The Importance of Jitter in High Performance Design

High Performance Serial Data Architectures bring the focus of clock oscillator jitter to the forefront

BY Dave Kenny, Vice President Research and Development, EE

Peter Rostotskyy, Senior Engineer, EE

Craig Taylor, Subject Matter Expert, EE

INTRODUCTION/OVERVIEW

High-speed serial bus architectures are the norm in today's high-performance designs. While parallel bus standards are undergoing some changes, serial buses are established across multiple markets and devices – computers, cell phones, entertainment systems, and more. Serial buses offer performance advantages and design simplification (fewer traces) in circuit and board layouts. Serial data links behave as arteries of today's informed world as they transfer data from one point to the other in a processing system. To ensure accurate delivery and receipt thereof, data in a digital system is governed by Clock and Data Recovery (CDR) circuitry, which then behaves as the handshake in the data system. The key to an accurate receipt and interpretation of data is specifically related to knowing precisely 'where' a clock edge is in any point in time.

Since the transmitting and receiving devices can be in any place -- from the same desktop to the other side of the world, there are influences in each distinct location or environment that can affect how a clock edge drifts from the time of sending data to the time a device receives and interprets the data. These influences are many and include temperature, physical movement/vibration, and even the architecture of clock signal origination. The net outcome is either having accurate data or not, and "not" is obviously never an option in any system. For end users this can mean a poor experiential quality and disruption to internet sessions and related services (be that poor voice quality, an uneven viewing experience of video content, or corrupted data file contents). The performance characteristic that is a measure of exactly how precise the clock edge *is* vs where it is *expected to be* is called 'Jitter'. There are three quantifications of jitter normally used in its measurement:

1. phase/RMS jitter—which can be considered as a 'fine-focus' measurement. This is commonly referred to as 'Absolute Jitter' which is the total and entire difference in the position of a clock's edge from where it would ideally be—usually revealed by measuring the phase noise of the signal with a network analyzer (Fig A);

2. peak jitter, and peak-to-peak jitter each of which can be thought of as a 'course' measurement and are broken down into two characteristics:

    a. Period jitter (aka cycle jitter) The difference between any one clock period and the ideal or average clock period—usually revealed by measuring the signal period with an oscilloscope (Fig B), and

    b. Cycle-to-cycle jitter -- The difference in duration of any two adjacent clock periods. It can be important for some types of clock generation circuitry used in microprocessors and RAM interfaces and also measured with an oscilloscope (Fig C)

19013 36th Ave. W Lynnwood, WA 98036 www.pletronics.com

The jitter performance/specification limits have been determined by standardization bodies like ITU-T, Telcordia and IEEE. The specifications and test methodologies for jitter on native Ethernet (IEEE) differ from those for SDH/ SONET/SyncE (ITU-T, Telcordia).

Phase jitter

Deviation from Ideal edge

Figure A

Period jitter

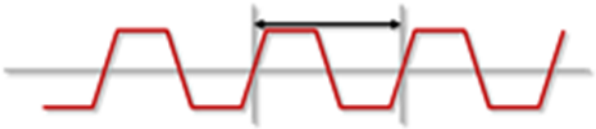

Deviation from Ideal period

Figure B

Cycle-to-cycle jitter

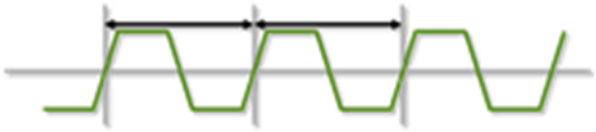

Deviation from zero difference

Figure C

JITTER GENERATION

As data rates for next-generation serial standards increase, analog anomalies have a greater impact on signal integrity and quality than ever before. Conductors in signal pathways, including circuit board traces, vias, connectors, and cabling, exhibit greater transmission line effects with return losses and reflections that degrade signal levels, induce skewing, and add noise---and therefore jitter. Everything however starts with the base System Clock signal (SYSCLK or Master Clock). Along with significant performance characteristics of a clock signal, the cost of creating the signal can vary significantly—over 10x—depending upon the architecture and design approach used. To aid in achieving a system design without excessive performance guard-bands (and therefore excessive cost), this article focuses on providing an update as to the different architectures used to create the clock signal compliant to each specific High-Speed Serial Data (HSSD) implementation scheme. Specific jitter types, definitions, and compliance testing methods are already well documented and will not be repeated here.

19013 36th Ave. W Lynnwood, WA 98036 www.pletronics.com

The prevalent foundation component that is utilized to create system clocks is a quartz crystal oscillator("XO"), a proven technology that has been used for years. Crystal oscillators themselves have an inherent jitter characteristic and their output jitter will vary based on the design/circuitry as will their unit price. Smart systems designers realize that the total cost of the system/product/design is itself a 'specification' that needs to be met. This article describes each method for signal creation together with a recommendation table to help potential users from incurring a higher component cost than necessary.

SYSCLK ORIGINATION METHODS

The basic 'no frills' crystal oscillator utilizes a quartz crystal and employs it together with simple circuitry to run on the crystal's fundamental mode and create a square wave output. This architecture provides the best performance for both peak-to-peak and RMS jitter and is generally the most cost-effective at frequencies up to 50Mhz. To reach higher frequencies with the absolute lowest jitter possible, a technology referred to as High Frequency Fundamental ("HFF") is used. A crystal can be made to vibrate at one of its overtone modes, which occur near odd-numbered multiples of the fundamental resonant frequency. Such a crystal is referred to as a 3rd, 5th, 7^{th}…etc overtone ("OT") crystal. To accomplish this, the oscillator circuit usually includes additional design elements to select the desired overtone. Relatedly, an architecture that makes the crystal operate on its 3^{rd} overtone to reach frequencies as high as 3x50Mhz = 150Mhz can be efficiently executed in a typical application.

Operation on higher overtones requires far more complicated circuitry, and several oscillator companies are working to increase quartz crystal Fundamental and Third Overtone resonance technology to support, for example, 10Gb Fibre Channel at 70.8333Mhz x 3 = 212.500Mhz. These efforts are focused on providing the lowest jitter Master Clock performance -- needed as data bus speeds continue to increase. Nonetheless, this technology is still in an advanced stage and not readily available from all crystal oscillator suppliers.

Another technology that has been successfully utilized is that of integrated integer multipliers. In these devices, the frequency is increased by locking the incoming signal to an integrated Voltage Control Oscillator running at a direct, integer multiplication (2x, 3x, 4x…etc) of the crystal frequency and then binarily divided back down to the required operating frequency. An additional method that can be employed is harmonic multiplication. This is similar in technique to crystal overtone utilization with the difference being the output signal from the crystal oscillator (not the crystal) is multiplied by an integer value. In addition to losses and other compromises made in circuit integration, the jitter performance is worsened by a factor of 20LogN (where N is the integer multiplication factor) over direct (i.e. crystal fundamental of overtone mode) frequency generation.

Thus, while fundamental, overtone, and/or harmonic frequency generation are possible, these technologies are often cost- and complexity-prohibitive compared to integrated integer multiplication which can meet, albeit with less margin, the jitter requirements. To avoid any unnecessary cost premiums, what a designer needs to concentrate upon during design margining is the specific bandwidth over which the output signal jitter is calculated.

A third technology that is utilized is referred to as integrated 'fractional N' multipliers. This is where the frequency of the input signal can be converted to practically any other frequency—integer related or not. For example, a 25Mhz crystal frequency could be converted to 644.53125Mhz by a fractional multiplication of 25.78125. For reasons that exceed the intended purpose and depth of this paper, this results in the highest amount of signal jitter. Just the same, it can be sufficient for some systems and is by far the most cost effective to employ at frequencies above 215Mhz.

CRYSTAL OSCILLATOR OUTPUT LOGIC

In the prior section, the method for *generating* the CLK and how it affects jitter performance was discussed. Regardless of whether the implementation architecture is Fundamental or Overtone Crystal, N-Multiplier or Fractional-N multiplier, the crystal oscillator also contains output drivers pursuant to existing logic technologies. The specific type of output logic compatibility can be Low Voltage CMOS (LVCMOS), Low Voltage, Positive Supply Emitter-Coupled Logic (LVPECL), Low Voltage Differential Signaling (LVDS), and/or High Speed Current Steering Logic (HCSL). The output logic type is primarily related to the output frequency and/or commonality of logic interface for the processing device(s) within a given application type. For example, the primary logic type for PCIe SYSCLKs is HCSL.
Crystal oscillator output logic compatibility typically lags development of processing device logic by 6 to 12 months, and sometimes even longer. Logic translators are used in the interim. An example of this is Transition-Minimized Differential Signaling (TMDS). TMDS is employed in some applications in a system design (ex: HDMI) but is presently not available as a choice for crystal oscillator output logic. The significance in output logic type is in the contribution of the 'interface jitter' introduced by connecting the crystal oscillator (and any additional output translation devices) to the processing device. Generally, the logic type with the fastest transition time (e.g. Rise/Fall time) through 'eye pattern' will result in the lowest interface jitter.

SELECTING THE OPTIMUM CLK ORIGINATION DEVICE

Regardless of performance specification, specification requirements, or specific PHY chipset/method of execution utilized, the most important specification of all is cost effectiveness of the implementation. All commercial and industrial system meeting all performance requirements but having a total cost above market requirements gives no value.
All quality crystal oscillator providers include an amount of guard-band in the jitter generation specification that they publish in their datasheets. As, with good reason, system designers also include an amount of guard-band in the specifications they request, doing so with a reputable crystal oscillator manufacturer may result in double guard-banding and therefore an undue higher cost of solution. To aid in the specifying of the proper crystal oscillator without adding excessive guard-band and cost, TABLE 1 shows the most popular data/communication applications of today.

At Pletronics, we offer solutions containing each of the technologies mentioned: high-frequency crystal fundamental, overtone, integer-N and fractional-N. Each product is executed to provide the Customer with a solution that is most cost- and performance-effective. Table 1 is a list of the technology utilized for each of the most popular serial data conventions of today. Table 2 contains the specific part number that can be called out on the system design Bill of Material for each. As with any Company that is market-leading, Pletronics products are continually evolving in an effort to enable our Customers to maintain a competitive edge.

TABLE 2 – PLETRONICS PARTNUMBER

Protocol	Purpose	Supply Voltage	Output Logic	PLE Partnumber	Operating Temp Range	Frequency Stability
5ppMPCIE Gen4	Chip-to-chip and module inter-connect	3.3V[4]	HCSL	HC5544DEV-100.0M	-40/+85C	±25ppM
222510G EPON	Systems Inter-connect		LVCMOS	SM5545JEV-161.1328M	-40/+85C	±50ppM[5]
10GbEthernet XAUI			LVDS	LV5545JEV-156.250M	-40/+85C	±50ppM
100GbEthernet (10X10)			LVCMOS, LVDS	SM5545JEV-161.1328M LV5545JEV-161.1328M	-40/+85C	±50ppM
100GbE (4x25Gb)			LVDS, LVPECL[2]	LC5545DEV-644.53125M PC5545DEV-644.53125M	-40/+85C	±50ppM
1Gb Fibre Chan-nel			LVDS	LV5545JEV-106.250M	-40/+85C	±50ppM
10Gb Fibre Channel			LVDS	LV5545JEV-212.5M LV5545JEV-159.375M	-40/+85C	±50ppM
SONET OC48	Long-Haul/Inter-Continental Transport		LVCMOS LVDS LVPECL[2]	SM5545JEV-155.520M LV5545JEV-155.520M PE5545JEV-155.520M	-40/+85C	±50ppM
SONET OC192			LVDS LVPECL[2]	LC5545DEV-622.080M PC5545DEV-622.080M	-40/+85C	±50ppM
SAS ATA Gen 3	Attached Stor-age Access		LVDS	LV5545JEV-150.0M	-40/+85C	±50ppM
SynchE3 Linecard	*Highspeed Data 1G:*		LVCMOS	OeD4281-25.0M	-40/+85C	±4.6ppM
	(1000BASE)10G (10GBASE 40GBASE 100GBASE):					
	10G (10GBASE):					
	25G (100GBASE):					

[4]Other supply voltages available
[5]25ppM frequency stability is available

19013 36th Ave. W Lynnwood, WA 98036 www.pletronics.com