variation

Process variation is a vital topic in the semiconductor process. It is the backbone to implement the physical unclonable function (PUF) and is also useful for fingerprinting. Because the AWS FPGA is manufactured with the advanced process, its process variation greatly influences the signal propagation of our TDCs. To show the process variation of FPGAs in AWS, we use the same image for different FPGAs. Figure 6 shows that individual TDCs in the same binary behave differently for FPGAs in different slots of the F1 16xlarge instance.

E. Analysis of Resolution and Precision

The number of usable TDCs is N, and there are M bits of TDC output. $N \times M$ is the resolution. The improvement of resolution per additional usable TDC is M. The smallest averaged TDC output value difference is $\frac{1}{N}$. Multiple TDCs can sense thermal information uniformly, and the impacts from the process variation, parasitic capacitance, and parasitic inductance are minimized.

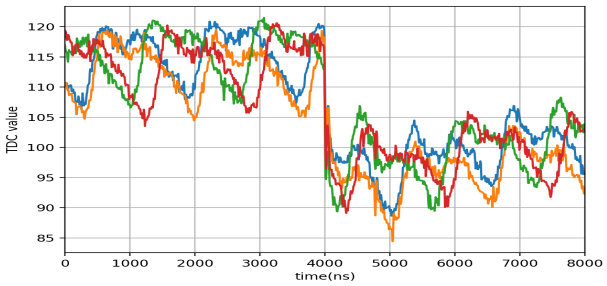

Fig. 7: Step Response

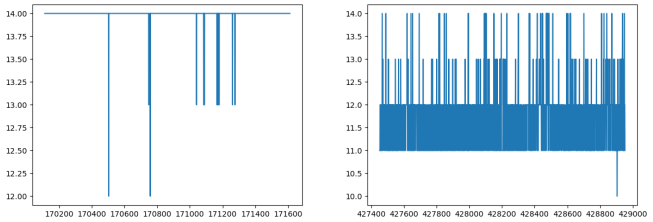

(a) Before heat up (b) After heat up

Fig. 8: x-axis:time step (10 ns), y-axis: TDC output

F. Transient Response

Transient response is an important way to test the power delivery system. Stressors are used to draw a large current to trigger step response. It could be any component of the FPGA as long as it can cause a voltage drop. Previous works show LUT ring oscillators, flip-flops, and programmable interconnect points (PIP)[15][19][20] [56] can be used as stressors. We use chains of LUTs as stressors. The switch of a number of LUTs in the AWS FPGA leads to the step response, and the undershoot and overshoot are visible in Figure 7. Temporal average also can recover them, but the information of switching frequency is lost in temporal average.

G. Local Experiments for Validity

Since we do not have any physical access to measure the temperatures inside AWS data centers as the ground truth, we conduct local experiments for proving the validity of our approach. We implement a TDC in the NEXYS A7 FPGA board and measure the temperature effect on the TDC output. We use the hairdryer to generate hot air and heat up the TDC for 1 minute. The results are shown in Figures 8a and 8b. The TDC outputs drop because the high temperature increases the delay, showing the effectiveness of our approach, which is applied to the cloud FPGA environment.

IV. Temperature Side Channel

Our precise FPGA side channel implemented with TDCs enables the collection of temperature information in real-time. We also compare the temperature measured by our TDC side channel with the DRAM side channel[51]. The principle of

the DRAM temperature side channel is that when temperature increases, it accelerates the decays of DRAM cells that hold 1s (The decays of DRAM cells vary due to the manufacture variation). In other words, we can deduce the temperature changes by counting the number of bits that flip from 1 to 0 in a fixed duration. We adopt Tian et al.'s[45] approach to implement the DRAM side channel in the AWS FPGA. We turn off the ECC and scrubber of DRAM by modifying the `ddr4_core_ddr4` module and `cl_dram_dma` example. The `cl_hello_world` example does not use the DRAM, so the DRAM stops refreshing when it is used. We use a modified version of the `cl_hello_world` example to avoid refreshing. The first step to implement the DRAM side channel is writing 1s to the selected DRAM region with the modified `cl_dram_dma` example. The second step is replacing the `cl_dram_dma` image with the modified `cl_hello_world` image so that the selected region will decay. Finally, after a certain time, we load the modified `cl_dram_dma` image back and read the selected region, and we can count the flipped bits to learn the temperature change.

Collecting temperature data in AWS data centers allows us to take a glance at temperature changes inside the data center. Adversaries can also exploit our FPGA side channel as thermal sensors to obtain temperature information for WSC data centers, and thus find optimal timings to mount power/thermal attacks.

A. Global-Scale Free Cooling

The free cooling technology has become the trend and supports various climates[37]. However, the local temperature is a physical constraint. In some cities with hot climates, the initial investment of free cooling can be high and its overall saving is small.

Moreover, inaccurate weather forecast and unexpected local high temperature deteriorate the complexity of workload management. In addition to seasonal/daily temperature changes, extreme weathers also should be considered. As global warming heats the earth, a data center should make sure its cooling system can maintain an appropriate temperature. The extremely low temperature should also be paid attention to because water is often used as a heat medium in the cooling system, and some computing equipment cannot run at a low temperature.

Since free cooling increases the predictability of cooling system, previous research on power/thermal attack strategies on a single data center can be extended to multiple data centers. Adversaries can use the FPGA side channel to verify that the data center uses free cooling. Then, they can choose a time when the outside temperature is high to attack.

B. Locations of Data Centers

Table I shows approximate locations of data centers in regions that support F1 instance based on the information of baxtel.com[8]. The locations of data centers are important because electricity price, data center market, weather, natural disaster, etc., are factors that should be considered to choose

TABLE I: AWS Data Centers Locations[8]

Region	***Approximate Location***	***Availability Zones***	***Time Offset (Sep. Oct.)***
Virginia (US)	Ashburn	6	UTC-4
Oregon (US)	Umatilla	4	UTC-7
Sydney (Asia)	Sydney	3	UTC+10
London (EU)	Didcot	3	UTC+1
Frankfurt (EU)	Frankfurt	3	UTC+2
Ireland (EU)	Dublin	3	UTC+1

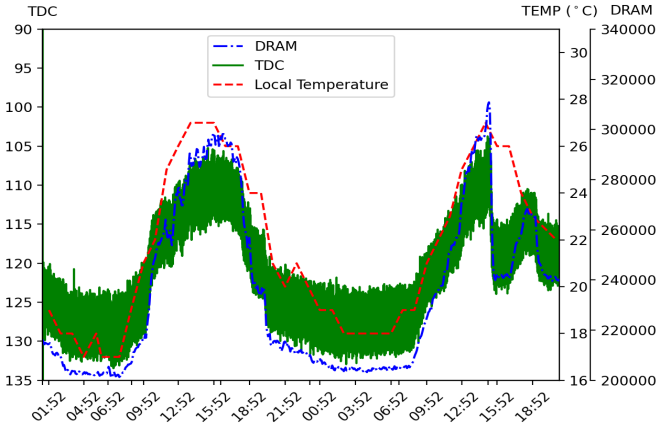

Fig. 9: The data collected with an F1 2xlarge instance in the subnet a, Northern Virginia region (local time: UTC - 4).

the location for a data center. Although addresses of AWS data centers can be found, we do not know how each availability zone corresponds to the data center. There are multiple AWS buildings in each region, and they are close to each other.

We need the local weather information to learn whether a data center uses free cooling or not. The local weather information of each region is based on its approximate locations shown in Table I. Furthermore, the locations of availability zones can be better approximated with latency shown in Table II. To know the physical location of each availability zone, one method is to use the FPGA-based temperature side channel as a receiver and the outside air as the transmitter to construct a covert channel because free cooling uses the outside air. Besides the temperature covert channel, the electromagnetic covert channel is also useful to localize the data center[41]. However, the city-level temperature is sufficient for investigating the free cooling scenario.

C. DRAM Side Channel versus FPGA Side Channel

Each F1 instance in AWS provides FPGA and DRAM that the FPGA image can access. We use them to implement temperature side channels. The DRAM side channel and TDC side channel are implemented with entirely different mechanisms. After the DRAM refresh is turned off, the time of memory retention for each bit depends on temperature. We write 1s to selected bits in DRAM, and then we read the selected bits after a certain amount of time to count the number of flipped bits[45][51]. If the temperature increases,

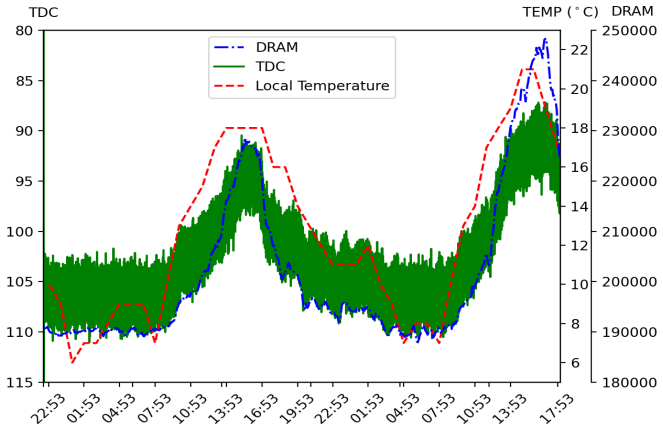

Fig. 10: The data collected with an F1 2xlarge instance in the subnet a, Oregon region (local time: UTC - 7).

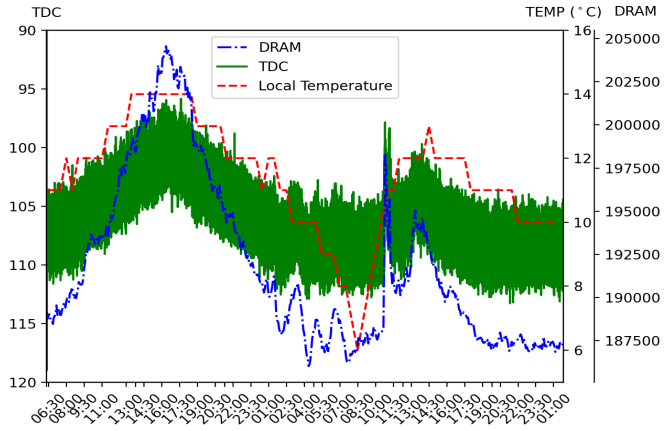

Fig. 12: The data collected with an F1 2xlarge instance in the subnet a, Ireland region (local time: UTC + 1).

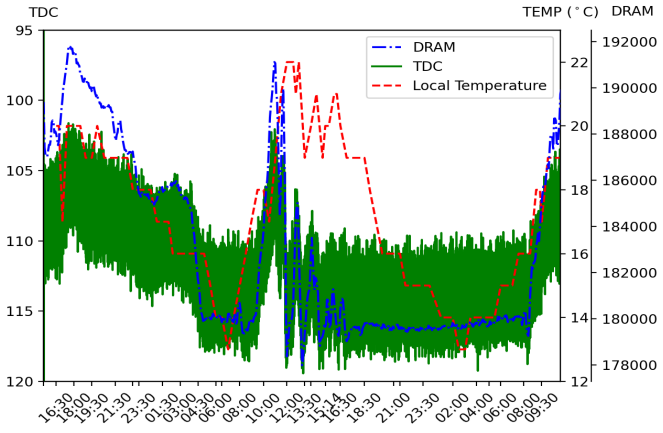

Fig. 11: The data collected with an F1 2xlarge instance in the subnet a, Sydney region (local time: UTC + 10).

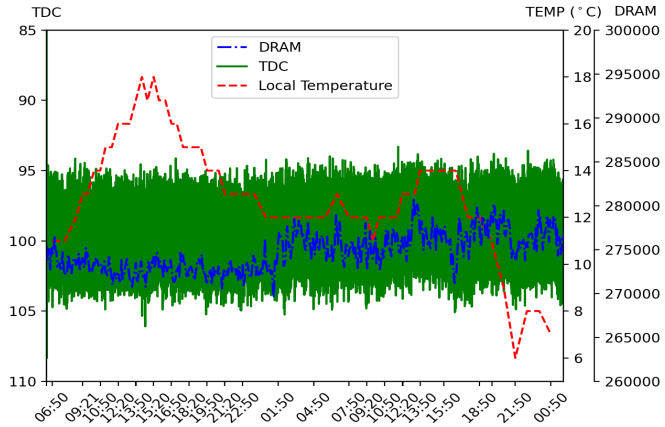

Fig. 13: The data collected with an F1 2xlarge instance in the subnet a, London region (local time: UTC + 1).

the number of flipped bits increases. Thus, we can deduce the temperature change. The TDC is implemented to measure the delay of signal propagation. It is well-known that the temperature affects the signal propagation in the digital chip. We observe that when the number of flipped bits increases, the TDC output decreases. Their temperature patterns are mostly consistent, as illustrated in Figures 9-21. AWS FPGA does not have the temperature inversion phenomenon. This phenomenon happens if the supply voltage in the digital chip is reduced[38][58]. The FPGA chip and DRAM are placed closely to each other, and it is worth noting that a large FPGA is equipped with a cooling fan. Its cooling functionality is slightly better than DRAM.

Accuracy Analysis. When the temperature increases, the signal delay in the FPGA increases close to linear, and according to Xiong et al.[51], the number of flipped bits in the DRAM increases superlinearly. Both the TDC side channel and DRAM side channel are non-linear functions on temperature. Because the temperature range is small (about 10 degrees), they can be considered approximately linear. We collect temperature data with both side channels simultaneously for double-checking.

Data Collection. To automate the data collection process, we write a script in Bash. It runs TDCs iteratively, and the DRAM side channel is run once in every ten iterations. The TDC data is stored in RAMs temporarily, and then the data is read with the interface and stored in a file. The spatial average of TDC values is also computed and stored. For the DRAM side channel, the script waits 20 seconds for the decay of selected bits. Then the number of flipped bits is counted and stored. We use the `date` command to record the data collection time in UTC (coordinated universal time) format. For the analysis of data, the average TDC values are concatenated, and then the high-frequency noise is removed with a low-pass filter for the data plot. The TDC values, numbers of DRAM flipped bits, and local temperatures are aligned and plotted based on the UTC from the `date` command. To show the local weather intuitively, we use the local time in Figures 9-21.

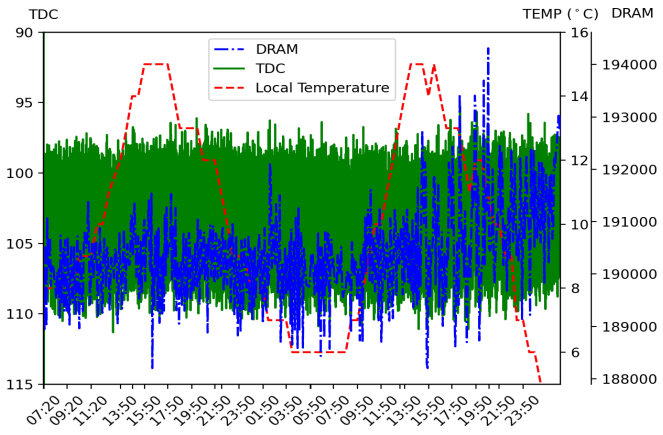

Fig. 14: The data collected with an F1 2xlarge instance in the subnet b, Frankfurt region (local time: UTC + 2).

D. Temperatures of Regions

In our AWS account, there are 21 regions to choose from. However, only Northern Virginia, Oregon, Ireland, Sydney, Frankfurt, and London regions support F1 instances. We believe the cooling systems of data centers in these regions are applicable for the rest of the regions. To analyze how local weather affects the temperature of a data center, we use the weather data in the website timeanddate.com[46]. It provides hour-by-hour past and forecast weather data in all cities around the world. Using the localization methods discussed in Section IV-B, we obtain the approximate local weather for data centers in each region. We launch an F1 2xlarge instance in the subnet a of every region that supports the F1 instance for measurement except the Frankfurt region. Only subnet b in the Frankfurt region supports F1 instances, and so we have to launch F1 instances in subnet b. The measurement in all regions starts approximately at 5am on October 14th, 2021 (UTC). The results from different regions are shown in Figures 9-14.

For F1 instances in Northern Virginia, Oregon, Ireland, and Sydney regions, the patterns of their DRAM side channel and TDC output shown in Figures 9-12 are consistent. Generally, their temperature patterns match the temperature changes of local weathers, and so we believe data centers in these regions adopt free cooling techniques. When the computing equipment heats up the data center, or/and the outdoor temperature in the daytime is hot, the cooling system cools down the computing equipment. The temperature drop caused by the cooling system is demonstrated in Figure 9. The temperature side channel in the Northern Virginia region shows a sudden temperature drop at about local time 3pm on October 15th, 2021. We can see similar phenomena in the Sydney and Ireland regions (Figures 11-12).

The timing that the cooling system acts is interesting. In Figure 12, we can observe a sudden temperature drop at 11am on October 15th, whereas there is no drop on October 14th. It is noticeable that the drop starts when the outside temperature is 11 °C. Similarly, for Figure 11, the temperature drops at about 11am on October 15th, when the outside temperature is 19 °C. We can see that the outside temperature is not the only factor for the operation of the cooling system, and there are probably temperature sensors in different locations inside the data center to measure the temperature of computing equipment more closely. Thus, the high-temperature spot inside the data center can trigger the cooling system to cool it down. The F1 instance in the Sydney region (Figure 11) shows that its temperature is constant during the nighttime on October 15th. We believe that the data center has a mechanism to maintain the constant temperature for the FPGA card for some duration.

For London region and Frankfurt region, temperatures of F1 instances are shown in Figures 13-14. The patterns of the DRAM side channel and TDC output demonstrate that the temperature changes are minimal. They are independent of the local temperatures, indicating that the data centers of these two regions do not use free cooling during our measurement.

Northern Virginia, Oregon, Ireland, and Sydney are in different time zones, and their local temperatures peak at different times. The computational tasks that require high computation power and less latency constraint should avoid the peak local temperature or be run in a different data center where the cool outside air and computing resources are available. However, current regions that support F1 instances are insufficient to allow sophisticated global resource allocation. It needs not only the weather forecast but also workload coordination among data centers. Xu et al.[52] formulated this problem and developed an algorithm to address it. AWS or other public clouds do not have much flexibility to distribute workloads, since cloud users usually request a service in a specific area. However, the price can be leveraged by cloud service providers to motivate users to use cloud resources with lower costs.

E. Temperatures of Availability Zones in Northen Virginia

TABLE II: Minimum Latencies among Availability Zones in Northern Virginia Region (Unit:ms)

src \ dst	a	b	c	d	e
Subnet a	0.237	0.567	0.656	1.278	0.341
Subnet b	0.570	0.072	0.413	0.433	0.630
Subnet c	0.570	0.072	0.413	0.433	0.630
Subnet d	1.272	0.438	0.346	0.139	0.656
Subnet e	0.348	0.620	0.613	0.659	0.109

There is a large data center market in Northern Virginia, and it continues boosting. AWS has invested many resources in this market, and so we study AWS data centers in the Northern Virginia region. AWS makes multiple availability zones in each region, and each availability zone guarantees that it has independent power and networking infrastructures[7]. More flexible choice of availability zone provides better latency service for cloud users. Power outages of data centers are not uncommon, and it happens every year[27]. Having multiple data centers can avoid the single point of failure from the

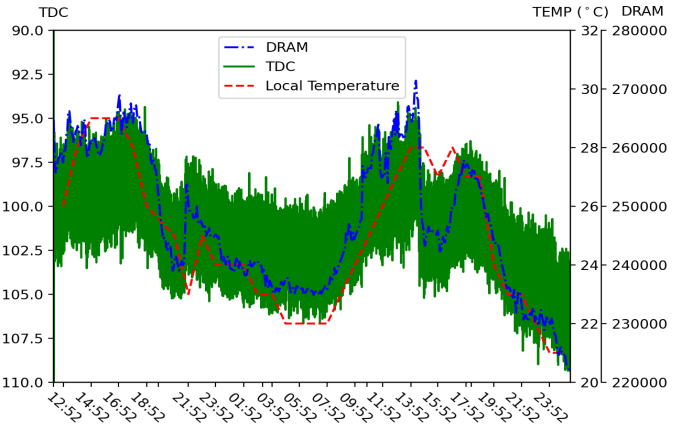

Fig. 15: The data collection with an F1 2xlarge instance in the subnet a, Northern Virginia region starts at 12pm on September 18th, 2021.

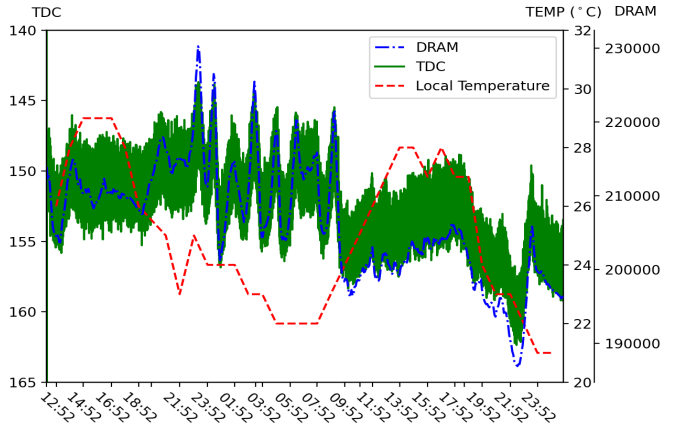

Fig. 17: The data collected with an F1 2xlarge instance in the subnet c, Northern Virginia region.

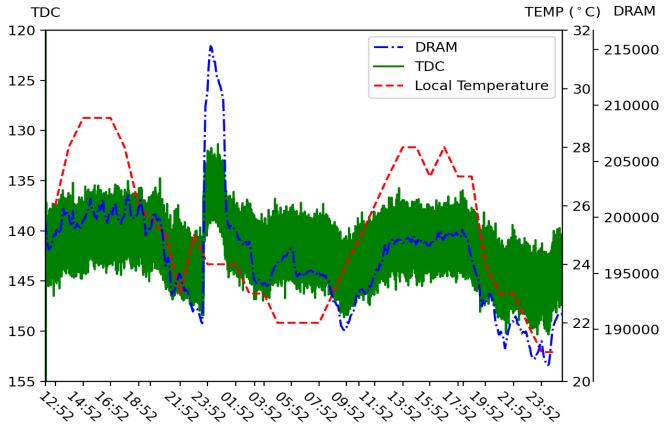

Fig. 16: The data collection with an F1 2xlarge instance in the subnet b, Northern Virginia region.

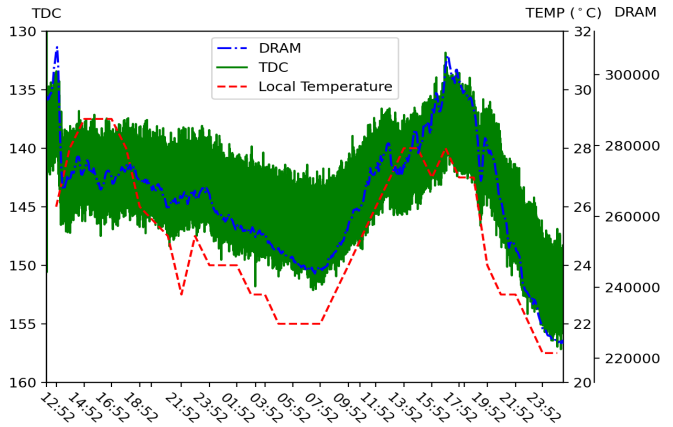

Fig. 18: The data collected with an F1 2xlarge instance in the subnet d, Northern Virginia region.

viewpoint of security engineering. We can find the locations of AWS data centers in baxtel.com[8], but we do not know the physical location of each availability zone. We use latency to estimate the physical distances among data centers. Moreover, latency is a metric to estimate the distances among physical machines[53]. Minimum latency among subnets shown in Table II are measured with the Linux command `ping`. From Table II, we can approximately know their relative distances among availability zones.

We measure the temperature of all subnets in the Northern Virginia region except subnet f because it does not support F1 instance so far. The measurement starts at around 4pm on September 18th, 2021 (UTC) or local Virginia time 12pm on September 18th. The local temperature is based on the weather in Ashburn, VA. Based on our data (Figures 15-19, we can verify that temperatures for F1 instances in different availability zones simultaneously synchronize well with the local temperature changes, except that at some occasions the temperatures rise due to high workloads. Therefore, we believe that these availability zones supporting F1 instances in Northern Virginia region adopt free cooling.

Since the data of Figures 15-19 are collected from different data centers, their operations of the cooling system and the heats from computing equipment are independent. From Figures 16-17, subnets b and c show high temperature periods during nighttime. It is due to the heat generated by their computing equipment, but we cannot determine if they could apply to the overall utilization of a data center. Each rise and drop of temperature lasts about 30 minutes. Since the efficiency of the data center is important, the data center always tries to improve its utilization and minimize idle time. Moreover, Figures 16-17 also demonstrate that the FPGA instance can be heated up higher than the peak temperature of the day. Figure 15 shows a 30-minute duration temperature drop at about 2:20 pm on September 19th even when the local temperature peaks. It shows that at some scenarios when too much heat generated by computing equipment, the cooling system has to work without using free air for quickly lowering the temperature. By contrast, Figures 16 - 19 do not show a

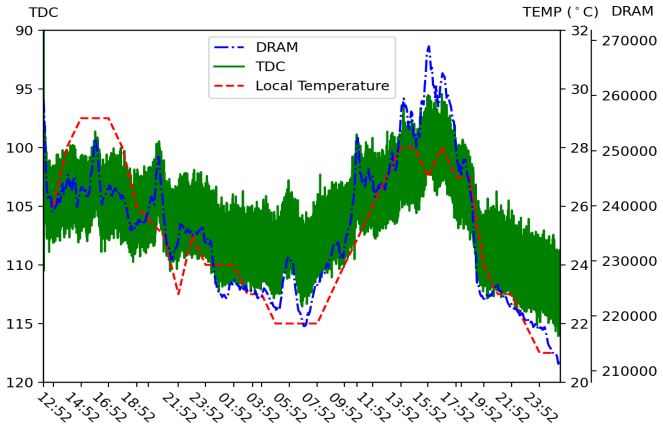

Fig. 19: The data collected with an F1 2xlarge instance in the subnet e, Northern Virginia region.

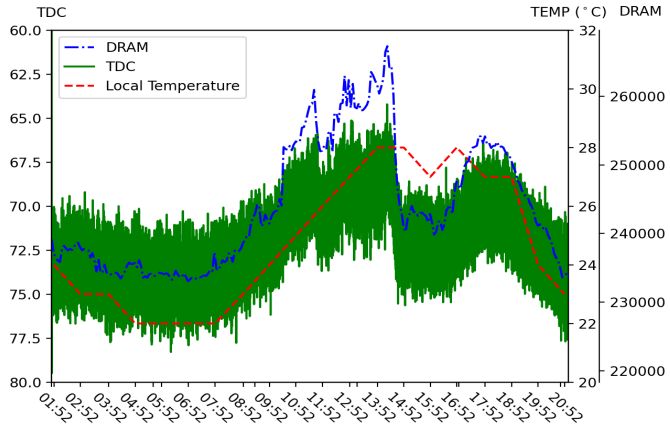

Fig. 21: The data collected with the additional F1 2xlarge instance 4 in the subnet a, Northern Virginia region.

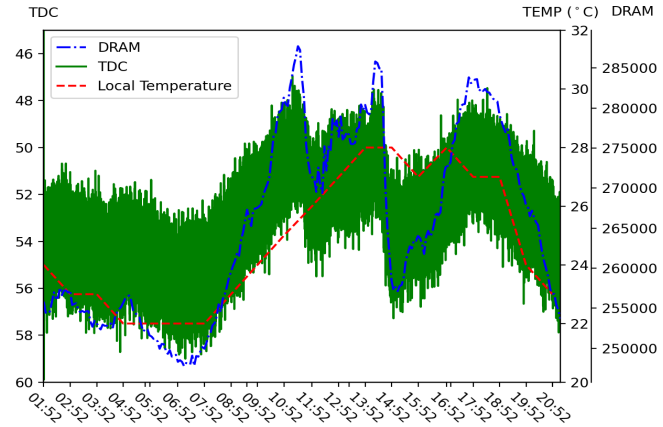

Fig. 20: The data collected with the additional F1 2xlarge instance 2 in the subnet a, Northern Virginia region.

sharp temperature drop at that time even the local temperature reaches peak (28 °C).

F. Temperatures in the Availability Zone A

To observe what types of information are useful for mounting data center level attacks, we collect data from 4 more instances in the subnet a of the Northern Virginia region. This measurement starts at local time 1:45am on September 19th, 2021. The F1 instance in Figure 15 is the instance 1, and the instance 2 and instance 4 at subnet a are shown in Figures 20 and 21, respectively. The temperature patterns of instance 3 and instance 5 are similar to these two instances, and thus their data are omitted.

Although AWS experienced a number of outages and problems, it claimed that outages never caused a loss of the entire data center[12]. AWS also revealed information about the relationship between the data center and the availability zone. It stated that "In every zone, you have at least one data center." In other words, it did not confirm two instances in the same availability zone are in the same data center, and it seems that AWS does a great job of isolating the fault effects.

However, from Figures 20 and 21, the overall temperature patterns of the two instances are similar. The effects of the local weather and cooling are identical, implying the physical locations of the two instances are possibly close to each other. Both temperatures dropped at about 2:40pm. Figure 15 also captures the temperature drop in the subnet a at that time. The subtle differences among them are noticeable. We can observe that temperature peaks at about 11:00am in Figure 20, which is comparable to the temperature when the local temperature peaks. Tian et al.[42] showed the probability that FPGAs share the same NUMA node is decent. Based on their similar temperature patterns, we speculate that these four F1 instances are in the same data center.

V. Discussion

In Section IV-E, we show all subnets a-e adopt free cooling. However, the climate in Northern Virginia does not always provide cool air. Since our measurements were conducted in the fall, these subnets simultaneously use the outside air of the same area for the cooling purposes. However, the usage of free cooling is highly dependent upon the physical location of a data center and its local weather. If the local weather is hot, all data centers in these subnets should avoid the usage of free cooling. Therefore, it is worth to pay attention to the physical location and current weather of data centers for studying their cooling systems and daily operations.

Since we do not have physical access to AWS data centers, we cannot know the exact physical locations of FPGAs and their relative relations with other computing equipment. However, from the perspective of an adversary, the information sources do not have to be reliable and accurate. In Section IV-F, we show similar temperature patterns of four F1 instances in subnet a, and they are likely in the same data center. The adversary can estimate the physical locations of FPGAs by observing their temperature pattern to launch more effective power/thermal attacks.

VI. Conclusion

In this paper, we conduct a measurement study on the cooling systems of AWS data centers through FPGA-based temperature side channels. We leverage TDCs to observe the process variation, step response, and temperature effect of AWS FPGAs. We also implement the DRAM temperature side channel. We use the DRAM side channel and TDCs to collect the temperature information leakage. We select 10 availability zones (10 or more AWS data centers) for the measurement study. Based on the collected data, we analyze the cooling system of AWS data centers and identify many of them adopt free cooling for the reduction of cooling cost. Our study not only reveals the temperature information leakage in data centers but also investigates the data center cooling system from security and privacy perspectives.

Acknowledgements

We are grateful to the anonymous reviewers for their insightful feedback. This work was supported in part by ARO grant W911NF-19-1-0049, NSF grant CNS-1815650, and the Commonwealth Cyber Initiative.

References

[1] Luiz André Barroso, Urs Hölzle, and Parthasarathy Ranganathan. *Data Center Basics: Building, Power, and Cooling*.

[2] AWS. *Amazon EC2 F1 Instance Expands to More Regions, Adds New Features, and Improves Development Tools*. https://aws.amazon.com/about-aws/whats-new/2018/10/.

[3] AWS. *Amazon EC2 Instance Types*. https://aws.amazon.com/ec2/instance-types/.

[4] AWS. *AWS/AWS-FPGA: Official Repository of the AWS EC2 FPGA hardware and software development kit*. https://github.com/aws/aws-fpga.

[5] AWS. *Our Controls*. https://aws.amazon.com/compliance/data-center/controls/.

[6] AWS. *Our Data Centers*. https://aws.amazon.com/compliance/data-center/data-centers/.

[7] AWS. *Regions and Availability Zones*. https://aws.amazon.com/about-aws/global-infrastructure/regions_az/.

[8] *Baxtel*. https://baxtel.com.

[9] BittWare. https://www.bittware.com/fpga/xup-p3r/.

[10] Jun Dai, Michael M. Ohadi, Diganta Das, and Michael G. Pecht. *Optimum Cooling of Data Centers*. 2014.

[11] Datacenterdynamics. *WikiLeaks publishes list of AWS data center locations, colo providers*. https://www.datacenterdynamics.com/en/news/wikileaks-publishes-list-aws-data-center-locations-colo-providers/.

[12] Datacenterknowledge. *AWS Says It's Never Seen a Whole Data Center Go Down*. https://www.datacenterknowledge.com/amazon/aws-says-it-s-never-seen-whole-data-center-go-down.

[13] Miyuru Dayarathna, Yonggang Wen, and Rui Fan. "Data Center Energy Consumption Modeling: A Survey". In: *IEEE Communications Surveys Tutorials* (2016).

[14] Xing Gao, Zhang Xu, Haining Wang, Li Li, and Xiaorui Wang. "Reduced Cooling Redundancy: A New Security Vulnerability in a Hot Data Center". In: *NDSS*. 2018.

[15] Ilias Giechaskiel, Kasper Bonne Rasmussen, and Jakub Szefer. "C3APSULe: Cross-FPGA Covert-Channel Attacks through Power Supply Unit Leakage". In: *IEEE Symposium on Security and Privacy*. 2020.

[16] Ilias Giechaskiel, Shanquan Tian, and Jakub Szefer. "Cross-VM Covert-and Side-Channel Attacks in Cloud FPGAs". In: *ACM TRETS* (2022).

[17] Ilias Giechaskiel, Shanquan Tian, and Jakub Szefer. "Cross-VM Information Leaks in FPGA-Accelerated Cloud Environments". In: *IEEE HOST*. 2021.

[18] Ognjen Glamočanin, Louis Coulon, Francesco Regazzoni, and Mirjana Stojilović. "Are Cloud FPGAs Really Vulnerable to Power Analysis Attacks?" In: *DATE*. 2020.

[19] Dennis R. E. Gnad, Fabian Oboril, Saman Kiamehr, and Mehdi B. Tahoori. "An Experimental Evaluation and Analysis of Transient Voltage Fluctuations in FPGAs". In: *IEEE VLSI* (2018).

[20] Dennis R.E. Gnad, Fabian Oboril, Saman Kiamehr, and Mehdi B. Tahoori. "Analysis of transient voltage fluctuations in FPGAs". In: *FPT*. 2016.

[21] Íñigo Goiri, Thu D. Nguyen, and Ricardo Bianchini. "CoolAir: Temperature- and Variation-Aware Management for Free-Cooled Datacenters". In: *ACM ASPLOS*. 2015.

[22] Google. *Google Data Centers*. https://www.google.com/about/datacenters/.

[23] 7 Series FPGALibraries Guide. https://www.xilinx.com/support/documentation/sw_manuals/xilinx2012_2/ug953-vivado-7series-libraries.pdf.

[24] Mordechai Guri, Matan Monitz, Yisroel Mirski, and Yuval Elovici. "BitWhisper: Covert Signaling Channel between Air-Gapped Computers Using Thermal Manipulations". In: *IEEE CSF*. 2015.

[25] Taras Iakymchuk, Maciej Nikodem, and Krzysztof Kepa. "Temperature-based covert channel in FPGA systems". In: *ReCoSoC*. 2011.

[26] Infotech. *Top 10 Energy-Saving Tips for a Greener Data Center*. http://static.infotech.com/downloads/samples/070411_premium_oo_greendc_top_10.pdf.

[27] Uptime Institute. *Uptime Institute data shows outages are common, costly, and preventable*. https://uptimeinstitute.com/data-center-outages-are-common-costly-and-preventable.

[28] Intel. *The State of Data Center Cooling*. http://www.ceclimited.com/sites/all/themes/creative/state-of-date-center-cooling.pdf.

[29] Mohammad A. Islam and Shaolei Ren. "Ohm's Law in Data Centers: A Voltage Side Channel for Timing Power Attacks". In: *ACM CCS*. 2018.

[30] Mohammad A. Islam, Shaolei Ren, and Adam Wierman. "Exploiting a Thermal Side Channel for Power Attacks in Multi-Tenant Data Centers". In: *ACM CCS*. 2017.

[31] Mohammad A. Islam, Luting Yang, Kiran Ranganath, and Shaolei Ren. "Why Some Like It Loud: Timing Power Attacks in Multi-Tenant Data Centers Using an Acoustic Side Channel". In: *Proc. ACM Meas. Anal. Comput. Syst.* (2018).

[32] Christos Kozyrakis. "Resource efficient computing for warehouse-scale datacenters". In: *DATE*. 2013.

[33] Kaspar Matas, Tuan Minh La, Khoa Dang Pham, and Dirk Koch. "Power-hammering through Glitch Amplification – Attacks and Mitigation". In: *IEEE FCCM*. 2020.

[34] meettechniek.info. *Accuracy, precision & resolution*. https://meettechniek.info/measurement/accuracy.html.

[35] Seyedeh Sharareh Mirzargar and Mirjana Stojilović. "Physical Side-Channel Attacks and Covert Communication on FPGAs: A Survey". In: *ACM FPL*. 2019.

[36] Shayan Moini, Shanquan Tian, Daniel Holcomb, Jakub Szefer, and Russell Tessier. "Remote Power Side-Channel Attacks on BNN Accelerators in FPGAs". In: *DATE*. 2021.

[37] NORTEK. *Free Cooling Concepts for Data Centers*. https://www.nortekair.com/wp-content/uploads/2017/01/Free-Cooling-Concepts-for-Data-Centers.pdf.

[38] Behzad Salami, Erhan Baturay Onural, Ismail Emir Yuksel, Fahrettin Koc, Oguz Ergin, Adrian Cristal Kestelman, Osman Unsal, Hamid Sarbazi-Azad, and Onur Mutlu. "An Experimental Study of Reduced-Voltage Operation in Modern FPGAs for Neural Network Acceleration". In: *IEEE/IFIP DSN*. 2020.

[39] Falk Schellenberg, Dennis R.E. Gnad, Amir Moradi, and Mehdi B. Tahoori. "An inside job: Remote power analysis attacks on FPGAs". In: *DATE*. 2018.

[40] Zhihui Shao, Mohammad A. Islam, and Shaolei Ren. "A First Look at Thermal Attacks in Multi-Tenant Data Centers". In: *SIGMETRICS Perform. Eval. Rev.* (2019).

[41] Cheng Shen, Tian Liu, Jun Huang, and Rui Tan. "When LoRa Meets EMR: Electromagnetic Covert Channels Can Be Super Resilient". In: *2021 IEEE Symposium on Security and Privacy*. 2021.

[42] Shanquan Tian, Ilias Giechaskiel, Wenjie Xiong, and Jakub Szefer. "Cloud FPGA Cartography using PCIe Contention". In: *IEEE FCCM*. 2021.

[43] Shanquan Tian, Shayan Moini, Adam Wolnikowski, Daniel Holcomb, Russell Tessier, and Jakub Szefer. "Remote Power Attacks on the Versatile Tensor Accelerator in Multi-Tenant FPGAs". In: *IEEE FCCM*. 2021.

[44] Shanquan Tian and Jakub Szefer. "Temporal Thermal Covert Channels in Cloud FPGAs". In: *ACM FPGA*. 2019.

[45] Shanquan Tian, Wenjie Xiong, Ilias Giechaskiel, Kasper Rasmussen, and Jakub Szefer. "Fingerprinting Cloud FPGA Infrastructures". In: *ACM FPGA*. 2020.

[46] Timeanddate. https://www.timeanddate.com/.

[47] Guosai Wang, Lifei Zhang, and Wei Xu. "What Can We Learn from Four Years of Data Center Hardware Failures?" In: *IEEE/IFIP DSN*. 2017.

[48] Wikileak. *Amazon Atlas*. https://wikileaks.org/amazon-atlas/.

[49] Xilinx. *Using Constraints*. https://www.xilinx.com/support/documentation/sw_manuals/xilinx2018_1/ug903-vivado-using-constraints.pdf.

[50] Xilinx. *Vivado*. https://www.xilinx.com/support/download/index.html/content/xilinx/en/downloadNav/vivado-design-tools/2020-2.html.

[51] Wenjie Xiong, Nikolaos Athanasios Anagnostopoulos, André Schaller, Stefan Katzenbeisser, and Jakub Szefer. "Spying on Temperature using DRAM". In: *DATE*. 2019.

[52] Hong Xu, Chen Feng, and Baochun Li. "Temperature Aware Workload Management in Geo-distributed Datacenters". In: *ICAC 13*. USENIX Association, 2013.

[53] Zhang Xu, Haining Wang, and Zhenyu Wu. "A Measurement Study on Co-residence Threat inside the Cloud". In: *USENIX Security*. 2015.

[54] Zhang Xu, Haining Wang, Zichen Xu, and Xiaorui Wang. "Power Attack: An Increasing Threat to Data Centers". In: *NDSS*. 2014.

[55] Mark Zhao and G. Edward Suh. "FPGA-Based Remote Power Side-Channel Attacks". In: *IEEE Symposium on Security and Privacy*. 2018.

[56] Kenneth M. Zick, Meeta Srivastav, Wei Zhang, and Matthew French. "Sensing Nanosecond-Scale Voltage Attacks and Natural Transients in FPGAs". In: *ACM FPGA*. 2013.

[57] Daniel Ziener, Florian Baueregger, and Jürgen Teich. "Using the Power Side Channel of FPGAs for Communication". In: *IEEE FCCM*. 2010.

[58] Yazhou Zu, Wei Huang, Indrani Paul, and Vijay Janapa Reddi. "Ti-states: Processor power management in the temperature inversion region". In: *IEEE/ACM MICRO*. 2016.