

# FPGAs Enable the Next Generation of Communication and Networking Solutions (WP021)



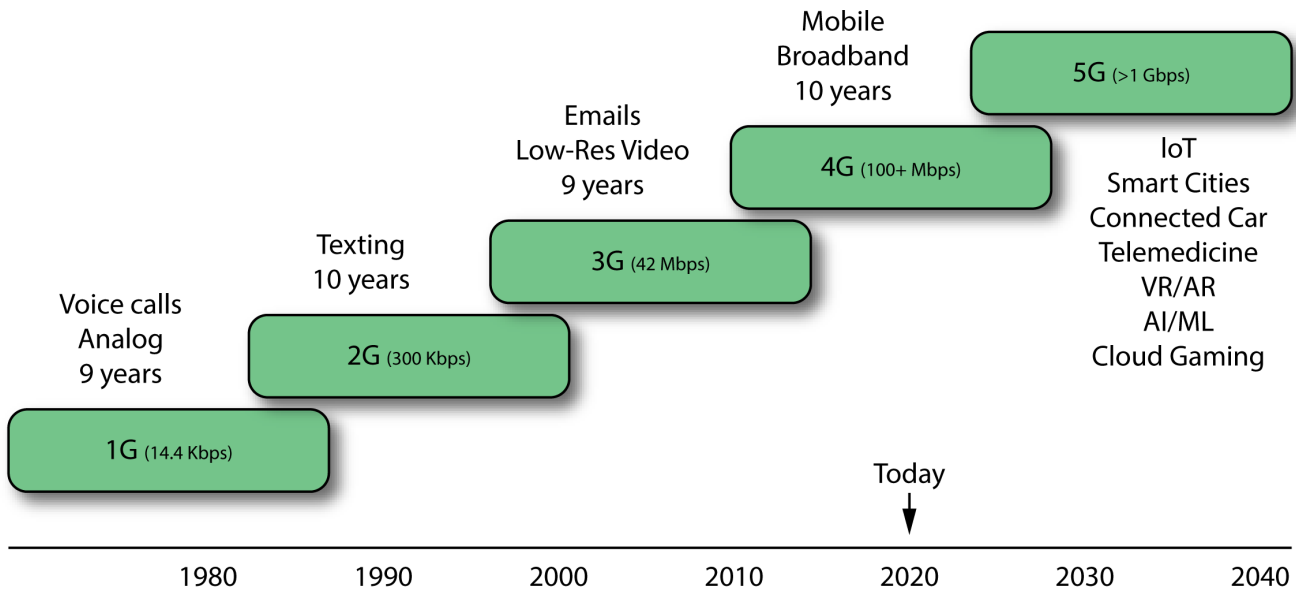
November 10, 2020

White Paper

## Executive Summary

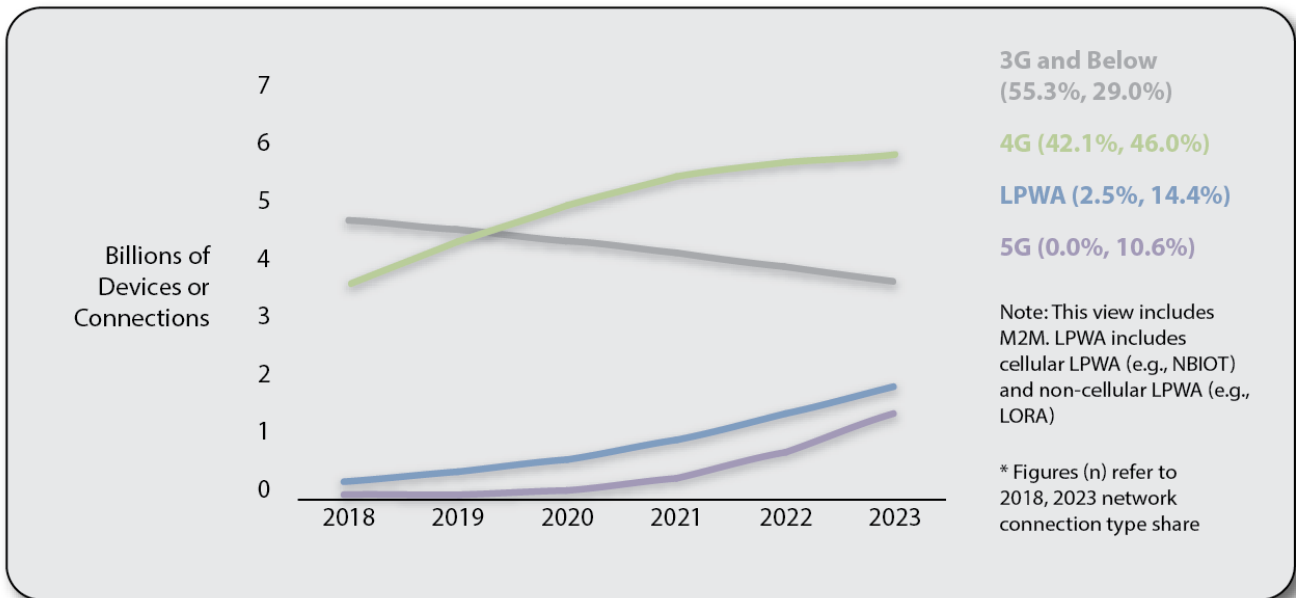
An easy way to comprehend the surge in network infrastructure capabilities is to look at the evolution over the last four decades (figure below). Innovations in cellular networks technology coupled with novel data storage and search techniques are becoming transformational in the industry's growth. It not only enables new use cases for companies and individuals alike, but also makes them think hard on leveraging technologies that otherwise would not have been part of their product mix. Perhaps the most telling change has been in new commercial models resulting in the value migration from infrastructure to services.

Connected devices are evolving (from 2G to 5G) to support the eruption in mobile applications and expanded connectivity to an ever increasing user base. A maturing industry calls for a competitive business model which translates to optimizing bandwidth management. It is estimated that the number of connected devices will be more than three times the global population by the year 2023 (figure following (see page 2)).



64721558-01.2020.30.04

**Figure 1: Network Infrastructure Evolution**

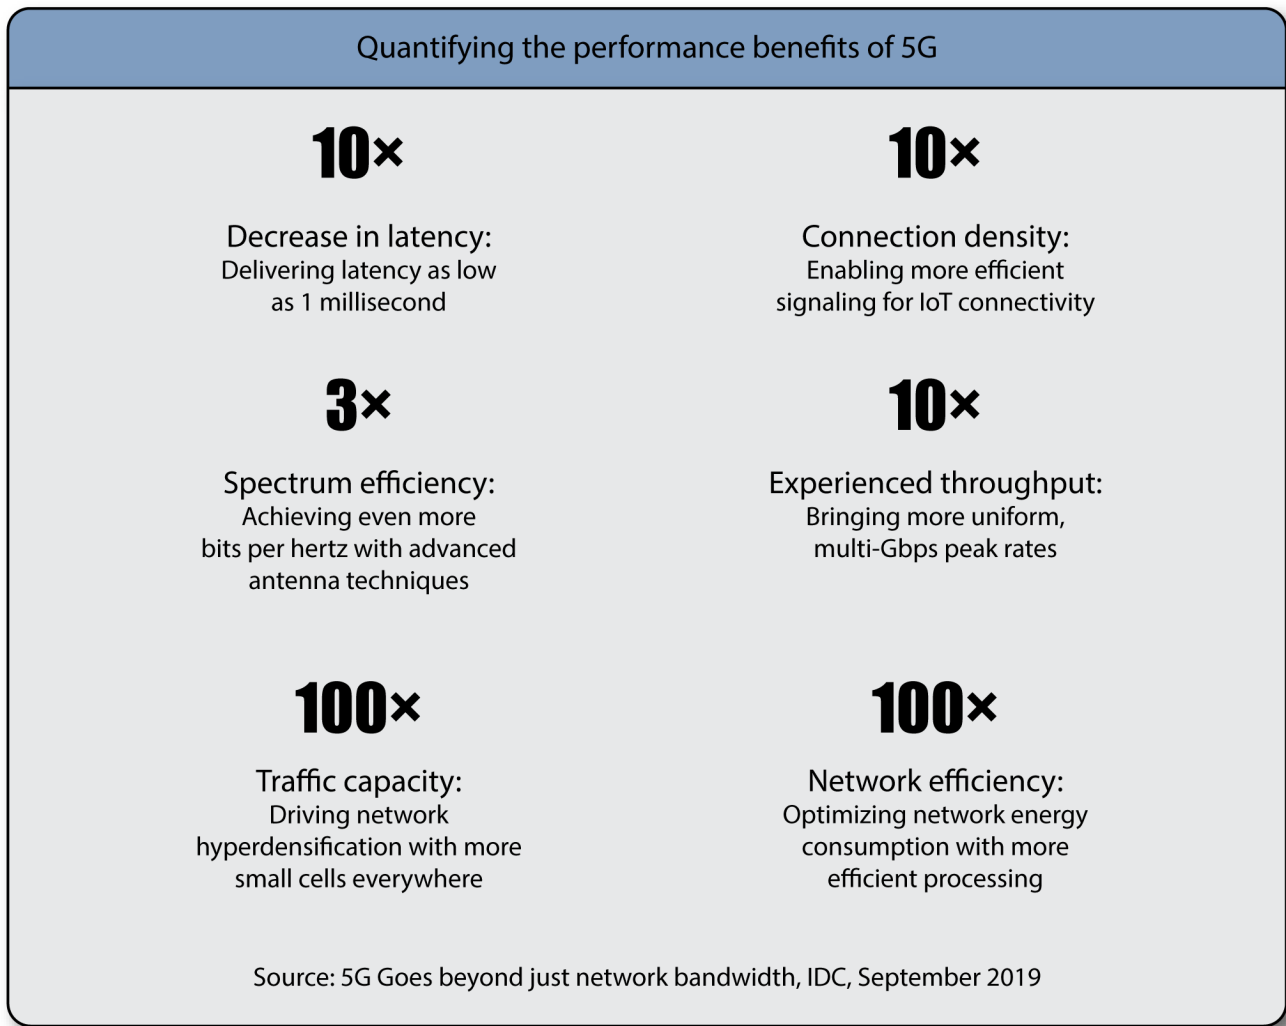


64721558-02.2020.06.30

**Figure 2: Global Mobile Device and Connection growth (Source Cisco)**

The rise of 5G and the re-architecting of data centers to better incorporate the expanding use of acceleration technologies is placing intense pressure on communications and networking designers to create systems that can process and forward terabits of data per second. These new systems not only have to be highly reliable, but also need to meet human-level response time to guarantee stringent performance guarantees (figure below), which necessitates new architectures.

Although programmable logic provides the best mixture of features to support the complex requirements of the new generation of communications and networking systems, conventional programmable silicon products cannot handle these demands. The entire FPGA architecture has to be rethought in order to balance on-chip processing, interconnect and external I/O. A state-of-the-art network on chip (NoC) and bus routing capabilities are needed to achieve the needed bandwidth and performance. An integrated NoC is the only conceivable way to architect systems supporting efficient compute, huge data throughput and deep memory hierarchy. Massive parallelism coupled with unique offloading and acceleration abilities of FPGAs ties well with the objective to deliver the highest performance per watt as well as highest performance per dollar.

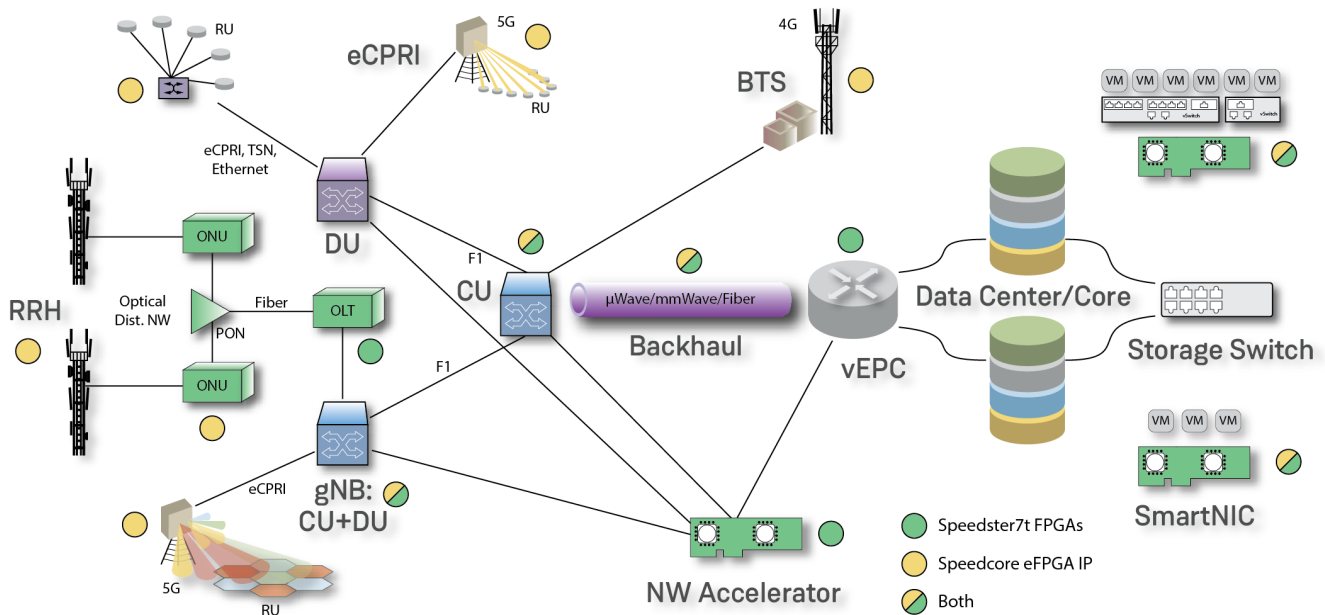


64721558-03.2020.30.04

**Figure 3: 5G Performance Metrics**

## The Changing Networking Landscape

Demand for advanced services delivered by high-bandwidth connections is reshaping the world of communications and networking. New applications in data centers, edge systems and access equipment are driving the need to deliver massive volumes of data but at the same time maintain strict latency requirements. FPGAs are becoming the centerpiece of all the network entities as shown below.



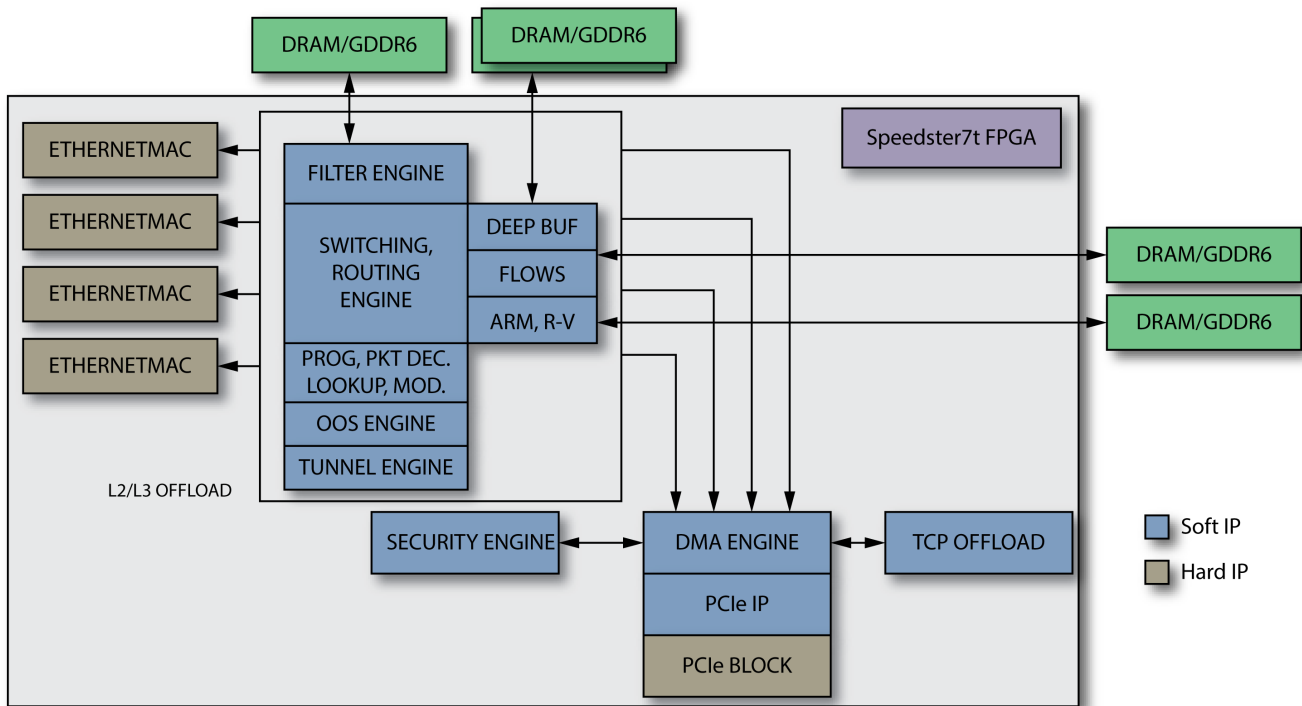
**Figure 4: FPGA in Different Network Entities**

For example, to support applications such as augmented reality and robotic control, 5G base stations and the network equipment behind them must guarantee extremely low latencies compared to prior cellular technologies. This requirement goes hand-in-hand with a demand for higher per-user throughput that leverage a number of different techniques, including multiple antennas, beamforming and the increased use of small cells as part of a network densification process. All of these factors lead to more intensive processing at centralized baseband units that are used to coordinate with multiple remote radio units over fiber links.

## The Rise of SmartNICs

Operators have embraced technologies such as software-driven networking (SDN) and network function virtualization (NFV) to improve the responsiveness of their systems. To run these services, data-center owners are adding smart network interface cards (SmartNICs) to their servers so that many of the network functions can be offloaded efficiently onto accelerators.

SmartNICs are able to deal with the bulk of the traffic entering and leaving the server, only requesting assistance from core server processors if unusual exceptions need to be handled. With sufficient acceleration features, such SmartNICs can perform a range of services at line speed. These services range from compression of data in-flight through detailed flow control to deep packet inspection applications that are able to detect anomalies and possible security breaches. As SmartNIC technology matures, increasingly advanced functions such as machine learning are now being considered to maximize the potential of flow and packet analysis. Some of the SmartNIC functions are shown in the figure below.



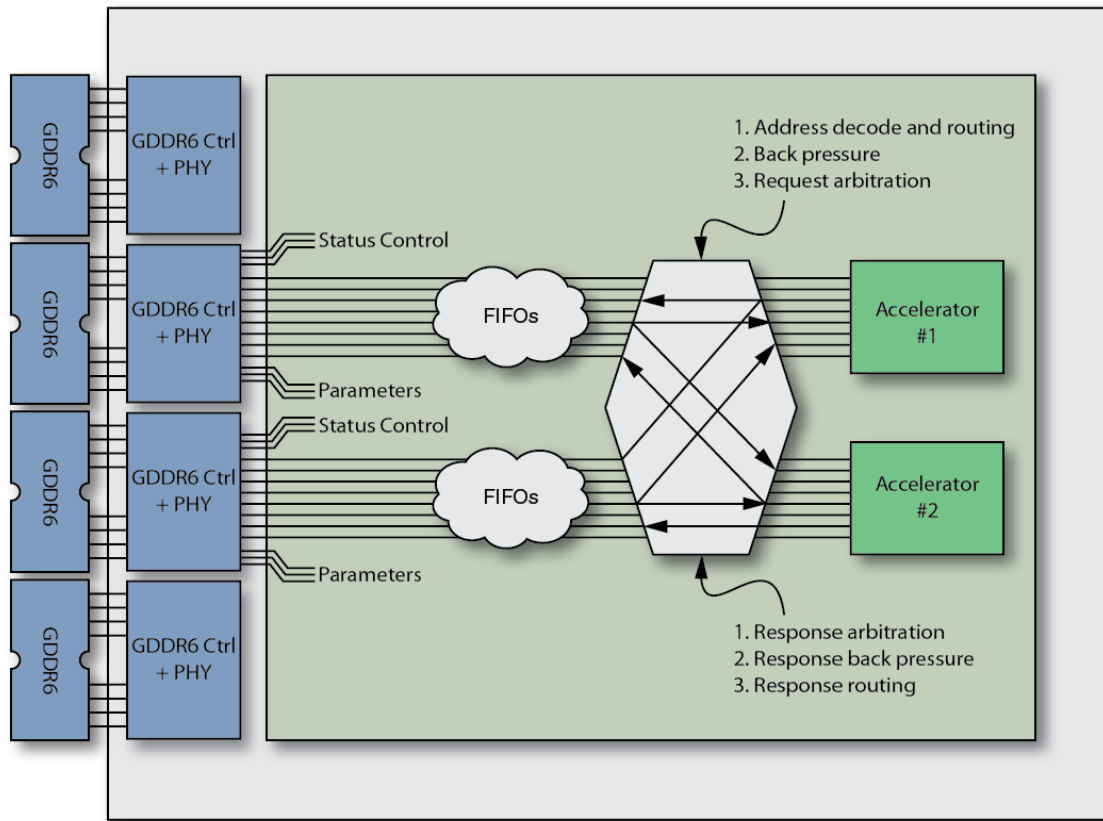
64721558-05.2020.30.04

**Figure 5: SmartNIC Block Diagram**

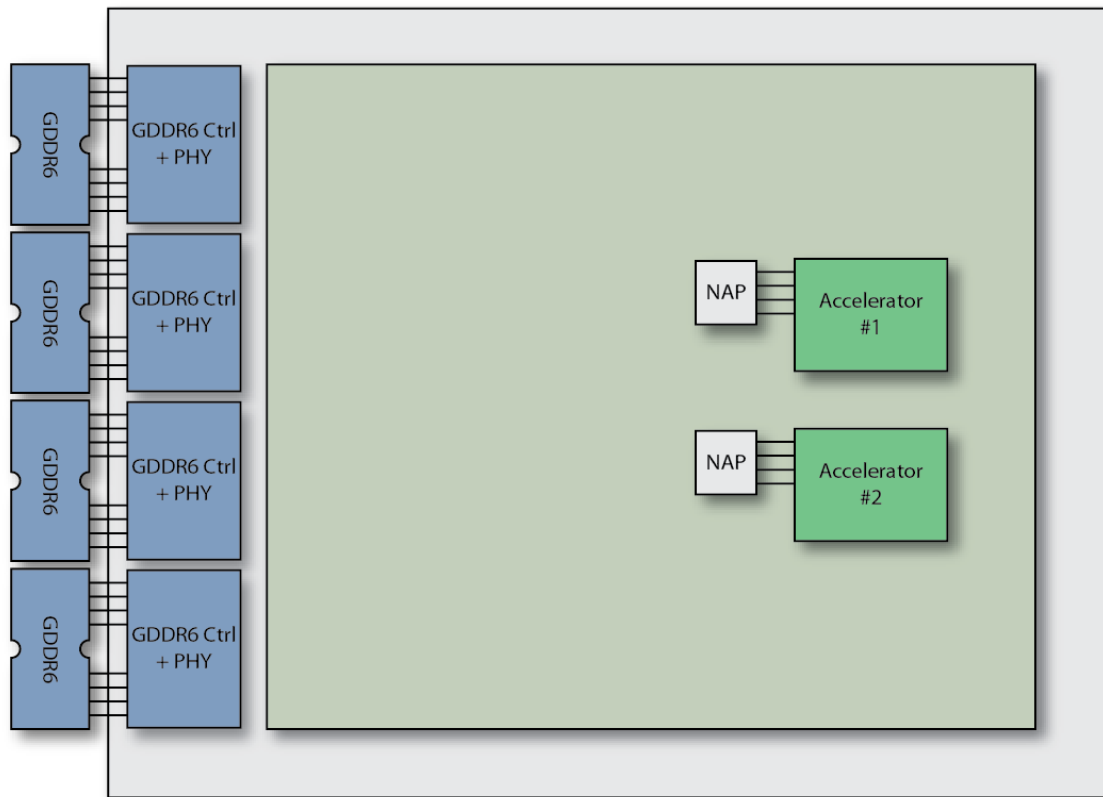
The need to convey high-speed data and respond quickly to changing conditions calls for systems that can handle a combination of high throughput and low latency. In conventional architectures, these two requirements are difficult to satisfy simultaneously. Microprocessor-based architectures now incorporate highly parallelized pipelines that are able to handle high-bandwidth data. But the need to continually move data in and out of a complex memory hierarchy makes delivering results at low latency extremely problematic. Even with dedicated offload processors, SmartNICs are being challenged by increasing data-rates and latency requirements.

## Responding to the Challenges of Designing SmartNICs

In a traditional FPGA architecture, the user needs to design the circuitry to connect the accelerators resulting in non-optimal placement and routing. The newer FPGA architectures use a network that streams data between processing elements within the logic array and the various on-chip high-speed interfaces and memory ports (figures below).

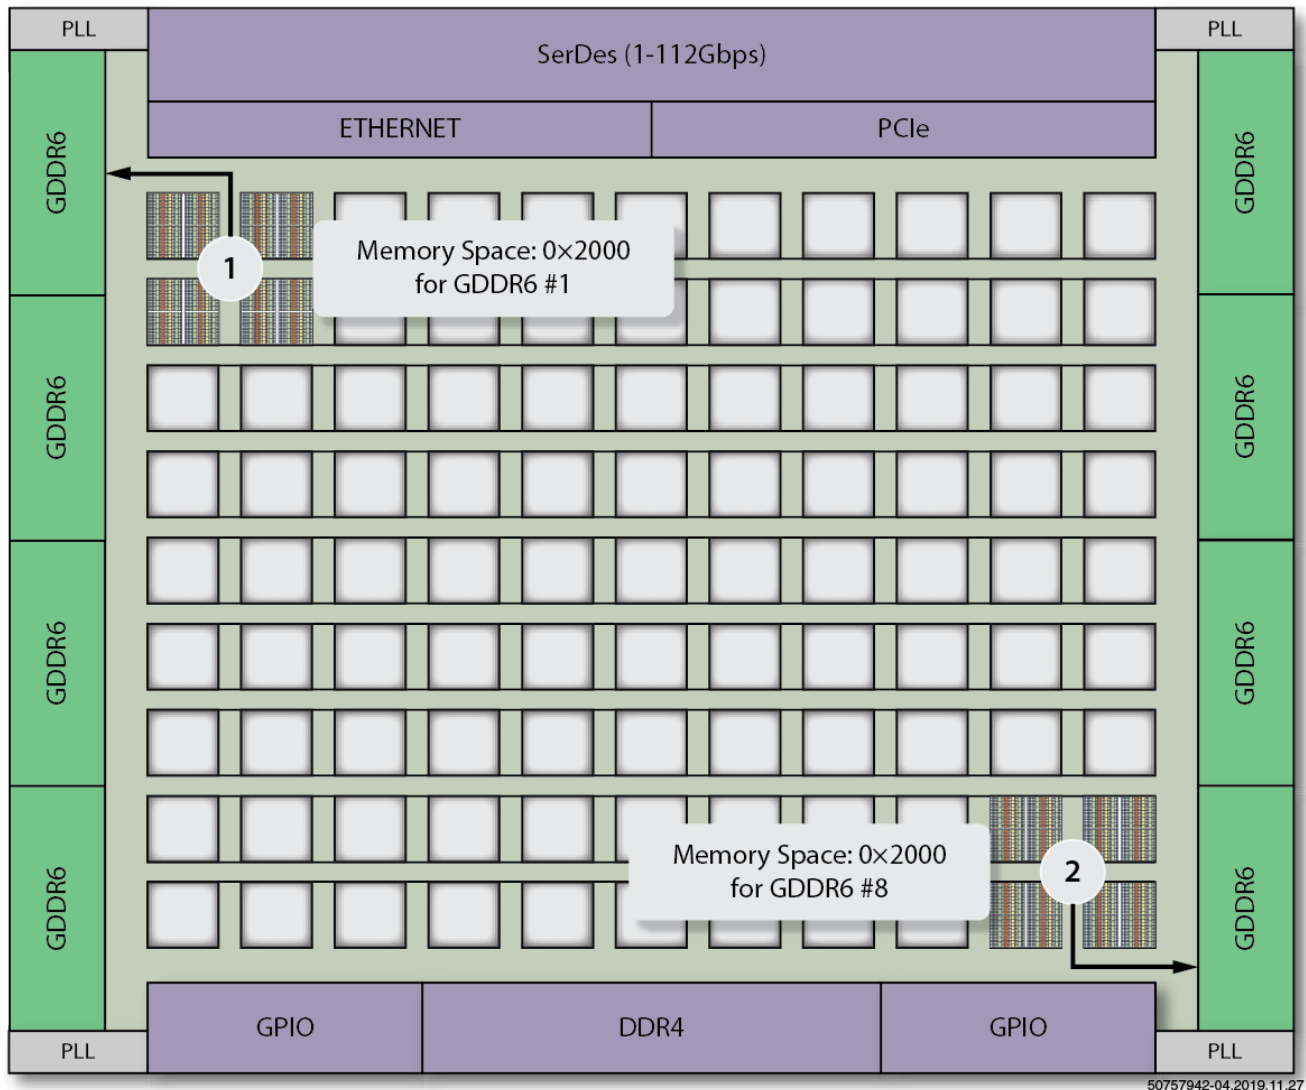


**Figure 6: Connecting Accelerators in a Traditional FPGA Architecture**



**Figure 7: Advanced FPGAs Reduce the Amount of Needed Circuitry**

Hardwired architectures greatly improve the latency and energy efficiency of processing but lack the flexibility to respond to changes in requirements. For applications such as data compression and encryption, data-center operators want to be able to embrace improvements in algorithms and react more easily to a changing threat landscape. The ability to (re)program accelerators to accommodate these changes is key a requirement. One way to achieve this reprogramming is through partial reconfiguration, taking advantage of the builtin address translation table to ease implementation (figure below).



**Figure 8: Address Translation Tables in Speedster7t Devices**

A programmable-logic architecture provides a solid foundation for implementing flexible control and dataflow structures that can deliver high throughput for many communications operations, such as packet processing. But conventional approaches in other FPGA architecture suffer from a number of limitations that make it difficult to attain the level of performance required for the next generation of 5G and data-center networking equipment.

Achronix Speedster<sup>®</sup>7t FPGAs overcomes those limitations through a balanced architecture that brings together major improvements in both compute density and data-movement capabilities.

The first device of Speedster7t FPGA, AC7t1500, delivers a range of high-speed interfaces that include fracturable Ethernet controllers (supporting rates up to 400G), PCI Gen 5 ports and up to 32 SerDes channels with speeds up to 112 Gbps. Accommodating the needs of communications systems that need to buffer bulk data at high speed and store large lookup tables, the AC7t1500 is the first FPGA to deploy multichannel GDDR6 memory interfaces. These peripherals are interconnected by a smart 2D network on chip (NoC), in addition to the bit-oriented routing structure employed by the programmable-logic fabric. As a result, the Speedster7t FPGA is the first to be able to implement terabit Ethernet (TbE) switch functionality — a critical enabling technology for data-center, networking and telecom infrastructure providers alike.



The architecture makes it possible to take networking designs even further. For example, the inclusion of matrix-oriented arithmetic units enables in-network machine learning. Using techniques such as deep learning or simpler statistical techniques, networking equipment can analyze traffic patterns to observe and enhance the flow of packets through the network and react quickly to changing conditions.

## Speedster7t Architecture — Optimized for Performance

---

The key requirement for any FPGA in communications and networking is to support the intensive I/O requirements of today's protocols. Speedster7t FPGAs accommodate this requirement with a full complement of hard I/O controllers implemented in the device's I/O ring, include 400G Ethernet, PCI Gen 5 and GDDR6 interfaces.

To avoid bottlenecks caused by the need to place some core functions into programmable logic, Speedster7t FPGAs provide complete 400 Gbps Ethernet MACs. These MACs handle forward error correction (FEC) with support for 4×100G and 8×50G options for 400G configuration. But to make full use of these features, an FPGA architecture needs more — an interconnect framework that unleashes their full performance.

Traditionally, FPGAs have used ultra-wide buses implemented using a programmable interconnect to match high-speed serial channels to the processing capabilities of the programmable logic in the core. The freely programmable nature of the interconnect matrix limits the speed at which data can be passed between logic modules. To overcome this speed penalty, FPGA users working on networking-oriented designs have frequently resorted to using extremely wide buses — often as wide as 1024 bits — synthesized from the bit-oriented interconnect matrix. For example, the bus widths needed for 400 Gbps in a conventional FPGA architecture, would require either 2048 bits running at 642 MHz or 1024 bits running at 724 MHz. Such wide buses are difficult to route because they consume a vast amounts of routing resources within the FPGA fabric. As a result, achieving timing closure for the clock rates needed to handle the incoming data is highly unlikely, even in the most advanced FPGAs.

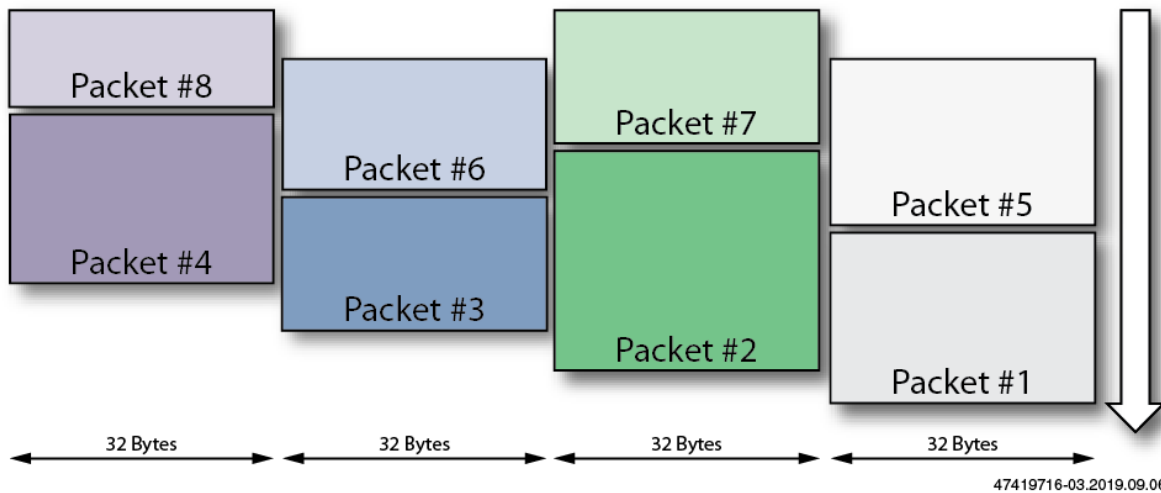
The Speedster7t architecture removes the bottleneck caused by the need to connect high-speed I/O channels directly to programmable logic which operates at much lower clock rates by providing a multi-level NoC hierarchy capable of an aggregate bandwidth of up to >20 Tbps. Not only does the NoC provide a huge upgrade in speed relative to the FPGA fabric interconnect, but the NoC is also able to move massive quantities of data without consuming any of the FPGA's programmable resources.. The internal NoC not only delivers increased bandwidth — smart connection mechanisms in Speedster7t FPGAs ease the task of getting data into the fabric from the NoC ports.

There are two main parts of NoC. The peripheral part of the NoC is responsible for data movement between the PCIe Gen 5 interfaces, memory controllers, and the core FPGA array. The other part of the NoC consists of rows and columns running over the top of the FPGA fabric. The NoC provides bidirectional, 256-bit wide horizontal and vertical channels that run between programmable clusters. Each NoC row or column can handle traffic at a rate of 512 Gbps simultaneously in opposite directions. To take maximum advantage of the infrastructure and its ability to distribute data rapidly across a Speedster7t device, the NoC is also connected directly to the on-chip 400G Ethernet controllers and uses smart traffic-distribution strategies to divide the data stream to parallelized groups of programmable-logic clusters along the NoC channel over easily implemented 256-bit wide interfaces.

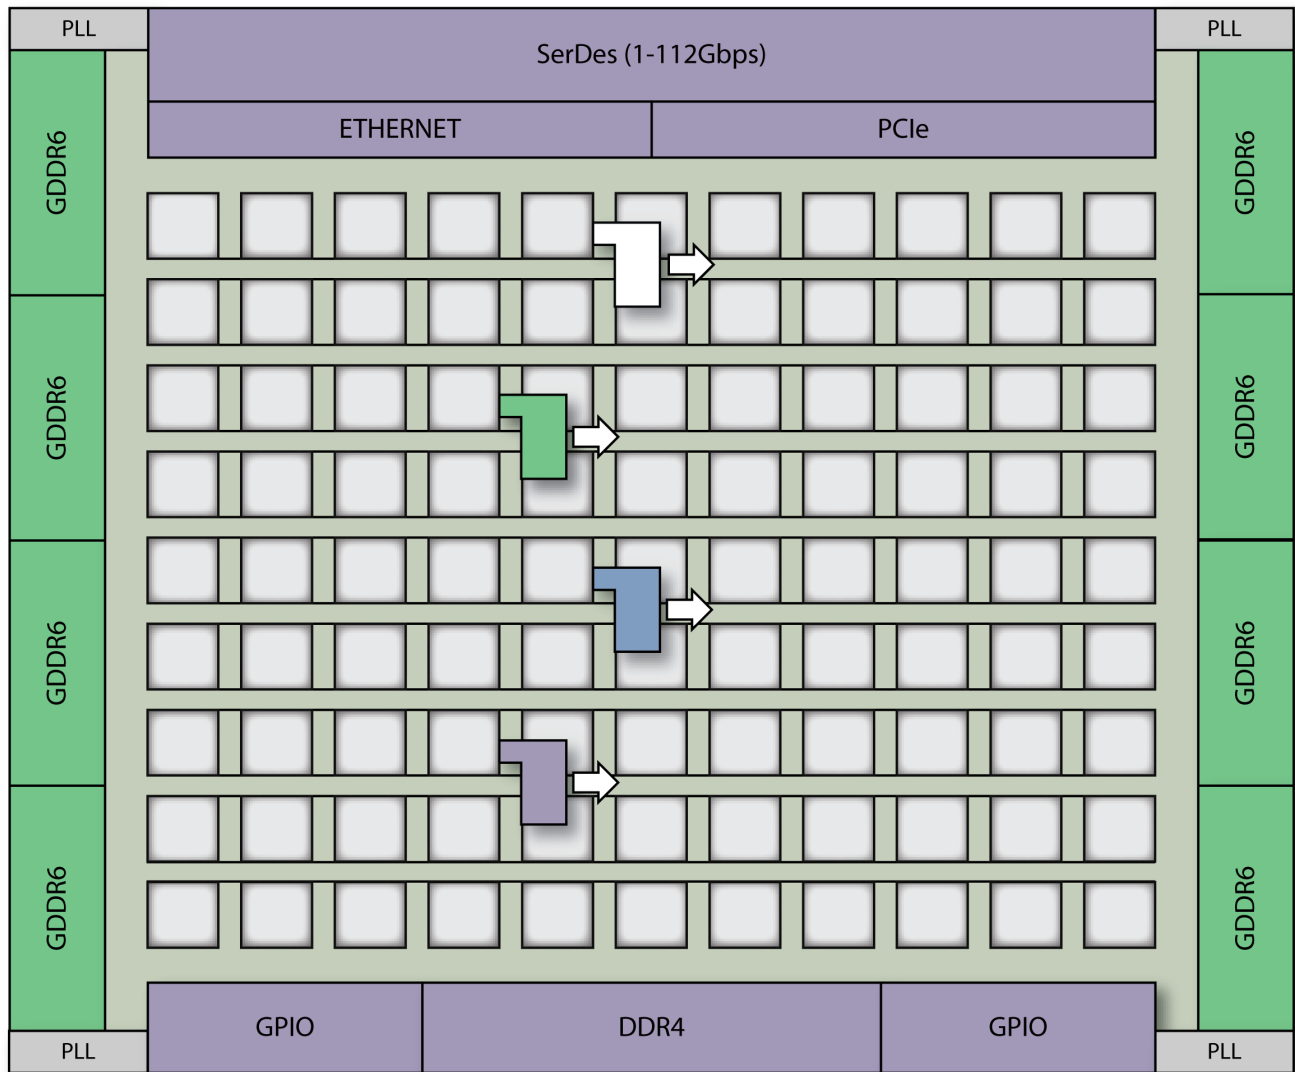
### NoC Data Modes

To achieve 400 Gbps performance, designers can use a new processing mode called packet mode where an incoming Ethernet stream is rearranged ([figure \(see page 10\)](#) below) into four narrower 32-byte packets on four independent 256-bit buses running at 506 MHz. The advantages of this mode include fewer wasted bytes when the packet ends as well as data can be streamed ([figure \(see page 11\)](#) following) in parallel without having to wait for the first packet to finish before starting the second packet transmission.

For typical networking applications operating on packetized data, each module can then classify and tag the packet headers it receives and direct payloads that need no further processing into buffer storage in external memory by calling upon the services of the NoC interfaces to off-chip GDDR6 or DDR4 memory. Once processing is complete for each packet, the necessary data is delivered to the relevant Ethernet egress ports by directing traffic from the external and internal buffers through the NoC. As a result, many operations do not need to call upon resources in the FPGA array and can take full advantage of the direct connection between the NoC and the Ethernet ports.



**Figure 9: Data Bus Rearrangement for Packet Mode**



47419716-04.2020.05.22

**Figure 10: 400 Gbps Ethernet Using Packet Mode**

The distribution of data over the NoC channels can also be done via non-packetized modes to support the widest possible range of protocols that are now being used on top of Ethernet, such as eCPRI in 5G systems, and help designers avoid the need to create ultra-wide buses in the fabric.

## High-Speed Memory Interfaces

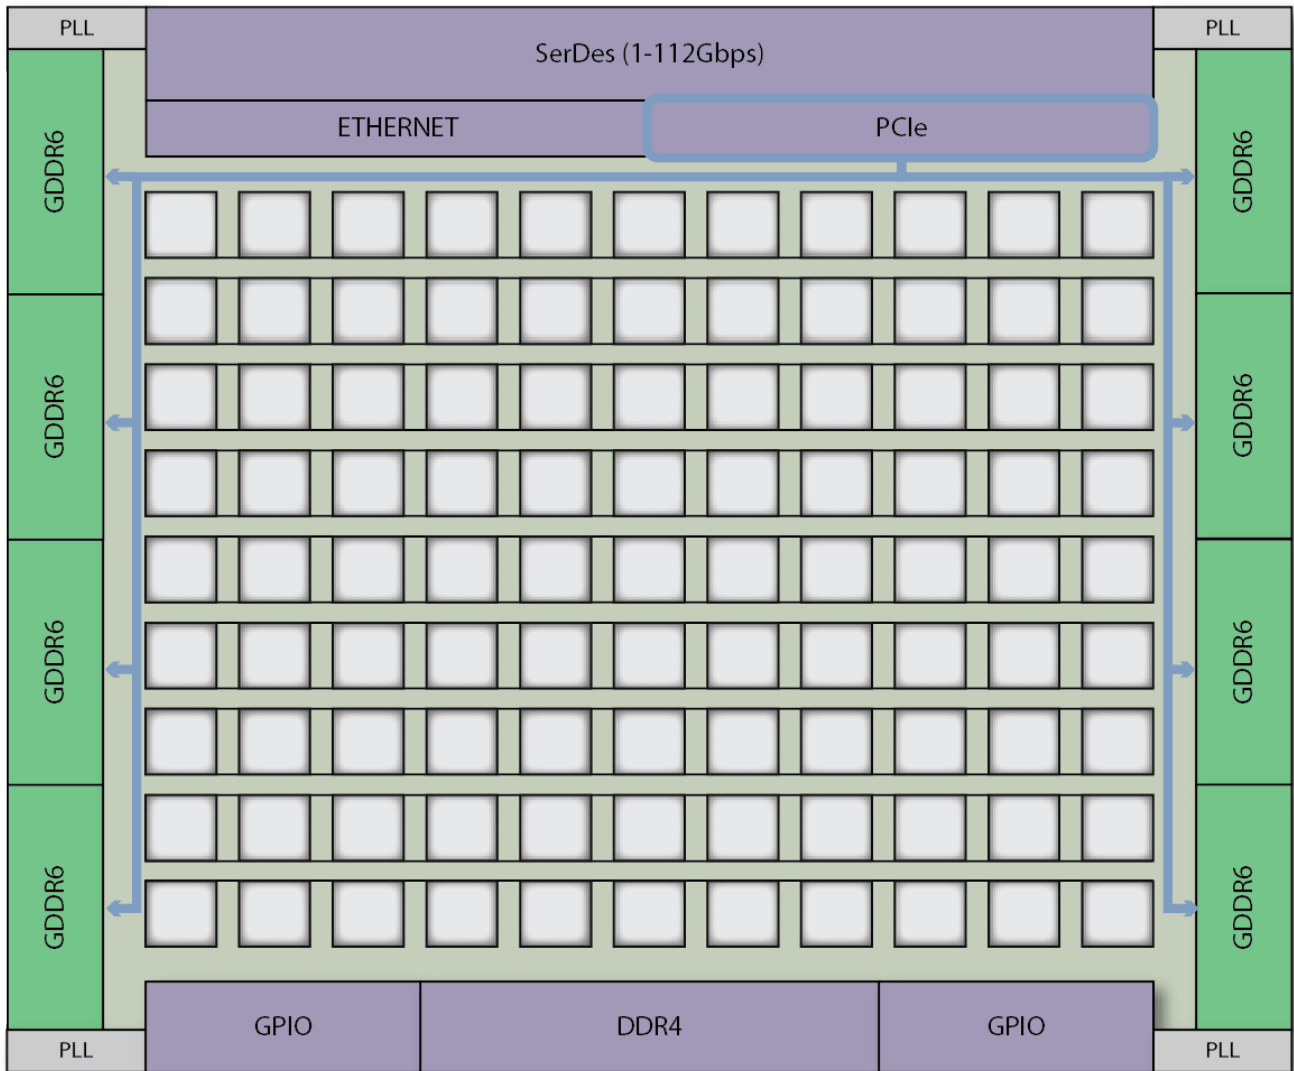
The choice of memory interfaces made by the Speedster7t architects reflects the massive capacity that the Ethernet and NoC connections provide. One possible approach would have been to design a family of products that employ the upcoming HBM2 interface. Although such an interface would deliver the level of performance needed, HBM2 is an expensive option and would force customers to wait for the necessary components and integration technology to become available.

The Speedster7t family instead employs the GDDR6 standard, which delivers the highest performance available for off-chip memories today. Speedster7t FPGAs are the first devices in the market to support this interface. Each on-chip GDDR6 memory controller can sustain 512 Gbps of bandwidth. With up to eight GDDR6 controllers in a single AC7t1500 device, aggregate memory bandwidth can total 4 Tbps.

## PCIe Gen 5 Support

Alongside the Ethernet and memory controllers, PCIe Gen 5 support on Speedster7t FPGAs allow tight integration with a host processor to support high-performance accelerator applications, such as sidecar SmartNIC designs. The PCI Gen 5 controller makes it possible to read and write data stored within the memory hierarchy of the FPGA, including the many block RAMs located within the fabric as well as external GDDR6 and DDR4 SRAMs devices attached to the FPGA's memory controllers. Data movement controllers, such as DMA engines, instantiated in the FPGA array can similarly access memory shared with the host processor over the PCIe Gen 5 bus. This high-bandwidth connection is achieved without consuming any FPGA fabric resources and nearly zero design time. The user only needs to enable PCIe and GDDR6 interfaces in order to send transactions via the NoC.

Direct connection between the PCIe subsystem and any of the GDDR6 or DDR4 memory interfaces is shown in the figure below.



50757942-03.2019.12.02

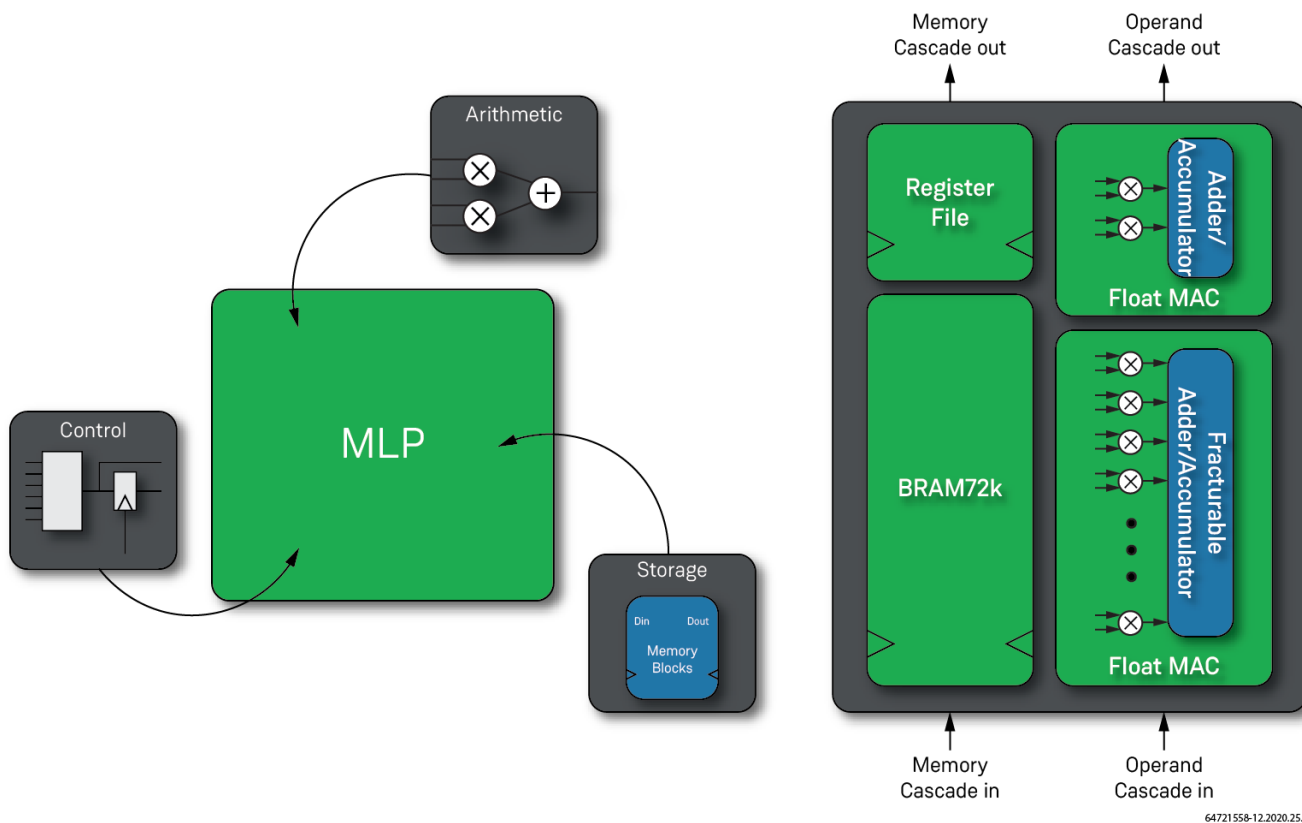
**Figure 11: Data Transfer Between PCIe and GDDR6 Without Fabric Intervention**

## 112-Gbps SerDes

Used by the 400G Ethernet channels for physical-layer access, the AC7t1500 provides up to 32 high-speed SerDes channels that can be employed for other standards that need data rates up to 112 Gbps, with full support for PAM4 signaling. These SerDes channels allow the implementation of extra-short reach (XSR) and ultrashort reach (USR) inter-device channels that are proving to be important for a range of communication systems. The flexibility of the SerDes implementation together with support for a variety of Ethernet speeds thanks to the inclusion of a fracturable controller provides ready support for designs that will work with any of the proposed CPRI and eCPRI formats (used in 5G front-haul design).

## Machine Learning Processor

For computationally intensive tasks, the Speedster7t machine learning processors (MLP), deployed across a Speedster7t FPGA, are flexible and factorable arithmetic units. MLPs are high-density multiplier arrays with floating-point and integer MAC blocks supporting multiple number formats. MLPs have integrated memory blocks that can perform operand and memory cascade functions without using FPGA resources. MLPs are suited for a range of matrix-math operations, ranging from beamforming calculations for 5G radio controllers to the acceleration of deep-learning applications, such as traffic pattern and packet content analysis.



**Figure 12: Machine Learning Processor Block Diagram**

## Conclusion

---

Communications and networking systems that extend from the edges of the 5G network to the switches inside data centers are placing extreme pressure on the ability of silicon to support the computational and data-transfer rates they require. Traditionally programmable logic provided the best mixture of flexibility and speed for these systems, but has been challenged in recent years by the increase in speed of protocols such as Ethernet to 100G and 400G. Through the inclusion of an innovative, multilevel network on chip that allows data to be streamed easily around the device without impacting the FPGA fabric, the Speedster7t architecture ensures all the included world-class I/O interfaces, such as 400G Ethernet, GDDR6, and PCI Gen 5, and the core programmable-logic fabric features are used to their full potential. Using an innovative architecture that leverages NoC technology and takes full advantage of 7nm technology to deploy the highest-performance controllers available, the Achronix Speedster7t architecture provides the elements that FPGAs have missed up to now. Designs targeting Speedster7t FPGAs can accept massive amounts of data from multiple high-speed sources, distribute that data to programmable on-chip algorithmic and processing units, and then retrieve those results with the lowest possible latency. The result is an innovative FPGA architecture that can support the next generation of 5G, software-defined networking and data-center systems that are now being designed. Speedster7t FPGAs can now drive communications and networking applications forward to the next generation.

# Achronix<sup>®</sup>

## Data Acceleration

Achronix Semiconductor Corporation

2903 Bunker Hill Lane  
Santa Clara, CA 95054  
USA

Website: [www.achronix.com](http://www.achronix.com)  
E-mail : [info@achronix.com](mailto:info@achronix.com)

---

Copyright © 2020 Achronix Semiconductor Corporation. All rights reserved. Achronix, Speedcore, Speedster, and ACE are trademarks of Achronix Semiconductor Corporation in the U.S. and/or other countries All other trademarks are the property of their respective owners. All specifications subject to change without notice.

### Notice of Disclaimer

The information given in this document is believed to be accurate and reliable. However, Achronix Semiconductor Corporation does not give any representations or warranties as to the completeness or accuracy of such information and shall have no liability for the use of the information contained herein. Achronix Semiconductor Corporation reserves the right to make changes to this document and the information contained herein at any time and without notice. All Achronix trademarks, registered trademarks, disclaimers and patents are listed at <http://www.achronix.com/legal>.