

Understanding clocking needs for high-speed 56G PAM-4 serial links

Madhu Balasubramanian

ABSTRACT

This application report outlines the advantages of using ultra-high performance clock synchronizers from Texas Instruments to generate reference clocks needed for high-speed serial links using 56G PAM-4 signaling, and even meets early requirements of 112G PAM-4 links. A methodology for deriving reference clock jitter requirements is described, and the advantages of clocking such a system with the LMK05318 are outlined.

Contents

1	Introduction	1
2	56G PAM-4 Standards for 400G Ethernet.....	2
3	56G PAM-4 SerDes Clocking With LMK05318	3
4	Summary.....	5

List of Figures

1	SerDes Interconnects of Varying Trace Lengths	2
2	156.25-MHz LVPECL Phase Noise of LMK05318	4

List of Tables

1	Reference Clock Requirements for 56G PAM-4 Standards.....	3
---	---	---

Trademarks

All trademarks are the property of their respective owners.

1 Introduction

Ethernet has been the standard, dominant way to connect computers, phones, routers, switches, and other internet devices on a network over a wired connection. This permits the creation of local area networks (LAN) and wide area networks (WAN), expansion of data centers to support increased needs for cloud computing, and innovative software services with high bandwidth levels and quality performance and reliability which allow connected devices to communicate with each other. With Ethernet networks growing over the years, service providers have been significantly enhancing and improving their network capacity to meet the soaring demand for next-generation video and additional multimedia applications. This application-focused need is driving the migration to 100-Gbps and 400-Gbps network speeds in order to accommodate faster data transmission with reduced latency. 100G and 400G Ethernet links support rapid expansion of cloud services, telecommunications, and high-bandwidth applications to unseen, futuristic efficiencies.

In a typical high-speed serial link, data transmission is through a pathological channel, typically a very short-reach PCB trace in a backplane or a long-reach fiber optic cable. Such channels impact signal integrity due to non-idealities like crosstalk, insertion loss, return loss, inter-symbol interference, and jitter. There are different ways of modulating the transmitted data.

1.1 NRZ Modulation

The traditional method of modulation is called Non-Return to Zero (NRZ) that 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.

1.2 PAM-4

Pulse Amplitude Modulation-4 (PAM-4) is a four-level modulation scheme used for data transmission greater than 28 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 traditional NRZ, PAM-4 achieves the double the data rate for a given bandwidth. There are three eyes in each unit interval with an equivalent height that is 1/3 of NRZ at a reduced signal-to-noise ratio (SNR) of 9.5 dB. As a result, PAM-4 data transmission is more susceptible to system noise including reference clock jitter.

2 56G PAM-4 Standards for 400G Ethernet

IEEE 802.3bs describes eight lanes of 50 Gbps per lane for 400G Ethernet using a PAM-4 encoding scheme. Several transmission media are supported for delivering 400 Gbps based on the distance. Similarly, OIF-CEI-56G-PAM4 describes data transmission of 56 Gbps in optical networks over different trace and cable lengths.

Figure 1 shows a Long Range (LR) connection between two chips, Very Short Range (VSR) connection between a Chip and a Module, and a Medium Range (MR) connection between two chips.




Figure 1. SerDes Interconnects of Varying Trace Lengths

2.1 Clocking Requirements for 56G PAM-4 Links

Table 1 lists the reference clock needs in 56G PAM-4 links for the different trace and cable lengths. It is to be noted that multiple switch ASIC vendors with 56G PAM-4 SerDes require a maximum reference clock jitter of 150 fs RMS over a frequency band of 12 kHz to 20 MHz.

Table 1. Reference Clock Requirements for 56G PAM-4 Standards

Standard	Data Rate (Gbps)	Specification	Max Ref Clock Jitter (fs RMS)	Offset Frequencies (MHz)	
				RX – Minimum CDR bandwidth	TX – Maximum PLL bandwidth
CEI-56G-LR-PAM4	57.8	0.023*UI (TX RMS)	250	4	20 (> Fref/2)
CEI-56G-MR-PAM4	57.8		250		
400GAUI-8 C2C	53.125		270		
CEI-56G-VSR-PAM4	57.8	0.265*UI (Eye Width)	240		
400GAUI-8 C2M	53.125	0.22*UI (Eye Width)	220		

The overall allowable jitter in a serial link is dictated by IEEE or CEI. For example, IEEE 802.3bs states in the 400GAUI-8 that for chip-to-chip (C2C) data transmission, the maximum transmit jitter (RMS) should be no more than $0.023 * UI$ where 1 UI is the inverse of 26.5625G. This equates to 865 fs RMS for the overall allowable transmit jitter. The jitter contributing elements are made up of the reference clock, driven from a device like LMK05318, the transmit medium, transmit driver etc. Only a portion of the overall allowable transmit jitter is allocated to the reference clock that could be as low as 10% for PAM-4 systems where SNR is typically lower. Therefore, the allowable reference clock jitter, for a 10% clock jitter budget, is 270 fs RMS.

Another example could be from the OIF-CEI-56G standard which states for very short range (VSR) data transmission that the minimum eye width should be $0.265 * UI$ where 1 UI is the inverse of 28.9G. This equates to 9.17 ps pp for the overall allowable transmit jitter. For both the CEI-56G-VSR-PAM4 and 400GAUI-8 C2M standards, it is assumed that the reference clock contributes up to 20% of the overall jitter and the target BER is $2.4 * 10^{-4}$. Therefore, the allowable reference clock jitter for the former is 240 fs RMS and the latter is 220 fs RMS.

3 56G PAM-4 SerDes Clocking With LMK05318

TI's LMK05318 device is an ultra-high performance clock generator, jitter cleaner, and clock synchronizer with advanced reference clock selection and hitless switching feature to meet the stringent requirements of communications infrastructure applications. The ultra-low jitter performance of this device minimizes bit error rates (BER) in applications involving high-speed serial links and is TI's recommended clocking solution for 56G PAM-4 SerDes. The device features a single channel synchronizer that can synchronize to one of two differential or single-ended reference clock inputs. The synchronizer is accompanied by one PLL domain that features TI's proprietary Bulk Acoustic Wave (BAW) resonator as a VCO and generates 58 fs rms typical jitter in the frequency band of 12 kHz to 20 MHz and 35 fs rms typical jitter in the frequency band of 4 MHz to 20 MHz for a 156.25 MHz output, both of which exceed the 56G PAM-4 reference clock needs as specified in the standards and for the switch ASICs. Figure 2 shows the phase noise plot of 156.25 MHz LVPECL output from the PLL featuring BAW VCO. There is an additional PLL domain that can be used to generate unrelated frequencies either locked to the reference clock input or the free-run XO input and generates 120 fs, rms jitter. The LMK05318 can generate up to eight high performance output clocks with up to six different frequencies.

With careful frequency planning and PLL optimization using clock design tools, such as TI's WEBENCH Clock Architect Tool, and on-chip LDOs to suppress supply noise, the LMK05318 is able to generate clock outputs that exceed the 56G PAM-4 requirements. This level of clock performance enables the serial link system to operate with fewer bit errors by providing the most margin among all available clocking solutions in the industry.




Figure 2. 156.25-MHz LVPECL Phase Noise of LMK05318

3.1 Advantages of LMK05318

LMK05318 offers the following advantages over competing solutions when used as a reference clock for 56G PAM-4 SerDes.

- **Jitter and associated component cost:** Competing solutions utilize LC VCOs with low quality factor (less than 20) and thus rely on an extremely low noise and high frequency external oscillator, that is fairly expensive and not commonly available, or a high frequency crystal, that is very unreliable, prone to failures and highly susceptible to shock and vibration events. In contrast, the LMK05318 utilizes a BAW VCO with a very high quality factor (greater than 1000) and thus has no dependency on the phase noise or frequency of the external oscillator. This minimizes the overall solution cost and makes use of commonly available components. With even the best available external components like the oscillator or crystal, competing solutions still offer inferior jitter performance against LMK05318.
- **Reliability:** Some clocking vendors are attempting to integrate a crystal and the clocking IC into a single package. While this offers a higher level of integration, the biggest drawback of such approaches is that the overall solution becomes highly unreliable and more prone to failures compared to external crystals placed on PCBs. When integrated with a higher power IC in a plastic package, crystals tend to experience large thermal gradients making them more vulnerable to failures as opposed to an oscillator to being on a PCB where proper precaution can be taken during PCB design.
- **Hitless Switching:** If Synchronous Ethernet is adopted, the LMK05318 offers the ability to support input switchover with industry leading phase transient of ± 50 ps using phase cancellation. The LMK05318 is also fully compliant to ITU-T G.8262.
- **PSRR:** The LMK05318 integrates multiple LDOs for all blocks that have a good level of dropout from the PCB supply that is provided to the supply pins. This in turn provides lower than -80 dBc PSRR on all supply domains in the device in the presence of 50 mV noise ripple that can be applied to any of its supply pins without degrading the jitter performance. The LMK05318 can meet the requirements for 400G Ethernet even when its power is supplied from a DC-DC converter and doesn't need an LDO for any of its supply pins.
- **Non Volatile Memory:** The LMK05318 has an integrated EEPROM that is 100 times re-writable. The EEPROM can be written either in the factory or in the field on the end-application's PCB without the need to apply a high voltage to program it. Competing solutions offer at best an OPT or two-time programmable that can only be written in the factory.

4 Summary

Ultra-high performance clock synchronizers from TI, like the LMK05318, can outperform competing solutions, which suffer from various drawbacks that affect the overall performance of high-speed serial link systems as outlined in this report. The ultra-low-jitter of the LMK05318, combined with its plethora of features like support for synchronization, frequency margining, simplify the overall system development, including design, prototyping and standards compliance. TI also offers WEBENCH Clock Architect Tool that helps the hardware engineers to select the appropriate settings for the LMK05318 that meet their requirements.

Revision History

NOTE: Page numbers for previous revisions may differ from page numbers in the current version.

Changes from Original (December 2018) to A Revision	Page
• Changed document title from: <i>Clocking High-Speed 56G PAM-4 Serial Links With LMK05318</i> to: <i>Understanding clocking needs for high-speed 56G PAM-4 serial links</i>	1

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES "AS IS" AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, or other requirements. These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you will fully indemnify TI and its representatives against, any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI's products are provided subject to TI's Terms of Sale (www.ti.com/legal/termsofsale.html) or other applicable terms available either on ti.com or provided in conjunction with such TI products. TI's provision of these resources does not expand or otherwise alter TI's applicable warranties or warranty disclaimers for TI products.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2019, Texas Instruments Incorporated