IOTLB-SC: An Accelerator-Independent Leakage Source in Modern Cloud Systems

Thore Tiemann
t.tiemann@uni-luebeck.de
University of Lübeck
Lübeck, SH, Germany

Zane Weissman
zweissman@wpi.edu
Worcester Polytechnic Institute
Worcester, MA, USA

Thomas Eisenbarth
thomas.eisenbarth@uni-luebeck.de
University of Lübeck
Lübeck, SH, Germany

Berk Sunar
sunar@wpi.edu
Worcester Polytechnic Institute
Worcester, MA, USA

ABSTRACT

Hardware peripherals such as GPUs and FPGAs are commonly available in server-grade computing to accelerate specific compute tasks, from database queries to machine learning. CSPs have integrated these accelerators into their infrastructure and let tenants combine and configure these components flexibly, based on their needs. Securing I/O interfaces is critical to ensure proper isolation between tenants in these highly complex, heterogeneous, yet shared server systems, especially in the cloud, where some peripherals may be under control of a malicious tenant.

In this work, we investigate the interfaces that connect peripheral hardware components to each other and the rest of the system. We show that the I/O memory management units (IOMMUs) — intended to ensure proper isolation of peripherals — are the source of a new attack surface: the I/O translation look-aside buffer (IOTLB). We show that by using an FPGA accelerator card one can gain precise information over IOTLB activity. That information can be used for covert communication between peripherals without bothering CPU or to directly extract leakage from neighboring accelerated compute jobs such as GPU-accelerated databases. We present the first qualitative and quantitative analysis of this newly uncovered attack surface before fine-grained channels become widely viable with the introduction of CXL and PCIe 5.0. In addition, we propose possible countermeasures that software developers, hardware designers, and system administrators can use to suppress the observed side-channel leakages and analyze their implicit costs.

CCS CONCEPTS

• **Security and privacy** → **Systems security**; **Side-channel analysis and countermeasures**; • **Computer systems organization** → *Heterogeneous (hybrid) systems*.

KEYWORDS

cloud, FPGA, side-channel, peripheral, IOMMU

1 INTRODUCTION

Modern server-grade computing infrastructures are becoming more heterogeneous: computational needs are spread over fast and flexible CPUs as well as powerful peripherals such as smart storage, GPUs, smart NICs and FPGAs. Major cloud service providers (CSPs) have started to shift tasks such as networking, memory management and VM management into more specialized hardware peripherals [2, 5, 14], freeing up precious CPU time that is rented to more tenants who share the same hardware. These multi-tenant, peripheral-heavy cloud systems rely on increasingly interlinked memory systems to provide high throughput for shared, scalable and parallelized cloud infrastructure. Technologies like VT-d, DDIO, and CXL allow peripherals to not only directly read and write to the memory of a virtual machine, but to also use a CPU's shared cache to speed up repeated reads and writes.

On a logic layer, input-output memory management units (IOMMUs) enforce memory isolation between these peripherals and guest VMs running on CPUs, making IOMMUs a key component for ensuring security of the cloud infrastructure [6, 28, 37]. The IOMMU ensures that accesses to virtual memory spaces are isolated and appropriately virtualized: e.g., devices may handle only I/O-specific virtual addresses and not the CPU-side virtual addresses or the underlying system's physical addresses; in addition, devices may only access memory with the appropriate permissions set.

However, when many tenants share the same hardware, side effects in these complex shared memory systems weaken the security promises of virtualization that make highly scalable multi-tenant cloud computing possible. These side effects of shared hardware are exploited by microarchitectural attacks, most prominently cache attacks. Cache attacks exploit the measurable difference in access times to the many tiers of modern caches to overcome the sophisticated memory isolation mechanisms that protect tenants' data and computation from each other. Besides cache attacks, which have been successfully applied in commercial cloud settings [25, 38], microarchitectural attacks like Meltdown [34] and related MDS

attacks [10, 56, 62] as well as Rowhammer attacks pose a real threat in shared cloud environments. One malicious tenant may, after successful co-location [54, 68], use these microarchitectural side effects to glean sensitive information from co-located VMs.

While Meltdown and other MDS-style attacks have mostly been patched with microcode updates [39, 40], they often also require CSPs to disable simultaneous multi-threading between separate security domains for full protection. Rowhammer attacks are significantly mitigated by usage of ECC memory and newer DDR4 and DDR5 architectures, even though vulnerability of both older ECC memory modules and newer ECC-free DDR4 memory to Rowhammer has been practically verified [12, 16, 50].

Cache attacks, however, are much more difficult to prevent, as the contention and timing differences that enable attacks such as Prime+Probe are inherent to modern cache architecture [46]. Many system-level solutions like cache partitioning have been proposed [55, 67], but have not been widely implemented in hardware and are costly in performance to implement in firmware or hypervisors. The most common way to prevent these attacks is through constant-time implementation of security-critical code; that is, rather than removing the leakage channel inherent to the cloud architecture, software developers must make sure their code does not leak sensitive data on timing-based channels [3, 8, 9, 65].

While most research in microarchitectural attacks has focused on attacks from core to core on CPUs, caches are no longer only accessible by CPUs. Intel's DDIO technology, present on all recent Intel server architectures, allows high speed peripherals to directly access a CPU's shared cache without interrupting CPU execution [26]. Cloud users may rent peripherals such as purpose-specific GPU or FPGA cloud instances for higher performance in particular workloads. In such heterogeneous compute environments, security is even more challenging, as tenants are no longer confined to virtual machines (VMs) on the CPU, but may additionally have control over peripherals. With CSPs renting instances that grant tenants full access to FPGAs designed specifically for heterogeneous computation [1, 4, 17], it becomes trivial for attackers to gain sufficient control over peripherals in the cloud that are more than capable of exploiting microarchitectural vulnerabilities.

First works have used peripherals like network cards [32] and FPGAs [52, 64] to target CPU caches as a powerful shared resource that is accessible by VMs and peripherals alike. These works indicate that not only are cache attacks mounted from peripherals possible; they can leak information about the private operations of both CPUs *and* peripherals. Furthermore, classical cache attack methods can become more powerful when the attacker controls peripherals in addition to a VM on the same machine. As of now, several other components that are also shared by peripherals such as the IOMMU, which is the main line of defense against compromised peripherals, remain unstudied and may open up new attack surfaces.

Our Contribution. This work exposes a vulnerability in an overlooked attack surface present in multi-tenant, peripheral-heavy cloud systems: the microarchitecture of the I/O Memory Management Unit (IOMMU). Knowing that the IOMMUs in modern CPUs have translation look-aside buffers (IOTLBs) to speed up repeated translations [6, 28, 43], we present a hardware design for an FPGA

acceleration card that uses memory access timing to reliably identify whether or not a translation is present in an IOTLB. With that design, we propose and evaluate an algorithm for IOTLB eviction set finding. With those eviction sets, we demonstrate the first two IOTLB-based covert channels. We use the FPGA to collect side-channel IOTLB traces from two other peripheral devices and analyze the viability and threat models of a full side-channel attack.

We show that the IOTLB is the source of side-channel vulnerability that CSPs are currently not aware of and thus do not protect against. We show that the IOTLB is an excellent source for constructing covert channels between co-located peripherals and can also be abused to extract information from neighboring peripherals such as GPU-accelerated databases. We provide comprehensive threat analysis of this vulnerability, in both the present and the near future, and present viable defenses and countermeasures. In summary, our main contributions are:

- We demonstrate a *previously ignored IOTLB timing side-channel* against PCIe peripherals *before* technologies such as CXL and PCIe 5.0 gain widespread adoption, and fine-grained attacks become viable on a large installation base.
- We develop a new algorithm that finds eviction sets without any prior assumptions of organization and demonstrate its advantages in finding IOTLB eviction sets over a similar eviction set finding algorithm.
- We use a custom FPGA hardware function to exploit the IOTLB timing side-channel and study traces collected from an SQL database acceleration library for a GPU.
- We leak IOTLB timing side-channel traces from a GPU-accelerated SQL database library and analyze the vulnerability of the library to a practical attack.
- We demonstrate the *first* two IOTLB covert channels, including a peripheral-to-peripheral channel with a generic application as the sender and our custom FPGA function as the receiver.
- We propose countermeasures for applications, cloud systems, and IOMMU implementations to counter the side-channel we identified.

2 BACKGROUND

When multiple hardware resources share data, it is often desirable to have direct memory access (DMA) from one resource to another. However, simply allowing any peripheral to read or write a host CPU's memory would be disastrous for security, especially in virtualized environments with multiple users sharing the CPU. AMD's AMD-Vi and Intel's VT-d features (present on both companies' performance desktop and server processors for the better part of a decade) allow for virtualized DMA with IOMMUs that dynamically map and translate virtual addresses used specifically by peripherals to access CPU memory. To speed up repeated access to the same memory location, IOMMUs often include translation look-aside buffers (TLBs, or IOTLBs when they are in IOMMUs) which cache recently translated I/O virtual addresses and their corresponding physical addresses to avoid the slow page-table walks otherwise required for translation. Like CPU caches and TLBs, which perform a similar function for CPU memory accesses and address translations, IOTLBs introduce a timing-based side-channel vulnerability.

2.1 Caches and TLBs

A cache stores data for faster access. A translation look-aside buffer (TLB) is technically just another cache, though rather than caching the data or instructions stored at an address, it caches an address translation. However, throughout this paper we will refer to memory caches as simply "caches". Intel's documentation[27] and several works reverse-engineering cache architectures [23, 29, 36, 45] and TLB architectures [18, 61] reveal that TLBs on modern Intel CPUs are organized very similarly to modern CPU memory caches. Modern TLBs and caches are typically organized into **sets** and **ways**. The number of ways is the number of entries each set can contain. For TLBs, each virtual address is mapped to one set, but can occupy any way within that set. When a set is full, old entries may be evicted to make room for new ones. A set of addresses which reliably causes the eviction of all other entries in a set when accessed is called an **eviction set**. A minimal eviction set contains as many addresses as there are ways in the cache/TLB and therefore fills an entire cache set when accessed [63].

2.2 Side-Channel Attacks

Timing side-channel attacks against the CPU's cache are widely studied and well understood: researchers have crafted several variants [13, 21, 36, 51, 66], used them as part of more complicated microarchitectural attacks [31, 34], and built defenses against them [20, 35, 67]. There are many cache side-channel strategies that work in different memory-sharing scenarios and have quite varied temporal and address resolutions. These are two of the most common and useful attack techniques:

Flush+Reload (F+R) [66] requires shared memory between the attacker and the victim and has three steps: **1)** The attacker flushes the cache line of interest. **2)** She then waits for the victim to execute. Later, **3)** she reloads the flushed line and measures the reload latency. If the latency is low, the cache line was served from the cache hierarchy, so the cache line was accessed by the victim.

Prime+Probe (P+P) does not require shared memory at the cost of a lower temporal resolution than F+R since the attacker checks the status of the cache by probing a whole cache set rather than flushing or reloading a single line. P+P has three steps: **1)** The attacker primes the cache set under surveillance with dummy data by accessing a proper eviction set, **2)** she waits for the victim to execute, **3)** she accesses the eviction set again and measures the access latency (probing). If the latency is above a certain threshold, some parts of the eviction set were evicted by the victim process, meaning that the victim accessed cache lines belonging to the cache set under surveillance [36].

2.3 Attacks on TLBs

In 1995, Silbert et al. remarked in a security analysis of Intel CPU architectures that "all 80x86 [now more commonly called x86] processors have a translation look-aside buffer (TLB) that [...] has potential for use as a covert timing channel" [59]. In 2013, Hund et al. [23] demonstrated that a TLB timing side-channel on then-modern Intel CPUs could reveal if a page was mapped by the operating system even if the user does not have permission to access the page directly. They demonstrated that this exploit could be used to identify the pages used by the kernel, even when the

addresses of the pages were randomized (a common defense against side-channel attacks of many types). In 2017, Gras et al. crafted an attack that uses a cache side-channel to identify TLB evictions. This was a robust attack that can be mounted even from JavaScript to de-randomize kernel pages [19]. Gras et al.'s "TLBleed" in 2018 [18] showed that TLBs in modern Intel CPUs were vulnerable to timing side-channel attacks of the sort that are typically used on CPU memory caches, and can be used for similarly complex attacks: with the help of some machine learning, the TLB side-channels on Skylake, Broadwell, and Coffeelake CPUs can be used to recover a key from an Edward-curve cryptographic function.

2.4 PCIe

Peripheral Component Interconnect Express (PCIe) [47] is the backbone of modern desktop and server systems. While often referred to as a bus, PCIe uses a high-speed point-to-point topology with devices being connected to switches or directly to a root port via serial links. The root complex connects the PCIe network to the CPU and the main memory. On a PCIe network, all devices can send memory requests to each other and to the main memory. An IOMMU can be used to virtualize addresses used by PCIe devices and to implement access restrictions. If supported, each root port of a root complex may define access rules for inter-device communication and implement them in the PCIe switches.

Two recent works [30, 60] describe covert- and side-channel attacks that rely on PCIe bus contention. A preliminary is that the two devices involved share the same PCIe switch. In contrast, our work assumes the two devices to share a PCIe root port. Our assumption is less restrictive as any two PCIe devices sharing a switch share a root port, but devices sharing a root port do not necessarily share a switch, as root ports can have many lanes to support multiple devices without sharing a physical bus [47].

Currently, PCIe 3.0 is the prevailing PCIe specification for commodity hardware. After a short period of CPUs supporting PCIe 4.0, PCIe specification 5.0 is the upcoming standard for the next generations of server-grade CPUs. CPUs supporting PCIe 5.0 are scheduled for November 2022 and January 2023, respectively [7, 58]. PCIe 5.0 doubles transfer rates compared to PCIe 4.0, making the interconnect compete with main memory speeds. As a result, PCIe 5.0 physical layer is also used by a new protocol named Compute eXpress Link (CXL) [57]. CXL supports three sub-protocols: *CXL.io* is based on PCIe and enables CXL devices to share the PCIe infrastructure with PCIe devices unaware of CXL. With *CXL.cache*, devices are enabled to cache data from main memory while maintaining coherency between the main memory, the CPU caches and the accelerator cached copy. *CXL.mem* is used by a host CPU to access CXL device memory and manage its coherent usage.

2.5 IOMMUs

Input-Output Memory Management Units (IOMMUs) are located between PCIe devices and the main memory. Usually, they are implemented as part of the root complex. Modern server systems feature one IOMMU per root port. Similar to MMUs in the CPU, IOMMUs provide address translation and protection for memory regions that are made accessible to PCIe devices [6, 28]. Address

virtualization allows to isolate or virtualize such devices. Also, it allows 32-bit peripherals to use memory regions above 4 GB.

The translation process of the IOMMU works very similar to the process in a CPU's MMU. Modern IOMMUs map PCIe devices to IOMMU groups or domains. The operating system, hypervisor, or VMM maintains a page table with all address mappings per group/domain. The page table is organized in a tree structure. Its depth depends on the width of the I/O virtual addresses (IOVAs) supported by the IOMMU. For IOVAs referencing 4 KB pages, the 12 least significant address bits (page offset) remain untranslated. Accordingly, the 21/30 least significant bits of IOVAs pointing to 2 MB/1 GB pages remain untranslated.

IOVAs are translated to physical addresses (PAs) by the IOMMU performing a page table walk. Since this is quite time consuming, modern IOMMUs feature a translation look-aside buffer called IOTLB. This cache is used to store translated IOVA→PA mappings and is shared by all devices managed by the IOMMU.

2.5.1 Attacks on IOMMUs. In the past, several attacks have been shown that circumvent the IOMMU to gain direct memory access or use the misconfiguration of the IOMMU to exploit device drivers through code injection or control-flow hijacking. However, the root cause always was a misconfigured IOMMU or a software vulnerability. We are not aware of any attacks that were made possible solely by the IOMMU hardware.

For example, a malicious peripheral can bypass the IOMMU by adding appropriate entries to the page table on startup before the IOMMU is activated by the BIOS [41, 42], or by exploiting PCIe address translation services (ATS), which allows a peripheral to mark any memory request as "translated" and bypass IOMMU translation and isolation [37]. Malicious devices may also exploit vulnerabilities in the kernel or device drivers. IOMMU address translation only works on a page-granular level, so memory that was never intended to be shared might be allocated to a shared page, leaking secret data or enabling code injection attacks that can compromise the whole system [37].

3 IDENTIFYING IOTLB SIDE-CHANNELS

In this section, we demonstrate two fundamental techniques for implementing IOTLB side-channel attacks on these or similar systems. We measure the latency difference between DMA accesses to addresses with cached and uncached translations in the IOMMU. We also demonstrate a new algorithm for reliably finding IOTLB eviction sets with no prior assumptions about size or organization. We have access to three different system setups that we will investigate throughout this work. Table 1 summarizes the key features of each. A detailed description of the setups is given next.

3.1 System Setup

For our experiments, we rely on three systems that are representative of modern cloud services featuring FPGA resources. The systems feature recent server-grade CPUs as well as FPGA extension cards based on Intel FPGAs. The FPGAs are managed by the Intel Acceleration Stack (IAS) which is designed to ease management of cloud deployments. The first system, *a10l*, is a system we have physical and administrative access to. The other two systems *a10v* and *s10v* are cloud-like systems that are accessible through

Table 1: Overview of the system setups used in this work

Name	*a10l*	*a10v*	*s10v*
CPU	2 Xeon Silver 4114	2 Xeon Platinum 8180	2 Xeon Platinum 8280
#PCIe RP	4 per CPU	4 per CPU	4 per CPU
#IOMMUs	4 per CPU	4 per CPU	4 per CPU
FPGA PAC	Arria 10	Arria 10	Stratix 10
OPAE ver.	1.1.2-1	2020-01-01	2020-01-01
Bitstream ver.	1.2.3	1.1.3	2.0.3
Root/phys. access	yes	no	no

the Intel Labs (IL) Academic Compute Environment (ACE)[1]. We operate the two IL ACE systems with user privileges only. This is why we evaluate our eviction set finding algorithm on all three systems but rely solely on *a10l* for the side- and covert channel experiments. More detailed information about the different systems is given in Table 1 and in the following paragraphs.

a10l: As our local setup, we use a Dell PowerEdge R740 server with two Intel Xeon Silver 4114 CPUs. Each CPU reports 4 PCIe root bridges with one IOMMU per root port. The system contains a Realtek PCIe ethernet network interface card (NIC). It is assigned to a dedicated IOMMU group. The NIC is passed-through to a virtual machine (VM) on the server. An ethernet cable connects the NIC with one of the on-board NICs. An NVIDIA Tesla T4 GPU is assigned to another dedicate IOMMU group that is managed by a different IOMMU than the NIC. Therefore, the NIC and the GPU do not share an IOTLB. An Intel Programmable Acceleration Card (PAC) with Intel Arria 10 GX FPGA shares the IOTLB with the NIC or the T4, depending on the experiment, by connecting it to PCIe slots that are managed by the IOMMU also managing the NIC or the GPU respectively. All other PCIe devices like the on-board NICs, memory controllers, etc. are connected to different IOMMUs and therefore cannot interfere with our measurements. The system has IAS 1.2 installed which contains OPAE version 1.1.2-1. Running `fpgainfo` reports bitstream id `0x123000200000185` and bitstream version `1.2.3`. We execute the GPU-accelerated database OmniSciDB[2] in version 5.10., which is the latest version at the time of writing. Additionally, CUDA version 11.4 and GPU driver version 470.57.02 are installed. The database consists of one table filled with the Meta Kaggle data set[3]. We have root access to this machine.

a10v: The IL ACE contains servers with two Intel Xeon Platinum 8180 CPUs. Each CPU reports 4 PCIe root bridges with IOMMU per root port. Two PCIe PACs with Arria 10 GX FPGAs are managed by two separate IOMMUs. All other PCIe devices are managed by other IOMMUs. The servers use IAS 1.1 and OPAE was installed on 01/01/2020 from the Git repository. Running `fpgainfo` reports bitstream id `0x113000200000177` and bitstream version `1.1.3`. We operate these machines with user privileges only.

s10v: The IL ACE features servers with two Intel Xeon Platinum 8280 CPUs. Each CPU reports 4 PCIe root bridges with one IOMMU

---

[1]https://wiki.intel-research.net/
[2]https://docs.omnisci.com/overview/overview#omniscidb
[3]https://www.kaggle.com/kaggle/meta-kaggle

per root port. An Intel FPGA PAC D5005 is connected via PCIe. All other PCIe devices are managed by other IOMMUs than the one managing the PAC. The servers use IAS 2.0 and OPAE was installed on 01/01/2020 from the Git repository. Running `fpgainfo` reports bitstream version `2.0.3` and bitstream id `0x203000200000339`. We operate these machines with user privileges only.

## 3.2 IOTLBs Cause Timing Behavior

During their PCIe performance benchmarking, Neugebauer et al. [43] found that an IOTLB miss results in a latency increase of 330 ns. Since the FPGAs in our systems are clocked at 200 MHz, the expected difference between fast and slow accesses is 66 clock cycles. Peglow's [49] work matches our expectation. With disabled IOMMU, the memory read latency for any address in main memory is distributed around 160 and 185 cycles. When the system is configured to use the IOMMU, this distribution shifts to 225 and 270 cycles for addresses that are accessed for the first time. Access times for subsequent accesses are distributed similarly to access times measured without IOMMU. Thus the measurable latency difference between accesses to addresses where the translation is present in or absent from the IOTLB lies between 65 and 85 clock cycles. We reproduced all values for the *a10l* system. On the IL ACE systems *a10v* and s10v, the latency difference between first accesses and subsequent accesses lies in the expected range. However, we cannot disable the IOMMU on the IL ACE systems to check whether the latency difference disappears.

## 3.3 Tools for Testing IOMMU Behavior

The IOMMU translates addresses for peripherals. Therefore, the CPU alone can only interact with the IOMMU in limited ways; we have to rely on a peripheral device to perform the experiments. For this purpose we used the PCIe PACs with DMA capabilities. We implement a hardware function for the FPGA that is programmable from software to capture the required measurements.

3.3.1 IOTLB Control from the CPU. To assist with these experiments, we also develop a kernel module that enables a program on the CPU to flush all entries from the IOTLB of a given IOMMU. When loaded, the kernel module uses a variety of functions and structures from the Linux kernel source, including those found in `<linux/pci.h>`, `<linux/iommu.h>`, and `<linux/dmar.h>` to find a PCIe device structure based on its vendor and device IDs, and from there find the device structure corresponding to the IOMMU that manages that PCIe device. That IOMMU device structure already contains a pointer to a function for flushing the IOMMU, so that function merely needs to be called. The kernel module uses a character file and ioctl as an interface by which user programs can call for the kernel module to flush the IOMMU. However, it takes root access to load a kernel module, since the module must read and write kernel memory. Therefore, we only tested algorithm 1 with the optional flush on our local system *a10l*.

3.3.2 Hardware Design. Our `iotlb_pnp` hardware module is designed against the Intel Acceleration Stack as would be the case in a cloud environment. The module is capable of performing memory accesses and timing the access latency. `iotlb_pnp` can be


(a) Model *1u*. Side-channel attacker with *user* privilege.

(b) Model *1k*. Side-channel attacker with *kernel* module.

(c) Model *2u*. Covert channel with *user* privilege.

(d) Model *2k*. Covert channel with *kernel* module.

Figure 1: Comparison of threat models. Dark red fills indicate functional units controlled by a malicious actor, and light green fills indicate functional units controlled by a victim. Diagonal lines indicate functional units that are only under coarse or indirect control, e.g., a simple network interface card or an accelerator that assists with certain applications but is not directly programmable. The dashed arrows indicate the flow of data through the channel.

programmed with up to 7 instructions. Currently, the design supports 5 instructions: `evset_prime`, `evset_probe`, `target_prime`, `target_probe`, and `wait`. Configuration and programming of the hardware module is performed via MMIO through OPAE. The prime instructions make the hardware module access a configured address (target) or set of addresses (eviction set). Probe instructions behave in the same way as the prime instructions but additionally count clock cycles. When probing an eviction set, the module can be configured to either measure the overall execution time of the instruction or time each memory access individually. The eviction sets used during priming and probing can be configured independent from each other, as is the case for the target instructions. The wait instruction simply makes the hardware module do nothing for a configured number of clock cycles.

3.3.3 Software. The software counterpart to the hardware module uses the OPAE C library to interact with the hardware design on the FPGA. This library allows us to control and observe the operation of the hardware module with memory-mapped I/O (MMIO) as well as — crucially for the work that this module must do — allocate shared pages of the system's main memory that the FPGA as well as the CPU can read and write.

3.4 Threat Models

We consider two general threat models with two variants each, as illustrated in Fig. 1. All four threat models include a malicious actor that can program and control a fast and programmable PCIe device (referred to in this section as the monitoring device) with

direct memory access (such as an FPGA or GPU) and an IOMMU providing address translation services for that device. Each model also includes a second peripheral (referred to in this section as the sending device) which also uses the same IOMMU for DMA address translation but does not need to be fast or directly programmable as part of the threat model. The monitoring device must be capable of timing memory accesses and reliably differentiate IOTLB hits from misses. The attacker must further be able to program the monitoring device directly to find eviction sets and execute Prime+Probes. The sending device only needs to have memory access patterns that can be triggered by a user, either by direct control, or triggerable through an application or system interface.

Models *1k* and *1u* are adversarial threat models for a *side-channel attack*, where a malicious user in control of the monitoring device exploits IOTLB contention to gain secret information from another user's application that triggers memory accesses in the sending device. Models *2k* and *2u* outline the requirements for a *covert channel* with cooperative sending and monitoring devices, where colluding malicious users in control of applications in separate security domains uses the IOTLB to transmit data covertly across the two devices. Models ending in *k* include kernel access alongside the monitoring device, and models ending in *u* do not. Kernel access is necessary to implement an IOTLB flush through a custom kernel module as outlined in Sec. 3.3. In Sec. 4 we show how fine-grained flushing control allows for more reliable eviction set construction. However, eviction set construction and Prime+Probe-based IOTLB side-channel attacks are still possible without flushing capabilities.

Whereas some side-channel attacks can be carried out with JavaScript from a web browser against a personal computer, we consider cloud environments as the primary site of IOTLB attacks, since the attacker must already have control of a peripheral. Renting a single GPU or FPGA in a cloud environment is easy; the primary logistical challenge of setting up a practical IOTLB side-channel or covert channel is IOMMU co-location – that is, ensuring that the monitoring device shares an IOMMU (and IOTLB) with the sending device. However, research into similar problems, like co-locating cloud instances for cache attacks, has yielded strategies for co-location that can be adapted to the IOTLB channel. İnci et al. [24] demonstrated two reliable co-location techniques for last-level caches that rely only on basic cache contention and so could be adapted to the IOTLB relatively easily. In a cooperative (covert channel) scenario, the sender instance sends a predetermined signal and the receiving instance searches the channel for a signal and attempts to match it with the agreed-upon signal. In an adversarial (attack) scenario, the attacker first chooses a target program and profiles it locally to learn to identify the traces it leaves. Then the attacker searches for such traces. For cache profiling, co-location is not necessary; the target program can be profiled within a single instance. In the case of IOTLB profiling, covert channel co-location may be used to first co-locate the cloud instance controlling the monitoring peripheral with another cloud instance that runs the target program which relies on a sending peripheral.

4 CONSTRUCTING EVICTION SETS

Initially, we hypothesized that the IOTLB would be organized like the CPU TLBs reverse-engineered in [18], with 2^s sets where s is an integer, some small number of ways per set, and a set mapping

Table 2: Notation used in algorithms

Symbol	Meaning
$A \leftarrow B$	A gets the value of B
$A \leftarrow_\in B$	A chosen randomly from B
$A \leftarrow_+ B$	B added to the set A
$A \leftarrow_- B$	B removed from the set A
$A \leftarrow_/ B$	Elements in B removed from A

```
1  Function evicts(target, evset)
       input  : target – address to be evicted
                evset – eviction set used for eviction attempt
       output : True, if 100 eviction attempts are successful
                False, otherwise
2      count ← 0 // # of contentions
3      for 0 ≤ i < 100 do
4          flush IOTLB // optional
5          target_prime()
6          evset_prime()
7          time ← target_probe()
8          if time > threshold then
9              count ← count + 1
10     return count == 100
```

Algorithm 1: The algorithm tests whether a given eviction set evicts a given target address from the IOTLB. The target_prime and evset_prime function calls have the FPGA access the respective set of addresses. The function call target_probe has the FPGA time the access time to the target address.

algorithm wherein the lowest s bits of the page address select the set number or some other combination of various bits of the page address forms the set number that the page is associated with. However, as we describe in Sec. A.1 in the appendix, this turned out to be false. So we set out to construct eviction sets for IOTLBs with an unknown architecture.

4.1 A New Approach to Eviction Set Construction

We developed a novel and platform-independent algorithm for finding eviction sets for any TLB or cache where the timing difference between a present entry and an evicted entry is known and measurable. Our approach is inspired by the baseline reduction algorithm in [63], which only reduces an already existing eviction set to its minimum necessary size, and the grow-split eviction set construction approach of Algorithm 1 in [36].

Like [36], our algorithm constructs eviction sets from a large pool of addresses by gathering candidates for an eviction set and then systematically discarding unnecessary ones; addresses not present in candidate eviction sets are used as test targets. The grow-split algorithm in [36] is specifically designed for a partitioned cache: it first constructs an eviction set for the entire cache, and then splits it into separate sets for each of the partitions. Our grow-reduce algorithm makes no assumption about cache organization, and uses a more generalized approach of building one eviction set at a time by adding addresses until evictions are reliable and then testing

```
1  Function constructEvset(target, pool)
      input    : target – target address to be evicted
                 pool – address pool
      output   : evset – an eviction set for target
2     evset ← ∅
3     count ← 0 // # of contentions
      // Grow
4     while count < 50 and |pool| > 0 do
5         page ←∈ pool; evset ←₊ page; pool ←₋ page
6         if evicts(target, evset) then
7             count ← count + 1

      // Reduce
8     foreach page in evset do
9         evset ←₋ page
10        if not evicts(target, evset) then
11            evset ←₊ page

12    return evset
```

Algorithm 2: The algorithm constructs an IOTLB eviction set for a given target address. The addresses for the eviction set are chosen from the given address pool.

```
1  Function evsetFinding(poolSize)
      input    : poolSize – number of addresses to be allocated
      output   : evsets – Eviction sets for the IOTLB
2     pool ← alloc(poolSize)
3     targets ← ∅
4     evsets ← ∅
5     while poolSize > 0 do
6         target ←∈ pool // Random page as target
7         pool ←₋ target
8         if evsets do not evict target then
9             targets ←₊ target
10            evsets ←₊ constructEvset(target, pool)
11            pool ←/ evsets
12        poolSize ← size(pool);
13    return evsets
```

Algorithm 3: This algorithm constructs as many eviction sets as needed to evict any target address from the IOTLB. The algorithm takes an integer as input that indicated the size of the address pool that is used to construct the eviction sets. A pool size of 4096 was used for the tests in this paper.

which addresses can be discarded without losing reliability. It aims to create an exhaustive set of eviction sets by searching the entire address pool; redundant sets are avoided by ensuring that potential test targets are not already reliably evicted by another set.

4.1.1 Grow-Reduce Algorithm.

The most basic function in our algorithm tests whether or not a hypothetical eviction set evicts a given target address (see algorithm 1). The software uses the hardware module described previously to perform a prime and probe test. First, the FPGA accesses the target followed by an access to each address in the eviction set. Then the target is accessed again and the access latency is measured. We define that an eviction set evicts a target if the latency of the second access to the target is above a certain threshold. We choose the threshold in the middle of the observed latency gap between fast and slow accesses observed on the different systems. Before each prime and probe test, we optionally cleared the IOTLB.

The construction of an eviction set for a fixed target address is given in algorithm 2. It takes a target address and a pool of addresses

as inputs. The eviction set is initialized as an empty set. During the "grow" step random addresses are chosen from the address pool and added to the eviction set until the eviction set contains enough addresses to evict the target. Obviously, the eviction set may contain unnecessary addresses at this point. This is why a reduction step follows where each address is tested for its necessity. If an address is not needed, it is removed from the eviction set and put back in the address pool.

At the highest level, our algorithm shown in algorithm 3 automatically constructs as many eviction sets as it can find. The program first allocates a pool of memory pages. For our experiments we used a pool size of 4096 addresses. The algorithm manages two sets: The *targets* set is used to store the different target addresses used during eviction set construction. The *evsets* set stores all eviction sets constructed by the algorithm. After this initialization step, the algorithm picks a random target address from the pool and removes it from the pool. If *evsets* does not contain an eviction set for the target address yet, a new eviction set is constructed. The target address and the new eviction set are added to their corresponding sets. All addresses in the newly constructed eviction set are then removed from the pool. This procedure is repeated until the pool does not contain any addresses anymore.

4.1.2 Evaluation of New Eviction Set Algorithm.
We found that the optional flushing of the IOTLB has an impact on the size and reliability of IOTLB eviction sets. [4] The major differences are laid out in Table 3, which enumerates general performance metrics of eviction sets constructed with our grow-reduce algorithm and [36]'s grow-split algorithm both with and without flushing. Enabling IOTLB flushes before the Prime+Probe step will make both algorithms return a single eviction set containing 118 addresses. The success rate of such eviction sets is 100% in every case we observed.

Without IOTLB flushes, neither algorithm produces such consistently sized or reliable eviction sets. This is likely due to a replacement policy that we were unable to deduce. In this scenario we can better see the advantage of our grow-reduce algorithm. It produces eviction sets that are both smaller and twice as reliable than those produced by the grow-split algorithm.

Fig. 2 visualizes in detail the results of further experimentation with small implementation tweaks in our algorithm. In these experiments we found that the size and number of eviction sets constructed were very similar on all tested systems, *a10l*, *a10v*, and *s10v*. We thus conclude that the IOTLB architecture on all tested systems is very similar in terms of IOTLB size, organization and replacement policy.

5 ANALYSIS OF SIDE-CHANNEL LEAKAGES

We now use the constructed eviction sets to further investigate the amount of leakage from PCIe devices observable in the IOMMU. Though we use the FPGA for channel monitoring outside of a virtualized environment for simplicity's sake, this channel still poses a threat from one virtual environment to another or from a virtual environment to hypervisor. Major cloud platforms like AWS and Alibaba Cloud now allow users to rent direct access to FPGAs with DMA capabilities, meaning that malicious tenants

---

[4]Flushing the IOTLB requires kernel access; see threat models *1k* and *2k* in Sec. 3.4. For this reason, Table 3 contains data only from experiments on the *a10l* system.

Table 3: Comparison of eviction set finding algorithms on the IOTLB of the *a10l* test system. All tests were conducted on the *a10l* system using pools of 4096 addresses, and repeated 40 times. Eviction set orders were randomized between prime and probe steps during testing.

| Flush	Algorithm	Number of sets	Set size	Useful sets per target	Average best eviction rate
enabled	Grow-Reduce (this work)	1.00	118.00	1.00	100.00 %
	Grow-Split ([36])	1.00	118.00	1.00	100.00 %
disabled	Grow-Reduce (this work)	32.08	110.05	0.98	82.23 %
	Grow-Split ([36])	10.70	50.69	0.98	28.00 %
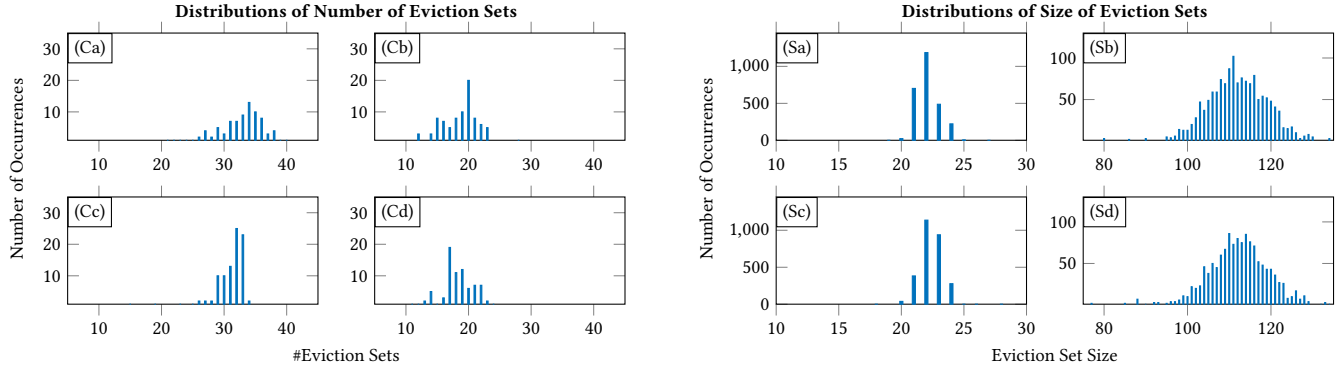
Figure 2: Number of eviction sets and the size of each constructed set needed to evict any target IOVA after running algorithm 3 for 100 times each. During eviction set construction, randomization of the eviction set was turned off for measurements (a) and (c) and turned on for (b) and (d). For measurements (c) and (d), the algorithm waited 100 ns between each eviction test. For measurements (a) and (b) this was not the case. If the order of accesses during the evset_prime() is static throughout one run of algorithm 3, the resulting eviction sets contain 20 to 25 addresses each. The average success rate is slightly below the average success rate of eviction sets constructed with randomized access order during evset_prime(). In turn, randomizing the access order yields on average slightly less but bigger sets. The success rate of these sets, with or without randomized access order, evict a target with probabilities above 90%.

could easily run hardware designs that monitor the IOMMU side-channel without root privileges. Any other PCIe devices that are co-located on the IOMMU with a malicious FPGA and using translated DMA (most modern devices use DMA, and virtualized DMA always requires translation if the IOMMU is shared) are sources of leakage and therefore potential attack targets. We focus our analysis on an in-memory SQL database accelerated by a graphics card.

## 5.1 GPU-Accelerated SQL Database Leakage

We now inspect the amount of IOTLB leakage observable from the FPGA when it is co-located with a GPU that runs an SQL database. For our tests, we co-locate the FPGA with an NVIDIA Tesla T4 GPU that runs the OmniSci SQL server on it. We wish to understand the data leakage patterns of the GPU-accelerated database application, so for these experiments we consider threat model *1k*, where the attacker has the most precise control over the channel. Fig. 3 shows a stack diagram of the setup on our *a10l* platform. The test application interacts with our hardware module on the FPGA to construct, prime and probe an eviction set for the IOTLB. Additionally, the application can issue SQL queries to the database which computes the result on the GPU.
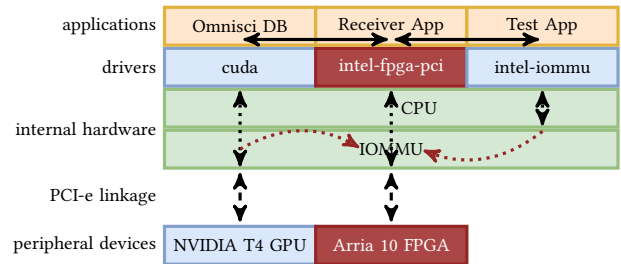
Figure 3: Stack diagram of the CPU to peripheral and peripheral to peripheral covert channel and side-channel tests.

After constructing an eviction set for the IOTLB, the test app primes the IOTLB. During the waiting phase, the app runs an SQL query on the GPU. The tested queries differ (significantly) in the size of the returned results. After the SQL result is returned to the test application, the FPGA probes the IOTLB and reports the access latency back to the application.
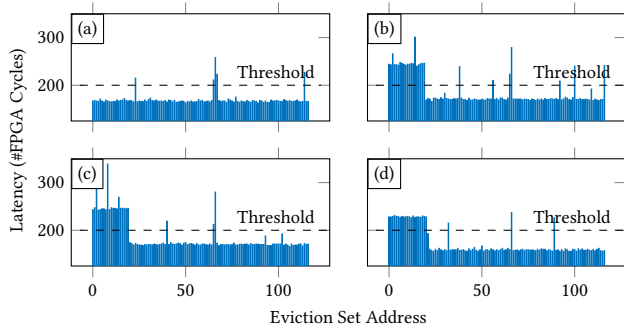
Figure 4: Measurements for the conducted experiments with the SQL database. During measurement (a), the test app did not run any query. The queries run in measurements (b) - (d) returned no, one and 409600 rows of data from the database. It is clearly visible that the SQL queries leave a footprint in the IOTLB.

Fig. 4 (b) - (d) show probe measurements for queries returning no, one and 409,600 rows[5] of data from the database. During the measurement shown in Fig. 4 (a), no query was executed on the GPU. The separate access times for each eviction set address are plotted along the x-axis. The y-axis shows the measured latency for this address. Clearly, the GPU leaves a footprint in the IOTLB when it computes an SQL query. But, there is no measurable difference between the queries even if their results significantly differ in size.

Changing the test app to probe the eviction set while the SQL query executes on the GPU shows that the observable activity in the IOTLB is similar for all queries over time, besides the fact that queries with larger results produce longer traces as it takes longer to compute the result. Interestingly, the activity in the IOTLB happens towards the beginning of the query's computation. At the time where the computed result is sent back to the CPU, there is no activity in the IOTLB. This is easily explained by the way CUDA realizes the data transfer of the result from the GPU to the CPU: it uses MMIO[6] instead of DMA[7]. We verified the explanation by inspecting the PCIe performance counters with the PCM tools[8]. The performance counters showed an increased amount of MMIO read requests that in total match the size of the returned result.

5.2 Side-Channel Impact

So far, the observed leakage introduced by the IOTLB is mostly limited to a single bit describing whether a neighboring accelerator is in use or not. This is caused by two facts:

(a) Controlling an accelerator via MMIO rather than through DMA is a common usage model and limits the attack surface for IOTLB-based side-channel attacks because the CPU performs the address translation in the CPU's MMU instead of the GPU translating addresses via the IOMMU.

(b) Current PCIe devices usually perform DMA as bulk transfers, thereby limiting the overall PCIe protocol overhead. Loading data in a bulk transfer into device memory, computing on the data locally

---

[5]One row in our case contains 36 bytes of data.
[6]The CPU initializes the data transfer.
[7]The peripheral initializes the transfer.
[8]pcm-pcie – https://github.com/opcm/pcm

and eventually transferring the result back to the main memory in a bulk transfer means that no data-dependent access patterns – which would leak information – are observable in general.

The two facts mentioned will likely change in the near future as PCIe 5.0 is rolled-out and Compute eXpress Link (CXL) is introduced. [9] PCIe 5.0 reaches transfer speeds that are comparable with CPU main memory accesses. This may lead device developers to include smaller memory on their devices and in turn access the main memory more often. Furthermore, CXL features a coherency protocol that streamlines caching between main memory and PCIe device memory. Again, this will lead device and driver developers to change from bulk transfers to more fine-grained data-dependent DMA accesses.

In addition, FPGA vendors keep pushing for FPGA devices being the first-class compute device in a system while the CPU is merely used to manage the system and provide the FPGA with (increasingly sensitive) data. Therefore, while the described side-channel is not yet very dangerous at the time of writing, it will become important in the near future. We highlight the side-channels existence and relevance *before* widespread deployment of CXL and PCIe 5.

6 COVERT CHANNELS

After identifying the IOTLB leakage and different ways to trigger and observe it, we now use our knowledge to construct two covert channels to prove the practicality of the channel with threat models *2u* and *2k*. The first channel is constructed between two peripherals and requires user privileges and DMA access to pages in main memory (model *2u*). This channel could be implemented between two virtual machines, each with control of a DMA-enabled peripheral, such as Amazon's F1 FPGA instances or various GPU-enabled EC2 instances, as long as the two instances' peripherals share an IOMMU. The performance of the covert channel can be improved if the receiver has root access on the host. The second channel is unidirectional from CPU to peripheral and requires the sender to have root access to the host machine (model *2k*), thereby mostly serving as a proof of concept. For both channels, the receiver must be able to measure time, e. g. through precise internal timers or high-speed network connection with external timers. This is the case for, e. g. GPUs [15], NICs and FPGAs. All experiments in this section were run on the a10l system.

6.1 Covert Channel between Peripherals

Another research question is whether two peripherals can use the IOTLB to construct a covert channel between each other. To answer this question, we co-locate the Arria 10 with the Tesla T4. Our goal is to use the footprint that an SQL query computed on the GPU leaves in the IOTLB to send information to the FPGA. Such a covert channel exists in a scenario where the sender uses a website that, depending on the actions performed on the website, runs SQL queries on a GPU-accelerated database. The sender can then exploit the website to send information to the co-located FPGA.

We prepare a10l as shown in Fig. 3. The sender encodes a one into running an SQL query and running no query encodes a zero. The receiver uses the `iotlb_pnp` hardware function on the FPGA to monitor the IOTLB using the Prime+Probe technique. Each SQL

---

[9]AMD CPUs and Intel FPGAs supporting CXL are already available. Intel plans rolling out compatible CPUs in the beginning of 2023 [7, 58].

Table 4: Throughput and error rate for the covert channels tested on the *a10l* system. For the peripheral-peripheral channel, sender and receiver are perfectly synchronous. The channel itself is very reliable which leads to nearly no errors. The throughput depends on the number of 1-bits in the message as each 1-bit is encoded into running a SQL-query on the sender peripheral which takes a rather long time of 0.3 seconds. For the CPU-peripheral channel, sender and receiver are not perfectly synchronous which leads to the rather high error rate. The throughput is limited by the speed of the CPU flushing the IOTLB. For both channels, plain bits were sent without encoding.

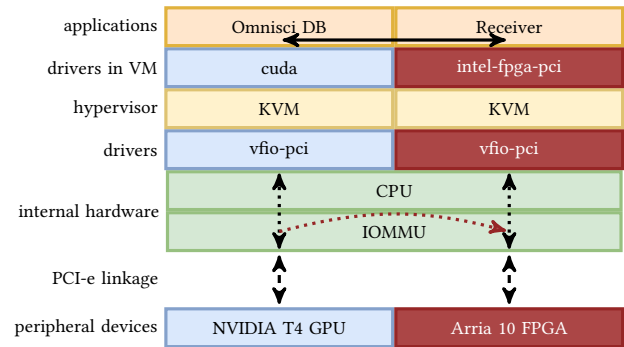|	Sender	Receiver	Method	Environment	Throughput	Error rate	Content of message
Sec. 6.1	Peripheral	Peripheral	Prime+Probe	Bare metal (cf. Fig. 3)	3.4 bps	0%	All 1s
					6.65 bps	0%	Even mix of 1s and 0s
					246.15 bps	0.1%	All 0s
					7.58 bps	0%	ASCII-encoded text
Sec. 6.2	CPU	Peripheral	Flush+Reload	Bare metal (cf. Fig. 3)	15023 bps	30.09%	*Performance not dependent on content*

(a) Scenario as in Fig. 3; big endian transmission. The FPGA uses an eviction set that was constructed using IOTLB flushes. This results in very reliable eviction sets and in turn a reliable transmission.

(b) Scenario as in Fig. 6; little endian transmission. The FPGA uses eviction sets constructed *without* IOTLB flushes. TEven though the transmission is free of errors, buit turns out to be more noisy.

Figure 5: Peripheral to peripheral covert channel transmissions of t. The message "Hello" was sent in big endian format.

Figure 6: Stack diagram of the GPU accelerated SQL database covert channel across virtual machines.

query evicts 18-20 entries of the receiver's eviction set (cf. Fig. 4 (b) - (d)). A plot of the number of IOTLB misses measured during message transmission is given in Fig. 5a. We found that basically no errors occur if sender and receiver are synchronized. This means that the channel is nearly free of bit-flip errors. If perfect synchronization is not achievable, the channel suffers from insertion and deletion errors. In this case techniques from [38] can be applied to overcome these errors. The channel's throughput highly depends on the number of one bits in the message. This is because the execution time of a single SQL query takes about 0.3 seconds. Table 4 shows more detailed measurements for different 0-1-ratios in the message that is transferred over the covert channel. Of course, a GPU application optimized for acting as a sender in this scenario would allow us to increase the bandwidth of the channel.

For the previous test, the eviction set used by the FPGA was constructed with IOTLB flushes to work with eviction sets of optimal reliability. As mentioned earlier, IOTLB flushes require kernel privileges on commodity host Linux systems. User-level receivers or receivers located in virtual machines have to use the less reliable eviction sets constructed without IOTLB flushes. As can be seen in Fig. 5b, this results in more noise in the measurements. The depicted transmission is still free of errors but some bits are at the edge of being falsely classified. To overcome potential bitflip errors, error detection mechanisms like CRC codes or error correction codes like Hadamard codes can be applied [38]. The presented covert channel works between any two peripherals that use DMA to access the main memory. For the receiver, the accessible memory region needs to be sufficiently large to allow for eviction set construction. Additionally, the receiver needs a mechanism to measure the memory access latency. Programmable or configurable peripherals like FPGAs or GPUs will meet both receiver requirements even

in the most stringent cloud environments if bare metal instances are available for rent. An FPGA or GPU sender has fine-grained control of the channel, but a more opaque sender like a smart NIC or PCIe-enabled storage device could work as a sender, albeit more likely to be noisy or unreliable.

Peripherals that manage secrets and perform DMAs depending on the value of the secret must be aware that neighboring devices connected to the same IOMMU may be able to observe their access patterns. This is especially true for peripherals where the programming model assumes unified memory that abstracts separate physical memory locations like device and system memory away from the developer as in this case the leaking DMA may occur without the knowledge of the developer. As of today, data-dependent DMA is used seldomly due to the overhead that renders it inefficient. But we expect this behavior to change with the introduction of PCIe 5.0 and CXL as mentioned in earlier sections.

6.2 Covert Channel from CPU to Peripheral

The CPU is very limited in interacting with the IOTLB directly. Because the IOMMU translates addresses for peripherals only, memory accesses from the CPU do not interfere with the IOTLB. The only way for the CPU to interfere with the IOTLB is by changing page table entries or instructing the IOMMU to flush certain (or all) entries in the IOTLB. Usually, only the OS, hypervisor or VMM issues page table changes or IOTLB flushes, which is why the Linux kernel does not provide an interface for flushing the IOTLB to userland. To overcome this problem, we load a self-developed kernel module that exposes a IOTLB flush API to our test application. An overview of our system setup for this covert channel is given in Fig. 3.

Since a peripheral can distinguish IOTLB hits from misses, flushing the IOTLB allows the CPU to send information covertly to peripherals. A global IOTLB flush takes $17\,\mu s$ on average. Flushing all entries from the IOTLB encodes a 1 and sleeping for $17\,\mu s$ encodes a 0. As the receiver we use the `iotlb_pnp` hardware module described in Sec. 3.3.2. The hardware function is programmed to continuously probe a fixed target address. Whenever a probe reports a slow access, a 1 is received. Otherwise, the hardware receives a 0. We implement the covert channel in a trivial way without applying any encoding for error correction or synchronization. Because a memory access from the FPGA running at 200 MHz only takes around $1\,\mu s$ we roughly synchronize the FPGA with the CPU by making the FPGA `wait` for a certain amount of cycles. We determined the number of cycles to wait by repeatedly flushing the IOTLB and increasing the number of wait cycles until all FPGA memory accesses are slow. After this very rough synchronization step, a message of $2^{16} - 1$ bits generated by a linear feedback-shift register is transmitted to measure throughput and error rates. The result is given in Table 4. As can be seen, this basic covert channel without further optimizations already achieves a throughput of around 15 kBit/s. The error rate is 30% which can be improved significantly by applying error-correction and error-handling techniques as e.g. described in [38].

Because so far the covert channel only offers communication in one direction, we tried to improve the channel to offer bi-directional

message transfer. To do so we checked the timing behavior of flushing the IOTLB. The `clflush` instruction on x86 CPUs has a data-dependent execution time [21]. In our case, a data-dependency of the flush time on IOTLB entries would allow us to construct the reverse covert channel. However, our experiments show no measurable timing behavior of the flush that can be related to the usage of the IOTLB; an IOTLB flush takes around $17\,\mu s$ independent of FPGA memory accesses before or even during the flush. The latency is also independent from whether only addresses of a certain peripheral or all entries of the IOTLB are flushed. However, peripheral-to-CPU covert channels based on the CPU cache do exist [64].

The demonstrated covert channel is reliable without applying special synchronization, error-correction, or error-detection techniques. However, only peripherals can act as the receiver while the CPU is limited to the role of the sender. Also, with the standard IOMMU drivers in Linux, the sending process is required to run kernel-level code to perform IOTLB flushes. A privileged device driver that flushes the IOTLB under certain circumstances may expose this flushing capability to an unprivileged user. Device drivers that make extensive use of IOTLB flushes may also be vulnerable to a side-channel attack from an untrusted peripheral device that monitors the IOTLB for flushes. For example, a driver developer may chose to include IOTLB flushes to remove traces of a trusted peripheral's activity for security; however, the timing between flushes could leak information about the operation of an application using that peripheral.

7 COUNTERMEASURES

Like many microarchitectural attacks, there are a variety of defenses against IOTLB side-channels that can be implemented at nearly any level of a system. We first present immediately available actions that can be taken by system administrators and cloud application developers, and then discuss defenses that can be built into future IOMMU architectures.

7.1 Securing Existing Systems

In cases where multiple users who do not trust each other may use the same machine, ensuring that no two users (or no one user and the hypervisor) have access to peripherals on the same IOMMU hardware is sufficient to protect against IOTLB side-channel attacks. On a Linux host, `/sys/class/iommu/` provides information on a system's IOMMU devices and the PCIe devices that use them [53]. Typically, systems have several IOMMU devices, each of which is linked to a few PCIe endpoints, which may be internal PCIe devices or external devices plugged into PCIe slots on the motherboard. Endpoints cannot be reassigned to new IOMMUs, so ensuring full isolation may limit scaling capacity. For example, a CSP could not use a motherboard with eight full-size, full-speed PCIe slots managed in pairs by four IOMMUs to provide eight fully isolated single-GPU cloud instances, even though eight GPUs fit in the PCIe slots of the system.

On the *application level*, code and hardware involved in data dependent computation can rely on constant time algorithms with constant memory access patterns, so no information about the operations is leaked through the IOTLB. For cryptographic implementations this is a common technique but for database systems constant memory access patterns and timings are not easily

achieved. Private Information Retrieval (PIR) protocols [11, 33] can be a solution, but modern implementations[10] usually only support index queries. Recent attempts [22] to also support range queries may still leak information about the response size.

A hypervisor can enable Address Translation Services (ATS) [48] for a peripheral to remove all of its traces from its IOTLB. Address Translation Services (ATS) allows a device to maintain and use a local on-device TLB for address translation and selectively bypass IOMMU translation. Since locally-translated requests are not translated by the IOMMU, they do not leave any trace in the IOTLB. However, devices must specifically support ATS to use it, and furthermore, allowing ATS for untrusted devices is not advisable. ATS allows a device to provide *any* physical address as part of a DMA request and mark it as "translated". Malicious devices may exploit ATS for unrestricted physical memory access [37]. Therefore, ATS must only be allowed for trusted devices.

Hypervisors can also achieve a separation of the IOTLB between mutually untrusted tenants by IOTLB partitioning. For set-associative IOTLBs, set partitioning can be done by the hypervisor in software by only allocating I/O virtual addresses of sets to each tenant [67]. However, set-based partitioning may not work with peripherals that rely on the address space being contiguous.

7.2 Securing Future IOMMUs

If *hardware modifications* are a viable option to implement countermeasures, then way-based partitioning is another option. It needs to be supported by the IOMMU hardware so that the hypervisor can map each address of a thread to a fixed number of ways like is possible with Intel CAT [35, 44] for CPU-internal caches.

Future IOMMUs could include support for flagging a page translation as uncacheable. This would ensure that it is never stored in the IOTLB and that the use of that page would never affect the IOTLB state, so it would be invisible to any side-channel attack. However, all accesses to that page would be as slow as IOTLB misses, increasing latency and likely reducing maximum throughput.

8 CONCLUSION

State-of-the-art cloud environments use direct memory access managed by IOMMUs to offer high speed, low latency, and isolated memory access to an increasingly wide variety of peripherals. These peripherals support and accelerate many types of applications and virtual hardware functions, including those that perform secure operations or handle sensitive data. In this paper we demonstrated a new side-channel attack against IOTLBs in such IOMMUs that works across virtual environments and threatens cloud tenants. We developed a new eviction set finding algorithm that works without prior assumptions of cache or TLB organization and a hardware module for an FPGA that implements the fundamentals necessary to exploit the IOTLB side-channel. We used these tools to record a side-channel trace from a GPU running a database acceleration library. The results prove that the IOTLB can be used as a side-channel to spy on co-located devices. We highlight this fact by showing a very reliable covert channel from the GPU to the FPGA where we use the database application running on the GPU to encode messages into the GPU's system memory access patterns. While we acknowledge

the limitations of the IOTLB channel with current hardware and applications, we argue that with the upcoming PCIe 5.0 and CXL standards, IOMMU usage patterns will change and fine-grained IOTLB side-channel attacks will become practical. To overcome the threat of the side-channel, we suggest a variety of countermeasures that can be implemented on different system levels ranging from hardware modifications up to the implementation of applications. Many of these countermeasures fully eliminate the threat of IOTLB side-channels, but at the same time reduce the speed of peripherals or scalability of the systems that host them. Therefore, when designing or choosing hardware for large-scale, high-performance, secure services, IOTLB threats must be acknowledged and IOTLB isolation measures must be carefully considered for the specific needs of the system. Furthermore, when designing security-critical peripherals or security-critical software or firmware that makes use of peripherals, timing leakages from peripheral memory accesses must be addressed with constant-time design practices.

REFERENCES

[1] Alibaba Cloud. 2019. FPGA-based compute-optimized instance families. https://www.alibabacloud.com/help/doc-detail/108504.html Access: 2019-10-15.
[2] Alibaba Cloud ECS. 2020. Introducing the Sixth Generation of Alibaba Cloud's Elastic Compute Service. https://www.alibabacloud.com/blog/introducing-the-sixth-generation-of-alibaba-clouds-elastic-compute-service_595716 Access: 2022-01-31.
[3] José Bacelar Almeida, Manuel Barbosa, Gilles Barthe, François Dupressoir, and Michael Emmi. 2016. Verifying Constant-Time Implementations. In *USENIX Security Symposium*. USENIX Association, 53–70.
[4] Amazon AWS. 2017. Amazon EC2 F1 Instances. https://aws.amazon.com/ec2/instance-types/f1/ Access: 2019-10-12.
[5] Amazon AWS. 2018. AWS Nitro System. https://aws.amazon.com/de/ec2/nitro/ Access: 2022-01-31.
[6] AMD 2021. *AMD I/O Virtualization Technology (IOMMU) Specification* (3.06-pub ed.). AMD.
[7] AMD. 2022. Offering Unmatched Performance, Leadership Energy Efficiency and Next-Generation Architecture, AMD Brings 4th Gen AMD EPYC Processors to The Modern Data Center. https://ir.amd.com/news-events/press-releases/detail/1100/offering-unmatched-performance-leadership-energy Access: 2022-12-04.
[8] Gilles Barthe, Gustavo Betarte, Juan Diego Campo, Carlos Daniel Luna, and David Pichardie. 2014. System-level Non-interference for Constant-time Cryptography. In *CCS*. ACM, 1267–1279.
[9] Sandrine Blazy, David Pichardie, and Alix Trieu. 2017. Verifying Constant-Time Implementations by Abstract Interpretation. In *ESORICS (1) (LNCS, Vol. 10492)*. Springer, 260–277.
[10] Pietro Borrello, Andreas Kogler, Martin Schwarzl, Moritz Lipp, Daniel Gruss, and Michael Schwarz. 2022. ÆPIC Leak: Architecturally Leaking Uninitialized Data from the Microarchitecture. In *USENIX Security Symposium*. USENIX Association, 3917–3934.
[11] Benny Chor, Eyal Kushilevitz, Oded Goldreich, and Madhu Sudan. 1998. Private Information Retrieval. *J. ACM* 45, 6 (1998), 965–981.
[12] Lucian Cojocar, Kaveh Razavi, Cristiano Giuffrida, and Herbert Bos. 2019. Exploiting Correcting Codes: On the Effectiveness of ECC Memory Against Rowhammer Attacks. In *IEEE S&P*. IEEE, 55–71.

---

[10] e. g. https://github.com/ReverseControl/MuchPIR

[13] Craig Disselkoen, David Kohlbrenner, Leo Porter, and Dean M. Tullsen. 2017. Prime+Abort: A Timer-Free High-Precision L3 Cache Attack using Intel TSX. In *USENIX Security Symposium*. USENIX Association, 51–67.

[14] Daniel Firestone, Andrew Putnam, Sambrama Mundkur, Derek Chiou, Alireza Dabagh, Mike Andrewartha, Hari Angepat, Vivek Bhanu, Adrian M. Caulfield, Eric S. Chung, Harish Kumar Chandrappa, Somesh Chaturmohta, Matt Humphrey, Jack Lavier, Norman Lam, Fengfen Liu, Kalin Ovtcharov, Jitu Padhye, Gautham Popuri, Shachar Raindel, Tejas Sapre, Mark Shaw, Gabriel Silva, Madhan Sivakumar, Nisheeth Srivastava, Anshuman Verma, Qasim Zuhair, Deepak Bansal, Doug Burger, Kushagra Vaid, David A. Maltz, and Albert G. Greenberg. 2018. Azure Accelerated Networking: SmartNICs in the Public Cloud. In *NSDI*. USENIX Association, 51–66.

[15] Pietro Frigo, Cristiano Giuffrida, Herbert Bos, and Kaveh Razavi. 2018. Grand Pwning Unit: Accelerating Microarchitectural Attacks with the GPU. In *IEEE S&P*. IEEE, 195–210.

[16] Pietro Frigo, Emanuele Vannacci, Hasan Hassan, Victor van der Veen, Onur Mutlu, Cristiano Giuffrida, Herbert Bos, and Kaveh Razavi. 2020. TRRespass: Exploiting the Many Sides of Target Row Refresh. In *IEEE S&P*. IEEE, 747–762.

[17] Silvia E. Gianelli. 2017. Xilinx Announces General Availability of Virtex UltraScale+ FPGAs in Amazon EC2 F1 Instances. https://www.xilinx.com/news/press/2017/xilinx-announces-general-availability-of-virtex-ultrascale-fpgas-in-amazon-ec2-f1-instances.html Access: 2019-10-15.

[18] Ben Gras, Kaveh Razavi, Herbert Bos, and Cristiano Giuffrida. 2018. Translation Leak-aside Buffer: Defeating Cache Side-channel Protections with TLB Attacks. In *USENIX Security Symposium*. USENIX Association, 955–972.

[19] Ben Gras, Kaveh Razavi, Erik Bosman, Herbert Bos, and Cristiano Giuffrida. 2017. ASLR on the Line: Practical Cache Attacks on the MMU. In *NDSS*. The Internet Society.

[20] Daniel Gruss, Julian Lettner, Felix Schuster, Olga Ohrimenko, István Haller, and Manuel Costa. 2017. Strong and Efficient Cache Side-Channel Protection using Hardware Transactional Memory. In *USENIX Security Symposium*. USENIX Association, 217–233.

[21] Daniel Gruss, Clémentine Maurice, Klaus Wagner, and Stefan Mangard. 2016. Flush+Flush: A Fast and Stealthy Cache Attack. In *DIMVA (LNCS, Vol. 9721)*. Springer, 279–299.

[22] Junichiro Hayata, Jacob C. N. Schuldt, Goichiro Hanaoka, and Kanta Matsuura. 2020. On Private Information Retrieval Supporting Range Queries. In *ESORICS (2) (LNCS, Vol. 12309)*. Springer, 674–694.

[23] Ralf Hund, Carsten Willems, and Thorsten Holz. 2013. Practical Timing Side Channel Attacks against Kernel Space ASLR. In *IEEE S&P*. IEEE, 191–205.

[24] Mehmet Sinan Inci, Berk Gülmezoglu, Thomas Eisenbarth, and Berk Sunar. 2016. Co-location Detection on the Cloud. In *COSADE (LNCS, Vol. 9689)*. Springer, 19–34.

[25] Mehmet Sinan Inci, Berk Gülmezoglu, Gorka Irazoqui, Thomas Eisenbarth, and Berk Sunar. 2016. Cache Attacks Enable Bulk Key Recovery on the Cloud. In *CHES (LNCS, Vol. 9813)*. Springer, 368–388.

[26] Intel Corporation 2012. *Intel Data Direct I/O Technology (Intel DDIO): A Primer* (1 ed.). Intel Corporation.

[27] Intel Corporation 2016. *Intel 64 and IA-32 Architectures Optimization Reference Manual*. Intel Corporation.

[28] Intel Corporation 2019. *Intel Virtualization Technology for Directed I/O – Architecture Specification* (3.1 ed.). Intel Corporation.

[29] Gorka Irazoqui, Thomas Eisenbarth, and Berk Sunar. 2015. Systematic Reverse Engineering of Cache Slice Selection in Intel Processors. In *DSD*. IEEE, 629–636.

[30] Salman Abdul Khaliq, Usman Ali, and Omer Khan. 2021. Timing-based side-channel attack and mitigation on PCIe connected distributed embedded systems. In *HPEC*. IEEE, 1–7.

[31] Paul Kocher, Jann Horn, Anders Fogh, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. 2019. Spectre Attacks: Exploiting Speculative Execution. In *IEEE S&P*. IEEE, 1–19.

[32] Michael Kurth, Ben Gras, Dennis Andriesse, Cristiano Giuffrida, Herbert Bos, and Kaveh Razavi. 2020. NetCAT: Practical Cache Attacks from the Network. In *IEEE S&P*. IEEE, 20–38.

[33] Eyal Kushilevitz and Rafail Ostrovsky. 1997. Replication is NOT Needed: SINGLE Database, Computationally-Private Information Retrieval. In *FOCS*. IEEE, 364–373.

[34] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom, and Mike Hamburg. 2018. Meltdown: Reading Kernel Memory from User Space. In *USENIX Security Symposium*. USENIX Association, 973–990.

[35] Fangfei Liu, Qian Ge, Yuval Yarom, Frank McKeen, Carlos V. Rozas, Gernot Heiser, and Ruby B. Lee. 2016. CATalyst: Defeating last-level cache side channel attacks in cloud computing. In *HPCA*. IEEE, 406–418.

[36] Fangfei Liu, Yuval Yarom, Qian Ge, Gernot Heiser, and Ruby B. Lee. 2015. Last-Level Cache Side-Channel Attacks are Practical. In *IEEE S&P*. IEEE, 605–622.

[37] A. Theodore Markettos, Colin Rothwell, Brett F. Gutstein, Allison Pearce, Peter G. Neumann, Simon W. Moore, and Robert N. M. Watson. 2019. Thunderclap: Exploring Vulnerabilities in Operating System IOMMU Protection via DMA from Untrustworthy Peripherals. In *NDSS*. The Internet Society.

[38] Clémentine Maurice, Manuel Weber, Michael Schwarz, Lukas Giner, Daniel Gruss, Carlo Alberto Boano, Stefan Mangard, and Kay Römer. 2017. Hello from the Other Side: SSH over Robust Cache Covert Channels in the Cloud. In *NDSS*. The Internet Society.

[39] MITRE. 2018. CVE-2018-12126. Available from MITRE, CVE-ID CVE-2018-12126. https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-12126 Access: 2022-09-01.

[40] MITRE. 2018. CVE-2018-12127. Available from MITRE, CVE-ID CVE-2018-12127. https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-12127 Access: 2022-09-01.

[41] Benoît Morgan, Eric Alata, Vincent Nicomette, and Mohamed Kaâniche. 2016. Bypassing IOMMU Protection against I/O Attacks. In *LADC*. IEEE, 145–150.

[42] Benoît Morgan, Eric Alata, Vincent Nicomette, and Mohamed Kaâniche. 2018. IOMMU protection against I/O attacks: a vulnerability and a proof of concept. *J. Braz. Comput. Soc.* 24, 1 (2018), 2:1–2:11.

[43] Rolf Neugebauer, Gianni Antichi, José Fernando Zazo, Yury Audzevich, Sergio López-Buedo, and Andrew W. Moore. 2018. Understanding PCIe performance for end host networking. In *SIGCOMM*. ACM, 327–341.

[44] Khang T. Nguyen. 2016. Usage Models for Cache Allocation Technology in the Intel Xeon Processor E5 v4 Family.

[45] Yossef Oren, Vasileios P. Kemerlis, Simha Sethumadhavan, and Angelos D. Keromytis. 2015. The Spy in the Sandbox: Practical Cache Attacks in JavaScript and their Implications. In *CCS*. ACM, 1406–1418.

[46] Dag Arne Osvik, Adi Shamir, and Eran Tromer. 2006. Cache Attacks and Countermeasures: The Case of AES. In *CT-RSA (LNCS, Vol. 3860)*. Springer, 1–20.

[47] PCI-SIG 2006. *PCI Express Base Specification*. PCI-SIG. Rev. 2.0.

[48] PCI-SIG 2009. *Address Translation Services*. PCI-SIG. Rev. 1.1.

[49] Christoph Peglow. 2020. *Security analysis of hybrid Intel CPU/FPGA platforms using IOMMUs against I/O attacks*. Master's thesis. University of Lübeck.

[50] Peter Pessl, Daniel Gruss, Clémentine Maurice, Michael Schwarz, and Stefan Mangard. 2016. DRAMA: Exploiting DRAM Addressing for Cross-CPU Attacks. In *USENIX Security Symposium*. USENIX Association, 565–581.

[51] Antoon Purnal, Furkan Turan, and Ingrid Verbauwhede. 2021. Prime+Scope: Overcoming the Observer Effect for High-Precision Cache Contention Attacks. In *CCS*. ACM, 2906–2920.

[52] Antoon Purnal, Furkan Turan, and Ingrid Verbauwhede. 2022. Double Trouble: Combined Heterogeneous Attacks on Non-Inclusive Cache Hierarchies. In *USENIX Security Symposium*. USENIX Association, 3647–3664.

[53] RedHat, Inc. 2014. */sys/class/iommu/<iommu>/devices/*. RedHat, Inc. https://www.kernel.org/doc/Documentation/ABI/testing/sysfs-class-iommu

[54] Thomas Ristenpart, Eran Tromer, Hovav Shacham, and Stefan Savage. 2009. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *CCS*. ACM, 199–212.

[55] Daniel Sánchez and Christos Kozyrakis. 2011. Vantage: scalable and efficient fine-grain cache partitioning. In *ISCA*. ACM, 57–68.

[56] Michael Schwarz, Moritz Lipp, Daniel Moghimi, Jo Van Bulck, Julian Stecklina, Thomas Prescher, and Daniel Gruss. 2019. ZombieLoad: Cross-Privilege-Boundary Data Sampling. In *CCS*. ACM, 753–768.

[57] Debendra Das Sharma. 2019. *Compute Express Link*. Whitepaper. Compute Express Link Consortium.

[58] Anton Shilov. 2022. Intel's Sapphire Rapids Formal Launch Date Revealed. https://www.tomshardware.com/news/intel-sapphire-rapids-launch-date-revealed Access: 2020-12-04.

[59] Olin Sibert, Phillip A. Porras, and Robert Lindell. 1995. The Intel 80×86 processor architecture: pitfalls for secure systems. In *IEEE S&P*. IEEE, 211–222.

[60] Mingtian Tan, Junpeng Wan, Zhe Zhou, and Zhou Li. 2021. Invisible Probe: Timing Attacks with PCIe Congestion Side-channel. In *IEEE S&P*. IEEE, 322–338.

[61] Andrei Tatar, Daniël Trujillo, Cristiano Giuffrida, and Herbert Bos. 2022. TLB;DR: Enhancing TLB-based Attacks with TLB Desynchronized Reverse Engineering. In *USENIX Security Symposium*. USENIX Association, 989–1007.

[62] Stephan van Schaik, Alyssa Milburn, Sebastian Österlund, Pietro Frigo, Giorgi Maisuradze, Kaveh Razavi, Herbert Bos, and Cristiano Giuffrida. 2019. RIDL: Rogue In-Flight Data Load. In *IEEE S&P*. IEEE, 88–105.

[63] Pepe Vila, Boris Köpf, and José F. Morales. 2019. Theory and Practice of Finding Eviction Sets. In *IEEE S&P*. IEEE, 39–54.

[64] Zane Weissman, Thore Tiemann, Daniel Moghimi, Evan Custodio, Thomas Eisenbarth, and Berk Sunar. 2020. JackHammer: Efficient Rowhammer on Heterogeneous FPGA-CPU Platforms. *IACR TCHES* 2020, 3 (2020), 169–195.

[65] Jan Wichelmann, Ahmad Moghimi, Thomas Eisenbarth, and Berk Sunar. 2018. MicroWalk: A Framework for Finding Side Channels in Binaries. In *ACSAC*. ACM, 161–173.

[66] Yuval Yarom and Katrina Falkner. 2014. FLUSH+RELOAD: A High Resolution, Low Noise, L3 Cache Side-Channel Attack. In *USENIX Security Symposium*. USENIX Association, 719–732.

[67] Ying Ye, Richard West, Zhuoqun Cheng, and Ye Li. 2014. COLORIS: a dynamic cache partitioning system using page coloring. In *PACT*. ACM, 381–392.
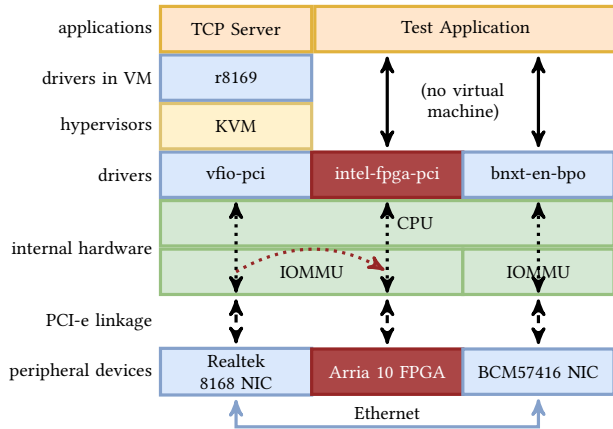
Figure 7: Stack diagram of the network card side-channel test. The Realtek network interface card (NIC) is "passed through" to a virtual machine with the VFIO driver. The test application exchanges packets with the TCP server in the virtual machine over the ethernet connection between the two network cards; meanwhile, the FPGA (connected to the same IOMMU as the VM's network card) probes the IOTLB for traces of network activity.

[68] Yinqian Zhang, Ari Juels, Michael K. Reiter, and Thomas Ristenpart. 2014. Cross-Tenant Side-Channel Attacks in PaaS Clouds. In *CCS*. ACM, 990–1003.

A APPENDIX

A.1 Initial IOTLB Organization Hypothesis

Initially, we hypothesized that the IOTLB would be organized like the CPU TLBs reverse-engineered in [18], with 2^s sets where s is an integer, some small number of ways per set, and a set mapping algorithm wherein the lowest s bits of the page address select the set number or some other combination of various bits of the page address forms the set number that the page is associated with. Initial experiments on all three systems showed that 128-address eviction sets of any randomly allocated pages reliably evicted any other single page, so we hypothesized that the IOTLB was organized with 128 sets and 1 way. We tested this hypothesized eviction set architecture in a scenario on *a10l* where the FPGA used Prime+Probe to monitor an IOTLB that it shared with a network card.

Fig. 7 shows the hardware and software setup for this test, an example of threat model *1u*. A virtual machine is configured with the IOMMU in a pass-through mode (Virtual Function I/O or VFIO) to allow a Realtek 8168 NIC direct access to the virtual environment, where it uses the standard r8169 drivers. The test application runs directly on the host, and uses the Broadcom BCM57416 NIC to exchange packets with the Realtek NIC over ethernet. The test application also manages our Prime+Probe hardware on the Arria 10 GX FPGA and uses it to collect IOTLB side-channel traces while the network is active. The eviction sets used in the Prime+Probe tests are constructed under the assumption that the IOTLB contains 128 sets of one way each.

Figure 8: Behavior is consistent after a reboot of the virtual machine shown in Fig. 7, but inconsistent between reboots; this graph shows the likelihood that an IOTLB entry will be consistently evicted by a Prime+Probe after a reboot of the system. Entries marked in red are evicted whether or not there is network activity and do not vary between reboots; entries in blue are those that are evicted when there is network activity but not when there is no activity and vary significantly.

Prime+Probe data from this experiment are visualized in Fig. 8. There was substantial variation of IOTLB activity after a reboot of the virtual machine operating the Realtek NIC, so results are plotted as means across many reboots. More evictions were detected in the probes of the Prime+Probe while the network was active, indicating a side-channel leakage in the IOTLB that originated from the Realtek NIC. There are two other phenomena of note that are observable in the data from this experiment. First, the excess evictions caused by the network activity (shown in blue in the figure) varied substantially in the number of sets they occupied. Whenever the virtual machine was rebooted, the number of sets that were evicted during network activity changed, but there were always evictions in one set (set 11). After examining the network driver source code, we found that it allocates the transaction buffers used by the network card by calling a kernel function `dma_map_single` on startup, and we verified that by unloading and reloading the network driver, we could reproduce the randomizing effect of rebooting the virtual machine. Second, sets 1-10 and 126-128 were always evicted in the probe, even absent any network activity or with the network drivers unloaded. This showed that the 128-page eviction sets, while effective in evicting IOTLB entries, were actually bigger than necessary, since they were evicting their own members.