APPLICATION NOTE

AN1432: Using the Si536x, Si540x, and SKY6310x in 112G and 224G SerDes Applications

Features

- Four-level Pulse Amplitude Modulation (PAM-4) 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 Skyworks Si536x, Si540x, and SKY6310x jitter attenuators is highlighted

Description

Ethernet speeds are continuing to increase, pushing to 800G and beyond due to demands for more data delivered at faster speeds. The latest switch processor chips with 112G and 224G PAM-4 interfaces enable new high-speed interconnects in the cloud and data center networks.

The challenge is to make the data pipeline faster while ensuring that networks remain reliable, easy to maintain, and cost-effective. This application note explains 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 Si536x, Si540x, and SKY6310x family of jitter attenuators are ideal for 112G and 224G SerDes designs.

1. Ethernet Interface Examples

Current 400G to 800G Ethernet connections commonly use 56G and 112G SerDes, which provide 50 Gbps and up to 100 Gbps per lane throughput. The inclusion of forward error correction increases overhead, reducing the 56 Gbps and 112 Gbps line rates to 50 Gbps and 100 Gbps data rates, respectively. 400 Gb to 800 Gb+ Ethernet switches are entering the market and starting to use 112G and 224G SerDes Phys to meet the requirements for the industry transition to 800 GbE.

Eventually, these SerDes Phys will need to scale further, as discussions about 448 Gbps lane rates have started. Supporting faster Ethernet speeds is a challenge for SerDes designers as the industry transitions from 56 Gbps to 112 Gbps, 224 Gbps, and eventually to 448 Gbps per single lane. The table below lists some examples of the standards for different 400 to 800 Gb Ethernet interfaces.

Table 1. 400 to 800 Gb Ethernet Interface Standards

Name	Medium	Media Count/Lanes	Gigabaud/Lane
200GBASE-KP4	Copper backplane	4	53.125, KP4-FEC
200GBASE-KR2	Copper backplane	2	53.125 using PAM4 modulation
400GBASE-CR4	Twin axial cable	4	53.125 (PAM4 modulation)
400GBASE-LR4	Single mode fiber WDM 1295.56 to 1709.14 nm	4	53.125 (PAM4 modulation)
800GBASE-SR16	Multimode fiber 850 nm laser	16	53.125
800GBASE-DR4	Single mode fiber 1304.5 to 1317.5 nm	4	106.25 (PAM4 modulation)
800GBASE-FR8	Single mode fiber WDM 1273.54 to 1309.14 nm	8	53.125 (PAM4 modulation)
800GBASE-LR8	Single mode fiber WDM 1273.54 to 1309.14 nm	8	53.125 (PAM4 modulation)

High speed 112G+ SerDes links are also important in 5G infrastructure. They 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, Intersymbol Interference (ISI), reflections, crosstalk, etc. There are also different methods for 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. Non-Return to Zero (NRZ) PAM-2 Modulation

The traditional method of modulation is called Non-Return to Zero (NRZ), which is also called Pulse Amplitude Modulation 2 (PAM-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 one 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 retain 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 $\log_{10}$ (1/3) = 9.5 dB. Thus, PAM-4 trades signal-to-noise ratio for bandwidth.

Figure 1 shows a comparison of NRZ versus 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 throughput for the same baud rate. PAM-4 has a lower SNR, as it is more susceptible to noise. The diagram shows a comparison of the different encoded signal levels and the corresponding eye diagrams for each.

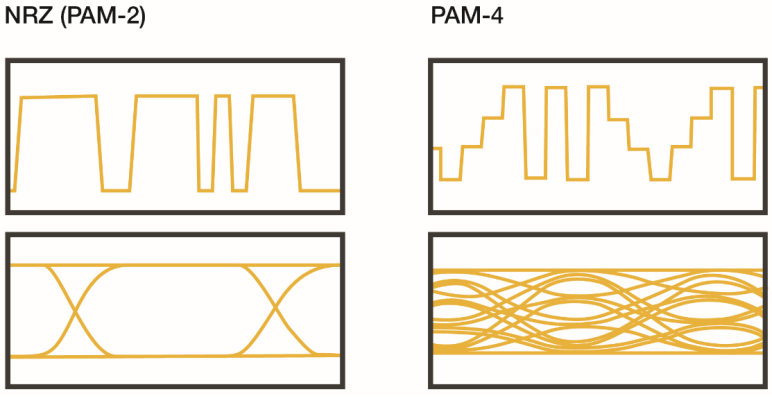

Figure 1. NRZ Compared to PAM-4 Signals

Figure 2 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 that 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.

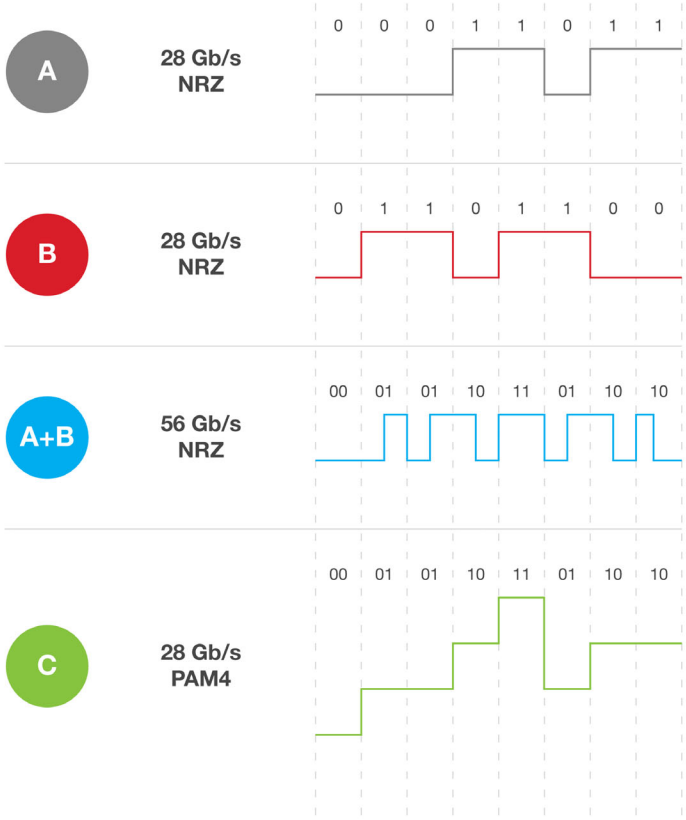

Figure 2. Signal Levels and Encoding of SNR versus PAM-4

2.3. Forward Error Correction

Forward error correction (FEC) is used to correct transmission errors for 100G and 200G 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 increases 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 200G lane rates (up to 224G 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 2018, an IEEE 802.3 task force developed a single-lane 100 Gbps Ethernet standard in 802.3ck to be used in 800 Gbps. The 802.3ck standard requires a PAM-4 encoding scheme, as described above, to reach 100 Gbps per channel throughput. There are several options for delivering 800 Gbps based on the distance requirements and multimode fiber versus single-mode fiber.

The Optical Internetworking Forum (OIF) also has standards published for 112G 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 112G 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-112G-PAM4 stands for Optical Internetworking Forum Common Electrical I/O for 112G PAM-4. Figure 3 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.

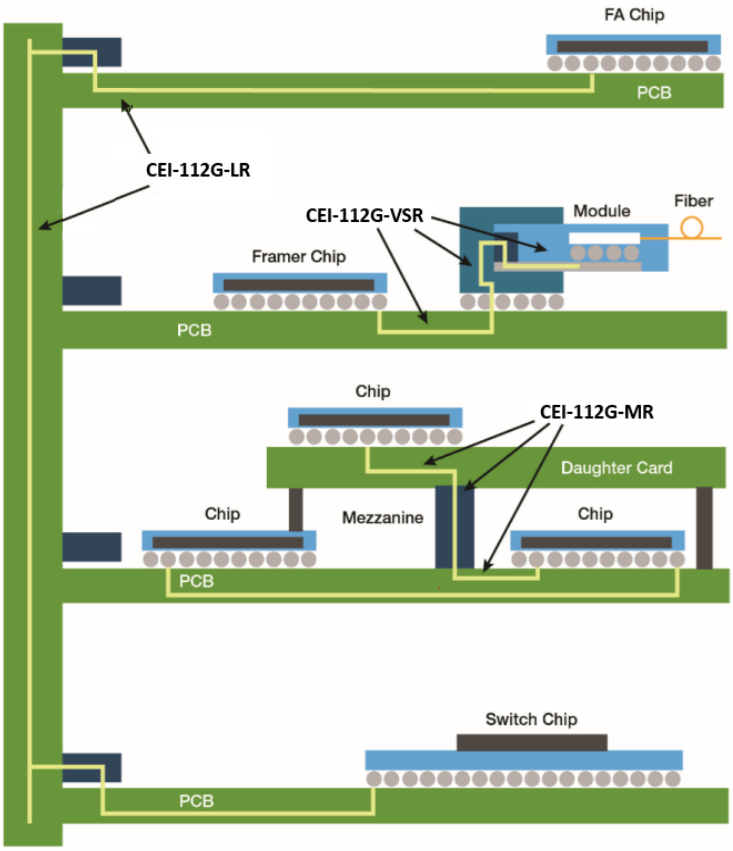

Figure 3. Different SerDes Interconnects of Varying Trace Lengths[1]

The VSR chip-to-module specification is <10 cm, one connector with a loss of up to 10 dB with pre-FEC BER < 1^{-6}. Figure 4 is a diagram from the OIF-CEI-05.0 specification describing the expected losses between the connections.

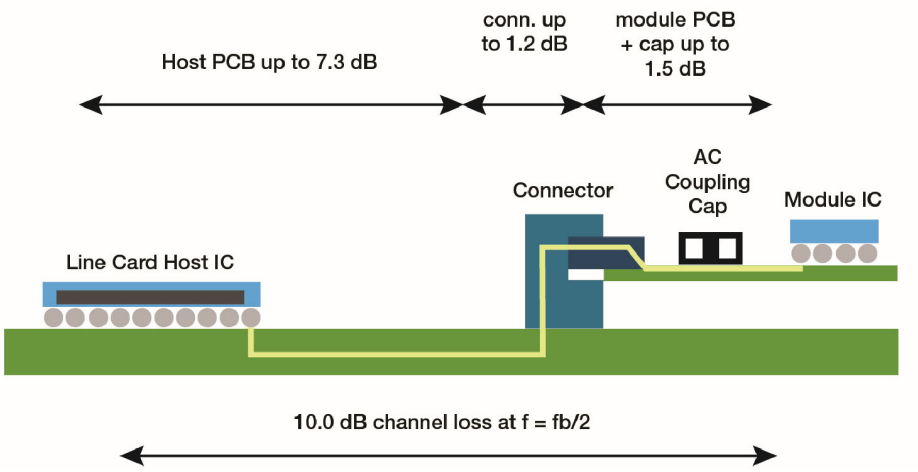

Figure 4. VEI-56G-VSR-PAM4 Channel Reference Model[1]

- Long reach has backplanes or copper cables and two connectors with loss of up to 30 dB at Nyquist and pre-FEC raw BER <$1e^{-4}$.
- 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 <$1e^{-6}$.

Table 2 is a summary of the standard and max data rates, and the spec given from the standard along with the translated jitter requirement for the reference clock. This is assuming a 10 percent budget for the reference clock with ~90 percent 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 budget calculation is done in the next section. 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 Si536x performance, shown in the right hand column of Table 2, provides significant margin to the requirements in the standards and is used as the base device in this application note. The Si540x and SKY6310x datasheet jitter performance is specified to be comparable or better than the Si536x. The typical jitter specified in the Si536x data sheet in the 12 kHz to 20 MHz range is 54 fs. Jitter in the 4 MHz to 20 MHz integration band is typically measured to be less than half of that value. An example phase noise plot is shown in Figure 7, “Si536x 312.5 MHz Phase Noise Plot RMS Jitter 12 kHz to 20 MHz, 4 MHz to 20 MHz, and 4 MHz High Pass Filtered 12 kHz to 20 MHz Integration Bands,” on page 11.

It is important to realize that the reference clock phase noise beyond Nyquist ($f_{ref}/2$) 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® 5052B or 5055A provides overly optimistic jitter numbers compared to what is actually seen at the interface due to these effects. The Keysight 5052B, for example, 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 780 MHz. It is assumed that at 780 MHz, the signal and reference clock receiver input bandwidth roll off and phase noise beyond that can be considered insignificant.

When folding the approximated (flat) phase noise back into the integration band, folds of 1.5F to 2.0F and 2.0F to 2.5F are added, which is the noise floor two times. In the case of a 312.5 MHz clock frequency, 2.5F is approximately 780 MHz. Since the estimated phase noise is conservatively assumed to be 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:

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

Equation 1. RMS Jitter Calculated From Flat Phase Noise Over ΔF

From the equation above: "f" is the carrier frequency of 312.5 MHz, "ΔF" is 16 MHz (4 MHz to 20 MHz) , and this is the integration band. "Phase_Noise" is the phase noise floor from the Si536x plot using the Keysight 5055A in Figure 7, “Si536x 312.5 MHz Phase Noise Plot RMS Jitter 12 kHz to 20 MHz, 4 MHz to 20 MHz, and 4 MHz High Pass Filtered 12 kHz to 20 MHz Integration Bands,” on page 11, which is approximately –164.5 dBc.

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

$$Aliased\_Jitter = \frac{\sqrt{2 * 2 * 16E6 * 10(-16.45)}}{1964E6} = 24\ fs$$

Equation 2. Aliased RMS Jitter Estimated as 2X Flat Phase Noise Over ΔF

For the integration band from 4 MHz to 20 MHz, the total jitter is sqrt (21^2 + 24^2) = 32 fs. Adding a 3 dB margin to the phase noise is 45 fs RMS total sampling jitter.

Table 2. Standards for Different Cable/Trace Lengths Comparing Specifications for Jitter and Integration Band

Standard	Max Rate	Specification	Budgeted REFCLK Jitter Including Aliased Noise (4 MHz – TX BW)	Si536x 312.5 MHz Jitter Including Aliased Noise (4 MHz to 20 MHz, 32 fs Typ)[1]
CEI-112G-LR-PAM4	116 Gbps	0.023 UI TX RMS	175 fs RMS	45 fs RMS Max
CEI-112G-MR-PAM4	116 Gbps	0.023 UI TX RMS	175 fs RMS	45 fs RMS Max
CEI-112G-VSR-PAM4	116 Gbps	0.20 UI eye width	125 fs RMS	45 fs RMS Max
400GBASE-FR4 Fiber	106 Gbps	3.4 dB TDECQ	115 fs RMS	45 fs RMS Max
800GAUI-8 C2Chip	106 Gbps	0.023 UI TX RMS	190 fs RMS	45 fs RMS Max
800GAUI-8 C2Module	106 Gbps	0.22 UI eye width	140 fs RMS	45 fs RMS Max

1. RMS total sampling jitter + 3 dB margin.

Multiple switch/ASIC vendors with 112G PAM-4 SerDes require 75 fs RMS max integrated over 12 k to 20 MHz. It is important to understand that this does not include the aliased jitter analysis from Table 2 and is just the measured jitter from the phase noise analyzer. The Si536x is typically well below that value at 54 fs per the data sheet from 12 kHz to 20 MHz. This means these products are better than these specifications, assuming the aliasing calculation is not included.

3.1. Example Calculation of Reference Clock Jitter from Standards

The reference clock jitter was calculated using the standards for each line item in Table 2. For example, the CEI-112G-LR-PAM4 standard specifies a maximum data rate of 116 Gbps. The uncorrelated jitter RMS (standard deviation of the probability distribution) is 0.023 UI TX RMS, measured at the output of the transmitter.

For PAM4, one unit interval, UI, is 1/(half the max bit rate). The bit rate is given at 116 Gbps. Therefore, 0.023/58G = 396 fs, rounded to 400 fs for simplicity.

Since the transmitter output jitter is a combination of the reference clock jitter and the intrinsic jitter of the transmitter, these are broken up into separate components. The reference clock jitter and the transmitter jitter are uncorrelated. Thus, Equation 3 is used to relate the reference clock jitter and the transmitter jitter, assuming that the jitter is random.

$$Tjrefclk = \sqrt{Tjtotal^2 - Tjtx^2}$$

Equation 3. Uncorrelated RMS Jitter Equation from the Standards

The total is 400 fs. If 10 percent is budgeted for the reference clock, the reference clock and transmitter jitter is calculated below:

$$\begin{aligned} Tjrefclk &= \sqrt{Tjtotal^2 - (0.9 * Tjtotal)^2} \\ &= \sqrt{400^2 - 360^2} \\ &= 175\ fs \end{aligned}$$

Equation 4. RMS Equation Calculations to Compute Budget Requirements

The assumption is that 10 percent is budgeted for the reference clock and 90 percent for the transmitter. This works out to 175 fs for the reference clock, leaving 360 fs for the transmitter.

The lower level of the integration band is determined based on the receiver jitter tolerance mask from the OIF-CEI 0.50, plotted in Figure 5.

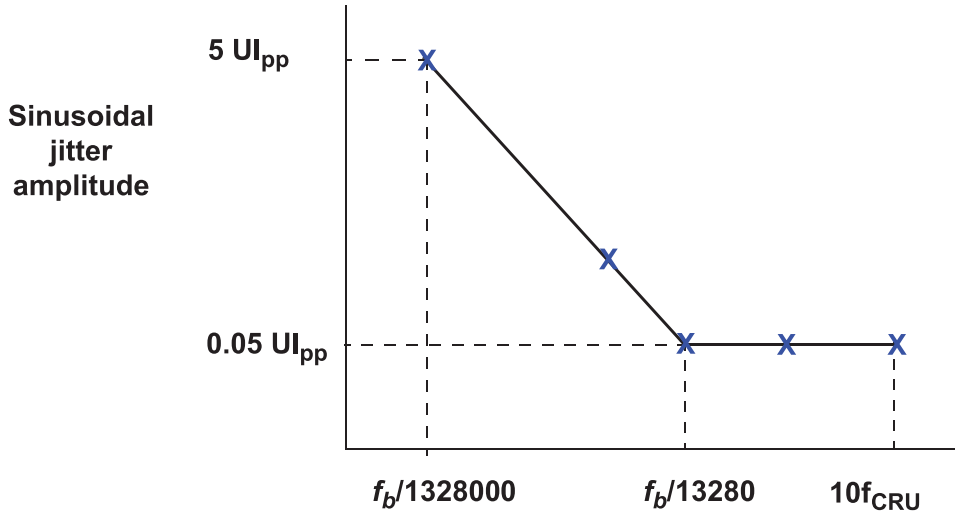

Figure 5. Receiver Jitter Tolerance from OIF-CEI-05.0 Standards[1]

Since the CDR tracks phase noise below its bandwidth, only phase noise above the CDR bandwidth contributes to eye closure. Using the mask from the above figure, the lower limit is fb/13280 = 56 G/13280 = 4.2 MHz. The upper integration limit is commonly Fref_min/10 to maintain SERDES PLL stability.

One of the common frequencies used in SerDes is 156.25 MHz; therefore, the upper limit is approximately 15.6 MHz. The lower limit is conservatively rounded down to 4 MHz, and the upper limit is conservatively rounded up to 16 MHz (the wider integration range conservatively includes more phase noise, hence more jitter).

3.2. Modeling the CDR BW Using a High Pass Filter

One way to model the effect of the CDR in tracking and effectively eliminating reference clock phase noise below 4 MHz is to start the jitter integration at 4 MHz. An alternative method is to leave the integration starting point at the traditional 12 kHz frequency, but to first apply a 4 MHz High Pass Filter (HPF) to the phase noise prior to performing the integration.

Some phase noise analyzers allow this to be done, including the Keysight E5052B and E5055A. An example using this method is shown in Figure 7, “Si536x 312.5 MHz Phase Noise Plot RMS Jitter 12 kHz to 20 MHz, 4 MHz to 20 MHz, and 4 MHz High Pass Filtered 12 kHz to 20 MHz Integration Bands,” on page 11. Some SerDes vendors have started specifying the reference clock jitter using this methodology and have established a 35 fs maximum jitter limit after applying a 4 MHz first order HPF.

As seen in Figure 7, the example Si536x Trace 2 phase noise plot yields a typical value of 22 fs. Product characterization provides a maximum jitter of 72 fs versus typical 54 fs jitter over temperature, voltage, and process for 12 kHz to 20 MHz. Assuming the same ratios apply suggests a max expected 4 MHz high-pass filtered jitter would be 22 fs RMS typ * (72/54) or 29 fs RMS max, which meets the 35 fs RMS max specification.

4. Si536x/Si540x/SKY6310x Jitter Attenuator Advantages for 112G and 224G PAM-4 SerDes

Figure 6 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.

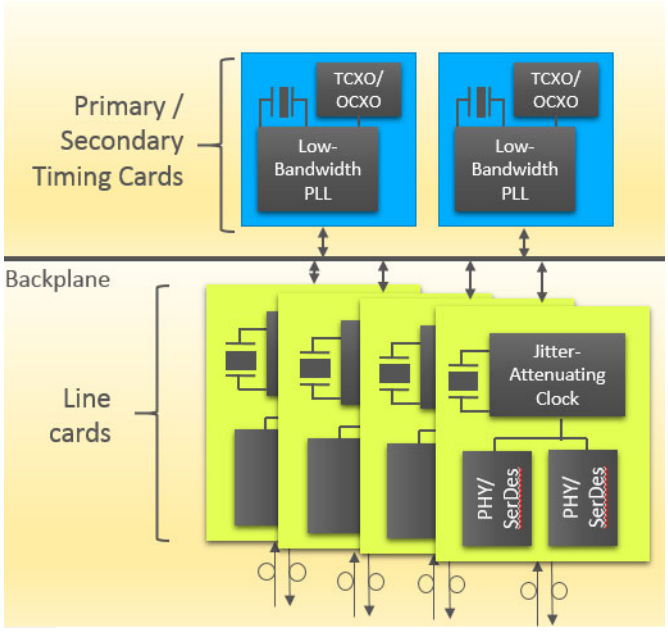

Figure 6. Line Card Showing Jitter Attenuator used for Signal Integrity of the 56G SerDes

The Si536x, Si540x, and SKY6310x family of jitter attenuators have a small footprint, with high quality internal power supply rejection (PSR) circuitry. 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 provides a bigger budget for other areas of the system.

Other important features of the Si536x, Si540x, and SKY6310x jitter attenuators include the following:

- Hitless redundant input clock switching with <150 ps maximum output phase transient
- Available in 18 or 12 output versions for flexible frequency planning
- Reprogrammable up to nine times
- Option to boot from external flash (SKY6310x only)

Figure 7 is an Si536x example phase noise plot measured with a Keysight E5055A signal source analyzer. Trace 1 shows the RMS jitter is 49 fs over the integration band from 12 kHz to 20MHz. This example is better than typical data sheet performance. Note that this is a wider integration band than the standards require, but often plots are taken in this integration band, so it provides a good comparison.

Trace 1 also shows the RMS jitter is 21 fs over the integration band from 4 MHz to 20 MHz, which is the band of interest for the SerDes requirement. Finally, the third integration band based on Trace 2 shows the RMS jitter is 22 fs after applying a first order 4 MHz HPF and integrating from 12 kHz to 20 MHz

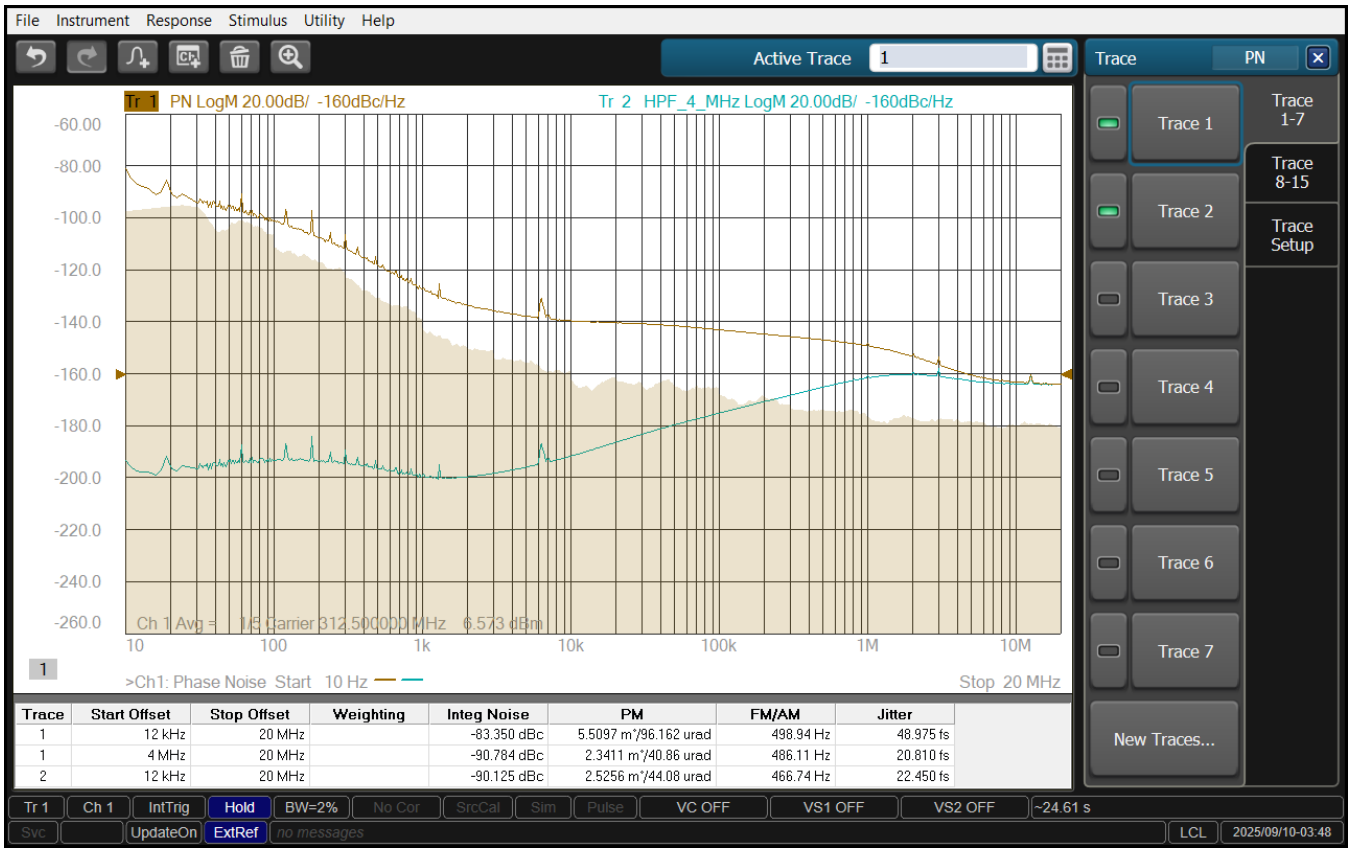

Figure 7. Si536x 312.5 MHz Phase Noise Plot RMS Jitter 12 kHz to 20 MHz, 4 MHz to 20 MHz, and 4 MHz High Pass Filtered 12 kHz to 20 MHz Integration Bands

5. Conclusion

Increasing data rates are pushing the industry to higher throughput speeds, such as the 112G PAM-4 SerDes and 224G PAM-4 SerDes. The 56 Gbaud speed started as the foundation for 400G by the 802.3ck working group and is now leveraged in many other projects.

Ethernet has introduced 112 Gbaud for a number of different physical layers. The Si536x, Si540x, and SKY6310x (which includes the SKY63104/5/6 12-output jitter attenuators, and the SKY63101/2/3 BAW family of jitter attenuators) are well positioned for this market and have 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.

6. References

1. © 2022 Optical Internetworking Forum, terms of use on page 4 of OIF Implementation Agreement, May 5, 2022 at https://www.oiforum.com/wp-content/uploads/OIF-CEI-5.0.pdf.
2. Si5361/Si5362/Si5363 18-Output, Any-Frequency, Any-Output Jitter Attenuators/Clock Multipliers with Ultra-Low Jitter data sheet.