Proposed Draft

Serial ATA International Organization

Version 4
17-Dec-2018

Serial ATA Revision 3.4 Technical Proposal # 083
Title : Gen3 Device Tx Emphasis Change

This is an internal working document of the Serial ATA International Organization. As such, this is not a completed standard and has not been approved. The Serial ATA International Organization may modify the contents at any time. This document is made available for review and comment only.

Author Information

Author Name	Company	Email address
Alvin Cox	Seagate	alvin.cox@seagate.com

Workgroup Chair Information

Workgroup (Phy, Digital, etc…)	Chairperson Name	Email address
PHY	Charles Hill	cphill@altaeng.com

Document History

Version	Date	Comments
0	19-Oct-2018	Initial release.
1	29-Oct-2018	Updated track changes coilors, Figure 152, and changed scope references to HBWS.
2	20-Nov-2018	Added backup material, changed alternate method to Peak-Mode, and made Peak-Mode the device TX Emphasis test method.
3	26-Nov-2018	Added back old method as option for device TX Emphasis measurement method in addition to the Peak-Mode method.
4	17-Dec-2018	Restored nominal value to oxisting method and lowered new method upper limit to match old method.

Introduction

The current SATA specification (Revision 3.4 Gold) has a normative specification for Transmitter Emphasis that uses a technique for measurement that has been found to exclude certain device Phy topologies. Instead of adopting a histogram method used by another standard in the industry, SATA chose to create a new means of identifying the transition bit peak focusing within the measurement interval 0.45 UI to 0.55 UI. When the Tx Emphasis specification was developed, the Phy topologies used to verify the measurement technique supported the method and no issues were expected moving forward.

There are compliant device Phy topologies that shift the timing of the transition bit peak (e.g., rise time, Tx emphasis profile, etc) outside the 0.45 to 0.55 UI range. The measurement methodology expecting the Tx Emphasis peak location to fall within 0.45 to 0.55 UI incorrectly calculates the effective emphasis for such device implementations. In most of these topologies, the Tx Emphasis is understated by the current measurement method. The effects of device transmitter over-emphasis have shown to include resonances in short reach systems, data errors, and interoperability issues.

To alleviate the issue with differing device Phy topologies, this TPR proposes adding a second Tx Emphasis measurement technique reliably used by other standards.

This proposal applies the new technique to only be used for Device Tx Emphasis measurements. This new technique has a fundamental value limit of greater than or equal to zero. It cannot be used for measuring Host Tx Emphasis because the Host Tx Emphasis range includes negative values.

To prevent disruption to the existing base of SATA compliant components, the current Device Tx Emphasis measurement technique is retained as an alternate measurement method and the new technique is added with the requirement that the device shall pass at least one of the test methods to achieve specification compliance for Device Tx Emphasis. There is precedence in the current specification that allows two means of specifying a value to allow compliance of components.

Technical Specification Changes

1.1 Gen3 DeviceTx Emphasis Measurement Change

Existing Table (next page):

Table 54 – Transmitted signal requirements
(part 1 of 4)

Parameter	Units	Limit	Electrical Specification Gen1i	Gen1m	Gen2i	Gen2m	Gen3i	Detail Cross-Ref Section	Measurement Cross-Ref Section
$V_{diffTxdevice}$, Tx Differential Device Output Voltage	mVppd	Min	400		400		-	7.4.3.3.2	7.6.7
		Min	-		-		240		7.6.5 7.6.5.4
		Nom	500		-		-		7.6.7
		Max	600		700		-		
		Max	-		-		900		7.6.5
$V_{diffTxhost}$, Tx Differential Host Output Voltage	mVppd	Min	400		400		-	7.4.3.3.2	7.6.7
		Min	-		-		200		7.6.5 7.6.5.4
		Nom	500		-		-		7.6.7
		Max	600		700		-		
		Max	-		-		900		7.6.5
UI_{VminTx}, Tx Minimum Voltage Measurement Interval	UI		0.45 to 0.55		0.45 to 0.55		-	7.4.3.3.3	7.6.7
			-		-		0.45 to 0.55		7.6.7.3
$V_{EmphasisDevice}$, Device TX Emphasis [a b c]	dB	Min	-		-		0.5	7.4.3.3.14	7.6.33
		Nom	-		-		2		
		Max	-		-		2.5		
$V_{EmphasisDevice}$, Peak-Mode Device TX Emphasis [a b c]	dB	Min	-		-		0.5	7.4.3.3.14	7.6.33
		Nom	-		-		-		
		Max	-		-		2.5		
$V_{EmphasisHost}$, Host TX Emphasis [a c d]	dB	Min	-		-		-2	7.4.3.3.14	7.6.33
		Nom	0	-	0	-	0		
		Max	-		-		1.5		

[a] TX emphasis is measured at the mated Serial ATA connector at the device. Unless otherwise specified, no CIC is used for this measurement.
[b] The TX emphasis requirement does not apply to the Internal 4-lane cable mated to a backplane or devices using the Internal LIF-SATA connector, the SATA MicroSSD interface, or the Internal M.2 connector. Tx emphasis for these cases is vendor specific.
[c] The Device shall pass at least one of the Device TX Emphasis test methods to achieve specification compliance for Device Tx Emphasis.
[c d] This measurement includes the channel between the host IC and the device connector (e.g., motherboard, backplane, cable, and connectors). For internal 1 m cabled host to device applications (see section 5.3.2), the CIC is substituted for the actual system cable.

7.6.33 Transmitter emphasis (Gen1i, Gen1u, Gen2i, Gen2u, Gen3i, Gen3u)

7.6.33.1 Transmitter emphasis overview

The transmitter emphasis values specified in Table 54 and Table 55 refer to the output signal from the UUT at the mated connector of the device. The host transmit emphasis is specified while attached to the system cable(s), connectors, motherboard, and/or backplane that are used when integrated with the device. For the alternate compliance method specified in Section 7.4.3.3.14.2, the measurement is made with the CIC rather than with the actual system cable.

The Device shall meet at least one of the Device Tx Emphasis specifications in Table 54 according to the measurement method defined in 7.6.33.3 or 7.6.33.4. Verification of compliance with both methods shall not be required. The measurement methods verify that Device Tx Emphasis is present and is bounded. It is up to the device manufacturer to determine the level of Device Tx Emphasis.

The Host shall meet the Host Tx Emphasis specifications in Table 54 and Table 55 according to the measurement method as defined in 7.6.33.3.

7.6.33.2 Transmitter emphasis measurement (Gen1i, Gen1u, Gen2i, Gen2u, Gen3i, Gen3u)

~~Transmitter emphasis is measured by comparing the mean differential voltage of the first bit of MFTP versus the mean differential voltage of the second bit of MFTP. Emphasis specifications in Table 54 and Table 55 shall be met according to the measurement method as defined in 7.6.33.3.~~

Figure 250 and Figure 251 show the test setups for measuring emphasis. The HBWS is the standard for measuring emphasis. The losses in the test connections may be significant so it is prudent to minimize and estimate these.

Several methods may be used to estimate the cabling losses:

a) use two cables of different lengths and compare the losses of each;
b) rely on published data for the cables; or
c) obtain a separate means for measuring the cable loss (e.g., characterization with a network analyzer or power meter).

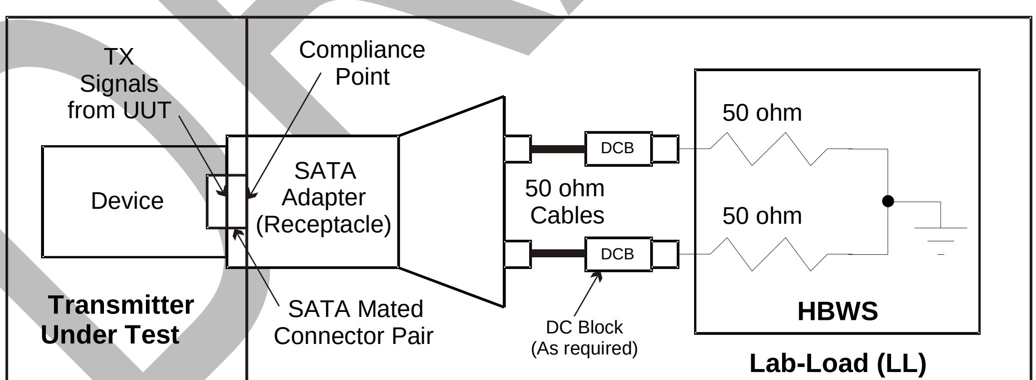

Figure 251 – Device transmit emphasis test with Lab-Load (LL)

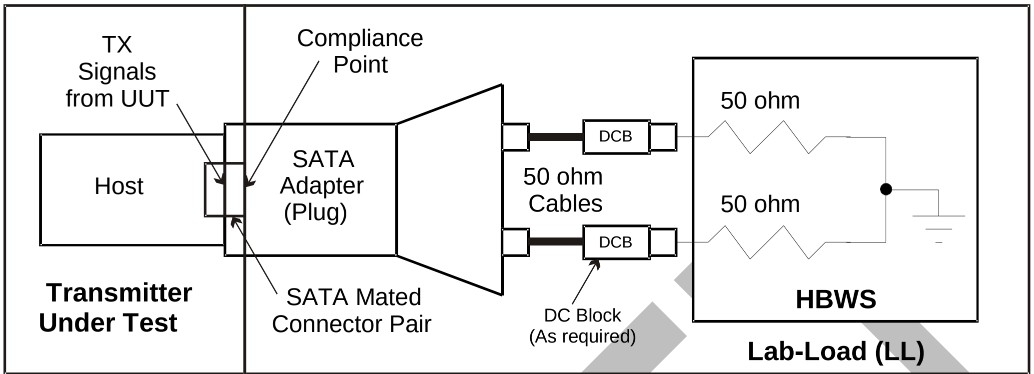

Figure 251 – Host transmit emphasis test with Lab-Load (LL)

This specification describes emphasis levels in terms of voltage amplitude ratio in dB while driving a test load of 100 ohm differential (i.e., lab-load) and 50 ohm single ended to ground.

7.6.33.3 Measurement of emphasis (Gen1i, Gen1u, Gen2i, Gen2u, Gen3i, Gen3u)

Transmitter emphasis is measured by comparing the mean differential voltage of the first bit of MFTP versus the mean differential voltage of the second bit of MFTP.

To test for emphasis, use the following steps:

Step 1, transmitting an MFTP pattern, for a UI corresponding to the first 1 bit, construct a histogram based on n samples collected in the waveform epoch [0.45 UI to 0.55 UI] for the UI. The number of samples in a histogram (n) for the UI shall be greater than or equal to 100 and shall meet the requirement that:

$$1\,537\,(s/\bar{x})^2 \le n$$

where:

- $\bar{x}$ is the mean of the voltage samples in the histogram that may be read from the HBWS in histogram measurement mode;
- s is the standard deviation of the voltage samples in the histogram that may also be read from the HBWS; and
- n is the number of samples that contribute to the histogram – this may also be read from the HBWS.

The inequality above is based on a requirement that enough samples are collected to define a confidence interval with at least 95 % probability and with a width no greater than 10 % of the sample mean.

Call the mean,

$$A = \bar{x}$$

Step 2, transmitting an MFTP pattern, for a UI corresponding to the first 0 bit, construct a histogram based on n samples collected in the waveform epoch [0.45 UI to 0.55 UI] for the UI. The number of samples in a histogram (n) for the UI shall be greater than or equal to 100 and shall meet the requirement that:

$$1\,537\,(s/\bar{x})^2 \le n$$

where:

- $\bar{x}$ is the mean of the voltage samples in the histogram that may be read from the HBWS in histogram measurement mode;
- s is the standard deviation of the voltage samples in the histogram that may also be read from the HBWS; and
- n is the number of samples that contribute to the histogram – this may also be read from the HBWS.

Call the mean,

$$B = \bar{x}$$

Step 3, transmitting an MFTP pattern, construct a histogram based on n samples collected in the waveform epoch [0.45 UI to 0.55 UI] for the UI of the last 1 bit. The number of samples in a histogram (n) for the UI shall be greater than or equal to 100 and shall meet the requirement that:

$$1\,537\,(s/\bar{x})^2 \le n$$

where:

- $\bar{x}$ is the mean of the voltage samples in the histogram that may be read from the HBWS in histogram measurement mode;
- s is the standard deviation of the voltage samples in the histogram that may also be read from the HBWS; and
- n is the number of samples that contribute to the histogram – this may also be read from the HBWS.

Call the mean,

$$C = \bar{x}$$

Step 4, transmitting an MFTP pattern, construct a histogram based on n samples collected in the waveform epoch [0.45 UI to 0.55 UI] for the UI of the last 0 bit. The number of samples in a histogram (n) for the UI shall be greater than or equal to 100 and shall meet the requirement that:

$$1\,537\,(s/\bar{x})^2 \le n$$

where:

- $\bar{x}$ is the mean of the voltage samples in the histogram that may be read from the HBWS in histogram measurement mode;

s   is the standard deviation of the voltage samples in the histogram that may also be read from the HBWS; and
n   is the number of samples that contribute to the histogram – this may also be read from the HBWS.

Call the mean,

$$D = \bar{x}$$

Step 5, from A, B, C, and D obtained in steps 1 to 4, compute:

$$V_{Emphasis} = 20 \times \log_{10}[(A - B)/(C - D)]$$

The test for minimum device emphasis is passed if:

$$V_{Emphasis} \geq V_{EmphasisDevice}(Min)$$

The test for maximum device emphasis is passed if:

$$V_{Emphasis} \leq V_{EmphasisDevice}(Max)$$

The test for minimum host emphasis is passed if:

$$V_{Emphasis} \geq V_{EmphasisHost}(Min)$$

The test for maximum host emphasis is passed if:

$$V_{Emphasis} \leq V_{EmphasisHost}(Max)$$

See Table 54 and Table 55, according to 7.4.2 for $V_{EmphasisDevice}$ and $V_{EmphasisHost}$, otherwise the test for emphasis has not been passed.

7.6.33.4 Peak-Mode Device TX Emphasis measurement (Gen1i, Gen2i, Gen3i)

The Peak-Mode device TX Emphasis measurement applies to devices only and shall be based on the following values:
a) VMA: a mode (i.e., the most frequent value of a set of data) measurement; and
b) V_{P-P}: a peak to peak measurement with a repeating LFTP pattern (i.e., D30.3).

The VMA and V_{P-P} measurements shall be made with the transmitter device terminated through the compliance point into a lab-load as shown in Figure 250.

The VMA and V_{P-P} measurements shall be made using a HBWS with a histogram function with the following or an equivalent procedure:
1) calibrate the HBWS for measurement of a 3 GHz signal; and
2) determine VMA and V_{P-P} as shown in Figure 252. A sample size of 1,000 minimum to 2,000 maximum histogram hits for VMA shall be used to determine the values. The histogram is a combination of two histograms (i.e., an upper histogram for TX+ and a lower histogram for TX-). The histograms on the left represent the test pattern signal displayed on the right. VMA and V_{P-P} are determined by adding the values measured for TX+ and TX-.

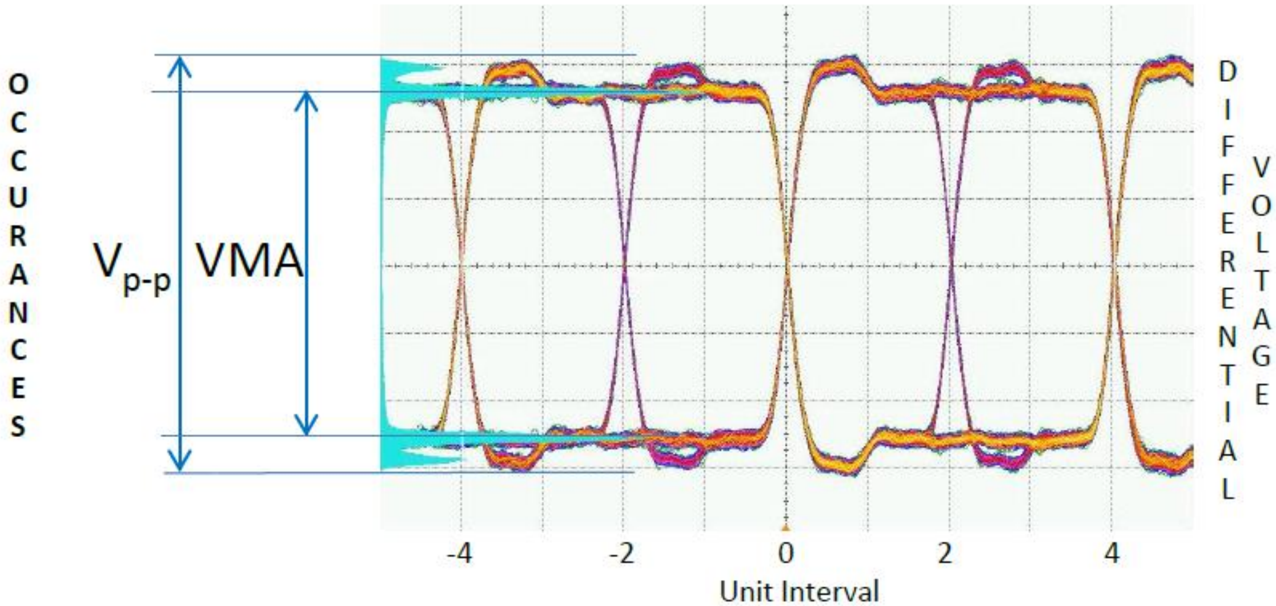

Figure 252 - Transmitter Emphasis Measurment (Alternate Method)

The following formula shall be used to calculate the transmitter equalization value:

Transmitter equalization = 20 × log(V_{P-P} / VMA) dB

where:

V_{P-P} is the peak to peak value; and

VMA is the mode value.

Backup Material

NOTE: In this material, the Tx Emphasis measurement method of 7.6.33.3 is referred to as TSG-17 (designation from the SATA compliance specification).

Introduction

- The purpose of Tx emphasis is to contain ISI through a channel, which can be achieved with various device topologies.
- The SATA spec does not define Tx emphasis topology (full-bit, partial-bit), but it does define the emphasis measurement technique in TSG-17.
- The SATA emphasis measurement technique defined in TSG-17 can be imprecise for transceivers in which the peak emphasis amplitude occurs after ½ UI point of the transition bit.
- Depending on the device topology, this can lead to SATA products that are compliant to the current SATA specification using TSG-17, but have interoperability issues due to over-emphasis.

Emphasis Measurement

- TSG-17 dictates that the two emphasis measurements occur at the ½ UI point of the first and second bit, as illustrated in the slide below (vertical blue lines).
- However, with some compliant transceivers, the 1st bit may not reach peak amplitude at the first ½ UI point. In addition, the amplitude may not drop to the plateau level at the ½ UI point of the 2nd bit. As seen in the figure below using an LFTP pattern, the emphasis value that is numerically derived by using the two ½ UI measurement points noticeably understates the emphasis on the signal (horizontal blue lines).
- The peak-mode technique for emphasis measurements used by other specifications reports a much higher emphasis value (horizontal orange lines).

Quantifying Emphasis Effectiveness

- Quantifying the effectiveness of the emphasis is key to understanding the differences between measurement techniques and PHY topologies.
- This was done by mapping measured emphasis to ISI.
- Hypothetical data is used in this section to illustrate how the plots in the Measurement section are created.
- A fixed Rj value is used for the ISI measurements. In the Measurement section, this Rj is measured on the near-end using a 1T (clock) pattern for each device tested.
- Keysight compliance applications for SATA (TSG-17) and SAS (Peak-Mode) are used for the emphasis measurements.

MFTP Pattern → Using TSG-17; LFTP Pattern → Using Peak-Mode; LBP Pattern → Change in ISI (ps)

Drive Emphasis Register Setting	Measured Emphasis on Near-End (dB)		Change in ISI (ps)	
	Using TSG-17	Using Peak-Mode	Near-End	Far-End
0	-1	1	0	0
1	0	2	2	-3
2	1	3	4	-6
3	2	4	6	-3
4	3	5	8	0

- Drive Emphasis Register Setting: The value in the Tx emphasis register field programmed in the PHY.
- Using TSG-17: The Tx emphasis as measured using TSG-17
- Using Peak-Mode: The Tx emphasis as measured using the SAS Peak-Mode technique
- Near-End / Far-End: The change in ISI at both the drive connector and through a CIC

Measuring emphasis as a function of drive emphasis settings.

x-axis	y-axis	y-axis		
Drive Emphasis Register Setting	Measured Emphasis on Near-End (dB)		Change in ISI (ps)	
	Using TSG-17	Using Peak-Mode	Near-End	Far-End
0	-1	1	0	0
1	0	2	2	-3
2	1	3	4	-6
3	2	4	6	-3
4	3	5	8	0

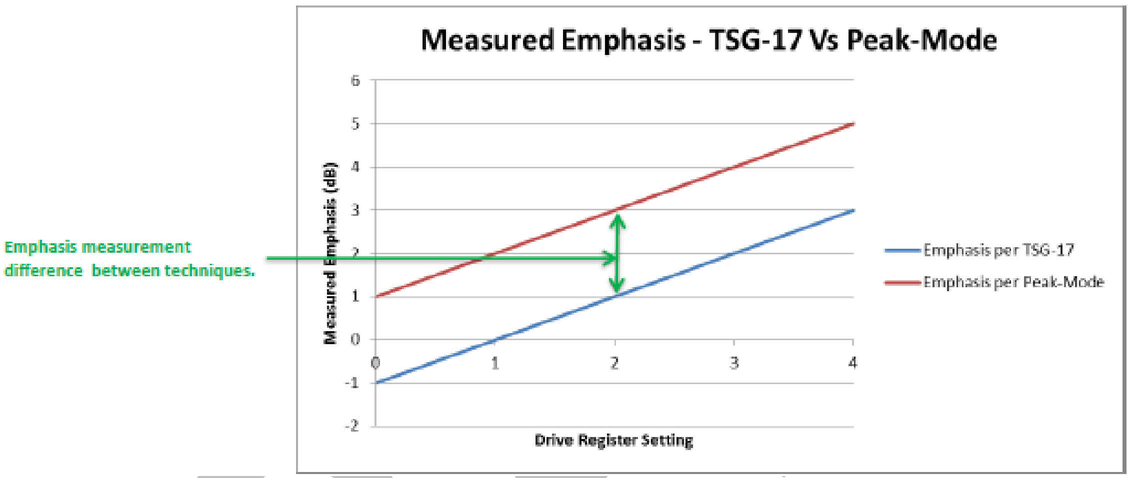

- Measuring the change in ISI as a function of drive emphasis setting.
- Starting at 0 emphasis, the base amount of ISI through the CIC is reduced with increasing emphasis until a minimum point is reached, then the ISI increases as more emphasis is applied.

x-axis				y-axis	
Drive Emphasis Register Setting	Measured Emphasis on Near-End (dB)		Change in ISI (ps)		
	Using TSG-17	Using Peak-Mode	Near-End	Far-End	
0	-1	1	0	0	ISI(Emphasis 0) - ISI(Emphasis 0)
1	0	2	2	-3	ISI(Emphasis 1) - ISI(Emphasis 0)
2	1	3	4	-6	ISI(Emphasis 2) - ISI(Emphasis 0)
3	2	4	6	-3	ISI(Emphasis 3) - ISI(Emphasis 0)
4	3	5	8	0	ISI(Emphasis 4) - ISI(Emphasis 0)

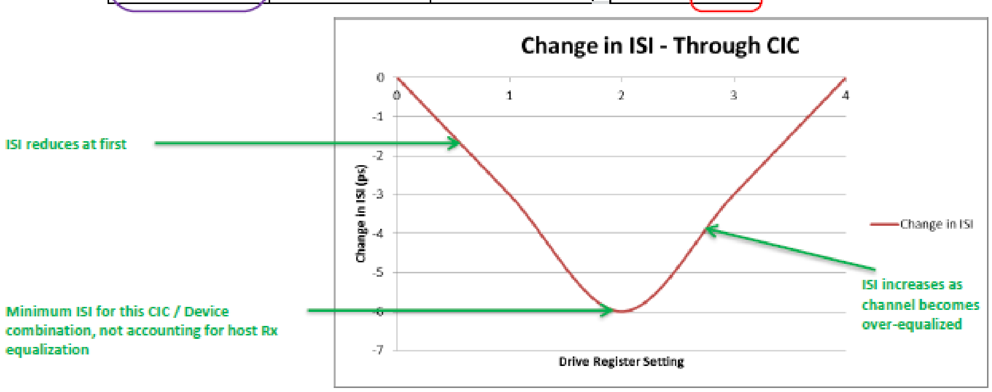

- These curves map measured emphasis to ISI by changing the x-axis from register settings to measured emphasis.
- For a given drive, the ISI plotted on the y-axis is the same since the register setting is the same, but note that the curves shift on the x-axis due to measurement technique.
- This allows us to compare different emphasis measurement techniques and devices.

	x-axis	x-axis		y-axis
Drive Emphasis Register Setting	Measured Emphasis on Near-End (dB)		Change in ISI (ps)	
	Using TSG-17	Using Peak-Mode	Near-End	Far-End
0	-1	1	0	0
1	0	2	2	-3
2	1	3	4	-6
3	2	4	6	-3
4	3	5	8	0

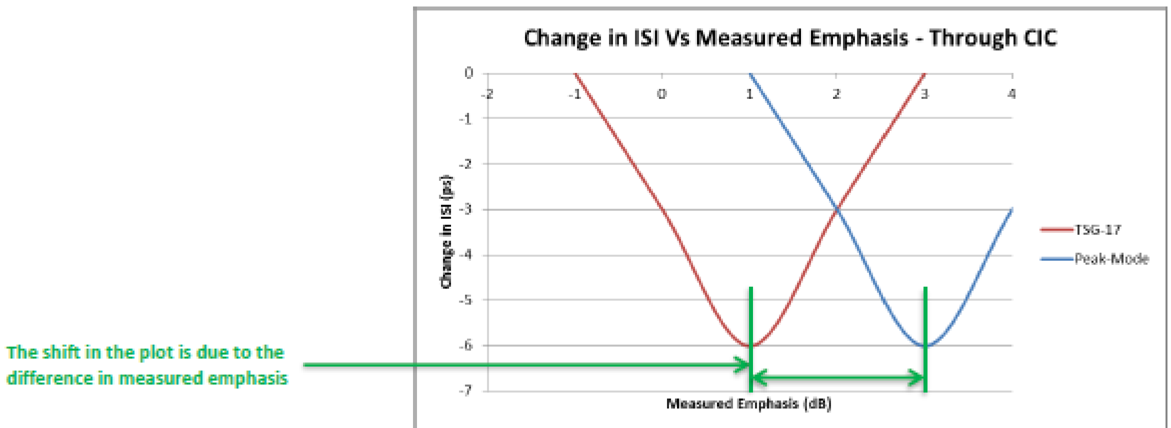

Measurements

Visual comparisons of Signals

- Signal shape comparison of the devices used in this testing.
- The Peak-Mode comparison shows the difference in devices when tuned to ~ 2.4 dB using peak-mode.
- The TSG-17 comparison shows the difference in devices when tuned to ~ 2.4 dB using TSG-17.

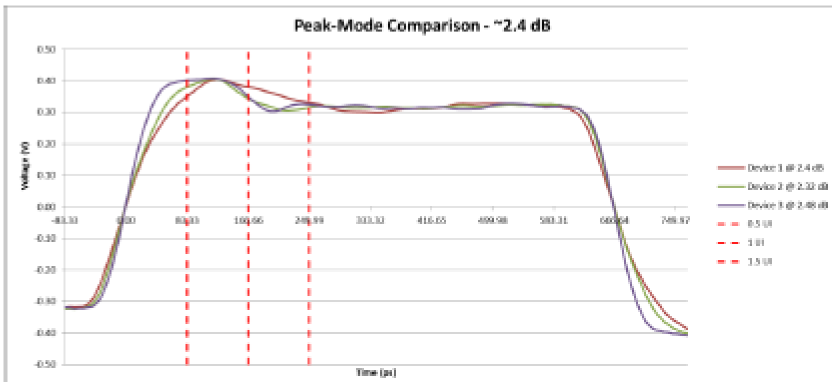

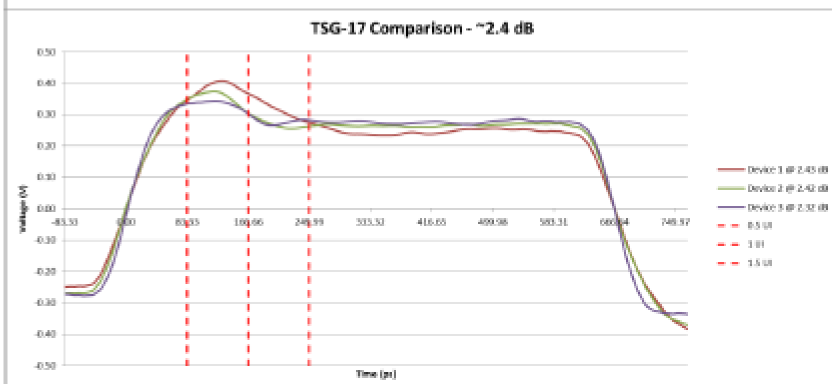

Device 1:
- 71 ps rise-time*

Device 2:
- 58 ps rise-time*

Device 3:
- 40.5 ps rise-time*

Slowest

Fastest

*20-80 rise-time measurements were taken with an HFTP pattern and normalized amplitude.

- These plots show the difference between the two emphasis measurement techniques when tuned to ~ 2.4 dB.
- The closer to "full-bit" the signal is (i.e., peak and shelf points not offset from respective measurement locations of TSG-17), the less difference seen between the two methods.

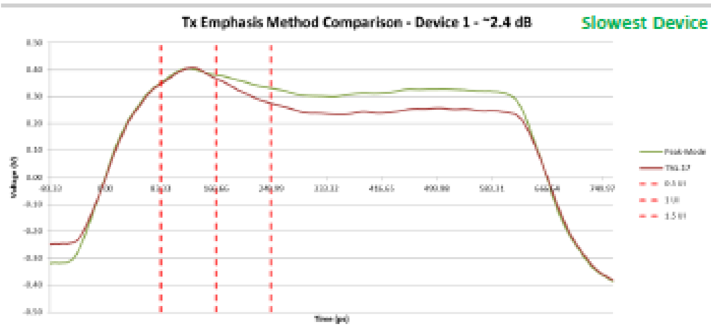

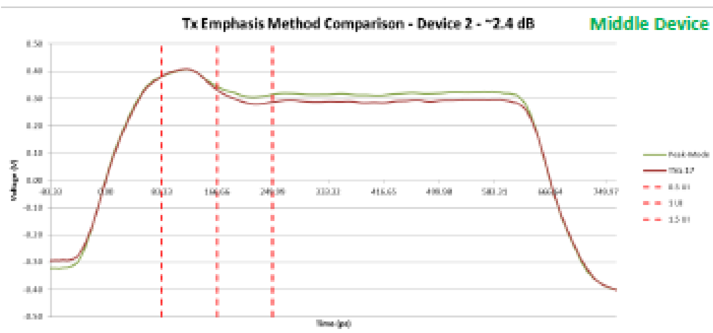

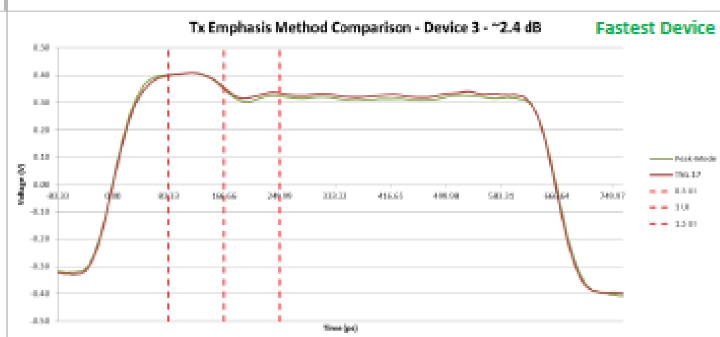

TX Emphasis versus Technique

- These plots show the difference between the two emphasis measurement techniques across device Tx emphasis register settings.
- Again, the closer to “full-bit” the signal is, the less difference seen between the two methods.

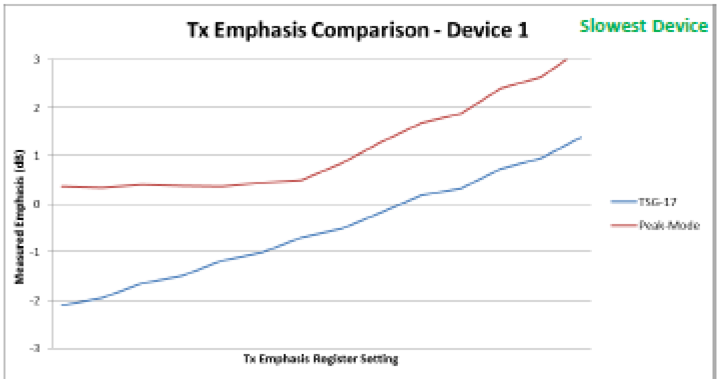

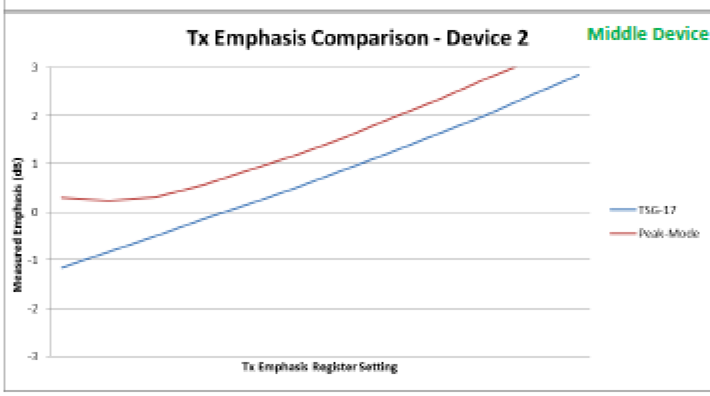

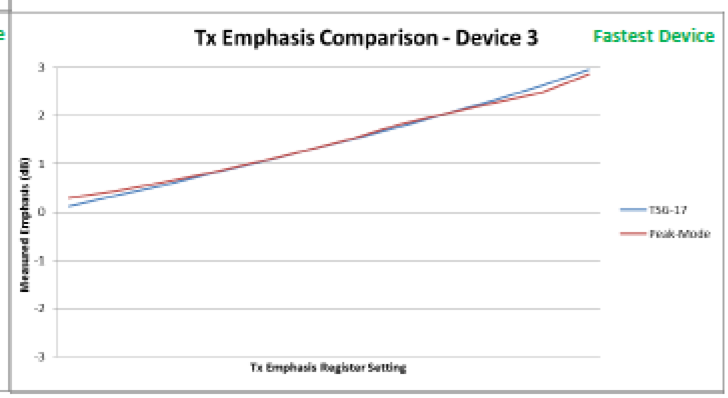

ISI versus TSG-17 Emphasis

- These graphs show the ISI vs. TSG-17 measured emphasis at the device connector and through a CIC.
- There is a significant difference in ISI performance between all three devices; even between Device 2 and Device 3.
- Device 1 is over-emphasizing the CIC with almost all TSG-17 emphasis settings within the compliant range.
- Device 1 is penalized at the connector (e.g., low loss channel applications) due to significant over-emphasis.
- Emphasis spread* between devices
    - 3 dB spread from Device 1 to Device 3
    - 1.3 dB spread from Device 2 to Device 3

*** Spread is determined by comparing the points where the ISI first hits the ISI floor through the CIC. The optimal setting is device dependent and to be determined by the device manufacturer.**

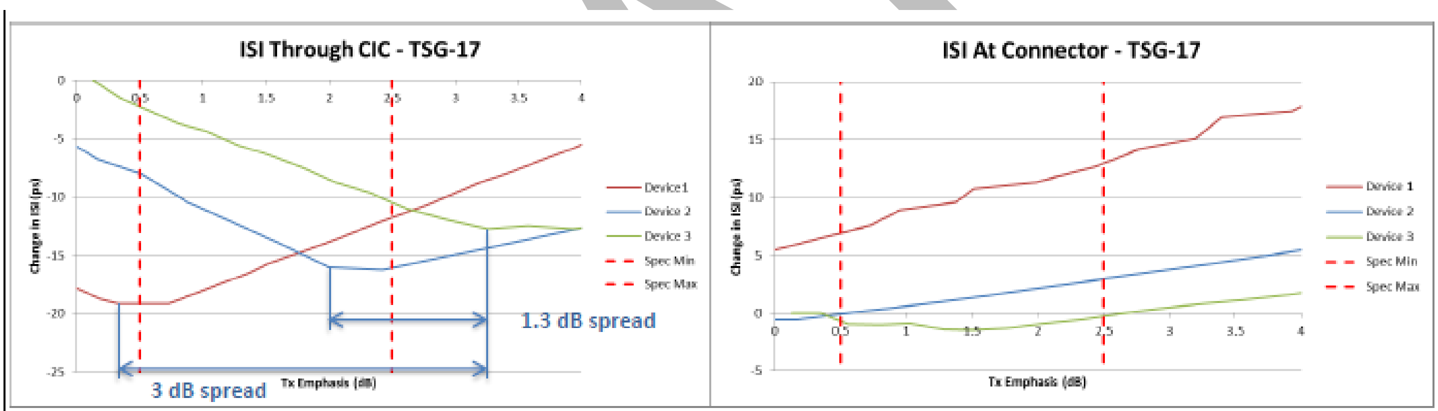

ISI versus Peak-Mode Emphasis

- These plots show the ISI vs. Peak-Mode measured emphasis at the device connector and through a CIC.
- The difference in ISI performance between the three devices is less when using the Peak-Mode technique.
- Devices 2 and 3 are very close.
- Device 1 is allowed more reasonable settings that don't over-emphasize the CIC channel and provides some relief at the connector.
- The Peak-Mode technique needs to maintain the lower limit of 0.5 dB defined for TSG-17, especially for devices with characteristics similar to Device 1.
- Emphasis spread* between devices
    - 1.2 dB spread from Device 1 to Device 3, down from 3 dB using TSG-17
    - 0.3 dB spread from Device 2 to Device 3, down from 1.3 dB using TSG-17

*** Spread is determined by comparing the points where the ISI first hits the ISI floor through the CIC. The optimal setting is device dependent and to be determined by the device manufacturer.**

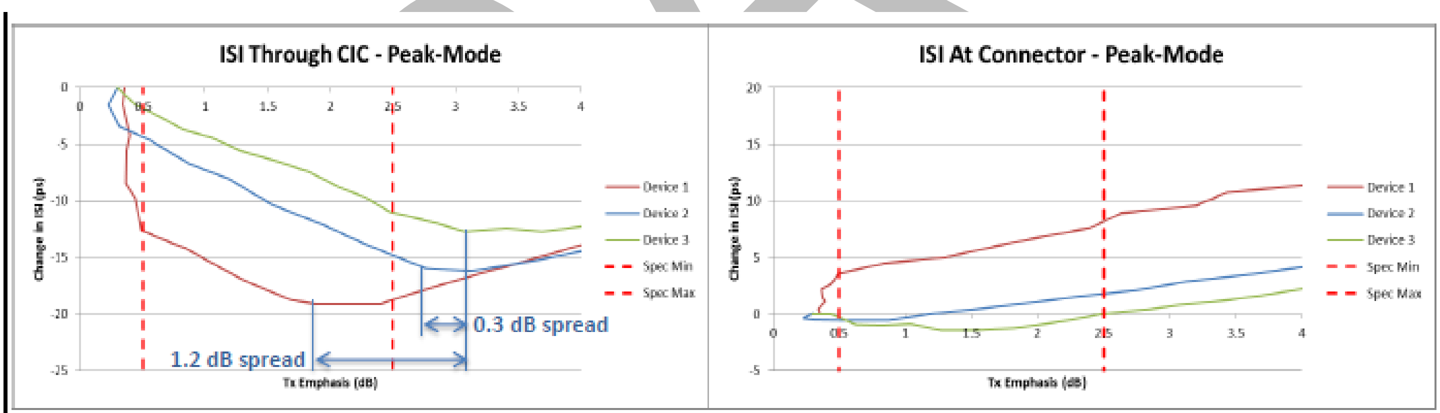

ISI versus Peak-Mode versus TSG-17 Emphasis

- These plots show the ISI vs. Peak-Mode emphasis and ISI vs. TSG-17 emphasis for each device.
- The Device 2 data supports setting the Peak-Mode emphasis limit to 3 dB in order to encompass the TSG-17 emphasis range.
- The Device 3 data also supports setting the Peak-Mode emphasis limit to 3 dB.

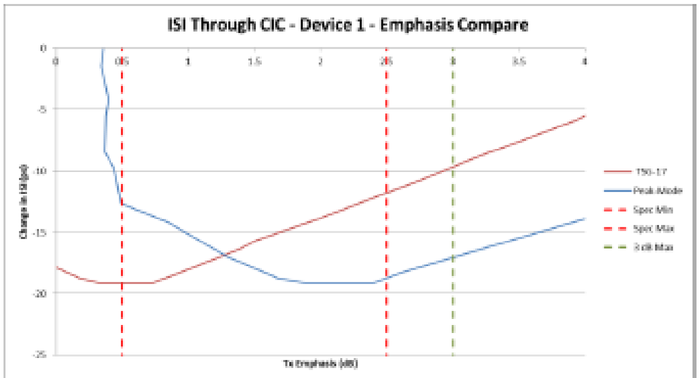

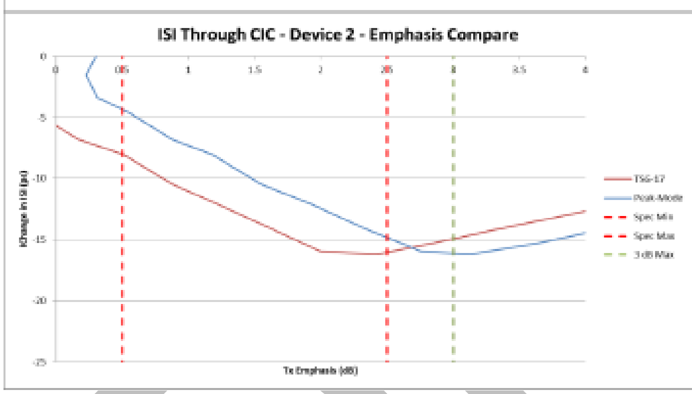

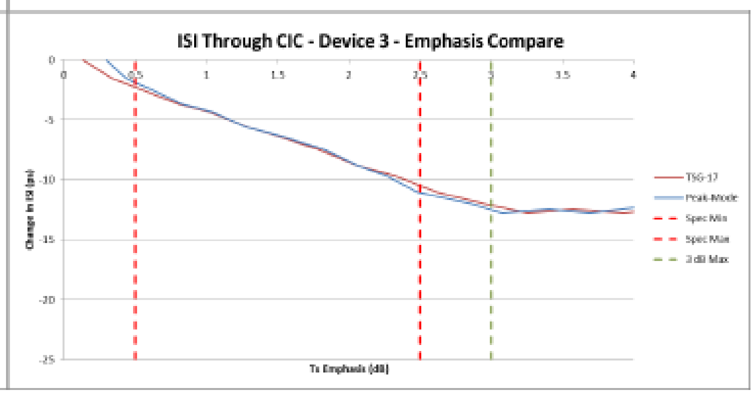

Conclusion

- The data supports having a Peak-Mode emphasis range of 0.5 dB to 3 dB.
  - The lower limit should stay at 0.5 dB to allow manufacturers to optimize emphasis for devices with characteristics similar to Device 1.
  - The upper limit needs to be set at 3 dB to support devices similar to Device 2 and Device 3.
- Considering that Peak-Mode better reflects effective emphasis with less spread across multiple implementations, the Peak-Mode method should be considered as the default option moving forward.
- This data illustrates the difficulty of dictating a nominal emphasis value across all devices. It's best left to the device manufacturer to pick the optimal emphasis value and the nominal value should not be stated for any device emphasis measurement method.