

The Future of FPGA Acceleration in Datacenters and the Cloud

CHRISTOPHE BOBDA, JOEL MANDEBI MBONGUE, University of Florida

PAUL CHOW, MOHAMMAD EWAIS, NAIF TARAFDAR, JUAN CAMILO VEGA, University of Toronto

KEN EGURO, Microsoft

DIRK KOCH, Manchester University

SURANGA HANDAGALA, MIRIAM LEESER, Northeastern University

MARTIN HERBORDT, HAFSAH SHAHZAD, Boston University

PETER HOFSTE, IBM POWER Systems Performance

BURKHARD RINGLEIN, IBM Research Europe

JAKUB SZEFER, Yale University

AHMED SANAULLAH, Red Hat, Inc

RUSSELL TESSIER, University of Massachusetts Amherst

In this manuscript, we survey existing academic and commercial efforts to provide Field-Programmable Gate Array (FPGA) acceleration in datacenters and the cloud. The goal is a critical review of existing systems and a discussion of their evolution from single workstations with PCI-attached FPGAs in the early days of reconfigurable computing to the integration of FPGA farms in large-scale computing infrastructures. From the lessons learned, we discuss the future of FPGAs in datacenters and the cloud and assess the challenges likely to be encountered along the way. The manuscript explores current architectures and discusses scalability and abstractions supported by operating systems, middleware and virtualization. Hardware and software security becomes critical when infrastructure is shared among tenants with disparate backgrounds. We review the vulnerabilities of current systems and possible attack scenarios, and discuss mitigation strategies, some of which impact FPGA architecture and technology. The viability of these architectures for popular applications is reviewed, with a particular focus on deep learning and scientific computing. This work draws from workshop discussions, panel sessions including the participation of experts in the reconfigurable computing field, and private discussions among these experts. These interactions have harmonized the terminology, taxonomy and the important topics covered in this manuscript.

CCS Concepts: • **Hardware** → **Reconfigurable logic and FPGAs**; • **Computer systems organization** → *Cloud computing*; • **Security and privacy** → *Systems security*.

Additional Key Words and Phrases: Cloud, Datacenter, FPGA, Virtualization, Security

Authors' addresses: Christophe Bobda, Joel Mandebi Mbongue University of Florida; Paul Chow, Mohammad Eweis, Naif Tarafdar, Juan Camilo Vega University of Toronto; Ken Eguro Microsoft; Dirk Koch Manchester University; Suranga Handagala, Miriam Leeser Northeastern University; Martin Herbordt, Hafsah Shahzad Boston University; Peter Hofste IBM POWER Systems Performance; Burkhard Ringlein IBM Research Europe; Jakub Szefer Yale University; Ahmed Sanaullah Red Hat, Inc; Russell Tessier University of Massachusetts Amherst.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

1936-7406/2022/2-ART \$15.00

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

ACM Reference Format:

Christophe Bobda, Joel Mandebi Mbongue, Paul Chow, Mohammad Ewais, Naif Tarafdar, Juan Camilo Vega, Ken Eguro, Dirk Koch, Suranga Handagala, Miriam Leeser, Martin Herbordt, Hafsah Shahzad, Peter Hofste, Burkhard Ringlein, Jakub Szefer, Ahmed Sanaullah, and Russell Tessier. 2022. The Future of FPGA Acceleration in Datacenters and the Cloud. *ACM Trans. Reconfig. Technol. Syst.* 1, 1 (February 2022), 44 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Since Microsoft published their work on Catapult in 2014 [183], Field-Programmable Gate Arrays (FPGAs) have become more than an exotic and niche technology that can only be tamed by the magicians of hardware. Major players such as Alibaba [35, 63], Amazon [22], Baidu [36], Huawei [53], and Tencent [46] now expose FPGAs to application developers in their datacenter infrastructures. Others use FPGAs to provide applications as a service (Microsoft [83] and Nimbix [43]) or for internally-developed applications. Furthermore, a growing number of projects are underway across the globe, in academia and other research organizations, to provide the benefit of acceleration and flexibility remotely to users. While FPGAs are increasingly available in datacenters, their long-term adoption in these venues is not guaranteed. Several technologies must be provided for FPGAs to compete with their peers – central processing units (CPUs), graphics processing units (GPUs) – and the flood of artificial intelligence (AI) specific processors currently deployed and in development.

Architecture and integration at the board and the system level must ensure that applications harness the strengths of reconfiguration and mitigate its weaknesses. The provisioning of hardware resources requires adequate middleware, hardware virtualization, and domain separation mechanisms. Efficient and flexible hardware resource provisioning increases the possibility of seamless integration, ideally with no redesign, in existing cloud management infrastructure. The spatial sharing of FPGAs increases the complexity of cloud schedulers and resource managers. Fortunately, the large body of work developed in the last decade can be leveraged in cloud task schedulers currently in use in datacenters, such as the list-scheduling used in Amazon cloud[1].

With multiple users temporally or spatially sharing FPGAs in datacenters, the security of designs must be guaranteed. It has recently been demonstrated that FPGAs can be used covertly for various types of cybersecurity attacks in the cloud [186, 187, 218, 219]. Mitigation strategies have been proposed [144, 181] and many more are currently in development. However, despite all best efforts, current FPGA architectures may still be resistant to protection against cloud-based attacks.

Without applications, any architecture incorporating FPGAs will be of no use. Programming languages and high-level design environments as well as efficient hardware/software mapping strategies are necessary to facilitate the transition for software engineers. In this manuscript, experts from academia and industry evaluate the road previously traveled, and the successes and failures of current solutions. We open a window on the future of FPGA-accelerated datacenters, identify opportunities and challenges, and discuss the path to success and broad adoption of FPGA technology for computing.

The remainder of the paper is organized as follows. The next section briefly describes the landscape of FPGA use in cloud infrastructure. A taxonomy of FPGA architecture integration is provided in Section 3 followed by a detailed discussion of current architectural development and trends. The integration of FPGA resources into operating systems (OS) and software is discussed in Section 4. Topics include shell interfaces, middleware, and virtualization. In Section 5, we explore challenges in hardware and system security in single-user (tenant) and multi-tenant cloud compute environments. A critical review of current FPGA architecture is provided and suggestions to increase resiliency are offered. In Section 6, we review applications that can benefit from cloud implementation. Deep learning is the main candidate, but image processing, sorting, and database

operations can also benefit. We conclude the paper in Section 7 with a summary of the opportunities and challenges of the paradigms discussed.

2 THE LANDSCAPE

After decades of single-thread driven performance growth for computing platforms, followed by more than a decade where the majority of performance growth resulted from increases in the numbers of hardware threads, we have now entered a period where the majority of hardware-based computational performance growth is expected to come from specializing hardware [108, 127, 196]. The most power-constrained platforms, such as cell phones, already have multiple types of specialized processing elements including GPUs, neural processors, and specialized signal processing, media, and security engines. While some think of ARM [33] as an architecture that has instruction set architecture (ISA)-driven efficiency, the main reason ARM is gaining traction beyond cell phones and embedded devices is that ARM provides a system-on-chip (SoC) ecosystem that is thus far unmatched [111], enabling the rapid construction of more specialized and efficient SoCs.

One of the earliest successes in using FPGAs in the cloud was the Catapult project at Microsoft being used for Bing searches [183]. This is an example of “provider application as a service” where the user is not aware that FPGAs are being used to accelerate their applications. Subsequently Microsoft provided acceleration for Azure customers by offloading machine learning and host networking to hardware [105]. In the last few years, FPGAs in the cloud have been used for a variety of different applications, see Section 6, including the acceleration of networking, privacy and security, and machine learning and data analytics. As processing becomes more heterogeneous, FPGAs stand out as accelerators that can process data directly from the network and provide benefits to users with or without their knowledge. Thus their use is likely to grow dramatically in the future.

Of course, specialization is challenging primarily because it intrinsically trades off flexibility for efficiency, primarily measured in reduced chip area and improved energy per computation at equivalent performance. This approach requires allocating hardware resources to the “right” priorities for the platform and also requires ensuring that the platform comes with a software development and runtime ecosystem that ensures the specialized elements are used to a sufficient degree to justify allocating the resources to these specialized rather than general-purpose functions. Platforms that provide more extensive and prescriptive application development environments will be in the lead in providing such specialization.

To date, most cloud providers have focused on more general-purpose platforms, renting out standardized scalable infrastructure. Offerings are primarily differentiated on the number of (virtual) cores, and the amount of memory and network bandwidth per CPU and storage. Offerings with more specialized hardware, primarily GPUs, are mostly separate specialized offerings, e.g., for high-performance computing (HPC), rather than elements that provide the broad underpinnings of a cloud platform architecture.

A number of cloud providers have started to use specialized hardware (including FPGAs) [22, 35, 36, 43, 46, 53, 83] and even started to develop their own SoCs [108, 196] rather than rely on vendor silicon, which opens the door to increased hardware-based platform differentiation. It is envisioned that as cloud infrastructure progresses from infrastructure as a service (IaaS) to virtualized platforms and then to microservices and function as a service (FaaS), cloud computing may also provide the software ecosystems that enable the increased introduction of specialized hardware.

In the creation of specialized hardware, reconfigurable logic occupies an interesting middle ground. Reconfigurable logic can be leveraged to create a wider set of specialized computational

elements from the same hardware that each can outperform and improve efficiency over a general-purpose processor for a specific task or set of tasks. Many systems combine reconfigurable offerings with highly flexible I/O, allowing a very large number of system configurations. At the same time, for any fixed configuration of the FPGA supporting a fixed set of tasks, hardware that is not (as) reconfigurable can be implemented significantly more efficiently.

Even if reconfigurable logic becomes ubiquitous, it still may not be programmed by a wide audience. This does not necessarily mean the logic is not regularly reprogrammed. The logic configuration is provided by the system vendor or cloud provider, and not by independent software vendors or end users. There are a number of well-known reasons for this: the difficulty of providing and maintaining a software ecosystem, difficulty of ensuring system integrity, difficulty of ensuring security, and the skills required to program the logic.

Before delving more into considerations specific to the deployment of reconfigurable logic in clouds and datacenters, we take stock of some of the different ways in which reconfigurable logic can be deployed and introduce some of the challenges within these contexts. Perhaps the tightest way to integrate reconfigurable logic in a processor is to use it as a reprogrammable execution unit sharing a register file with the remainder of the core. While not a common use of reconfigurable logic today, contemplating such a use illustrates some of the challenges. For example, a unit could be configured to implement new instructions before they are directly supported in custom hardware. In this case the software ecosystem is essentially the same as for any other new (set of) instruction(s). A more ambitious step would be to allow reprogramming and changing of the performance profile but not the set of supported instructions. This approach requires a much more complex software ecosystem, as compilers tend to be highly tuned to performance specifics. This approach is also likely to require OS-level coordination, as schedulers would likely have to become aware of these profile changes to ensure the right tasks are scheduled to the processors with the appropriate performance profile. A next step might include support for different profiles with different (vendor-defined) sets of instructions. This step now puts an even larger burden on the software ecosystem but still leaves the responsibility with the hardware vendor. A final step would be to allow cloud provider or end-user reprogramming of the embedded reconfigurable logic to create user-defined instructions. Architectures such as RISC-V [8] and OpenPOWER [7] explicitly allow for the introduction of implementation-specific instructions. It is not difficult to contemplate a cloud-specific use of such a capability, for example to improve the performance of a specific encryption or decryption algorithm that does not already have direct hardware support.

A first approach might be to treat a processor with a customer-specific instruction as a customer-specific processor to be verified in its entirety by the customer as if they had built a new processor with a customer-specific instruction. Such an approach might not be satisfactory, but with cloud providers now building their own SoCs it is perhaps not unrealistic. Beyond the software ecosystem and system functionality, system integrity and system security are also issues, even for this very restricted way of introducing reprogrammable functions. For example, one will want to ensure that reprogramming cannot result in physical damage to the system or affect reliability. The issue is not far-fetched, as FPGAs can be configured with collections of ring oscillators with local power densities that exceed normal design constraints [181]. Security could also be a concern, as one would have to be very careful that a custom execution unit does not observe (and store) any information from another user context or a context at a different privilege level. Clearly, much work remains before general-purpose user-reprogramming of a reconfigurable unit integrated at this level can be commonplace.

Next, we consider a reconfigurable unit that shares memory with other processors in the system, either physically integrated in the same chip or on the same module or provided separately. Bus protocols range from architecture-specific (QPI or UPI for example) to open standards such as

OpenCAPI [174], AXI [32], CCIX [206] and CXL [92]. Standards at the protocol level have to be accompanied by software standards at the system level. One common approach is for the native OS (or hypervisor) to retain control over system resources and manage memory access to shared memory. This provides a clear foundation for system organization, scheduling and security, as it can build on the knowledge from other systems with heterogeneous processing elements. However, there are special considerations for processing elements built from reconfigurable logic. To maintain system integrity, the address translation and/or bus interface units typically must remain the responsibility of the processor or system vendor to ensure the accelerator operates in a virtual (effective) address space instead of having physical addressing capabilities. This goal can be achieved if the interfaces to the reconfigurable logic operate in a virtual (effective) address space, or any reconfigurable logic provided by someone not responsible for the physical infrastructure is combined (and verified to be combined) with logic that is. This can also raise an issue of trust. For example, if the only way to guarantee integrity is to allow the system provider to do extensive checking of the logic during the build process, an end user or independent software vendor who wants to differentiate based on the reconfigurable logic they provide may be forced to expose that logic to the system provider.

Another level is integration as a device, either for computational acceleration or to enhance network or storage capabilities. A standard device does not have access to all system resources, and it is easier to delineate the responsibilities. CXL supports device and shared memory accelerators by providing a variety of protocol profiles. A typical use of reconfigurable logic in this context might be the creation of “smart” storage or network devices, providing functions like security, data compression or decompression or filtering. Smart networking or smart storage devices that allow the combination of user-created reconfigurable logic within a network or storage controller do again raise a set of concerns that must be addressed.

Finally, perhaps the most ambitious approach would be to develop a datacenter that consists of (essentially) only FPGAs. Such an approach is made easier because FPGAs integrate increasingly powerful CPU cores, (stacked) DRAM memory, and high-speed I/O. The software foundation for such an infrastructure would build on the conventional notions underpinning conventional system design.

3 FPGA CLOUD ARCHITECTURES

The computational capability and flexibility of a reconfigurable fabric, coupled with the performance of ASICs for high speed I/O (e.g., SERDES units), allows FPGA-based systems to be deployed almost anywhere in the cloud to accelerate compute, networking and storage. An example of this versatility is illustrated in Figure 1, which shows two common, contrasting deployment models. However, while the potential for arbitrary deployment certainly exists, FPGAs cannot be randomly placed as part of any cloud infrastructure. The need for cost-effectiveness leads to an emphasis on size, power, cooling, compatibility, and in-place upgradability - all while ensuring that the specific performance, memory, server capability, and network connectivity needs of cloud workloads are met. Thus, based on a system’s requirements, certain types of cloud FPGA architectures can have substantially more benefit than others.

This section discusses architectural choices for the integration of FPGAs into cloud computing infrastructures. Current architectures in use in academia and industry have distinctive characteristics, the most important of which are listed below, that have helped shape them over the past two decades.

- Very large scale, with (potentially) millions of nodes
- Mature, having evolved through many generations

- Highly competitive, with several vendors offering similar products, yet with high value in differentiation
- Diverse evolving workloads, but with some emerging high performance applications
- System upgrades are always partial placing a premium on compatibility through several generations
- Heterogeneity, with different sets of nodes for different types of users

We hypothesize that future organizations will only slightly evolve from existing architectures, depending on the devices and integration technology. A radical break is not expected. In the next subsection, we begin with a taxonomy of system organizations, followed by a review and discussion of existing architectures.

3.1 Taxonomy

This subsection presents a set of high level attributes that define cloud FPGA systems. These taxonomic categories provide an effective means to not only analyse the relationships between design choices and system constraints; but also to describe and compare different cloud FPGA architectures. We continue with an overview of mainstream FPGA cloud systems, both in production and research, and conclude by analysing architectural trends in these systems in the context of the proposed taxonomy.

The taxonomy is based on the critical aspects of any cloud FPGA system: **A) Type of FPGA boards**, **B) Placement of FPGAs in the system**, **C) Network Connectivity**, **D) Intra-node Connectivity**, and **E) Use Cases**. Note that these taxonomic categories are neither mutually exclusive nor comprehensive: it is possible for system to have multiple sub-categories, and for new sub-categories to be added later to encapsulate future innovation.

Figure 2 illustrates common architectures. These architectures, and a sample benefit, are as follows:

- (1) **Bump-in-the-Wire**: Large scale compute, network and storage acceleration
- (2) **Co Processor**: Local compute acceleration
- (3) **Storage attached**: Local storage acceleration
- (4) **Back-end**: Ultra low latency, rack scale FPGA-FPGA communication
- (5) **Smart NIC**: Local network acceleration
- (6) **Network HW**: Flexible routing/switching protocols
- (7) **Local Cluster**: Multi-accelerator system
- (8) **Shared Memory**: Cache coherent acceleration
- (9) **Disaggregated**: High infrastructure utilization

3.1.1 Type of FPGA boards. Given that cloud providers do not currently create their own FPGAs, the smallest unit of differentiation is the FPGA board; both **a) Off-the-shelf** and **b) Custom** are possible. Economic advantages depend on the scale of deployment and provider development infrastructure. Given the latter, custom boards still have higher start-up and upgrade costs, but may be cheaper in large quantities. But the advantage of scale also affects off-the-shelf economics as the provider has the market power to affect price and board features.

With custom FPGA boards, just about any attribute can be varied, such as number/types/bandwidths of I/O ports, FPGA family, off-chip memory type and size, form factor, and other on-board devices.




Fig. 1. Example FPGA placement in datacenters

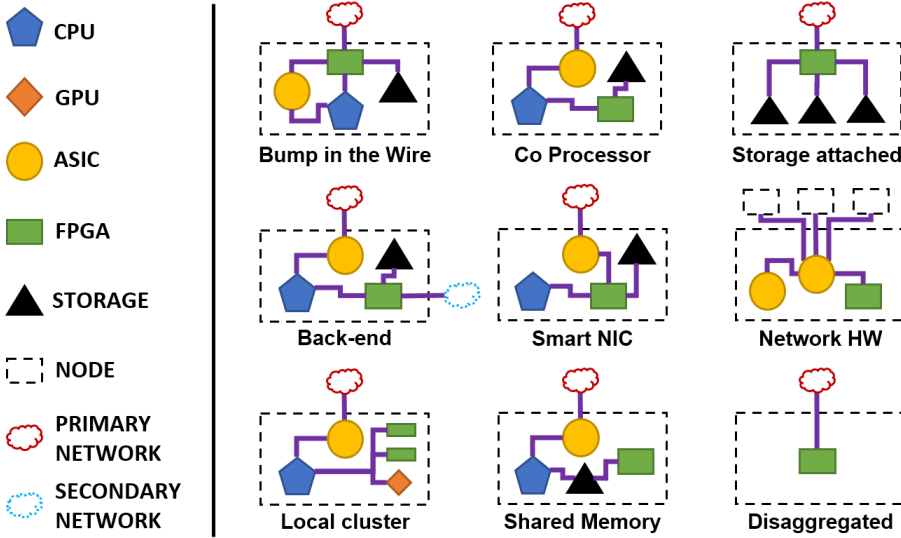


Fig. 2. Common FPGA architectures. Note that Accelerator Model and Peer Model from Figure 1b are illustrated here as Co Processor and Disaggregated, respectively

This ensures that the boards closely match the specifications/requirements of the target system, from computation to cooling. None-the-less, off-the-shelf boards are available for every (currently) sizable usage domain, including SoCs [42, 49], node-level networking [61, 62, 65, 132, 165, 169, 202], NoCs [26], data center switches [40, 41, 171, 198] and storage [37, 45, 48, 56, 98].

3.1.2 Placement of FPGAs. FPGAs can be placed in either **a) distributed** or **b) centralized** manner. Having a distributed FPGA placement means that compute/storage nodes have their own FPGAs, and thus do not have to compete for the resource. This leads to more offload capability, greater reliability since FPGA failure does not affect other compute/storage nodes, and reduces security concerns since offloads for different nodes can be isolated. It is also possible to place FPGAs in a centralized manner, typically inside the networking nodes (e.g. in switches as ASIC-FPGA, CPU-FPGA, or FPGA only circuits). Substantially fewer FPGAs are needed for such a deployment; this typically translates to easier management, lower power consumption, lower total cost of ownership (TCO), smaller average node sizes, and potentially higher performance since expensive high-end FPGAs can be used (and upgraded more frequently).

3.1.3 Network connectivity. Within each node, it is possible for FPGAs to be **a) not connected to any network** or connected to **b) the primary network** and/or **c) a secondary network**. Being connected to the primary data center network enables FPGAs to intercept/accelerate network traffic to the node, as well as achieve data-center-wide scalability for FPGA workloads since multiple FPGAs can directly communicate with each other. However, the circuitry needed to support this FPGA position can consume a significant portion of FPGA resources. This includes circuits to support high resiliency since the FPGAs can be a single point-of-failure: an entire node can become unstable in the case of an FPGA error. In the case of secondary network connectivity, FPGAs can communicate across nodes with significantly more flexibility in the topology used (e.g. mesh, torus, switched), as well as the communication protocol, all of which can lead to ultra low latencies. However, using a custom network configuration means that complex router hardware, routing algorithms, and switch arbitration policies may need to be implemented on each FPGA. Moreover,

complex cabling may be required, which can add a significant burden to the overall data center architecture [83].

3.1.4 Intra-node connectivity. FPGAs within a node can be **a) not connected to any other significant device** (i.e., be a *Disaggregated* resource) or connected to one or more devices: **b) connected to CPUs**, e.g., through PCIe and possibly with cache coherence using interconnects such as CCIX [206], CXL [92], or CAPI [174]; **c) connected to other FPGAs**, e.g. through a PCIe switch and/or using direct and programmable interconnects; **d) connected to GPUs**, e.g. through a PCIe switch; **e) connected to ASICs**, e.g., through multiple potential forms of connectivity depending on the ASIC and nature of coupling such as a NIC or tensor processor; **f) connected to storage devices** through the device-specific interface, e.g., SPI for flash and DDR for SDRAM.

3.1.5 Use cases. Use cases have a substantial impact on architecture, since cloud providers must ensure workload requirements are met (e.g., performance) without compromising on core aspects (e.g., security, reliability). Here we look at some common cloud FPGA use cases. **a) Customer applications:** Customers can develop, simulate, debug and compile their custom FPGA logic, as well as scale their infrastructure and change resources according to their workload demands. A wide pool of applications can be deployed, e.g., in genomics, financial analytics, computational fluid dynamics, video processing, transcoding, and security. Several development environments are available so users do not need to write their own HDL code [69]. **b) Provider Application as a Service (AaaS):** The cloud provider supports a limited set of customer applications by developing the FPGA design themselves: only the necessary APIs and high level design parameters are exposed. This model ensures high performance and resilience at the expense of reducing customer access to the entire FPGA. **c) Provider applications:** In this case cloud providers use FPGAs to accelerate their internal workloads, e.g., software defined networking (SDN), as well as save CPU resources that can then be rented to the customer.

3.2 Production Architectures

In this subsection we discuss what can be referred to as *core* or *mainstream* production cloud FPGA systems that are either in widespread or large-scale use.

Perhaps the most unique and widely deployed production system is Microsoft's Catapult v2 [83], which has FPGAs in most Azure and Bing stock keeping units (SKUs) in a *bump-in-the-wire* configuration: an FPGA sits between the top-of-rack switch (TOR), network interface card (NIC), ASIC and CPU, hence enabling data-center wide communication within tens of microseconds of latency. These hundreds of thousands of FPGAs (or more) are used for both internal (e.g. network packet processing [166] and Bing search [89]) and external workloads (e.g., machine learning inference as a service [89]).

Another type of a widespread production system is the single node accelerator model, which leverages FPGAs in either a *Co-Processor* configuration, or as a *Local Cluster* where devices are connected either via a PCIe switch or using direct FPGA-FPGA interconnects. A number of cloud providers such as AWS [22], Huawei [53], Baidu [36], Tencent [46], Nimble [43] and Alibaba [35] use this model. These systems are used by customers to run a wide pool of cloud native applications such as genomics, financial analytics, data acquisition, computational fluid dynamics, video processing, image processing, transcoding, security, and AI workloads [50, 57, 59, 60, 63, 220]. There are also examples of these FPGAs being used by providers for their own workloads. Baidu uses FPGAs to accelerate its cloud based storage, SQL queries, data security, search engine, and AI workloads [50, 55]. FPGA-based AI chips—such as Baidu's Kunlun for AI, Alibaba's Ouroboros for speech recognition, and Alibaba's Hanguang 800 for inference operations—are deployed in their cloud data centers [101]. Alibaba has reported 75% savings in TCO by using FPGAs to oversee product images

on its e-commerce site [235]. In 2018 it reported over \$30 billion retail on its website in a single day (compared to \$5 billion on all US online and in-store retail on Black Friday 2017); this was possible with its data center FPGAs being used to accelerate transactions and provide recommendations to users [106].

There are also systems that are widely deployed, but where there is insufficient publicly available information for analysis. Amazon has announced AQUA (Advanced Query Accelerator) nodes for its Redshift data warehouse, available through the RA3.16XL and RA3.4XL instances. These nodes use FPGAs to accelerate dataset filtering and aggregation [38, 58]. Baidu uses Smart NICs to improve virtualisation and workload performance [66]. OVHCloud also uses Smart NICs, but for network packet processing to mitigate distributed-denial-of-service (DDoS) attacks in its cloud traffic [44, 51]. Scaleflux CSD2000 SSDs are deployed by over 40 data centers globally [39]. An example is the Alibaba cloud, which uses Scaleflux CSD20004 in place of traditional solid state drives (SSDs) on their storage nodes to accelerate applications such as MySQL, Aerospike, Oracle, and PostgreSQL [54]. Samsung Smart SSDs [48] are deployed in the Nimble cloud where they accelerate Apache Spark, running queries up to 6× faster when using software from Bigstream [47]. Eideticom's computational storage processor [98] has been implemented in Bareflye G2 servers on Rackspace [180].

3.3 Research Architectures

In this subsection we discuss systems that are presently in research and development and represent the most technologically plausible candidates for widespread future deployment.

One of the most commonly used research architectures is the cluster of *Back-End* tightly coupled FPGAs that deploy a secondary network using direct and programmable interconnects to connect FPGAs across nodes. Microsoft's Catapult v1 was a back-end system that connected multiple nodes in 6x8 tori [183]. It was demonstrated using Microsoft's Bing workloads; it is not clear whether it was ever part of a production cloud. Other research examples include Maxwell [70], Novo-G# [112], Noctua system at the Paderborn Center for Parallel Computing [64, 176, 179], and Albireo nodes of the Cygnus supercomputer system at University of Tsukuba [2]. Although this approach can substantially reduce FPGA-FPGA latency, it is difficult to scale beyond a single rack due to wiring requirements; in the general case it also requires each FPGA to implement a router to support the communication. Currently no such example can be found operating in the production cloud.

Another research area proposed in [109, 141] involves *Channel-over-Ethernet* (CoE); a back-end, inter-FPGA Ethernet communication network using the OpenCL kernel programming. The results demonstrate the feasibility of such a configuration as the system achieves a latency of 0.99μs for inter-FPGA communication via the secondary Ethernet switch as compared to 29.03μs via the host CPU. A drawback is that data are sent as packets and that there is additional overhead, such as IP addresses and flags, that reduce the effective data rate [226].

Other research architectures include systems that support a *Local Cluster*, but where the communication scaling via direct interconnects is limited to a single node. An example includes Novo-G (a former version of Novo-G#) [113]. Other examples include the research systems currently deployed at the IBM SuperVessel Cloud [4] and the IBM Power8+CAPI cluster at the University of Texas, Austin [52] that use a *Shared Memory* cache coherency model.

A different approach is to directly connect FPGAs to the datacenter network as a standalone resource. Each FPGA can be accessed by a CPU or another FPGA resulting in good scalability. CloudFPGA by IBM Zurich Research Lab is one example [190, 193, 231–233]. The authors have built a prototype with multiple chassis for data center scale, capable of hosting 1024 FPGAs per rack [25]. A drawback of such an architecture may be that FPGA-CPU communication is necessarily among

separate nodes and has high latency. Another consideration is the increase in the number of TOR connections.

The University of Toronto SAVI testbed connects FPGAs to the primary network [212]. The authors in [82, 211, 213] have demonstrated that virtualising FPGA resources on the SAVI testbed enables multiple regions within an FPGA device to support different designs using APIs such as OpenStack. Enzian [3] at ETH Zurich employs an FPGA as a node connected to the network on one end and coherently attached to a large server-class SoC on another node. Unlike Microsoft's *bump-in-the-wire*, this system allows CPUs to either connect directly to the network or via the FPGA. Unlike other cache coherent systems, it allows the FPGA side of the cache coherency protocol to be extended and tailored [29]. The Gator Reconfigurable Cloud infrastructure (GatorRecc) provides a platform to explore multi-tenancy FPGA usage in cloud applications, using virtual instances (CPU+FPGA) and managed by Openstack[13, 161].

3.4 Architectures Trends

Table 1 classifies production and research cloud FPGA architectures based on the taxonomy in Section 3.1. As we can see, research systems generally explore different varieties of cloud architecture options. While production systems are bounded by factors such as total-cost-of-ownership (TCO), power-usage-effectiveness (PUE), performance, resilience, modularity, scalability and security; research systems tend to enjoy greater degrees of freedom. To effectively analyse these architectures we highlight trends in their architectures based on the following categories: 1) Boards, 2) Placement, 3) Network Connectivity and Use Cases, and 4) Intra-node Connectivity.

Table 1. Classification for production and research FPGA cloud architectures based on the taxonomic categories discussed in Section 3.1. Though all systems place FPGAs in a 'Distributed' manner, the placement column is still shown to highlight this trend. It is further discussed in Section 3.4.2.

	Board Type	Placement	Network Connectivity	Intra-node Connectivity	Use Case
PRODUCTION SYSTEMS					
Alibaba [35]	Custom	Distributed	None	FPGA, CPU, Storage	Customer, Provider
Baidu[36]	Custom	Distributed	None	CPU, Storage	Customer, Provider
Microsoft Catapult v2[83]	Custom	Distributed	Primary	CPU, ASIC, Storage	AaaS, Provider
Amazon AWS F1[22]	Custom	Distributed	None	FPGA, CPU, Storage	Customer
Huawei[53]	Custom	Distributed	None	FPGA, CPU, Storage	Customer, Provider
Nimbix[43]	Off-the-shelf	Distributed	None	CPU, Storage	AaaS
Tencent[46]	Off-the-shelf	Distributed	None	CPU, Storage	Customer
RESEARCH SYSTEMS					
Microsoft Catapult v1[183]	Custom	Distributed	Secondary	CPU, Storage	Research
Enzian[3]	Custom	Distributed	Secondary	CPU, Storage	Research
Cygnus[2]	Off-the-shelf	Distributed	Secondary	FPGA, CPU, GPU, Storage	Research
IBM CloudFPGA[25]	Custom	Distributed	Primary	Storage	Research
Maxwell[70]	Off-the-shelf	Distributed	Secondary	CPU, Storage	Research
NARC[91]	Off-the-shelf	Distributed	None	CPU, Storage	Research
Noctua[176]	Off-the-shelf	Distributed	Secondary	FPGA, CPU, Storage	Research
Novo-G[113]	Off-the-shelf	Distributed	None	FPGA, CPU, Storage	Research
Novo-G#[112]	Off-the-shelf	Distributed	Secondary	FPGA, CPU, Storage	Research
IBM Power8 + CAPI[52]	Off-the-shelf	Distributed	None	FPGA, CPU, GPU, Storage	Research
IBM SuperVessel[4]	Off-the-shelf	Distributed	None	CPU, Storage	Research
SAVI[147]	Off-the-shelf	Distributed	Primary	CPU, Storage	Research

3.4.1 Boards. Table 1 shows that a majority of production cloud vendors have used custom boards in their deployments. For Microsoft in particular, this was necessary since requirements for placing FPGAs in special HPC SKUs "constrained power to 35W, the physical size to roughly a half-height half-length PCIe expansion card (80mm x 140 mm), and tolerance to an inlet air temperature of

70°C at 160 lfm airflow" [83]. Custom boards are not required, however: Nimbix and Tencent both use off-the-shelf boards. For research systems, custom boards are preferred if the proposed systems are *Disaggregated*, network attached (e.g., Enzian), and like IBM CloudFPGA. Also, for earlier *Back-end* systems, like Catapult v1 and Novo-G#, a customised board allowed the system to increase transceiver count. However, we can see that recent *Back-end* and *Local Cluster* systems mostly use off-the-shelf boards. Systems with no inter-node communication network almost always use off-the-shelf boards.

3.4.2 Placement. Table 1 shows that all these systems deploy FPGAs in a distributed fashion. This is because: i) FPGA resources are easier to orchestrate, ii) FPGAs can be offered as bare-metal resources, which simplifies the tooling needed, and iii) FPGA failure only affects local resources, as opposed to potentially millions of nodes.

3.4.3 Network Connectivity - Use cases. Table 1 shows two important trends. First, none of the production systems uses a secondary network. This is likely because of: i) the cost and complexity of wiring a second network for potentially millions of nodes and additional networking hardware, ii) the potentially limited scalability if direct FPGA-FPGA connectivity is supported, and iii) high chip resource usage for building routers and securing the system. The second important trend is the relationship between network connectivity and use cases. Specifically, due to security and reliability constraints, systems that allow customers to offload their own applications do not support any direct network connectivity. Rather, this connectivity is only available if workloads are either internal, or if the offering is an application where only a limited set of APIs are exposed to the customer. Research systems are distributed evenly across the different network connectivity options. We also note that newer research systems almost always have network connectivity, either primary or secondary. This helps scale the application across multiple FPGAs and achieve lower latency.

3.4.4 Intra-node Connectivity. Table 1 shows four major trends. First, in all systems with the exception of IBM CloudFPGA, FPGAs communicate with the CPU over the PCIe slot. This emphasises the role of the CPU as being the core computational resource, whereas the FPGA is a complexity offload engine managed by the CPU. Second, FPGAs are almost always connected to some form of off-chip storage; typically a DDR memory chip on the same board. Third, no production system currently offers instances with FPGA-GPU connectivity. To the best of our knowledge, none of the cloud providers has placed GPUs and FPGAs within the same node. For research systems, GPUs are being employed on the same node as FPGAs, especially for highly parallel, SIMD-like workloads and communication occurs over a PCIe switch [2]. Fourth, in terms of FPGA-ASIC connectivity, only Microsoft v2 supports this approach since the FPGA must transparently process packets for the traditional NIC. None of the research architectures connect an ASIC with an FPGA on the same node.

3.5 Potential Future Innovation

We identify areas of potential novelty that can be derived by traversing the categories in the taxonomy, and by comparing different sub-categories with what is already present in Table 1.

Type of Boards: Potential novelty here is with modular boards that lie at the intersection of custom and commodity. Similar to what is commonly done with micro-controllers, semi-custom boards can be built by buying and connecting together off-the-shelf modules for different FPGA chips, memory chips, and interfaces (QSFP+, PCIe etc). This would allow providers to tailor boards to their specific requirements, reduce the penalties of designing a custom board (development costs, upgrade costs, probability of failure, time to market), and easily replace specific modules as needed (due to hardware failure or for regular upgrades).

Placement of FPGAs: While FPGAs have been used in high end network switches [40], their role is typically limited to providing the performance and flexibility needed to support changing protocols. However, there is currently no system that leverages TOR switches where FPGAs are responsible for implementing the entire switch hardware [41, 171, 198].

Supporting such an architecture has a number of benefits. i) Customer offloads: Customers could use these TOR FPGAs to compute in the network e.g. for doing collective operations such MPI All-Reduce and Broadcast. ii) Provider offloads: Providers could leverage these FPGAs to implement services such as metering, accounting, analytics, and packet filtering. iii) Flexible networking: By combining FPGA based TORs with Bump-in-the-Wire FPGAs, a data-center-wide network could be created that does not rely on a standard protocol for communication. As a result, the communication latency could be reduced substantially. Alternatively, it may be possible to dynamically switch between different standard protocols based on the target workload.

Network Connectivity: A potential novelty here would be to support both Primary and Secondary network connectivity, either within the same FPGA, or through multiple tightly coupled FPGAs within the same node. This would effectively combine key benefits of Microsoft's Catapult v1 and v2, i.e. having ultra low latency for rack scale communications through custom interconnects and still supporting data-center scale FPGA-FPGA connectivity.

Intra-node connectivity and Use cases: The connectivity between FPGAs and CPUs is typically done using the PCIe bus. This is because existing use cases define the role of the FPGA as an offload engine for the CPU. However, a potential novelty here is supporting sufficient low-level electrical coupling, such as the FPGA has read-modify-write access to the CPUs Baseboard Management Controller and firmware. This would effectively turn the FPGA into a management and security controller for the CPU, and enable new *system administrator* use cases such as CPU firmware attestation.

4 USABILITY, SCALABILITY, AND PORTABILITY

Datacenters have become the enablers of many technologies and services. Search engines, personal assistants, streaming, video conferences, 5G and telecommunications, along with the newly popular infrastructure as a service (IaaS) are all examples of such services. With the scale of these applications and services, deployment, provisioning of resources, and isolation all become tricky, especially with the loads of these services becoming more dynamic. For example, in cellular networks there may be a significant increase in peer-to-peer telephone or SMS traffic during holidays, but then during large sporting events there may be a spike in video streaming and cellular data use. Even if the total load is the same, different types of traffic must be processed differently making the load on individual functions dynamic. Furthermore, the principle of elastic computing assumes that instantaneous load is not constant over time [152] creating opportunities to scale down compute resources, putting excess resources in a low power state, or quickly scale up when high usage resumes. Similar examples of dynamic usage profiles exist in most datacenter use cases, and have led to multiple studies into load balancing, traffic control, and usage forecasting to improve performance and minimize energy consumption [68, 136, 173, 201].

We start this section with an overview of the singularity of integrating FPGAs for elastic computing in datacenters. We briefly survey existing commercial systems, then dive into the details of the tools needed to program and deploy FPGA in the cloud. We discuss programability of single node, integration of design in vendor cloud using shell, then virtualization. We then survey current development and discuss the outlook.

4.0.1 Resource Provisioning and Reconfigurability. The problem of datacenter provisioning has motivated deployment and orchestration platforms that allow datacenter managers to monitor their

network and modify in real time the amount of compute resources given to each function. They make use of platforms such as Openstack [199] and Kubernetes [72] that instantiate, provision, and reconfigure datacenter resources in real time. These tools orchestrate the entire datacenter, but require individual components on each node to provision its resources. These tools implement resource provisioning in the cloud in several steps, including admission control, capacity allocation, load balancing, energy optimization and quality of service guarantees. Unfortunately, these tools are currently only limited to provisioning CPU resources. For FPGA-accelerated clouds, they must be extended or re-engineered to include FPGA resource management. Besides the spatial nature of computing in FPGAs, reconfiguration along with bitstream management must be taken into account.

FPGAs are becoming more and more popular within the datacenter, as discussed earlier. But a critical aspect of FPGAs is their configurability, which could play a major role in making them suitable for accelerating such dynamic tasks, as their reconfigurability could naturally allow for the deployment and scaling these applications within the datacenter. Extending the typical datacenter orchestration tools with compatibility to FPGAs is thus natural, as it minimizes the required changes datacenter operators have to sustain. However, the bulk of the work needed for this support will happen on the FPGA virtualization side to make FPGAs compatible with orchestration tools, especially to support FPGAs as primary resources in the datacenter like CPUs, rather than just accelerators. This includes the ability to remotely provision the FPGA resources, split and isolate them between tenants, configure their network interfaces, and virtualize these network connections as needed.

FPGA programming is typically performed via an attached CPU. This is suitable with the current programming model of FPGAs as function accelerators for CPUs but is not suitable for standalone FPGA compute nodes as we envision in a datacenter. Depending on the FPGA type, the current programming interface is usually either JTAG, popular amongst PCIe attached FPGAs, or through an AXI memory mapped interface, as is typically the case in Xilinx MPSoC configurations. If FPGAs are to be primary compute resources in the datacenter, their programming model requires more flexibility.

In a datacenter architecture where part of, or entire, racks are made up of FPGA boards, while others are made up of CPU based servers (Figure 1), it should be possible for independent FPGAs to receive partial or full bitstreams over the network. This ability is not new and has been previously proposed [128, 193, 224], though this would come with, among other things, security challenges that have not yet been overcome, and many of these alternate programming approaches come at the cost of area and or performance.

FPGA network configuration and virtualization can be achieved using a variety of ways, each with its benefits and downsides. The simplest solution would be to statically plan the network connections between FPGAs in a deployment, along with assigning their correct network configurations (i.e., IPs and MAC addresses). Static network configuration is used by various current datacenter FPGA implementations such as Microsoft's Catapult [83, 183] or Galapagos [209, 212, 214]. This requires prior knowledge about the deployment, and lacks the ability to scale the deployment up or down if needed, and the flexibility of deploying these applications alongside others, which may cause conflicts between the network configurations, for example. Another possible solution is to use some common chaining protocols like VXLAN [155], MPLS [67], and Segment Routing (SR) [24]. These protocols may or may not (in the case of VXLAN) require the active participation of the switch as well as requiring modifications to the network interfaces at the compute nodes. These protocols can create virtual network connections over-top the existing physical network architecture. This approach is reconfigurable, allowing a change of the virtual network whenever needed. However, a few drawbacks of these approaches are that they apply additional routing restrictions, latency

overheads, and they require larger packet headers affecting throughput. These protocols could be merged into the underlying FPGA Shells used in the datacenter, but allowing these Shells (along with the switches, if needed) to receive and apply the virtual network configuration will be required. A third approach is similar to that followed by orchestration tools with CPUs today, which configures the network interfaces in a non-virtualized way, but instead has to keep each resource informed about the network interfaces of other resources (i.e., using a DNS table) and update them in the case of redeployment or scaling. This approach is not restrictive, and allows any resource to connect to any other resource within its deployment. However, this approach is more complex in terms of implementation and integration in an FPGA Shell. Regardless of the approach adopted, MPSoC-based FPGAs are very appealing candidates, since the CPU on the MPSoC can act as an accelerator to the FPGA itself, offloading most of the network configuration tasks, along with interfacing with the central datacenter orchestration.

Existing datacenters circumvent these issues in a variety of ways. Microsoft Catapult, for instance, rebuilt the database infrastructure from scratch fitting FPGAs into the picture [83, 183]. However, this system lacks the quick reconfigurability that is possible in other scenarios, and complete rebuilding of datacenters, as they have done, is extremely costly and may be infeasible. Also, the wiring is built and optimized for their physical layout with no virtualization in the node functions or their networking. Instead, it would be ideal to be able to add FPGAs to existing infrastructures and reap their benefits without the upfront cost of a datacenter overhaul for widespread FPGA adoption to work. While this showed significant performance benefits without the orchestration or virtualization, this improvement is limited to one type of application, and can cost more in the long run for changing the supported applications.

Amazon's AWS [22] adds FPGAs to an existing infrastructure. It does not directly connect FPGAs to their datacenter network. Packets instead travel through the software layers of a PCIe-attached CPU before being forwarded to the network, and incoming packets need to traverse this stack in reverse. Similar frameworks are utilized by Huawei [53] and Alibaba cloud infrastructures [229]. This approach resolves all the networking and virtual routing issues, as well as some security issues with network attached FPGAs, although this practice heavily limits throughput and increments latency. Even in full pass-through mode where the CPU is not asked to perform any processing or encapsulation, these architectures have network throughputs capped at about 40 Gbps [229]. However it is possible to reconfigure what each FPGA compute node is capable of doing and can even offer FPGAs as a service by temporarily lending FPGA nodes.

4.0.2 FPGA Programming. Designing for FPGA consists of creating hardware circuits that will operate on the FPGA, using basic resources available on the device. The process uses Hardware Description Languages (HDL) such as VHDL and Verilog. It involves substantial hand and tedious work to implement all of the I/O functionality for the target board for data exchange. The migration of a design to a different FPGA board requires a complete redesign of the I/O handling.

There have been many early attempts to create software-like programming environments that were commercially available circa 2000 such as *Handel-C* [15] and *Impulse C* [215]. These environments used a subset of C with some language extensions, data types and functions that could be synthesized to hardware running on supported FPGA boards. They provided an abstraction that hides most of the hardware complexities from the application developer. Maxeler [11] developed its Data Flow Engine (DFE) data flow programming model based on an extended Java called MaxJ and provided a full stack from the hardware layer to a software runtime and application layer interfaces. Again, the goal was to provide a software development flow that runs on CPU and FPGA platforms. The FPGA vendors have now developed their own programming environments, such as those using OpenCL [16, 20] and are now moving towards more integrated environments that

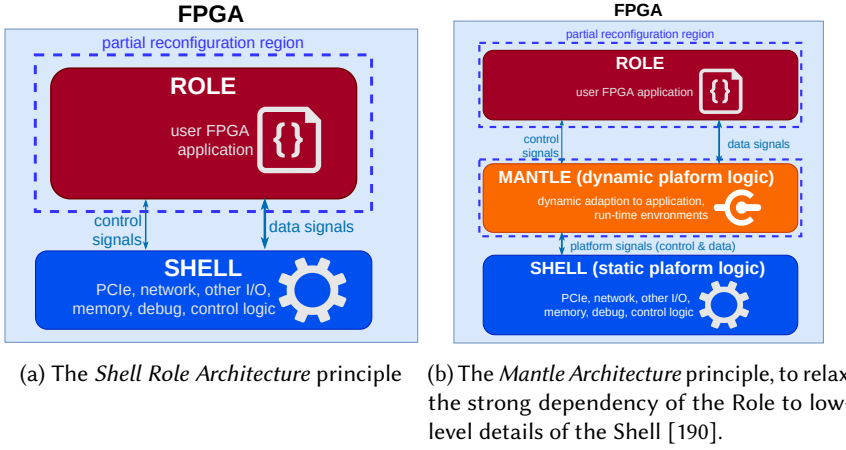


Fig. 3. Shell Role and Mantle Architecture Principle

provide support for several languages and libraries for application spaces such as Vitis [21] from Xilinx and Intel's oneAPI Toolkit [17]. There are also efforts to build more of an open middleware stack such as the Intel Open FPGA Stack and the Open Programmable Acceleration Engine (OPAE) Technology [18].

Supporting computing environments for FPGAs and making them more usable for application developers are important steps, but they still do not enable the use of FPGAs at the scale of clouds and datacenters.

4.0.3 Design Deployment: The Shell-Role Architecture. The emergence of FPGAs as compute accelerators in the Cloud and other multi-user environments inevitably led to a split of the FPGA design into a user application programmed by the developers and a platform-specific part controlled by the infrastructure provider. This split of FPGA logic into a cloud service provider (CSP) controlled *Shell* and a user-controlled *Role*¹ allows the necessary introduction of different privilege levels within an FPGA design and potentially improves re-usability of user applications [193].

The general principle of the Shell Role Architecture (SRA) is depicted in Figure 3a. Usually, the Shell contains all necessary I/O and control logic for data exchange between the user's hardware design and the rest of the system. On the one hand, this abstraction prevents the user from dealing with cumbersome I/O details. On the other hand, the central control of the infrastructure by the CSP simplifies the management and allows the CSP to implement resource allocation, isolation, and necessary security guarantees. Consequently, SRAs allow the CSP to stay in control of the Shell while application logic is controlled by the user [193]. This SRA pattern is widely used across different FPGA platforms [22, 83, 102, 193], each with its own level of details.

However, despite their simplicity and wide-usage, SRAs generate new problems when used at scale: The important character of the interface between Shell and Role "freezes" this interface after its release. To tackle this strong dependency while preserving the advantages of SRAs, recent research developed the notion of a three-layer architecture [190]. The general idea is to introduce an "adapter" layer between Shell and Role to provide forward and backward compatibility between

¹The terms *Shell* and *Role* were first introduced by Putnam et al. [183] and are now commonly used. Sometimes, the term *Hypervisor* is used instead of *Shell* because of a loose correspondence to the functionality of hypervisors in the software systems stack.

Shell and Role. The authors call this middle layer *Mantle* and their overall idea is sketched in Figure 3b.

4.1 FPGA Resource Virtualization

While the shell-role architecture can be seen as a bare metal deployment of design in FPGAs, the integration of resource management, including enforcement of elasticity in datacenter requires some form of resource virtualization, particularly FPGA resources. This section discusses the issue of virtualization in datacenters that use FPGA. We discuss virtualization from two different perspectives. The first one is the hardware resource virtualization on FPGA devices and the second one is the integration in hypervisors used to span new virtual instances in the cloud.

Virtualization essentially consists of creating abstraction layers over hardware components such as processors, memory, storage, etc, with the goal of enabling resource sharing at the software level. While the introduced abstraction layers may decrease native performance, they hide hardware complexity, improve the flexibility and programmability of the underlying hardware and make the integration in cloud manage tool.

Processor virtualization is well investigated and currently relies either on instruction set translation or hardware support with technologies such as Intel VT [170]. Memory and disk virtualization essentially map the virtual space to physical locations in the physical space [203]. FPGAs differ from the previous types of hardware in that they do not execute sequential programs, instructions after instruction, but rather implement parallel circuits. Traditional virtualization techniques used in software are therefore not applicable to FPGAs. Abstracting FPGA resources at higher levels requires a new approach. FPGAs consist of configurable computing resources, storage, and interconnect, all immersed in a programmable interconnect infrastructure. Programming is achieved with HDLs. The long synthesis time makes it impossible to achieve binary translation needed in most virtualization frameworks today.

Abstractions are provided in FPGAs using two different approaches: Slot-based allocation and FPGA overlays. Slot-based allocation divides the FPGA into blocks where pre-synthesized circuits can be mapped at run-time, while overlays use processor-like configuration, which makes them more convenient for traditional virtualization.

4.1.1 Slot-Based Allocation. Slot-based FPGA resource management [77] has been around since the early days of partial reconfiguration. The goal is to allocate FPGA area for the execution of pre-synthesized circuits at run-time through partial reconfiguration. The FPGA is divided into fixed areas, also called virtual regions, that will accommodate pre-synthesized circuits at run-time through partial reconfiguration. The slot-based allocation is extended for cloud FPGA resource management in [27, 93, 214, 242]. These works leverage partial reconfiguration at the FPGA level and provide packages for integration in cloud management infrastructure. Table 2 lists some of the recently published research work in this direction. Companies such as VMAccell [9] and Inaccel [5] provide commercial frameworks to achieve the same goal.

The table classifies the architectures based on the shell area overhead, the number of virtualized regions per FPGA, whether spatial sharing is enabled or not, the presence or absence of on-chip communication support, the data width, the maximum frequency, and the access method to the virtual resources on the FPGA. For example, Chen et al. [85] divide each FPGA into four locations or “virtual FPGAs”. The architecture multiplexes FPGAs in space, and provisions hardware resources over PCIe. However, this architecture is limited in that it only allows the use of pre-built hardware functions without support for direct on-chip communication between accelerators. This results in middleware copy overhead for data movement between the accelerators of a user.

Table 2. Study of recent research in FPGA virtualization architectures for cloud infrastructure

	Shell Area Overhead	# FPGA regions	Spatial Sharing	On-chip Comm. Support	Data Width (bits)	Fmax (MHz)	Access Method	PCIe Version	Network Specification
Fahmy et al. [103]	7%	4	Yes	No	256	250	PCIe	Gen 3x8	-
Vesper et al. [225]	18%	4	Yes	Yes	32-256	variable	PCIe	Gen 3x8	-
Weerasinghe et al [233]	21.7%	1	No	No	64	156.25	Network	-	10GbE
Tarafdar et al. [212]	27%	1	No	No	43	125	Network	-	10GbE
Mbongue et al. [159]	1%	4	Yes	Yes	-	227	PCIe	Gen3x16	-
Catapult [183]	23%	1	No	No	16-48	200	PCIe & Network	-	10Gb SAS
Byma et. al [82]	19%	-	Yes	No	256	160	Network	-	10GbE
Feniks [241]	13%	-	Yes	No	-	-	PCIe	Gen3x8	-
Mandebi et al. [164]	0.1%	6	Yes	Yes	32-256	1500	PCIe & Network	-	Fast Ethernet
Chen et al. [85]	6.46%	4	Yes	No	-	100	PCIe	Gen2x8	-
Asiatici et al. [34]	-	3	Yes	No	-	-	PCIe	-	-

To minimize the data movement overhead, an on-chip interconnect can be used between virtualized hardware regions as a support to hardware elasticity [159, 164]. It increases the on-chip throughput and enables pipelined processing within the hardware domain of cloud and datacenter applications. [162] evaluates the IO bottleneck in multi-tenant cloud FPGAs. Weerasinghe et al. [233] observe that network-attached FPGAs may offer lower latency and higher throughput compared to accessing accelerators over PCIe. Catapult [183], Tarafdar [212], and Weerasinghe [233] do not support spatial sharing of FPGA components between the workloads of multiple cloud users. They focus on FPGA time sharing. Table 2 also shows that the hardware footprint of the resource overhead that supports the virtualization is less than 30%. Since Shells occupy a small area on the FPGA, and considering the increasing size of FPGA devices, this overhead will become negligible.

4.1.2 FPGA Overlays. FPGA overlays have been developed as a promising way to increase FPGA programmability and productivity [204]. The approach commonly consists of designing a layer above the fine-grained FPGAs resources, thus hiding the complexity associated with programming low-level hardware components such as look-up tables (LUTs), flip flops (FFs), signal processing blocks (DSPs), etc. Moreover, it considerably increases productivity as it mitigates the long compilation time inherent in FPGA design flows. In general, FPGA overlays use coarse-grained reconfigurable arrays (CGRA) of processors that are programmed at run-time through software-level function calls. At the architecture level, FPGA overlays usually implement interconnect topologies allowing parallel processing and data exchange among processing cores [81, 124, 138, 146, 156, 158]. The software programmability of the coarse-grained processors makes it possible to develop efficient compilers for automatic mapping of sequential applications to overlays [23, 153, 160, 234]. Beyond standalone utilization of FPGAs and FPGA overlays, cloud and datacenter architectures offer the opportunity to exploit FPGAs at a higher level of abstraction. FPGA overlays are most often immersed in the Shell to provision reconfigurable resources to cloud users.

4.1.3 Operating System Extensions for FPGA Virtualization. The integration of hardware resource, including those on FPGAs into virtual instances assumes some sort of operating system support resources and protocols. There is a sizeable body on FPGA and integration in conventional operating system. While these work don't directly address cloud and datacenter, the low-level protocols developed would facilitate the integration of FPGA resources in hypervisor and virtual machines. Ma et. al [151] propose OPTIMUS, a hypervisor for a scalable shared system between cloud FPGAs and host CPUs. It implements efficient virtual DMA isolation via page table slicing, resulting in up to 7× improved throughput. FPGAVirt uses VirtIO for efficient communication between virtual

machine and FPGA through in-kernel network stack [159]. Korolija et al. explore the application of traditional operating system resource abstraction to FPGAs [143]. The authors implement Coyote, an open source, portable, and configurable shell for FPGAs. It supports secure spatial and temporal FPGA multiplexing, virtual memory, communication, and memory management inside an uniform execution environment. Hategekimana et al. extends Security-Enhanced Linux (SELinux) security context to hardware IPs on multi-tenant cloud FPGAs [126]. This capability allows system security policies to propagate access control privileges expressed at the hypervisor level down to individual FPGA accelerators at run-time. Though operating system-level concerns such as abstraction and scheduling are being investigated in the literature, there is still need for operating system support for FPGA multi-tenancy in the cloud. Therefore, Khawaja et al. detail AmorphOS as a solution [139]. It enable cross-domain protection, and replaces fixed slot-based allocation with elastic resource management to increase utilization and throughput. Moreover, recent research also addresses the compilation and execution of FPGA accelerators in cloud-based systems. As example, Landgraf et al. present SYNERGY, an FPGA compiler tool capable of generating controls to software execution, necessary to support core virtualization primitives such as suspend, resume, and program migration on FPGA [145]. Fumero et al. introduce TornadoVM, a virtual machine for applications acceleration on heterogeneous hardware at run-time [110]. It relies on JIT compilation to map kernels to adequate hardware accelerators. Using computer vision workloads, the authors show that TornadoVM allows achieving $7.7\times$ speedup in an heterogeneous platform provisioning CPUs, FPGAs, and GPUs.

4.2 Using and Scaling of FPGA Clusters

Besides providing easy-to-use abstractions or virtualization of a single FPGA in a datacenter environment, the problem of programming and deploying multiple FPGAs is important as well for making FPGAs a scaling-up solution. The first problem of programming a cluster of FPGAs includes the question of how to abstract and use the communication links between the FPGAs and also how to link the kernels on multiple FPGAs together. The second problem of deploying an FPGA cluster contains the question of how to distribute, orchestrate and manage an multi-FPGA application once it is build. In the following, we summarize major academic and commercial approaches tackling this problems.

4.2.1 The Galapagos Research Project: Providing a Middleware for communication, build, and deployment. The Shell – Role interface simplifies using a single FPGA. At a cloud-level, IP core placement within and across devices is needed. The Galapagos infrastructure [209, 212, 214] provides multiple layers of abstraction including a Middleware. The responsibility of the Middleware is to provide an easy abstraction for users to deploy an application across many FPGAs. This entails providing an abstraction for users to describe clusters, scale and replicate IP cores, as well as to make the routing simple between IP cores, agnostic to their placement. The agnostic routing removes the headache the user would have when designing IP cores, and also gives the Middleware layer the flexibility to place IP cores anywhere. This flexibility is leveraged to provide reliability as IP cores can be migrated between different physical locations, scalability as these IP cores can be replicated, and even ease of development as users can implement their IP core (that is part of a large multi-FPGA deployment) in software and incrementally migrate parts of their multi-FPGA application from software to hardware.

The Galapagos Middleware standardizes the interface between IP cores through AXI-Stream [32], a standard streaming protocol used by ARM and Xilinx IP cores. This streaming protocol is typically used within single FPGA systems in the Xilinx environment, with an optional bit-width configurable destination side-channel. This side-channel is used to route packets between on-chip AXI-stream IP cores. Galapagos Middleware takes this established streaming protocol for on-chip IP cores

and expands this to multiple FPGAs. By leveraging standard AXI-stream, IP-cores do not need to be aware of the location of their destination IP cores. Figure 4 represents the user's view of the cluster versus what is actually implemented. The user-provided IP cores are unaware of the physical placement of these kernels.




Fig. 4. The logical view of the cluster from the user's perspective is transformed with the Galapagos Middleware into a multi-FPGA implementation [212].

The logical to physical translation is implemented by first performing a partitioning of IP cores across different FPGAs. The first iteration of this partitioning is simple as a greedy algorithm is used. It attempts to fit as many IP cores on one FPGA before allocating another FPGA. This partitioner is a first implementation and can be improved independent of the rest of the Galapagos stack. Once a partitioning of all IP cores in the cluster has been established, the middleware generates the routers and bridges to connect all the IP cores within the cluster. Figure 5 shows the bridges used to establish on-chip and off-chip connections between IP cores. Since the formation of the bridges is automated through the middleware, user-provided IP cores can remain agnostic to the placement of all the other IP cores within the cluster. The autogenerated router on each FPGA is aware of the placement of all IP cores in the cluster and can route packets to an IP core on the same FPGA or to a different FPGA through a network bridge. A network bridge translates AXI-stream packets to off-chip network packets. The network protocol is configurable to be any off-chip protocol the user wants as long as there is a bridge available to translate AXI-stream packets into network packets. Currently bridges are available for Layer 2 Ethernet, UDP and TCP/IP. However, as long as the user supplies the bridge this can support any off-chip protocol.

To support heterogeneous interaction between CPUs and FPGAs, a software library, called *libGalapagos* [210] has been created. The library provides software interfaces that exactly match the streaming interfaces used in hardware Galapagos kernels. Any software kernel that uses these interfaces can interact directly with any hardware Galapagos kernels making it easy for a software kernel on a CPU to communicate with a hardware kernel on an FPGA. One benefit of this capability is that if a Galapagos kernel is described using HLS-synthesizable C++, that HLS code can be wrapped using *libGalapagos* so that it can be run purely in software. This enables a co-simulation environment where all the hardware kernels can be first tested in software to debug functionality. Once functionality is achieved, the HLS kernels can then be synthesized and run as hardware without any code changes. Whether such a kernel is run in software or in hardware is determined by specifying the version in a Galapagos configuration file.




Fig. 5. Bridges and routers generated by the Middleware placed on each FPGA.

4.2.2 The ZRLMPI prototype: Providing a programming model for FPGA clusters. Besides abstracting and managing the communication between different FPGA kernels, the behaviour of this type of concurrency needs to be programmed as well, so a programming model for FPGA clusters is needed. In the past decade, there have been numerous approaches for this question (among others [94, 102, 112]), but there is no convergence to a common standard yet.

The trend towards network-attached FPGAs and thus the development of the FPGA from a co-processor to an "equal" stand-alone node, caused the community to revisit this search for a programming model for FPGA clusters. The ZRLMPI prototype [191, 192] tries to end this search by proposing to port MPI, the de-facto standard for HPC, to FPGA clusters.

Bringing MPI to a heterogeneous cluster of FPGAs and CPUs involves two steps: First, a compiler is needed to compile a given program to different hardware. Second, a run-time environment is required, which implements the MPI APIs, start, stop, and synchronizes the execution between the nodes. ZRLMPI provides both. At the one hand, it provides a transpiler – or cross-compiler – that splits and optimizes an input MPI program to multiple programs, one for each physical node. This split breaks with the "Single Program, Multiple Data" (SPMD) notion of MPI, but this is necessary due to two reasons: First, in opposition to software, unused program parts for individual nodes results in wasted logic for FPGAs. Second, CPU and FPGA parts of the program need different compilers to build the binaries. At the other hand, ZRLMPI implements a run-time environment for FPGA and CPU nodes that ensures the synchronous message processing independent of the executing hardware. Hence, the CPU-FPGA cluster can be programmed and managed like a pure CPU cluster and existing programs could be ported easily to such heterogeneous clusters.

4.2.3 Commercial programming model: Maxeler's Dataflow Engines. Maxeler, a pioneer in FPGA-based high-performance computing, provides an entire infrastructure consisting of hardware and software for dataflow applications [6]. Its MaxCloud provides cloud implementation of a high-performance dataflow computing system. MaxCloud runs on Maxeler MPT compute nodes that consist of multi-core x86 CPUs with multiple Maxeler Dataflow Engines (DFEs), large memory systems, and fast disks. It runs an industry standard Linux distribution. The DFEs are FPGA-boards attached to the CPU through PCIe.

4.2.4 Commercial deployment platforms: InAccel and VMAccel. InAccel [5] offers a general-purpose accelerator orchestrator capable of integrating FPGAs from various vendors to simplify the deployment, scaling and resource management of FPGA clusters. Their Coral software contains high-level APIs in C/C++, Java and Python, and a unified engine that supports every heterogeneous multi-accelerator platform. InAccel also provides a runtime specification that vendors can use to advertise system hardware resources to Coral. It aims to specify the configuration, and execution interface for the efficient management of any FPGA hardware resource, without customizing the code for Coral itself. In Coral, client applications may call accelerators on a local FPGA or remotely with their arguments and configuration parameters. Coral orchestrates the execution on a variety of execution environments, from bare-metal Linux servers to containers inside a Kubernetes cluster.

VMAccel [9], a company in stealth before January 2021, operates a large, multi-tenant FPGA IAAS cloud platform based on Openstack, Docker, and Kubernetes. This platform supports virtual machines, bare metal, and containerized FPGA acceleration options. VMAccel's cluster consists of over 1000 FPGAs and is designed from the ground up with FPGAs in mind. VMAccel removes restrictions on hardware and offers direct access to JTAG and QSFP interfaces for clients; their offering is designed to be customizable and flexible to accommodate varying user needs and workloads. The company plans to implement disaggregation technology by Q4 2021 to offer users even more flexibility in hardware composition.

4.3 Cloud FPGA usability – Current Needs and Outlook

Considering the usability, scalability and portability, we notice the progress that has been made in the past years to decrease the barriers of leveraging FPGAs in datacenter and cloud environments, but we also see the need to further improvements. In particular, we identify two areas: First, to enable "seamless" portability, the application interfaces within the FPGAs need to be standardized. Second, FPGAs should be integrated further into classical software stacks and debugging of (distributed) FPGA applications has to become simpler. Both areas are addressed below.

4.3.1 Enabling Portability: Standardizing the Application Interface. From a general point of view, FPGA applications do not exist in empty space: they are part of a more or less complex communication schema with other FPGAs or CPUs, as discussed in Sections 3 and 4.0.2. If we analyze all the earlier mentioned SRAs from this perspective, we find some general patterns, besides the pure split of logic into a Shell and a Role. This high-level generic design of SRA platforms is shown in Figure 6. Every Role is controlled by some configuration and control registers or signals such as start and stop, node_id or virtual

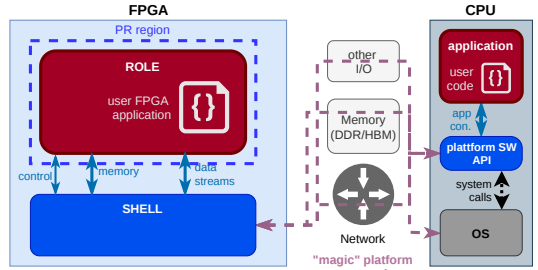


Fig. 6. Shell Role Architecture share many common concepts.

I/O pins. Besides this, two types of communication channels can be found:

- (1) Address based communication: Every data belongs to an address in an address space. This address space is at least partially shared with the CPU. So, the content of the data can be interpreted based on the address the data belongs to. This communication type is used by PCIe-attached or memory-synchronized FPGAs such as openCAPI systems [174] or OpenCL platforms.
- (2) Stream based communication: Here, every data belongs to a stream of data without further meta data such as addresses. In some cases, the data stream is accompanied by other streams that indicate the origin or destination of this stream. Consequently, the data can be interpreted only with the knowledge of the specific physical connection the data came from and to some extent, the content of parallel meta data streams. This type of communication is used by network-attached FPGAs – like the IBM cloudFPGA project [25] or the Galapagos cluster [102] – or for some fabric abstractions.

The synchronization between these FPGAs and the other devices, e.g. the CPU is up to the “magic” protocol of this platform, as sketched in the middle of Figure 6. “Magic” in the sense that the details of this synchronization are not important for FPGA application developers. FPGA application developers are more concerned with the high-level capabilities (what can be done) as opposed to the low-level implementation details (how the synchronization between FPGAs and CPUs is done).

Therefore, since most FPGA cloud platforms share the same architectural pattern with the SRA and due to the observation that interfaces of the data path are more important to users of cloud FPGAs than the details of how the data path ends up in the FPGA, we think it is possible to describe a standard interface for cloud FPGA applications at this level and leave the details of the “magic” to each individual platform. The SRA architecture has been around since the early days of FPGA integration in desktop systems. Examples include the ESM [77], Walder and Platzner OS [227] and the Celoxica series 1000PP and 2000 [78], and many more. This standard interface can also apply to

extensions like the Mantle architecture [190], because the adapter could be "below" or "above" a common interface. There is no reason to doubt that progress will continue in this direction and that future systems will rely on a more portable, standardized SRA as interface for seamless deployment of design in the cloud.

Every system designer has implemented a variation with a library on the hardware and software side to facilitate the development of hardware and software systems. A semblance of standardization has been developed over the years by FPGA vendors in the form of board support packages and libraries. While these packages simplify the design and integration of a single FPGA in a desktop or standalone mode, extension to cloud computing platforms has yet to be achieved. Increased adoption of FPGA infrastructure requirements in the cloud will help achieve this goal. In this direction, we are starting to see new CSP players such as VMAccell [9] and Inaccel [5] that provide tools and architectures to seamlessly expose FPGAs to users in numerous possible cloud computing paradigms. This progress will continue with companies offering more flexible, efficient solutions that increase resource usage, with single or multiple tenants. The Cyborg project discussed in section 4.3.2 is an effort by the opensource community to facilitate the integration of non-conventional processors such as FPGA, GPUs in datacenters. We expect these developments to gain traction and lead to better design tools and efficient management frameworks that will accelerate adoption of FPGA in datacenters.

4.3.2 Cyborg Project: Managing FPGA Accelerators in OpenStack Clouds. OpenStack is a free and open standard for cloud computing platforms. It allows easy deployment and management of infrastructure-as-a-service for public and private clouds. The development of OpenStack is organized into projects. The Cyborg project specifically addresses the integration and general management of hardware accelerators such as FPGAs, GPUs, SoCs, etc. in OpenStack, enabling the provision of hardware acceleration as a service [14]. Cyborg provides REST APIs to list, create, update, and delete devices hardware devices in a cloud infrastructure [19]. Enabling FPGA acceleration in OpenStack opens opportunities for researchers to explore use cases of FPGA technology in the cloud. In addition, it enables cloud providers with the ability to provision heterogeneous architectures with minimal effort and cost. However, services still need to be developed to program devices at run-time, manage bitstreams, update and migrate shells, enable debugging designs with the insertion of probes, monitor on-chip sensors, etc [12].

4.3.3 Simplifying Usability: Increase FPGA Integration and DevOps support. Besides the limited portability of FPGA applications, which leads to strong dependencies to specific hardware and cloud vendors [190], FPGA applications need still be developed from scratch, most of the time. Usually, one need to write a new FPGA application in a tool that depends on the used FPGA platform. In addition, debugging is often still done with debug-probes at "wire-level". This is in strong contrast to the usability of other accelerator platforms, where in the best case, the application developer doesn't notice the usage of heterogeneous hardware by the toolchain. For example, if using today's deep-learning frameworks, GPUs are seamlessly integrated and can be used with one or two lines of code [121, 184, 216]. The usage of one or many FPGAs must become as simple as the usage of GPUs in this examples.

As discussed in above sections, there are many efforts that explore this research direction. For example, it is possible to deploy a given FPGA binary on a multitude of cloud FPGA platforms with a single line of code [137] or to compile and deploy an application of a given framework on a specific platform [191]. However, to reach the desired level of simplicity, both categories need to be combined. It is not sufficient to deploy a fixed functionality on a flexible platform or a flexible functionality on a fixed platform. For FPGAs, to be of use for complex scientific and commercial applications, it must be possible to develop, debug and deploy complex distributed

applications easily and independently of specific platforms. This includes the development of vertically integrated design environments [122, 167, 172, 178] and supporting more debugging possibilities [185]. We notice that this problem is subject of many active research projects and we appreciate that this focus is starting to be reflected in academic conferences and workshops.

5 SECURITY OF FPGA ACCELERATION IN CLOUD AND DATACENTER ENVIRONMENTS

The availability of FPGAs in cloud datacenters has opened up unprecedented levels of application flexibility and performance, although security is an important consideration. The development of Spectre [142] and Meltdown [148] demonstrated how hardware vulnerabilities can be leveraged at the software level to successfully launch attacks on infrastructure. FPGAs add another infrastructure resource to the cloud that must be free of vulnerabilities to prevent malicious activities. Unfortunately, a user's ability to implement any logic function provides unique avenues for malicious attacks on other cloud users' applications and data, and the cloud infrastructure itself [134, 222, 239]. In this section, current and future threats to cloud FPGA security are discussed with an eye towards prevention of existing and future potential attacks.

5.1 Cloud FPGA Security Overview

Many of the threats facing cloud FPGAs are similar to threats that have long been prevalent in microprocessor-based cloud computing [90, 175, 208]. Multi-tenancy can lead to side channels in which confidential user data is stolen [186, 243], or covert channels in which otherwise logically isolated users are able to transmit sensitive information between each other [116]. Software and interface manipulation can also result in a denial of service for cloud users or even the infrastructure. Malicious bitstreams uploaded to the FPGAs in the data center can lead to faults or incorrect results being generated [120, 181]. Unlike previous software-based attacks, cloud FPGA attacks often include digital hardware manipulation whereby the FPGA logic is configured in a malicious way that can lead to accelerated device wear-out, short circuits [31, 71], or performance degradation. In addition to the existing single-tenant deployments already available from vendors such as AWS, multi-tenant cloud FPGAs are being actively researched. In multi-tenant FPGAs, different users share the same FPGA, making attacks between users more dangerous. The use of multi-tenant cloud FPGAs in the future will expand attack surfaces, necessitating more comprehensive security plans.

The threats associated with cloud FPGAs include malicious FPGA tenants in a single-tenant setting, co-tenants in a multi-tenant setting, nefarious cloud service providers, malicious and compromised tools, or intellectual property cores that may try to attack or steal information [134]. The most striking of these threats is related to malicious co-tenants in the multi-tenant setting.

In this section, we review cloud and datacenter FPGA security threats. Vulnerabilities broadly target the FPGA chip, the interfaces between the FPGA chip and other components on the FPGA board or the server where the FPGA boards are located, and the system software used to manage the FPGA bitstreams (e.g. loading designs onto the FPGA) and access data generated by the FPGA. Other infrastructure components related to FPGAs, such as the shared power distribution within the server, can also be a source of security vulnerabilities. Current and potential future attacks are reviewed in this section in addition to existing and potential countermeasures.

5.2 Cloud FPGA Usage Models

In this subsection, we categorize FPGA usage models in the context of the architectures and interfaces introduced in Sections 3, 4.0.2, and 4.3.1. In many respects, the level of security available to cloud FPGAs is dependent on the architectural model used and the associated interfaces.

Unlike in software-based virtual machine (VM) instances, with FPGA-accelerated instances, users naturally engage with the hardware at a lower level. In software-based VMs, a hypervisor provides an encapsulated environment for user execution that largely consists of virtualized CPUs, memory, storage, and network interfaces. This encapsulation facilitates memory protection and network isolation. Today's commercial single-tenant cloud-based FPGAs do not abstract or virtualize fundamental logical resources such as LUTs, block memory, or multiply-accumulate units. This limitation is currently present for two reasons. First, abstracting the underlying FPGA fabric with a template incurs non-trivial overhead in terms of usable FPGA logic density. Since many of the performance benefits of using an FPGA hinge upon exploiting maximal parallel processing, the efficient use of the available logic is very important. Second, FPGA virtualization is difficult. For example, adding logic to stop an arbitrary circuit mid-computation, save the intermediate state, and restore it at a later point presents design challenges and is expensive to implement in terms of FPGA resources and performance. For example, although block memory and flip flop contents in Xilinx UltraScale+ FPGAs can be streamed out of the device at 6.4 Gbps using a configuration interface, the largest devices contain nearly 500 Mbit of RAM and 3.5 million flip flops [10, 236]. Also, pipeline registers in multiply-accumulate blocks cannot be easily retrieved via the interface.

Difficulties in implementing FPGA virtualization affect how they can be offered in the cloud. For example, unlike traditional software-based servers, the fine-grain context switching of user Roles is not practical. Rather, today's commercial cloud providers assign an entire FPGA or more than one entire FPGA to a user for the duration of their computation, forming a single-tenant setting. For example, in the case of an AWS F1 instance, a user is assigned a VM on a given server that has exclusive access to N FPGAs, with the VM resources and number of FPGAs scaled to the size of the purchased instance [22]. The user first programs their circuit(s) onto the FPGA(s) and then can utilize them at will, within the constraints of what the Shell will allow. When the user's instance terminates, or the cloud provider elects to reclaim the instance, the software-side and FPGA resources will be returned to the allocation pool to be re-assigned to a different user, who will then gain exclusive access to the FPGA(s) for themselves. The provider is expected to properly erase all the resources before allocating them to the next user. For instance, the DDR memories provided on the FPGA boards in AWS F1 instances get overwritten before an instance is allocated to a new user.

In the near future, it is likely that large FPGAs that could be shared by multiple users will be installed in datacenters. Sharing of the same FPGA by multiple users in the multi-tenant setting is an example of the multi-tenant spatial model. In this case, FPGAs will have enough resources to subdivide each device into multiple regions that can be assigned to independent users and a single Shell would manage multiple Role regions, using dynamic partial reconfiguration to install a given user's Role at runtime. As a result, a single device would need to connect to multiple different VMs and communication to the device would need to be multiplexed from a specific VM to a particular Role. Since the circuits for different users reside on the same FPGA simultaneously in different Role regions, significant new security issues may arise (e.g., side-channel attacks, fault injection attacks, attacks on shared memory and I/O infrastructure [134]).

In addition to accelerator-focused deployments, such as AWS F1 instances, in which users access the FPGAs via PCIe as a co-processor, there are other models of intra-node FPGA communication. As discussed in Section 3, bump-in-the-wire architectures directly connect FPGAs to a network so that all network communication flows through FPGAs. This gives faster access to network data, but, effectively, data from many users may pass through an FPGA instance controlled by a malicious user.

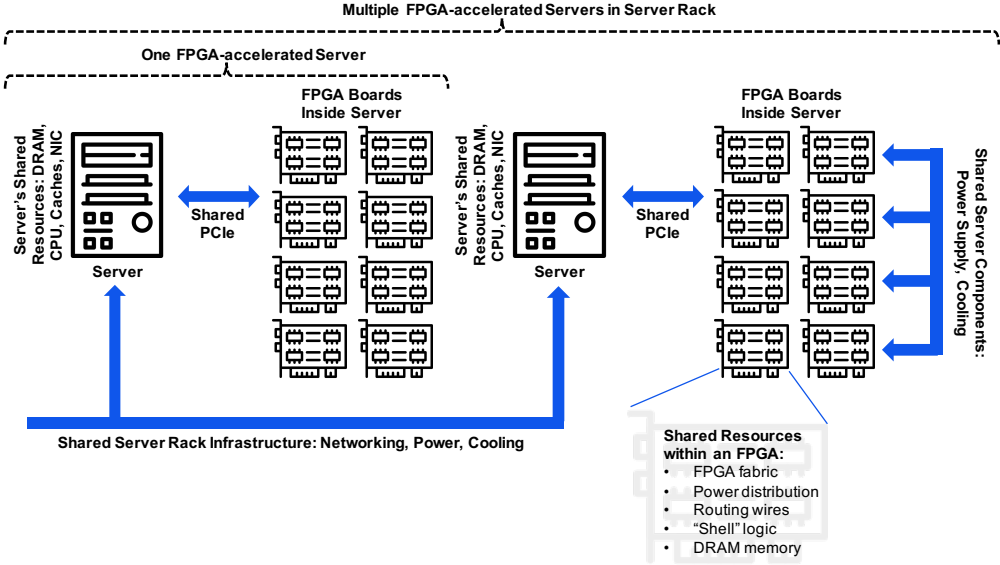


Fig. 7. Diagram of hardware resources that are shared in FPGA-accelerated servers and among servers in a server rack; each shared resource could be a source of a possible attack or information leak.

A further possible deployment includes the CPU and FPGA together on the server, possibly as different chips on the same motherboard or with the CPU and FPGA sharing the same system-on-chip (SoC). With the acquisitions of Altera by Intel and Xilinx by AMD, FPGAs may become directly integrated into large servers as fundamental computers, not just as accelerator cards. It is also expected that FPGAs will become parts of the processor dies themselves. This opens up even further security threats, as, for example, the FPGA may have access to the memory bus and the coherence messages being sent between CPUs.

5.3 Cloud and Datacenter FPGA Attack Categories

The co-location of multiple tenants inside cloud-based FPGAs, or having multiple dedicated FPGAs per user but sharing a server between different users, leads to numerous resource sharing threats. Resource sharing can occur within an FPGA in a multi-tenant setting, between an FPGA and associated host CPU and DRAM, or between different FPGA boards within the server in a back-end configuration. Figure 7 provides a summary of resources that can be shared. The vulnerable resources include the system and individual chip power distribution network (PDN), FPGA logic and wiring resources, Shell logic provided by the cloud provider, DRAM memory on the FPGA board, and the host CPU. Further, a system bus in a storage-attached system may be shared with the host server, which can include vulnerable DRAM memory that can be accessed by multiple users via direct memory access (DMA).

The following subsections detail current and future threats that are rooted in the shared resources. These threats have been broadly grouped into the following categories:

- (1) Single-tenant attacks
- (2) Multi-tenant attacks
- (3) Node-level attacks
- (4) Datacenter system-level attacks
- (5) Bump-in-the-wire attacks

(6) FPGA and CPU on same die, motherboard, or SoC attacks

5.4 Single-Tenant Attacks

These attacks affect FPGA computation in which one user at a time is using the FPGA. Today's cloud-based FPGA deployments (e.g., [22]) exclusively follow such a single-tenant model.

5.4.1 Single-Tenant Temporal Channels. Allowing users to load custom FPGA designs opens up the opportunity for malicious users to sense the environment on or around an FPGA. In particular, temperature affects operation of many electronic circuits: it affects signal delay in wires, transistor switching, and capacitor decay (used in DRAM memories, for example).

Previous work [218] has shown that it is possible to create thermal covert channels for communication between users. In particular, even in a single-tenant FPGA setting, users share the FPGAs in time: once a user is done using an FPGA, it is assigned to a different user. Other temporal attacks may be possible, for example one FPGA user may generate substantial PCIe traffic, which reduces the bandwidth quota assigned to the FPGA. When the next user design is loaded onto the FPGA, he or she may still experience a reduced bandwidth quota from the prior user. Any system component that maintains state could be abused for temporal attacks if the state is not properly erased. The state could be physical (e.g., thermal state) or logical (e.g., PCIe bandwidth allocation). For example, old DRAM values may be read by subsequent users. In practice, this attack is not currently feasible since vendors (e.g., AWS F1) actively clear DRAM between users.

5.5 Multi-Tenant Attacks

Multi-tenant attacks impact FPGA computation in which multiple users access the FPGA or shared resources (e.g., memory) at the same time. Most current research focuses on this setting as resource sharing can lead to a diverse set of attacks.

5.5.1 Multi-Tenant Voltage Attacks. FPGAs in the multi-tenant spatial model are particularly vulnerable to supply voltage attacks. All logic implemented in a commercial FPGA shares a power distribution network (PDN). As a result, the PDN can be manipulated to perform a broad spectrum of attacks including the implementation of covert channels [116], the side channel theft of encryption keys [197, 243] or other information, and fault injection [120, 181]. In extreme cases, large power waster circuits can consume sufficient power to drive the FPGA board into reset [182]. A full summary of multi-tenant voltage attacks is available [134].

Although multi-tenant voltage attacks have been widely explored, there are several additional vectors of interest. It may be possible to localize failures to other specific devices on a board through clever on-FPGA voltage manipulation. Remediation to prevent such attacks is also needed. AWS F1 includes a power monitor that will shut down an entire FPGA that consumes excessive power [119]. It has been shown that voltage sensor networks instantiated in an FPGA can quickly detect attacks [181]. This information could be used to quickly suppress offending attack circuitry. Finally, it may be possible to architect FPGAs so that PDN regions are isolated or a global "kill" signal can quickly suppress an FPGA region suspected of malicious activity.

5.5.2 Multi-Tenant Attacks on Shared Block RAMs. Allowing a user to intentionally cause write collisions in FPGA dual-port block RAMs can induce voltage and temperature fluctuations and result in circuit faults [28]. The faults may affect other users. Although unexplored, it may be possible to abuse other modules (e.g., communication blocks, DSP blocks) within an FPGA to cause faults or unintended behavior for other users.

5.5.3 Multi-Tenant Crosstalk Attacks. The crosstalk of routing channel wires in multi-tenant FPGAs has been extensively explored [118, 186]. In particular, static values on routing wires (logical 1 vs.

0, which corresponds to a charged or discharged routing wire) affect the propagation delay on immediately adjacent channel wires. It is possible to sense the state of the adjacent wire and use the information as a covert channel or side channel. Crosstalk has been used to steal AES [186] and RSA [187] cryptographic keys.

Crosstalk could also be used to target Shell logic, or wires associated with communication buses, such as the AXI bus within the FPGA used to connect different modules. The shared internal FPGA bus could potentially be snooped to extract data from different users at different times. The existence of crosstalk necessitates physical separation so that different users' modules or the Shell are spatially far apart to prevent crosstalk. User-level control of logic placement may need to be prevented in multi-tenant scenarios so that attackers cannot choose custom locations for their logic, e.g., next to the Shell's communication buses.

5.5.4 Multi-Tenant Attacks on DRAM. Each cloud FPGA board typically has its own attached DRAM that augments the shared DRAM on the server. Both sets of DRAM storage could be susceptible to Rowhammer-style attacks [140] in which one user attempts to modify the DRAM of another user. This issue is particularly worrisome for multi-tenant FPGAs in which multiple spatial co-tenants have immediate hardware access to the same DRAM modules, although this has not yet been demonstrated in practice.

Since cloud FPGAs contain multiple DRAM chips (e.g., AWS F1 includes four DRAM modules, or chips, per FPGA), it may be possible to assign data from different tenants to different DRAM chips, mitigating Rowhammer attacks by tenants. If more tenants than DRAM chips exist in a bank, tenants could be assigned to non-adjacent DRAM locations to prevent attacks. The development of a DRAM controller and DRAM allocation and access policies for multi-tenant FPGA settings could also be considered.

5.6 Node-level Attacks

These attacks influence computing components located outside of the FPGA. Many of these attacks are possible already in single-tenant cloud settings.

5.6.1 Node-level Attacks on Server's DRAM. An FPGA Role could influence memory access for software components that are not running on the FPGA. An FPGA Role or compromised Shell could induce Rowhammer-style attacks on the server's DRAM since the FPGA has DMA access to the memory on the server. FPGA accelerators typically use memory mapped buses to access DRAM. Cloud providers must implement ad-hoc memory protection units to prevent malicious accesses, especially when resources are shared by another user.

While current FPGA node implementations typically include accelerator cards, FPGAs will eventually be more tightly integrated into the server. Once integrated, FPGAs may have direct access to coherence messages and the ability to read and write DRAM, bypassing the PCIe and DMA used today. This opens up the possibility of attacks on the coherence messages or the modification of the caches' contents since FPGA operations will affect the CPU's memory.

5.6.2 Node-level Attacks on a Server's Power Distribution Network. Shared power distribution networks are not only a concern for FPGA co-tenants, but also for components that share a system-level power supply. Significant changes in FPGA power usage on a data center board can be sensed by an FPGA on a different board if both boards use the same power supply [117]. While covert communication using a shared power supply has been demonstrated in a lab setting, it has not been shown using commercial public cloud hardware. A natural extension may be the creation of a covert channel between servers or between server racks, although these channels may be blocked by power supply filtering circuits.

It has also been shown that the encryption key of a cryptographic algorithm running on a microprocessor can be extracted using an FPGA if both devices share a supply voltage [243]. This approach could also be used to create a covert communication channel. Similar channels could be developed from CPU to FPGA or from GPU to FPGA, as previously demonstrated [117]. In both cases the FPGA could be the receiver as it is able to measure voltage changes.

5.6.3 Node-level Bus Attacks. The Shell could be attacked by an FPGA Role to stall the bus connecting the FPGA to the host server, creating a denial-of-service attack. This attack could potentially be achieved by manipulating messages sent to the bus controller on the FPGA side, or by causing voltage attacks which in turn create faults in the bus controller that stall the bus.

5.7 Datacenter System-Level Attacks

These attacks affect more than one node in the datacenter.

5.7.1 FPGA Instance Fingerprinting. Users can create malicious FPGA designs that steal cloud infrastructure information. For example, it is possible to fingerprint FPGA instances using DRAM physical uncloneable functions (PUFs) [219]. Given the assumption that a given physical DRAM module will not be moved from one FPGA board to another, fingerprinting DRAM modules is equivalent to fingerprinting the attached FPGAs. An FPGA may also be fingerprinted using dedicated PUF circuits inside the device [223]. Any circuit with a delay path, such as time-to-digital converter (TDC) circuit or existing hard intellectual property (IP) modules for digital signal processing (DSP), could potentially be used as a PUF. Recent work [217] has shown that PCIe bus contention can also be used to fingerprint cloud infrastructures. Such contention-based attacks could be extended to also cover network cards, which are attached to the bus.

Fingerprinting attacks allow users to learn the number of FPGAs that are available and their physical co-location within servers. Although fingerprinting attacks have focused on individual FPGAs and servers, they could be extended to fingerprint server racks or whole data centers, possibly using server thermal or voltage fluctuation patterns.

5.7.2 Potentially Destructive Attacks. Any attack that can physically damage the data center hardware is very serious. Most destructive attacks would likely use excessive power, although, as mentioned earlier in this section, power usage is already checked by cloud providers such as AWS, and a user's design is stopped if an FPGA crosses a power usage threshold. In general, destructive power attacks may not have an instantaneous effect. Accelerated hardware degradation via less aggressive power attacks may take longer, but can be equally destructive and harder to detect.

5.8 Bump-in-the-wire Attacks

Risks associated with cloud FPGAs impact system architecture and user access models. Bump-in-the-wire architectures have unique security challenges since FPGAs are directly connected to the network and process network packets from many users. Exposing FPGA resources in this manner allows reconfigurable hardware application developers to exploit the natural ability of the hardware to process at network line rates. Furthermore, bump-in-the-wire (or Smart-NIC) architectures can support new capabilities while minimizing the impact of host software.

Malicious hardware in a bump-in-the-wire FPGA can easily and quickly flood the network with nuisance packets. Although modern datacenters have systems to both monitor for problems and take corrective actions, which limit the potential scope of issues that do arise, there is an inherent delay between the start of a potential problem and the implementation of a solution. Thus, network-connected FPGAs can, at the very least, create transitory local service issues. In a datacenter where network-connected FPGAs are ubiquitous, a well-coordinated attack could have a systemic impact.

This issue is further complicated in bump-in-the-wire architectures since blocking traffic from an ill-behaved FPGA also blocks access to the host server.

Future research in bump-in-the-wire architectures could consider the moderation of network packet injection and the careful scanning of network requests. Effectively, a broad class of network security approaches could be applied to this class of cloud FPGA architecture.

5.9 Comparison with Existing CPU- and GPU-based Threats in Cloud Computing

Most hardware-related CPU threats faced by cloud providers involve information leakage and side-channel or covert-channel attacks. The information leakage is usually due to resource sharing (e.g., via shared processor caches) [207]. These sharing-based attacks have some similarity to possible shared bus [217] and DRAM attacks involving multiple FPGA tenants. For users that demand security, cloud providers can allocate full CPU cores (to avoid level 1 cache sharing and related attacks, for example) or even allocate full servers to a single user. This limits any sharing of CPU resources, but also clearly prevents resource sharing that could lower costs or improve utilization. Researchers have also demonstrated GPU-based cache attacks that can leak information between users due to shared resources such as caches [96]. FPGAs also suffer from resource sharing and contention, but also have new types of FPGA threats due to ability to create malicious FPGA circuits that can measure temperature, measure voltage changes, observe decay of the DRAM modules, etc.

5.10 Cloud FPGA Security – Current Countermeasures and Needs

Looking forward, security for cloud FPGAs can be categorized into four axes of protection: FPGA intellectual property protection, system-level security, run-time user support, and FPGA device architecture improvements. Each of these axes is addressed below.

5.10.1 FPGA Intellectual Property Protection. Although FPGA design disclosure to the cloud vendor helps limit the loading of malicious bitstreams to the cloud, it exposes a user's intellectual property to the cloud vendor. AWS and Microsoft, for example, do not allow users to directly create FPGA bitstreams for use in the cloud. Zeitouni et al. [240] have developed a multi-step procedure to protect cloud infrastructure from damage by a malicious FPGA without disclosing the user's netlist to the cloud provider. This goal is accomplished using a trusted execution environment, trusted Shell, and mutual authentication using a PUF for key generation. This approach could be extended by considering FPGA architecture enhancements for trusted execution and key storage.

5.10.2 Preventing Attacks Using System-Level Security. Protection for FPGAs, memory, transport buses, and OS support are needed to create a secure system. These aspects of cloud FPGAs use could be potentially be vulnerable and a security approach must consider how to protect the whole system. System level integration increases the difficulty of securing an entire system in the face of remote hardware attacks.

The integration of FPGAs in cloud systems induces temporal and/or spatial interaction among cloud tenants. Without guarantees of isolated execution, this sharing can lead to scenarios where shared accelerators act as potential covert channels among software guests which reside in different security contexts. Remote attacks have been demonstrated by manipulating FPGA hardware (e.g., the PDN), as explained earlier in this section. Domain isolation in hardware can be addressed from a software perspective using an operating system. Emphasis can be placed on a system's ability to limit privileges associated with executing processes in order to contain the scope of damage that can result from the exploitation of application vulnerabilities [149, 194]. Access control used in SoC and FPGA-based embedded systems have assumed the availability of *reference monitors* to manage hardware access [79, 130, 188]. Similar approaches could be provided in a multi-tenant cloud FPGA scenario.

A comprehensive solution to domain isolation that cuts across hardware and software has been proposed [125, 144, 157, 163] in which an FPGA is shared among virtual machines. The isolation framework guarantees that hardware modules execute and reside in the same security context as the “caller” VM. Similar operating systems and hypervisor approaches could form the basis of multi-tenant isolation.

5.10.3 Run-time User Support. The Shell, cloud software, and cloud management systems used for cloud FPGAs will need to significantly change in the coming years to support the multi-tenant spatial model. This prospect raises multiple new FPGA-specific questions for cloud providers. Starting at the beginning of the work flow, a user’s Role must be dynamically programmed to a specific Role region. However, the cloud provider may not elect to expose an API for the user’s VM to perform this programming for itself. Doing so may unnecessarily expose a vital part of the system, the dynamic partial reconfiguration port, to potential abuse (e.g., a denial of service attack that prevents or slows other Roles from being programmed onto the device). Rather, the cloud provider may modify their management system to program the Role region in advance, only starting the user’s VM after this operation is complete. This model also simplifies the Shell, software drivers, and VM facilities that must be provided, since no safeguards are needed to ensure that a given VM can only reprogram a specific Role region. In this case, the cloud management software will send commands to the hypervisor, similar to what occurs when a VM is allocated. This model may depart from the experience that some users may expect or need. For example, disallowing a user from reprogramming the FPGA directly changes the potential use-cases that can be supported and complicates the recovery process when a Role that has become unresponsive due to a bug in the user’s code.

Once the user’s Role is active on the device, the user’s VM must have exclusive access to the Role. Although the VM, driver, and Shell modifications needed to implement this access may be similar to SR-IOV for devices that use PCIe, this approach may require different techniques for devices that do not, e.g., network-attached standalone FPGAs. Furthermore, the Shell has responsibilities that extend beyond communication with a user’s software - it must virtualize all external resources. In the case of a resource like external memory, the Shell not only needs to segment the address space for different Roles, it also must make fairness or quality of service considerations between Roles. In the case of a resource like a network connection, this action likely requires support for virtual networking to enforce isolation across the network and quality of service considerations from both Roles and upstream switches.

Put together, the basic requirements needed to support the multi-tenant spatial model of execution indicate that all parts of the cloud eco-system will need to become more sophisticated. However, this added complexity presents new issues. For example, adding complexity naturally also increases the potential for security vulnerabilities. In the case of the Shell, the necessary changes also affect utilization. For example, the value proposition of the multi-tenant spatial model is the availability of hardware for different users, thereby minimizing idle FPGA space. However, if significant additional logic must be added to the Shell to support the multi-tenant spatial model, any potential advantage may be nullified.

5.10.4 FPGA Architecture Changes. In general, commercial FPGAs were not designed for simultaneous use by multiple independent users (the multi-tenant spatial model). To prevent attacks, additional security measures may need to be added to FPGA devices. One possible architectural enhancement is the isolation of PDNs to FPGA regions to avoid on-FPGA voltage attacks. Such isolation, which is common in multi-core microprocessors, may also allow for more flexible voltage scaling on a per-region basis. FPGAs could also be instrumented with additional hardened voltage and temperature monitors that can be read quickly. This information could then be used to isolate

and suppress voltage, thermal, and other attacks before they can cause damage. Finally, new techniques to dynamically reconfigure FPGA regions and save state more quickly may be needed to fully support rapid on-demand FPGA use in the cloud.

Improved bitstream security may spur interest in FPGA use as a root-of-trust in the data center. Although modern FPGAs contain PUFs used in key management, and specialized blocks for hashing and decryption of bitstreams, full breaks of modern FPGA bitstreams are commonplace [99]. Increased bitstream security will allow for a wider role for FPGAs in boot-strapping systems, incorporating roles which have historically used trusted platform modules (TPMs) to validate firmware.

6 APPLICATIONS

In this section we focus on customer applications that in most cases are run on public cloud services. We divide these applications into two categories, namely infrastructure-based and other applications.

6.1 Infrastructure

FPGAs have traditionally been used within networking equipment such as routers and switches. In addition to the base functionality provided by the equipment vendor, more flexible switching and routing capabilities were later introduced with protocols such as OpenFlow. An OpenFlow switch [168] was implemented using NetFPGA hardware [230] to handle traffic going through the electrical engineering and computer science building at Stanford University. Subsequently, NetFPGA SUME was introduced for prototyping 10, 40 and 100 Gbps applications [244]. FPGAs have also been used as disaggregated computing resources by directly connecting to the datacenter network as standalone resources [25]. The infrastructure of this cluster can be separated into privileged and non-privileged regions using partial reconfiguration [193]. In recent years, computing requirements on the data plane have increased significantly with the emergence of applications such as server virtualization and overlay tunneling. Performing such functions in software causes significant CPU overhead, and SmartNICs were used to implement highly complex server side data plane functions. Microsoft introduced FPGA-based SmartNICs to provide accelerated cloud services for Azure customers by offloading host networking to hardware [105]. In this architecture, an FPGA is placed in series with the network connection of the blade server (bump-in-the-wire) that enables pre- and post-processing of data. Several other cloud service providers have also started using SmartNICs in their acceleration applications [150, 237]. A framework [100] has been proposed that provides a software abstraction and an FPGA-based hardware runtime for accelerating general purpose applications in multi-tenant systems.

6.2 Other Applications

Numerous data center applications can benefit from using FPGA acceleration in areas such as database acceleration, big data, machine learning, security and privacy, some of which are described below.

6.2.1 Data Analytics. The core functionality of Memcached, a powerful open-source in-memory key-value store for small chunks of data, has been implemented on an FPGA [84]. This architecture has tightly integrated compute, network and memory resources so that it achieves significant benefits in terms of power consumption and performance compared to a baseline server. Xilinx has also built two versions of Memcached on FPGA [74, 75]. A methodology has been presented for integrating FPGAs into Apache Spark [87], an open source analytics engine used for big data processing. The cluster used in this work consists of a master node and six worker nodes each

of which is equipped with a PCIe-attached Alpha Data ADM-PCIE-7V3 FPGA board. It has been shown that the performance of Spark-FPGA integration scales well up to six worker nodes. A challenge of using Java-based analytics platforms like Spark is being able to efficiently share the application data managed by Java with the custom accelerator inside the FPGA [114, 115]. Samsung has proposed a near storage accelerator for database sort using a SmartSSD which consists of an integrated FPGA that has a direct connection (P2P) to the storage device [195]. Having a direct connection obviates the need to transfer data through the host memory, which incurs significant latency in addition to occupying host resources. The implementation of a k -means algorithm using multiple FPGAs in a heterogeneous computing cluster has been presented [88]. Researchers have shown the ability to sort large amounts of data on an FPGA to support non-volatile memory [135], an application that could be useful for cloud architectures in the future. Similarly, researchers have shown that FPGAs can accelerate analytics for relational databases using a single node [154], an approach that could be implemented in the cloud. Network attached FPGAs have been shown to significantly accelerate functions useful for database transactions, such as atomic broadcast in hardware, which has been demonstrated on a distributed key-value store implementation [133] that demonstrates low latency and high throughput. FLOEM is a system for accelerating network applications where reconfiguration is part of the NiC. FLOEM also uses key-value stores as a demonstrator of its efficacy [177].

6.2.2 Deep Learning. Computational requirements for processing AI workloads have grown significantly over the past few years, and cloud operators have opted to use FPGAs in applications involving DNNs because of their high power efficiency and performance compared to CPUs and GPUs. For some of the implementations, they were motivated by research outcomes in the field of machine learning. DNNWeaver [200] generates custom, synthesizable accelerators for deep neural networks that best match the needs of the DNN while providing high performance and efficiency gains for the target FPGA. Notably, DNNWeaver is FPGA agnostic, and has been demonstrated on both Intel and Xilinx FPGAs. The automatically generated accelerators deliver superior performance and efficiency compared to multicore CPU and GPU implementations. The redundancy associated with floating point number representation in CNNs used by CPUs and GPUs can be removed by quantizing weights and activations and retraining the neural network. FINN is a framework developed by Xilinx Research Labs to design flexible quantized neural network inference accelerators on FPGAs [76]. This work shows that simple arithmetic operations in binary neural networks (BNNs) are well suited for energy-efficient inference by fitting parameters entirely in on-chip memory (OCM), which enables high computational performance. Microsoft's Brainwave [89, 107] is a hardware architecture developed to perform real time AI calculations using a soft-core neural processing unit (NPU) implemented on an Intel Stratix 10 FPGA. This NPU uses reduced precision proprietary floating point formats, and enables competitive levels of performance and energy efficiency compared to hard NPUs without compromising the accuracy of the model. Brainwave was used as a cloud service [95] to accelerate ResNet-50 image classification for particle physics computing applications. An enhancement has been made to the Brainwave's NPU [80] to efficiently use the tensor blocks in an FPGA, and its performance has been compared against Nvidia's T4 and V100 GPUs. An FPGA acceleration platform using a supertile-based design method for general-purpose CNNs and image/video inference applications has also been introduced [238]. An interleaved task dispatching method has been used to scale up supertile units (SUs) to efficiently perform different types of convolutions. Training CNNs using an FPDeep framework has also been demonstrated [228]. Although the proposed architecture is not cloud-ready yet, the authors have provided some insights onto a possible future cloud deployment using a scalable FPGA cluster. The use of multiple FPGAs to implement more complex 3D CNNs is of research interest. Such

implementations have demonstrated higher inference throughput while preserving the end-to-end latency of a single FPGA implementation [73, 205].

6.2.3 Security and Privacy. An important application area for FPGAs in the cloud is the security and privacy of user data undergoing computation. Early research identified the necessity of having FPGAs secure client data from attackers and untrusted system administrators [97]. A virtualization framework was proposed and prototyped [85] that allows for the abstraction of cloud FPGA resources using accelerator pools (APs). Garbled circuits (GC) [104] provide a promising approach for the hardware acceleration of privacy preserving computation. This work demonstrated the implementation of a generic reconfigurable coarse-grained FPGA overlay architecture for secure function evaluation. GC has been deployed using AWS [129] and applied to privacy preserving machine learning on cloud servers using FPGAs [131]. A garbler circuit [129] can achieve a 15× speedup compared to a software implementation. The acceleration of homomorphic encryption is also being investigated [221]. Here, a domain-specific co-processor architecture is used to accelerate homomorphic function evaluation on encrypted data using AWS F1 FPGA instances. Researchers at Microsoft have demonstrated homomorphic encryption on FPGAs and demonstrated over 100x speedup compared to previous implementations [189].

6.3 Future Applications

FPGAs are increasingly being made available in user programmable cloud systems including those from Alibaba [245], AWS [30], and the Open Cloud Testbed [123]. The trend is to have these FPGAs be directly attached to the network to remove the latency incurred when data is transferred to and from the host. At the same time, FPGAs continue to grow in capacity and to incorporate larger and diverse memory elements. The applications of the future will take advantage of these properties to deliver graph processing [86], machine learning, security and privacy applications on large data as well as accelerating a host of scientific applications and applications not yet imagined.

7 CONCLUSIONS

This manuscript has discussed the recent and potential future evolution of FPGAs in cloud computing platforms and datacenters from the perspective of experts in this field of research. The manuscript focuses on the architectural organization, resource management, security, and application deployment of FPGAs used in cloud and datacenter environments. These disciplines must be well understood for cloud service providers (CSP) to effectively provide resources to users, and for users to efficiently use the resources. The evolution of these FPGA systems has been chronicled in the manuscript and current research provides insight into future development trajectories.

Towards GPUs and CPUs, FPGAs have the advantage of low latency, high throughput and energy efficient. While FPGAs will not replace CPU and GPUs, they will be a complementary alternative in the cloud in applications like computer vision, AI, security, crypto processing and more.

We believe that future FPGA-inclusive cloud systems and datacenters will consolidate the vast selection of choices available today into a smaller range of architectures, programming models, and security-enhanced deployments. This consolidation may mirror the integration of CPUs and FPGAs in desktop and embedded systems that has taken place over the last two decades. With increased available FPGA resources per device, many applications in the future will use only a portion of an entire FPGA, leading to a need for resource sharing and system multi-tenancy. As a result, FPGA architectures for cloud computing will require physical resource and protocol isolation to ensure user security across different tenants, and the containment of faults and attacks that can disrupt system operation. Enhancements in architecture, programming models, and security will lead to a

broader range of applications for FPGA-based cloud deployment and growing market share in the years to come.

The main obstacle for a broad adoption of FPGAs in datacenters is the lack of the software stack that allows easy deployment, management and scaling of FPGAs on the cloud. Providing a software stack similar to the one that has led to the success of CPUs/GPUs will increase architecture diversification in the cloud, with FPGAs operating as equal partners on the side of CPUs and GPUs. The foundations for broad adoption of FPGA in datacenters are currently being developed. Innovative companies such as VMACcel, Inaccel as well as the opensource initiatives such as the Cyborg project whose goal is the seamless integration of accelerator in OpenStack are paving the way.

REFERENCES

- [1] AWS Batch, User Guide. <https://docs.aws.amazon.com/batch/latest/userguide/what-is-batch.html>.
- [2] Cygnus. <https://www.ccs.tsukuba.ac.jp/wp-content/uploads/sites/14/2018/12/About-Cygnus.pdf> [Last accessed: April 29, 2021].
- [3] Enzian. <http://enzian.systems/> [Last accessed: April 29, 2021].
- [4] IBM SuperVessel, OpenPower Cloud. <https://www.research.ibm.com/labs/china/supervessel.html> [Last accessed: April 29, 2021].
- [5] Inaccel website. <https://inaccel.com/>.
- [6] Maxeller website. <https://www.maxeler.com/products/maxcloud/>.
- [7] Open power foundation. <https://openpowerfoundation.org/>.
- [8] Riscv website. <https://riscv.org/>.
- [9] VMACcel – cloud service provider for fpga-based computing. <https://www.vmaccel.com/>.
- [10] Xilinx Virtex UltraScale+. <https://www.xilinx.com/products/silicon-devices/fpga/virtex-ultrascale-plus.html>.
- [11] Programming MPC Systems. <https://www.maxeler.com/media/documents/MaxelerWhitePaperProgramming.pdf> [Last accessed: January 26, 2021], June 2013.
- [12] Cyborg FPGA Programming Service Proposal. <https://docs.openstack.org/cyborg/stein/specs/rocky/approved/cyborg-fpga-programming-proposal.html> [Last accessed: October 25, 2021], 2018.
- [13] Gator reconfigurable cloud computing. <https://smartsystems.ece.ufl.edu/research/projects/gatorrecc/> [Last accessed: June 2, 2021], 2018.
- [14] Cyborg. <https://wiki.openstack.org/wiki/Cyborg> [Last accessed: October 25, 2021], 2021.
- [15] Handel-C. <https://en.wikipedia.org/wiki/Handel-C/> [Last accessed: October 22, 2021], 2021.
- [16] Intel FPGA SDK for OpenCL Software Technology. <https://www.intel.com/content/www/us/en/software/programmable/sdk-for-openccl/overview.html> [Last accessed: January 26, 2021], 2021.
- [17] Intel oneAPI. <https://www.intel.com/content/www/us/en/developer/tools/oneapi/overview.html> [Last Accessed: October 22, 2021], 2021.
- [18] Intel Open FPGA Stack. <https://www.intel.com/content/www/us/en/programmable/solutions/acceleration-hub/acceleration-stack.html> [Last Accessed: January 26, 2021], 2021.
- [19] OpenStack Acceleration APIs. <https://docs.openstack.org/api-ref/accelerator/> [Last accessed: October 25, 2021], 2021.
- [20] SDAccel. <https://www.xilinx.com/products/design-tools/software-zone/sdaccel.html> [Last accessed: January 26, 2021], 2021.
- [21] Vitis unified software platform. <https://www.xilinx.com/products/design-tools/vitis/vitis-platform.html> [Last Accessed: January 26, 2021], 2021.
- [22] AWS. EC2 F1 Instances. <https://aws.amazon.com/ec2/instance-types/f1/> [Last accessed: April 29, 2021].
- [23] Mohamed S Abdelfattah, David Han, Andrew Bitar, Roberto DiCecco, Shane O’Connell, Nitika Shanker, Joseph Chu, Ian Prins, Joshua Fender, Andrew C Ling, et al. Dla: Compiler and fpga overlay for neural network inference acceleration. In *2018 28th International Conference on Field Programmable Logic and Applications (FPL)*, pages 411–4117. IEEE, 2018.
- [24] A. Abdelsalam, F. Clad, C. Filsfils, S. Salsano, G. Siracusano, and L. Veltri. Implementation of virtual network function chaining through segment routing in a Linux-based NFV infrastructure. In *2017 IEEE Conference on Network Softwarization (NetSoft)*, pages 1–5, 2017.
- [25] Francois Abel, Jagath Weerasinghe, Christoph Hagleitner, Beat Weiss, and Stephan Paredes. An FPGA platform for hyperscalers. *Proceedings - 2017 IEEE 25th Annual Symposium on High-Performance Interconnects, HOTI 2017*, pages 29–32, 2017.

- [26] Achronix. Speedster 7t FPGAs. <https://www.achronix.com/product/speedster7t-fpgas> [Last accessed: April 29, 2021].
- [27] Amran A Al-Aghbari and Muhammad ES Elrabaa. Cloud-based fpga custom computing machines for streaming applications. *IEEE Access*, 7:38009–38019, 2019.
- [28] Md Mahbub Alam, Shahin Tajik, Fatemeh Ganji, Mark Tehranipoor, and Domenic Forte. RAM-Jam: Remote temperature and voltage fault attack on FPGAs using memory collisions. In *Workshop on Fault Diagnosis and Tolerance in Cryptography*, pages 48–55, 2019.
- [29] Gustavo Alonso, Timothy Roscoe, David Cock, Muhsen Owaid, Kaan Kara, Dario Korolija, Zeke Wang, et al. Tackling hardware/software co-design from a database perspective. In *Proceedings of the 6th biennial Conference on Innovative Data Systems Research (CIDR), Amsterdam, Netherlands, January 2020.*, 2020.
- [30] Amazon.com, Inc. Amazon EC2 F1 Instances, 2021.
- [31] R. Amerson, R. J. Carter, W. B. Culbertson, P. Kuekes, and G. Snider. Teramac-Configurable Custom Computing. In *IEEE Symposium on FPGAs for Custom Computing Machines*, pages 32–38, 1995.
- [32] ARM. AMBA 4 AXI4-Stream Protocol Specification. Technical report, ARM, 2010.
- [33] ARM. AMBA AXI and ACE Protocol Specification. Technical report, ARM, 2011.
- [34] Mikhail Asiatici, Nithin George, Kizheppatt Vipin, Suhaib A Fahmy, and Paolo Ienne. Virtualized execution runtime for fpga accelerators in the cloud. *Ieee Access*, 5:1900–1910, 2017.
- [35] Alibaba. Alibaba FPGA Cloud. <https://www.alibabacloud.com/help/doc-detail/108504.htm> [Last accessed: April 29, 2021].
- [36] Baidu. FPGA Cloud. <https://cloud.baidu.com/product/fpga.html> [Last accessed: April 29, 2021].
- [37] Bittware. Computational Storage. <https://www.bittware.com/fpga/storage/> [Last accessed: April 29, 2021].
- [38] Blocks&Files. A brief look at AWS Redshift’s AQUA acceleration hardware. <https://blocksandfiles.com/2019/12/05/amazon-aqua-data-warehouse-acceleration-hardware/> [Last accessed: April 29, 2021].
- [39] Blocks&Files. Our Computational Storage Drives are Bigger, Faster and Cheaper than Ordinary SSDs. <https://blocksandfiles.com/2021/02/22/scaleflux-ceo-hao-zhong-interview/> [Last accessed: April 29, 2021].
- [40] Cisco. Nexus 3000 Switch Architecture. <https://www.ciscolive.com/c/dam/r/ciscolive/us/docs/2018/pdf/BRKDCN-3734.pdf> [Last accessed: April 29, 2021].
- [41] Huawei. CloudEngine Is the Foundation of the Intent-driven Network. <https://actfor.net.com/ueditor/php/upload/file/20191206/1575567949680852.pdf> [Last accessed: April 29, 2021].
- [42] Intel. SoC FPGAs. <https://www.intel.com/content/www/us/en/products/programmable/soc.html> [Last accessed: April 29, 2021].
- [43] Nimbix. FPGA Cloud. <https://www.nimbix.net/what-is-an-fpga> [Last accessed: April 29, 2021].
- [44] OVH. Acceleration-as-a-Service leveraging Intel PAC. https://www.ovh.com/world/news/press/cp2541.ovh_launches_acceleration-as-a-service_leveraging_the_new_intel_programmable_acceleration_card_and_app_store_from_fpga_acceleration_partner_accelize [Last accessed: April 29, 2021].
- [45] Scaleflux. Computational Storage. <https://www.scaleflux.com/> [Last accessed: April 29, 2021].
- [46] Tencent. FPGA Cloud Server. <https://cloud.tencent.com/product/fpga> [Last accessed: April 29, 2021].
- [47] TheNextPlatform. Computational Storage Winds its Way Towards the Mainstream. <https://www.nextplatform.com/2020/02/25/computational-storage-winds-its-way-towards-the-mainstream/> [Last accessed: April 29, 2021].
- [48] Xilinx. SmartSSD Computational Storage Drive. <https://www.xilinx.com/applications/data-center/computational-storage/smartssd.html> [Last accessed: April 29, 2021].
- [49] Xilinx. SoCs with Hardware and Software Programmability. <https://www.xilinx.com/products/silicon-devices/soc/zynq-7000.html> [Last accessed: April 29, 2021].
- [50] Baidu. Baidu ABC Platform. https://www.exascale.org/bdec/sites/www.exascale.org/bdec/files/Baidu_ABC_Platform.pdf [Last accessed: April 29, 2021].
- [51] Bittware. How OVHcloud Uses FPGAs to Mitigate DDoS Attacks. <https://www.bittware.com/resources/case-study-ovh/> [Last accessed: April 29, 2021].
- [52] FaBRIC. FPGA Accelerator Research Infrastructure Cloud (FABRIC). <https://wikis.utexas.edu/display/fabric/Home> [Last accessed: April 29, 2021].
- [53] Huawei. FPGA Cloud. <https://www.huaweicloud.com/en-us/product/fcs.html> [Last accessed: April 29, 2021].
- [54] Scaleflux. Applications. <https://www.scaleflux.com/news.html> [Last accessed: April 29, 2021].
- [55] TheNextPlatform. Baidu Takes FPGA Approach to Accelerating SQL at Scale. <https://www.nextplatform.com/2016/08/24/baidu-takes-fpga-approach-accelerating-big-sql/> [Last accessed: April 29, 2021].
- [56] Xilinx. FPGAs: The Key to Accelerating High-Speed Storage Systems. https://www.flashmemorysummit.com/Proceedings2019/08-07-Wednesday/20190807_Keynote11_Xilinx_Raje.pdf [Last accessed: April 29, 2021].
- [57] AWS. Accelerate applications using Amazon EC2 F1 FPGA instances. https://d1.awsstatic.com/events/reinvent/2019/Accelerate_applications_using_Amazon_EC2_F1_FPGA_instances_CMP314.pdf [Last accessed: April 29, 2021].

- [58] AWS. AQUA (Advanced Query Accelerator) for Amazon Redshift. https://pages.awscloud.com/AQUA_Preview.html [Last accessed: April 29, 2021].
- [59] Baidu. AI Cloud. <https://intl.cloud.baidu.com/product/abc-stack.html> [Last accessed: April 29, 2021].
- [60] Huawei. FPGA as a Service in the Cloud. https://indico.cern.ch/event/669648/contributions/2838181/attachments/1581893/2500031/Huawei_Cloud_FPGA_as_a_Service_CERN_openlab.pdf [Last accessed: April 29, 2021].
- [61] Intel. FPGA Programmable Acceleration Card D5005. https://www.intel.com/content/www/us/en/programmable/products/boards_and_kits/dev-kits/altera/intel-fpga-pac-d5005/overview.html [Last accessed: April 29, 2021].
- [62] Xilinx. Alveo SmartNIC Accelerator Card. <https://www.xilinx.com/products/boards-and-kits/alveo.html> [Last accessed: April 29, 2021].
- [63] Intel. Intel FPGAs Power Acceleration-as-a-Service for Alibaba Cloud. <https://newsroom.intel.com/news/intel-fpgas-power-acceleration-as-a-service-alibaba-cloud/#gs.uijhu> [Last accessed: April 29, 2021].
- [64] TheNextPlatform. Another Step towards FPGAs in Supercomputing. <https://www.nextplatform.com/2018/04/04/another-step-toward-fpgas-in-supercomputing/> [Last accessed: April 29, 2021].
- [65] Xilinx. Alveo SN1000 Accelerator Card. <https://www.xilinx.com/applications/data-center/network-acceleration/alveo-sn1000.html> [Last accessed: April 29, 2021].
- [66] Intel. Intelligent Infrastructure Transformation. <https://newsroom.intel.com/news/intel-baidu-drive-intelligent-infrastructure-transformation> [Last accessed: April 29, 2021].
- [67] Daniel O Awduche. Mpls and traffic engineering in ip networks. *IEEE communications Magazine*, 37(12):42–47, 1999.
- [68] Hitesh Ballani, Paolo Costa, Thomas Karagiannis, and Ant Rowstron. Towards predictable datacenter networks. In *Proceedings of the ACM SIGCOMM 2011 Conference, SIGCOMM '11*, page 242–253, New York, NY, USA, 2011. Association for Computing Machinery.
- [69] AWS. Marketplace. <https://aws.amazon.com/marketplace/search/results?x=0&y=0&searchTerms=fpga> [Last accessed: April 29, 2021].
- [70] R. Baxter, S. Booth, M. Bull, G. Cawood, J. Perry, M. Parsons, A. Simpson, A. Trew, A. McCormick, G. Smart, R. Smart, A. Cattle, R. Chamberlain, and G. Genest. Maxwell - a 64 FPGA Supercomputer. In *Second NASA/ESA Conference on Adaptive Hardware and Systems (AHS 2007)*, pages 287–294, 2007.
- [71] Christian Beckhoff, Dirk Koch, and Jim Torresen. Short-circuits on fpgas caused by partial runtime reconfiguration. In *International Conference on Field Programmable Logic and Applications*, pages 596–601, 2010.
- [72] D. Bernstein. Containers and cloud: From lxc to docker to kubernetes. *IEEE Cloud Computing*, 1(3):81–84, 2014.
- [73] Saman Biookaghazadeh, Pravin Kumar Ravi, and Ming Zhao. Toward multi-FPGA acceleration of the neural networks. *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 17(2):1–23, 2021.
- [74] Michaela Blott, Kimon Karras, Ling Liu, Kees Vissers, Jeremia Bär, and Zsolt István. Achieving 10gbps line-rate key-value stores with fpgas. In *5th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 13)*, San Jose, CA, June 2013. USENIX Association.
- [75] Michaela Blott, Ling Liu, Kimon Karras, and Kees Vissers. Scaling out to a single-node 80gbps memcached server with 40terabytes of memory. In *7th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 15)*, Santa Clara, CA, July 2015. USENIX Association.
- [76] Michaela Blott, Thomas B Preußner, Nicholas J Fraser, Giulio Gambardella, Kenneth O'brien, Yaman Umuroglu, Miriam Leaser, and Kees Vissers. FINN-R: An end-to-end deep-learning framework for fast exploration of quantized neural networks. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, 11(3):1–23, 2018.
- [77] C. Bobda, A. Majer, A. Ahmadinia, T. Haller, A. Linarth, and J. Teich. The erlangen slot machine: increasing flexibility in fpga-based reconfigurable platforms. In *Proceedings. 2005 IEEE International Conference on Field-Programmable Technology*, 2005., pages 37–42, 2005.
- [78] Christophe Bobda. *Introduction to Reconfigurable Computing: Architectures, Algorithms, and Applications*. Springer Publishing Company, Incorporated, 1st edition, 2007.
- [79] Marc Boulé and Zeljko Zilic. Automata-based assertion-checker synthesis of PSL properties. *ACM Trans. Des. Autom. Electron. Syst.*, 13(1):4:1–4:21, February 2008.
- [80] Andrew Boutros, Eriko Nurvitadhi, Rui Ma, Sergey Gribok, Zhipeng Zhao, James C. Hoe, Vaughn Betz, and Martin Langhammer. Beyond peak performance: Comparing the real performance of AI-Optimized FPGAs and GPUs. In *2020 International Conference on Field-Programmable Technology (ICFPT)*, pages 10–19, 2020.
- [81] Alexander Brant and Guy GF Lemieux. Zuma: An open fpga overlay architecture. In *Field-Programmable Custom Computing Machines (FCCM), 2012 IEEE 20th Annual International Symposium on*, pages 93–96. IEEE, 2012.
- [82] Stuart Byma, J Gregory Steffan, Hadi Bannazadeh, Alberto Leon Garcia, and Paul Chow. Fpgas in the cloud: Booting virtualized hardware accelerators with openstack. In *2014 IEEE 22nd Annual International Symposium on Field-Programmable Custom Computing Machines*, pages 109–116. IEEE, 2014.
- [83] A.M. Caulfield, E.S. Chung, A. Putnam, H. Angepat, Jeremy Fowers, Michael Haselman, Stephen Heil, Matt Humphrey, Puneet Kaur, Joo-Young Kim, Daniel Lo, Todd Massengill, Kalin Ovtcharov, Michael Papamichael, Lisa Woods,

- Sitaram Lanka, Derek Chiou, and Doug Burger. A cloud-scale acceleration architecture. In *49th IEEE/ACM Int. Symp. Microarchitecture*, pages 1–13, 2016.
- [84] Sai Rahul Chalamalasetti, Kevin Lim, Mitch Wright, Alvin AuYoung, Parthasarathy Ranganathan, and Martin Margala. An FPGA memcached appliance. In *Proceedings of the ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, page 245–254, 2013.
- [85] Fei Chen, Yi Shan, Yu Zhang, Yu Wang, Hubertus Franke, Xiaotao Chang, and Kun Wang. Enabling fpgas in the cloud. In *Proceedings of the 11th ACM Conference on Computing Frontiers*, page 3. ACM, 2014.
- [86] Xinyu Chen, Hongshi Tan, Yao Chen, Bingsheng He, Weng-Fai Wong, and Deming Chen. Thundergp: Hls-based graph processing framework on fpgas. In *The 2021 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 69–80, 2021.
- [87] Yu-Ting Chen, Jason Cong, Zhenman Fang, Jie Lei, and Peng Wei. When spark meets FPGAs: A case study for next-generation DNA sequencing acceleration. In *8th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 16)*, 2016.
- [88] Yuk-Ming Choi and Hayden Kwok-Hay So. Map-reduce processing of k-means algorithm with FPGA-accelerated computer cluster. In *2014 IEEE 25th International Conference on Application-Specific Systems, Architectures and Processors*, pages 9–16, 2014.
- [89] Chung, Eric and Powers, Jeremy and Ovtcharov, Kalin and Papamichael, Michael and Caulfield, Adrian and Massengill, Todd and Liu, Ming and Lo, Daniel and Alkalay, Shlomi and Haselman, Michael and others. Serving DNNs in Real Time at Datacenter Scale with Project Brainwave. *IEEE Micro*, 38(2):8–20, 2018.
- [90] Catalin Cimpano. Vast majority of cyber-attacks on cloud servers aim to mine cryptocurrency. <https://www.zdnet.com/article/vast-majority-of-cyber-attacks-on-cloud-servers-aim-to-mine-cryptocurrency/>, August 2020.
- [91] Chris Conger, Ian Troxel, D Espinoza, Vikas Aggarwal, and A George. NARC: Network Attached Reconfigurable Computing for High Performance, Network Based Applications. In *Proceedings of the Eighth Annual International Conference on Military and Aerospace Programmable Logic Devices (MAPLD'05)*, 2005.
- [92] CXL. CPU-to-Device Interconnect. <https://www.computeexpresslink.org/about-cxl> [Last accessed: April 29, 2021].
- [93] Guohao Dai, Yi Shan, Fei Chen, Yu Wang, Kun Wang, and Huazhong Yang. Online scheduling for fpga computation in the cloud. In *2014 International Conference on Field-Programmable Technology (FPT)*, pages 330–333. IEEE, 2014.
- [94] Tiziano De Matteis, Johannes de Fine Licht, Jakub Beránek, and Torsten Hoeffler. Streaming Message Interface: High-Performance Distributed Memory Programming on Reconfigurable Hardware. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–33, 2019.
- [95] Javier Duarte, Philip Harris, Scott Hauck, Burt Holzman, Shih-Chieh Hsu, Sergio Jindariani, Suffian Khan, Benjamin Kreis, Brian Lee, Mia Liu, et al. FPGA-accelerated machine learning inference as a service for particle physics computing. *Computing and Software for Big Science*, 3(1):1–15, 2019.
- [96] Sankha Baran Dutta, Hoda Naghibijouybari, Nael Abu-Ghazaleh, Andres Marquez, and Kevin Barker. Leaky buddies: Cross-component covert channels on integrated cpu-gpu systems. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*, pages 972–984. IEEE, 2021.
- [97] Ken Eguro and Ramarathnam Venkatesan. FPGAs for trusted cloud computing. In *22nd International Conference on Field Programmable Logic and Applications (FPL)*, pages 63–70, 2012.
- [98] Eideticomm. NVMe Computational Storage. <https://www.eideticom.com/> [Last accessed: April 29, 2021].
- [99] Maik Ender, Amir Moradi, and Christof Paar. The unpatchable silicon: A full break of the bitstream encryption of Xilinx 7-Series FPGAs. In *29th USENIX Security Symposium*, 2020.
- [100] Haggai Eran, Lior Zeno, Maroun Tork, Gabi Malka, and Mark Silberstein. NICA: An infrastructure for inline acceleration of network applications. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*, pages 345–362, 2019.
- [101] Dieter Ernst. Competing in Artificial Intelligence Chips: China’s Challenge amid Technology War. *Centre for International Governance Innovation, Special Report*, 2020.
- [102] Nariman Eskandari, Naif Tarafdar, Daniel Ly-Ma, and Paul Chow. A modular heterogeneous stack for deploying fpgas and cpus in the data center. In *Proceedings of the 2019 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays, FPGA ’19*, pages 262–271, New York, NY, USA, 2019. ACM.
- [103] Suhaib A Fahmy, Kizheppatt Vipin, and Shanker Shreejith. Virtualized fpga accelerators for efficient cloud computing. In *2015 IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom)*, pages 430–435. IEEE, 2015.
- [104] Xin Fang, Stratis Ioannidis, and Miriam Leeser. Secure function evaluation using an FPGA overlay architecture. In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 257–266, 2017.
- [105] Daniel Firestone, Andrew Putnam, Sambhrama Mundkur, Derek Chiou, Alireza Dabagh, Mike Andrewartha, Hari Angepat, Vivek Bhanu, Adrian Caulfield, Eric Chung, et al. Azure accelerated networking: SmartNICs in the public

- cloud. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 51–66, 2018.
- [106] Forbes. Xilinx FPGAs: The Chip Behind Alibaba’s Singles Day. <https://www.forbes.com/sites/moorinsights/2018/11/29/xilinx-fpgas-the-chip-behind-alibabas-singles-day/?sh=5f2294e27e3b> [Last accessed: April 29, 2021].
- [107] Jeremy Fowers, Kalin Ovtcharov, Michael Papamichael, Todd Massengill, Ming Liu, Daniel Lo, Shlomi Alkalay, Michael Haselman, Logan Adams, Mahdi Ghandi, Stephen Heil, Prerak Patel, Adam Sapek, Gabriel Weisz, Lisa Woods, Sitaram Lanka, Steven K. Reinhardt, Adrian M. Caulfield, Eric S. Chung, and Doug Burger. A configurable cloud-scale DNN processor for real-time AI. In *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*, pages 1–14, 2018.
- [108] Karl Freund. Google cloud tpu: Strategic implications for google, nvidia and the machine learning industry. <https://www.forbes.com/sites/moorinsights/2017/05/22/google-cloud-tpu-strategic-implications-for-google-nvidia-and-the-machine-learning-industry/?sh=69d2f5a13af7>, May 2017.
- [109] Norihisa Fujita, Ryohei Kobayashi, Yoshiaki Yamaguchi, and Taisuke Boku. Parallel Processing on FPGA Combining Computation and Communication in OpenCL Programming. In *2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 479–488. IEEE, 2019.
- [110] Juan Fumero, Michail Papadimitriou, Foivos S Zakkak, Maria Xekalaki, James Clarkson, and Christos Kotselidis. Dynamic application reconfiguration on heterogeneous hardware. In *Proceedings of the 15th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments*, pages 165–178, 2019.
- [111] William Gayde. How arm came to dominate the mobile market and it’s coming for more, much more. <https://www.techspot.com/article/1989-arm-inside>, dec 2020.
- [112] A. D. George, M. C. Herbordt, H. Lam, A. G. Lawande, J. Sheng, and C. Yang. Novo-G#: Large-Scale Reconfigurable Computing with Direct and Programmable Interconnects. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7, 2016.
- [113] Alan George, Herman Lam, and Greg Stitt. Novo-G: At the Forefront of Scalable Reconfigurable Supercomputing. *Computing in Science & Engineering*, 13(1):82–86, 2010.
- [114] Ehsan Ghasemi and Paul Chow. Accelerating apache spark big data analysis with fpgas. In *2016 Intl IEEE Conferences on Ubiquitous Intelligence Computing, Advanced and Trusted Computing, Scalable Computing and Communications, Cloud and Big Data Computing, Internet of People, and Smart World Congress (UIC/ATC/ScalCom/CBDCoM/IoP/SmartWorld)*, pages 737–744, 2016.
- [115] Ehsan Ghasemi and Paul Chow. Accelerating apache spark with fpgas. *Concurrency and Computation: Practice and Experience*, 31(2):e4222, 2019.
- [116] Ilias Giechaskiel, Kasper Rasmussen, and Jakub Szefer. Reading between the dies: Cross-SLR covert channels on multi-tenant cloud FPGAs. In *IEEE International Conference on Computer Design*, pages 1–10, 2019.
- [117] Ilias Giechaskiel, Kasper Rasmussen, and Jakub Szefer. CAPSULE: Cross-FPGA covert-channel attacks through power supply unit leakage. In *IEEE Symposium on Security and Privacy*, pages 1728–1741, May 2020.
- [118] Ilias Giechaskiel, Kasper B Rasmussen, and Ken Eguro. Leaky wires: Information leakage and covert communication between FPGA long wires. In *Asia Conference on Computer and Communications Security (ASIACCS)*, pages 15–27, 2018.
- [119] AWS GitHub. AFI Power. https://github.com/aws/aws-fpga/blob/master/hdk/docs/afi_power.md, 2020.
- [120] Dennis RE Gnad, Fabian Oboril, and Mehdi B Tahoori. Voltage drop-based fault attacks on FPGAs using valid bitstreams. In *International Conference on Field Programmable Logic and Applications (FPL)*, pages 1–7, 2017.
- [121] Google, Inc. Use a GPU. <https://www.tensorflow.org/guide/gpu> [Last accessed: October 19, 2021].
- [122] Christoph Hagleitner, Dionysios Diamantopoulos, Burkhard Ringlein, Constantinos Evangelinos, Charles Johns, Rong N. Chang, Bruce D’Amora, James A. Kahle, James Sexton, Michael Johnston, Edward Pyzer-Knapp, and Chris Ward. Heterogeneous computing systems for complex scientific discovery workflows. In *2021 Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 13–18, February 2021.
- [123] S Handagala, M Herbordt, and M Leeser. Oct: The open cloud fpga testbed. In *31st International Conference on Field Programmable Logic and Applications (FPL)*, 2021.
- [124] Reiner Hartenstein. Coarse grain reconfigurable architecture (embedded tutorial). In *Proceedings of the 2001 Asia and South Pacific Design Automation Conference*, pages 564–570. ACM, 2001.
- [125] Festus Hategekimana, Joel Mandebi Mbongue, Md Jubaer Hossain Pantho, and Christophe Bobda. Inheriting software security policies within hardware IP components. In *IEEE International Symposium on Field-Programmable Custom Computing Machines*, pages 53–56, 2018.
- [126] Festus Hategekimana, Joel Mandebi Mbongue, Md Jubaer Hossain Pantho, and Christophe Bobda. Secure hardware kernels execution in cpu+ fpga heterogeneous cloud. In *2018 International Conference on Field-Programmable Technology (FPT)*, pages 182–189. IEEE, 2018.
- [127] Nicole Hemsothd. First in-depth look at google’s new second-generation tpu. <https://www.nextplatform.com/2017/05/17/first-depth-look-googles-new-second-generation-tpu/>, May 2017.

- [128] Edson L. Horta, John W. Lockwood, David E. Taylor, and David Parlour. Dynamic hardware plugins in an fpga with partial run-time reconfiguration. In *Proceedings of the 39th Annual Design Automation Conference, DAC '02*, page 343–348, New York, NY, USA, 2002. Association for Computing Machinery.
- [129] Kai Huang, Mehmet Gungor, Xin Fang, Stratis Ioannidis, and Miriam Leeser. Garbled circuits in the cloud using FPGA enabled nodes. In *2019 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6, 2019.
- [130] Ted Huffmire, Brett Brotherton, Nick Callegari, Jonathan Valamehr, Jeff White, Ryan Kastner, and Tim Sherwood. Designing secure systems on reconfigurable hardware. *ACM Trans. Des. Autom. Electron. Syst.*, 13(3):44:1–44:24, July 2008.
- [131] Siam U Hussain, Bitu Darvish Rouhani, Mohammad Ghasemzadeh, and Farinaz Koushanfar. Maxelerator: FPGA accelerator for privacy preserving multiply-accumulate (MAC) on cloud servers. In *Proceedings of the 55th Annual Design Automation Conference*, pages 1–6, 2018.
- [132] Inventec. FPGA SmartNIC C5020X. <https://ebg.inventec.com/en/product/Accessories/Smart%20NIC%20Card/Inventec%20FPGA%20SmartNIC%20C5020X> [Last accessed: April 29, 2021].
- [133] Zsolt István, David Sidler, Gustavo Alonso, and Marko Vukolic. Consensus in a box: Inexpensive coordination in hardware. In *13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16)*, pages 425–438, 2016.
- [134] Chenglu Jin, Vasudev Gohil, Ramesh Karri, and Jeyavijayan Rajendran. Security of cloud FPGAs: A survey. *arxiv*, arXiv:2005.04867, 2020.
- [135] Sang-Woo Jun, Shuotao Xu, and Arvind. Terabyte sort on FPGA-accelerated flash storage. In *2017 IEEE 25th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pages 17–24, 2017.
- [136] Abdul Kabbani, Balajee Vamanan, Jahangir Hasan, and Fabien Duchene. Flowbender: Flow-level adaptive routing for improved latency and throughput in datacenter networks. In *Proceedings of the 10th ACM International on Conference on Emerging Networking Experiments and Technologies, CoNEXT '14*, page 149–160, New York, NY, USA, 2014. Association for Computing Machinery.
- [137] Chris Kachris, Elias Koromilas, and Ioannis Stamelos. Single-line deployment of Cloud FPGAs. https://cfdevops.github.io/cfDevOps21/slides/cfDevOps21_Kachris.pdf [Last accessed: October 19, 2021].
- [138] Nachiket Kapre and Jan Gray. Hoplite: Building austere overlay nocs for fpgas. In *2015 25th International Conference on Field Programmable Logic and Applications (FPL)*, pages 1–8. IEEE, 2015.
- [139] Ahmed Khawaja, Joshua Landgraf, Rohith Prakash, Michael Wei, Eric Schkufza, and Christopher J Roszbach. Sharing, protection, and compatibility for reconfigurable fabric with amorphos. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, pages 107–127, 2018.
- [140] Yeongu Kim, Ross Daly, Jeremie S. Kim, Chris Fallin, Jihye Lee, Donghyuk Lee, Chris B Wilkerson, Konrad K Lai, and Onur Mutlu. Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors. In *International Symposium on Computer Architecture*, pages 361–372, 2014.
- [141] Ryohei Kobayashi, Yuma Oobata, Norihisa Fujita, Yoshiaki Yamaguchi, and Taisuke Boku. OpenCL-Ready High Speed FPGA Network for Reconfigurable High Performance Computing. In *Proceedings of the International Conference on High Performance Computing in Asia-Pacific Region*, pages 192–201, 2018.
- [142] 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. Spectre attacks: Exploiting speculative execution. In *40th IEEE Symposium on Security and Privacy (S&P'19)*, 2019.
- [143] Dario Korolija, Timothy Roscoe, and Gustavo Alonso. Do {OS} abstractions make sense on fpgas? In *14th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 20)*, pages 991–1010, 2020.
- [144] S. Kumar Saha and C. Bobda. Fpga accelerated embedded system security through hardware isolation. In *2020 Asian Hardware Oriented Security and Trust Symposium (AsianHOST)*, pages 1–6, 2020.
- [145] Joshua Landgraf, Tiffany Yang, Will Lin, Christopher J Roszbach, and Eric Schkufza. Compiler-driven fpga virtualization with synergy. In *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 818–831, 2021.
- [146] Xiangwei Li, Abhishek Jain, Douglas Maskell, and Suhaib A Fahmy. An area-efficient fpga overlay using dsp block based time-multiplexed functional units. *arXiv preprint arXiv:1606.06460*, 2016.
- [147] Thomas Lin, Byungchul Park, Hadi Bannazadeh, and Alberto Leon-Garcia. Savi Testbed Architecture and Federation. In *Future Access Enablers of Ubiquitous and Intelligent Infrastructures*, pages 3–10. Springer, 2015.
- [148] 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. Meltdown: Reading kernel memory from user space. In *27th USENIX Security Symposium (USENIX Security 18)*, 2018.
- [149] Peter Loscocco and Stephen Smalley. Meeting critical security objectives with security-enhanced Linux. In *Ottawa Linux Symposium*, pages 115–134, 2001.
- [150] Layong Larry Luo and TEG Tencent. Towards converged SmartNIC architecture for bare metal & public clouds, 2018.

- [151] Jiacheng Ma, Gefei Zuo, Kevin Loughlin, Xiaohe Cheng, Yanqiang Liu, Abel Mulugeta Eneyew, Zhengwei Qi, and Baris Kasikci. A hypervisor for shared-memory fpga platforms. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 827–844, 2020.
- [152] M. Ma and V. W. S. Wong. An optimal peak hour content server cache update scheduling algorithm for 5g hetnets. In *ICC 2019 - 2019 IEEE International Conference on Communications (ICC)*, pages 1–6, 2019.
- [153] Sen Ma, Zeyad Aklah, and David Andrews. A run time interpretation approach for creating custom accelerators. In *2015 25th International Conference on Field Programmable Logic and Applications (FPL)*, pages 1–4. IEEE, 2015.
- [154] Divya Mahajan, Joon Kyung Kim, Jacob Sacks, Adel Ardalan, Arun Kumar, and Hadi Esmaeilzadeh. In-RDBMS Hardware Acceleration of Advanced Analytics. *Proceedings of the VLDB Endowment*, 11(11):1317–1331, July 2018.
- [155] Mallik Mahalingam, Dinesh G Dutt, Kenneth Duda, Puneet Agarwal, Lawrence Kreeger, T Sridhar, Mike Bursell, and Chris Wright. Virtual extensible local area network (vxlan): A framework for overlaying virtualized layer 2 networks over layer 3 networks. *RFC*, 7348:1–22, 2014.
- [156] Pongstorn Maidee, Alireza Kaviani, and Kevin Zeng. Linkblaze: Efficient global data movement for fpgas. In *2017 International Conference on ReConfigurable Computing and FPGAs (ReConFig)*, pages 1–8. IEEE, 2017.
- [157] Joel Mandebi Mbongue, Sujan Kumar Saha, and Christophe Bobda. Domain isolation in fpga-accelerated cloud and data center applications. In *Proceedings of the 2021 on Great Lakes Symposium on VLSI*, pages 283–288, 2021.
- [158] Joel Mandebi Mbongue, Danielle Tchuinkou Kwadjo, and Christophe Bobda. Flexitask: A flexible fpga overlay for efficient multitasking. In *Proceedings of the 2018 on Great Lakes Symposium on VLSI*, pages 483–486. ACM, 2018.
- [159] Joel Mbongue, Festus Hategekimana, Danielle Tchuinkou Kwadjo, David Andrews, and Christophe Bobda. Fpgavirt: A novel virtualization framework for fpgas in the cloud. In *2018 IEEE 11th International Conference on Cloud Computing (CLOUD)*, pages 862–865. IEEE, 2018.
- [160] Joel Mandebi Mbongue, Danielle Tchuinkou Kwadjo, and Christophe Bobda. Automatic generation of application-specific fpga overlays with rapidwright. In *2019 International Conference on Field-Programmable Technology (ICFPT)*, pages 303–306. IEEE, 2019.
- [161] Joel Mandebi Mbongue, Danielle Tchuinkou Kwadjo, Alex Shuping, and Christophe Bobda. Deploying multi-tenant fpgas within linux-based cloud infrastructure. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, 15(2):1–31, 2021.
- [162] Joel Mandebi Mbongue, Sujan Kumar Saha, and Christophe Bobda. Performance study of multi-tenant cloud fpgas. In *2021 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 168–171. IEEE, 2021.
- [163] Joel Mandebi Mbongue, Sujan Kumar Saha, and Christophe Bobda. A security architecture for domain isolation in multi-tenant cloud fpgas. In *2021 IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, pages 290–295. IEEE, 2021.
- [164] Joel Mandebi Mbongue, Alex Shuping, Pankaj Bhowmik, and Christophe Bobda. Architecture support for fpga multi-tenancy in the cloud. In *2020 IEEE 31st International Conference on Application-specific Systems, Architectures and Processors (ASAP)*, pages 125–132. IEEE, 2020.
- [165] Mellanox. Innova-2 Flex Open Programmable SmartNIC. <https://www.mellanox.com/files/doc-2020/pb-innova-2-flex.pdf> [Last accessed: April 29, 2021].
- [166] Microsoft Azure. Deploy ML models to FPGAs with Azure Machine Learning. <https://azure.microsoft.com/en-us/pricing/details/virtual-machines/windows/> [Last accessed: April 29, 2021].
- [167] Thierry Moreau, Tianqi Chen, Luis Vega, Jared Roesch, Eddie Yan, Lianmin Zheng, Josh Fromm, Ziheng Jiang, Luis Ceze, Carlos Guestrin, and Arvind Krishnamurthy. A hardware–software blueprint for flexible deep learning specialization. *IEEE Micro*, 39(5):8–16, 2019.
- [168] Jad Naous, David Erickson, G Adam Covington, Guido Appenzeller, and Nick McKeown. Implementing an OpenFlow switch on the NetFPGA platform. In *Proceedings of the 4th ACM/IEEE Symposium on Architectures for Networking and Communications Systems*, pages 1–9, 2008.
- [169] Napatech. FPGA acceleration cards. <https://www.napatech.com/products/> [Last accessed: April 29, 2021].
- [170] Gil Neiger, Amy Santoni, Felix Leung, Dion Rodgers, and Rich Uhlig. Intel virtualization technology: Hardware support for empy processor virtualization. *Intel Technology Journal*, 10(3), 2006.
- [171] NewWaveDV. 32-Port Programmable Switch. <http://newwavedv.com/wordpress/wp-content/uploads/2019/04/32-Port-Programmable-Switch-Datasheet.pdf> [Last accessed: April 29, 2021].
- [172] Jonas Ney, Dominik Loroch, Vladimir Rybalkin, Nico Weber, Jens Krüger, and Norbert Wehn. HALF: Holistic Auto Machine Learning for FPGAs. In *Proceedings of the 31st IEEE International Conference on Field-Programmable Logic and Applications (FPL)*, pages –. IEEE, 2021.
- [173] M. Noormohammadpour and C. S. Raghavendra. Datacenter traffic control: Understanding techniques and tradeoffs. *IEEE Communications Surveys Tutorials*, 20(2):1492–1525, 2018.

- [174] OpenCAPI. A New Standard for High Performance Memory, Acceleration and Networks. <https://opencapi.org/2017/04/opencapi-new-standard-high-performance-memory-acceleration-networks/> [Last accessed: April 29, 2021].
- [175] Opeyemi Osanaiye, Kim-Kwang Raymond Choo, and Mqhele Dlodlo. Distributed denial of service (DDoS) resilience in cloud: Review and conceptual cloud DDoS mitigation framework. *Journal of Network and Computer Applications*, 67:147–165, May 2016.
- [176] Paderborn Center for Parallel Computing, University of Paderborn. Noctua. <https://pc2.uni-paderborn.de/hpc-services/available-systems/noctua> [Last accessed: April 29, 2021].
- [177] Phitchaya Mangpo Phothilimthana, Ming Liu, Antoine Kaufmann, Simon Peter, Rastislav Bodik, and Thomas Anderson. Floem: A programming system for nic-accelerated network applications. In *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, pages 663–679, 2018.
- [178] Christian Pilato, Stanislav Bohm, Fabien Brocheton, Jeronimo Castrillon, Riccardo Cevasco, Vojtech Cima, Radim Cmar, Dionysios Diamantopoulos, Fabrizio Ferrandi, Jan Martinovic, Gianluca Palermo, Michele Paolino, Antonio Parodi, Lorenzo Pittaluga, Daniel Raho, Francesco Regazzoni, Katerina Slaninova, and Christoph Hagleitner. Everest: A design environment for extreme-scale big data analytics on heterogeneous platforms. In *2021 Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 1320–1325, February 2021.
- [179] Christian Plessl. Bringing FPGAs to HPC Production Systems and Codes. In *H2RC’18 workshop at Supercomputing (SC’18)*, 2018. doi: 10.13140/RG.2.2.34327.42407.
- [180] PRNewsWire. NVMe Production Ready System. <https://www.prnewswire.com/news-releases/eideticom-ibm-rackspace-and-xilinx-demonstrate-worlds-first-pcie-gen4-nvm-express-production-ready-system-676532203.html> [Last accessed: April 29, 2021].
- [181] George Provelengios, Daniel Holcomb, and Russell Tessier. Characterizing power distribution attacks in multi-user FPGA environments. In *International Conference on Field Programmable Logic and Applications (FPL)*, pages 194–201, 2019.
- [182] George Provelengios, Daniel Holcomb, and Russell Tessier. Power wasting circuits for cloud FPGA attacks. In *International Conference on Field Programmable Logic and Applications (FPL)*, 2020.
- [183] A. Putnam, A. M. Caulfield, E. S. Chung, D. Chiou, K. Constantinides, J. Demme, H. Esmailzadeh, J. Fowers, G. P. Gopal, J. Gray, M. Haselman, S. Hauck, S. Heil, A. Hormati, J. Kim, S. Lanka, J. Larus, E. Peterson, S. Pope, A. Smith, J. Thong, P. Y. Xiao, and D. Burger. A Reconfigurable Fabric for Accelerating Large-Scale Data Center Services. In *2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA)*, pages 13–24, 2014.
- [184] PyTorch community. Training a Classifier on GPU. https://pytorch.org/tutorials/beginner/blitz/cifar10_tutorial.html#training-on-gpu [Last accessed: October 19, 2021].
- [185] Arzhang Rafii, Welson Sun, and Paul Chow. Pharos: a multi-fpga performance monitor. In *2021 31st International Conference on Field-Programmable Logic and Applications (FPL)*, pages 257–262, 2021.
- [186] Chethan Ramesh, Shivukumar B Patil, Siva Nishok Dhanuskodi, George Provelengios, Sébastien Pillement, Daniel Holcomb, and Russell Tessier. FPGA side channel attacks without physical access. In *IEEE 26th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pages 45–52, 2018.
- [187] Kasper Rasmussen, Ilias Giechaskiel, and Ken Eguro. Leakier wires: Exploiting FPGA long wires for covert-and side-channel attacks. *ACM Transactions on Reconfigurable Technology and Systems*, 12(3):11:1–11:29, 2019.
- [188] Sandip Ray and Yier Jin. Security policy enforcement in modern SoC designs. In *Proceedings of the IEEE/ACM International Conference on Computer-Aided Design*, pages 345–350, 2015.
- [189] M Sadegh Riazi, Kim Laine, Blake Pelton, and Wei Dai. Heax: An architecture for computing on encrypted data. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 1295–1309, 2020.
- [190] B. Ringlein, F. Abel, , D. Diamantopoulos, B. Weiss, C. Hagleitner, M. Reichenbach, and D. Fey. A Case for Function-as-a-Service with Disaggregated FPGAs. In *Proceedings of the 2021 IEEE 14th International Conference on Cloud Computing (CLOUD 2021)*, pages 333–344, Virtual Conference, September 2021. IEEE.
- [191] B. Ringlein, F. Abel, A. Ditter, B. Weiss, C. Hagleitner, and D. Fey. Programming Reconfigurable Heterogeneous Computing Clusters Using MPI With Transpilation. In *2020 IEEE/ACM International Workshop on Heterogeneous High-performance Reconfigurable Computing (H2RC)*, pages 1–9. IEEE, November 2020.
- [192] B. Ringlein, F. Abel, A. Ditter, B. Weiss, C. Hagleitner, and D. Fey. ZRLMPI: A Unified Programming Model for Reconfigurable Heterogeneous Computing Clusters. In *2020 IEEE 28th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, page 220. IEEE, May 2020.
- [193] Burkhard Ringlein, Francois Abel, Alexander Ditter, Beat Weiss, Christoph Hagleitner, and Dietmar Fey. System Architecture for Network-Attached FPGAs in the Cloud using Partial Reconfiguration. In *2019 29th International Conference on Field Programmable Logic and Applications (FPL)*, pages 293–300. IEEE, 2019.
- [194] John M Rushby. Proof of separability a verification technique for a class of security kernels. In *International Symposium on Programming*, pages 352–367, 1982.

- [195] Sahand Salamat, Armin Haj Aboutalebi, Behnam Khaleghi, Joo Hwan Lee, Yang Seok Ki, and Tajana Rosing. NASCENT: Near-storage acceleration of database sort on SmartSSD. In *The 2021 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, FPGA '21, page 262–272, 2021.
- [196] Grigory Sapunov. Will asic chips become the next big thing in ai? <https://moorinsightsstrategy.com/will-asic-chips-become-the-next-big-thing-in-ai/>, jan 2021.
- [197] Falk Schellenberg, Dennis RE Gnad, Amir Moradi, and Mehdi B Tahoori. An inside job: Remote power analysis attacks on FPGAs. In *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1111–1116, 2018.
- [198] Sdxcentral. Broadcom Sharpens Tomahawk Switch Chips, Versatile SmartToR. <https://www.sdxcentral.com/articles/news/broadcom-sharpens-tomahawk-switch-chips-versatile-smarttor/2020/12/> [Last accessed: April 29, 2021].
- [199] Omar Sefraoui, Mohammed Aissaoui, and Mohsine Eleuldj. Openstack: toward an open-source solution for cloud computing. *International Journal of Computer Applications*, 55(3):38–42, 2012.
- [200] Hardik Sharma, Jongse Park, Divya Mahajan, Emmanuel Amaro, Joon Kyung Kim, Chenkai Shao, Asit Mishra, and Hadi Esmaeilzadeh. From high-level deep neural models to fpgas. In *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 1–12. IEEE, 2016.
- [201] S. B. Shaw, C. Kumar, and A. K. Singh. Use of time-series based forecasting technique for balancing load and reducing consumption of energy in a cloud data center. In *2017 International Conference on Intelligent Computing and Control (I2C2)*, pages 1–6, 2017.
- [202] Silicom. FPGA SmartNIC N5010 Series. https://www.silicom-usa.com/pr/fpga-based-cards/fpga-intel-based/fpga-intel-stratix-based/silicom-fpga-smartnic-n5010_series/ [Last accessed: April 29, 2021].
- [203] Jim Smith and Ravi Nair. *Virtual machines: versatile platforms for systems and processes*. Elsevier, 2005.
- [204] Hayden Kwok-Hay So and Cheng Liu. Fpga overlays. In *FPGAs for Software Programmers*, pages 285–305. Springer, 2016.
- [205] Mengshu Sun, Pu Zhao, Mehmet Gungor, Massoud Pedram, Miriam Leeser, and Xue Lin. 3D CNN acceleration on FPGA using hardware-aware pruning. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6, 2020.
- [206] Synopsys. An Introduction to CCIX. <https://www.synopsys.com/designware-ip/technical-bulletin/introduction-ccix-2017q3.html> [Last accessed: April 29, 2021].
- [207] Jakub Szefer. Survey of microarchitectural side and covert channels, attacks, and defenses. *Journal of Hardware and Systems Security*, 3(3):219–234, Sept. 2019.
- [208] David Talbot. Vulnerability seen in Amazon’s cloud-computing. <https://www.technologyreview.com/2009/10/23/208662/vulnerability-seen-in-amazons-cloud-computing/>, October 2009.
- [209] N. Tarafdar, N. Eskandari, V. Sharma, C. Lo, and P. Chow. Galapagos: A full stack approach to fpga integration in the cloud. *IEEE Micro*, 38(06):18–24, nov 2018.
- [210] Naif Tarafdar and Paul Chow. libGalapagos: A Software Environment for Prototyping and Creating Heterogeneous FPGA and CPU Applications. In *Sixth International Workshop on FPGAs for Software Programmers (FSP 2019)*, September 2019.
- [211] Naif Tarafdar, Nariman Eskandari, Thomas Lin, and Paul Chow. Designing for FPGAs in the Cloud. *IEEE Design and Test*, 35(1):23–29, 2017.
- [212] Naif Tarafdar, Thomas Lin, Eric Fukuda, Hadi Bannazadeh, Alberto Leon-Garcia, and Paul Chow. Enabling flexible network fpga clusters in a heterogeneous cloud data center. In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 237–246, 2017.
- [213] Naif Tarafdar, Thomas Lin, Daniel Ly-Ma, Daniel Rozhko, Alberto Leon-Garcia, and Paul Chow. Building the Infrastructure for Deploying FPGAs in the Cloud. In *Hardware Accelerators in Data Centers*, pages 9–33. Springer, 2019.
- [214] Naif Tarafdar, Thomas Lin, Daniel Ly-Ma, Daniel Rozhko, Alberto Leon-Garcia, and Paul Chow. Building the infrastructure for deploying fpgas in the cloud. In *Hardware Accelerators in Data Centers*, pages 9–33. Springer, 2019.
- [215] Impulse Accelerated Technologies. <https://web.archive.org/web/20110904033728/http://www.impulseaccelerated.com/> [Last accessed: October 22, 2021], 2021.
- [216] The Apache Software Foundation / TVM community. Quick Start Tutorial for Compiling Deep Learning Models. https://tvm.apache.org/docs/tutorial/relay_quick_start.html [Last accessed: October 19, 2021].
- [217] Shanquan Tian, Ilias Giechaskiel, Wenjie Xiong, and Jakub Szefer. Cloud FPGA cartography using PCIe contention. In *IEEE International Symposium on Field-Programmable Custom Computing Machines*, May 2021.
- [218] Shanquan Tian and Jakub Szefer. Temporal thermal covert channels in cloud fpgas. In *Proceedings of the International Symposium on Field-Programmable Gate Arrays*, FPGA, February 2019.
- [219] Shanquan Tian, Wenjie Xiong, Ilias Giechaskiel, Kasper Rasmussen, and Jakub Szefer. Fingerprinting cloud fpga infrastructures. In *Proceedings of the International Symposium on Field-Programmable Gate Arrays*, FPGA, February 2020.

- [220] TIBCO. Customers See Success with TIBCO on AWS. <https://www.tibco.com/blog/2018/11/26/tibco-customers-see-success-with-tibco-on-aws/> [Last accessed: April 29, 2021].
- [221] Furkan Turan, Sujoy Sinha Roy, and Ingrid Verbauwhede. HEAWS: An accelerator for homomorphic encryption on the amazon AWS FPGA. *IEEE Transactions on Computers*, 69(8):1185–1196, 2020.
- [222] Furkan Turan and Ingrid Verbauwhede. Trust in FPGA-accelerated cloud computing. *ACM Computing Surveys*, 53(6):28:1–28:128, December 2020.
- [223] Mohammad Usmani, Shahrzad Keshavarz, Eric Matthews, Lesley Shannon, Russell Tessier, and Daniel E Holcomb. Efficient PUF-based key generation in FPGAs using per-device configuration. *IEEE Transactions on VLSI Systems*, 27(2):364–375, February 2019.
- [224] Juan Camilo Vega, Qianfeng Clark Shen, Alberto Leon-Garcia, and Paul Chow. Introducing recpri: A field reconfigurable protocol for backhaul communication in a radio access network. In *2019 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, pages 329–336, 2019.
- [225] Malte Vesper, Dirk Kocha, and Khoa Phama. Pciehls: an opencl hls framework. In *FSP 2017; Fourth International Workshop on FPGAs for Software Programmers*, pages 1–6. VDE, 2017.
- [226] Hasitha Muthumala Waidyasooriya and Masanori Hariyama. Multi-FPGA Accelerator Architecture for Stencil Computation Exploiting Spatial and Temporal Scalability. *IEEE Access*, 7:53188–53201, 2019.
- [227] Herbert Walder and Marco Platzner. A runtime environment for reconfigurable hardware operating systems. In Jürgen Becker, Marco Platzner, and Serge Vernalde, editors, *Field Programmable Logic and Application*, pages 831–835, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.
- [228] Tianqi Wang, Tong Geng, Ang Li, Xi Jin, and Martin Herbordt. FPDeep: Scalable acceleration of CNN training on deeply-pipelined FPGA clusters. *IEEE Transactions on Computers*, 69(8):1143–1158, 2020.
- [229] X. Wang, Y. Niu, F. Liu, and Z. Xu. When fpga meets cloud: A first look at performance. *IEEE Transactions on Cloud Computing*, pages 1–1, 2020.
- [230] Greg Watson, Nick McKeown, and Martin Casado. NetFPGA: A tool for network research and education. In *2nd Workshop on Architectural Research using FPGA Platforms (WARFP)*, volume 3, 2006.
- [231] J. Weerasinghe, F. Abel, C. Hagleitner, and A. Herkersdorf. Enabling FPGAs in Hyperscale Data Centers. In *2015 IEEE 12th Intl Conf on Ubiquitous Intelligence and Computing and 2015 IEEE 12th Intl Conf on Autonomic and Trusted Computing and 2015 IEEE 15th Intl Conf on Scalable Computing and Communications and Its Associated Workshops (UIC-ATC-ScalCom)*, pages 1078–1086, 2015.
- [232] J. Weerasinghe, F. Abel, C. Hagleitner, and A. Herkersdorf. Disaggregated FPGAs: Network Performance Comparison against Bare-Metal Servers, Virtual Machines and Linux Containers. In *2016 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, pages 9–17, 2016.
- [233] J. Weerasinghe, R. Polig, F. Abel, and C. Hagleitner. Network-Attached FPGAs for Data Center Applications. In *2016 International Conference on Field-Programmable Technology (FPT)*, pages 36–43, 2016.
- [234] David Wilson and Greg Stitt. Seiba: An fpga overlay-based approach to rapid application development. In *2019 International Conference on ReConfigurable Computing and FPGAs (ReConFig)*, pages 1–8. IEEE, 2019.
- [235] Xilinx Case Study. Xilinx Powers Alibaba Cloud FaaS with AI Acceleration Solution for E-Commerce Business. <https://www.xilinx.com/publications/powered-by-xilinx/xilinx-alibaba-case-study.pdf> [Last accessed: April 29, 2021].
- [236] Xilinx Corporation. *Virtex UltraScale+ FPGA Data Sheet: DC and AC Switching Characteristics*, 2021.
- [237] Wang Xu. Hardware acceleration over NFV in China Mobile. *OPNFV Plugfest*. (Accessed: Jan. 2019). June, 2018.
- [238] Xiaoyu Yu, Yuwei Wang, Jie Miao, Ephrem Wu, Heng Zhang, Yu Meng, Bo Zhang, Biao Min, Dewei Chen, and Jianlin Gao. A data-center FPGA acceleration platform for convolutional neural networks. In *2019 29th International Conference on Field Programmable Logic and Applications (FPL)*, pages 151–158, 2019.
- [239] Shaza Zeitouni, Ghada Dessouky, and Ahmad-Reza Sadeghi. SoK: On the security challenges and risks of multi-tenant FPGAs in the cloud. *arxiv*, arXiv:2009.13914, 2020.
- [240] Shaza Zeitouni, Jo Vliegen, Tommaso Frassetto, Dirk Koch, Ahmad-Reza Sadeghi, and Nele Mentens. Trusted configuration in cloud FPGAs. In *IEEE International Symposium on Field-Programmable Custom Computing Machines*, 2021.
- [241] Jiansong Zhang, Yongqiang Xiong, Ningyi Xu, Ran Shu, Bojie Li, Peng Cheng, Guo Chen, and Thomas Moscibroda. The feniks fpga operating system for cloud computing. In *Proceedings of the 8th Asia-Pacific Workshop on Systems*, pages 1–7, 2017.
- [242] Ke Zhang, Yisong Chang, Mingyu Chen, Yungang Bao, and Zhiwei Xu. Computer organization and design course with fpga cloud. In *Proceedings of the 50th ACM Technical Symposium on Computer Science Education*, pages 927–933. ACM, 2019.
- [243] Mark Zhao and G Edward Suh. FPGA-based remote power side-channel attacks. In *IEEE Symposium on Security and Privacy (S&P)*, pages 229–244, 2018.

- [244] Noa Zilberman, Yury Audzevich, G Adam Covington, and Andrew W Moore. NetFPGA SUME: Toward 100 Gbps as research commodity. *IEEE Micro*, 34(5):32–41, 2014.
- [245] Xiantao Zxt, Zhengxiao Zx, and Justin Song. High-density multi-tenant bare-metal cloud with memory expansion soc and power management. In *2020 IEEE Hot Chips 32 Symposium (HCS)*, pages 1–18, 2020.