PCIe® Data Transmission Overview

An Introduction to Timing Applications

MICROCHIP

microchip.com

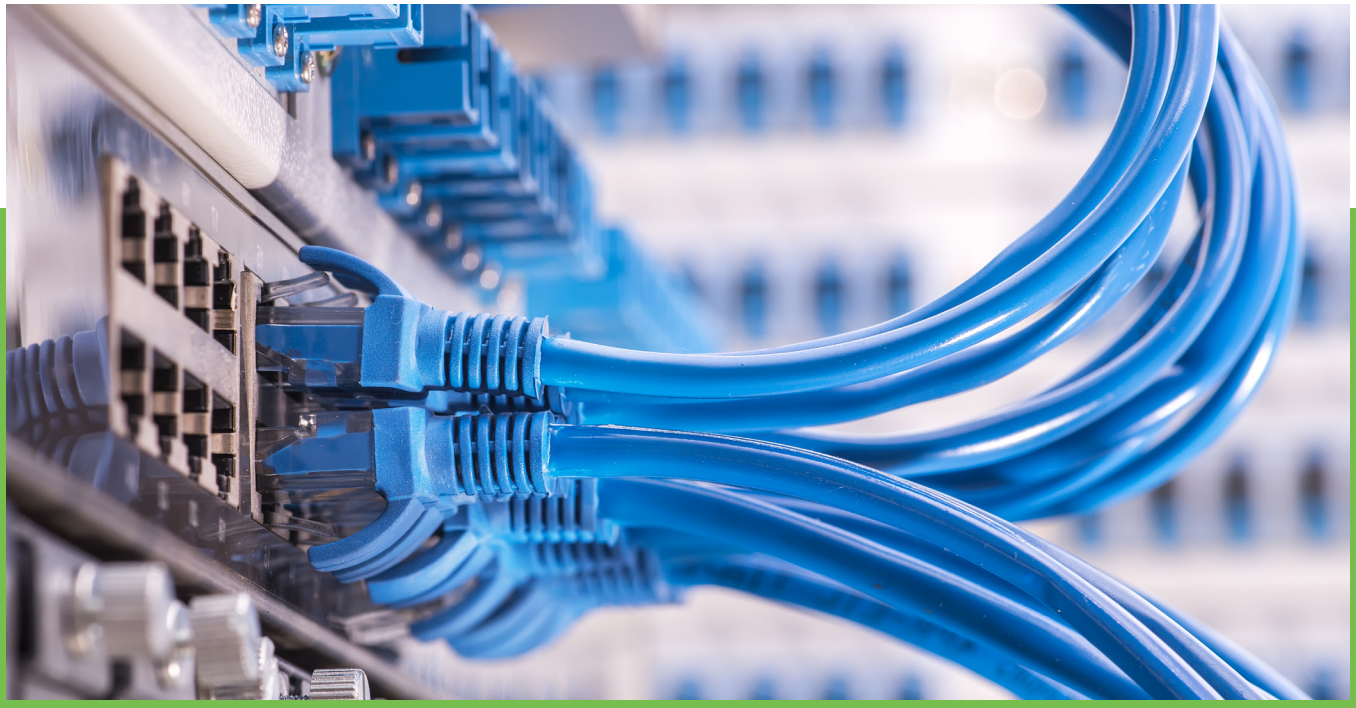

The PCI Express® (PCIe) bus standard has gained broad acceptance across many markets in a large diversity of applications. These documents provide a general overview of PCIe data communication and then focuses on the specific issues related to timing (clocks) in PCIe systems. Other topics include Bus architecture and spread-spectrum clocking as well as the unique requirements of automotive applications and data centers.

Chapter 1

PCIe: Data and Clocks Overview

This chapter provides a high-level overview of the PCIe bus standard.

PCIe is a high-speed serial bus standard that features:

I. Point-to-point communication

II. Bi-directional communication

III. Scalable architecture to accommodate varying bandwidth needs

IV. Backwards-compatibility with previous versions of the PCIe standard

V. Wide adoption across many markets

1. Point-to-Point Communication

Point-to-point communication means that there is one participant at each endpoint, such as a telephone call. This contrasts with a multi-drop bus, where multiple components are connected to a bus and signaling protocols determine which component is transmitting data and which is receiving data.

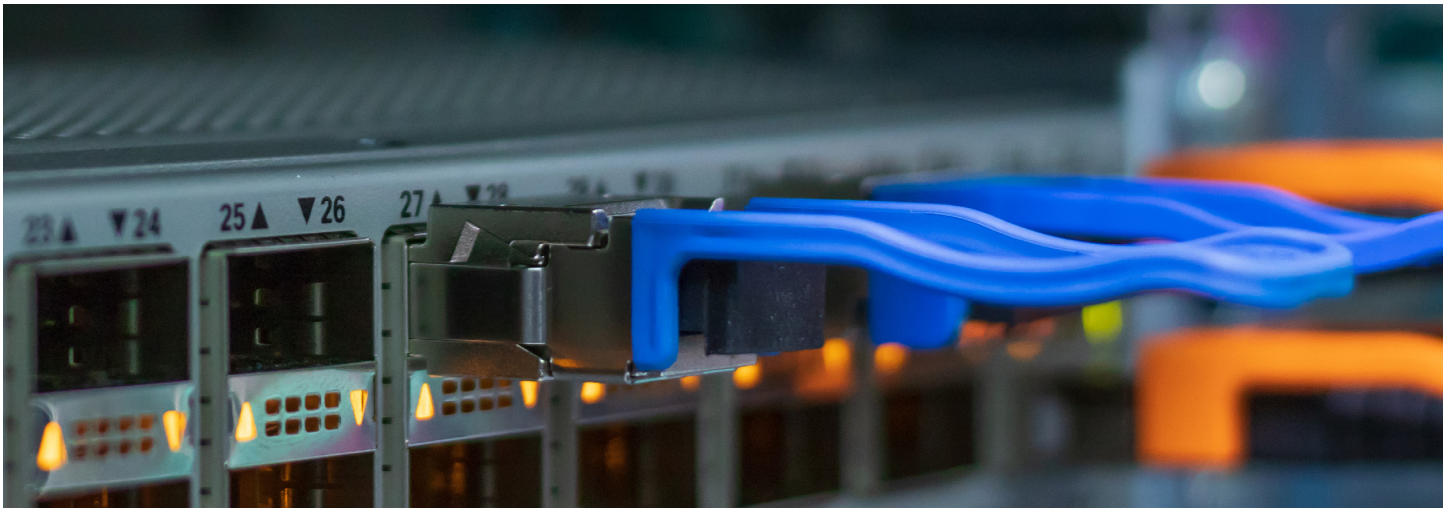

2. Bi-Directional Communication

Bi-directional communication means that each participant can both transmit and receive data. PCIe uses dual simplex as the implementation for this simultaneous two-way communication.

Simplex = One-way communication on the physical connection

Dual = Two

Therefore, dual simplex is two physical one-way connections, to/from each endpoint. Each dual simplex connection is called a lane.

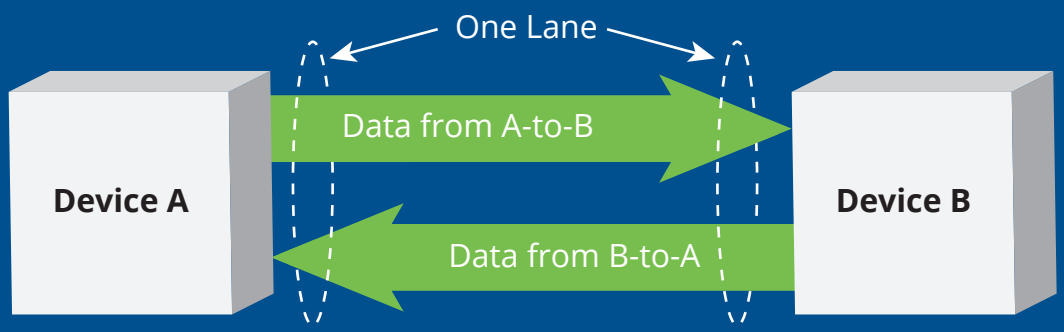

3. Scalability Depending on Bandwidth Needs

Just as a multi-lane road can carry more vehicle traffic than a one-lane road, a data bus that contains multiple copies of the same lane can carry more data traffic than only one lane. It is usually physically impractical to have more than 16 lanes connecting two endpoints, although the current PCIe 6.0 standard allows for up to 32 lanes. The connection between PCIe endpoints is called a link, no matter how many lanes it contains.

As PCIe evolved, the transfer speeds of each lane have increased, thereby increasing the data throughput of each link.

PCIe® Version (Year Introduced)	Per-Lane One-Way bit Rate [Gbit/sec] [GT/sec]*	Per-Lane One-Way Data Rate ~ a/10 (PCIe 1.0 and 2.0) ~ a/8 (PCIe 3.0+) [GByte/sec]	Per-Link (16-Lane) Maximum One-Way Data Rate ~ b × 16 [GByte/sec]
1.0 (2003)	2.5	~ 0.25	~ 4
2.0 (2007)	5	~ 0.5	~ 8
3.0 (2010)	8	~ 1	~ 16
4.0 (2017)	16	~ 2	~ 32
5.0 (2019)	32	~ 4	~ 64
6.0 (2021)	64	~ 8	~ 128

* Gigabit/sec (Gbit/sec) = Giga Transfer/sec (GT/sec)

4. Backwards Compatibility

Each PCIe version is backwards compatible to previous versions. When endpoints contain devices that are from different PCIe versions, the lanes automatically function at the slower data rate of the earlier PCIe version.

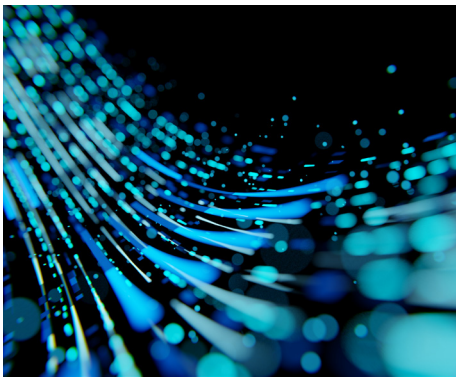

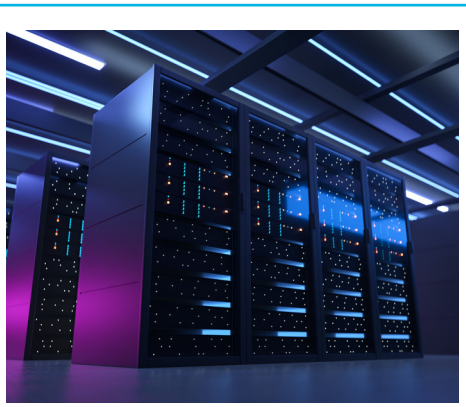

5. Wide Adoption Across Many Markets

Although the PCIe bus was originally focused primarily on the communication between a Central Processing Unit (CPU) and cards/devices on a personal computer's motherboard, its applications have expanded far beyond the PC market. These applications include:

- Data centers
- Storage systems
- Servers
- Vehicle-based systems (telematics and ADAS)
- Communication networks and infrastructure (such as Ethernet-based systems and switches)
- Home and office external peripheral connections (such as the PCIe data link that is part of the Thunderbolt™ hardware interface used to connect external peripherals to a computer)
- Instrumentation
- Consumer devices

Chapter 2

PCIe: Clock/Data Architecture and Requirements

All PCIe devices run off of a clock, and this clock needs to adhere to certain performance specifications. The common reference clock architecture, also known as the Common Clock (CC) architecture, is the most commonly implemented method for clock distribution among PCIe devices. This document focuses on the CC architecture.

Both devices connected by a PCIe data link share the same reference clock (Refclk). In other words, they have their reference clock in common. For PCIe 1.0 through PCIe 4.0, this clock is specified to be ±300 ppm; for PCIe 5.0 and PCIe 6.0 it is ±100 ppm.

Let us now focus on one lane in the PCIe data link, and let us narrow our focus further to one simplex (one-way) path in the bi-directional lane. There is a transmit (Tx) device and a receive (Rx) device. The Tx and Rx devices can be modeled as shown on the next page.

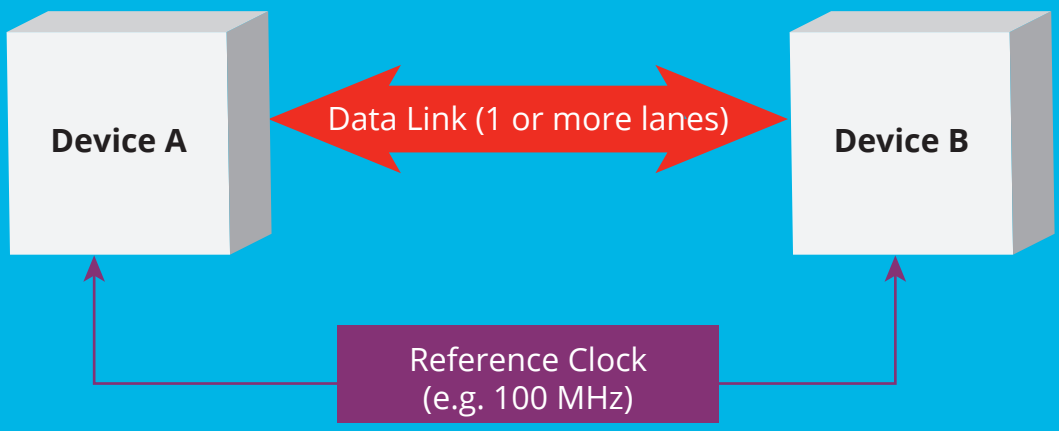

Ideally, the clock driving the Rx Latch is an exact duplicate in frequency and phase of the clock driving the Tx Latch. If this is not the case, may cause errors. The jitter of the Rx Latch clock relative to the Tx Latch clock.

In real applications, of course, the reference clock is altered as it progresses to the Tx and Rx latches. Any jitter present on the Refclk is low-pass filtered by the Tx and Rx PLLs. In assessing the performance of the reference clock, the PCIe-defined jitter is not equal to the traditional 12 kHz–20 MHz integrated phase noise. The PCIe specification post-filters the phase noise to calculate PCIe jitter values that are seen in many wired communications applications. The Clock and Data Recovery block high pass filters its incoming clock. The clock filtering effects of the PLLs and the CDR block are modeled by the PCIe standards as transfer functions which are different for each of the three aforementioned blocks' incoming clock. Each PCIe version expresses these individual transfer functions $H_{TxPLL}(s)$, $H_{RxPLL}(s)$, and $H_{CDR}(s)$ differently. The general jitter transfer function that acts on the Refclk, irrespective of how the individual transfer functions are modeled, is expressed in the relevant PCIe specifications as:

$$H_{jitter}(s) = [H_{TxPLL}(s)e^{-sT} - H_{RxPLL}(s)]\, H_{CDR}(s)$$

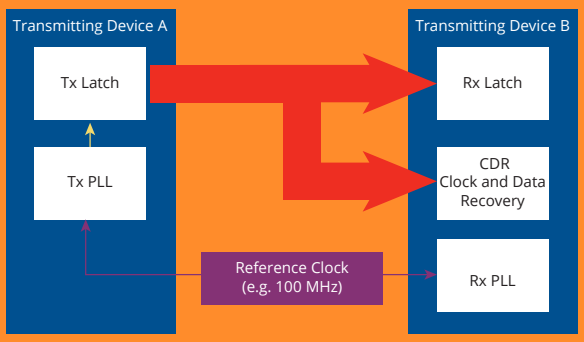

Each PCIe version specification further provides multiple possibilities for HTxPLL(s), HRxPLL(s) and $H_{CDR}(s)$ to model various PCIe use cases. Consequently, there are many calculated $H_{jitter}(s)$ for each PCIe version.

As PCIe lane speeds have increased with each new PCIe version, the magnitude of the allowed rms jitter—calculated after $H_{jitter}(s)$ is applied to the Refclk—has decreased to maintain acceptable bit error rates.

Increasing	PCIe® Version (Year Introduced)	Maximum Allowed PCIe Post-Filtered Jitter Measured in the [fs, rms]	Decreasing
	2.0 (2007)	3100	
	3.0 (2010)	1000	
	4.0 (2017)	500	
	5.0 (2019)	150	
	6.0 (2021)	100	

But how can we determine whether a Refclk source we have selected meets the specification for maximum rms jitter? We must save a Refclk waveform set of points from many periods captured by an oscilloscope or (as of PCIe 5.0) we can use the data provided by a phase noise analyzer. Microchip's PCIe jitter calculator takes the Refclk input data and produces a detailed report showing the calculated rms jitter after applying the many possible $H_{jitter}(s)$ transfer functions. It also clearly shows whether the Refclk data passes or fails the PCIe jitter specification. The PCIe jitter calculator is located at: microchip.com/PCIeCalculator

Note that if there are "N" possible $H_{jitter}(s)$ transfer functions for a given PCIe version, there are 2N calculated maximum rms jitter results because $H_{jitter}(s)$ must be applied to a Refclk with spread-spectrum on as well as to a Refclk with spread-spectrum off. (Spread spectrum modulation is explained in Chapter 3).

Below, we see a portion of the report produced by the Microchip PCIe Jitter Calculator.

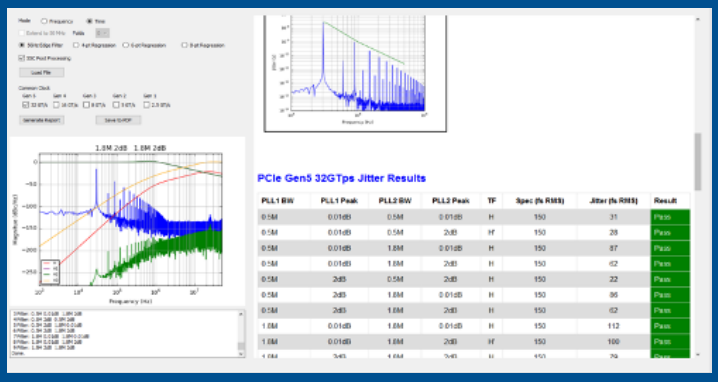

Chapter 3

PCIe: Spread-Spectrum Clocks

When a clock is generated in any system, it will radiate a certain amount of its energy into its surroundings. Sometimes, as is the case for a broadcast radio antenna, this is the desired result. The maximum amount of power emitted as waves in the RF spectrum is desired so as to reach the farthest possible distance. Other times, however, maximum power is not desired. If a radio station is transmitting with so much power that it interferes with another radio station of the same frequency some distance away, a receiver produces an undesired mixture of the two radio stations when it converts the radio transmissions to sound.

This is why in the United States, for example, the more powerful AM radio stations are required to reduce their transmitting power during nighttime hours: The characteristics of the earth's atmosphere tend to result in the RF signals travelling thousands of km farther with enough power to be detected by a typical radio receiver than these same RF signals travel during daylight hours. In other words, one electromagnetic wave is interfering with another; hence the name Electromagnetic Interference (EMI). A powerful AM radio station "A" reduces its power at night so that a typical receiver that should be demodulating the signal of radio station "B" a certain distance away cannot detect the signal from "A". Note that the signal from station "A" needs only to be below the threshold of a radio receiver's ability to detect it when that receiver is tuning in to station "B". The power of signal "A" at the location of the receiver is not required to be zero.

There are times when a PCIe clock may interfere with nearby clocks because of the power the PCIe clock radiates at, for example, 100 MHz. Another clock may incorporate the energy of that PCIe clock, thereby corrupting that clock. The solution is to reduce the power at that 100 MHz carrier so that it is low enough not to corrupt the clocks in nearby systems. This is the reason that the PCIe standards provide Spread-Spectrum Clock (SSC) specifications; SSC modulation is used to intentionally modulate the ideal position of the clock edge so that the resulting signal's spectrum is spread around the ideal frequency of the clock. This reduces the peak amplitude of the EMI associated with the fundamental frequency of the signal. An SSC, however, has much higher jitter than the un-modulated clock. In fact, SSC introduces significantly more noise power into the system. The trade off for this extra noise power is that the peaks are at levels low enough not to interfere with the reception of other clocks.

Below we see the spectrum of a 100 MHz square wave clock signal generated by a Microchip ZL30251 clock synthesizer. This plot shows a narrow band of the spectrum, 2.5 MHz above and below the carrier. The carrier frequency is 100 MHz, and power amplitude at that frequency is 0.5 dBm. The next highest power peaks are located approximately ±0.75 MHz from the carrier. Because they are significantly lower, these smaller power peaks can be ignored for the purposes of this discussion (i.e., they are below the detectable threshold of other nearby systems).

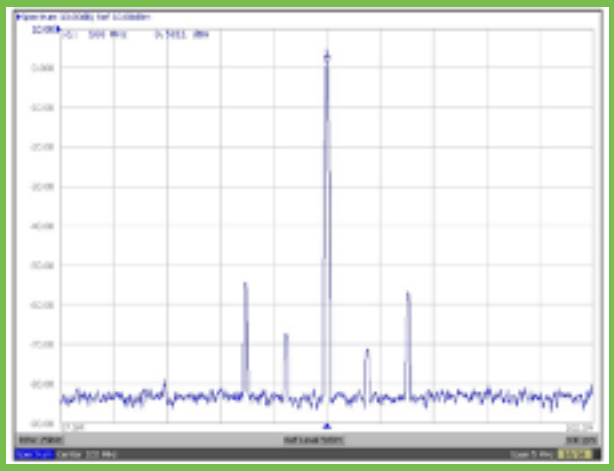

The high power of the carrier signal can result in radiated emissions and cause EMI if the circuit traces carrying the signals are not perfectly balanced and terminated. By modulating a clock with spread-spectrum, high power in a single narrow frequency band is replaced with lower power spread over a wider band. This reduces the EMI associated with the signal.

To generate a PCIe SSC, a square wave carrier (usually 100 MHz) is frequency modulated by a triangle wave with a frequency in the range of 30–33 kHz. The amplitude of the triangle-wave modulation is selected to result in a spread amplitude of 0.5% of the nominal carrier frequency.

The SSC in the image below (overlaid with the unmodulated clock shown above) demonstrates the "flattening" of the power spectrum. In this example, a 30 kHz triangle wave with 0.5% down-spread is used. The term down-spread implies that the carrier is modulated to lower frequencies, not higher, so the maximum frequency of the spread-spectrum signal equals the nominal clock frequency without spread. The lowest frequency component is 0.5% below the carrier, at 99.5 MHz in this example. An alternative to down-spread is center-spread where the carrier is modulated both lower and higher than the nominal clock frequency by a ±0.25%.

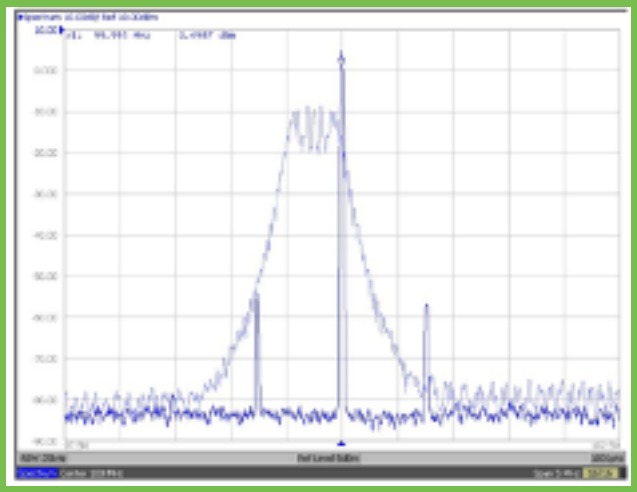

In the 10 kHz–50 MHz PCIe phase-offset-from-carrier integration band, a 30 kHz–33 kHz modulation creates many spurs (i.e., deterministic jitter) within the integration band that can increase the rms jitter by two or three orders of magnitude. The modulation of the clock edges in the most widely adopted PCIe architecture (Common Clock) provides the identical clock edge to both the transmitting (Tx) and receiving (Rx) device. Therefore, the relative jitter (described in Chapter 2) is only very slightly impacted. The Tx device clock edges and the Rx device edges move in phase (i.e. are modulated) identically, resulting in the difference between an Rx device clock phase and a Tx device clock phase being virtually unchanged when compared to this same Rx/TX phase difference of an unmodulated clock at any given instant.

Microchip provides a large portfolio of PCIe-compliant clock generators with a variety of features and specifications. For more information, please see microchip.com/PCIeTiming

Chapter 4

PCIe Clocks in Automobiles

Automotive electronic content continues to increase. As a result, the communications bus between various systems must support increasing data loads. Automotive Ethernet significantly improved data transmission rates compared to older standards such as the CAN bus and, therefore, has become the parallel bus of choice. Similarly, the burgeoning bandwidth requirements resulting from increasingly complex infotainment systems, Advanced Driver-Assistance Systems (ADAS) and connectedness to the Internet of Things (IoT) is leading vehicle manufacturers to embrace the highly versatile PCIe bus standard as the serial bus of choice for high-speed automotive applications.

As described in the Automotive Electronics Council AEC-Q100 specification, integrated circuit components used in automobiles must:

- Operate across an extended temperature range
- Withstand significant shock and vibration
- Operate with very noisy power supplies
- Fit within very limited spaces
- Have high reliability

Microchip provides AEC-Q100 and PCIe 1.0–PCIe 5.0 compliant MEMS oscillators, buffers and clock generators.

Automotive MEMS Oscillators

Micro-Electro-Mechanical Systems (MEMS) oscillators are uniquely suited to the harsh environment of a vehicle. For example, Microchip's AEC-100-qualified DSA12x3 family of PCIe 1.0–PCIe 5.0 compliant MEMS oscillators:

- Operates from −40˚C to +125˚C
- Withstands 50K G Shock and 70K G vibration
- Has high power supply noise rejection (−50 dBc)
- Has a small package size (2.5 mm × 2.0 mm footprint)
- Is qualified to MIL-STD-883 with a mean time between failures (MTBF) 20 times longer than quartz oscillators

Below, we see the significant stability over temperature that a Microchip MEMS oscillator provides compared to a traditional quartz-based oscillator.

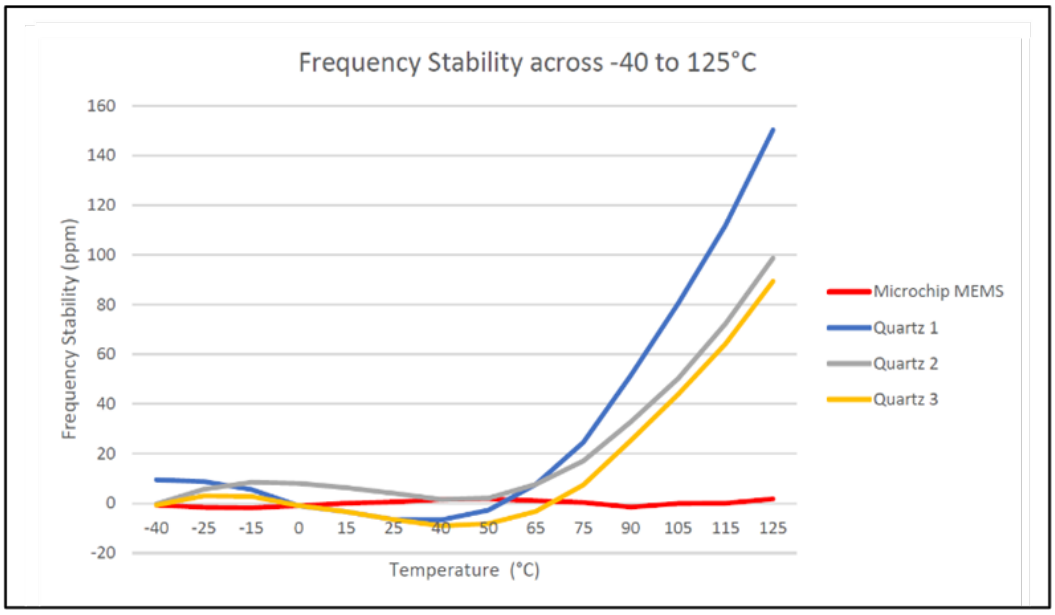

This excellent stability over temperature can be a crucial system requirement, especially with safety-critical systems such as ADAS. Microchip provides the following selection of AEC-100 and PCIe compliant MEMS oscillators:

Device	Package Footprint (mm)	Maximum Temperature Range (˚C)	PCIe Compliance Generations	Available Output Formats
DSA12x2	2.5 × 2.0	−40 to +105	1–5	LVPECL
DSA12x3	2.5 × 2.0	−40 to +125	1–5	LVDS
DSA12x4	2.5 × 2.0	−40 to +105	1–5	HCSL
DSA1124	2.5 × 2.0	−40 to +105	1–4	HCSL

Automotive Buffers

Oftentimes, the output of an oscillator needs to be buffered to create more copies of that clock; consequently, automotive-grade buffers are required. Complementing these MEMS oscillators are Microchip's AEC-100 and PICe 1.0–PCIe 5.0 compliant Low-Power HCSL (LPHCSL) output format buffers:

Device	Number of Outputs	Package Footprint (mm)	Maximum Temp Range (˚C)	Output Format
SY75602TWL	2	1.4 × 1.6	−40 to +85	LPHCSL
SY75603TWL/ SYA75603TWLVAO	2	3.0 × 3.0	−40 to +85/−40 to +105	LPHCSL
SY75604TWL/ SYA75604TWLVAO	4	3.0 × 3.0	−40 to +85/−40 to +105	LPHCSL
SY75608TWL/ SYA75608TWLVAO	8	6.0 × 6.0	−40 to +85/−40 to +105	LPHCSL
SY75612TWL/ SYA75612TWLVAO	12	8.0 × 8.0	−40 to +85/−40 to +105	LPHCSL

Automotive Clock Generators

Multiple outputs can be created when a device combines a MEMS oscillator with buffers and phase-locked-loops. These outputs need not be the same frequency as the MEMS oscillator output clock. Microchip provides PCIe 1.0–PCIe 5.0 compliant and AEC-100 qualified clock generators designed specifically for the automotive market.

Device	Number of Outputs	Minimum Package Footprint (mm)	Maximum Temperature Range (˚C)	Available Output Formats
DSA557-05	4	5 × 3.2	–40 to +105	Fixed 100 MHz PCIe clocks HCSL, LVDS, CMOS
DSA557-04	3	5 × 3.2	–40 to +105	Fixed 100 MHz PCIe clocks HCSL, LVDS, CMOS
DSA557-03	2	3.2 × 2.5	–40 to +105	Fixed 100 MHz PCIe clocks HCSL, LVDS, CMOS
DSA400	4	5.0 × 3.2	–40 to +105	General-purpose clock generator CMOS, HCSL, LVDS, LVPECL
DSA2000	2	3.2 × 2.5	–40 to +105	General-purpose clock generator CMOS, HCSL, LVDS, LVPECL

Chapter 5

PCIe Clocks in Data Centers

Owing to its high data rates and ability to customize throughput to meet varying bandwidth requirements, PCIe has become the most widely implemented data bus in data centers around the world. As cloud-computing has increased and applications related to the IoT have grown, so have the demands on data center transfer bandwidth between servers and peripherals such as storage devices and input/output devices. Higher data rates have resulted in increasingly tighter jitter requirements, as evidenced in the evolution of PCIe jitter specifications (described in Chapter 1).

Microchip provides a total system solution for PCIe data communications between a server and peripherals. Microchip PCIe-compliant oscillators, clocks, and buffers have been qualified and validated in data centers worldwide, as well as being selected as the vendor-preferred clock source on numerous vendor designs. See, for example, Microchip Application Note AN3604.

Data centers not only require low-jitter clocks for internal communication between servers and peripherals, but also for external communication via, for example, physical layer communication devices (commonly known as PHYs). One such implementation is shown in this example. Here, we show how the ZL3026x family of devices can provide multiple output frequencies. Note also that, if needed, the ZL3026x can simultaneously generate some or all of its output clocks with spread-spectrum enabled and others with no spread-spectrum enabled modulation.

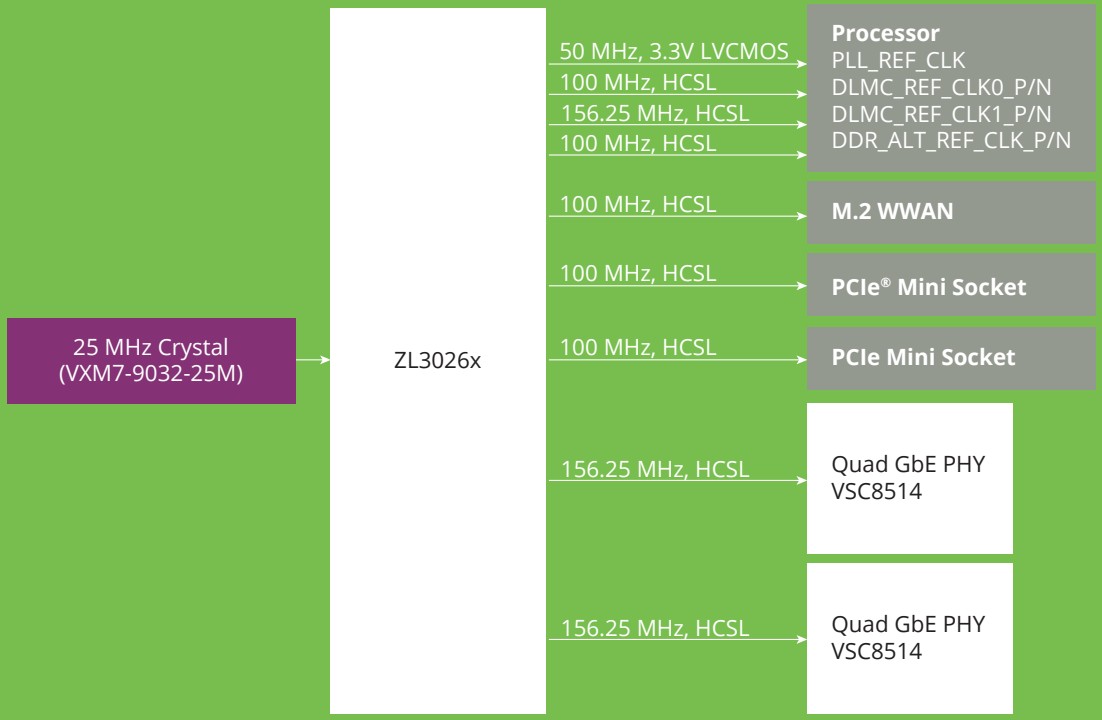

Here we see another example of Microchip crystals, clock generators, and buffers providing a complete solution for a PCIe switch. The symbols with the letters EP on them represent end points that can be Network Processing Units (NPUs), Graphical Processing Units (GPUs), FPGAs, Solid State Drives (SSDs) and so forth.

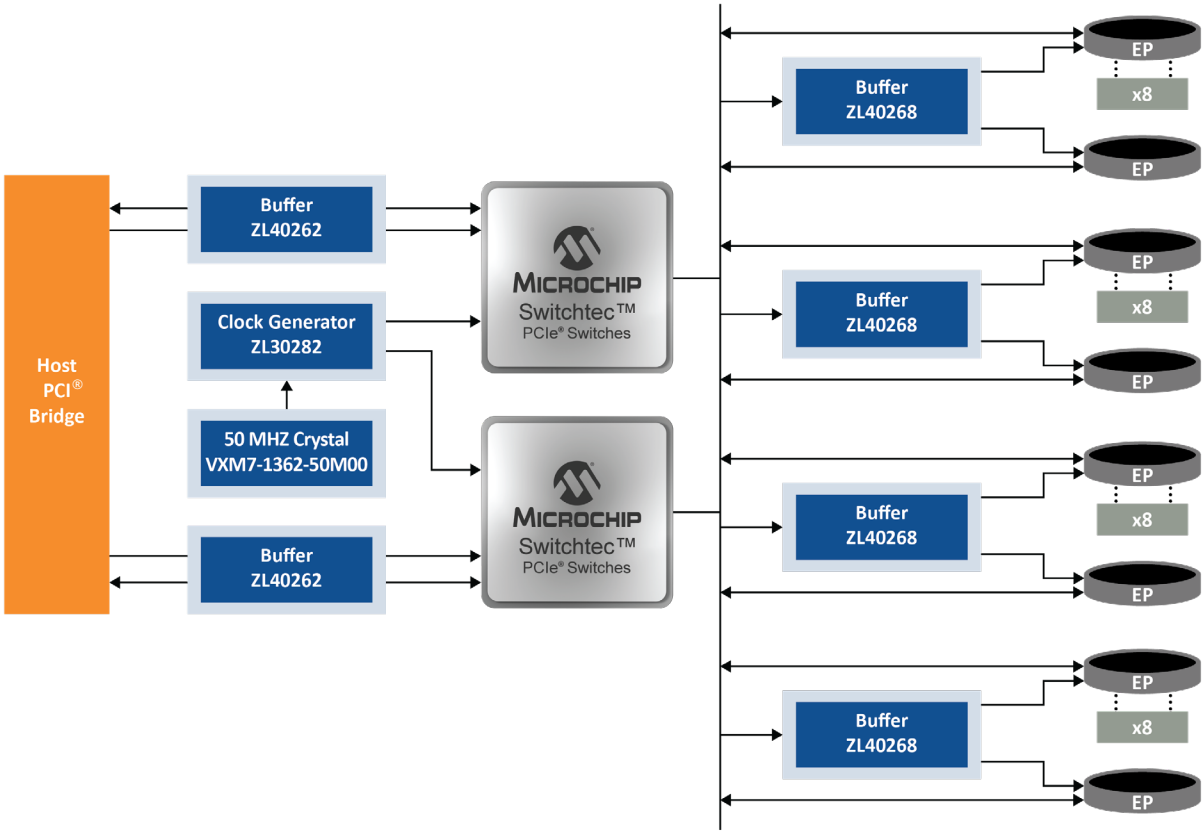

In this example, we see a ZL30263 Microchip clock generator providing the PCIe clocks for SSDs and a PCIe-to-PCIe bridge.

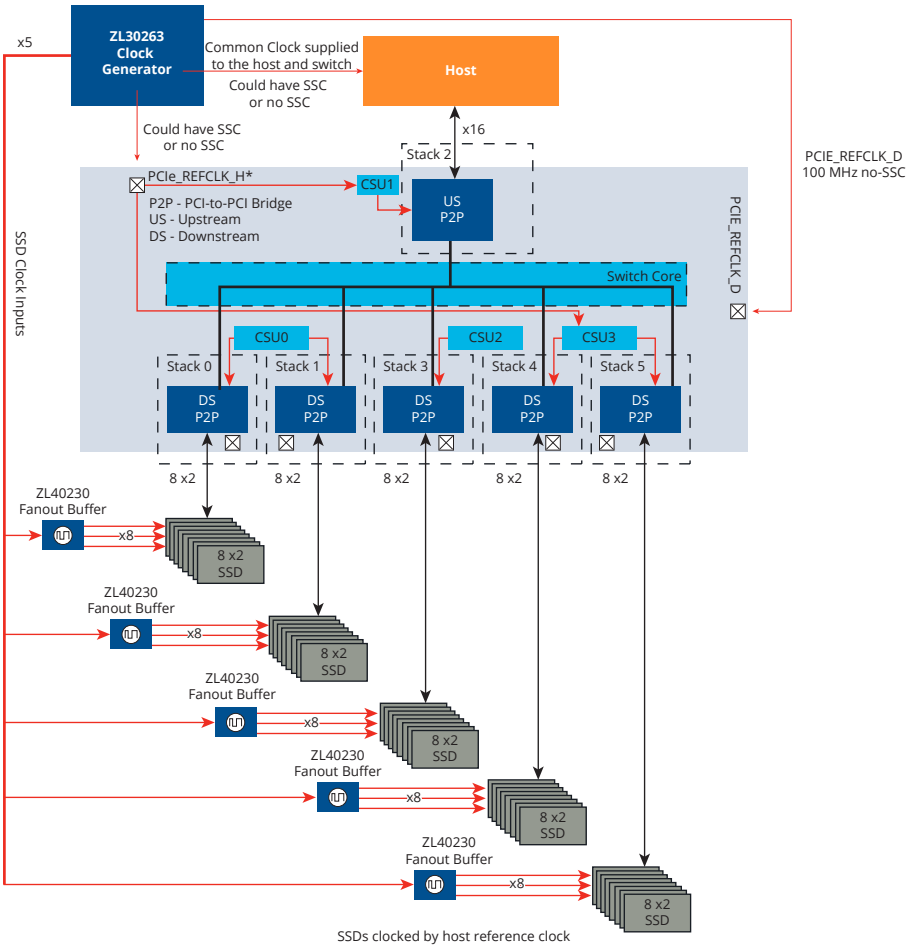

As shown in the following table, Microchip offers a wide variety of PCIe-compliant oscillators, buffers, and clock generators that can be tailored to the specific data center implementation requirements.

Clock Generators							
Part Name	**PCIe® Gen**	**# or Outputs**	**Input**	**Output Logic**	**Voltage (V)**	**Package Dimensions**	**Temp. Range**
PL602-21SC	Gen 1/2	1	25 MHz Crystal or Ref	HCSL	2.5–3.3	8-pin SOP 5 × 4 mm	–40°C+ 85°C
SM802379UMG	Gen 1/2/3/4/5	2	25 MHz Crystal or Ref	HCSL	2.5–3.3	24-pin QFN 4 × 4 mm	–40°C+ 85°C
ZL30281LDF1	Gen 1/2/3/4/5	2	25 MHz Crystal or Ref	CML	2.5–3.3	32-pin QFN 4 × 4 MM	–40°C+ 85°C
SM806033UMG	Gen 1/2/3/4/5	4	50 MHz Crystal	HCSL	2.5–3.3	48-pin QFN 7 × 7 mm	–40°C+ 85°C
MX875BB0022	Gen 1/2/3/4/5	4	Crystal integrated	HCSL	2.5–3.3	48-pin QFN 7 × 7 mm	–40°C+ 85°C
MX875BB0020	Gen 1/2/3/4/5	4	Crystal integrated	LVDS	2.5–3.3	48-pin QFN 7 × 7 mm	–40°C+ 85°C
MX852EB0027	Gen 1/2/3/4/5	5	Crystal integrated	HCSL	2.5–3.3	38-pin LGA 7.0 × 5.0 mm	–40°C+ 85°C
SM806034UMG	Gen 1/2/3/4/5	6	50 MHz Crystal	HCSL	2.5–3.3	48-pin QFN 7 × 7 mm	–40°C+ 85°C
MX875BB0023	Gen 1/2/3/4/5	6	Crystal integrated	HCSL	2.5–3.3	48-pin QFN 7 × 7 mm	–40°C+ 85°C
DSC557-0344FI1	Gen 1/2/3/4	2	MEMS Integrated	HCSL	2.5–3.3	14-pin QFN 3.2 × 2.5 mm	–40°C+ 85°C (2)
DSC557-0333FI1	Gen 1/2/3/4	2	MEMS Integrated	LVDS	2.5–3.3	14-pin QFN 3.2 × 2.5 mm	–40°C+ 85°C (2)
DSC557-04444KI1	Gen 1/2/3/4	3	MEMS Integrated	HCSL(1)	2.5–3.3	20-pin QFN 5.0 × 3.2 mm	–40°C+ 85°C (2)
DSC557-0544444KI1	Gen 1/2/3/4	4	MEMS Integrated	HCSL(1)	2.5–3.3	20-pin QFN 5.0 × 3.2 mm	–40°C+ 85°C (2)
Clock Generator with Spread Spectrum							
ZL30282	Gen 1/2/3/4/5	6	50 MHz Crystal or Ref	HCSL(1)	2.5–3.3	56-pin QFN 8 × 8 mm	–40°C+ 85°C
ZL30265LDG1Q062	Gen 1/2/3/4/5	6	50 MHz Crystal or Ref	HCSL (1)	2.5–3.3	56-pin QFN 8 × 8 mm	–40°C+ 85°C
ZL30265LDG1Q062	Gen 1/2/3/4/5	6	50 MHz Crystal or Ref	HCSL(1)	2.5–3.3	56-pin QFN 8 × 8 mm	–40°C+ 85°C
ZL30267	Gen 1/2/3/4/5	10	50 MHz Crystal or Ref	HCSL(1)	2.5–3.3	56-pin QFN 8 × 8 mm	–40°C+ 85°C
ZL30267LDG1Q06Y	Gen 1/2/3/4/5	10	50 MHz Crystal or Ref	HCSL(1)	2.5–3.3	56-pin QFN 8 × 8 mm	–40°C+ 85°C
Oscillators							
MX55ABD100M000	Gen 1/2/3/4/5	1	Crystal integrated	HCSL	2.5–3.3	6-pin 5 × 3.2 mm	–40°C+ 85°C
MX55ABB100M000	Gen 1/2/3/4/5	1	Crystal integrated	LVDS	2.5–3.3	6-pin 5 × 3.2 mm	–40°C+ 85°C
VC-820-9005-100M000000	Gen 1/2/3/4/5	1	Crystal integrated	CMOS	1.8	4-pin 3.2 × 2.5 mm	–40°C+ 85°C
VC-820-9005-100M000000	Gen 1/2/3/4/5	1	Crystal integrated	LVDS	1.8	6 pin 3.2 × 2.5 mm	–40°C+ 85°C (2)
DSC1124DI1-100.0000	Gen 1/2/3/4	1	MEMS Integrated	HCSL	2.5–3.3	6-pin 5 × 3.2 mm	–40°C+ 85°C (2)
DSC1224DI1-100M0000	Gen 1/2/3/4/5	1	MEMS Integrated	HCSL	2.5–3.3	6-pin 5 × 3.2 mm	–40°C+ 85°C (2)
DSC1223DI1-100M0000	Gen 1/2/3/4/5	1	MEMS Integrated	LVDS	2.5–3.3	6-pin 5 × 3.2 mm	–40°C+ 85°C (2)
Buffers							
ZL40262	Gen 1/2/3/4/5	2	HCSL, PECL, LVDS, CMOS	HCSL	2.5–3.3	20-pin QFN 4 × 4 mm	–40°C+ 85°C
ZL40264	Gen 1/2/3/4/5	4	HCSL, PECL, LVDS, CMOS	HCSL	2.5–3.3	20-pin QFN 4 × 4 mm	–40°C+ 85°C
ZL40268	Gen 1/2/3/4/5	8	HCSL, PECL, LVDS, CMOS	HCSL	2.5–3.3	48-pin QFN 7 × 7 mm	–40°C+ 85°C
ZL40272	Gen 1/2/3/4/5	12	HCSL, PECL, LVDS, CMOS	HCSL	2.5–3.3	56-pin QFN 8 × 8 mm	–40°C+ 85°C
SY75602(3)	Gen 1/2/3/4/5	2	HCSL, PECL, LVDS, CMOS	LPHCSL	1.8–3.3	8-pin VDFN 1.4 × 1.6 mm	–40°C+ 85°C
SY75603(3)	Gen 1/2/3/4/5	2	HCSL, PECL, LVDS, CMOS	LPHCSL	1.8–3.3	16-pin QFN 3 × 3 mm	–40°C+ 85°C (2)
SY75604(3)	Gen 1/2/3/4/5	4	HCSL, PECL, LVDS, CMOS	LPHCSL	1.8–3.3	16-pin QFN 3 × 3 mm	–40°C+ 85°C (2)
SY75608(3)	Gen 1/2/3/4/5	8	HCSL, PECL, LVDS, CMOS	LPHCSL	1.8–3.3	48-pin QFN 7 × 7 mm	–40°C+ 85°C (2)
SY75612(3)	Gen 1/2/3/4/5	12	HCSL, PECL, LVDS, CMOS	LPHCSL	1.8–3.3	56-pin QFN 8 × 8 mm	–40°C+ 85°C (2)
SY75572L	Gen 1/2/3/4	2	LVDS/HCSL	HCSL	3.3	16-pin QFN 3 × 3 mm	–40°C+ 85°C
SY75576L	Gen 1/2/3/4	4	LVDS/HCSL	HCSL	3.3	20-pin TSSOP 6.4 × 6.5 mm	–40°C+ 85°C
SY75578L	Gen 1/2/3/4	8	LVDS/HCSL	HCSL	3.3	32-pin QFN 5 × 5 mm	–40°C+ 85°C

1. LVDS and LVCMOS output logic options available 2. –40°C+ 105°C options available 3. Sampling, production release

Microchip Technology Inc. | 2355 W. Chandler Blvd. | Chandler AZ, 85224-6199 | microchip.com

DS00003818A