

A guide to multi-channel synchronization for MIMO systems

A Guide to Multi-Channel Synchronization for MIMO Systems



1: Introduction

On the modern battlefield, dominating the electromagnetic spectrum is increasingly important. In order to do so, radar, SIGINT and electronic warfare systems must be able to cover wider and wider bands of the spectrum. For phased array systems, this means that more channels and higher sample rates are being deployed in order to maintain dominance of the spectrum. As our adversaries rapidly ramp up their capabilities, the urgency to develop more sophisticated radar, SIGINT, and EW systems is only accelerating.

As channel count and sampling rates increase in these systems, there are some significant data processing and data movement challenges. With MIMO systems with very high channel counts, a notable challenge faced by systems engineers is channel-to-channel synchronization.

For these applications, precise channel synchronization within several degrees is essential. Channel-to-channel synchronization is often one of the first steps when developing a radar, SIGINT, or EW system - but it can take months to achieve.

In this paper, we describe the steps necessary to synchronize multiple Gigasample JESD204B ADCs using the Abaco SRS6000 as a case study.

The SRS6000 can synchronize 32 ADC channels out of the box, with the ability to daisy chain up to eight systems for a total of 256 synchronized channels. The case study targets specific Abaco hardware, but the concepts presented can be applied to any system leveraging JESD204B-based converters.



A Guide to Multi-Channel Synchronization for MIMO Systems




Figure 1.

2: JESD204B Subclass 1 Overview

Before discussing synchronizing JESD204B devices, a quick review of JESD204B Subclass 1 is in order. In a JESD204B Subclass 1 system, a common clock source distributes a device clock and SYSREF signal to all transmitters and receivers in the system, as depicted in **Figure 1**. The device clock is used to generate sample clocks as well as frame and local multi-frame clocks (LMFCs) within the transmitters and receivers. They SYSREF signal is used to align frame and local multi-frame clocks across all transmitters and receivers and it can be a single pulse or a periodic signal, but it must be synchronous to the device clock. If SYSREF is periodic, it should be a subharmonic of the lowest frequency clock in the system - typically, the local multi-frame clock. LMFC alignment is critical to achieving deterministic latency because it serves as a common low-frequency time reference for all devices.

A JESD link between a receiver and one or more transmitters starts with code group synchronization (CGS) and is initiated by the receiver issuing a synchronization request to the transmitters by asserting the SYNC~ signal. The transmitters latch SYNC~ and, on the next LMFC boundary, begin transmitting the comma character /K/ = /K28.5/ repeatedly.

The receiver de-asserts SYNC~ on an LMFC edge after correctly receiving four consecutive /K/ characters. The transmitters latch the SYNC~ de-assertion and begin transmitting the initial lane alignment (ILA) sequence, consisting of exactly four multi-frames.

Each multi-frame in the ILA starts with the comma character /R/ = /K28.0/ and ends with the /A/ = /K28.3/ character. The receiver uses these characters to align data to a multi-frame boundary. Beginning with the /R/ character, the received data is then fed into a FIFO - referred to as an elastic buffer. Data is buffered for each serial lane in order to absorb any misalignment between lanes. After a release buffer delay, which is some number of frame clocks after an LMFC edge, data frames are read from the buffer, resulting in a deterministic latency between the JESD transmitter and the receiver.



A Guide to Multi-Channel Synchronization for MIMO Systems




Figure 2- JESD204B timing diagram. Taken from [1].

After four /A/ characters have been received on a serial lane, marking four complete multi-frames, valid transmitter data is released from the core. A timing diagram of this process is depicted in **Figure 2**.

Notice that in this timing diagram, the serial data arrives misaligned by four frame clocks. However: because the serial data is buffered for each lane and data is only released from the buffers once all data has arrived on all serial lanes, the data leaves the elastic buffers aligned and with a deterministic latency.

To summarize, JESD204B Subclass 1 achieves deterministic latency between transmitters and receivers in a system by first aligning the clocks in all the devices to SYSREF signals that are coherently distributed to all devices in the system. Once all the clocks in the system are aligned, the JESD204B link that is established after code group synchronization and initial lane alignment results in deterministic latency, provided a suitable release buffer delay is used. For more information about JESD204B, consult the JESD204B standard [1] or any of the numerous JESD204B overviews available online [2]-[5].



A Guide to Multi-Channel Synchronization for MIMO Systems




Figure 3 - JESD204B system of multiple FPGAs and ADCs with a common clock and trigger generator.

3: Synchronizing JESD204B ADCs

In the last section, we reviewed JESD204B Subclass 1 and showed that, provided certain conditions were met, the latency in a JESD link is deterministic. Deterministic latency between transmitters and receivers in a system is necessary in order to synchronize the system. In this section, we will review the steps required to synchronize multiple JESD204B Subclass 1 ADCs (transmitters) to several FPGAs.

Figure 3 illustrates a system consisting of N of the subsystems presented in **Figure 1** with a common clock source to distribute a system clock, synchronization pulse, and acquisition trigger to all subsystems.



A Guide to Multi-Channel Synchronization for MIMO Systems




Figure 4 - Valid assertion window for SYSREF within a device clock period. Taken from [5].

Extending the JESD204B system that was discussed in the last section to now include multiple FPGAs and ADCs, we can achieve deterministic latency for all ADC and FPGA links by ensuring that all devices receive a common device clock and phase coherent SYSREF signal. This can be achieved by using a master clock generator to distribute a common system clock and synchronization pulse to each subsystem.

The clock generator in each subsystem buffers the external system clock and uses it to generate device clocks for the FPGA and ADCs. The synchronization pulse is used by the clock generators in each subsystem to reset all clock dividers used to create device clocks and SYSREF. This way, all device clocks and SYSREF signals in the system are aligned to the external synchronization pulses and each other.

However: in a practical system, path length differences between the system clock generator and each subsystem, as well as variations between subsystems, will result in skew between device clocks and the synchronization pulses across subsystems. Both SYSREF and the synchronization pulse are synchronous to the device clock, so setup and hold time violations can lead to non-deterministic clock relationships which would result in a loss of deterministic latency and synchronization. The likelihood of setup and hold time violations increases as the frequency of the device clock increases.

For example, assuming a 1.0 GHz device clock, all subsystems must register the synchronization pulse within the same 1.0 ns clock period and satisfy setup and hold time. The same requirement applies to each subsystem, only with SYSREF instead of the synchronization pulse. It is therefore typically necessary to be able to adjust the skew of both the synchronization pulse and the SYSREF signals relative to a device clock.

We will assume the synchronization pulses and SYSREF signals have been skewed such that they are registered on the same clock period across all devices and that JESD204B links have been established after performing code group synchronization and initial lane alignment. At this point, data in the FPGA is time-aligned and is available to be stored in memory. All that is required is to issue a simultaneous trigger to all FPGAs that will result in data being stored in memory.

As with the synchronization pulse and SYSREF, the acquisition trigger should be synchronous to the FPGA device clock and must satisfy setup and hold time. If an FPGA does not register the acquisition trigger on the same FPGA device clock edge as the other FPGAs, or if there is a setup and hold time violation, the acquired data will be skewed relative to data acquired on other FPGAs. It is therefore typically necessary to be able to adjust the skew of acquisition trigger relative to a device clock.

To summarize: a system consisting of multiple subsystems of an FPGA connected to multiple ADCs through a JESD interface with a common clock generator can be synchronized with a system clock generator. The system clock generator distributes a master clock to all devices from which local device clocks are generated; synchronization pulses to each subsystem clock generator to align SYSREF across all subsystems; and a synchronous trigger to initiate simultaneous data acquisition on all FPGAs.





A Guide to Multi-Channel Synchronization for MIMO Systems




Figure 6-Block diagram of the SRS6000 showing the routing of signals required for multi-channel synchronization.



A Guide to Multi-Channel Synchronization for MIMO Systems

4.2 Clock Alignment

Now that we have an overview of the system connections, we can follow the synchronization procedure outlined in Section 2. The first step was to align all clocks with a deterministic phase relationship between them. The FMC123 LMK04828s are programmed to use the trigger/sync signal from the FMC407 to reset its clock dividers. One complication is that the synchronization pulse is routed to an AC-coupled input, ClkIn0, of the LMK04828. This requires the synchronization pulse to be a periodic signal. Since all clocks generated by the LMK04828 as well as within the ADCs are multiples of SYSREF, the synchronization pulse should be a sub-harmonic of SYSREF.

The synchronization pulses arriving at each FMC123 LMK04828 are synchronous to the 1 GHz system clock but, due to differing propagation delays from the FMC407 to the FMC123 modules, the synchronization pulse could be registered on a different system clock period on some FMC123 LMK04828s than others. This would cause the clock dividers of the affected FMC123 LMK04828s to reset on a different clock cycle, resulting in a 1 ns skew in the data. Fortunately, the FMC407 can delay the trigger/sync outputs relative to the system clock output such that the LMK04828 on all FMC123s reset on the same device clock period. This process is discussed in more detail in the System Calibration section.

Once all clocks are aligned, the synchronization feature in the FMC123 LMK04828s is disabled so that the trigger/sync signals generated by the FMC407 can now be repurposed as data acquisition triggers.

4.3 JESD204B ILA for Deterministic Latency

With all the clocks aligned in the system, the FMC123 ADCs and the FPGA JESD204B receivers can establish a JESD link. In order to operate the ADCs at the maximum sample rate of 1 GS/s, two JESD lanes running at 10.0 Gb/s per ADC channel are required. This amounts to eight JESD lanes per FMC123. Since there are two FMC123s connected to each FPGA, the FPGA has two copies of the JESD204 receiver core. The FPGA initiates code group synchronization and initial lane alignment with the ADCs in each FMC123 by asserting the SYNC~ signals. Once ILA is complete, data is read from the elastic buffers and output from the JESD204B core. As discussed in Section 2, the release buffer delay is a critical parameter for achieving deterministic latency in a serial link. For this reason, the JESD204B receiver distributed as part of the FMC123 board support package uses a calibrated release buffer delay value that maximizes the timing margin for serial lane misalignment.

Note: the JESD204B link alignment is initiated in software so each FMC123 establishes a JESD link sequentially. However, since all clocks are aligned and there is a uniform, deterministic latency between ADCs in an FMC123 and the JESD204B receiver running in the FPGA, data leaving the elastic buffers from the JESD204B FPGA receivers are synchronized.

4.4 Synchronous Data Acquisition

At this point, all that remains is to ensure that all FPGAs begin capturing data simultaneously. This is achieved with a hardware acquisition trigger that is distributed by the FMC407, repurposing the FMC123 trigger/sync input from a synchronization trigger to an acquisition trigger.

However, similar to the process of aligning SYSREF across all FMC123s and FPGA JESD receivers, the acquisition trigger must be registered on the same data clock period and satisfy setup and hold time across all devices. The timing requirements are considerably easier in this scenario, however, since the FPGA data clocks are running at 250 MHz as opposed to 1 GHz in the LMK04828s and in the ADCs. This results in a 4.0 ns window in which all acquisition triggers must be registered. If this requirement is not satisfied, data captured from the affected FPGA data clock domain will be skewed by a multiple of the 4.0 ns data clock period. As discussed previously, the FMC407 can adjust the skew of the trigger signals such that the acquisition trigger is registered in the same data clock period for all devices. This process is discussed in more detail in the next section.

4.5 System Calibration

So far in this discussion, two occasions in which signals issued from the FMC407 need to arrive within the same clock period and satisfy setup and hold time across all devices have been identified. The first time was during clock alignment where a timing violation would result in a 1.0 ns skew. The second instance was during the data acquisition in the FPGA, where a timing violation would result in a 4.0 ns skew in the data. Further complicating matters is that, if a timing violation is marginal, the skew will not be deterministic but could occur sporadically. A calibration routine is therefore needed to perform a large number of data acquisitions and adjust the FMC407 trigger delays based on the measured skew.



A Guide to Multi-Channel Synchronization for MIMO Systems

The calibration routine works as follows. The first source of skew and non-determinism that is corrected is the acquisition trigger, which results in some channels being skewed by 4 ns relative to the other ADC channels. This is accomplished by first aligning all clocks and establishing JESD204B links. Some channels could be skewed by +/- 1.0 ns in addition to the 4.0 ns skew introduced by the misaligned acquisition trigger, but this will be corrected later. Next, a data acquisition is triggered on all channels. One channel is chosen as the reference against which the skew on all other channels is measured.

Channels that are skewed with a magnitude greater than 1.0 ns have their acquisition trigger delayed or advanced accordingly. Another triggered data acquisition is performed and the process is repeated until all channels are aligned within +/- 1 ns for N_CAL_VALIDATION runs. The last requirement is to ensure that no triggers are arriving near a clock edge.

At this point, any remaining skew should be entirely due to the synchronization pulses being registered on different FMC123 LMK04828 1 GHz clock edges, resulting in skews that are within +/-1 ns. To remove this residual skew, clocks are once again re-aligned, JESD204B links are established, and a data acquisition is triggered on all channels.

Synchronization triggers corresponding to FMC123 channels that are skewed relative to the reference channel are advanced or delayed in time to compensate for the skew. This process is repeated until all channels are aligned for N_CAL_VALIDATION runs. The calibration process is summarized below:

1 Synchronize all devices.

2 Loop

- 2.1 Acquire records on all devices.
- 2.2 Measure channel-to-channel skew.
- 2.3 Adjust the acquisition trigger delays associated with channels that are skewed relative to the reference channel.
- 2.4 If N_CAL_VALIDATION consecutive runs complete with no trigger delay adjustments, exit loop and save acquisition trigger delay values.

3 Loop

- 3.1 Synchronize all devices.
- 3.2 Acquire records on all devices.
- 3.3 Measure channel-to-channel skew.
- 3.4 Adjust the synchronization pulse delays associated with channels that are skewed relative to the reference channel.
- 3.5 If N_CAL_VALIDATION consecutively successful runs, exit loop and save synchronization trigger delay values.

4.6 System Validation

In order to validate the SRS6000, the following test setup is used. An RF signal generator generates a continuous wave (CW) test stimulus which is fed into a splitter network to distribute a phase-coherent stimulus to all Abaco FMC123 ADC channels. The splitter network consists of a cascade of a 1x4 splitter and two 1x16 splitters to create a 1x32 splitter.

The outputs of the 1x16 splitters are cabled to the Abaco FMC123 ADC inputs through equal length cables. **Figure 7** shows the complete test setup. The PC821s with the FMC123s are in the left side of the chassis and the PC821 with the FMC407 is in the right side. The FMC407 distributes clocks and triggers to the FMC123s and the FMC123s are connected to the splitter network.



A Guide to Multi-Channel Synchronization for MIMO Systems




Figure 7- SRS6000 connected to a 1x32 splitter network for system validation.

In order to validate the repeatability of the ADC synchronization, 1,000 test runs were performed consisting of 100 synchronization steps with 10 triggered data acquisitions per synchronization. A synchronization step consists of aligning all clocks with the synchronization pulse and performing JESD CGS and ILA. The triggered data acquisition was performed by generating a trigger from the FMC407, as described in the Synchronous Data Acquisition section.

The channel-to-channel mean skew and standard deviation were measured for the 1000 test runs.

4.7 Results

Figure 8 shows the channel-to-channel skew data corresponding to 1,000 test runs with channel 0 chosen as the reference channel. The data for each run is plotted in the figure as a separate trace and, as can be seen, the traces are virtually superimposed, indicating there is very little jitter in the channel-to-channel skew. The maximum skew between channels was measured to be 101.3 ps; however, the skew introduced by the splitter network was not accounted for, so this result should be interpreted as an upper bound for channel-to-channel skew.

Eliminating the channel-to-channel skew is ideal, but there will generally be some residual offset which can be compensated for either by further adjustment of the clock and trigger delays, or with digital filters in the PC821 FPGAs. In order to calibrate the residual skew, it must be constant. Therefore, the more critical metric is the repeatability of the channel-to-channel skew. To quantify this metric, the standard deviation of the skew for a given channel relative to channel 0 was measured over the 1,000 test runs. The results are plotted in Figure 9, which shows the standard deviation of the skew to be less than 1.0 ps for all channels. The data are summarized in Table 1.

Maximum Skew	39.5 ps
Minimum Skew	-65.4 ps
Maximum Standard Deviation of Skew	0.8 ps

Table 1- Summary of synchronization results with a 10.00 MHz stimulus.



A Guide to Multi-Channel Synchronization for MIMO Systems




Figure 8- Channel skew relative to channel 0 for 1,000 test runs. Each run is plotted as a separate trace, which shows the repeatability of the channel skew.




Figure 9- Standard deviation of the skew for a channel relative to channel 0 over 1,000 test runs.



A Guide to Multi-Channel Synchronization for MIMO Systems




Figure 10- Multiple SRS6000 systems can be synchronized to achieve up to 256 synchronized ADC channels by distributing the system clock and trigger to the FMC407 in each SRS6000 from a common clock generator.

4.8 Channel Count Extension

The SRS6000 provides 32 synchronized 1Gs/s ADC channels but, depending on the end application, 32 channels may not be sufficient. For this reason, the SRS6000 was architected in such a way that multiple SRS6000s can be synchronized together, yielding up to 256 synchronized ADC channels.

This is accomplished by using the FMC407 in the SRS6000 to distribute an externally sourced clock and trigger to the data acquisition modules (PC821s + FMC123s) as opposed to synthesizing its own [8]. The original requirement that all JESD devices have a phase-aligned device clock and SYSREF can be

satisfied by using another FMC407 to distribute the new system clock and trigger to the FMC407s in each SRS6000.

Currently, there is an unused FMC site on the same PC821 FPGA carrier that the SRS6000 FMC407 is connected to, so the master FMC407 can be placed on this FMC site. This configuration is illustrated in **Figure 10**.



A Guide to Multi-Channel Synchronization for MIMO Systems

5: Conclusion

Achieving channel-to-channel synchronization across a number of ADC converters connected to multiple FPGAs can be a challenge in any MIMO system – whether radar, SIGINT, or EW. JESD204B Subclass 1 provides a mechanism for achieving deterministic latency between the data converter and the FPGA but, as described above, there are a number of potential complications that the system designer must be aware of in order to synchronize multiple ADCs and FPGAs.

The Abaco SRS6000 addresses these complications and provides users with 256 synchronized ADC channels sampling at 1 GHz.

The race toward electromagnetic dominance on the battlefield will not slow any time soon, so finding small ways to get the latest cutting edge technology into the hands of the warfighter is paramount – and Abaco is here to help.

6: References

- [1] JEDEC Standard No.204B Serial Interface for Data Converters
- [2] Harris, Jonathan. "What is JESD204B and why should we pay attention to it?".
https://www.eetimes.com/document.asp?doc_id=1279796
- [3] Guibord, Matt. "JESD204B multi-device synchronization: Breaking down the requirements".
<http://www.ti.com/lit/an/slyt628/slyt628.pdf>
- [4] Carnes, Joshua. "JESD204B Link Latency Design Using a High Speed ADC and FPGA".
<http://www.ti.com/lit/ug/tidu171/tidu171.pdf>
- [5] Understanding JESD204B Subclasses and Deterministic Latency.
<http://www.ti.com/lit/ug/snau140/snau140.pdf>
- [6] Abaco PC821 FPGA Carrier User Manual.
<https://www.abaco.com/download/pc821-user-manual>
- [7] Abaco FMC123 FMC Card User Manual.
<https://www.abaco.com/download/fmc123-user-manual>
- [8] Abaco FMC407 FMC Card User Manual.
<https://www.abaco.com/download/fmc407-user-manual>

WE INNOVATE. WE DELIVER. YOU SUCCEED.

Americas: 866-OK-ABACO or +1-866-652-2226 Asia & Oceania: +81-3-5544-3973

Europe, Africa, & Middle East: +44 (0) 1327-359444

Locate an Abaco Systems Sales Representative visit: abaco.com/products/sales

abaco.com | [@AbacoSys](https://twitter.com/AbacoSys)

©2018 Abaco Systems. All Rights Reserved. All other brands, names or trademarks are property of their respective owners. Specifications are subject to change without notice.

