Selecting the Optimum PCI Express Clock Source

PCI Express (PCIe) is a serial point-to-point interconnect standard developed by the Peripheral Component Interconnect Special Interest Group (PCI-SIG). Although originally designed for desktop personal computers, the PCIe standard has been widely adopted in a broad range of applications including blade servers, storage, embedded computing, and networking and communications. Not only is the PCIe interface supported by a wide base of commercially available devices, it is also becoming more readily available in FPGAs and SoCs, providing designers with flexible solutions for transferring data within their systems. One of the key advantages of using PCIe is its scalable data bandwidth and flexible clocking solutions. Let's explore some of the standard clocking architectures for PCIe and consider their benefits for typical system applications.

The PCIe Link

Before considering clocking architectures, let's examine a PCIe data link. It consists of one or more lanes that provide a transmit (Tx) and receive (Rx) differential pair. Figure 1 shows two devices that need to transfer data. One of the key advantages of PCIe is its bandwidth scalability enabling up to 32 lanes to be configured on a single link, but, due to space limitations, most of the commercial PCIe links work with 16 lanes or less. With the recent introduction of PCIe Gen 5, each lane can accommodate up to 32G transactions per second (GT/s) for a maximum throughput close to 128 GB/s with 16 lanes. Applications that need less data bandwidth can simply be configured with fewer lanes. Choosing a PCIe standard with higher data rates ultimately means using less lanes or connection wires between devices, but it also places additional requirements on the clocking performance. We will examine these requirements in the following sections.

	Raw Bit Rate	Data Throughput Per Lane Per Direction	Max Data Throughput (16 Lanes Duplex)	Year Released
PCIe 1.1	2.5 GT/s	250 MBytes/s	8 GBytes/s	2003
PCIe 2.1	5.0 GT/s	500 MBytes/s	16 GBytes/s	2007
PCIe 3.1	8.0 GT/s	~1 GBytes/s	~32 GBytes/s	2010
PCIe 4.0	16 GT/s	~2 GBytes/s	~64 GBytes/s	2017
PCIe 5.0	32 GT/s	~4 GBytes/s	~128 GBytes/s	2019

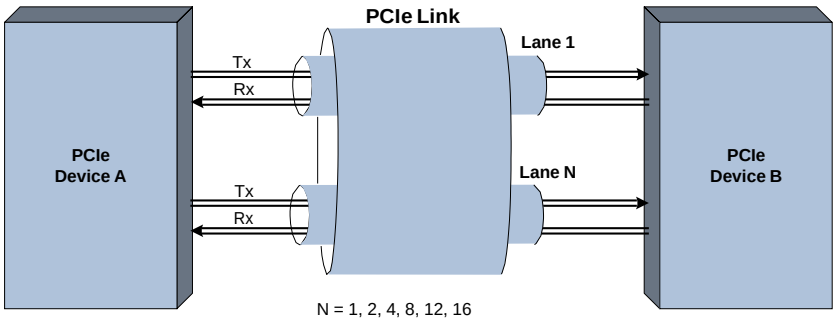

Figure 1. The PCI Express Link

PCIe Applications

Because of the popularity of PCIe, growing numbers of application-specific devices (e.g., ASICs and SoCs) are adopting the PCIe interface as a common interconnect with other devices that are readily available on the market. Even today's field-programmable gate arrays (FPGAs) offer built- in PCIe protocol stacks and physical layer interfaces to help simplify system-level design. Figure 2 illustrates a few examples of system-level solutions using PCIe interconnects. It is important to note that reliable data transmission over a PCIe link requires a stable clock reference at both the transmitting and receiving ends.

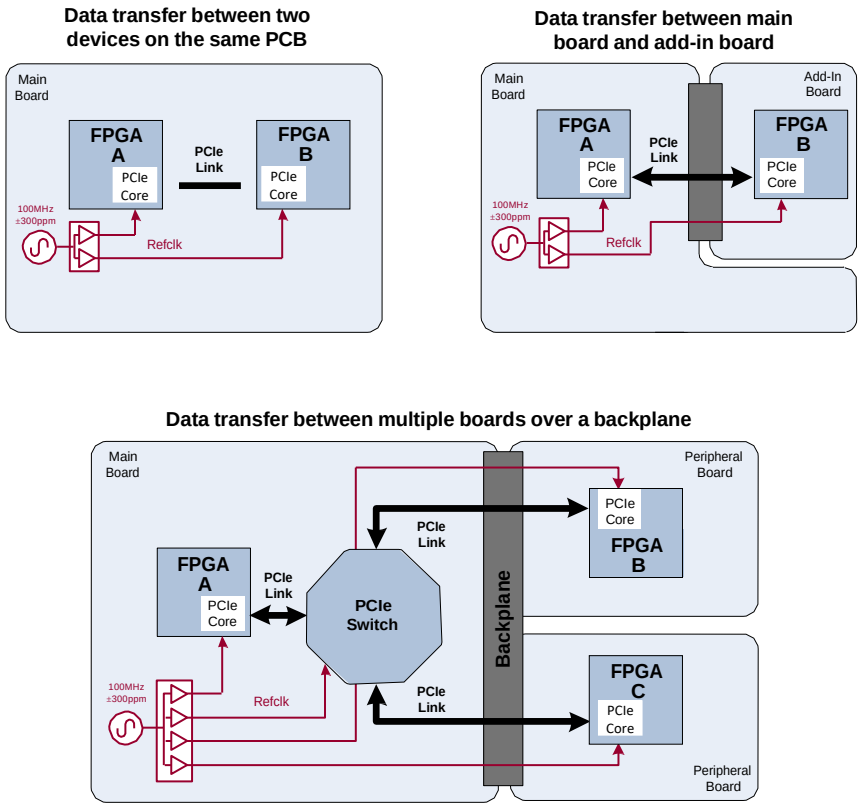

Figure 2. System Applications of PCIe Interconnects

PCIe Clocking Architectures

The PCIe standard specifies a 100 MHz clock (Refclk) with at least ±300 ppm frequency stability for Gen 1, 2, 3 and 4, and at least ±100 ppm frequency stability for Gen 5, at both the transmitting and receiving devices. It also specifies support for different clocking architectures: Common Clock, Data Clock, Separate Reference Clocks with No Spread-Spectrum Clocking (SRNS), and Separate Reference Clocks with Independent Spread-Spectrum Clocking (SRIS). Figure 3 shows a block diagram of each architecture. Of all the architectures, the Common Clock is the most widely supported clocking method used by commercially available devices. An advantage of this clocking architecture is that it supports spread spectrum clocking (SSC) which can be very useful in reducing electromagnetic interference (EMI) with less stringent reference clock requirements than SRIS. A disadvantage is that the same clock source must be distributed to every PCIe device while keeping the clock-to-clock skew to less than 12 ns between devices. This can be a challenge for large circuit boards or when crossing over a backplane connector to another circuit board. Examples of applications using the Common Clock architecture are also shown in Figure 2.

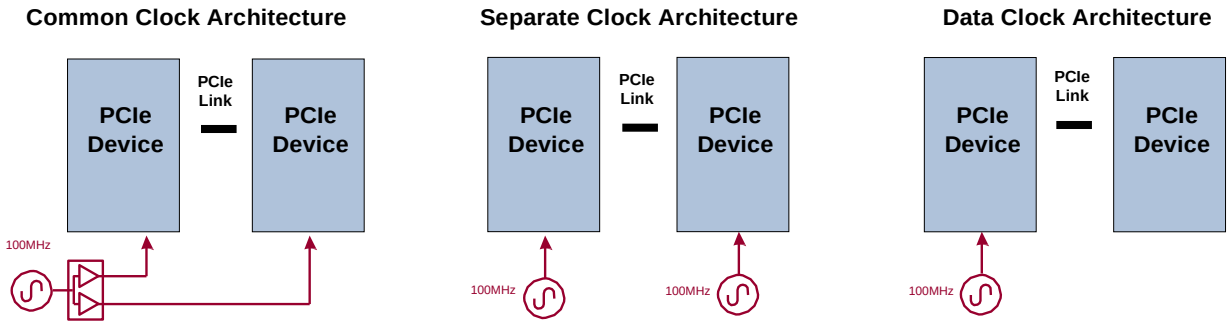

Figure 3. PCIe Clocking Architectures

Another clocking architecture is the *Separate Reference architecture* where a different clock source is used at each end of the PCIe link. The advantage of this architecture is that tightly-controlled reference clock distribution is no longer required over connectors and backplanes. The *Data Clock* architecture is the simplest to implement since it only requires one clock source located at the transmitter. In this case, the receiver simply extracts the embedded clock from the transmitter. The data clock architecture is only supported in PCIe Gen 2 and Gen 3, so its raw bit rate is limited to 8 GT/s.

Refclk Frequency and Jitter Requirements

The industry-standard reference clock frequency used for devices supporting PCIe 1.1, 2.1 and 3.1 is 100 MHz (±300 ppm generated using an HCSL signal format). It is common for embedded processors, system controllers and SoC-based designs to use 100 MHz HCSL format as the reference clock for the PCIe bus interface circuitry.

However, in applications that use FPGAs, the PCIe reference clock requirements can deviate from the standard 100 MHz HCSL format to other frequencies and formats such as 125 MHz, 200 MHz or 250 MHz in LVCMOS, LVDS or LVPECL. A typical example is an FPGA that supports both PCIe and Ethernet functions. Using a common reference clock frequency of 125 MHz to clock both functions helps to reduce clocking domains or timing islands in the FPGA. Internally the FPGA multiplies this reference clock to the required PCIe lane rate (e.g. 125 MHz x64 for PCIe 3.0). Depending on the mix of ICs used in a design, the PCIe clock strategy may vary from one of generating multiple 100 MHz HCSL clocks to generating a mix of different frequencies and output formats.

The jitter performance of the reference clock is also an important consideration. While the higher data throughput of PCIe Gen 5 is intended to reduce the number of interconnect wires between two devices, it also requires reference clocks with low jitter. Clock sources that meet the jitter requirements of previous PCIe specifications may not meet the needs of a device that supports PCIe Gen 5.

Table 1 summarizes the jitter requirements for all PCI Express standards. Silicon Labs provides a free software, the PCIe Clock Jitter Tool, which allows for quick and easy characterization of the reference clock across all the PCIe specifications and architectures, including PCIe Gen 5, according to the most recent PCI-SIG releases to ensure standards compliance for the system design.

Table 1. PCIe Clocking Architectures

	Description	Symbol	Limit	Units
Common Clock Architecture				
Gen 1	Random Jitter	Rj	4.7	ps pk-pk
	Deterministic Jitter	Dj	41.9	ps RMS
Gen 2	High Frequency RMS Jitter Measured from 10 kHz to 50MHz	J_{RMS-HF}	3.1	ps RMS
Gen 3	High Frequency RMS Jitter Measured from 10 kHz to 50MHz	J_{RMS-HF}	1	ps RMS
Gen 4	High Frequency RMS Jitter Measured from 10 kHz to 50MHz	J_{RMS-HF}	0.5	ps RMS
Gen 5	High Frequency RMS Jitter Measured from 10 kHz to 50MHz	J_{RMS-HF}	0.15	ps RMS
Data Clock Architecture				
Gen 2	High Frequency RMS Jitter Measured from 1.5 MHz to 50MHz	J_{RMS-HF}	4	ps RMS
	Low Frequency RMS Jitter Measured from 10 kHz to 1.5 MHz	J_{RMS-LF}	7.5	ps RMS
Gen 3	High Frequency RMS Jitter Measured from 10 kHz to 50MHz	J_{RMS-HF}	1	ps RMS

The tight requirements of PCIe Gen 5 push the clock source performance close to the noise floor of some measurement equipment, making it challenging to characterize. A detailed discussion of how to take precise PCIe Gen 4 and Gen 5 measurements can be found in AN1104 available at http://www.silabs.com/products/timing/pci-express-learning-center .

Spread Spectrum Clocking (SSC)

The use of spread spectrum clocking is an important consideration when selecting a PCIe clock source since it is an effective method of lowering the amount of radiated electromagnetic interference (EMI) that is generated from high speed clock and datapath signals. Spread spectrum clocks use low frequency modulation of the carrier frequency to spread out the radiated energy across a broader range of frequencies. Radiation from data lines transmitted with a device using a spread spectrum clock reference will also benefit from the same EMI reduction. PCIe Gen 1 to Gen 5 devices are specified to reliably transmit data when using a Refclk with a spread spectrum modulation rate of 30 to 33 kHz and modulation amplitude up to 0.5%.

Separate Reference Architectures – SRNS and SRIS

The simplifications that come with the PCIe Separate Reference architectures when compared to Common Clock architectures (that need a central clock source to be distributed to all devices) is an increase in the performance requirements of the reference clocks. Having different clock sources for each device means that now there are two jitter sources for the system, as shown in Figure 4. For most PCIe clock sources, the dominant jitter is uncorrelated and random, so we can find the total jitter of the two system clocks by taking the root sum square of the RMS jitter of each of them. It is important to notice that currently, the PCIe Gen4 and Gen 5 specifications do not define the jitter limits for SRNS or SRIS. But, using a conservative approach that assumes the same type of clock is used on both devices, and that they contribute equally to the total jitter of the system, we can calculate a Separate Reference clock jitter limit for all PCIe compliances based on the Common Clock jitter requirements, as listed in Table 2.

Note that the total reference clock jitter budget of the system is larger than the RMS clock component jitter specified by PCI-SIG. For example, the Gen 5 specification budgets 250 fs RMS for clock jitter (also known as the system simulation budget) but specifies 150 fs RMS for the reference clock component. This is to allow for noise coupling into the clock since the clock trace will often route long distances to connect to both the transmitter and receiver. However, in the Separate Reference clock architecture one clock source will be located near the transmitter and a separate clock source will be located near the receiver which eliminates the long trace and associated noise coupling. Therefore, when calculating the allowed component jitter for the Separate Reference architecture the entire system simulation budget is available for the clock components.

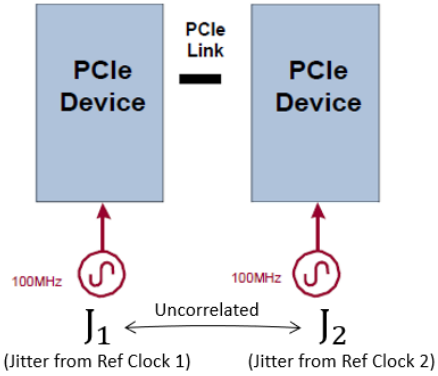

J_T (Total jitter of the system)

$$J_T = \sqrt{J_1^2 + J_2^2}$$

Assuming $J_1 = J_2$, we can set a max spec for SRNS and SRIS:

$$J_{SRNS,SRIS_{MAX}} = \frac{J_{CC_{Syst.Sim}}}{\sqrt{2}}$$

$J_{CC_{Syst.Sim}}$ (Common Clock System Simulation Max jitter spec)

	System Simulation Max Jitter
Gen 2	3.1 ps RMS
Gen 3	1 ps RMS
Gen 4	0.7 ps RMS
Gen 5	0.25 ps RMS

Figure 4. Separate Reference Architecture Jitter Distribution

Table 2. SRNS and SRIS Jitter Limits

	Description	Symbol	Limit	Units
SRNS and SRIS Jitter Limits				
Gen 2	High Frequency RMS Jitter Measured from 10 kHz to 50MHz	J_{RMS-HF}	2.192	ps RMS
Gen 3	High Frequency RMS Jitter Measured from 10 kHz to 50MHz	J_{RMS-HF}	0.707	ps RMS
Gen 4	High Frequency RMS Jitter Measured from 10 kHz to 50MHz	J_{RMS-HF}	0.495	ps RMS
Gen 5	High Frequency RMS Jitter Measured from 10 kHz to 50MHz	J_{RMS-HF}	0.177	ps RMS

Choosing the Optimum PCIe Clock Source

As we have seen, several factors should be considered when choosing the optimal PCIe clock source Ideally, the clock should provide performance sufficient to meet all PCIe specifications for the desired data rate and also support features like spread spectrum clocking. Standard off- the-shelf PCIe clock devices with 100 MHz HCSL format typically can address a majority of PCIe applications. In some cases, flexible output frequencies and flexible output signal formats are required. Single clock oscillators and buffers have traditionally been used for simple PCIe clocking applications, but frequency and format-flexible clock generators are widely used in more complex timing applications requiring reference clock generation of different frequencies and output formats such as LVDS and LVPECL.

Figure 5 shows two PCIe clocking solutions. One uses Silicon Labs' Si52204 clock generator as a single-chip PCIe clock tree solution. The Si522xx PCIe clock generators and Si532xx PCIe Buffers deliver the highest performance, lowest-footprint, lowest-power PCIe Gen 1 to Gen 5 clocking solutions. They are ideal for off-the-shelf PCIe devices that require 100 MHz HCSL clocks. Both the Si522xx and the Si532xx outputs are designed with Low-Power Push-Pull HCSL drivers that do not require any external resistors (Rs or Rt). The clock generators provide an additional pin-controlled, modulation amplitude-adjustable spread spectrum feature for SSC during electromagnetic interference (EMI) compliance testing.

Another solution using Silicon Labs' Si5332 low-jitter, any-frequency, any-output clock generator provides greater flexibility for clock trees for FPGA or custom ASIC and SoC designs. The Si5332 provides output-to-output skew adjustments to compensate for large differences in PCB trace lengths and supports spread spectrum on a per output basis, enabling developers to target EMI reduction in areas of the board where it is most needed. The flexibility of features in Si5332 not only simplifies system design and reduces component count and Bill of Materials cost, the in-system programmability of the Si5332 enables developers to make changes when prototyping in the lab and helps guarantee a successful first-pass design.

Si522xx and Si532xx PCI Express HCSL Clock Generators & Buffers

- High clock density – 2 to 12 output HCSL outputs
- PCIe Gen 1 to Gen 5 compliant
- Up to 50% lower power than competing solutions
- Dedicated output enable pins
- Pin selectable Spread Spectrum with 0.5% or 0.25% spread amplitude
- Low-power push-pull HCSL drivers that do not require any external resistors

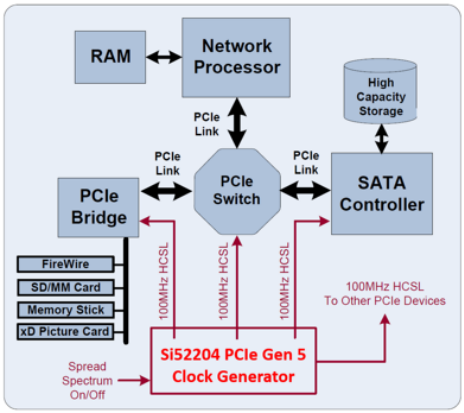

Pre-Configured Fixed Frequency PCIe Clock Solution

Si5332 low-jitter any-frequency any-output Clock Generator

- Up to 12 any-format outputs -HCSL, LVDS, LVPECL, LVCMOS
- Any-frequency output up to 333.33 MHz
- PCIe Gen 1 to Gen 5 compliant
- PCIe Gen 5 jitter: 25 fs (typ)
- Support 2 different Spread Spectrum domains at the same time
- Typical jitter (12 kHz to 20 MHz):<250 fs

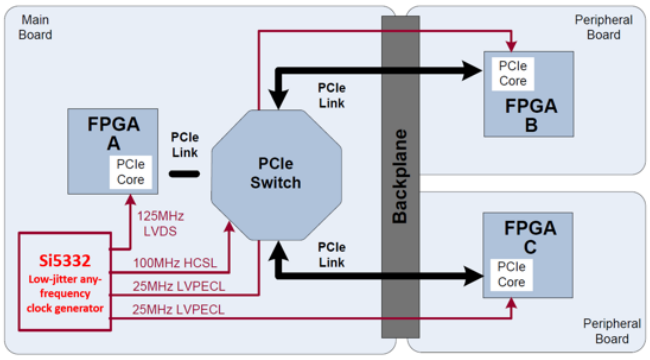

Flexible PCIe Clock Solution Using Web-Customizable Clock Generator

Figure 5. PCI Express Clock Generation Solutions

Summary

To ensure proper compliance with the PCIe standard, systems that use the PCIe interface require careful attention to the timing subsystem and architecture. Designers must consider which of the three PCIe specified reference clock architectures – Common Clock, Data Clock, SRNS, SRIS – will meet their application's functional and performance goals. Due to the diversity of ICs such as FPGAs, processors and switches that integrate PCIe cores, developers may want to use timing solutions that support multiple I/O voltage and formats, as well as spread spectrum clocking technology. Silicon Labs provides up to 12 outputs PCIe Clock generators, buffers and a wide portfolio of any-frequency programmable, multi-output clock generators that meet the requirements of today's most demanding PCIe Gen 1 to Gen 5 applications.
For much more information on Silicon Labs' PCIe products, PCIe measurement techniques and supporting software, sample schematics, layouts and reference manuals, please visit our PCIe Learning Center at http://www.silabs.com/products/timing/pci-express-learning-center .

# # #