

JESD204B Subclasses— Part 2: Subclass 1 vs. Subclass 2 System Considerations

Del Jones, Staff Applications Engineer

Introduction

In “*JESD204B Subclasses (Part 1): An Introduction to JESD204B Subclasses and Deterministic Latency*,” a summary of the JESD204B subclasses and deterministic latency was given along with details regarding an application layer solution for multichip synchronization in a subclass 0 system. Part 2 of the series takes a closer look at the differences between subclass 1 and subclass 2. In particular, we will look at the challenges to meeting the deterministic latency-related timing requirements, device clock speed limitations in subclass 2, and guidelines on which subclass is best for a given system application.

Subclass 1

In a subclass 1 system, the accuracy of the deterministic latency depends on the timing relation between device clock and SYSREF, and the distribution skews of these signals within the system. In addition to the setup time (t_{SU}) and hold time (t_{HOLD}) requirements for SYSREF, the application's tolerance for deterministic latency uncertainty will be critical in defining the application's distribution skew requirements for SYSREF and device clock.

Capturing SYSREF Accurately

Converters employing the JESD204B interface sample data at a high rate. In order to reduce phase noise in the system, it is common for these converters to use a reference clock, which is the same as the JESD204 device clock, that is at or above the sample rate. In many cases, this clock is in the GHz range. At these speeds, meeting the setup and hold time requirements become very challenging. To ease the system design, it may be necessary for the phase offset of SYSREF and/or device clock to be programmable for each device that is part of the JESD204B system.

One advantage that subclass 1 has over subclass 2 is that it uses source synchronous clocking. A subclass 2 system uses system synchronous clocking and will encounter frequency limitations sooner than that of one that uses source synchronous clocking. This will be demonstrated when we take a closer look at specific subclass 1 and subclass 2 timing examples.

Deterministic Latency Uncertainty

Deterministic latency uncertainty (DLU) is the local multiframe clock (LMFC) skew in the JESD204B system and is determined by the difference between the earliest and latest possible capture of SYSREF in the system. Figure 1 illustrates the worst case DLU that occurs when setup and hold time requirements for SYSREF capture are not met at every device in the system.¹ This occurs when the distribution skew of the device clocks in the system are not controlled and creates up to one device clock (DCLK) of uncertainty. This is added to the SYSREF distribution skew (DS_{SYSREF}) to produce the total DLU.

$$DLU = DS_{SYSREF} + t_{DCLK}$$

DS_{SYSREF} is the difference in the arrival time of the earliest arriving SYSREF in the system (across all devices in the system) and the last arriving SYSREF. In the illustration, t_{SU} is one-half t_{DCLK} and t_{HOLD} is one-quarter t_{DCLK} . The earliest arriving SYSREF, SYSREF (A), is captured at the earliest possible time ($DCLK_A$ just meets the setup time requirement) while the last arriving SYSREF, SYSREF (N), is captured at the latest possible time ($DCLK_N$ just misses the setup time requirement). So the corresponding LMFCs are misaligned by $DS_{SYSREF} + t_{DCLK}$.




Figure 1. Worst-case deterministic latency uncertainty.

¹ In order to keep the illustration of the concept of DLU simple, clock jitter and variations due to process, voltage, and temperature (PVT) are not considered here.

In many applications, the requirements for DLU are such that this worst case scenario is acceptable. For these applications, it may not be necessary to tightly control the device clock distribution skew. Ensuring the pulse width of SYSREF is $\geq (2 \times t_{DCLK})$ and controlling the SYSREF distribution skew to meet the system timing requirement should be adequate.

In applications where the additional device clock of uncertainty is not acceptable, then the device clock distribution skew must be tightly controlled to ensure that the timing requirements for SYSREF are met at each device in the system. This case is illustrated in Figure 2 and the uncertainty is given by the equation:

$$DLU = DS_{SYSREF} + t_{Valid Window}$$

$$\text{where } t_{Valid Window} = t_{DCLK} - (t_{SU} + t_{HOLD})$$




Figure 2. DLU while meeting setup and hold time for SYSREF.

Minimizing Deterministic Latency Uncertainty

As implied by the equation for DLU, DLU can be minimized by ensuring that setup and hold time are met for each SYSREF/DCLK pair and by minimizing the interpair distribution skew.

To meet the setup and hold time requirements, each device in the JESD204B system should have its own SYSREF/DCLK pair. Within each of these pairs, trace length matching can be employed to ensure timing. The limit for matching trace lengths is determined by the valid window time for SYSREF switching. Also, SYSREF should be output with the capture edge of DCLK and the SYSREF length must be greater than the DCLK length as determined by the hold time requirement (if t_{HOLD} is 0, then the lengths can be equal).

Since trace length matching is employed, minimizing the interpair distribution skew is effectively the same as minimizing the SYSREF distribution skew. The limit for this distribution skew is the DLU limit minus the valid window time and can also be managed through trace length matching. The DLU limit is set by the requirements of the application.

These methods for minimizing DLU are illustrated in Figure 3. Since each device in the JESD204B system has its own SYSREF/DCLK pair, meeting the timing requirements for capturing SYSREF is just like any system that uses source synchronous clocking. Timing margins at each device are considered independently of the other devices in the system.




Figure 3. SYSREF/DCLK routing for a 3-device JESD204B system.

SYSREF Timing Example Using the AD9250

The AD9250 is a 14-bit, 250 MSPS dual ADC with JESD204B serial data output specified at 5 Gbps. To maximize PLL performance, the AD9250 can accept device clock speeds as high as 1.5 GHz. This provides an excellent example to use to demonstrate how to meet SYSREF timing using trace length matching under the most stringent system DLU requirement.² Here are the conditions for the example:

- ▶ DCLK = 1.5 GHz (667 ps period)
- ▶ $t_{SU} = 500$ ps and $t_{HOLD} = 0$ ps
- ▶ For the example, the system's $DLU_{MAX} = 1$ DCLK (667 ps)

Intrapair Trace Length Matching to Meet SYSREF Timing

Based on specifications for the example, the valid window for meeting setup and hold time is 167 ps ($667 \text{ ps } t_{DCLK} - 500 \text{ ps } t_{SU}$). Travel time is the time from the time when a signal leaves the source to when it arrives at the sink. The travel time of the SYSREF minus the travel time of DCLK needs to be less than 167 ps to meet the setup time and more than 0 ps to meet the hold time. To roughly convert this travel time difference into inches, we estimate the travel time through 1 inch of FR-4 material to be 167 ps per inch. So, for each SYSREF/DCLK pair in the system, the following routing requirement would need to be met:

$$\begin{aligned} DCLK \text{ trace length} &< SYSREF \text{ trace length} \\ &< DCLK \text{ trace length} + 1 \text{ inch} \end{aligned}$$

Meeting this requirement will ensure that the SYSREF transition will occur within the valid window as illustrated in Figure 4.




Figure 4. Meeting the SYSREF/DCLK timing requirement.

Interpair Trace Length Matching to Meet DLU Limit

Since the DLU limit has been set to 667 ps and we know the relationship between the DLU limit and the interpair (or SYSREF) distribution skew (DS_{SYSREF}), it is a straightforward derivation to find the trace length match limit:

$$DS_{SYSREF} = DLU - t_{Valid Window} = 667 \text{ ps} - 167 \text{ ps} = 500 \text{ ps}$$

² Having the DLU requirement equal to the device clock is the worst case for meeting timing on SYSREF.

So, the interpair distribution skew across all SYSREF/DCLK pairs must be within³:

$$500 \text{ ps}/167 \text{ ps per inch} = 3 \text{ inches}$$

Figure 5 illustrates the timing for this example. The best-case distribution skew (DS_{SYSREF}) refers to the case that would allow for a less stringent trace length matching requirement.




Figure 5. Meeting interpair distribution skew requirement.

Advanced Solutions to Meet SYSREF Timing and DLU Limit

Of course, the issue can be solved by using a slower device clock, which would make the length matching easier. This would come at the cost of sacrificing some system phase noise performance. A similar solution to this is to relax the DLU requirement, which would retain the advantage of the improved system phase noise performance. How the DLU requirement is established depends on the application. This will be discussed in the context of deterministic latency accuracy. If the increased phase noise performance is required and the DLU requirement cannot be relaxed, it may prove too difficult to meet the routing requirement for SYSREF/DCLK intradevice skew and interdevice skew (1 inch and 3 inches, respectively). In this case, adjustable phase delay for the device clock and/or SYSREF is required. The resolution of the adjustment must be less than the valid window based on the setup and hold time. From the example, the valid window is 167 ps.

Some FPGAs may have difficulty meeting small adjustment resolution requirements. However, the AD9528 meets this requirement since it is capable of adjusting SYSREF phase delay in 60 ps steps with less than 50 ps variability across all outputs. Figure 6 illustrates how SYSREF can be delayed to meet the timing requirements. In the illustration, SYSREF is delayed in 60 ps increments. Selecting a phase setting that places the SYSREF edge near the middle of the valid window is recommended. In the illustration, green edges indicate good phase settings and red edges indicate bad settings. The phase setting of 3 is in the middle of the valid window and should be used in this case.




Figure 6. Programmable phase delay for SYSREF to meet timing.

In addition to the 60 ps phase steps available on the SYSREF outputs, the AD9528's device clock outputs can be phase delayed by one-half of a device clock cycle. This feature is also helpful in meeting the SYSREF timing requirements.

SYSREF Setup and Hold Time Monitoring

ADI's AD9680 implements SYSREF setup- and hold-time monitor circuits to help in the adjustment of the relative timing between SYSREF and device clock. By monitoring these two registers, the user is able to determine if there is a danger in violating the timing requirements for capturing SYSREF. If either of these registers give an indication that the timing margins are insufficient, the user knows he needs to adjust the relative position of SYSREF vs. the device clock. This adjustment would be made by either adjusting the phase of SYSREF relative to device clock (using the AD9528, for example) or by adjusting the trace length of the SYSREF and/or device clock signal.

Deterministic Latency Accuracy

To better understand how a system's deterministic latency uncertainty is set, an understanding of the application is required. Most systems requiring deterministic latency need to know exactly which sample in time marks the beginning of the data of interest. A common use for deterministic latency is for synchronizing multiple converters in a system. This is referred to as multichip synchronization. In these systems, sample alignment is needed across all converters. Therefore, deterministic latency must be sample accurate. In these systems, the DLU needs to be $\pm$ one-half the sample clock. An advantage of having a device clock that is a multiple of the sample clock is that it simplifies the task of capturing SYSREF in such a way as to be sample accurate. In the AD9250 example, the device clock is 6 $\times$ the sample clock. In order to be sample accurate, the DLU requirements of $\pm$ one-half the sample clock translates to be ± 3 device clocks. This is illustrated Figure 7. Our example for the AD9250 showed that even the most stringent DLU requirement can be met easily with the ability to adjust the phase of the SYSREF at each device. When the device clock is a multiple of the sample clock, capturing SYSREF for sample accuracy is greatly simplified. As the sample rates for converters increase to 1 Gbps and beyond, the ability to phase delay the SYSREF and device clocks will become essential.




Figure 7. Sample accurate requirement for SYSREF capture.

Potential Issues with SYSREF Capture

Other than meeting the SYSREF setup- and hold-time requirements and the DLU requirements, there are other potential issues that could occur related to SYSREF capture. For example, upon initial power-up of the system, it is possible that the SYSREF becomes active prior to the system clocks settling. This can happen when using a continuous SYSREF signal. This is resolved by including programmability in the JESD204B interface that allows the device to wait a certain number of edges before synchronizing clocks. Another

³ 500 ps is the worst-case skew for SYSREF and should be used to determine the trace length match limit.




Figure 8. SYSREF monitor window.

programmability option could allow the user to arm the SYSREF capture when a valid edge is expected. This provides control over when to synchronize with a continuous SYSREF. Many ADI converter devices employing the JESD204B interface, including the [AD9625](#) and AD9680, have these features.

Another example is that small variations in SYSREF may be able to cause an unnecessary resynchronization. This is resolved by including programmability in the JESD204B interface that allows the user to specify the valid window around the LMFC for the SYSREF edge. If the SYSREF occurs within this valid window, then the system is still considered to be in-sync. This is a very useful feature since many applications monitor a continuous SYSREF signal to determine the state of the link. The LMFC boundary is compared to SYSREF to determine the state of synchronization in this case. ADI's AD9680 implements this feature as illustrated in Figure 8.

Other useful features for helping with SYSREF capture are the ability to change which edge of the device clock is used for SYSREF capture and which edge of SYSREF is used to align the LMFC. Many ADI converter devices employing the JESD204B interface include these features.

Subclass 2

In a subclass 2 system, the accuracy of the deterministic latency depends on the timing relation between device clock and SYNC~ signals and a variety of items that will consume the timing budget that will be described forthcoming. As with subclass 1, the application's tolerance for deterministic latency uncertainty will be critical in defining the application's trace length matching requirements for SYNC~ and device clock.

Capturing and Launching SYNC~ Accurately

The challenge for meeting the timing requirements for capturing SYNC~ accurately is essentially the same challenge presented in the subclass 1 discussion on capturing SYSREF. However, since the clocking scheme in subclass 2 is system synchronous, you can no longer perform the timing analysis at each capture device independently from the others and it becomes increasingly difficult in a multiconverter application. Not only this, but you must also account for the uncertainty associated with the launching of the SYNC~ signal. Each device in a system that uses system synchronous clocking will consume a portion of the timing budget. Among the items consuming the timing budget are clock distribution skew (DS_{DCLK}), SYNC~ distribution skew ($DS_{SYNC~}$) for multiconverter systems, the propagation delay of the SYNC~ signal, setup- and hold-time requirements for each JESD204B transmitter, and the clock-to-SYNC~ output delay at each JESD204B receiver's SYNC~ output.

Upper Limit for Device Clock in Subclass 2

The JESD204B standard acknowledges that a subclass 2 implementation will impose a limit on the device clock rate due to the system synchronous clocking scheme employed. Annex B in the standard suggests that this limit is 500 MHz:

"Since SYSREF is a source synchronous signal which can be generated in an accurate phase aligned manner with the device clock, it is expected that a system designer aiming to operate at device clock rates higher than 500 MHz would prefer to use a Subclass 1 approach."

Let's explore a detailed timing example in order to illustrate why such a limit exists.

A Subclass 2 Multiple DAC Timing Example

Let's consider a transmitter application employing two subclass 2 DAC devices connected to a single logic device, as illustrated in Figure 9.




Figure 9. Subclass 2 multiple DAC application.

For the example, a 500 MHz device clock is used. The SYNC~ and DCLK signals have the PCB skews⁴ listed below.

- ▶ Clock to FPGA = 300 ps
- ▶ Clock to DAC1 = 600 ps
- ▶ Clock to DAC2 = 720 ps
- ▶ SYNC~1 to FPGA = 660 ps
- ▶ SYNC~2 to FPGA = 750 ps

⁴ 300 ps $\approx$ 1.8 inches of PCB trace.

Which Subclass Is Best for Your Application?

Choosing which subclass to use for your JESD204B system depends on your need for deterministic latency, how accurate it needs to be if you need it, and the device clock requirements for your system.

Subclass 0 is easiest to implement and can be used when there is no need for deterministic latency. Even if your multiconverter system requires synchronization of the samples from all (or some) of the converters, this can be realized using the time stamp method supported by the AD9625 and AD9680.

Given the ability of subclass 1 to support ultrahigh device clock rates and that are used on high sample rate converters, it is the lowest risk solution for systems that require these high rates. Subclass 1 devices can be used at lower rates as well. If using a device clock rate below 500 MHz, meeting the timing requirements are fairly straightforward without having adjustability in the phase of the clock.

Subclass 2 devices can also be used below 500 MHz. The small advantage in using subclass 2 at lower rates is that it reduces the IO count on the logic device and eliminates the need to route SYSREF to each of the JESD204B devices.

References

[JEDEC Standard JESD204B](#). JEDEC Solid State Technology Association, July 2011.

About the Author

Del Jones is an applications engineer for the High Speed Converters Team in Greensboro, North Carolina. He has worked for ADI since 2000, supporting ADCs, DACs, and serial interfaces. Prior to ADI, he worked as a board and FPGA design engineer in the telecommunications industry. Del earned his bachelor's degree in electrical engineering from the University of Texas at Dallas. He can be reached at del.jones@analog.com.

Online Support Community



Engage with the Analog Devices technology experts in our online support community. Ask your tough design questions, browse FAQs, or join a conversation.

[Visit ez.analog.com](http://ez.analog.com)



AHEAD OF WHAT'S POSSIBLE™

Analog Devices, Inc.
Visit analog.com

For regional headquarters, sales, and distributors or to contact customer service and technical support, visit analog.com/contact.

Ask our ADI technology experts tough questions, browse FAQs, or join a conversation at the EngineerZone Online Support Community. Visit ez.analog.com.

©2019 Analog Devices, Inc. All rights reserved. Trademarks and registered trademarks are the property of their respective owners.

Ahead of What's Possible is a trademark of Analog Devices.

TA12595-10/19