AN1151: Using the Si539x in 56G SerDes Applications

Ethernet speeds are continuing to increase, pushing to 400G and beyond due to demands for more data delivered at faster speeds. The latest switch processor chips with 56G PAM-4 interfaces have allowed for new high-speed interconnects in the cloud and data center networks. The challenge is to make the data pipeline faster while ensuring the networks remain reliable, easy to maintain, and cost effective. This application note strives to explain some of the design challenges and changes along with the jitter requirements for the timing chip in the data path. This application note also explains why the Si539x family and particularly the Si5395P is ideal for 56G and future 112G SerDes designs.

KEY FEATURES

- PAM-4 Modulation is driving the need for high performance timing
- Standards are given for different trace/cable lengths
- Total Jitter performance and reference clock jitter is calculated for different cables/trace standards
- Jitter performance of the Silicon Labs Si5395P Jitter Attenuator is highlighted

1. Example Ethernet Interfaces

Current 40G to 100G Ethernet are commonly using 10G and 25G/28G SerDes, which provide 10 Gbps up to 25 Gbps per lane throughput. In some cases, forward error correction is included, which adds the extra overhead. 100 Gb to 400 Gb+ Ethernet switches are entering the market and starting to use 56G SerDes Phys to meet the requirements for the industry transition to 400 GbE. Eventually, these SerDes Phys will need to scale further, as discussions about 112 Gbps lane rates have started. Supporting faster Ethernet speeds is a challenge for SerDes designers as the industry transitions from 28 Gbps to 56 Gbps and eventually to 112 Gbps per single lane. The table below lists some examples of the standards for different 100 Gb to 400 Gb Ethernet interfaces.

Table 1.1. 100–400 Gb Ethernet Interface Standards

Name	Medium	Media Count/Lanes	Gigabaud per Lane
100GBASE-SR10	Multi-mode fiber, 850 nm	10	10.3125
100GBASE-KP4	Copper backplane	4	25.78125, RS-FEC
100GBASE-KR2	Copper backplane	2	26.5625 using PAM4 modulation
100GBASE-CLR4 MSA	Single-mode fiber, WDM: 1271 nm, 1291 nm, 1311 nm, 1331 nm	4	25.78125, optional RS-FEC
100GBASE-DR	Single-mode fiber, 1304.5 to 1317.5 nm	1	53.125 using PAM4 modulation
200GBASE-CR4	Twin axial cable	4	26.5625 (PAM4 modulation)
200GBASE-LR4	Single mode fiber WDM 1295.56-1709.14 nm	4	26.5625 (PAM4 modulation)
400GBASE-SR16	Multimode fiber 850 nm laser	16	26.5625
400GBASE-DR4	Single mode fiber 1304.5 to 1317.5 nm	4	53.125 (PAM4 modulation)
400GBASE-FR8	Single mode fiber WDM 1273.54-1309.14 nm	8	26.5625 (PAM4 modulation)
400GBASE-LR8	Single mode fiber WDM 1273.54-1309.14 nm	8	26.5625 (PAM4 modulation)

High speed 56G+ SerDes links are also important in 5G infrastructure. They will play an important role in the development of wireless networking infrastructure while optimizing tradeoffs between serial link speeds and operational efficiency and performance.

2. Higher Order Modulation is Driving the Need for High-Performance Timing Solutions

There is a need to send data faster, so what are the techniques to accomplish that? What technique is traditionally used to send data through a channel?

In a typical high-speed serial link, data is transmitted from the transmitter to the receiver through a channel. This channel can be as long as 1 m for a backplane and 5 m for copper cable channels. Along the way, signal integrity can be affected by jitter, intra-pair skews, frequency dependent attenuation, ISI, reflections, crosstalk, etc. There are different ways of modulating the data to send it from one point to another. There are tradeoffs associated with different schemes depending on the interface and the distance the data must travel.

2.1 NRZ PAM-2 Modulation

The traditional method of modulation is called Non-Return to Zero (NRZ), which is also called PAM-2 (Pulse Amplitude Modulation 2). It has two different voltage levels to represent a zero and a one. The voltage level remains constant through the bit interval. The symbol is equal to the bit, and there is one eye in each unit interval. For serial data at a rate of 56 Gbps, 1 UI = 1/56e9 = 17.857 ps. The Nyquist frequency is 56 Gbps/2 = 28 GHz.

2.2 PAM-4 Modulation

Pulse Amplitude Modulation-4 (PAM-4) is a four-level modulation scheme used in high speed systems, particularly for signaling speeds greater than 26 Gbps. PAM-4 encodes two bits into 1 symbol. There are, therefore, four signal levels, as two bits have four unique combinations. Compared to the binary modulation NRZ, PAM-4 achieves the same data rate with half the bandwidth. It splits the eye diagram into four levels and prevents having to increase the bandwidth of RX/TX channels. The trade-off is that the signal-to-noise ratio degrades to keep the same bandwidth. This requires tighter clock jitter requirements for this higher throughput. The Nyquist frequency for PAM-4 is 56 Gbps/4 = 14 GHz. The eye height for PAM-4 is 1/3 of NRZ and, therefore, the SNR loss is at least 20 x log10 (1/3) = 9.5 dB. Thus, PAM-4 trades signal to noise ratio for bandwidth.

The figure below shows a comparison of NRZ vs PAM-4. NRZ, also referred to as PAM-2, has two amplitude levels (0 or 1). There is one bit of information in every symbol. PAM-4 contains four amplitude levels (0, 1, 2, 3) with two bits of information in every symbol and, therefore, two times the througput for the same baud rate. PAM-4 has a lower SNR, as it is more succeptible to noise. The diagram shows a comparison of the different encoded signal levels and the corresponding eye diagrams for each.

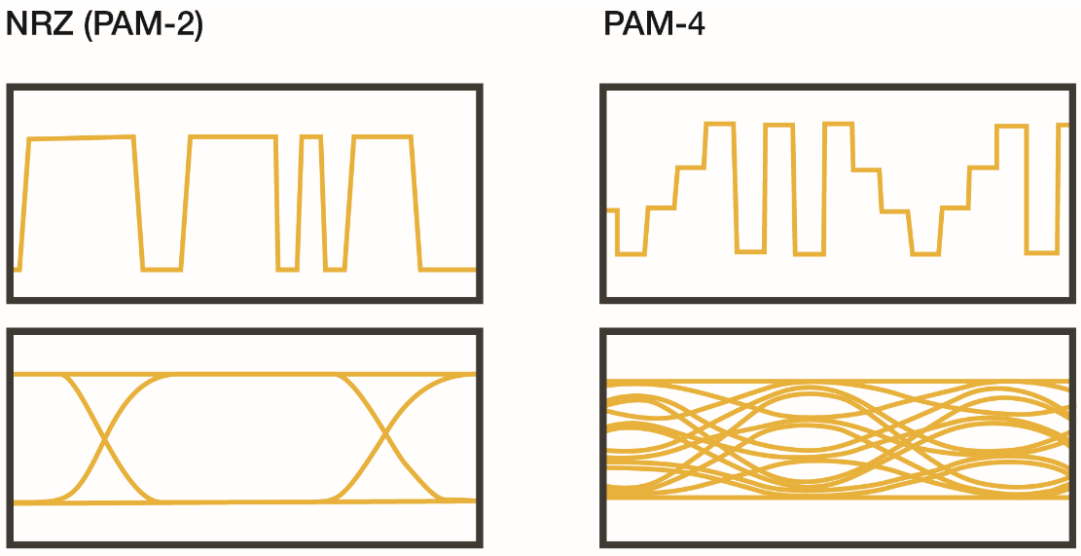

Figure 2.1. NRZ Compared to PAM-4 Signals

The figure below shows the encoding sequence bit values where the NRZ “A” and NRZ “B” sequences of data can be combined. Sequence C is a PAM-4 modulation scheme, which contains both the top two NRZ data sequences sent at the same bandwidth of 28 GHz. This is done by using four amplitude levels. PAM 4 signaling has 12 symbol transitions, six different rise times, six different fall times, and three eye diagrams.

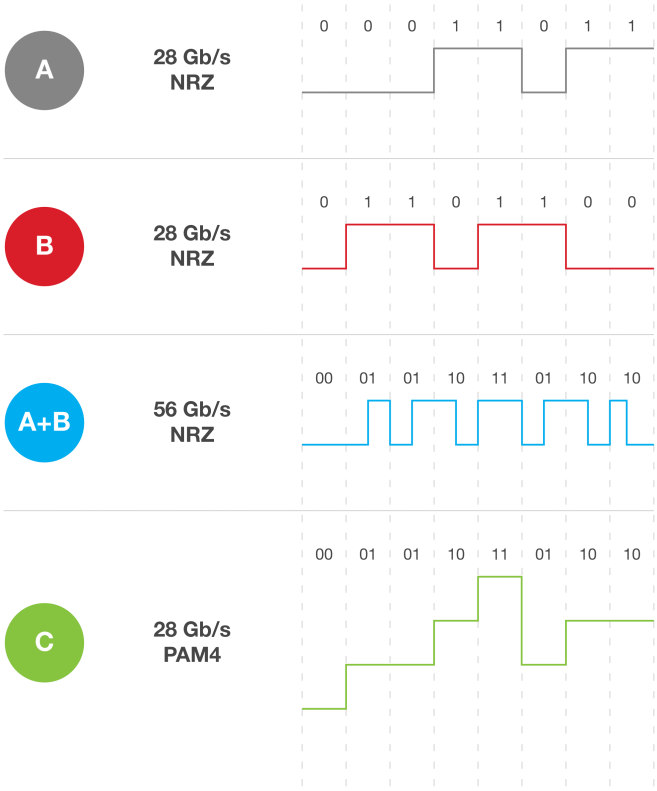

Figure 2.2. Signal Levels and Encoding of SNR vs. PAM-4

2.3 Forward Error Correction

Forward error correction (FEC) is used to guarantee end-to-end error free performance for 25G, 50G, and 100G lane rates. This ultimately helps relax these demanding requirements and, therefore, the reference clock jitter requirements coming from the jitter attenuator. Forward error correction requires more data to be sent, which is used for error detection and correction purposes only. This data adds some overhead. The forward error correction is part of the standards and helps with the degrading SNR from PAM-4 modulation. As the speed increases to 100G lane rates (up to 112G including the FEC overhead), a low jitter reference clock combined with forward error correction becomes invaluable for data reliability.

3. Standards for Different Trace and Cable Lengths

In 2016, an IEEE 802.3 task force developed a single-lane 50 Gbps Ethernet standard in 802.3bs to be used in 400 Gbps. The 802.3bs standard requires a PAM-4 encoding scheme, as described above, to reach 50 Gbps per channel throughput. There are several options for delivering 400 Gbps based on the distance requirements and multimode fiber vs. single-mode fiber.

The Optical Internetworking Forum also has standards published for 56G PAM-4 SerDes for each of the different trace and cable lengths. The goal of these committees is to promote the development of implementation agreements (IAs) for optical networking products and component technologies including SerDes. This creates a universal compatibility in the marketplace among products. The standards published specify 56G per channel, and PAM-4 modulation is used with forward error correction. The standards for each of the different trace and cable lengths specify the maximum data rate, the total jitter, and the integration band.

OIF-CEI-56G-PAM4 stands for Optical Internetworking Forum Common Electrical I/O for 56G PAM-4. The figure below shows the Long Range (LR) connection, Very Short Range (VSR) with a Chip to module trace, and a MR medium range with chip to chip trace. These are examples of the different interfaces seen in the standards.

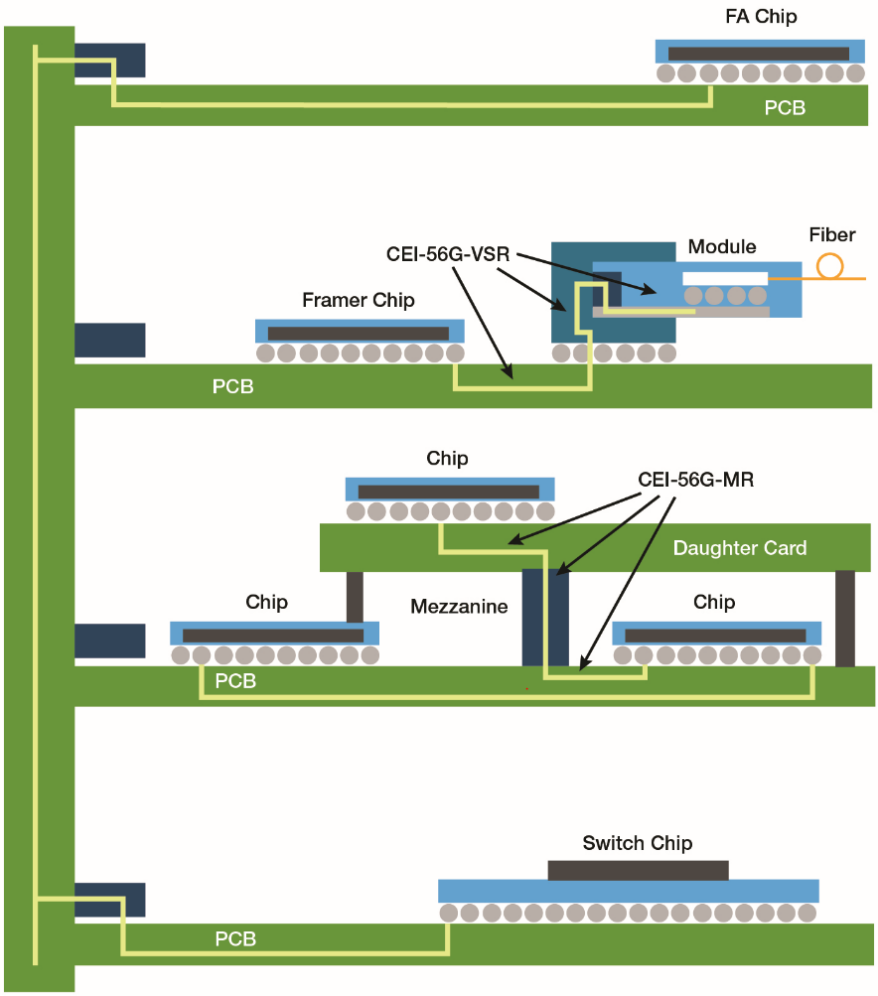

Figure 3.1. Different SerDes Interconnects of Varying Trace Lengths[1]

The Very Short Reach Chip to module specification is <10 cm, one connector with a loss of up to 10 dB with pre-FEC BER < 1^{-6}. The figure below is a diagram from Figure 16-15 out of the OIF-CEI-04.0 specification describing the expected losses between the connections.

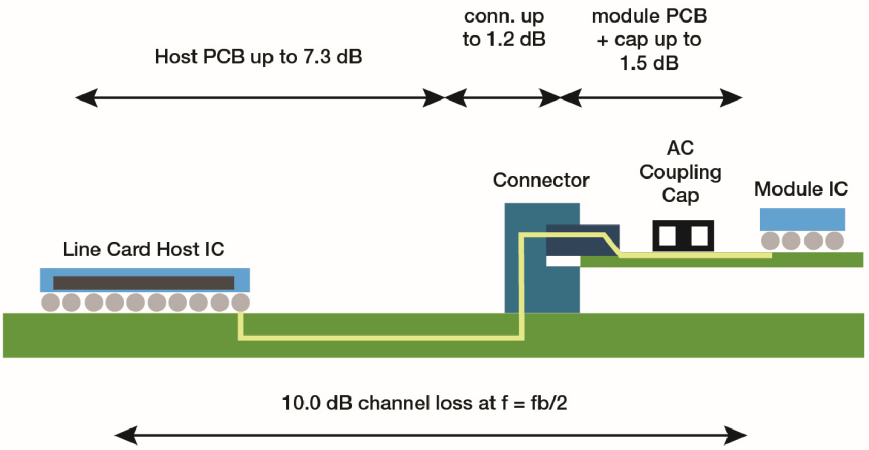

Figure 3.2. VEI-56G-VSR-PAM4 Channel Reference Model[1]

- Long Reach has backplanes or copper cables and two connectors with loss up to 30 dB and pre-FEC raw BER $<1^{-4}$.
- Medium Reach: Chip-to-chip for mid-range backplanes with less than 50 cm one connector, has loss up to 20 dB and pre-FEC BER $<1e^{-6}$.

The table below is a summary of the standard, the max data rate, the spec given from the standard along with the translated jitter requirement for the reference clock. This is assuming a 20% budget for the reference clock which works out in the calculation to ~90% of the total jitter remaining for the SerDes. The integration band is also given assuming a common Tx bandwidth of < $f_{ref}/10$. An example calculation is done in the next section showing how the conclusion of 90% of the total jitter is budgeted for the transmitter. The standards specify the requirements differently based on the trace/cable length. In some cases, they are specified based on T_J RMS in UI RMS. In some cases, the eye width and bit error rate is given, and, in other cases, the eye height, transmitter, and dispersion eye closure are given. All of these specifications have been converted to the reference clock jitter spec for the jitter attenuator with the integration band at which it should be measured.

The Si5395P performance, shown in the right hand column of the table below, provides significant margin to the requirements in the standards. The typical Si5395P jitter measured from the Keysight 5052 in the 12kHz-20MHz range is 69 fs, while jitter in the 4MHz to 20MHz integration band is 38 fs. The phase noise plots are shown in Figure 4.2 Phase Noise Plot Si5395P RMS Jitter 12 kHz–20 MHz Integration Band on page 10 and Figure 4.3 Phase Noise Plot Si5395P RMS Jitter 4 MHz–20 MHz Integration Band on page 10. It is important to realize that the reference clock phase noise beyond Nyquist (fref/2) is filtered by the receiver reference clock input path and aliased by the transmitting PLL phase detector input. This filtering and aliasing combined with the reference clock phase noise roll off at higher frequencies can be approximated by RMS addition of three to four times the reference clock noise floor, folding back into the integration band. The typical phase noise integration by the Keysight 5052 provides overly optimistic jitter numbers compared to what is actually seen at the interface due to these effects. The Keysight 5052 measures out to 40 MHz, but it is important to also consider the phase noise beyond that, assuming a straight line continuing from 40 MHz to 550 MHz. It is assumed that at 550 MHz the signal rolls off and can be considered insignificant. When folding the approximated (flat) phase noise back into the integration band folds of 1.5F to 2.0F, 2.0F to 2.5F, 2.5F to 3.0F and then 3.0F to 3.5F are added, which is the noise floor four times. In the case of a 156.25 MHz clock frequency, 3.5F is approximately 550MHz. Since the estimated phase noise is flat, the integral is easy to calculate and the conversion of phase noise to jitter for white (flat) noise is given in the following equation.

$$Jitter\_RMS = \frac{\sqrt{2 \times \Delta F \times 10^{(\text{Phase Noise} / 10)}}}{2 \times \pi \times f}$$

From the equation above: "f" is the carrier frequency of 156.25MHz, "ΔF" is 16 MHz (20 MHz to 4 MHz) and this is the integration band. "Phase_Noise" is the phase noise floor from the Si5395P plot using the Keysight 5052 in Figure 4.3 Phase Noise Plot Si5395P RMS Jitter 4 MHz–20 MHz Integration Band on page 10, which is approximately -165dBc .

We need to add this in 4 times so there is a factor of 4 to compute the aliased jitter, which is provided in the following equation using the actual values.

$$Aliased\_Jitter = \frac{\sqrt{4 \times 2 \times 16E6 \times 10^{(-16.5)}}}{982E6} = 65 \text{ fs}$$

The total jitter is the inclusion of both the calculated aliased jitter and the measured jitter from the Keysight 5052.

For the integration band from 4 MHz to 20 MHz, the total jitter is sqrt ($38^2 + 65^2$) = 76 fs. Adding 3dB margin to the phase noise is 107 fs RMS total sampling jitter. The integration band from 12 kHz to 20 MHz following the same analysis above, the total jitter is sqrt ($69^2 + 73^2$) = 100 fs. Adding 3 dB margin works out to 140 fs RMS total sampling jitter max.

Table 3.1. Standards for Different Cable/Trace Lengths Comparing Specifications for Jitter and Integration Band

Standard	Max Rate	Specification	REFCLK Jitter	Integration Band	Si5395P[2] Jitter 12 kHz–20 MHz (69 fs typ) ---------- 4 MHz–20 MHz3 (38 fs typ)
CEI-56G-LR-PAM4	58 Gbps	0.023 UI TX RMS	350 fs RMS	4 MHz – TX BW	140 fs (RMS total sampling jitter + 3 dB margin) ---------- 107 fs (RMS total sampling jitter + 3 dB margin)
CEI-56G-MR-PAM4	58 Gbps	0.023 UI TX RMS	350 fs RMS	4 MHz – TX BW	140 fs (RMS total sampling jitter + 3 dB margin) ---------- 107 fs (RMS total sampling jitter + 3 dB margin)
CEI-56G-VSR-PAM4	58 Gbps	0.265 UI Eye Width	240 fs RMS	4 MHz – TX BW	140 fs (RMS total sampling jitter + 3 dB margin) ---------- 107 fs (RMS total sampling jitter + 3 dB margin)
400GBASE-LR8 Fiber	53 Gbps	3.3 dB TDECQ	250 fs RMS	4 MHz – TX BW	140 fs (RMS total sampling jitter + 3 dB margin) ---------- 107 fs (RMS total sampling jitter + 3 dB margin)
400GAUI-8 C2Chip	53 Gbps	0.023 UI TX RMS	380 fs RMS	4 MHz – TX BW	140 fs (RMS total sampling jitter + 3 dB margin) ---------- 107 fs (RMS total sampling jitter + 3 dB margin)
400GAUI-8 C2Module	53 Gbps	0.22 UI Eye Width	275 fs RMS	4 MHz – TX BW	140 fs (RMS total sampling jitter + 3 dB margin) ---------- 107 fs (RMS total sampling jitter + 3 dB margin)

Multiple switch/ASIC vendors with 56G PAM-4 SerDes require 150 fs RMS max integrated over 12 k–20 MHz. It is important to understand whether this includes the aliased jitter analysis from above or just the measured jitter from the phase noise anlyzer. The Si5395P is below half of that value at 69 fs measured on the phase noise analyzer from 12 kHz - 20 MHz and 38 fs jitter integrated over 4 MHz to 20 MHz. This means Si5395P is > 6x to 9x better than the specification, assuming this doesn't include the aliasing calculation. With the inclusion of aliased jitter the Si5395P max jitter spec is 140 fs integrated from 12 k-20 MHz and 107 fs when integrated from 4 MHz to 20 MHz, which is still under that 150 fs spec. The Si5395P provides a large amount of margin to ensure the performance in a SerDes device based on any of the different combinations of trace and cable lengths.

3.1 Example Calculation of Reference Clock Jitter from Standards

The reference clock jitter was calculated using the standards for each line item in XREF Table 2. For example, the CEI-56G-LR-PAM4 standard specifies a maximum data rate of 58 Gbps. The uncorrelated jitter RMS (standard deviation of the probability distribution) is 0.023 UI TX RMS. This translates to 800 fs of jitter.

The spec shows 0.023 unit intervals. One unit interval is 1/ (half the max bit rate). The bit rate is given at 58 Gbps, which is measured at the output of the transmitter. Therefore, 0.023/29G = 793 fs, rounded to 800 fs for simplicity. Since the transmitter output jitter is a combination of the reference clock jitter and the transmitter's intrinsic jitter, these are broken up into their components to compute the allowable transmitter jitter. The reference clock jitter and the transmitter jitter are uncorrelated. Therefore, Equation 1 is used to relate the reference clock jitter and the transmitter jitter, assuming that the jitter is random.

$$T_jRefclk = \sqrt{Total^2 - T_{jtx}^2}$$

Equation 1. Uncorrelated RMS Jitter Equation from the Standards

The total is 800 fs. If 20% is budgeted for the reference clock, then, by rearranging the equation as shown below, the reference clock and transmitter jitter is calculated:

$$Tj\ \text{Reflck}^2 + T_{jtx}^2 = Total^2$$

$$Tj\ \text{Reflck}^2 + 4 x Tj\ \text{Reflck}^2 = Total^2 = 800^2$$

$$Tj\ \text{Reflck} = \frac{800}{\sqrt{5}} = 358\ fs$$

Equation 2. RMS Equation Calculations to Compute Budget Requirements

The assumption is that 20% is budgeted for the reference clock and 80% for transmitter. This works out to 350 fs for the reference clock, leaving 700 fs for the transmitter. Therefore, there is approximately 90% (700/800) x 100% fs of the total 800 fs budget for the transmitter. This is how the 350 fs number has been computed in XREF Table 2.

The integration band is calculated based on the receiver jitter tolerance mask from the OIF-CEI 0.40, plotted in Figure 5.

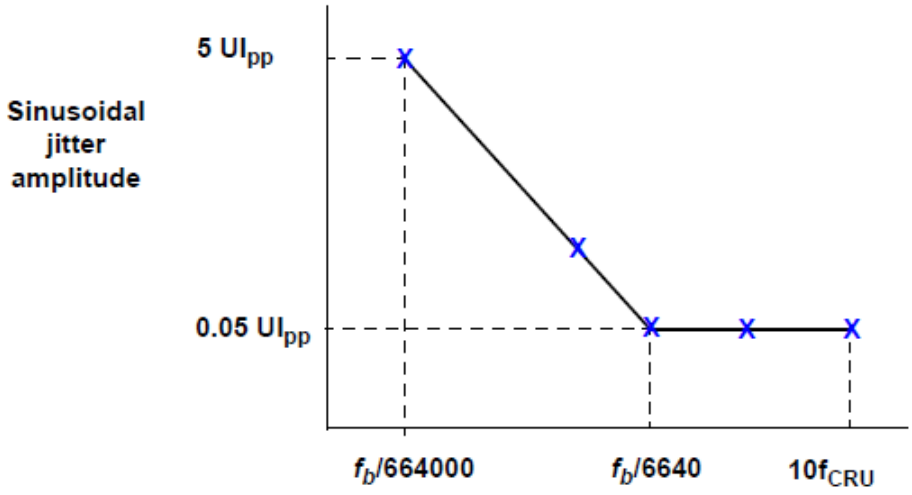

Figure 3.3. Receiver Jitter Tolerance from OIF-CEI-04.0 Standards[1]

Only phase noise above the CDR bandwidth contributes to eye closure. Using the mask from the above figure, the lower limit is fb/6640 = 29 G/6640 = 4.3 MHz. The upper limit is commonly Fref_min/10. One of the common frequencies used in SerDes is 156.25 MHz; therefore, the upper limit is approximately 15.6 MHz.

4. Si539x Jitter Attenuator Advantages for 56G PAM-4 SerDes

The figure below is a picture of the jitter attenuator used in a backplane on a line card containing a SerDes. The jitter attenuator is used to clean the recovered timing signal before using it as a reference for the SerDes. This shows where the Si539x jitter attenuator is used in the system.

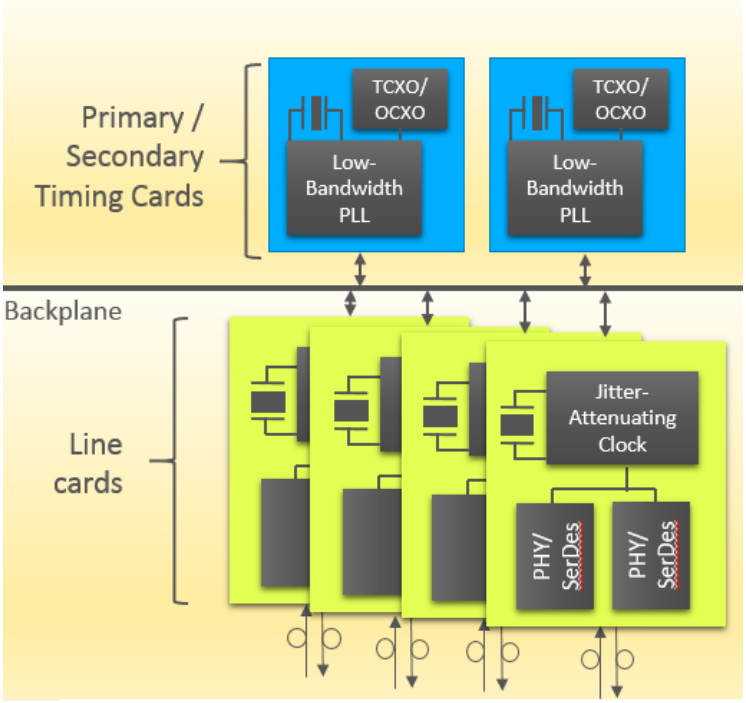

Figure 4.1. Line Card Showing Jitter Attenuator used for Signal Integrity of the 56G SerDes

The Si539x family of jitter attenuators has a small footprint, with high quality internal power supply rejection (PSR) circuitry, which eliminates the need for additional bulky supply decoupling capacitors on the board. The Si5395/94/92 P-grade products have been jitter performance optimized for common SerDes application frequencies (156.25 MHz, 312.5MHz). In a low jitter application that demands strict specifications for high performance it is critical that the jitter attenuator minimizes the jitter as much as possible. This will provide a bigger budget for other areas of the system.

Other important features of the Si539x jitter attenuators include the following:

- Hitless redundant input clock switching with <1ns output phase transient
- Up to 12 outputs for flexible frequency planning
- Integrated crystal option

The following figures are phase noise plots taken from a Keysight E5052B Signal Source Analyzer of the Si5395P. The first figure shows the RMS Jitter is 69 fs over the integration band from 12 kHz to 20MHz. This is typical performance for the P grading option. Note that this is the wider integration band than the standards require, but often plots are taken in this integration band, so it provides a good comparison. The second figure is a plot of the Si5395P showing RMS jitter is 38 fs over the integration band from 4 MHz to 20 MHz, which is the band of interest for the SerDes requirement.

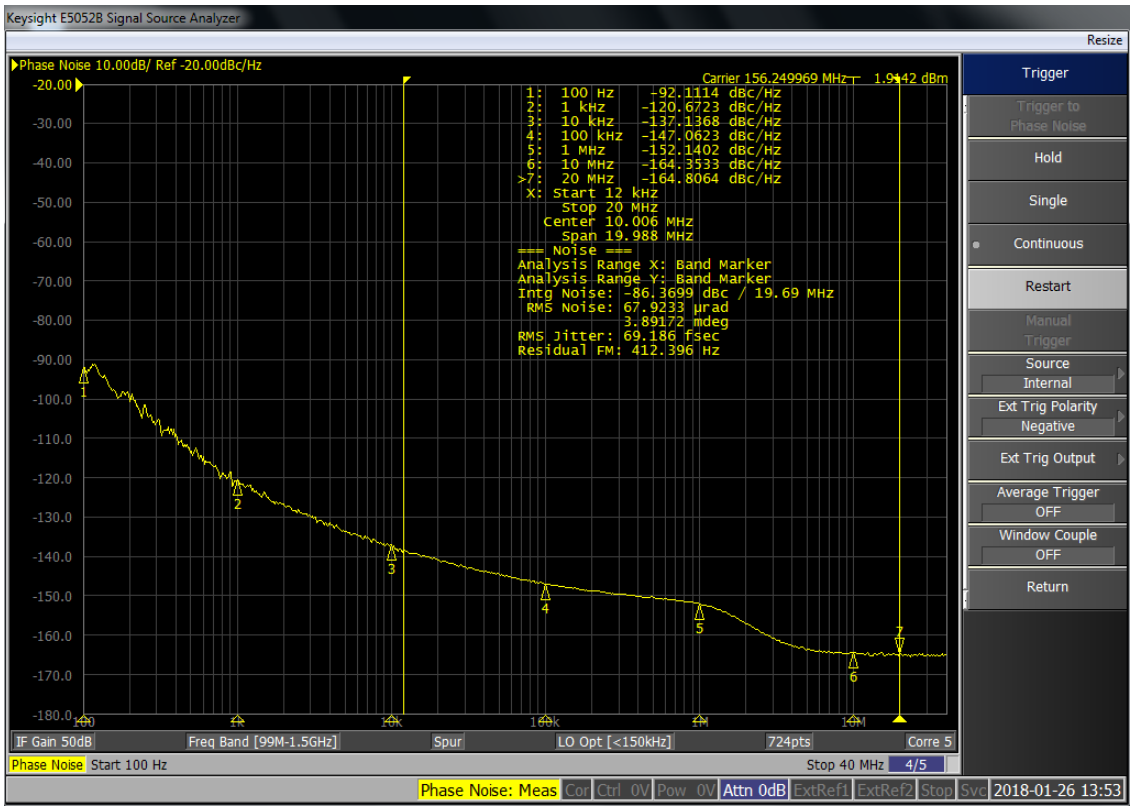

Figure 4.2. Phase Noise Plot Si5395P RMS Jitter 12 kHz–20 MHz Integration Band

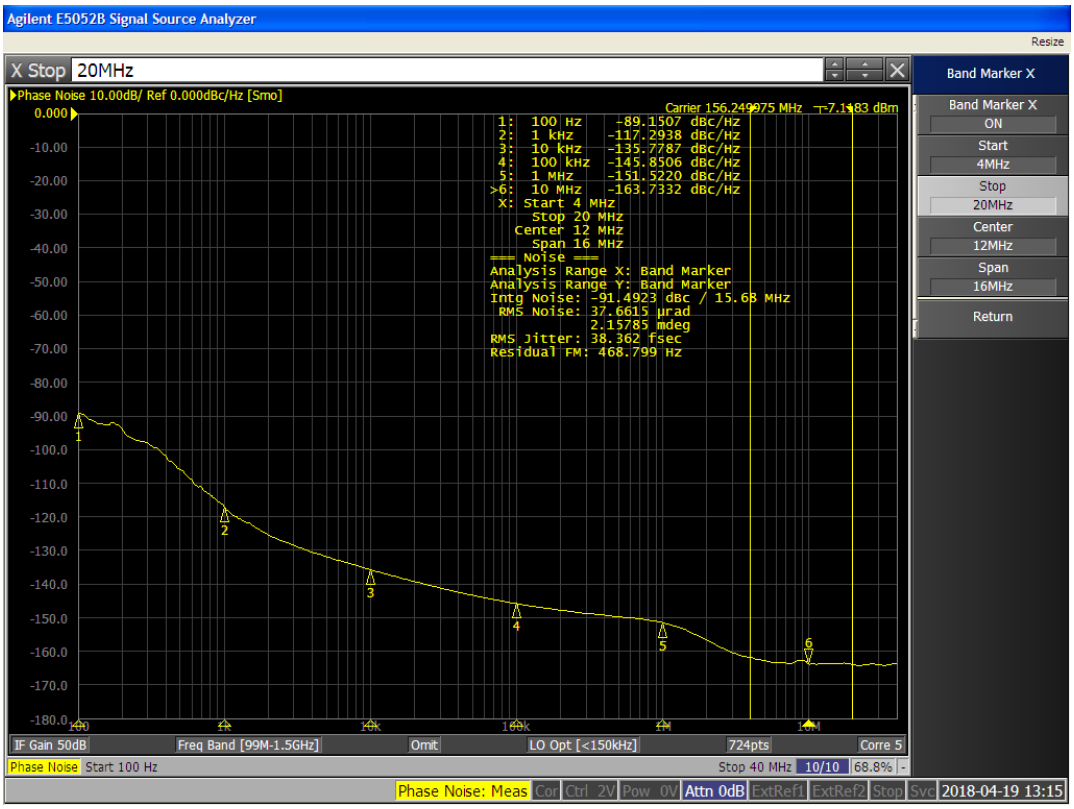

Figure 4.3. Phase Noise Plot Si5395P RMS Jitter 4 MHz–20 MHz Integration Band

5. Conclusions

Increasing data rates are pushing the industry to higher throughput speeds, such as the soon-to-be-common 56G PAM-4 SerDes and 112G PAM-4 SerDes. The 56 Gbaud speed started out as the foundation for 400G by the 802.3bs working group and has now been leveraged to many other projects. Ethernet is introducing 56 Gbaud for a number of different physical layers. The Si539x family of jitter attenuators is well positioned for this market and has been optimized to provide ultra-low jitter for common SerDes frequencies. This will create reliable communication links by ensuring that these systems operate well within their budget requirements. While the 112Gbps (2 x 56 GBaud) standards do not yet exist, one could extrapolate that the direct scaling of the 56 GBaud spec would cut the clock jitter in Table 3.1 Standards for Different Cable/Trace Lengths Comparing Specifications for Jitter and Integration Band on page 7 in half. That would take the CEI-56G-VSR-PAM4 from 250 fs RMS max for 56Gbps systems to 125 fs RMS max for 112Gbps systems, which is still higher than the 107 fs max value (which includes the measured plus aliased jitter calculation) of the Silicon Labs Si5395P. This further justifies the need for the Si5395P going forward for 112 Gbps systems.

6. References

1. © 2017 Optical Internetworking Forum, terms of use on page 4 of OIF Implementation Agreement, December 29, 2017 at http://www.oiforum.com/wp-content/uploads/OIF-CEI-04.0.pdf
2. https://www.silabs.com/documents/login/data-sheets/si5395-94-92-a-datasheet.pdf

ClockBuilder Pro

One-click access to Timing tools, documentation, software, source code libraries & more. Available for Windows and iOS (CBGo only).

www.silabs.com/CBPro

Timing Portfolio
www.silabs.com/timing

SW/HW
www.silabs.com/CBPro

Quality
www.silabs.com/quality

Support and Community
community.silabs.com

SILICON LABS®

Silicon Laboratories Inc.
400 West Cesar Chavez
Austin, TX 78701
USA

http://www.silabs.com