

112G and 224G PAM-4 SerDes Clocking for Rapid Data Center Switches



Jennifer Bernal, Kumarpal Mandoth

Clocks and Timing Solutions

ABSTRACT

Hyper-scale data centers and telecommunication market sectors are currently driving the need for high speed serial links using 112G and 224G Pulse Amplitude Modulation with 4-Levels Serializer and Deserializer (PAM-4 SerDes). The higher data speeds lower the jitter budget for the reference clock to less than 100 fs RMS.

TI offers a complete clocking design for high-speed serial links using 112G and future 224G PAM-4 SerDes.

The design includes the [LMK5B33216](#) and [LMK5B33414](#) network synchronizers, single-ended and differential buffers, BAW-based oscillators and reference-less clock generators.

TI's Bulk Acoustic Wave (BAW) technology offers industry-leading, ultra-low jitter clocks critical for 112G and 224G PAM-4 SerDes. The LMK5B33216 achieves 42fs RMS typical for 312.5MHz (12kHz to 20MHz) thanks to the low-noise performance of the Voltage-Controlled BAW Oscillator (VCBO). In addition, the LMK5B33216 meets the jitter, rise or fall time, hitless switching, and holdover requirements of 800G reference clocks.

Table of Contents

1 Introduction	2
1.1 800G Market Trend.....	2
2 LMK5B33216 for SerDes Applications	2
2.1 BAW Technology in LMK5B33216.....	5
3 LMK5B33216 for Ethernet Applications	5
3.1 Frequency and Phase Adjustments.....	6
3.2 Input Reference Switching.....	6
3.3 Holdover.....	7
3.4 Zero-Delay Mode.....	8
4 LMK5B33216 Performance	8
4.1 Phase Noise Profile.....	9
4.2 RMS Jitter.....	10
5 Summary	11
6 References	11

Trademarks

All trademarks are the property of their respective owners.

1 Introduction

1.1 800G Market Trend

112G PAM-4 SerDes technology is designed to transmit data at 112 gigabits per second (Gbps). 800G switch technology has made significant leaps forward in data networking by leveraging 112G PAM-4 SerDes. For more details on PAM-4 applications, refer to [Understanding Clocking Needs for High-Speed 56G PAM-4 Serial Links](#).

The 800G high-speed switches are engineered to meet increasing data center and telecommunication demands. The 800G switches have a port speed of 800 Gbps that provides bandwidth for rapid data transmission which reduces network congestion for data processing, storage, and distribution. The 800G switches have fast transmission rates, often measured in terabits per second, that handle immense volumes of data generated by virtualization and high-performance computing.

Clocking in high-speed SerDes systems plays an important role in data synchronization and communication reliability. A SerDes has an integrated PLL that clocks out high-speed serial data while maintaining phase-lock to a reference clock. A high-speed SerDes relies on a precise clocking mechanism to maintain proper timing and accurate serialization and deserialization of data. Therefore, providing a low-jitter clock to the integrated SerDes PLL is necessary.

Incorporating a low-jitter clock is achievable by using an external-network synchronizer or jitter cleaner to reduce the noise propagated onto the serial data. A network synchronizer allows the flexibility of switching between various inputs of different frequencies and phase domains without disrupting the SerDes reference.

TI network synchronizers and jitter cleaners with integrated BAW, such as the [LMK5B33216](#), are designed to provide the reference clocks to the SerDes and to support 800G or higher throughputs of data in real-time communication, AI, or IOT applications.

[Figure 1-1](#) demonstrates a simplified block diagram of a 112G and 224G PAM-4 SerDes clock input sourced from the LMK5B33216.

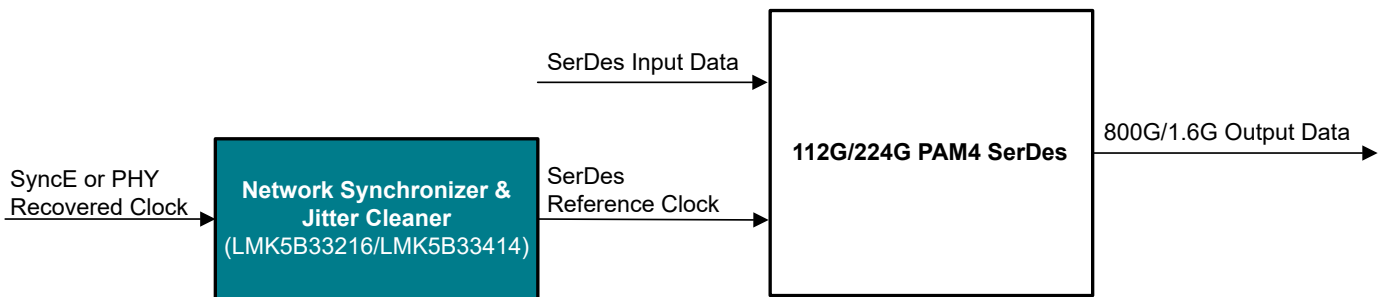


Figure 1-1. 112G and 224G PAM-4 SerDes Simplified Block Diagram

2 LMK5B33216 for SerDes Applications

Designed for SerDes applications, the LMK5B33216 is a high-performance network synchronizer that meets the stringent Ethernet-based networking requirements for jitter, rise or fall time, hitless switching, and holdover.

[Table 2-1](#) compares the performance of the LMK5B33216 with an example SerDes core reference clock used for 800G switch applications. The LMK5B33216 meets and exceeds the requirements of the reference example clock.

Table 2-1. Comparison Between an Example SerDes Core Reference Clock and LMK5B33216 Output

Parameter	Example SerDes Core Reference Clock			LMK5B33216 ⁽¹⁾			UNIT
	MIN	TYP	MAX	MIN	TYP	MAX	
Frequency	—	312.5	—	—	312.5	—	MHz
RMS jitter (12kHz to 20MHz)	—	—	90	—	42 ⁽²⁾	60 ⁽²⁾	fs
Reference clock phase noise at 100kHz	—	—	-137	—	-146	—	dBc/Hz
Reference clock phase noise at 1MHz	—	—	-143	—	-159	—	dBc/Hz
Reference clock phase noise at 10MHz	—	—	-158	—	-160	—	dBc/Hz
Differential peak-to-peak voltage swing	500	—	1800	670	—	2300	mVpp
Rise or fall time (20% to 80%)	—	300	400	175	230	300	ps
Operating ambient temperature	—	—	—	-40	—	85	°C

- (1) The measurements are taken using an XO input frequency of 48MHz with all LMK5B33216 outputs set to the same frequency and sourced from the BAW APLL (APLL3).
- (2) The APLL3 post divider is 8 and the output swing (VOD) is ≥ 800mV.

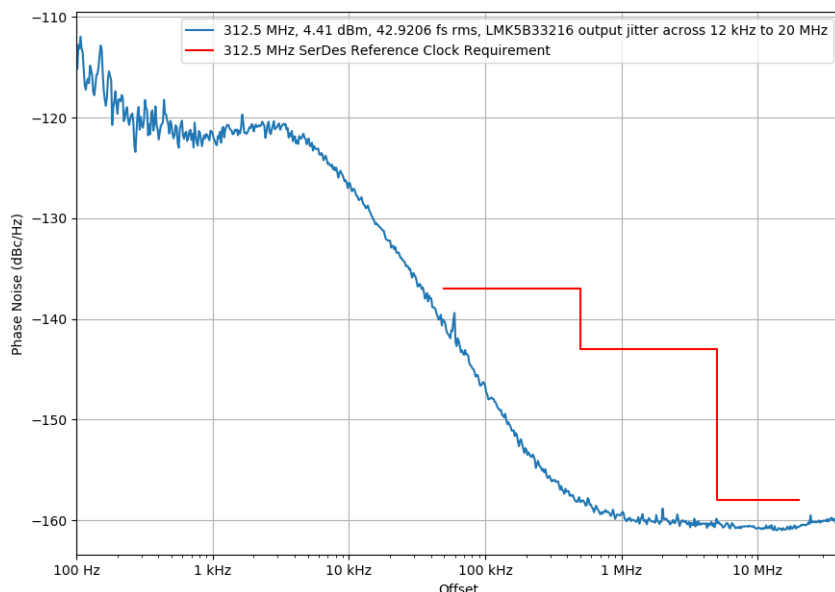


Figure 2-1. LMK5B33216 Output Phase Noise and Jitter for 312.5MHz (using a 48MHz XO input)

Bit error rates (BER) are minimized in applications involving high-speed SerDes links with the excellent jitter performance of the LMK5B33216. The LMK5B33216, along with the TI PTP Stack option, achieves better than Class D accuracy (with sub 5ns timing accuracy). This accuracy meets ITU-T standards like G8275.1 and G8275.2 for full timing and partial timing support. For more information on TI's complete IEEE-1588 PTP and SyncE clocking options, visit TI's [Clocks and Timing](#) page.

TI's network synchronizers leverage a Digital Phase-Locked Loop (DPLL) and an Analog Phase-Locked Loop (APLL) combination to maintain phase lock with a reference clock. The DPLL steers the clock output phase by making continuous updates to the APLL numerator as the DPLL tracks the phase of the reference clock. The LMK5B33216 contains three pairs of DPLL + APLL to support up to three synchronization and frequency domains, as illustrated in the block diagram in [Figure 2-2](#).

The clock outputs can phase-lock between two inputs (LMK5B33216) or four inputs (LMK5B33414) when using the DPLL + APLL. Alternatively, the outputs can lock to a single reference, or oscillator, when using only the APLL. The device can operate in three modes: normal operation with the DPLL active, holdover when the input clocks are unavailable, and free-run where the DPLL is turned off and only the internal APLL is functioning.

Additionally, the LMK5B33216 provides input clock detection and monitoring, wander and jitter filtering, hitless switching, holdover, and Zero-Delay Mode (ZDM) functionality. Each DLL includes a programmable loop bandwidth (LBW) to maximize flexibility and ease of use, which can be set between 1MHz to 4kHz without external loop filter components. The APLL loop filter is fully integrated except for one external capacitor.

The LMK5B33216 clock outputs are configurable from several output formats: 1.8V or 2.65V LVCMOS, HCSL, and LVDS/HSDS (High Swing Differential Signaling). The output swing and common-mode voltage are programmable to meet various receiver requirements with the HSDS format.

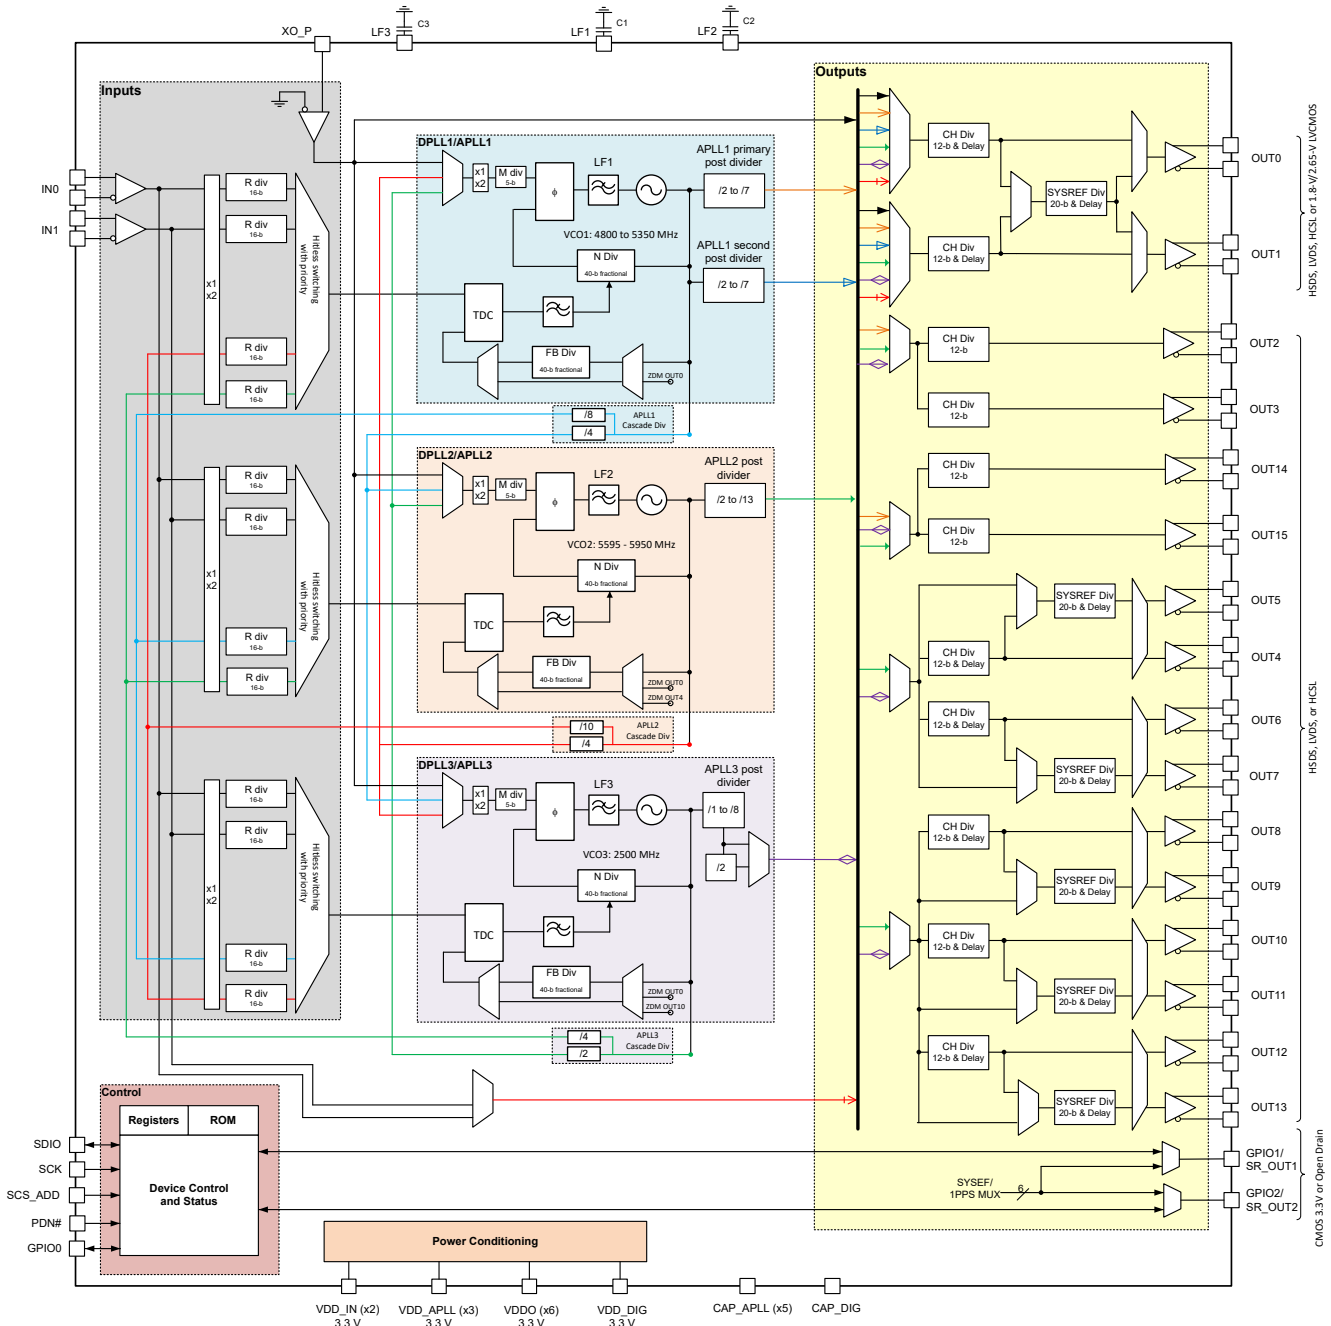


Figure 2-2. LMK5B33216 Functional Block Diagram

2.1 BAW Technology in LMK5B33216

The LMK5B33216 integrates a VCBO into APLL3 (the BAW APLL) to achieve ultra-low phase noise and jitter. The VCBO is a high-Q BAW resonator that is co-packaged with the LMK5B33216 silicon IC and divided down to generate low-jitter Ethernet clocks.

Figure 2-3 shows the composition of the BAW resonator technology. The structure includes a thin layer of piezoelectric film between two metal films and other layers that confine the mechanical energy. The BAW utilizes this piezoelectric transduction to generate a vibration, therefore, outputting a clock frequency. For more information on the BAW, refer to our list of [BAW Application Notes](#).




Figure 2-3. Basic Structure of a Bulk Acoustic Wave Resonator

BAW technology greatly reduces the phase noise and RMS jitter of the synchronized output clock, thereby meeting and exceeding the challenging system requirements for 400 Gbps and 800 Gbps systems by a comfortable margin.

3 LMK5B33216 for Ethernet Applications

The LMK5B33216 is used to jitter-clean SyncE or PHY recovered clocks and provide synchronized, low-jitter, outputs to the ASIC and CPU. For SyncE input clocks, TI recommends setting the DPLL LBW between 1Hz and 3Hz to filter out SyncE transient noise.

Figure 3-1 shows the full system design for a 400G and 800G switch using the LMK5B33216. Pairing the LMK6Cx (TI's BAW-based LVCMOS oscillator family) with the LMK5B33216 yields a low-cost option for the XO input. Additional clocks can be fanned out to the ASIC through a 4, 8, 12, or 16 output, low additive jitter, and clock buffer from the LMK1Dxxxx family, such as the LMK1D1204.

The LMK3H0102 is a reference-less, BAW-based, clock generator used to clock up to two PCIe Gen 1 to PCIe Gen 6 compliant outputs. Each LMK3H0102 output is capable of generating any frequency between 2.5MHz and 400MHz by dividing down from two fractional output dividers (FODs).

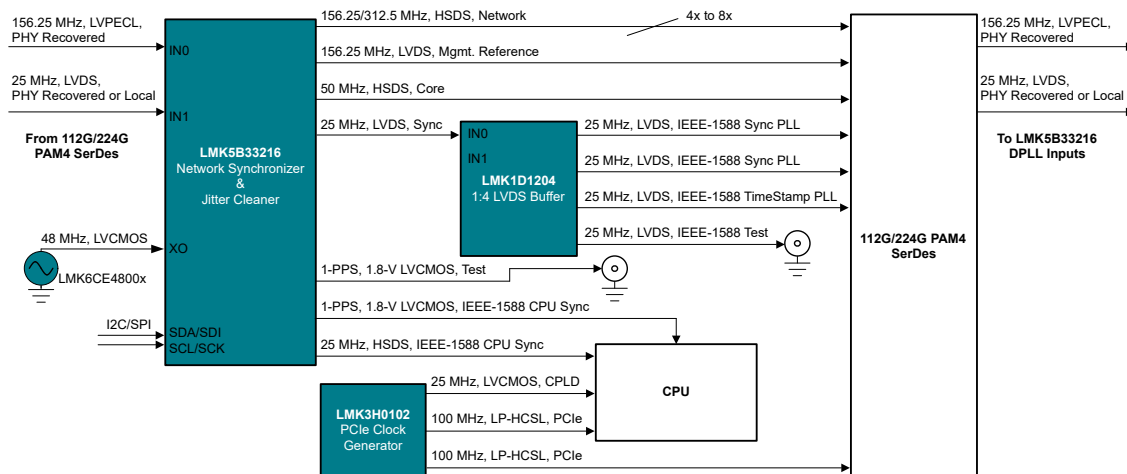


Figure 3-1. Full System Design for 400G and 800G Switch Applications

3.1 Frequency and Phase Adjustments

To support IEEE-1588 PTP or other clock steering applications, each DPLL allows precise frequency and phase adjustments through register, or pin control, by using a Digitally-Controlled Oscillator (DCO), as shown in Figure 3-2. Adjustments with less than 1ppt (part per trillion) frequency resolution are supported by the DPLL DCO. The DPLL DCO feature allows increments and decrements to the numerator of the DPLL fractional N-divider. Such frequency adjustments are effectively propagated through the APLL domain and onto the output clocks or any cascaded DPLL/APLL domains.

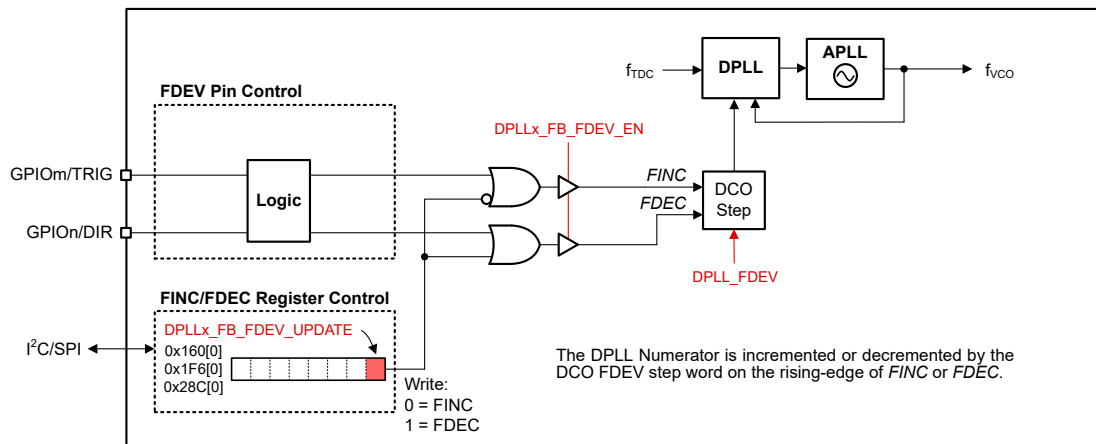
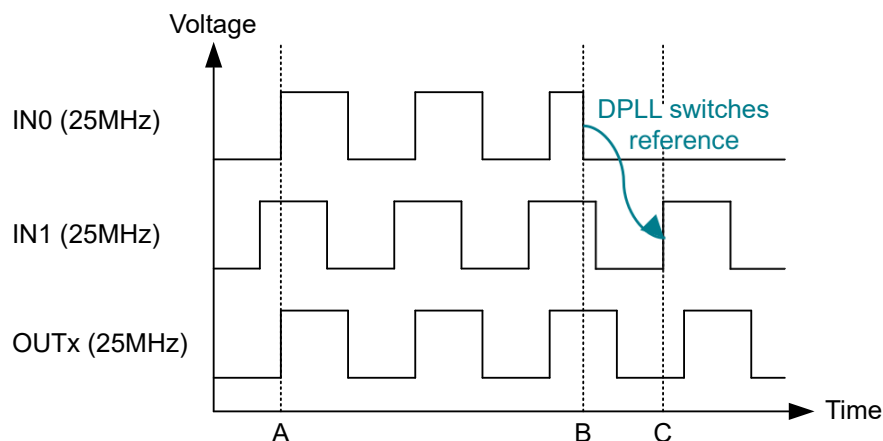


Figure 3-2. Overview of the DPLL DCO

3.2 Input Reference Switching

Each DPLL supports hitless reference switching with a minimal phase hit through a phase cancellation scheme (also known as phase buildout) and offers an optional phase slew control feature.

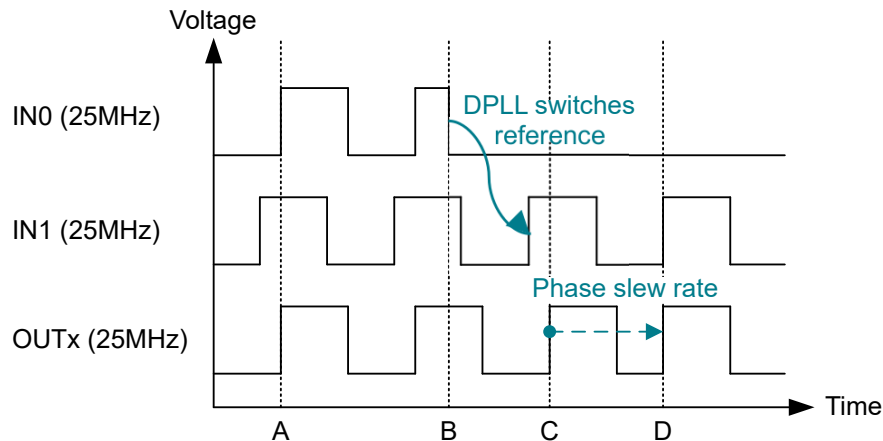
The outputs maintain the same phase with hitless switching, with only minor disturbance from before the reference switch, to after the reference switch, as illustrated in Figure 3-3. Without hitless switching, there is a risk of the outputs getting phase hits, which can propagate to the downstream clocks and cause misalignment across data packets. The phase hit is equal to the phase offset between the two reference clocks.



- A: OUTx is locked to IN0 with the same phase.
- B: IN0 is lost and the DPLL switches reference from IN0 to IN1 without affecting the phase on OUTx.
- C: OUTx is locked to IN1 without changing the phase.

Figure 3-3. Phase Cancellation Enabled for Switching Reference With Minor Phase Hits

The outputs gradually adjust phase after a hitless switch when using the phase slew control feature. The phase transitions from the original to the new reference at a rate defined by the programmed phase slew rate. This feature is demonstrated in Figure 3-4.



- A. OUTx is locked to IN0 with the same phase.
- B. IN0 is lost and the DPLL switches reference from IN0 to IN1.
- C. The phase of OUTx gradually changes at a rate defined by the programmed phase slew rate.
- D. OUTx is locked to IN1 with the same phase.

Figure 3-4. Phase Slew Control Enabled to Provide a Steady Output Phase Change When Switching Inputs.

3.3 Holdover

Holdover occurs when the DPLL reference clocks become unavailable, as seen in Figure 3-5. During holdover, the APLL numerator is set by the DPLL tuning word history, which determines the output frequency accuracy upon entry into holdover—any error on entry is considered a short-term holdover error. The tuning word history can be configured as either the accumulated averaged reference phase history, a user-specified value, or the last APLL numerator value.

The quality of the external oscillator (provided to XO input) determines the long-term frequency stability and accuracy of the output frequency. Over time, temperature fluctuations impact the frequency accuracy of the output clocks. Hence, the decision to choose between an XO, TCXO, or OCXO lies primarily on the long-term holdover requirements of the system.




Figure 3-5. Holdover Functionality

- A. Reference failure is detected by the device as the DPLL references (INx) are lost.
- B. The DPLL enters holdover and outputs remain phase or frequency locked with a small frequency delta.
- C. The output frequency drifts; the long-term holdover stability is based on the XO/TCXO/OCXO input accuracy.

For applications requiring output frequency correction during long-term holdover, the output frequency can be corrected by making APLL DCO adjustments through software to compensate for temperature variations, as shown in Figure 3-6.

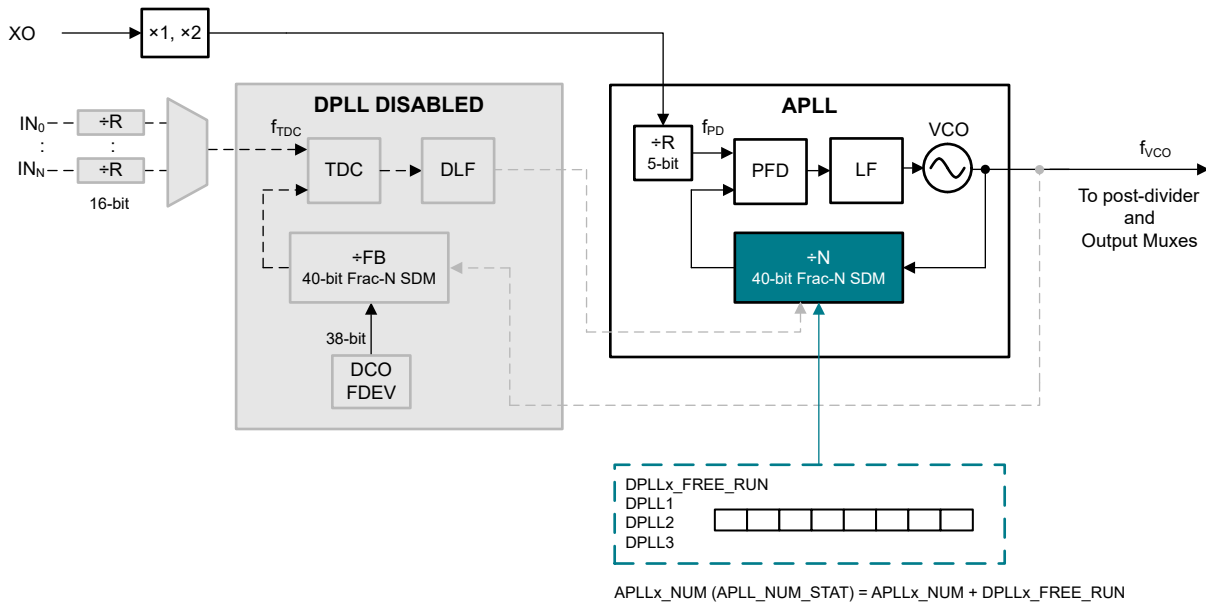


Figure 3-6. APLL DCO Operation During Holdover

3.4 Zero-Delay Mode

Each DPLL supports Zero Delay Mode (ZDM) to achieve a deterministic phase relationship between the DPLL reference clock and the ZDM feedback output clock at every boot-up or software reset. All output clocks sourced from a ZDM-configured DPLL become phase aligned through the synchronization (SYNC) feature. A zero-phase delay is attainable across all clocks by inserting analog or digital delays available on the LMK5B33216.

Figure 3-7 demonstrates how select outputs, such as OUT0, can internally feed back to any DPLL as a zero-delay output clock. For additional details on ZDM theory, refer to [Multi-Clock Synchronization](#).




Figure 3-7. DPLL ZDM Synchronization Between Reference Input and OUT0

4 LMK5B33216 Performance

The LMK5B33216 meets the challenging requirements of advanced high-speed communication networks such as 800 Gbps systems.

Output clocks generated from the BAW APLL achieve excellent output RMS jitter and phase noise performance for these industries. The VCBO of the BAW APLL operates at 2500MHz (± 100 ppm) and can be divided down to output 312.5MHz clocks with 42fs typical and 60fs maximum RMS jitter (12kHz to 20MHz) regardless of the DPLL reference clock frequency and jitter characteristics. The remaining APLLs are LC-VCOs and can be used for additional frequency domains if other phase domains are required outside the Ethernet domain, for example, when the frequency cannot be sourced from the BAW APLL.

4.1 Phase Noise Profile

The output phase noise profile of a network synchronizer like the LMK5B33216 can be broken down into four dominating regions:

- Below the DPLL LBW
- Between the DPLL and APLL LBWs
- Above the APLL LBW
- Output noise floor

The quality of the DPLL reference and XO, TCXO, and OCXO inputs impact the close-in phase noise of an output clock. The DPLL LBW is a low-pass filter to the DPLL reference clocks. The DPLL reference primarily impacts the output noise profile at carrier offsets less than the DPLL LBW. The DPLL can be configured with a narrower LBW (such as 10Hz) for applications where close-in phase noise matters and noisy references are used as the DPLL inputs. Otherwise, either 10Hz or 100Hz are common settings that can be used since the DPLL reference does not impact the jitter at carrier offsets from 12kHz to 20MHz. [Figure 4-2](#) and [Figure 4-2](#) are examples of phase noise plots with outputs sourced from the VCBO. These plots demonstrate the negligible effect of jitter across 12kHz to 20MHz with different DPLL LBWs.

Above the DPLL LBW, the DPLL reference is attenuated and a combination of the XO/TCXO/OCXO input and the APLL noise dominates up until the APLL LBW. To take advantage of the excellent VCBO performance, the BAW APLL is configured with a narrow LBW, typically 3kHz to 5kHz; as a result, the VCBO dominates the phase noise in the carrier offset range of 8kHz to around 400kHz. Lastly, the noise floor is set by the output buffer, beginning at approximately 1MHz carrier offset.

Thanks to the ultra-low noise VCBO, a low frequency XO (such as 12.8MHz) can be used for the LMK5B33216 without significantly impacting the total RMS jitter, thereby lowering the total design cost.

[Figure 4-1](#) provides a summarized illustration of how each region affects the phase noise of an output clock.




Figure 4-1. General Phase Noise Plot

- DPLL loop bandwidth can be set between 1mHz to 4kHz.
- APLL loop bandwidth for VCBO can be set between 1kHz and 10kHz.
- APLL loop bandwidth for LC VCO can be set between 100kHz and 1MHz.

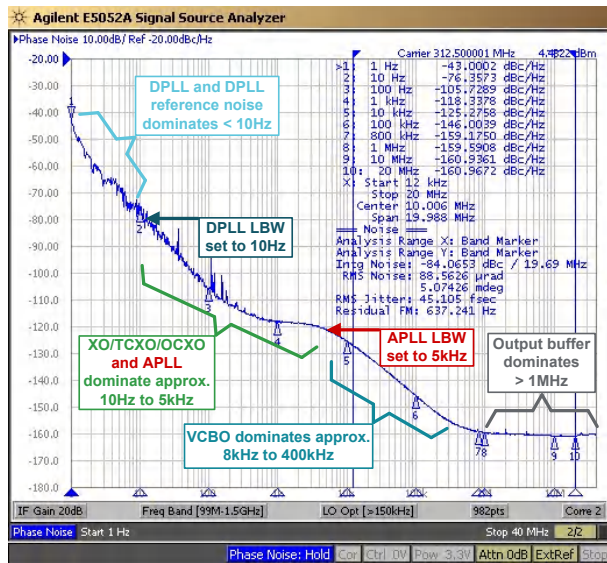


Figure 4-2. Phase Noise Breakdown of an LMK5B33216 Output Using a 10Hz DPLL LBW




Figure 4-3. Phase Noise Breakdown of an LMK5B33216 Output Using a 100Hz DPLL LBW

4.2 RMS Jitter

Refer to [Table 4-1](#) for additional RMS jitter data across various integration ranges.

Table 4-1. Jitter Performance of the LMK5B33216 Across Different Integration Ranges

Output Clock Frequency [MHz]	RMS Jitter (typ.) [fs] ⁽¹⁾	Jitter Integration Range [MHz]
156.25 ⁽²⁾	47	0.012 to 20
	26	0.100 to 5
	31	0.750 to 10
	71	0.200 to 50
	15	2 to 4
312.5 ⁽³⁾	39	4 to 20
	42	0.012 to 20
	19	0.100 to 5
	21	0.750 to 10
625 ⁽⁴⁾	47	0.200 to 50
	10	2 to 4
	26	4 to 20
	35	0.012 to 20
	12	0.100 to 5
	11	0.750 to 10
	25	0.200 to 50
	5	2 to 4
	13	4 to 20

- (1) The measurements were taken using an XO input frequency of 48MHz with all LMK5B33216 outputs set to the same frequency with an output swing (VOD) ≥ 800mV.
- (2) The APLL3 post divider is 16.
- (3) The APLL3 post divider is 8.
- (4) The APLL3 post divider is 4.

5 Summary

The role of clocks are critical in supporting the high-speed data links using the 112G PAM-4 SerDes. Clocks must deliver low jitter and phase noise signals to maintain integrity of the PAM-4 symbols and minimize errors. TI's BAW technology enables network synchronizer designs with greater system noise margins (primarily the LMK5B33216) which meets the stringent SerDes jitter and phase noise requirements.

Precise frequency and phase adjustments are possible with integrated DCOs and this supports IEEE-1588 PTP designs. The network synchronizer outputs can be configured to lock between different domains (including SyncE and PTP) and hitless switching can occur between references with minimal disturbance to the downstream clocks. The LMK5B33216 leverages a DPLL history feature which sets frequency accuracy to minimize the short-term holdover errors in the event reference clocks are lost.

6 References

- Texas Instruments, [LMK5B33216](#), product page
- Texas Instruments, [LMK5B33414](#), product page
- Texas Instruments, [LMK6Cx](#), product page
- Texas Instruments, [LMK1D1204](#), product page
- Texas Instruments, [LMK3H0102](#), product page
- Texas Instruments, [Understanding Clocking Needs for High-Speed 56G PAM-4 Serial Links](#), application note
- Texas Instruments, [TI BAW Technology Enables Ultra-Low Jitter Clocks for High-Speed Networks](#), white paper
- Texas Instruments, [BAW Oscillator Solutions for Network Interface Card](#), application brief
- Texas Instruments, [Multi-Clock Synchronization](#), application note

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATA SHEETS), 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, regulatory 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](#) or other applicable terms available either on [ti.com](https://www.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.

TI objects to and rejects any additional or different terms you may have proposed.

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2024, Texas Instruments Incorporated