

What to Know About the Differences Between JESD204B and JESD204C



ABSTRACT

This article originally appeared in [All About Circuits](#) and is republished with permission.

As data converter speeds continue to increase, the digital interface is often the limiting factor for throughput. The Joint Electron Device Engineering Council's (JEDEC) current serial standard, JESD204B, is limited by lane speed as well as inefficient 8B/10B coding. JESD204 revision C addresses this, but what are the implications for board layout and protocol implementations? This white paper explains the differences between the JESD204 B and C standards and the impact those changes have on engineers working on high-speed data converter board designs.

Table of Contents

1 Introduction	2
2 Background	2
3 What's New in JESD204C	4
4 Summary	5
5 References	5

Trademarks

All trademarks are the property of their respective owners.

1 Introduction

The JESD204 interface standard was born out of the need to develop a common method for serializing data-converter digital data and reduce the number of interconnects between mixed-signal devices and a processing element such as a field-programmable gate array (FPGA). The standard has gone through several iterations – the B version being one of the most recognized and employed by device developers. The current version is starting to lose steam, however, given the requirements for even faster data converters. New device releases such as the Texas Instrument's [ADC12DJ5200RF](#) are breaking the 10-GSPS barrier and are adopting the C standard to keep up with the required throughput.

JESD204 may seem complicated for those that have read the JEDEC documents, but the standard was designed to manage the fantastic amounts of data being provided by modern mixed-signal devices such as analog-to-digital converters (ADC), digital-to-analog converters (DAC) and analog front ends. For those new to the standard, more information is available on the [Texas Instruments training portal](#).

2 Background

JESD204B was the successor to the A standard, which lacked the means to synchronize multiple collocated devices. In other words, if you had two identical data converters on a circuit card assembly and wanted the inputs (or outputs) to be phase-coherent, then you could not use the A standard – it simply did not support any means to accomplish this. The B standard introduced the concept of sub-classification, which provided a backward-compatibility mode (subclass 0) along with two new classes (1 and 2) that provided deterministic latency between the sample time and the data arrival in the consumer (such as an FPGA).

JESD204A was much slower than the B revision. The original standard had a maximum lane rate of 3.125 Gbps, while the B standard was capable of up to 12.5 Gbps. As these lane rates increased, it introduced issues that are common with high-speed serial links: signal integrity, clock recovery and base line wander. Base line wander is a phenomenon caused by an unequal number of 1s and 0s on a transmission line, causing the dielectric to charge and the DC center point of the transmission line to move. This movement can cause two significant issues: it can introduce jitter as well as coding errors.

Base line wander most commonly occurs in systems that are AC-coupled. The AC-coupling is used to mitigate DC offsets between the transmitter and the receiver.

To mitigate base line wander and provide a means for AC coupling, the original standard used 8B/10B symbol encoding. For those not familiar with this type of coding, it is a simple lookup table that returns a 10-bit “symbol” for every octet (8-bit word). 8B/10B symbol encoding also introduces other clever techniques to overcome the running disparity of 1s and 0s and maintain a fixed DC point on the line, while always providing transitions to bridge-coupling capacitors.

Some additional symbols are not used for data but rather for in-band signaling to the receiver. For example, the K28.5 (comma) symbol is used for lane alignment, so at startup, the receiver can establish where in the stream the bits begin (for example, to find the start of a symbol frame). Once the receiver establishes the connection, it can signal back (through the SYNC signal) that the receiver is now matching the encoder, and data traffic can begin.

8B/10B coding is an excellent method for the task of serialization and deserialization, but it adds a 20% penalty to the capacity of each lane (channel). So for every 8 bits of data provided by the sender, 10 bits are actually transmitted (see [Figure 2-1](#) as an example). As data capacity increases, 8B/10B coding becomes greatly inefficient.

Another side effect of serialization and extremely fast lanes is noise injection, which is a byproduct of 8B/10B encoding along with the high-speed transmission lines. Repetitive symbols that couple to the input (or output) of data converters may appear as spectral spurs, which can affect system performance. JESD204 provides a scrambling polynomial that can be turned on once the link is established to spread the energy and lower the spurious content.

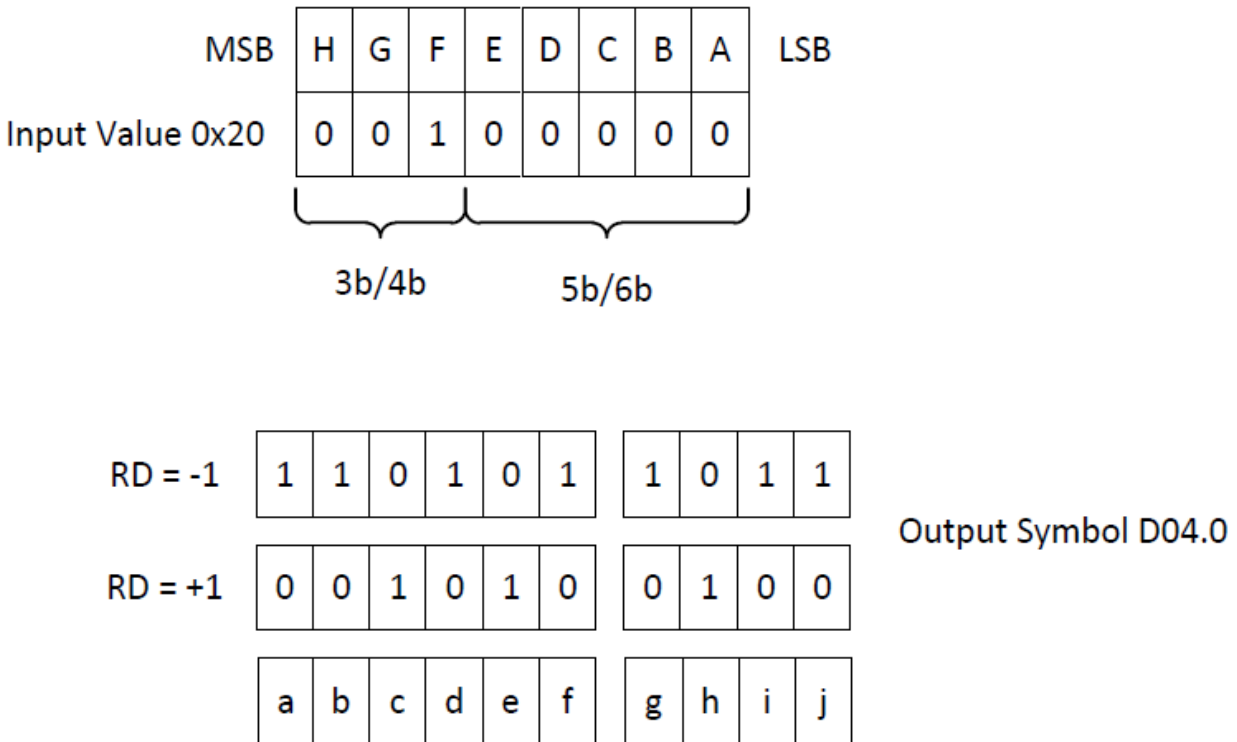


Figure 2-1. Shown is an Example of 8B/10B Coding in Which 8-Bit Value 0x20 is Encoded to the D.04.0 10-Bit Symbol

Note

Running Disparity [RD] is the difference of 1's and 0's in the output of the encoder. Depending on the state, RD can yield multiple values for the same symbol.

As with any serialization methodology, data from multiple sources can be multiplexed into several lanes and reconstructed on the other end. In the International Organization for Standardization Open System Interconnection (ISO-OSI) model, the interface is broken into several layers: the transport, data link and physical layers. The data-link layer performs much of the aforementioned encoding and decoding, but data multiplexing takes place in the transport layer before encoding, which is one of the most complex blocks in the JESD204 standard. This block fundamentally allows the transporting of multiple converter channels across multiple serialized lanes. Multiple channels on individual data converters will be phase coherent, but the delay of the transport layer will vary for various reasons. JESD204B added a mechanism to provide deterministic latency so that multiple devices can be made phase coherent which is very important for electronically steered arrays (ESA).

3 What's New in JESD204C

There are many enhancements in the C revision of the standard; many of the enhancements improve coding efficiency and overall throughput. JESD204C is backward-compatible with the A and B standards, but with some limitations in subclass-0 operation. Designers familiar with the JESD204B revision will see compatibility based on the coding scheme and recommendations for higher throughput, using various enhancements to the standard.

The most obvious change is the addition of denser 64B/66B and 64B/80B coding schemes; these methods are far more efficient in symbol coding. 64B/66B has only a 3.125% coding overhead versus the 20% overhead on 8B/10B. The downside is that it initially takes longer to encode due to the longer symbol length, which will result in a certain coding latency since most data converters provide 8 to 16 bits of data per sample.

For 8B/10B, not much has changed from the B revision. Subclass 0, 1 and 2 are all supported. As a refresher, subclass 0 is the A revision's backward-compatibility mode, used for the lowest possible link delay without deterministic latency. Subclasses 1 and 2 establish deterministic channel latency and multi-device phase alignment. The transport still uses the SYNC interface and initial lane alignment sequence (ILAS) to establish both data alignment and channel latency.

The 64B/66B and 64B/80B coding standards used in JESD204C are a bit more complicated and no longer use the ILAS and SYNC interface to establish proper frame alignment. Instead, these modes use a SYNC header embedded as the first 2 bits of every frame, which are concatenated into a 32-bit SYNC message. Every run of 66 bits starts with the sync header symbol (2 bits – 01 or 10 are valid sync symbols, 00 and 11 are illegal values) encoded into the SYNC header stream. This stream always contains a pilot signal used for sync alignment to the 66-bit frame. Multiple frames form blocks, and groups of blocks form extended multi-blocks. Once the receiver identifies where in the stream the SYNC resides, it then moves on to identify where the frames are aligned by using the last frame indicator in the stream to sync to the frame boundaries.

Beyond the coding, there are several new classes of data transmitters and receivers that help the transmission speed and lower the interface power. These classes are shown in [Table 3-1](#).

Table 3-1. There are Several New Classes of Data Transmitters and Receivers Based on the B and C JESD204 Revisions and Related Data Rates

Device Class Category	Device Class	Comments	Supported Data Rates
B	B-3		0.3125-3.125 Gbps
	B-6		0.3125-6.375 Gbps
	B-12		6.375-12.5 Gbps
C	C-S	Short-reach channel class	6.375-32 Gbps
	C-M	Medium-reach channel class	6.375-32 Gbps
	C-R	Reflective channel class	6.375-32 Gbps

The three new revision C classes allow device manufacturers to provide driver/receiver pairs that have varying amounts of signal integrity processing to reduce power in shorter channels. For example, the C-S class requires that the receiver's continuous time linear equalizer (CTLE) have only 6 dB of gain and no decision feedback equalization (DFE). Class C-R, however, requires that the CTLE have a minimum of 12 dB of gain and a 14-tap DFE. So manufacturers can provide modes that support all classes and give designers the flexibility to choose the best mode for the channel.

In addition to using SYSREF in subclass 1 for phase alignment, the latest revision introduces a new signal called MULTIREF. Unlike the SYSREF signal, which originates in clocking devices, MULTIREF is an output from the data converter and is fed to the SYSREF input of the next data converter. Because the signal originates from the data converter, it may be difficult to meet the setup and hold times of some devices. Where applicable, however, the MULTIREF signal can greatly simplify the clocking of multiple converters where frame alignment is required but deterministic latency is not.

Table 3-2 lists the most significant differences between the two standards. Higher data rates are a significant difference; to better support them, there are two new coding schemes. The methods for synchronization vary by these coding schemes, and there is some additional complexity to this process.

Table 3-2. Shown is a Comparison of the Major Differences Between JESD204B and JESD204C

Parameters	JESD204B	JESD204C
Raw serial bit rate	Up to 12.5 Gbps	Up to 32 Gbps
Support for deterministic latency	Yes	Yes
Transceiver classes	No	Yes
Transport layer coding	8B/10B	8B/10B, 64B/66B, 64B/80B
Phase synchronization	Local multiframe clock	LMFC, Local extended multiblock clock
Phase synchronization clock options	SYSREF (subclass 1) SYNC (subclass 2)	SYSREF (subclass 1: 8B/10B, 64B/66B, 64B/80B) SYNC (subclass 2: 8B/10B)
Lane data alignment	SYNC interface	SYNC headers
Maximum K frames in multiframe	32	256
Programmable ILAS length	Yes	No (fixed at 4)
Cross-standard compatibility	A and B	A (limited), B and C
Subclass sync support	Classes 1 and 2 (8B/10B)	Classes 1 and 2 (8B/10B) Class 1 only (64B/66B, 64B/80B)

4 Summary

Engineers designing with JESD204C interfaces have several new modes of operation to both synchronize multiple converters as well as increase the speed of the digital interface. For those familiar with JESD204B, the additional steps to create a C-compliant interface should not be overwhelming; if the signals are used, a bring-up mode in 8B/10B might simplify development. Once the interface is operational in 8B/10B, then you can apply the higher coding and lane rates to get the maximum performance from the data converter.

5 References

- Texas Instruments: [System Design Considerations When Upgrading From JESD204B to JESD204C](#)
- [JESD204 Rapid Design IP](#) (royalty-free IP for FPGAs connected to TI high-speed data converters)
- [High-Speed Signal Chain University](#) (training series)

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 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 © 2022, Texas Instruments Incorporated