

Successful SAS/SATA Equipment Design and Development: What You Need to Know for 6 Gb/s and Beyond

Abstract:

The new SAS-2 and SATA Gen-3 system protocols enable 6 Gb/s link speeds between storage units, disk drives, optical and tape drives, and protocol host bus adapters (HBAs). Carrying such high signal signals several meters over copper cables, however, tests the limits of signaling technology. As a consequence, the majority of design concerns are expected to arise at the physical layer.

This whitepaper will describe the technical challenges in maintaining signal integrity at 6 Gb/s and how improper test setups can inadvertently degrade signals during product development and testing. Specifically, developers and systems engineers will learn how to deploy test equipment to minimize its impact on signal integrity. In this way, developers can avoid the time-consuming process of attempting to resolve signal integrity issues that are ultimately the result of improper testing practices.

The introduction of 6 Gb/s SAS-2 and SATA Gen-3 promises new levels of performance for networks. At these higher speeds, however, signal integrity becomes a significantly more important design concern for equipment designers and network engineers than it was at 3 Gb/s as tolerances drop to the point where test equipment can adversely affect signal integrity. For example, a test setup that was already at the performance edge for 3 Gb/s will cause undesirable and misleading failures at 6 Gb/s.

Apart from strictly adhering to each standard's specifications, the key behind successful SAS/SATA product development and network debugging will be an understanding of the tighter tolerances at 6 Gb/s and the simple steps that one can take to minimize the impact of test equipment on the device under test. By understanding how attenuation and jitter impact signal integrity, developers and systems engineers can adjust test setups to minimize their impact during testing.

Attenuation and Jitter

The higher frequency signals used by SAS-2 and SATA Gen-3 have increased sensitivity to attenuation and jitter. Higher frequencies attenuate faster than lower frequencies over distance (see Figure 1A). Additionally, higher frequencies are more susceptible to jitter as jitter remains constant even while the signal period decreases (see Figure 1B). When attenuation and jitter become too pronounced, it becomes impossible to accurately sample and decode signals on the receive-side.

The SAS-2 and SATA Gen-3 standards take different approaches to resolving attenuation and jitter issues. SAS-2 utilizes de-emphasis and equalization techniques to minimize the impact of attenuation and jitter on signal integrity (see description below). Alternatively,

SATA Gen-3 employs neither, thus offering a lower-cost link technology for applications that don't require these capabilities. The lack of these capabilities, however, makes SATA Gen-3 more susceptible to attenuation and jitter. As a result, it is even more important to carefully implement the suggestions in this whitepaper when working with SATA Gen-3 and its higher data rate. Specifically, test SATA Gen-3's 6 Gb/s signals with even shorter cables than were used to test 3 Gb/s systems.

De-emphasis and Decision Feedback Equalization

The presence of de-emphasis and equalization not only impacts signal quality, it can affect test setup and operation. On the transmit-side, SAS-2 de-emphasizes lower frequency components to compensate for the expected attenuation of these components (see Figure 2A). On the receive side, SAS-2 selectively boosts higher frequency components using Decision Feedback Equalization (DFE) to even out signal amplitude (see Figure 2B). DFE includes a highly complex state machine and differs from traditional equalization through its ability to reduce jitter.




Figure 1: 6 Gb/s SAS/SATA signals have increased sensitivity to attenuation and jitter. Higher frequencies attenuate faster than low frequencies over distance (A) and are more susceptible to jitter (B). Proper test setups reduce the additional attenuation or jitter introduced into a system during testing.




Figure 2: 6 Gb/s SAS-2 overcomes attenuation effects by de-emphasizing lower frequency components at the transmitter (see Figure 2A) as well as by selectively boosting higher frequency components on the receive side (see Figure 2B) using an equalization method called Decision Feedback Equalization (DFE). Developers designing or troubleshooting SATA Gen-3 equipment must resolve attenuation and jitter issues differently as SATA Gen-3 does not support either de-emphasis or equalization.

If the amount of jitter is too high, the advantages of using a buffered approach are lost. Additionally, by placing an electrical circuit in a signal pathway, buffering can mask channel issues, such as reflections.

Digital Retiming is a method where test equipment operates as a network device at the link layer. At this layer, the test equipment receives signals, decodes them, re-encodes them, and then resends signals on to their destination (see Figure 3B). (For those not familiar with the network stack, devices at the link layer do not receive entire files, for example, but rather receive and resend network traffic frame-by-frame.) It is also important to note that digital retiming can add latency as well as alter clock-alignment commands at the link layer. For example, SAS/SATA ALIGN characters may be utilized to overcome clock skew between the tester and both the host and target. The tester drops or adds ALIGN characters to maintain clock alignment with the devices-under-test. Whether this subtle change to network traffic affects testing depends upon the specific test's goals.

Because analog passthrough tends to have the least impact on the electrical characteristics of a signal, it gives users the most accurate real-world representation of network signals. For SAS systems where de-emphasis and equalization manage attenuation, analog passthrough is the most commonly used connection method.

In some cases, however, analog passthrough fails to maintain sufficient signal integrity. If attenuation issues arise regardless of the shortness of cable length (i.e., if induced attenuation is discovered to be an issue), a buffered approach may provide better results. Likewise, if long cables are necessary, a buffered connection may eliminate attenuation concerns.

For systems where designers have a high degree of confidence in the physical layer of the network and whose concerns reside primarily in the protocol domain, a digitally retimed connection may provide the best results. Digital retiming is also appropriate when a physical setup requires long cables. For 6 Gb/s SATA applications, where the channel model specifies a maximum 1 meter cable length, a passive method such as analog pass-through is not an option since the introduced attenuation would go over the link budget. SATA requires either a buffered or retimed signal.

Since the SAS-2 specification—and the PHY layer in particular—have been under development for close to 3 years, it is perhaps not surprising that most design and test issues relating to SAS-2 are related to the introduction of these new, sophisticated de-emphasis and DFE features. It is critical to follow the specification exactly and set de-emphasis and DFE variables correctly. Failure to do so will not only make testing more difficult but also potentially lead to interoperability difficulties with other SAS-2 equipment.

Connection Methods

While a necessary debugging aide, inserting test equipment between devices-under-test introduces electrical discontinuities into the signal path which induce both jitter and attenuation that can adversely affect signal integrity. There are several methods available for connecting test equipment that reduce or compensate for these effects to varying degrees. Users should be aware of the particular advantages and disadvantages of each in order to select the method best suited for their situation.

Analog Passthrough achieves the lowest impact to signal integrity of the available options by passing the signal through a solid-state switch (see Figure 3A). This method also keeps induced jitter to a minimum. Its primary disadvantage, however, is that an analog passthrough creates a discontinuity in the link, similar in effect to using a connector splice to join two cables. As a result, it attenuates the signal.

Buffered or Re-amplified connections reduce the attenuation induced by test equipment by electrically amplifying signals (see Figure 3C), thus enabling the use of longer cables and providing the best signal integrity in terms of amplitude, i.e., decreased attenuation. The primary disadvantage of buffering a signal, however, is that it introduces non-deterministical jitter.

Pretesting

An important element of testing is working from a stable foundation; if the network infrastructure has unresolved integrity issues, these may incorrectly appear to be caused by the device under test, complicating or delaying problem resolution. Pretesting physical infrastructure (i.e., cables and connectors) confirms the suitability and reliability of a test setup without the device under test present.

Pretesting is essential when first moving to 6 Gb/s because most test setups created for 3 Gb/s applications will simply not perform well enough when signal rates are doubled. These previous test setups may operate at just within the limits of what 3 Gb/s systems can tolerate and so will fail if they are not updated to meet 6 Gb/s requirements. By first characterizing the physical infrastructure between end points as speed-capable, developers can more confidently assume problems found are with the device under test rather than the test setup.

JDSU recommends using a system test suite, such as its Medusa Lab Test Tools, to pretest physical infrastructure. This test suite allow users to generate test patterns specifically designed to test the signal integrity limits of a network. Running a broad spectrum of stress-inducing traffic types across links tests both attenuation and jitter tolerances, revealing whether the physical layer will work at full line rates.

Cabling Issues

A critical element of any test setup is the quality of cables used and how they are connected to the analyzer and device under test. Problems arise when cables fail to meet the standard specifications or when multiple cables are connected such that they introduce discontinuities (i.e., impedance mismatches in the cable).

When standards are being developed, a common assumption used in the mathematical models is the use of single cable, one with no connectors or other impedance discontinuities between

its endpoints. However, in practice, users will often make use of whatever cables they have on hand in the lab, cables which may be longer than necessary, are unshielded, or worse yet, contain multiple cables strung together. Alternatively, developers may assume the test setups used to test the previous generation of devices will work for the latest line rates. In these cases, the test setup may yield attenuation too great for 6 Gb/s devices (i.e., cable length acceptable for 3 Gb/s may be too long, and thus induce too much attenuation, for 6 Gb/s).

Consider the test setup shown in Figure 4. In this example, the Xgig Analyzer generates a signal out to the device under test, which is then captured by the analyzer. At the same time, the original signal is fed back into the analyzer so that the sent and received signals can be compared.

There are several problems with this test setup. The 4-link Hydra cable coming out of the back of the Xgig Analyzer sends the generated signal right back to the Analyzer across a male-to-male cable segment or “splice”. As a result, the electrical characteristics of the link tend to attenuate the signal and make it look like an extended length of cable that is much longer than it actually is. Additionally, both the Hydra cable and splice are shielded much less than higher-quality cables, such as Mini SAS cables, and jitter could be further aggravated by this setup.

Compare this to the test setup in Figure 5. A high-quality, shielded Mini SAS 4-lane cable connects the Xgig Analyzer to the device under test. This cable individually shields each lane as well as the 4-lanes together. It is also a reasonable length and has no discontinuities.

Regarding feeding the original generated signal back to the Analyzer, this is accomplished using a Loopback Plug. Rather than using two cables and a splice to create the signal feedback loop, the Loopback Plug keeps the loop as short as possible and minimizes discontinuities in the link. Internally, the Loopback Plug passes the signal over just a 1 mm trace of copper.




Figure 3: The method used to connect test equipment to the device under test can have a significant impact on signal integrity. Analog Passthrough (A) achieves the lowest impact to the signal's electrical characteristics of the available options by passing signals through a solid-state switch. While keeping induced jitter to a minimum, analog passthrough creates a discontinuity in the link and, as a result, attenuates the signal. Digital retiming (B) mimics a network device at the link layer, receiving signals, decoding them, re-encoding, and then resending to their destination. A buffered or re-amplified connection (C) reduces attenuation to provide the best signal integrity in terms of amplitude, thus enabling the use of longer cables; however, it introduces jitter that can overcome the advantages of using a buffered approach. As it places an electrical circuit in the middle of a link, it may mask electrical channel issues such as reflections.

Other ways developers can improve signal integrity include:

- **Co-locate the analyzer, host, and device under test:** This is one of the simplest ways to reduce cable length. Rather than run long cables between workbenches, use shorter lengths of cable.
- **Use high-quality shielded cables:** High-quality shielded cables are more expensive than others; however, unshielded cables can induce jitter through Electro-Magnetic Interference (EMI), reducing signal integrity and potentially leading developers to attempt to resolve phantom problems that would otherwise not be present with a shielded cable. It is important to note that the shielding of even standard cables may be less than ideal. For example, standard SATA cables have shielding around individual conductors but not around the cable itself. Higher-quality cables can actually accelerate development—and quickly pay for themselves—by avoiding the problems caused by low-quality cabling.
- **Use only 6 Gb/s rated cables:** The 3 Gb/s cables sitting in the lab, such as cables with SAS 4x connectors, were not designed to carry 6 Gb/s within the higher data-rate’s specification. Using inappropriate cables, even if they should work theoretically, could significantly delay development with ‘red-herring’ signal integrity issues.
- **Eliminate all discontinuities:** Every discontinuity (connectors, splices, adapters, etc.) reduces signal integrity by causing reflections that attenuate signals by sending signal energy back in the opposite direction. Remove all but necessary discontinuities. In many cases, this may mean custom test cabling or introducing test fixtures such as JBODs.
- **Use the shortest reasonable cable:** The longer the cable, the greater the attenuation. The best way to reduce attenuation is by shortening cables. This is especially true for SATA, where the absence of the SAS specification’s de-emphasis and equalization mean that the link’s characteristics completely control signal attenuation and jitter.

JDSU provides a complete tool-set for SAS/SATA product design, development, SQA, and manufacturing. Whether your needs are traffic generation, error-injection, or analysis—JDSU has you covered.

- **Don’t use too short a cable for SAS-2:** While this may seem to contradict the previous point, SAS-2 employs de-emphasis and equalization to attempt to overcome the effects of attenuation and jitter. However, as endpoints don’t communicate as to how much de-emphasis and equalization there should be, these values are fixed within a range with expectation that they will work over the specification’s supported cable lengths. If a cable is too short, however, de-emphasis and equalization may actually end up overcompensating for a signal, actually reducing signal integrity. While the specification supposedly can work for any range of cable length, if you encounter issues with short cables working with SAS-2, try experimenting with longer lengths to determine if this is the problem. Note that short cable issues do not apply to SATA as SATA does not support de-emphasis or equalization.
- **Remember the 3 S’es:** use Short, Shielded, and Single cables wherever possible.

Moving to 6 Gb/s increases the difficulty of maintaining signal integrity between network devices. Many developers will struggle with signal integrity issues that unfortunately will arise from their own carelessness in how they manage their test setups. Those developers and systems engineers who respect the tighter tolerances of operating at 6 Gb/s—taking care to use the appropriate connection method, pretesting their systems, and observing proper cable use—will find themselves free to accurately identify and resolve protocol system issues more quickly and painlessly.




Figure 4: Test Setup to Avoid: Signals are 1) sent over a splice, 2) the extended double cable is longer than required, and 3) the cable used is unshielded.



Figure 5: Appropriate Test Setup: 1) Signals are sent over a single cable, 2) the cable is a minimum length (short), and 3) the cable is properly shielded and rated for 6 Gb/s use.

Test & Measurement Regional Sales

NORTH AMERICA TEL: 1 888 746 6484 sales-snt@jdsu.com	ASIA PACIFIC apacsales-snt@jdsu.com	EMEA emeasales-snt@jdsu.com	WEBSITE: www.jdsu.com/snt
---	---	---------------------------------------	--