

TI Designs High Speed JESD204B Link Latency Design Using a High Speed ADC and FPGA



TI Designs High Speed

TI Designs High Speed designs are analog solutions created by TI's analog experts. Reference Designs offer the theory, part selection, simulation, complete PCB schematic & layout, bill of materials, and measured performance of useful circuits. Circuit modifications that help to meet alternate design goals are also discussed.

Design Resources


- [Design Files](#)
- [OPA2333](#)
- [OPA2188](#)

- Design Folder
- Product Folder
- Product Folder

JESD204B Interface Link Latency

JESD204B links are the latest trend in data-converter digital interfaces. These links take advantage of high-speed serial-digital technology to offer many compelling benefits including improved channel densities. This reference design addresses one of the challenges of adopting the new interface: understanding and designing the link latency. Understanding this discussion and example will enable a system designer to:

- Ensure deterministic latency across the JESD204B link
- Understand the tradeoff between link latency and tolerance to link delay variation
- Use a formulaic and procedure-based approach to design the link latency
- Implement a JESD204B link using Texas Instrument's ADC16DX370 or LM97937 ADC and a Xilinx Kintex 7 FPGA


[ASK Our Analog Experts](#)
[WebBench Calculator Tools](#)

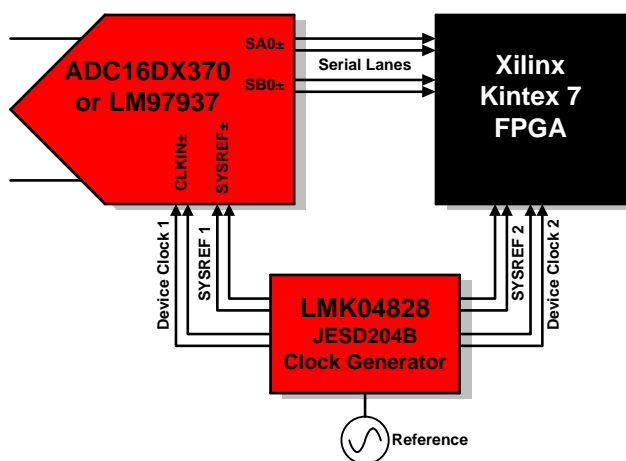


Figure 1. JESD204B Interface Implementation With the LM97937 and LMK04828



An IMPORTANT NOTICE at the end of this TI reference design addresses authorized use, intellectual property matters and other important disclaimers and information.

1 Introduction

The latency through a signal path describes the length of time for the signal to propagate between two points in the path. Latency can be a critical design detail in applications that are sensitive to long latencies or latency variation. Applications involving gain control loops can become unstable if the latency is not well controlled. Systems requiring the synchronization of multiple data converter devices can have conversion-phase misalignments which result in sub-optimal signal processing if the latency is not properly designed.

Serialized data-converter interfaces offer a tremendous benefit by simplifying the hardware interface to a logic device, such as an FPGA or ASIC, and by saving PCB space. However, these interfaces typically have longer latencies than a parallel-type interface and the latencies are potentially unpredictable. With the introduction of the most recent revision of the JESD204 standard (JESD204B), system architecture provisions are included that achieve known and constant (deterministic) latency across the serial link when certain requirements are met. This reference design describes the requirements for achieving deterministic latency with a Subclass 1 device. The configuration of the link and impact of the link on the total latency are also discussed. This document also outlines a procedure for achieving a desired latency with the ADC16DX370, a 16-bit 370-MSPS analog-to-digital converter (ADC) with a JESD204B serial interface transferring data at 7.4 Gb/s/lane as shown in the example implementation of [Figure 1](#). The analysis also applies to the LM97937, a 9-bit 370-MSPS ADC with a similar JESD204B interface, as well as to other data converters from Texas Instruments in a general sense.

2 Latency in the Serial Link

The term *latency* can be used in many ways to describe different interrelated concepts in systems that include serial links. These concepts include the following list that are represented on the simplified data-flow block diagram shown in [Figure 2](#).

Sample-to-Serial Out (S2SO) Latency—S2SO is the latency from when the signal is sampled at the ADC input until the sample appears in the serial stream at the ADC (TX) output.

Sample-to-Parallel Out (S2PO) Latency—S2PO is the latency from when the signal is sampled at the ADC input until the sample is released on the parallel bus at the receiving device (RX).

Link Latency—Link latency is the latency from when the sampled parallel data is input to the serializer at the transmitter (ADC) until the same data is available in parallel form at the receiver. This simplified description is more accurately described relative to local multi-frame clock (LMFC) boundaries in the JESD204B standard.

Deterministic Latency—Deterministic latency is the concept that the link latency of a JESD204B link is always known and constant across varying system conditions including supply voltage, temperature, and system reset.

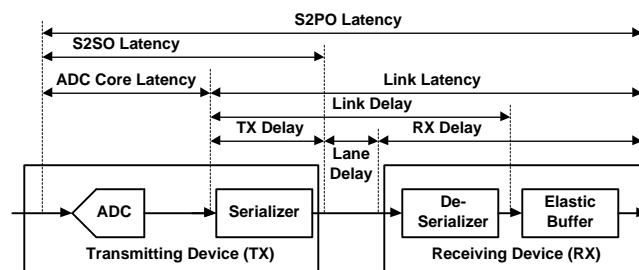


Figure 2. Block Diagram of Simplified Latencies in a System Using the LM97937

Ultimately, the S2PO latency is most important at the system level because it affects the processing of data through the signal path. The S2SO Latency is not a useful value in the JESD204B system design because it can vary under different conditions. S2SO latency is mentioned only because the sample-to-data-out latency is a familiar concept for ADCs with parallel output data busses. The link latency is important as a serial-link design parameter to ensure that deterministic latency is achieved.

The JESD204B standard defines multiple subclasses that ensure that the latency across the serial data link is always deterministic. Most Texas Instruments data-converter devices with JESD204B interfaces implement Subclass 1 because of the relative ease of meeting critical timing requirements that must be met to synchronize the serial data receivers (RX) and transmitters (TX) together and achieve deterministic latency.

3 Requirements for Achieving Deterministic Latency

As extracted from the JESD204B standard, the serial link must meet the following requirements across supported conditions to ensure deterministic latency of a Subclass-1 link.

1. The internal LMFCs of the TX and RX must have a deterministic phase relationship between each other.
2. Both the time length of a multi-frame ($K \times T_{\text{FRAME}}$) and the quantity $\text{RBD} \times T_{\text{FRAME}}$ must be larger than the maximum possible delay across any link.
3. The elastic buffer at the receiver must accommodate a buffering time greater than the delays in the link. The buffering time ($\text{RBD} \times T_{\text{FRAME}}$) can be between T_{FRAME} and $K \times T_{\text{FRAME}}$.

3.1 Requirement 1: Aligning the LMFCs and Meeting SYSREF Setup and Hold Requirements

The phase of the SYSREF signals distributed to each device determines the phase alignment of the LMFC internal to each device. The first requirement is that the phase relationships between the LMFCs of each device do not change over time, nor do they differ each time the link must be re-initialized. To achieve the first requirement, the SYSREF signals distributed to each device must meet setup and hold requirements relative to their companion device clocks. This is particularly critical when very high frequency device clocks are distributed.

Techniques for meeting the timing requirements include replication of the routing for the device clocks and the SYSREF clocks to their destinations, and using duplicate drivers and load circuits with similar electrical characteristics.

3.2 Requirement 2: Ensure the LMFC Period is Longer than the Link Delay

In the simplest terms, the TX, RX, and the channel in between can have individual latencies that vary across conditions. However, the sum of all the latencies must be constant and known. A constant sum of all the latencies is achieved using an adaptable (elastic) buffer at the receiver and a very low frequency reference signal, namely the LMFC clock, to avoid any ambiguity in the alignment of the data frames output on the serial lanes. To avoid phase ambiguity, the LMFC clock period must be longer than sum of all the potential delays across the link. Therefore the K parameter (integer value representing the number of frames per multi-frame in the serial stream) must be set large enough so that the LMFC period is longer than the maximum sum of delays across the link.

Larger K values result in increased total latency through the link when deterministic latency is achieved. Therefore, minimizing the value of this parameter can be beneficial. For system simplicity, the desired K value can also be influenced by the ability to generate SYSREF pulses, and therefore choosing 2^N values such as 16 or 32 is more convenient.

3.3 Requirement 3: Ensure the RX Elastic Buffer is Large Enough to Buffer the Data

The elastic buffer must be large enough to ensure that all time-synchronized critical data is received at the various serial lanes and buffered before releasing data from the buffer. The release-buffer delay (RBD) parameter controls the time at which data is released from the buffer at the receiver.

In the simplest system implementation, RBD is made equal to K which causes the total S2PO link latency to equal $K \times T_{\text{FRAME}}$. If the decided K value results in a latency length that is unacceptable, then the RBD parameter can be reduced from the value K to a lesser value to reduce the amount of time the data spends in the elastic buffer which causes the total S2PO link latency to equal $\text{RBD} \times T_{\text{FRAME}}$.

4 Dissecting the Link Latency and Link Delay

The *link latency* is the sum of the TX delay, the lane propagation delay, the RX delay, and the RX-TX LMFC phase skew. The TX delay (t_{TX_SER}) is the time from when data appears at the serializer input on an LMFC boundary until the data appears on the serial lanes. The lane delay (t_{LANE}) is the propagation time period from the transmitter to the receiver across the channel. The RX delay is the delay from when the data arrives at the receiver until the data is available as a parallel word and is composed of two components ($t_{RX_BUFF} + t_{RX_DESER}$):

1. The buffering delay of the elastic buffer of the receiver
2. All other receiver delays required to de-serialize and output the data as a parallel word

The RX-TX LMFC phase skew is the absolute time difference ($t_{TX_LMFC} - t_{RX_LMFC}$) where t_{TX_LMFC} and t_{RX_LMFC} represent the SYSREF to LMFC delays in the paths from the common clock reference to the transmitter and receiver respectively.

The *link delay* includes the TX delay, lane propagation delay, and the RX delay excluding the RX buffering delay. The RX Buffering delay is an intentional addition used to counteract the variation in the link delay and make the link latency constant.

This document has a particular use of the terms *latency* and *delay* in regards to the JESD204 link. Latency is used to describe the propagation time between two points in the system and always includes the entire link. Delay is used to describe the propagation time across only a portion of the serial link. This distinction is made to avoid confusion about whether RX elastic buffering is included in the quantity.

The deterministic-latency of [Requirement 2](#) and [Requirement 3](#) is calculated with [Equation 1](#).

$$t_{TX_SER} + t_{LANE} + t_{RX_DESER} + (t_{TX_LMFC} - t_{RX_LMFC}) < RBD \times T_{frame} \leq K \times T_{frame} \quad (1)$$

When all of the requirements are satisfied and the parameters are selected, the link latency can be calculated. The *relative link latency* through the link is calculated with [Equation 2](#). Relative link latency is described in units of frame-clock cycles and is referred to in the standard as the *delay across the link*, although this phrase can lead to confusion about whether RX buffering is included. This value is relative to LMFC boundaries and does not account for skew between the RX and TX LMFCs.

$$t_{LINK_LAT_REL} = t_{TX_SER} + t_{LANE} + t_{RX_DESER} + t_{RX_BUFF} = RBD \times T_{frame} \quad (2)$$

The *absolute link latency* through the link includes the skew between the RX and TX boundaries calculated with Equation 3. Note that additional delay in the transmitter SYSREF to LMFC path reduces the absolute latency but increases the required RBD at the receiver. Figure 3 shows the involved latency parameters.

$$t_{\text{LINK_LAT_ABS}} = \text{RBD} \times T_{\text{frame}} + (t_{\text{RX_LMFC}} - t_{\text{TX_LMFC}}) \quad (3)$$

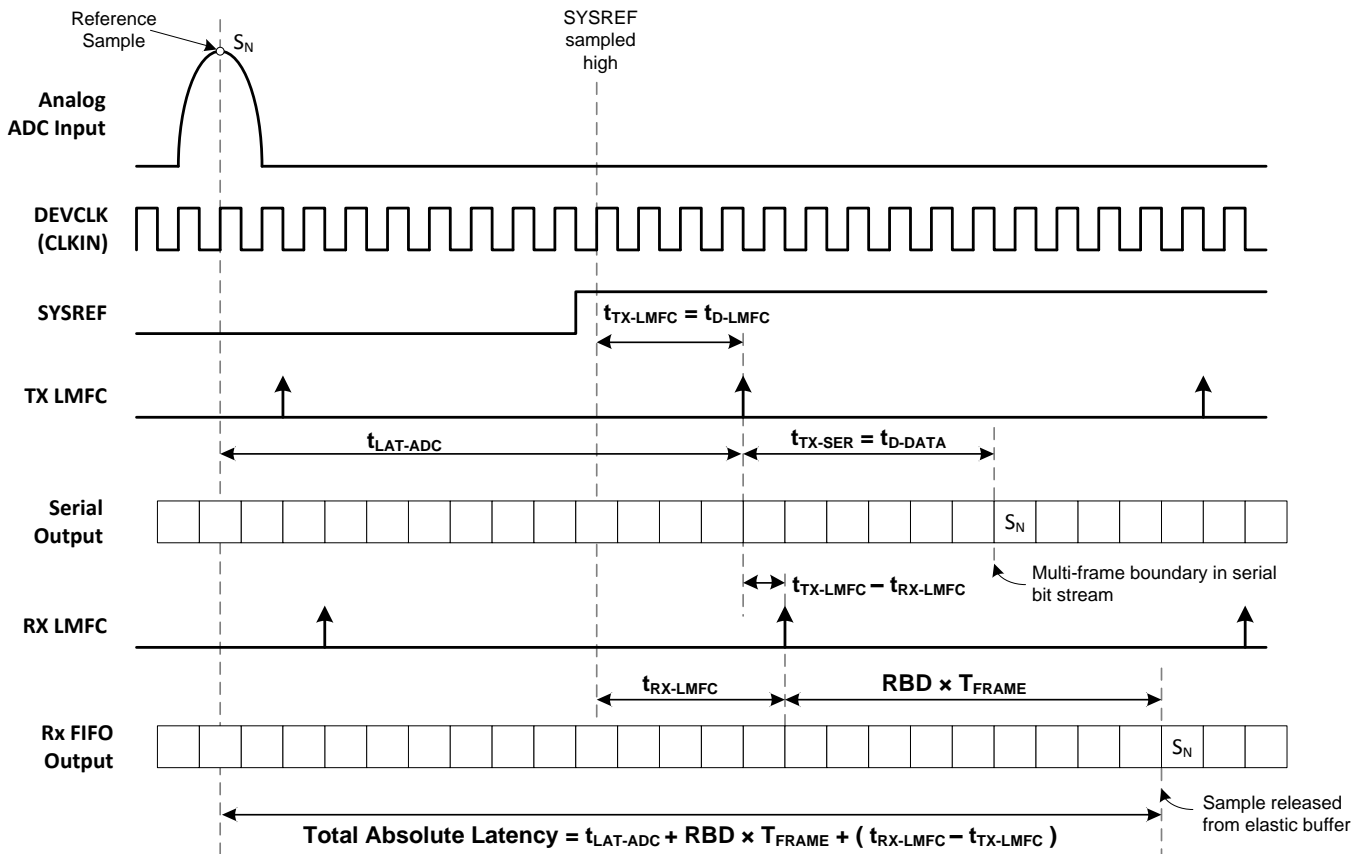


Figure 3. Latency Timing Diagram

A tradeoff is presented when the value of K is set. A large value of K allows for a large amount of delay in various parts of the link. However, a large value of K also increases the total latency through the link if RBD equals K in the simplest case. Setting RBD to a value less than K allows the total latency to be fine-tuned and independent of the K value. Ultimately, the minimum latency must be slightly greater than the link delay.

The following lists the design procedure for setting K and RBD, and calculating the system latency:

1. Determine the maximum values of the delays $t_{\text{TX-SER}}$, t_{LANE} , $t_{\text{RX-DESER}}$.
2. Determine the skew between the RX and TX SYSREF-to-LMFC paths ($t_{\text{TX_LMFC}} - t_{\text{RX_LMFC}}$).
3. Set K to satisfy Equation 1.
4. Set RBD equal to K. To reduce the total link latency, RBD can also be set to a value less than K.
5. Calculate the absolute link latency.
6. Add the absolute link latency to the latency of the ADC core ($t_{\text{LAT_ADC}}$) to determine the total latency from sampling instant to parallel data available at the receiver.

4.1 Latency Design Example

The following example demonstrates the latency design of a JESD204B link between the LM97937 and Xilinx Kintex 7 FPGA.

The previous design procedure can be followed to determine the latency across a link that includes an ADC transmitter and logic device receiver, although the naming convention between the generalized name and the datasheet parameter may differ. Table 1 lists the required parameters for the LM97937 and Kintex 7 FPGA latency design. Note that equations given throughout this text are equated in units of seconds, but they are evaluated in units of frame clock cycles which do not depend on the frame rate. When evaluating the equations, the T_{frame} parameter is factored out.

Table 1. Delay Parameters for the LM97937 and Kintex 7 FPGA

DEVICE	GENERALIZED NAME	DATASHEET PARAMETER	VALUE [Frame Clock Cycles]
LM97937 or ADC16DX370 (TX)	t_{TX_SER}	t_{D_DATA}	6 ± 1
	t_{TX_LMFC}	t_{D_LMFC}	3.5
	—	t_{LAT_ADC}	14.5 (LM97937) 12.5 (ADC16DX370)
Xilinx Kintex 7 (JESD204 IP Core v5.1)	t_{RX_DESER}	T_{RXIN}	92 ± 2
	t_{RX_LMFC}	T_{RXLMFC}	28
Other	t_{LANE}	—	± 1

In this case, the link delay ($t_{TX_SER} + t_{LANE} + t_{RX_DESER}$) is very large and exceeds the standard maximum for the K parameter, $K = 32$. The designer has two options, either increasing the K value to satisfy Requirement 2 or choosing the *release moment* of the data to be used in a later multi-frame period. The first option ensures deterministic latency for a very wide range of delay variation but $K > 32$ is not supported by all devices, including the LM97937. In this example, the release moment is positioned in a later multi-frame period as shown in Figure 4.

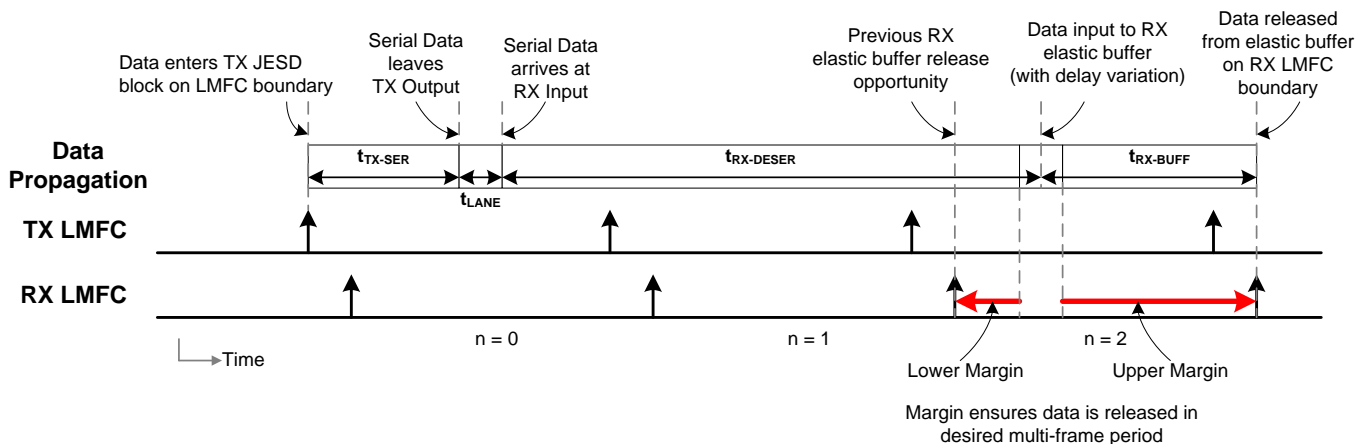


Figure 4. Elastic-Buffer Release Moment in Later Multi-Frame Period When Link Delay is Greater Than K

Releasing data in a later multi-frame period changes the inequality of Equation 1 to the inequality bounds of Equation 4 and Equation 5 where n is an integer value greater than 0 and describes the later multi-frame period in which the data is released.

$$((n - 1) \times K + RBD) \times T_{\text{frame}} \leq n \times K \times T_{\text{frame}} < t_{\text{TX_SER}} + t_{\text{LANE}} + t_{\text{RX_DESER}} + (t_{\text{TX_LMFC}} - t_{\text{RX_LMFC}}) \quad (4)$$

$$t_{\text{TX_SER}} + t_{\text{LANE}} + t_{\text{RX_DESER}} + (t_{\text{TX_LMFC}} - t_{\text{RX_LMFC}}) < (n \times K + RBD) \times T_{\text{frame}} \leq (n + 1) \times K \times T_{\text{frame}} \quad (5)$$


Applying the given values to the inequalities, assuming that K is set to 32 to maximize the tolerable delay variation, and assuming the data is released on an LMFC boundary ($RBD = K$) results in Equation 6 and Equation 7.

$$n \times 32 < (6 \pm 1) + (\pm 1) + (92 \pm 2) + (3.5 - 28) = 73.5 \pm 4 \quad (6)$$

$$(6 \pm 1) + (\pm 1) + (92 \pm 2) + (3.5 - 28) = 73.5 \pm 4 < (n + 1) \times 32 \quad (7)$$

The solution $n = 2$ is an acceptable solution in this case, but some cases may not have a solution because the link delay variation violates one of the bounds. The available methods to achieve a solution when the bounds are violated include changing the K value, modifying RBD to be less than K , or changing $SYSREF$ phases and the relative alignment of the LMFCs. Modifying K to be a different value is a straightforward method but is not always an available system-level option and is not described further here.

Changing RBD to a value less than K shifts the release moment to an earlier point in time. In this example where the release moment is in a later multi-frame boundary and RBD is set less than K , both the intended release moment and the previous release moment are shifted earlier in time as shown in Figure 5. This shift allows the designer to improve the margin between the time the data enters the elastic buffer and the LMFC boundary times.



Set $RBD < K$ to shift the release moment and improve the time margin relative to the RX LMFC boundary

Figure 5. Shifting the Release Moment to Improve Timing Margin using $RBD < K$

Assume that RBD is reduced from 32 down to 24. The inequalities for this case are shown in Equation 8 and Equation 9 and have plenty of margin in the presence of the delay variation for $RBD = 24$ and $n = 2$. One can select this RBD value to maximize the upper and lower margin although RBD can be reduced further to reduce the total latency.

$$((n - 1) \times K + RBD) = ((n - 1) \times 32 + 24) < 73.5 \pm 4 \quad (8)$$

$$73.5 \pm 4 < (n \times K + RBD) = (n \times 32 + 24) \quad (9)$$

Manipulating the phase of the $SYSREF$ signals delivered to the receiver and transmitter can also be used to shift the elastic buffer release time to present a solution with adequate margins. This method changes the alignment of the RX and TX LMFCs and is useful in systems where a device may not have the flexibility (or feature) to set RBD to a value less than K . This is the recommended solution for the systems involving the Xilinx JESD204 IP Core. This method changes the analysis by adding delay to the $t_{\text{TX_LMFC}}$ or $t_{\text{RX_LMFC}}$ parameters mentioned above. Figure 6 demonstrates how adding delay in the TX LMFC path shifts the relationship between the LMFCs and improves the margin similar to the method of setting $RBD < K$. Ultimately, this influences the total latency by modifying the quantity $(t_{\text{TX_LMFC}} - t_{\text{RX_LMFC}})$.

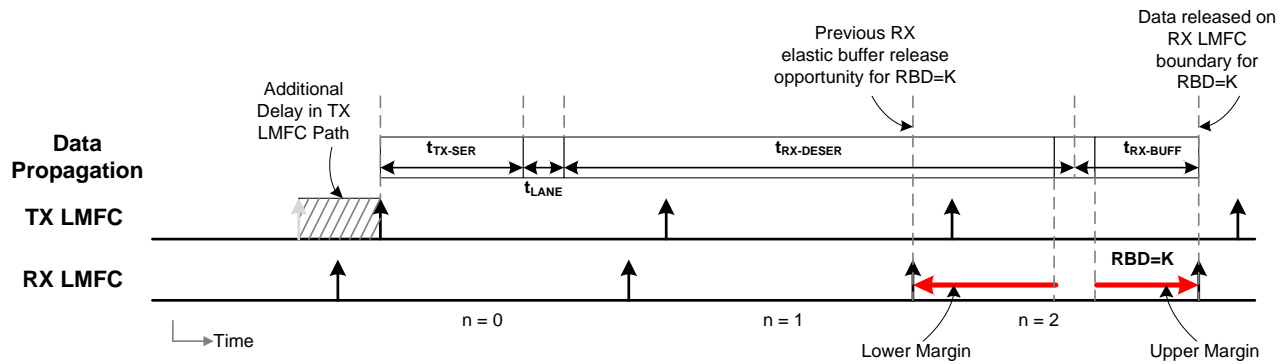


Figure 6. Shifting the Release Moment to Improve Timing Margin Using SYSREF Delays

Continuing with the $RBD < K$ method and the link-delay inequalities satisfied for $n = 2$, the total link delay can be calculated according to a modified version of Equation 3 to account for releasing in a later multi-frame boundary which is calculated with Equation 10. The results are calculated in Equation 11 as 112.5 frame-clock cycles.

$$t_{LINK_LAT_ABS} = (n \times K + RBD) \times T_{frame} + (t_{RX_LMFC} - t_{TX_LMFC}) \quad (10)$$

$$t_{LINK_LAT_ABS} = 2 \times 32 + 24 + (28 - 3.5) = 112.5 \quad (11)$$

The absolute link latency only accounts for the latency of the link. To find the latency from the ADC sampling instant to the output of the link, the latency of the ADC core ($t_{LAT_ADC} = 12.5$ for the ADC16DX370) is added to the link latency for a total S2PO latency of $112.5 + 12.5 = 125$ frame-clock cycles.

4.2 Tips for Implementing the JESD204B Link

Achieving deterministic latency and accurately calculating the total latency requires great attention to detail when clocking the device and when implementing the Xilinx JESD204B core. The following tips can be used to ensure a robust, accurate design. The tips related to the Xilinx JESD204 core apply to version 5.1.

- Pay close attention to the synchronization of the various device clocks and SYSREF signals distributed from the common clock generator (such as Texas Instrument's LMK04828) and any delays that may have been added to the signals. The phase relationship between these signals may create addition skew between the RX and TX LMFC's.
- The Xilinx JESD204 logic core registers the SYSREF input on the falling edge of the device clock. This influences the alignment of the RX and TX LMFC.
- The Xilinx JESD204 logic core has timing restrictions that may limit the frequency of the device clock going to the FPGA that is used to sample the SYSREF signal. Dividing the device clock down within the FPGA is not an acceptable solution unless the divided clock phase can be deterministic.
- Advance clocking schemes, different from the default condition of the Xilinx JESD204 core, may be required to ensure that SYSREF meets setup and hold requirements relative to the input device clock.
- Unpacking the data samples from the octets output from the Xilinx JESD204 core requires additional clock cycles and may require changing clock domains. These additional cycles must be accounted for in the total latency.
- The Xilinx JESD204 core outputs four (4) octets during a clock cycle. The earliest arriving octet has more added latency than the latest arriving octet.
- The 'RX Buffer Delay' parameter in the Xilinx JESD204 core provides a function similar to the RBD parameter in this text and described in the JESD204B standard, but their meanings are different. RX Buffer Delay represents the number of frame clock cycles that the data should be released before the LMFC boundary. Therefore, setting 'RX Buffer Delay' = 4 is equivalent to RBD=28 for a value of K=32.

4.3 Multi-Device Synchronization

Multi-device synchronization refers to matching the latency through multiple ADC signal paths, aligning the sampling instant between ADCs, or both.

Matching latencies in systems with multiple LM97937 devices requires deterministic latency and knowledge of the S2PO latency as previously discussed. If the requirements needed to achieve deterministic latency are satisfied, then the RBD values of each link can be adjusted or additional buffering can be added in logic to match the latencies between links. When multiple links are synchronized together, the K and RBD values must accommodate the link with the longest link delay and all the devices in the synchronized system must share the same LMFC frequency. Note that if the elastic-buffer release moment is designed to release in the multi-frame period later than the first instance, care must be taken to ensure that all similar devices also release in the later multi-frame period.

Synchronizing the sampling instants between multiple LM97937 devices requires matching the distribution of the critical device clocks to each device CLKIN input such that the delay from the common clock reference to the CLKIN input is exactly the same and does not vary significantly. Note also that if the link latency is not matched between the synchronized devices, the phase aligned samples become offset in time because of the latency mismatch.

Appropriate device clock and SYSREF pairs must be generated for JESD204B devices in the system to ensure synchronization. Texas Instrument's LMK04828 low-noise clock generator and jitter cleaner accommodates up to seven different JESD204B devices. All device clock and SYSREF pair outputs can be clocked synchronously with programmable phases to achieve well-controlled distribution of the SYSREF signals and fine tuning of the sampling instants across multiple similar ADC devices. Features of this clocking chip also allow for synchronization of multiple LMK04828 devices together so that applications with a very large channel counts can be realized.

5 Conclusion

The robustness of a JESD204B high-speed serial link against latency variation is dependent upon meeting the deterministic latency criteria. This reference design outlines the criteria in detail, provides equations to quantify the conditions, and outlines procedures in a step-by-step guide to ensure that the criteria are met. The example interface between the LM97937 and Kintex 7 FPGA explained in this reference design demonstrates identification of the individual delays, calculation of the total link delay, appropriate selection of the K and RBD parameters, and final calculation of the total link latency.

IMPORTANT NOTICE FOR TI REFERENCE DESIGNS

Texas Instruments Incorporated ("TI") reference designs are solely intended to assist designers ("Buyers") who are developing systems that incorporate TI semiconductor products (also referred to herein as "components"). Buyer understands and agrees that Buyer remains responsible for using its independent analysis, evaluation and judgment in designing Buyer's systems and products.

TI reference designs have been created using standard laboratory conditions and engineering practices. **TI has not conducted any testing other than that specifically described in the published documentation for a particular reference design.** TI may make corrections, enhancements, improvements and other changes to its reference designs.

Buyers are authorized to use TI reference designs with the TI component(s) identified in each particular reference design and to modify the reference design in the development of their end products. HOWEVER, NO OTHER LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE TO ANY OTHER TI INTELLECTUAL PROPERTY RIGHT, AND NO LICENSE TO ANY THIRD PARTY TECHNOLOGY OR INTELLECTUAL PROPERTY RIGHT, IS GRANTED HEREIN, including but not limited to any patent right, copyright, mask work right, or other intellectual property right relating to any combination, machine, or process in which TI components or services are used. Information published by TI regarding third-party products or services does not constitute a license to use such products or services, or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

TI REFERENCE DESIGNS ARE PROVIDED "AS IS". TI MAKES NO WARRANTIES OR REPRESENTATIONS WITH REGARD TO THE REFERENCE DESIGNS OR USE OF THE REFERENCE DESIGNS, EXPRESS, IMPLIED OR STATUTORY, INCLUDING ACCURACY OR COMPLETENESS. TI DISCLAIMS ANY WARRANTY OF TITLE AND ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, QUIET ENJOYMENT, QUIET POSSESSION, AND NON-INFRINGEMENT OF ANY THIRD PARTY INTELLECTUAL PROPERTY RIGHTS WITH REGARD TO TI REFERENCE DESIGNS OR USE THEREOF. TI SHALL NOT BE LIABLE FOR AND SHALL NOT DEFEND OR INDEMNIFY BUYERS AGAINST ANY THIRD PARTY INFRINGEMENT CLAIM THAT RELATES TO OR IS BASED ON A COMBINATION OF COMPONENTS PROVIDED IN A TI REFERENCE DESIGN. IN NO EVENT SHALL TI BE LIABLE FOR ANY ACTUAL, SPECIAL, INCIDENTAL, CONSEQUENTIAL OR INDIRECT DAMAGES, HOWEVER CAUSED, ON ANY THEORY OF LIABILITY AND WHETHER OR NOT TI HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES, ARISING IN ANY WAY OUT OF TI REFERENCE DESIGNS OR BUYER'S USE OF TI REFERENCE DESIGNS.

TI reserves the right to make corrections, enhancements, improvements and other changes to its semiconductor products and services per JESD46, latest issue, and to discontinue any product or service per JESD48, latest issue. Buyers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All semiconductor products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its components to the specifications applicable at the time of sale, in accordance with the warranty in TI's terms and conditions of sale of semiconductor products. Testing and other quality control techniques for TI components are used to the extent TI deems necessary to support this warranty. Except where mandated by applicable law, testing of all parameters of each component is not necessarily performed.

TI assumes no liability for applications assistance or the design of Buyers' products. Buyers are responsible for their products and applications using TI components. To minimize the risks associated with Buyers' products and applications, Buyers should provide adequate design and operating safeguards.

Reproduction of significant portions of TI information in TI data books, data sheets or reference designs is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

Buyer acknowledges and agrees that it is solely responsible for compliance with all legal, regulatory and safety-related requirements concerning its products, and any use of TI components in its applications, notwithstanding any applications-related information or support that may be provided by TI. Buyer represents and agrees that it has all the necessary expertise to create and implement safeguards that anticipate dangerous failures, monitor failures and their consequences, lessen the likelihood of dangerous failures and take appropriate remedial actions. Buyer will fully indemnify TI and its representatives against any damages arising out of the use of any TI components in Buyer's safety-critical applications.

In some cases, TI components may be promoted specifically to facilitate safety-related applications. With such components, TI's goal is to help enable customers to design and create their own end-product solutions that meet applicable functional safety standards and requirements. Nonetheless, such components are subject to these terms.

No TI components are authorized for use in FDA Class III (or similar life-critical medical equipment) unless authorized officers of the parties have executed an agreement specifically governing such use.

Only those TI components that TI has specifically designated as military grade or "enhanced plastic" are designed and intended for use in military/aerospace applications or environments. Buyer acknowledges and agrees that any military or aerospace use of TI components that have **not** been so designated is solely at Buyer's risk, and Buyer is solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI has specifically designated certain components as meeting ISO/TS16949 requirements, mainly for automotive use. In any case of use of non-designated products, TI will not be responsible for any failure to meet ISO/TS16949.