PCI Express Timing Solutions in Automotive

Introduction

The growing demand for AD (automated driving) and ADAS (advanced driver assistance system) features in modern vehicles is driving an increase in the number and complexity of automotive electronic systems. These systems feature numerous sensors that generate a large amount of data and require increased computing power. Coprocessors are often used since a single SoC may not be sufficient to handle the computing requirements. PCI Express® (PCIe) is one of the interfaces commonly used to connect these components.

Automotive application, such as AD/ADAS, require real-time processing with ultra-low latency. Reliable and robust timing soltions are needed to support the data transfer between processors via PCIe. This paper will discuss the key considerations of PCIe clocking including clock tree architecture, jitter, signalling type, EMI reduction, and stability.

PCIe in Automotive

PCI Express is a point-to-point serial interface created in 2003, originally for the computing industry. It is optimized for closed highly integrated systems such as AD/ADAS, domain controllers or zonal controllers found in today's vehicles. The PCIe standard has become one of the preferred high-speed interfaces for automotive and has a strong, well-established eco-system.

PCIe is a bidirectional bus, based on a pair of unidirectional lanes (one in each direction). Up to 16 lanes can be aggregated in parallel to increase transfer rate. The bandwidth of PCIe has doubled with each generation. The transfer rate per lane evolved from 2.5 GT/S (GTransfer/s) with a rate of 250 MB/s per lane in PCIe Gen 1, to 64 GT/s with a rate of 7.56 GB/s in PCIe Gen 6. PCIe Gen 4 is widely used in automotive systems currently. Gen 4 features 16 GT/s, with a rate of 1.97 GB/s per lane.

PCIe Clocking

Figure 1 shows an example of an automotive ECU (electronic control unit), featuring one PCIe interface. The PCIe interface requires a 100-MHz clock at each end of a bus. In this this block diagram, several SiTime MEMS oscillators are recommended including an ultra-low jitter SiT9396 oscillator at each end of the PCIe bus.

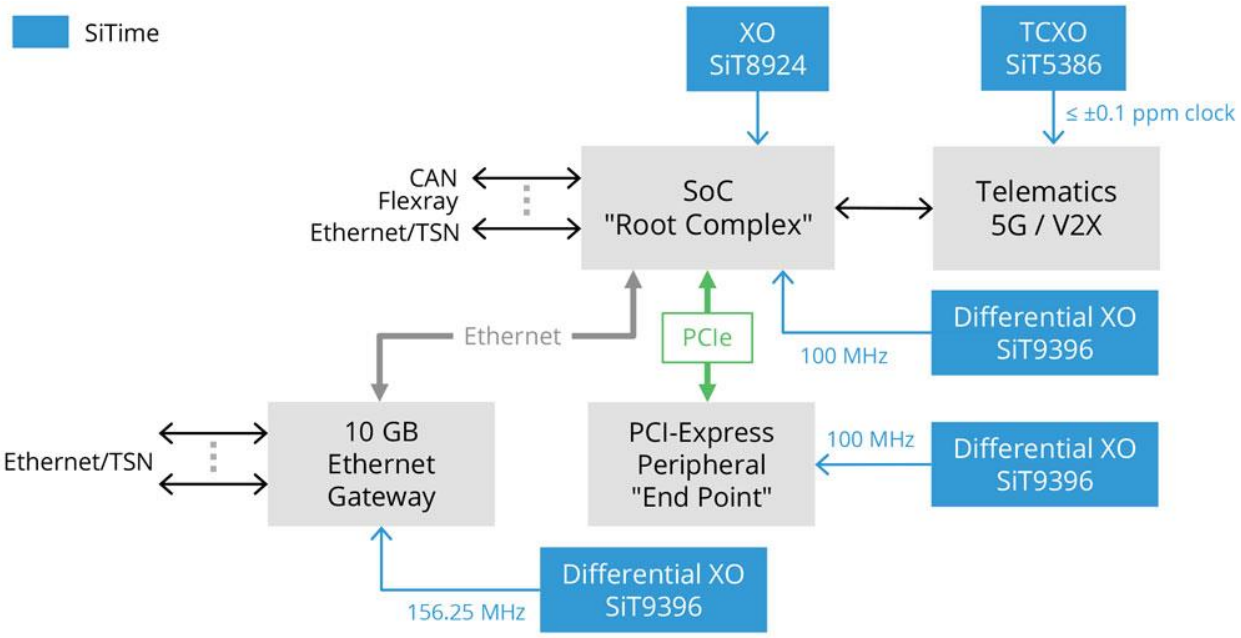

Figure 1: Block diagram of an automotive ECU (electronic control unit) with one PCIe interface

Listed below are several parameters that need to be considered in PCIe clocks. Each will be discussed in the following sections.

- Jitter, depending on the PCIe generation
- Signaling type: HCSL or LP-HCSL
- Frequency accuracy and stability
- Spread spectrum for EMI reduction
- Clock tree architecture: common clock or separate references

Jitter

Clock jitter is an important factor to consider when clocking any high-speed serial interface. Jitter is defined as a deviation from true periodicity of a presumably periodic signal.[1] In other words, jitter causes the rising and falling edges of the clock signal (which is periodic) to rise or fall too early or too late. Jitter has many causes such as thermal noise, crosstalk, or simply a poorly performing clock source.

Excessive jitter closes the eye-diagram of the PCIe link, resulting in transmission errors on the bus. As the data rate in each generation of PCIe increases, lower jitter is required as shown in Table 1. Clock sources must be chosen with a jitter level below the allowable limit for each PCIe generation.

PCIe Generation	Jitter Limit (Common Clock Architecture)
1.0	108 ps peak-to-peak
2.0	3.1 ps_{RMS}
3.0	1.0 ps_{RMS}
4.0	500 ps_{RMS}
5.0	150 ps_{RMS}
6.0	100 ps_{RMS}

Table 1: Common clock architecture jitter limits per PCIe generation

SiTime PCIe-compatible clock sources, such as the SiT9396 oscillator, support PCIe Gen 1 to Gen 6.

For more information on jitter

Refer to the Jitter Fundamentals and Clock Jitter in High-Speed Serial Links courses on the Timing Essentials Learning Hub

See application note: AN10007 Clock Jitter Definitions and Measurement Methods

Signaling Type: HCSL and LP-HCSL

The PCI Express standard, developed by PCI-SIG, originally specified HCSL signaling. HCSL outputs are driven by a current generator, which requires a 4-resistor termination to create the desired voltage at the destination stage.

Modern systems uses LP-HCSL (low-power HCSL). The outputs are voltage-driven, therefore terminations are no longer required. In addition to saving saving power, HP-HCSL saves 2 resistors per lane and clock, reducing board space and BOM cost. Furthermore, when the bus is stopped, an LP-HCSL output can be switched off entirely, bringing the power down to virtually 0 mA. In contrast, HCSL continuously consumes current, regardless of whether the bus is active or stopped.

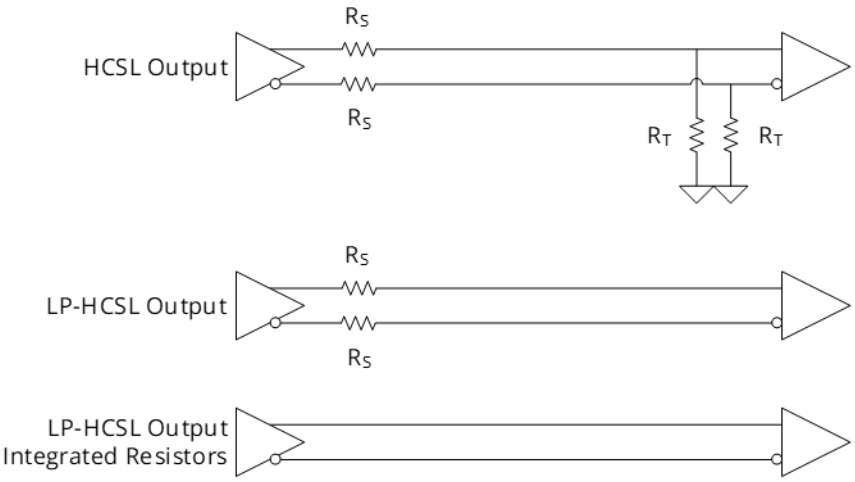

HCSL requires external resistors:
Series resistors R_S = 33 Ω
Termination resistors R_T = 50 Ω

LP-HCSL does not require termination resistors. Furthermore, series resistors R_S = 33 Ω are usually integrated on the die (bottom). SiTime devices feature on-die internal resistors.

Figure 2: HCSL, LP-HCSL, and LP-HSCL output with integrated resistors

Frequency Accuracy and Stability

The frequency accuracy of a clock is expressed in parts per million (ppm) or parts per billion (ppb) and can be divided into two parts:

- **Initial accuracy**: The deviation from the target frequency, at 25°C. Temperature effects and aging are not taken into account.
- **Stability over temperature**: The deviation from the target frequency over the entire operational temperature range.

The total accuracy of a clock includes initial accuracy, stability over temperature, and other possible effects that impact accuracy such as aging.

PCIe Gen 1 to 4 require a total accuracy of ±300 ppm. PCIe Gen 5 and above require ±100 ppm. In the remainder of this document, we will use the PCIe Gen 1 to Gen 4 requirement of ±300 ppm in our examples.

Spread Spectrum for EMI Reduction

The PCI-SIG standards foresee the use of spread spectrum clocking (SSC). The purpose of SSC is to reduce electromagnetic emissions to comply with EMC (electromagnetic compatibility) standards. SSC modulates the frequency of the 100-MHz PCIe clock with the following parameters.

- Allowable modulation frequency range (per PCI-SIG spec) is 30 kHz to 33 kHz [2]
- Downspread, meaning the modulated frequency is always lower than the 100-MHz carrier frequency, which is different than center spread that modulates symmetrically around the carrier frequency
- Max modulation amplitude of 0.5%

Refer to SiTime Application Note AN10005 for more information on spread spectrum.

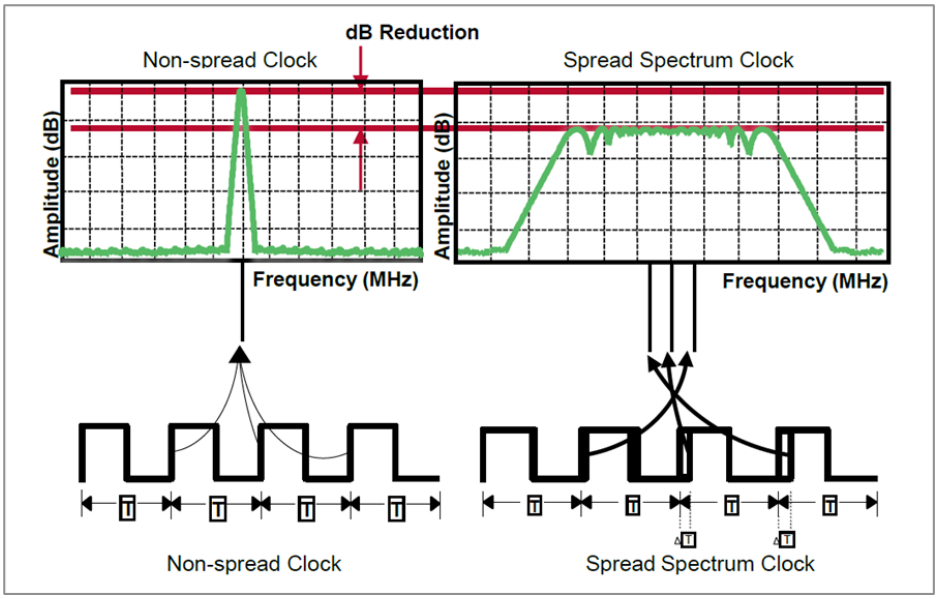

Figure 3: Non-spreading clock and a spread spectrum clock measured by a spectrum snalyzer with a 100 kHz resolution bandwidth [3]

PCIe Clock Tree Architectures

Common Clock Architecture

Common clock architecture is possibly the easiest PCIe clocking architecture and is by far the most frequently used. In this architecture, the root complex as well as each end point (there might be more than one) receive their clock signal from the same clock generator. In simple terms, a clock generator is like an oscillator with more than one output.

Because all clock signals come from the same source, spread spectrum modulation on each clock are in phase with each other. The frequency is therefore:

- 100 MHz ±300 ppm (SSC inactive)
- 100 MHz +300 ppm, down to -5300 ppm (SSC active with max possible modulation)
- 0 ppm frequency difference between all clocks

Common clock architecture is simple and elegant. One device clocks the entire PCIe system; SSC can be activated at will; and board space is reduced to a minimum.

Please contact SiTime for information on automotive clock generators.

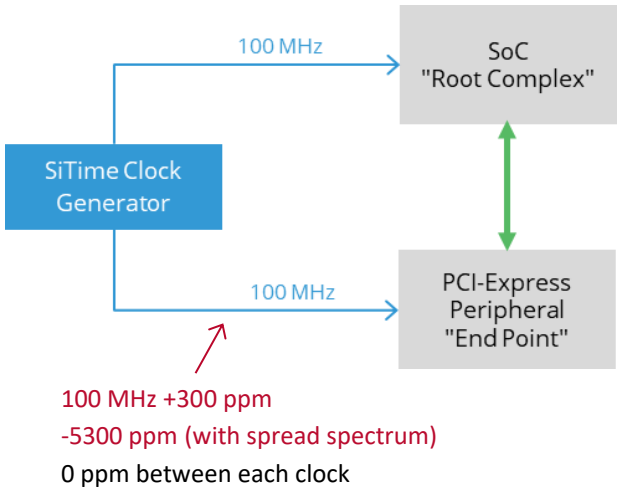

Figure 4: Common Clock Architecture

Oscillator + Buffer (Works well, but is not the best architecture)

The reason PCIe clock buffers exist is to distribute an incoming clock comiong from a main board to add-on cards or mezzanine cards in computers or servers. As automotive systems typically use PCIe within one PCB board, distributing an external clock is rarely, if ever, needed.

However, a frequent implementation of the PCIe common clock architecture consists of an oscillator, followed by a clock buffer. Several points need to be considered:

- Two components are needed instead of one, increasing the Bill-Of-Materials (BOM).
- A clock buffer increases jitter. Care must be taken that the jitter at the output of the buffer meets the jitter requirement mentioned above. In practice, this is rarely an issue, though.
- Some PCIe clock buffers have an internal PLL, whose aim is to attenuate the jitter of the incoming clock. This feature is motivated by the use case of distributing external clocks, whose jitter might have degraded over long traces or cables. PLL-based clock buffers will not pass any spread spectrum modulation. If spread spectrum is needed, be sure to choose a clock buffer without a PLL.

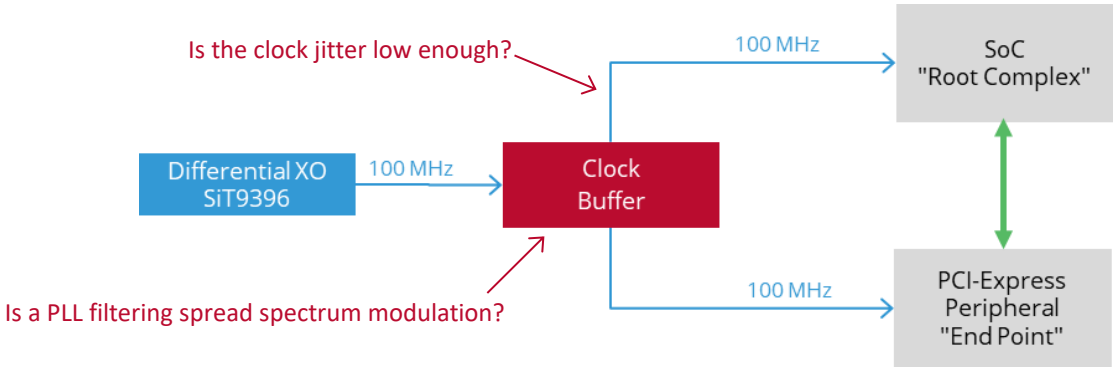

Figure 5: Oscillator + buffer architecture

This architecture works, but is not the best solution. Use of clock generators are much more elegant. In particular, SiTime clock generators which include an integrated MEMS resonator, clock generation, clock distribution and more features within one small QFN package. Please contact SiTime for information on automotive clock generators.

Separate Clock References, No Spread Spectrum (SRNS)

In the SRNS (separate reference, no spread) architecture, different clock sources, such as oscillators, are used to clock the root complex and each end point.

Per the PCI-SIG specification, the frequency of the clock must be 100 MHz with a maximum allowed deviation of ±300 ppm (including initial accuracy + stability over temperature) or ±100 ppm (PCIe Gen 5 and above).

For example, with two oscillators (PCIe Gen 1 to Gen 4), the maximum allowable frequency deviation is ±600 ppm, respectively ±300 ppm for each oscillator.

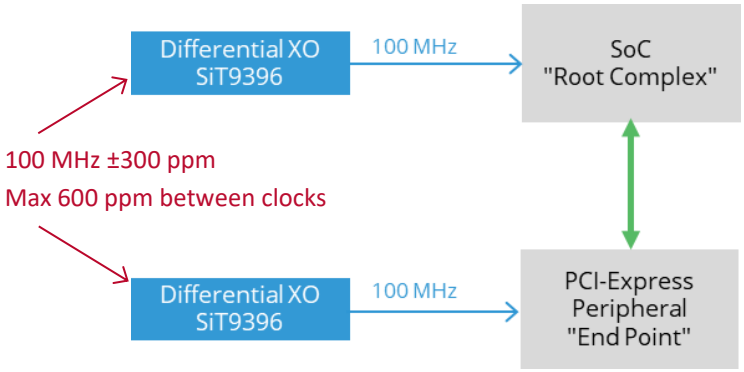

Figure 6: SRNS architecture

This is the architecture of choice when oscillators are preferred to clock generators in the design. With low jitter and a stability of ±50 ppm over -40°C to +125°C, SiTime SiT9396 support SRNS with PCIe Gen 1 to Gen 6.

Separate Clock References, with Independent Spread Spectrum (SRIS)

Beginning with PCIe Gen 4, the PCI-SIG specification includes the provision for SRIS (separate reference, independent spread). The main reason for introducing this architecture is to avoid having a cable for the clock, when end points are distant from the root complex. In most automotive systems however, root complex and end points are on the same board.

In this architecture, using different clock sources (typically oscillators) can enable spread spectrum at the same time. The max allowed frequency difference is therefore ±5600 ppm. PCIe transmitters insert dummy data into the link to deal with the difference in clock frequency (in other words, link speed) between transmitter and receiver.

Note that either the root complex, or the end point, can have SSC active, but not both. The SRIS architecture introduces additional constraints on clock jitter compared to common clock architecture. SiTime devices support the SRIS architecture.

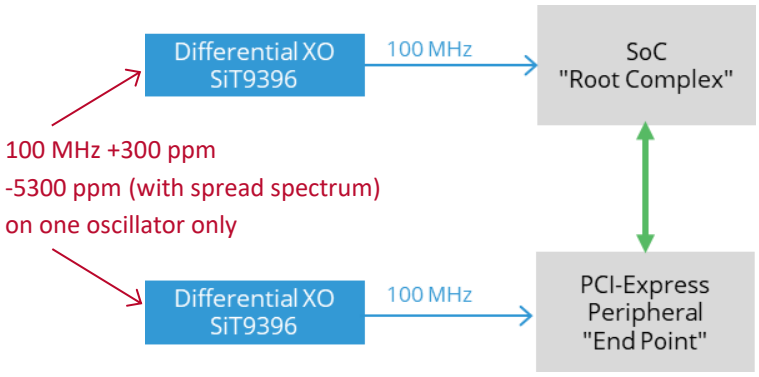

Figure 7: SRIS architecture

SiTime Advantages

SiTime timing solutions support PCIe Gen 1 to 6 and support common clock, SRNS, and SRIS architectures, while offering several advantages that are particularly important for automotive and functional safety applications. They provide higher reliability, better robustness, and small footprints.

- SiTime clock generators feature an internal MEMS resonator and do not require an external reference, removing the "weak link" represented by crystal resonators.
- Excellent reliability (< 0.5 FIT, < 0.1 DPPM) – up to 50x more reliable than crystal-based devices.
- Up to 10x better resilience to shock and vibration than crystal-based devices. Shock and particularly vibration in crystal resonators can increase the jitter of generated clocks, therefore increasing the bit error rate (BER) on a PCIe link. Shock and vibration in crystals can also create frequency micro jumps and activity dips, which degrade BER. SiTime devices are free of these problems.

For more information

See application webpage: PCI Express in Automotive Applications

References

[1] Wikipedia: https://en.wikipedia.org/wiki/Jitter

[2] PCI-SIG specifications: https://pcisig.com/specifications

[3] AN10005 Spread Spectrum Clock Oscillators: https://www.sitime.com/support/resource-library/application-notes/an10005-spread-spectrum-clock-oscillators

SiTime Corporation | 5451 Patrick Henry Drive, Santa Clara, CA – USA | Phone: +1-408-328-4400