

Factors Affecting Cable Length in 3G systems — It's Not Just About Return Loss

By Peter Krywiak
Senior Design Engineer

The typical broadcast environment interconnects a multitude of specialized audio/video equipment — from cameras to routers, both analog and digital, and from vendors from around the world. Each piece of equipment must be able to reliably interconnect at expected cable lengths to its corresponding signal type. To facilitate this interoperability, SMPTE has been tasked with providing guidelines to the broadcast industry.

Interoperability requires that the following two parameters are satisfied: the content (timing and format) and the physical layer. This paper will look at the physical layer, and 3G-SDI signals in particular. The discussion will focus on how the physical layer is measured and what these measurements mean.

The key physical layer properties are:

- Rise/Fall Time — Noise vs. System Performance: General compromises...
- Overshoot Measurements — Fast Scopes vs. WFM: What do we really want?
- Jitter — How many reentries can we support? What helps us get more?
- Return Loss vs. Cable EQ — Minimizing Bit Error Rate

To simplify the discussion, we will look at a signal chain that begins at a source and ends at a receiver, and is restricted to baseband SDI signals. A source is where the signal originates in its baseband format, and the receiver is where it ends in its baseband format.

In the analog domain, the physical layer properties will have a direct effect on the timing and quality of the video/audio received. In the digital domain, however, we only care about receiving the signal error free. If the signal is received error free, the original content can be perfectly extracted.

To examine all the physical layer properties except return loss, we need an instrument that can create an eye diagram from an externally applied SDI signal. Let us examine the creation of an eye diagram.




Figure 1. Sequence of bits representing the number 1101011001111




Figure 2. Same sequence as above, but showing finite rise time of signal between bits




Figure 3. Creating an Eye Diagram

Figure 1 may be familiar to many readers as a bitstream. Although it may not be obvious to some, these are NRZI-encoded bits, where 1's are represented by a transition, and 0's are represented by no transition. In reality, signals don't transition from high to low and low to high instantly. It takes time for signals to transition, and Figure 2 shows this.

Figure 2 is a snapshot in time of a bitstream. Sampling oscilloscopes and WFMs will create an eye diagram (Figure 3) by overlaying sequential portions of a stream one on top of the other. By having a fixed frame of reference, it is easy to see the cumulative effect of things like jitter, overshoot, etc.

Rise/Fall Time — Noise vs. System Performance: General compromises...

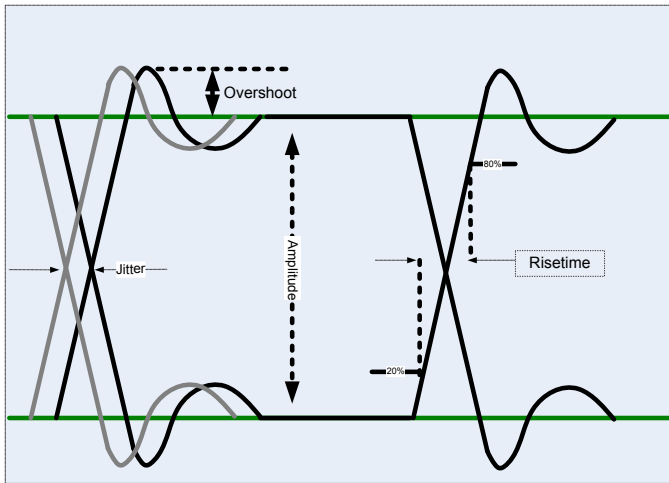


Figure 4. Key physical layer properties that can be observed and measured in an eye diagram

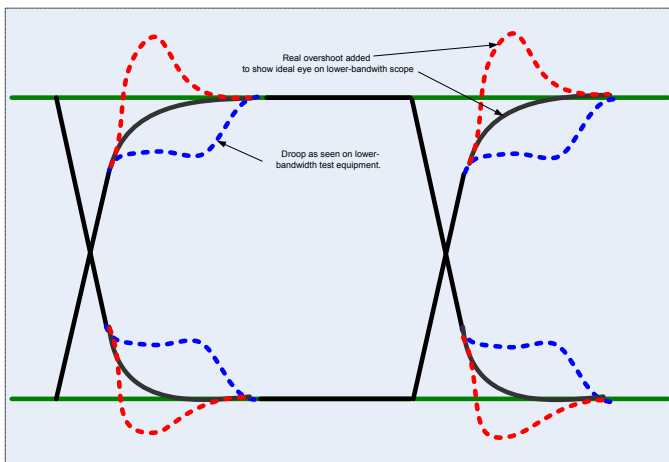


Figure 4b. The bandwidth of test equipment can determine the look of your eye.

Figure 4 shows the eye diagram close-up, enabling us to see the imperfection of the signal. SMPTE has provided guidance as to what a 3G-SDI signal should look like coming out of a source measured across a 75-ohm load. The load is internal to a scope, so it does not have to be added externally. The signal should have a peak-to-peak swing of 800 mV +/- 10%; the rise/fall time should be between 50 and 135 ps; and the overshoot should not be more than 10%.

A signal adhering to these limits creates spectrum that most EQs in receivers are tuned to. An EQ uses this reference information to determine the amount of cable (loss) in front of it and apply the correct amount of gain to equalize it. Designers have some control over the rise/fall time. Reducing it reduces the amount of EMI, but decreases the effective bit-period — thus reducing jitter margin.

Overshoot Measurements — Fast Scopes vs. WFM: What do we really want?

Manufacturers can adjust matching networks on their PCB to optimize the output return loss and/or improve the “look” of the eye on a lower-band scope like WFM. This may negatively impact the optimal signal spectrum and reduce the maximum cable length by introducing excessive overshoot.

Some manufacturers have access to expensive, high-speed, high-bandwidth scopes. These scopes will likely see a good eye, where a lower-bandwidth scope will see a droop in the rise time. To give customers what they want to see, some manufacturers have had to add real overshoot to the signal, so that the signal can be viewed as ideal on a WFM. See figure 4b.

Jitter — How many reentries can we support? What helps us get more?

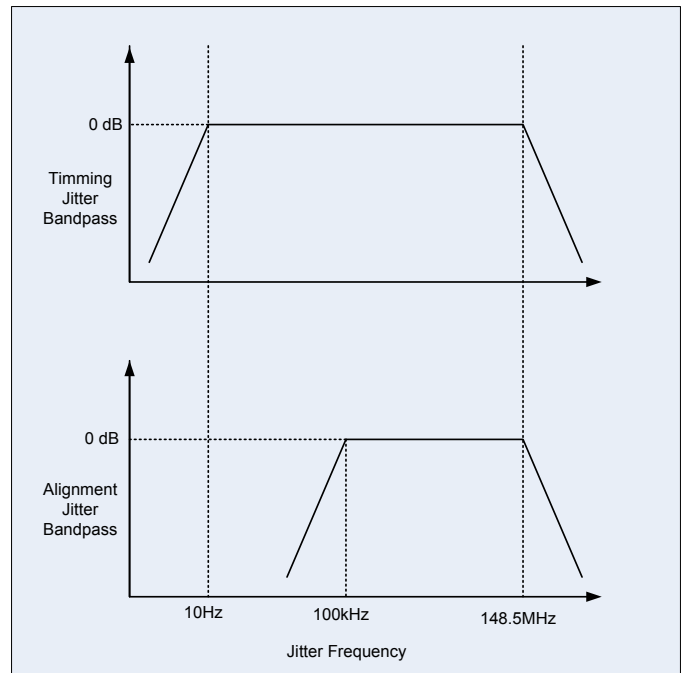


Figure 5. 3G SDI jitter bandwidths

SMPTE defines the allowable jitter out of a source with the pass-band filters shown in Figure 5. For a 3G-SDI signal, the timing jitter cannot exceed 2 UI. 2 UI is equal to two bit periods, which at 3G translates to 637 ps. Timing jitter includes low-frequency and high-frequency jitter. Alignment jitter is the high-frequency subset of timing jitter. In a simple source-to-destination setup, most receivers can handle many times more timing jitter than 2 UI, assuming that most of it is low-frequency jitter (alignment jitter is small). In multi-pass configuration, such as signals passing through multiple routers and DAs, low-frequency jitter can not be cleaned by reclockers, and it accumulates; therefore, output jitter should be kept as low as possible to allow enough margins for multiple passes.

Alignment jitter, which is essentially the high-frequency portion of timing jitter, is cleaned by reclockers in routers and DAs; however, if it is too high, it will cause errors. The alignment jitter can be measured by WFMs, and most real-time oscilloscopes with the required sampling rate and upper bandwidth. However, real-time scopes are not able to accurately measure timing jitter because their bandwidth lower limit will only go down to a few kHz. Equivalent time scopes such as most WFMs can help us view this band of jitter. Jitter is one of the most difficult physical layer properties to control cost-effectively in routers.




Figure 5b. Eye with 0.5 UI of jitter

Figure 5b shows an eye diagram half closed due to 0.5 UI of jitter. As jitter closes the eye horizontally and Noise closes the eye vertically, the sampling area gets squeezed, increasing the chance to sample incorrectly and get bit errors.

Return Loss vs. Cable EQ — Minimizing Bit Error Rate

The minimum return loss measured at the end of a 1 m cable is given in Figure 6 below.

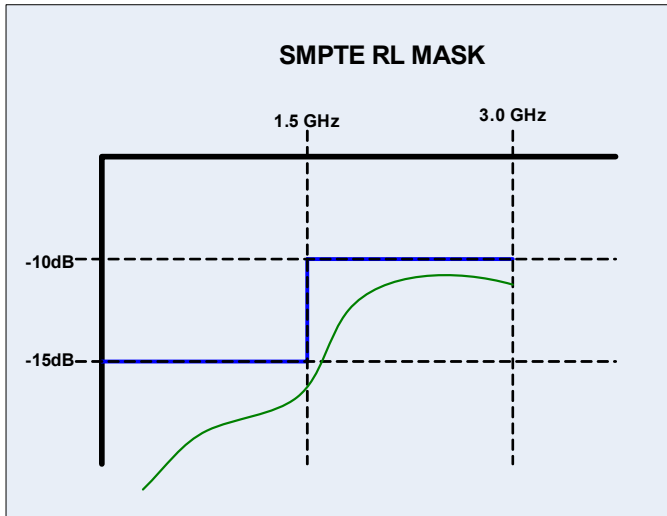


Figure 6. Return Loss template specification (blue) and a compliant return loss measurement (green)

Figure 6 shows the RL mask. The blue line is the required return loss limit; the green swerve is the hypothetical performance of a system exceeding the spec.

Most, if not all, sources use a SMPTE-compliant cable driver. There are several IC vendors that offer these devices, and they come with the required application circuit to create a voltage swing and rise/time compliant to the SMPTE specification. They are also designed to be impedance-matched to a 75-ohm BNC through a short PCB trace in order to meet the required output return loss specification.

Most, if not all, receivers use a SMPTE-compliant equalizer. They are offered by the same IC vendors as the cable drivers and come with the required application circuit for these devices to recover the coax cable loss between the source and receiver. They are also designed to be impedance-matched through a short PCB trace to a 75-ohm BNC in order to meet the required input return loss specification.

With the availability of standardized interface ICs, one would expect that all equipment would perform similarly; however, maximum cable length and return loss varies significantly between different pieces of equipment, even by the same manufacturer. Why is this so?

The answer lies in the breadth of equipment configuration offered. The form factor and density offered in the industry varies greatly, as manufacturers try to offer more I/O and more processing in less space. They also want to make cards easy to replace, without having to disconnect cables from behind densely packed racks.

This has made it impossible to follow the simple application circuit from the IC vendors, which requires the IC to be placed very close and be directly connected to the ideal BNC.

In reality, there are several factors that complicate the signal path. These are:

- The BNC has to meet physical space requirements, so its connection may not be ideal.
- The BNC may be on a different board from that of the IC, in which case it must pass through an additional connector.
- The density of I/O may force some signals to travel further across the PCB and through layers.
- Design for manufacturing constraints may prevent the use of known optimization techniques such as back-drilling vias.

Between every change in the signal path, there is some impedance mismatch. Each impedance mismatch causes part of the signal to be reflected back, thus diminishing return loss. See figure 7.

Despite these challenges, system designers skilled at signal integrity can still achieve acceptable cable lengths and return loss.

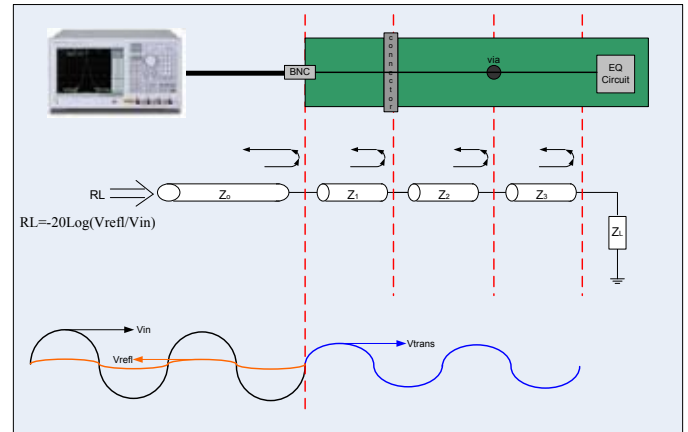


Figure 7. What is return loss?

In Figure 7, we see a vector network analyzer (VNA) connected through a 1 m cable to the BNC of a receiver. The VNA plots the return loss of a load over a specified frequency range. The VNA sweeps the desired range with bursts of sine waves, then measures the signal reflected back. Using equation $RL = -20 \cdot \text{Log}(V_{\text{refl}}/V_{\text{in}})$, the return loss is calculated and plotted as the green line in Figure 6.

As a point of reference, a return loss of 10 dB means 32% of the signal was reflected back, and 15 dB means 18% of the signal is reflected back.

By telling us how much signal is reflected, return loss also gives us an indication of the amount of signal transmitted. If we read a return loss plot and see that $RL = 15\text{dB}$ at 1.5 GHz, it means that 18% of the signal was reflected — so no more than 82% of the signal was transmitted to the receiver IC. Therefore, the return loss tells us roughly what fraction of the input signal at the BNC is actually transmitted to the receiver chip.

A big challenge is to match the load (IC circuit) to the PCB traces and connectors leading to the BNC. This can also potentially have the largest impact on the return loss measurement.

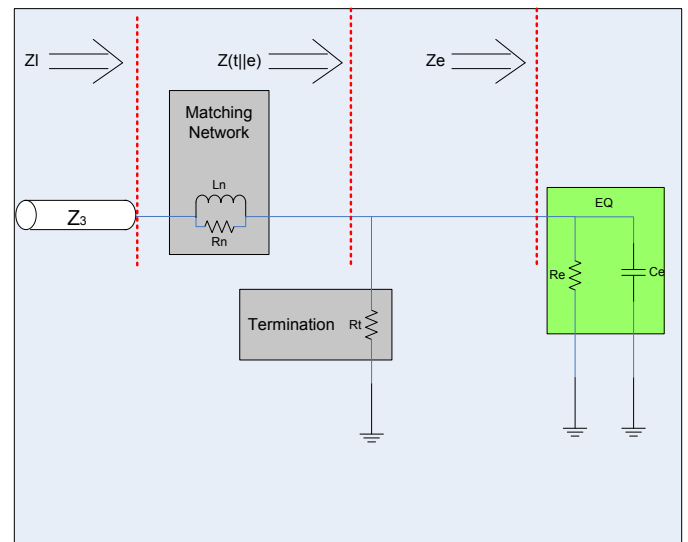
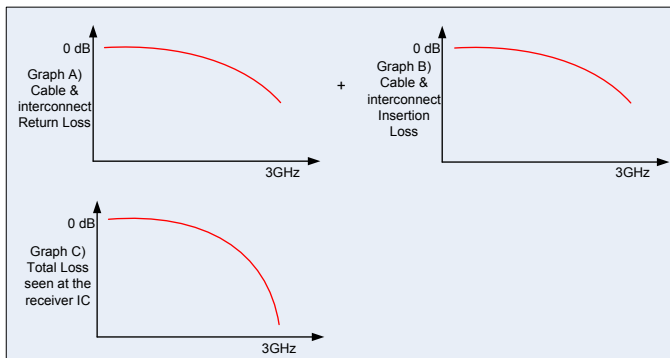


Figure 8. Simplified receiver circuit showing matching network, termination and an equalizer


Figure 8 shows a simplified receiver circuit. The IC has an input impedance Z_e , which varies with frequency inversely with a time constant $Re||C_e$. Alone, the IC could not match to the PCB trace impedance Z_3 . The return loss and cable performance would be very bad. To help match the IC to the PCB trace Z_3 , with a characteristic impedance of 75 ohms, we can add a 75-ohm termination resistor R_t close to the EQ input. Now the impedance looking into the circuit at the termination resistor is $Z(t||e)$, is close to 75 ohms at low frequency, but goes to zero as the frequency rises above 1 GHz.

The circuit needs to be compensated at high frequency as well. This is accomplished with the matching network in series with the signal. The matching network consists of a resistor in parallel with the inductor, and its impedance varies inversely with the internal impedance of the EQ made up of a capacitor C_e in parallel with resistor R_e . Theoretically, these two frequency-dependant impedances can be designed to cancel each other out, leaving a purely resistive 75-ohm load impedance of R_t , which would cause no reflections at all.

The downside is that this matching network is now a low-pass filter in front of the EQ. So to the outside world, the return loss is better because less energy is being reflected back from the BNC; however, the insertion loss is greater because the energy is not reaching the EQ — but instead being absorbed by the matching network.



- Graph A) shows the reflection loss over frequency due to the reflections coming back from the receiver. Higher frequency components are reflected more.
- Graph B) shows the insertion loss over frequency of the whole channel (including cable). Higher frequency components are absorbed more.
- Graph C) shows the combined loss suffered by the signal by the time it arrives at the receiver IC.




- Graph C) shows total loss applied to the signal through insertion and return loss.
- Graph D) shows the gain applied to a signal by the EQ.
- Graph E) shows net loss to the signal after equalization.

In an ideal world, Graph E would show a flat line; however, in the real world, there is a limit to how well the EQ can recover the loss of a channel, so there is some residual loss.

Graphs A) and B) are interrelated, and there is an optimal design that enables the lowest attenuation and greatest cable length.

As loss increases, the jitter coming out of the EQ increases. As the loss of the system exceeds the equalization capability of the EQ, the jitter out of the EQ increases exponentially.



- Graph F) shows the relationship between loss into and jitter out of an EQ.
- With the signal coming from a SMPTE-compliant source. The green zone is a safe amount of loss.

If the receiver is something that needs to process, display or store the video contained in the SDI stream, then it will have a de-serialiser that converts the SDI to parallel data. The de-serialiser has a Phase Locked Loop (PLL) to lock to the incoming SDI stream. A PLL has a certain input jitter tolerance (ability to lock to a signal in the presence of jitter). This tolerance will then limit the total cable length due to the relationship between loss and jitter. De-serialiser implementations vary from design to design, and some de-serialisers use PLLs with greater input jitter tolerance than others.

If the receiver is a router or DA, it will most likely have a reclocker to clean the jitter before sending it out again. A reclocker also has a PLL with input jitter tolerance. As was mentioned previously, jitter accumulates through multiple passes; therefore, with a single pass, you may be able to achieve a maximum cable length that you might not be able to achieve with multiple passes. It may be necessary to scale back the cable length to achieve error-free multi-pass.

Conclusion

There is no doubt that return loss is a very important physical layer parameter, and that low return loss does not mean longer cable length; however, has specmanship of this parameter created a point of diminishing return, and even negative return? As with any engineering decision, there is a tradeoff: Users of broadcast equipment need to ask themselves if they want the highest return loss a piece of equipment can offer, or the return loss that maximizes the cable performance.

For more information, please visit www.broadcast.harris.com.

Harris is a registered trademark of Harris Corporation. Trademarks and tradenames are the property of their respective companies.