



Achieving deterministic latency is the topic of discussion in much contemporary systems design today. Past efforts focused on increasing data transmission speeds and bandwidths but increasingly, modern applications now set an equally high value on determinism - the requirement that a data packet be delivered at a precise and repeatable moment in time.

This article considers determinism, at device level and expands on the topic of how ultra-fast data conversion and signal processing systems can be designed to guarantee deterministic latency.

Three factors determine how determinism is achieved as follows:

1. Action is taken to mitigate against metastable events occurring in digital design elements.
2. Latency of the digital backend is calculated to ensure alignment of data across multiple data link lanes (e.g., across HSSLs).
3. Time delay margins are sized to ensure that indeterminism does not appear inadvertently owing to PVT variations.

Specifically, we consider the influence of metastability, its mitigation in synchronous systems and show how to maintain determinism at the interface between analog and the digital signal processing domain.

The ability to manage latencies across a data converter array in ultra-fast systems is critical in complex systems spanning digital beam-steered radar to beam-formed, multi-carrier communications. Latency degrades performance. The engineer's desire is to bound latency to a known maximum.

LATENCY - A DEFINITION

A lazy definition is that latency is a measure of the time delay between an action and a response. In the case of sampled data systems, we are most often concerned about maximum latencies. For this hardware focused article, the sources of indeterminism and how those sources are managed is key. Determinism is the simple

Glossary

ADC - Analog to Digital Converter
 CDC - Clock Domain Crossing
 CLK - Sample Clock
 CMU - Clock Management Unit
 ESIstream - Efficient Serial Interface
 ESS - ESIstream Synchronization Sequence
 FPGA - Field Programmable Gate Array
 GT - Gigabit Transceiver
 HSSLs - High Speed Serial Lane
 LD - Logic Device (e.g. an FPGA or ASIC)
 LMFC - Local Multi-frame Clock
 MZ - Metastable Zone
 PVT - Process, Voltage & Temperature
 SSO - Slow Synchronization Output

Two IC data interfaces do just that: license-free ESIstream and the Industry standard JESD204B (sub-classes 1 & 2). Both are widely used, connecting data converters to logic devices (LD) such as FPGAs and ASICs. Both promise determinism, each differ in specific implementation. The conclusion will show that today, designers have a choice between ultimate flexibility or the low-overhead simplicity and reduced absolute latency of ESIstream.

requirement that a system yields the same result for a given set of inputs. Irrespective of environmental changes, or start-up conditions, the outcome is predictable; random factors are eliminated. In essence, a deterministic system provides a bounded response.



CHALLENGES TO ACHIEVING DETERMINISTIC BEHAVIOUR

Identifying the sources of indeterminism is not intuitive, especially at gigahertz sampling frequencies. Figure 1 identifies several sources for the simple case of a single ADC, the EV12AQ600, connected to a Logic Device (LD). Indeterminism arises due to metastability (see sidebar), a factor arising in synchronous logic systems. It is exacerbated by three additional factors:

- Clock domain crossing (CDC) leading to potentially unequal length data paths made worse through any physical signal trace length differences.
- Latency differences acquired at the LD output buffers arising from data alignment across multiple HSSLs
- PVT (process technology, voltage, and temperature) effects

Metastability

Metastability is logic state uncertainty arising in synchronous systems during state transitions due to finite set-up and hold times. Metastability is avoided by creating state sample points backed off from the MZ (figure 5).

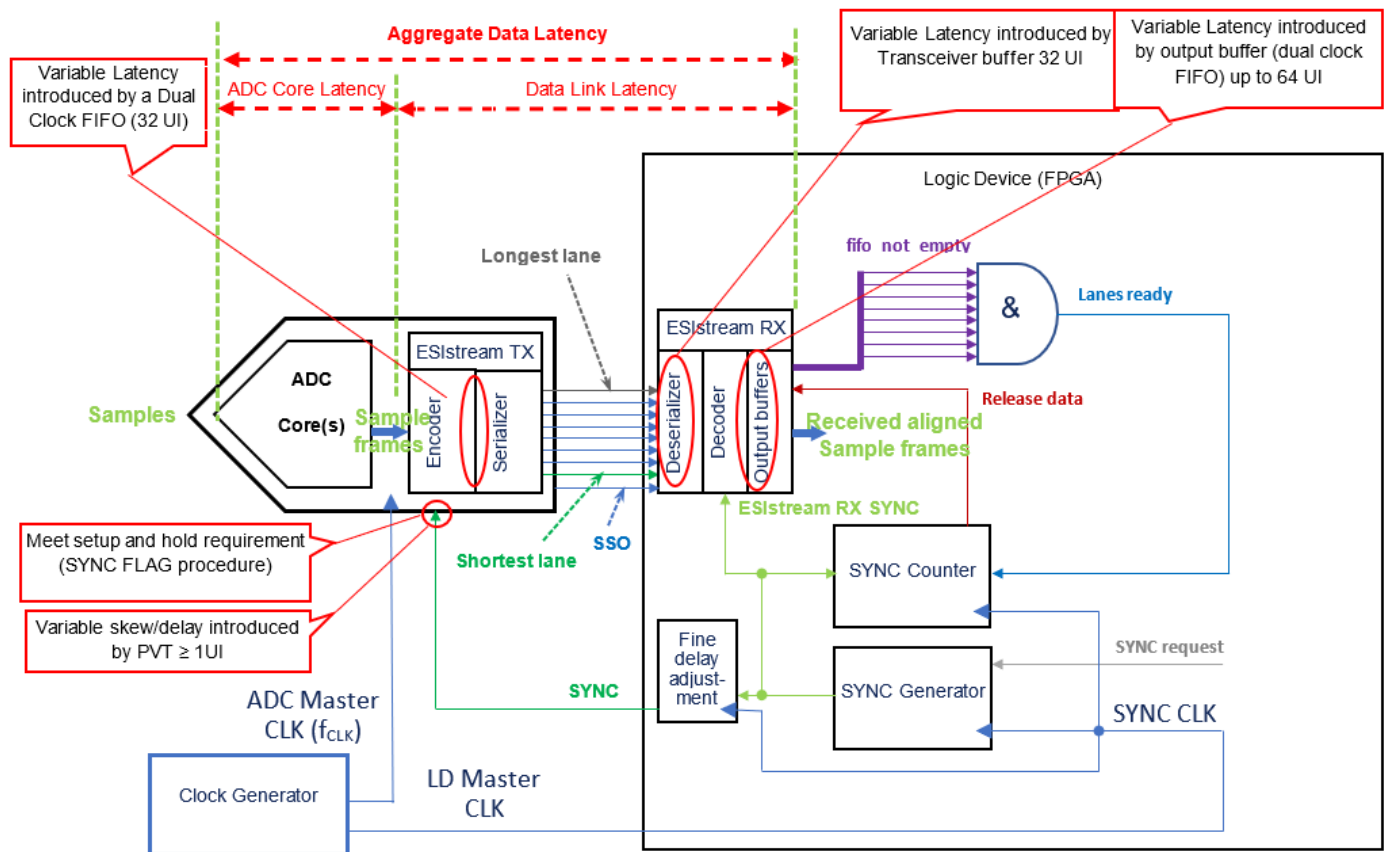


Figure 1 Sources of indeterminism and latency accumulation within the EV12AQ600.



AVOIDING METASTABILITY

It is important to emphasize that a generated SYNC (synchronization) signal for the system must be sampled outside of the metastable zone (MZ). Furthermore, this should always occur on the same ADC master clock (fCLK) edge to ensure deterministic latency throughout a multi-channel sampling system.

CLOCK DOMAIN CROSSING (CDC)

Both the data converter and the attached logic device (an FPGA here) are complex synchronous sub-systems with associated hierarchical clock structures enforcing local determinism. An external low-jitter master clock must be applied to synchronize the two domains.

In the ADC, a variable latency occurs when transporting data from the encoder clock domain to the transmitter/serializer clock domain using a dual clock FIFO. In the FPGA, a variable latency occurs when transporting data from the receiver/deserializer to the decoder using the transceiver buffer and from the decoder to the user application using the output buffer. Output data from the EV12AQ600 is transported over quad pairs of ES1stream serial lanes. Individual lanes have their own slightly different latency due to the CDC. The latency for each lane at EV12AQ600 data output can vary between 126 to 142 clock cycles (32 UI of variable latency). In addition, the physical distance between the ADC and the receiving decoder delay data transfer. Any difference in PCB trace lengths across parallel lanes adds further delay or skew to the link.

De-skewing (see figure 2) and re-aligning data frames at the receiver end to account for the modest arrival time

PHYSICAL SIGNAL SKEW

On conventional PC boards, delay may be incurred as a 6 GHz sampled system typically incurs a 6.5 ps/mm propagation delay in a 50 ohms microstrip line (i.e., across copper traces). Any length variability between

differences incurred, demands flexible data buffering in the LD output buffer. De-skewing enables the correct lane alignment at the receive end. As will be shown shortly, this occurs using a time of travel counter - having trained the system to establish a latency limit. Once this limit is known it can flag a 'release data' event.

EV12AQ600 Highlights

- 12-bit quad core ADC
- Up to 6.4 Gbps
- Up to 6.5 GHz bandwidth
- Integrated cross-point switch

EV12AQ600 Synchronous Clocks

- fCLK & fSSO
- fCLKMAX = 6.4 GHz (fserial = 2 x fCLK)
- fSSO = fCLK/32

data lanes thus injects additional transport delay. The LD de-skew buffers therefore should be sized to account for this factor as well.



P, V, T EFFECTS

Process (as in semiconductor process), voltage and temperature differences over time affect operating points for electronic systems. This explains why components undergo a thorough qualification to map their performance – establishing PVT boundary conditions. Any system designed to offer deterministic latency must be robust enough to avoid changes in P, V or T impacting determinism. This almost certainly requires some control mechanism to allow for initial system calibration as well as a second order method to monitor performance over time. This is a topic to return to shortly.

Taking all the above factors into account, latency of the system is deterministic if the delay between a SYNC pulse and receive output buffer 'valid data' assertion is a fixed constant (figure 2: Release data). Moreover, this is a robust event if it is repeatable after multiple power-up & reset cycles.

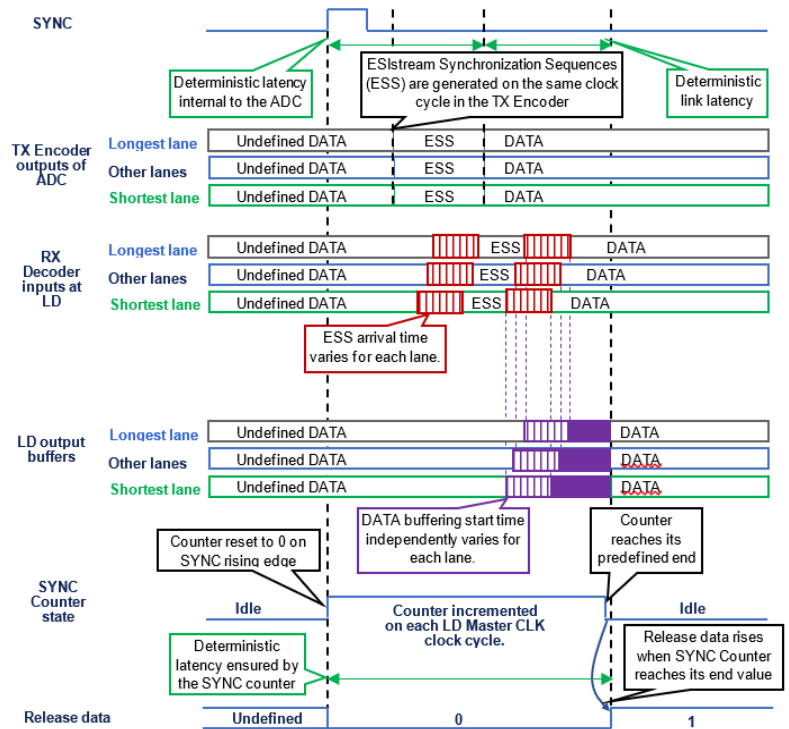


Figure 2 De-skewing ADC output data in the LD.

DEFEAT METASTABILITY WITHIN THE ADC USING THE SYNC FLAG PROCEDURE

To avoid metastability, gated time delays relative to the master clock are introduced as depicted in figure 3. This avoidance practice is essentially a re-timing approach. Figure 3 Sync pulse delay to avoid metastable zone.

Synchronizing the four cores of the EV12AQ600 requires precise clocking to enable precision core interleaving. That is the task of the ADC Clock Management Unit (CMU) which also features a metastability mitigation device exercised via the SYNC_CTRL register (0x000C). Initially, the ADC indicates metastability by asserting the SYNC_FLAG bit (0x000D = 1). Once asserted, the SYNC_CTRL register allows user programming of the ADC sampling edge (figure 3). Verifying that metastability is

avoided is simply confirmed by checking if SYNC_FLAG is reasserted. If all is well, the SYNC_FLAG remains low (SYNC FLAG procedure available in the [EV12AQ600 datasheet](#)).

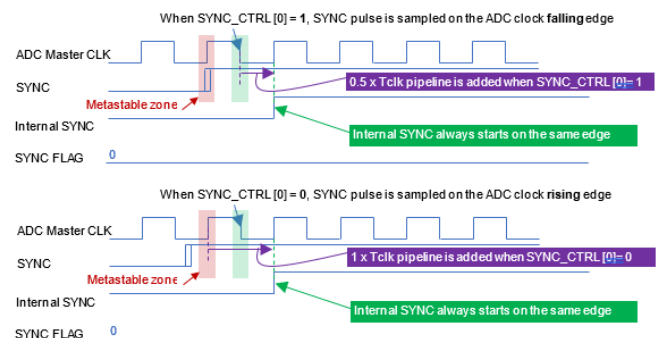


Figure 3 Sync pulse delay to avoid metastable zone.



SYNC CHAINING – A SIMPLE ROUTE TO MULTI-CHANNEL DETERMINISM

The EV12AQ600's CMU provides controls to defeat internal metastability. Advantageously, the EV12AQ600 facilitates sync chaining through its Sync output signal (SYNCO). This output may be daisy-chained to other ADCs within an extended system. In this way, deterministic, phase coherent sampling is maintained throughout. This is a huge benefit in systems where phase information is critical such as in digital beam

forming applications e.g. synthetic aperture Radar (SAR). Though this approach extends deterministic sampling across multi-channel systems, it only impacts the analog front end. It does nothing to ensure that output data transmitted to the LD is deterministic. Therefore, in the digital domain, further mitigation steps are needed.

ENSURING DIGITAL BACK-END DETERMINISM

Figure 2 previously showed that individual ESSs suffered varying arrival times. An obvious low overhead approach to de-skewing these lanes is to establish a delay counter – easily implemented within the LD see figure 4.

The counter accumulates the number of clock cycles passed since the initial SYNC pulse at the ADC and the receipt of the slowest ESS by the LD. At that point, the 'release data' event is asserted indicating received data deserialization is complete. Having trained the system, the SYNC delay quantifies the link latency of the slowest ESStream lane which includes contributions from both the link layer as well as the physical copper interconnect.

The counter delay allows for the subsequent alignment of all the receive buffer data. Obviously in large distributed systems, the data link delay will vary for each converter and needs to be established in an initial training phase. Fortunately in ESStream systems, enforcing deterministic sampling is assisted by SYNC chaining. Offsetting data ready assertion from the SYNC event, with a delay and appropriate margining accounting for the slowest lane, extends deterministic latency across distributed systems.

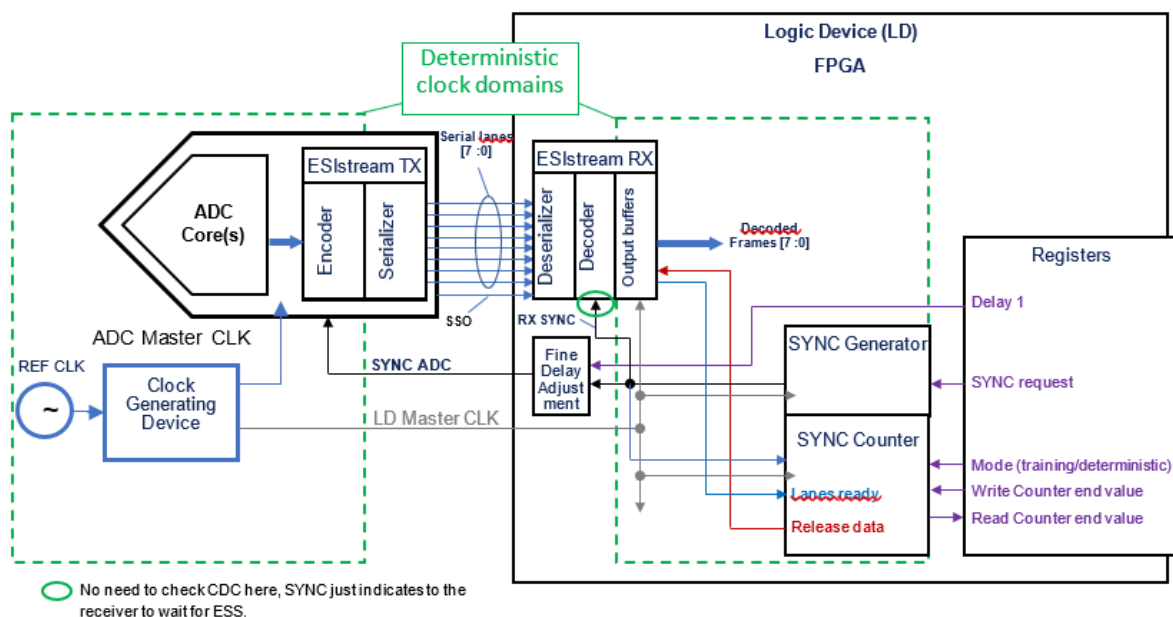


Figure 4 A Sync counter loop delays 'data lanes ready' assertion until the slowest lane is ready.



MANAGING PVT IMPACTS ON DETERMINISM

As sample frequency increases and especially as sampling nears the upper limit for the EV12AQ600 at 6.4 GHz, clock signal skew introduced by temperature variation can cause the system to wander out of deterministic operation and needs to be guarded against.

Two countermeasures are suggested by Teledyne e2v:

- Characterizing the system over temperature to

establish its functional operating limits

- Developing a dynamic, fine adjustment algorithm for setting SYNC pulse edge position

Obviously the later, more complex approach offers more resiliency over lifetime, all be it with the added development cost.

THERMAL CHARACTERIZATION

Here the aim is to establish a safe mid-temperature operating point that assures determinism, then to adjust temperature over the operating range and monitor individual ADC SYNC_FLAGS for metastable zones (MZ). With the resultant MZ map it is possible to determine the best SYNC_EDGE value (0: rising or 1: falling) at specific temperatures giving the best operating margins. Armed with this information held in a local look-up table, the system can respond to temperature changes with appropriate SYNC_EDGE changes.

Careful MZ mapping, should help avoid metastability. One area where this approach may have limitations is in the case of age induced changes. It is tougher to characterize lifetime performance and establish a time dependent MZ map. In this case, an alternative approach might help.

THERMAL CONTROL ALGORITHM

An algorithm to dynamically adjust the SYNC pulse phase offset (versus the master clock) can be conceived. This may be implemented as an extra time delay block within the LD – for instance, the ODELAY module within Xilinx FPGAs.

- Setting the SYNC pulse with maximum phase margin
- Or dynamically adjusting the phase to avoid metastability

As previously outlined, establish a mid-temperature deterministic operating point. Then, using the SYNC_FLAG process, adjust the phase of SYNC relative to the master clock across the full phase range (0 to 360 °) and monitor for individual ADC SYNC_FLAG assertion events. This process establishes the range of synchronization phase margins. Armed with this information, deterministic operation can be maintained either by:

Adopting either one of these tactics still requires careful system consideration. At high clock frequencies, phase margins are always under pressure, highlighted by figure 5. Depending upon trade-offs and layout concerns, it may be necessary to introduce fine SYNC phase adjustment control provisioned from an external time delay IC co-located with individual ADCs.

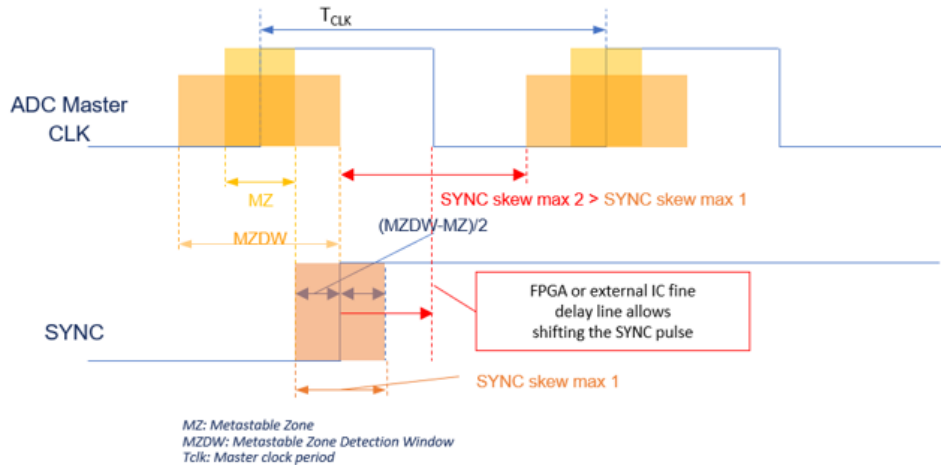


Figure 5 This fine delay approach may be required in the most challenging system environments.

ABSOLUTE LATENCY IS DETERMINED BY DIGITAL DESIGN CHOICES

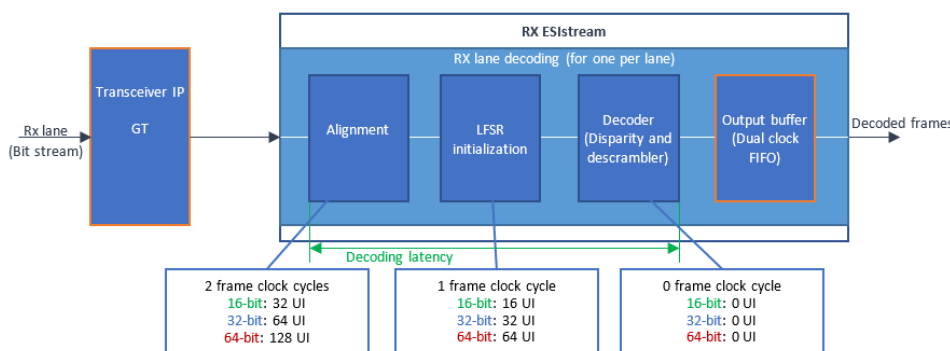
One factor determines overall latency – the choice of data frame length. This impacts upon logic device design. Table 2 quantifies the impact of choosing one-, two- or three-word length frames.

Frame length	No. of HSSLs	Maximum data rate (Gbps)	FPGA frames frequency	FPGA resources	
				KU060 (LUTs)	KU060 Flip-flops
16-bit	8	6.25	390.625	1803	1854
32-bit	8	12.5	390.625	3551	2514
64-bit	8	12.5	195.31	5291	4222

Table 1 Frame length selection determines logic resources & data rates.

Contemporary FPGAs can decode lane data rates up to between 400-500 MHz. However economic considerations will impact their selection. In some applications, perhaps a slower frame rate is desirable. This can be achieved by using longer frames (table 1).

However, this choice affects the complexity of the digital resources needed and implicitly increases aggregate absolute latency (figure 6).



The Unit Interval (UI)

- The time to transmit a single symbol
- UI at 6.4 GHz = 156 ps

Figure 6 Aggregate system latency expressed in terms of Unit Interval (UI).



A QUICK WORD ON ESISTREAM VERSUS JESD204B/C

Whilst JESD204B/C lends itself to ultimate reconfigurability, there is no doubt that the signal processing industry is wary of its hidden complexity. This sense of foreboding is hinted at by one vendor, whose technical document proclaims: 'The JESD204 Survival guide'. The mystery arises from the multiple clock domains and complex transport layer. High-level features of the two alternatives are summarized below (table 2).

ESIstream eliminates the transport layer coding complexity of JESD204 with the extra upside that it is easily described in a concise 12-pages specification. Moreover, deployment is made easier for several reasons:

- Elimination of the local multi-frame clock (LMFC) simplifies frame structure, aiding debug.
- No need to consider PCB trace length adaptation for the SYNC signal as it is re-timed (within individual converters) to the master clock at SYNCO output.
- Elimination of the external SYSREF signal, thus ESIstream often requires no added hardware to achieve determinism.
- Deterministic latency is derived during a one-time training process. Once latency parameters are established, they are fixed for a given design. ESIstream is thus easier to scale into production.

Parametric Consideration	JEDEC JESD204B/C	ESIstream
Maximum Raw Data Throughput (per lane)	12.5 Gbps (B) 32.5 Gbps (C)	12.8 Gbps
Specification body	JEDEC (Industry standard)	Teledyne e2v (license-free)
Data coding	Octets using 8b/10b (B) 64b Words 64b/66b or 64b/80b (C)	14b Words using 14b/16b
Status	Off-the-shelf licensable IP	Off-the-shelf IP (license free)
Primary user benefit	Ultimate flexibility (adds overhead/power)	Substantially lower resource & protocol overheads
Deterministic latency?	Yes (needs extra hardware)	Yes
Noteworthy design considerations	<ol style="list-style-type: none"> 1. Embedded complexity 2. Transport layer complexity 3. Multiple clock domains 4. Trace length matching becomes a critical design feature 	<ol style="list-style-type: none"> 1. Direct integer sample & data relationship (CLK:SSO) No added PLL, minimal hardware 2. Lower link latency 3. No IP license fees 4. Lowest power

Table 2 Summary of features JESD204B/C and ESIstream.



CONCLUSION

Managing system design to ensure deterministic latency is critical in many advanced applications. Absolute latency is rarely the key performance determinant. Rather it is the certainty of a fixed (bounded) latency that matters. Achieving this is progressively challenging in ultra-fast systems since timing margins are stressed more. Fortunately, there is much that specialist component suppliers do to ease these headaches.

In the case of the EV12AQ600, several techniques have been shown:

- From a hierarchical perspective, the simplest aid is the metastability indicator (SYNC_FLAG) which, coupled with SYNC edge control allows adjustment of SYNC phase to avoid disallowed states.
- Next up, the provision of the synchronization output signal (SYNCO). This re-timed SYNC signal can be daisy-chained through a series of ADCs to ensure coherent sampling phase across the extended system.

- Finally, using a master clock combined with a SYNC delay counter/generator logic block offers a simple method to eliminate varying data lane arrival times at the LD.

Controversially, we argue that license-free ESStream, with its simplified data link layer offers a clear advantage in complex systems. JESD204B/C (sub-classes 1 & 2) also provides mechanisms to assure determinism, however it is reported to be harder to establish robust link operation. Many of its technical challenges derive from transport layer complexity arising from the diversity of the operational profiles it supports.

For more information on low overhead ESStream, checkout the online resource centre at ESStream.com. Whichever approach chosen; deterministic latency remains an achievable system design goal.



For further information, please contact:

Romain Pilard,
Applications Engineer,
Signal Processing Solutions
romain.pilard@teledyne.com



For further information, please contact:

Stéphane Breyse,
Applications Engineer,
Signal Processing Solutions
stephane.breyse@teledyne.com



For further information, please contact:

Nicolas Chantier,
Marketing Director
Signal and Data Processing Solutions
nicolas.chantier@teledyne.com



For further information, please contact:

Jane Rohou,
MarCom Manager.
jane.rohou@teledyne.com

