White Paper

June 2020

PCIe® – What can it do for me?

By Kay Annamalai, Senior Marketing Director, Diodes Incorporated

The rate at which we now generate raw data is immense but, perhaps, not that staggering, at least not to people working in an engineering environment. There are predictions that by 2021 there will be more bytes of data floating around our networks, sitting in our storage devices and moving through the air than there are stars in our observable universe. The unit of measure being used is zettabytes, which is one Kbyte raised to the power of 7. A Gbyte is a Kbyte raised to the power of just 3, by comparison.

A great deal of this data will be coming from 'things' and, for the purposes of analysis, we can include people in that. The IoT does not discriminate between mechanical, organic or inanimate 'things'; data lives everywhere but it does not remain static. The purpose of data is to convey information and it is up to the receiver to understand the meaning of that information. Pulling actionable insights out of raw information is both compute- and data-intensive, and because of this, the process of moving data around systems has had to keep pace with the increased rate at which data is created.

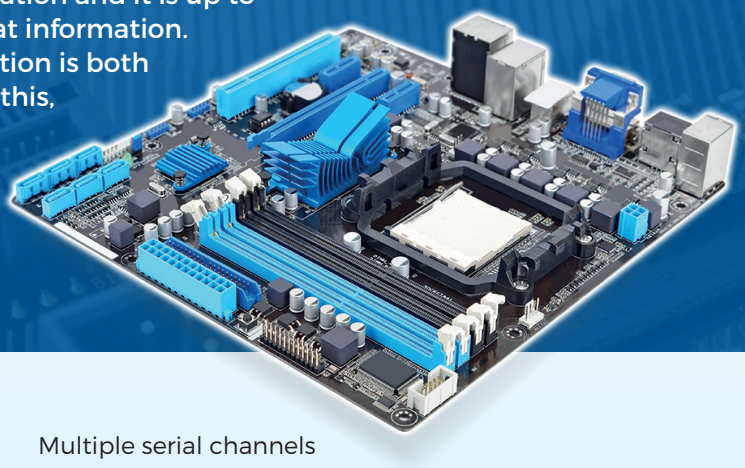

Data buses have oscillated somewhat between parallel and serial topologies in recent history. Logic might indicate that a wide parallel bus should be able to transfer more data per second than a two-wire serial bus, and for a while this was the case. However, signal integrity, power consumption and physical space have conspired to direct the industry towards serial bus topologies. More pertinently, parallel buses are invariably shared between multiple devices, which is inherently limiting. A serial, point-to-point bus overcomes this limitation and allows two devices to communicate over a dedicated channel, with less channel management overhead.

The potential for scaling a serial bus through parallelism isn't ignored, however.

Multiple serial channels can be joined together to create wider connections and further increase bandwidth. This succinctly describes the evolution of the Peripheral Component Interconnect, or PCITM bus from a parallel topology to PCI Express® (PCIe®), which is serial. As demand for higher transfer rates has increased, so too has the speed of PCI Express, such that it is now on its fifth iteration.

Backwards-compatibility has been a major factor in this evolution, so while PCIe 5.0 is still largely nascent it will support migration from PCIe 4.0, which is now used in a host of application areas and will continue to be the right solution for many long after PCIe 5.0 reaches its zenith.

FIGURE 1. THE EVOLUTION OF PCIE

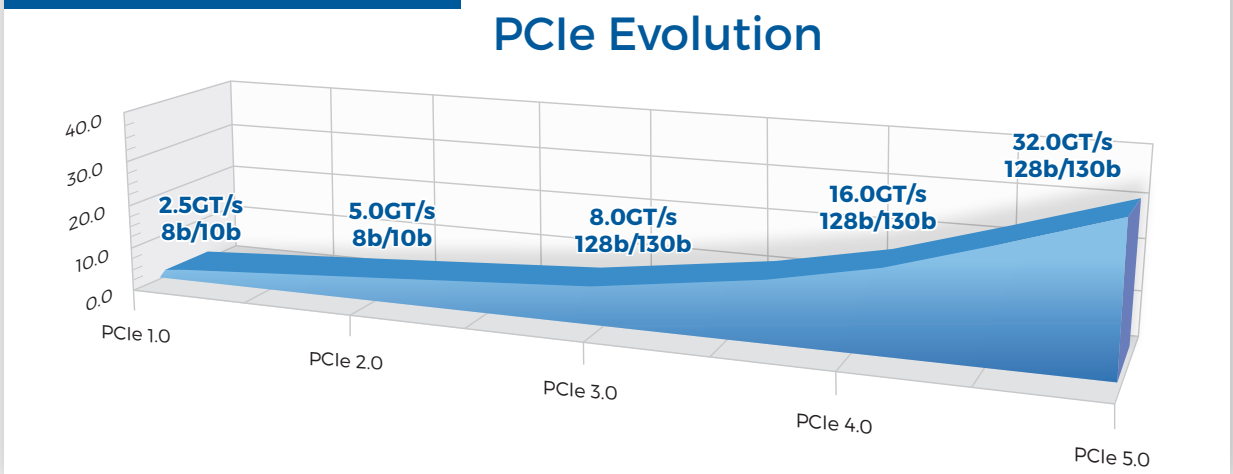

Meeting data center demands

The IoT, 5G communications and the applications these are enabling, such as smart cities and autonomous vehicles, all rely on data, the vast majority of which will pass into, around and out of a data center during its working life.

The real-time nature of these applications will drive up demand for technologies that can deliver the speeds needed to support that. This includes higher speed memory devices, processors and interfaces. Today, data centers are using 100Gbit Ethernet connections, but this is likely to increase to 400Gbit Ethernet in the near term.

All other parts of the system will be expected to provide a commensurate increase in performance.

The rate at which the bandwidth of PCIe and Ethernet has increased is not coincidental; the two technologies are complementary, particularly in data centers. At times, PCIe has been the limiting factor and, at these points in time, other technologies have emerged. These include CCIX, Gen-Z and OpenCAPI. These can be broadly compared to PCIe in that they are also point-to-point protocols but they also include cache coherency, which PCIe does not.

FIGURE 2. ALTERNATIVE AND COMPLEMENTARY TECHNOLOGIES TO PCIE

Standard	Physical Layer	Topology	Unidirectional Bandwidth	Mechanicals	Coherence
CCIX	PCIe PHY	p2p and switched	32-50 GB/S x16	PCIe	Full cache coherency between processors and accelerators
GenZ	IEEE 802.3 Short and Long Haul PHY	p2p and switched	Signalling Rates: 16, 25, 28, 56 GTs Multiple link widths: 1 to 256 lanes	Supports existing PCIe mechanicals/form factors Will develop new, Gen-Z specific mechanicals/form factors	Does not specify cache coherent agent operations, but does specify protocols that support cache coherent agents
OpenCAPI 3.0	BlueLink 25Gbs PHY Used for OpenCAPI and NVLINK	p2p	25 GB/s x8	In definition, see Zaius design for a possible approach	Coherent access to memory Cache coherence not supported until v4.0

High-speed interfaces

It is not difficult to appreciate how faster interfaces promote higher workloads but it is worth exploring how the changing shape of those workloads is sustaining demand for more speed.

If we consider the enterprise to be the origin of PCIe then networking is the first branch on its lineage. Network processors continue to form the foundation of networks and are constantly being redeveloped to deliver higher performance. This invariably means faster processing cores, but it also extends to hardcoded accelerators for specific functions, such as firewalls. In these cases, it is common to integrate multiple PCIe lanes to keep the hardware fed with data.

At the enterprise level, accesses to and from storage are another potential performance limitation. Here, the complementary technology known as NVMe (non-volatile memory express) provides system processors and other ICs with high-speed access to solid-state memory through the PCIe bus. It offers an interoperable way of accessing solid-state devices from various vendors, while being a practical alternative to the SATA and SAS interfaces, both of which were developed when moving media were still the norm. Conversely, NVMe was developed after the dawn of solid-state technology and is therefore more suited to faster storage media. Hardware acceleration and wide data pipes are also common requirements in HPC (high performance computing) and cloud computing.

Other emerging applications where PCIe is being implemented include artificial intelligence (AI). Here the need for high volumes of data is most apparent on the algorithm training side where SoCs are being developed to deliver massive increases in throughput. These high-performance processors, naturally, require high bandwidth data interfaces, which is where PCIe plays its part. Streaming media is also becoming prevalent and relies heavily on bandwidth. Augmented and virtual reality headsets are examples of where a compact, high-speed interface is essential.

Perhaps the biggest market for PCIe outside of the enterprise is in automotive. Autonomy, advanced driver assistance systems (ADAS), and conventional infotainment equipment all depend on high volumes of time-sensitive data, which brings another dimension to faster interfaces. PCIe has proven itself capable of meeting the demands coming from new applications, even while it continues to evolve. This will ensure its future, as covered in the next section.

The evolution of PCIe

Since its inception in 2003, the evolution of PCIe has been measured largely in terms of the throughput each generation can deliver.

PCIe 1.0 offered 2.5 Giga Transfers per second (GT/s) and in 2006 this doubled to 5GT/s with PCIe 2.0. In 2010, PCIe 3.0 was introduced, which reached 8GT/s.

FIGURE 3. AGGREGATE SPEEDS OFFERED BY PCIE VERSIONS

PCIe Generations	Raw Bit Rate	Interconnect Bandwidth	Bandwidth Lane Direction	Total Bandwidth for x16 link
PCIe 1.0	2.5 GT/s	2 Gb/s	~250 MB/s	~8 GB/s
PCIe 2.0	5.0 GT/s	4 Gb/s	~500 MB/s	~16 GB/s
PCIe 3.0	8.0 GT/s	8 Gb/s	~1 GB/s	~32 GB/s
PCIe 4.0	16.0 GT/s	16 Gb/s	~2 GB/s	~64 GB/s
PCIe 5.0	32.0 GT/s	32 Gb/s	~4 GB/s	~128 GB/s

The evolution of PCIe (Cont)

Although this was followed by a seven-year break, during which other technologies (see figure 2) started to be developed, with 2017 came PCIe 4.0, which saw throughput double again to 16GT/s, an increase that has been hugely influential in the embedded space.

The popularity of PCIe 4.0 is likely to continue for some time, thanks to its strong offering, but progress knows no bounds and in 2019 the PCI-SIG® outlined its intention to release PCIe 5.0. When silicon becomes widely available it will offer throughputs of 32GT/s, re-doubling the performance of PCIe 4.0. In terms of bandwidth, PCIe 5.0 will support bidirectional data traffic at up to 128Gbyte/s in a 16-lane system.

Every new generation of the PCIe specification has introduced new opportunities but, with these, come increased design challenges. These challenges include the signal integrity issues that arise as clock and data frequencies increase. This issue needs to be managed at every point in the signal chain, including the physical, data and transaction layers. Other design considerations exist around jitter and crosstalk, as well as inter-symbol interference. The PCI-SIG has over a hundred recommended tests that should be applied when developing a PCIe interface.

Future-proofing your design with PCIe

One of the great strengths and a fundamental part of the PCIe standard is its backward-compatibility. It means that systems based on PCIe will benefit from higher transfer speeds as new silicon becomes available, with minimum design changes.

However, while a high-speed interface is becoming necessary or beneficial in more systems, PCIe was never intended to operate in isolation.

Bridging to and from PCIe, between legacy or alternate interfaces, is a crucial part of the product offering.

This is where technologies and solutions such as clock generators and clock buffers, controllers, packet switches/bridges, ReDriver™ chips, and high-speed multiplexers are essential. Diodes Incorporated has product offerings for all of these functions, some of which are outlined below.

FIGURE 4. AN EXAMPLE OF THE PCIE SOLUTIONS OFFERED BY DIODES

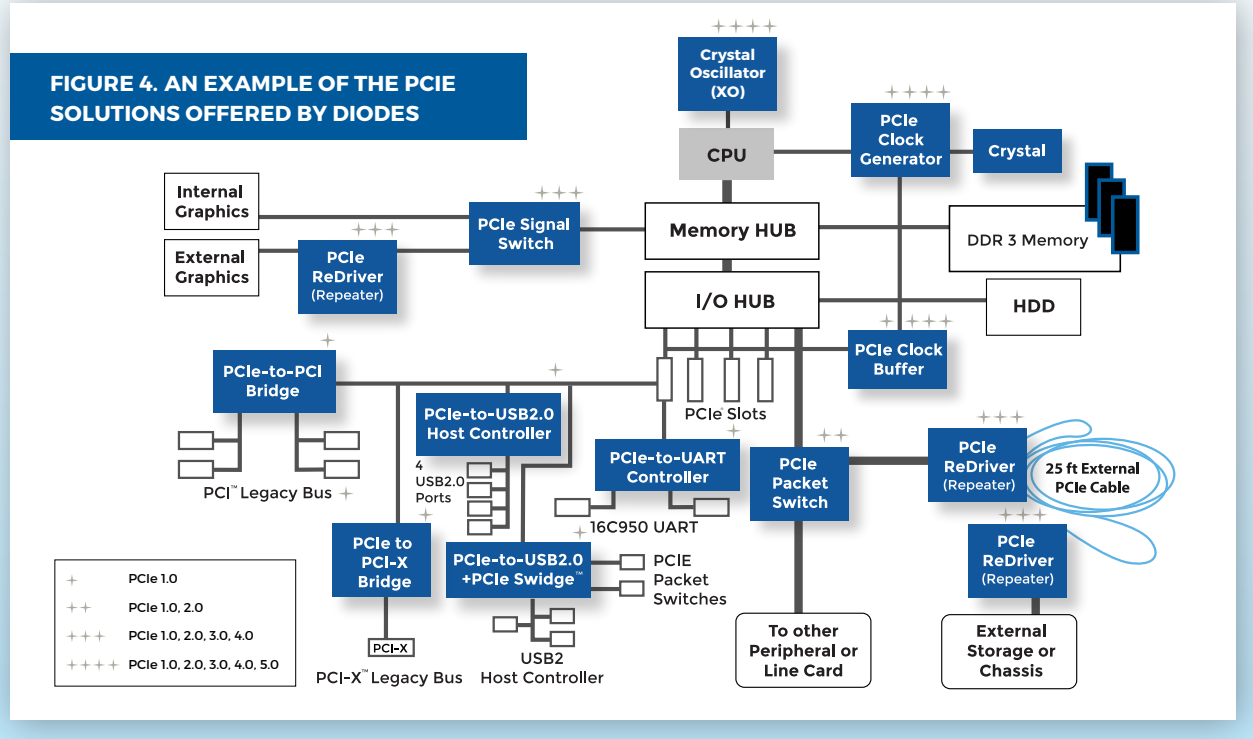

Future-proofing your design with PCIe (Cont)

In some instances, instead of ReDrivers, retimers might be used where the loss in the transmission line cannot be handled with a ReDriver. A ReDriver is effectively an equalizer compensating for the loss of the transmission line and an output driver implemented with very low latency compared to a retimer. Both ReDriver and retimer devices allow a signal to travel further on a PCB track without attenuation but the ReDriver is generally easier to implement and more cost-effective.

A packet bridge will typically provide an interface between two layers in the OSI reference model, or between two protocols. For example, a bridge may be used to connect PCIe to USB, UART or even other versions of the PCI bus (such as PCI-X). Packet switches are multiport/multilane devices. They are used where there is a single root complex that needs to be expanded to several ports with multiple lanes for accessing other peer systems. Diodes offers packet bridges and switches with various port configurations and translation capabilities.

Passive signal multiplexers are used for routing signals over the same system for bandwidth expansion for graphics or computation. They can also be used for enabling connections from a single multi-protocol interface.

Clock buffers can typically take a single reference signal as an input and produce multiple outputs with very low additive jitter for wider distribution around a PCB. Diodes offers both 1.8V and 3.3V PCIe buffer devices, up to 20 outputs. All these buffer devices offer low power HCSL outputs with on-chip termination, which removes the need for two additional resistors per pair, along with output enable control for each clock output. These can also be used in collaboration with clock generators, or synthesizers. Clock generators are able to generate a clock signal at a specific frequency with very low output jitter, making them suitable for PCIe as well as other system clocks.

The PCIe continuum

The timeline for PCIe is undeniably variable; the relatively long delay between PCIe 3.0 and PCIe 4.0, followed by the very short time between PCIe 4.0 and PCIe 5.0 means that the latest version of the specification will face some competition from its own predecessor. That may feel like a counterintuitive move by the PCI-SIG, but in reality there is likely to be significant crossover between them.

In fact, there is no expiry date on any of the PCIe versions; all five are able to coexist in the same industries, even in the same applications.

By understanding this it becomes clear that, unlike other technologies that evolve and change, PCIe will be a valuable technology in all its iterations for many years to come. For this reason, it is important that semiconductor manufacturers continue to invest in the technology, both by developing devices that meet the latest version of the specification and, equally, by supporting and improving product families that address all previous (yet current) versions.