Transmitter Tuning Guide
for Untrained 25/50/100G NRZ Links

Application Note

Ethernet Products Group (EPG)

March 2024

Revision 1.1
647280-001

Revision History

Revision	Date	Comments
1.1[1]	March 22, 2024	Initial public release.

1. There are no previous publicly-available versions of this document.

Contents

NOTE: ***This page intentionally left blank.***

1.0 Introduction

Intel Ethernet network controllers and adapters support a wide range of link modes over different media interfaces. This document provides guidelines on transmitter tuning for Chip-to-Chip (C2C) and Chip-to-Module (C2M) links on Intel Ethernet network controllers supporting 25G-based single and multi-lane link modes.

The Physical Media Dependent (PMD) segment function supporting C2C and C2M links does not use link training (in contrast to backplane and direct-attach copper media, which do). Therefore, for link robustness, manual transmitter equalization adjustment is needed, which requires knowledge of the transmitter equalization parameter ranges and channel loss between transmitter output and receiver input pins.

Included in this application note are examples of common link topologies requiring transmitter equalization tuning along with a heuristic approach for rapid adjustment.

Note: Intel Ethernet adapters are not addressed in this document since they are supplied with port-specific transmitter setting from the Intel factory.

1.1 Audience

This document is meant for use by Intel internal teams and external customers under CNDA.

1.2 Product Coverage

This document covers the Intel products and link modes listed in Table 1. These guidelines are to be used for interfacing Intel Ethernet controllers with optical modules or co-located link partner ASIC either directly or through a retimer.

Table 1. Products and Link Modes

Ethernet Link Mode	Media	Intel Ethernet Device	Reference
25G AUI C2C	PC Trace	XXV710-AM1/AM2 E810-XXVAM2/CAM1/CAM2	IEEE Std. 802.3-2018 Annex 109A
25G AUI C2M	Fiber	C827-AM/C827-IM/XL827-AM FM10420/FM10840	IEEE Std. 802.3-2018 Annex 109B
50G LAUI-2 C2C	PC Trace	C827-AM/C827-IM/XL827-AM E810-XXVAM2/CAM1/CAM2	IEEE Std. 802.3-2018 Annex 135B
50G LAUI-2 C2M	Fiber	C827-AM/C827-IM/XL827-AM E810-XXVAM2/CAM1/CAM2	IEEE Std. 802.3-2018 Annex 135C
100G CAUI-4 C2C	PC Trace	C827-AM/C827-IM/XL827-AM E810-XXVAM2/CAM1/CAM2 FM10420/FM10840	IEEE Std. 802.3-2018 Annex 83D
100G CAUI-4 C2M	Fiber	C827-AM/C827-IM/XL827-AM E810-XXVAM2/CAM1/CAM2 FM10420/FM10840	IEEE Std. 802.3-2018 Annex 83E

1.3 Related Documentation

The documents listed in Table 2 can be accessed from the Intel Resource & Design Center at:

https://www.intel.com/content/www/us/en/design/resource-design-center.html

Table 2. Related Documents

Document Title	Document ID
Intel® Ethernet Controllers Product Guide	639793
HSS Physical Layer Conformance Tests Application Note	631745
Intel® 25G Ethernet Media Guide	630739

2.0 Common Topologies for Untrained Links

Optical links are achieved using C2M link modes, which do not use link training and constitute most of the applications requiring transmitter equalization adjustment. C2C links can be achieved using backplane link modes (KR, KR2, KR4), the Intel recommended option when both link partners are capable of supporting auto-negotiation and link training. However, when capabilities preclude the use of the appropriate backplane link mode, the C2C option without link training is the fallback.

The maximum channel insertion loss for a C2C link is roughly twice that of C2M link, but the same considerations apply for both regarding transmitter tuning. Common link topologies and associated specification limits for maximum channel insertion loss and transmitter amplitude are shown in Figure 1 and Table 3, respectively.

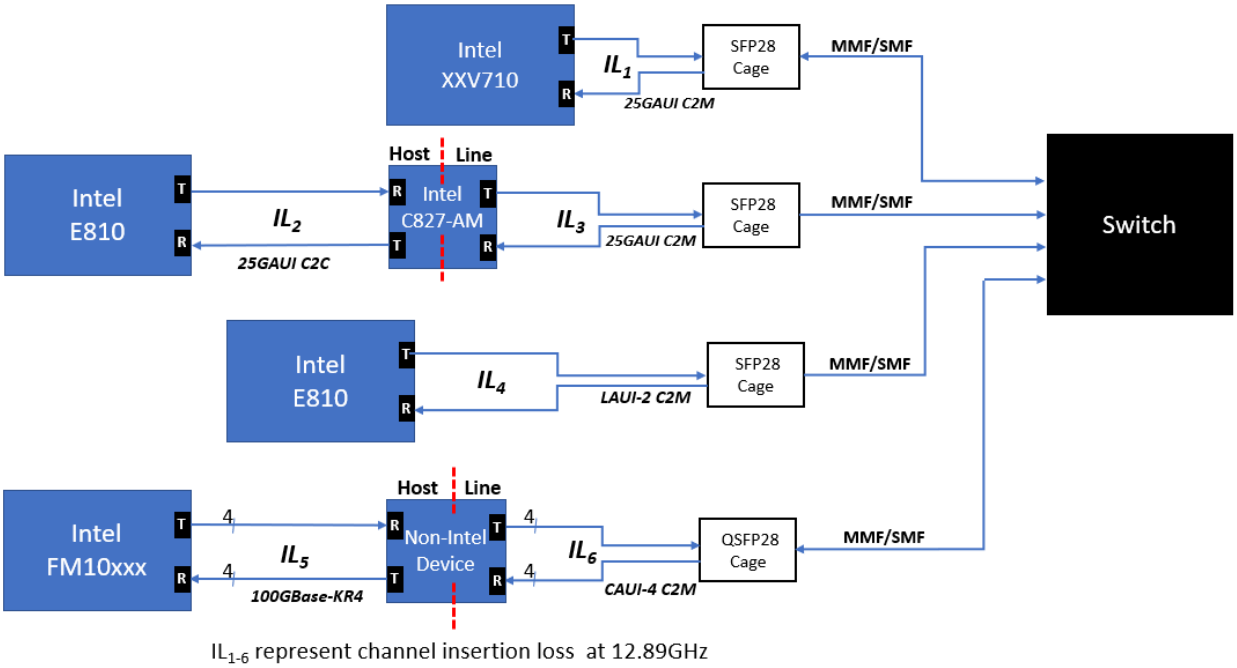

Figure 1. Link Topologies Requiring Transmitter Tuning

Table 3. Specification Limits for C2C and C2M Links

Link Segment	Insertion Loss	Maximum Amplitude	Maximum IL @ 12.89G
XXV710 to SFP28 Cage	IL1	900 mVp-p Diff.	10 dB
E810 to C827AM	IL2	1200 mVp-p Diff.	20 dB
C827-AM to SFP28 Cage	IL3	900 mVp-p Diff.	10 dB
E810 to QSFP28 Cage	IL4	900 mVp-p Diff.	10 dB
FM10xxx to 3rd-party Device	IL5	1200 mVp-p Diff.	35 dB[1]
3rd-party Device to QSFP28 Cage	IL6	1200 mVp-p Diff.	20 dB

1. 100GBASE-KR4 channel loss limit is listed here to include all segments shown in Figure 1.

3.0 Recommended Transmitter Tap Settings

Table 4 through Table 7 provide recommended transmitter settings differentiated by channel loss for each product listed in Table 3. These settings can be used as a good starting point for further adjustment when paired with link partner ASICs (C2C) or optical modules (C2M), if necessary. When adjusting, it is important to note that a transmitter equalization setting suitable for a stable link is often not unique.

Table 4. Intel Ethernet 700 Series Controllers

Applicable Link Modes	Channel Loss (dB) @ 12.89G	Recommended Transmitter Settings[1,2]		
		C-1	C0	C+1
25G AUI C2C/C2M	< 5	1	14	5
	> 5 and up to maximum loss	2	10	8

1. A negative sign is implied for C-1 and C+1 setting.
2. The sum {C-1+C0+C+1} must not exceed 32.

Table 5. Intel Ethernet 800 Series Controllers

Applicable Link Modes	Channel Loss (dB) @ 12.89G	Recommended Transmitter Settings[1,2]		
		C-1	C0	C+1
25G AUI C2C/C2M 50G LAUI-2 C2C/C2M 100G CAUI-4 C2C/C2M	< 5	0	4	0
	> 5 and up to maximum loss	4	2	18

1. A negative sign is implied for C-1 and C+1 setting.
2. The sum {C-1+C0+C+1} must not exceed 32.

Table 6. Intel Multi-Host Controllers

Applicable Link Modes	Channel Loss (dB) @ 12.89G	Recommended Transmitter Settings[1,2]		
		C-1	C0	C+1
25G AUI C2C/C2M 100G CAUI-4 C2M	< 5	1	14	5
	> 5 and up to maximum loss	4	2	18

1. A negative sign is implied for C-1 and C+1 setting.
2. The sum {C-1+C0+C+1} must not exceed 32.

Table 7. Intel Retimers

Applicable Link Modes	Channel Loss (dB) @ 12.89G	Recommended Transmitter Settings[1,2]		
		C-1	C0	C+1
25G AUI C2C/C2M 50G LAUI-2 C2C/C2M 100G CAUI-4 C2C/C2M	< 5	1	14	5
	> 5 and up to maximum loss	4	2	18

1. A negative sign is implied for C-1 and C+1 setting.
2. The sum {C-1+C0+C+1} must not exceed 32.

For additional guidelines on tuning based on channel insertion loss, refer to Appendix C, "Heuristic Method for Equalization Tuning".

A good practice for tuning is to use the following as targets for Eye Height and Eye Width:

- **Vertical:** 70% of maximum amplitude specification
- **Horizontal:** 0.7 UI Eye Width

Example: 25GAUI C2M (IEEE Std. 802.3-2018, Table 83E-1)

Maximum Amplitude specification: 900 mVp-p Diff.

Eye Width Specification: 0.46UI (minimum)

Targets:

Vertical: 70% maximum amplitude = 630 mVp-p Diff. (margin of 270 mVp-p Diff.)

Horizontal: 0.7 UI (margin of 0.24 UI)

4.0 Methodology and Tools for Transmitter Tuning

4.1 Tuning and Device Configuration

The following steps describe the flow for determination and automatic configuration of port-specific transmitter settings:

1. Customer identifies all ports that use either C2C or C2M links.
2. Customer tunes each port using guidelines in this document or preferred procedure. For a description of available common methods, refer to Appendix B, "Background on Transmitter Equalization".
3. Customer provides to Intel the set of transmitter settings determined for each port.
4. Intel generates an NVM image with transmitter settings provided by customer.
5. Customer uses the NVM Update Tool to program the EEPROM with the Intel-provided NVM image.
6. Customer verifies that the transmitter settings have taken effect.

Note: Step 2 and Step 6 can be performed with Intel® Ethernet Inspector (Intel® EI). Consult an Intel representative for tool documentation and assistance with installation and use with a specific operating system.

4.2 Transmitter Tuning Using Intel® Ethernet Inspector

The steps in the following sections show the locations of transmitter equalization settings in Intel® EI for each product listed in Table 1.

4.2.1 Intel® Ethernet 700 Series Controllers

1. On the **General** tab, uncheck the boxes for **Link Firmware Management** and **Admin Queue ON**.

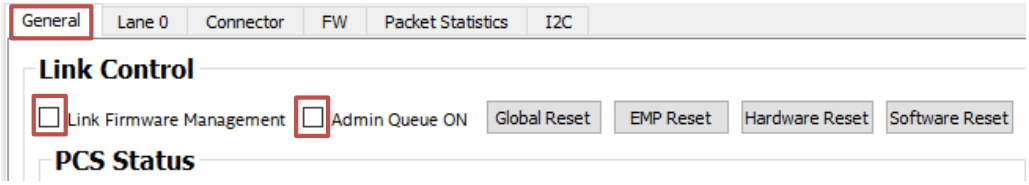

2. On the **Lane 0** tab, write (Force) the desired transmitter settings.

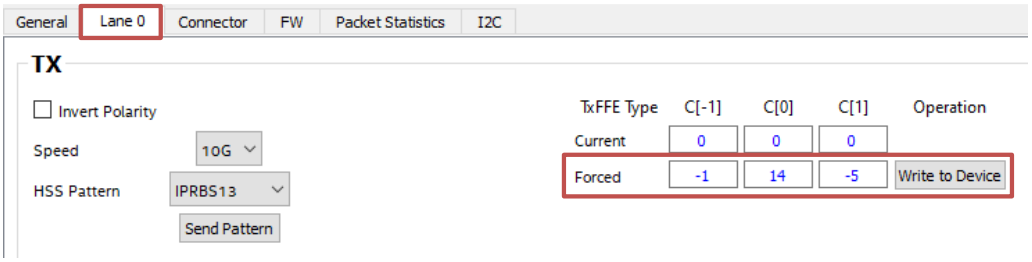

4.2.2 Intel® Ethernet 800 Series Controllers

1. On the **Lane 0** tab, click on the **Disable Link Management** button.
2. On the **Lane 0** tab, write (Force) the desired transmitter settings.

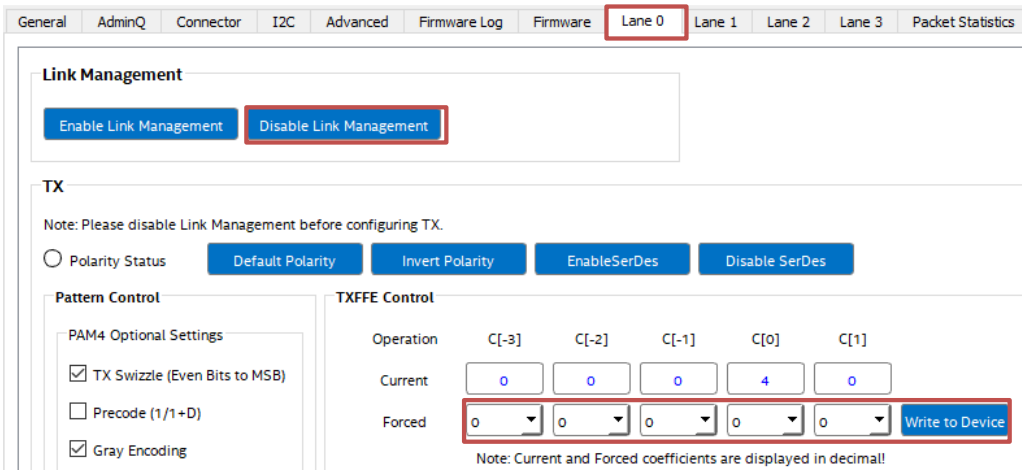

4.2.3 Intel Multi-Host Controllers

Intel Ethernet Multi-Host Controllers FM104xx and FM108xx are commonly paired with a retimer device as outermost PHY when interfaced to optical modules. the FM104xx and FM108xx are controlled via IES-API accessed through the TestPoint Command Line Interface (CLI). Please refer to the multi-host controller Datasheet for further information on accessing API functions through CLI.

Following is an example of CLI commands for controlling transmitter settings.

```
<0>% expert
<expert>%fmPlatformRetimerGetTxFFE(0,1,0,0,0)         /* Enter Expert mode */
Tx FFE values for lane 0:                             /* Get TxFFE value from Port 1, Retimer Lane 0 */
preCursor: 1
Cursor: 17
postCursor: 6
<expert>%fmPlatformRetimerSetTxFFE(0,1,0,0,0,0,4,10) /* Set TxFFE (0,4,10) to Port 1, Retimer Lane 0 */
<expert>%fmPlatformRetimerGetTxFFE(0,1,0,0,0)         /* Read new TxFFE value from Port 1, Retimer Lane 0 */
Tx FFE values for lane 0:                             /* Get TxFFE value from Port 1, Retimer Lane 0 */
preCursor: 0
Cursor: 4
postCursor: 10
<expert>% exit                                        /* Exit Expert mode */
```

4.2.4 Intel Retimers

Retimer devices C827-AM1/IM1 and XL827-AM1 are controlled via two-wire serial interface, MDIO, or I^2C, which are selectable. The registers for manipulating transmitter equalization settings are provided in Table 8. These retimer devices are split into Line and Host sides with separate registers for transmitter settings for either side. Please refer to the retimer Datasheet for further information on register access through two-wire interface and register description.

Table 8. Intel Retimer Device Registers for Modifying Transmitter Settings

Interface Side	Device ID	Port	Register Address		
			C-1	C0	C+1
Line	3	0	0xF217	0xF218	0xF219
		1	0xF297	0xF298	0xF299
		2	0xF318	0xF318	0xF219
		3	0xF397	0xF398	0xF399
Host	4	0	0xF217	0xF218	0xF219
		1	0xF297	0xF298	0xF299
		2	0xF318	0xF318	0xF219
		3	0xF397	0xF398	0xF399

Following is an example showing register writes required for applying transmitter setting C-1 = -1, C0 = 14, C+1 = 5 to line side transmitter:

Device 3, Register 0xF217 = 0x100 (C-1)

Device 3, Register 0xF218 = 0xE00 (C0)

Device 3, Register 0xF219 = 0x500 (C+1)

5.0 Summary

The following points are essential for ensuring link robustness when working with Intel Ethernet controllers:

1. Transmitter tuning is mandatory for applications using untrained link modes (C2C or C2M).
2. The result of tuning is the identification of a set of transmitter equalization settings that are port-specific, thus requiring each port to be tuned separately.
3. These settings are applied to production systems either by software control after boot to OS, or loaded from an EEPROM containing controller firmware and configuration settings during boot phase, where the latter is most prevalent and the Intel recommended method. Contact your Intel representative for guidance on tuning procedures and applying tuned settings during device configuration.

Appendix A Programmable Transmit Equalization Feature

IEEE Std. 802.3-2018 Clause 92.8.3.5 stipulates that a programmable equalization feature be implemented as part of the transmit function. A 3-tap FIR filter structure shown in Figure 2 is suggested for providing transmit equalization where the tap ranges, which can vary with implementation, are governed by waveform amplitude and step size specification limits.

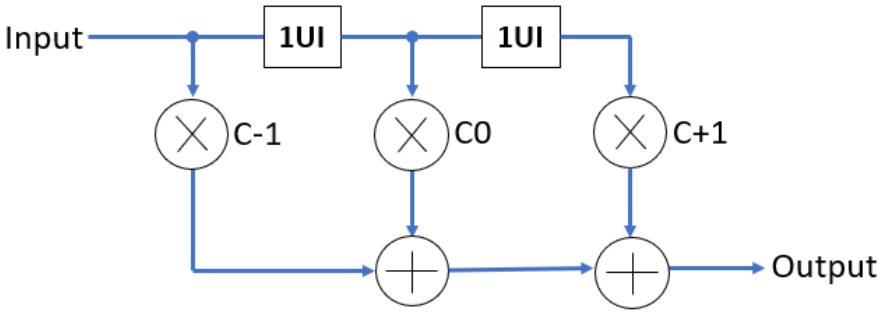

Figure 2. 3-Tap FIR Filter-Based Programmable Transmit Equalizer

The salient properties of the programmable transmit equalization function are:

- Transmitter equalization is fully characterized by waveform measurements at the following locations respectively:
  - **C2C (IEEE 802.3-Std. 2018 Annex 109A and 83D):** Inferred at the Transmit (TX) pin by applying a linear pulse fit technique defined in Cl. 93.8.1.5.1 on the waveform captured using a test fixture. Channel loss between TX pin and the test fixture must be between 1.2 and 1.6 dB at 12.89G.

    Note: Linear pulse fit technique is applicable for channels up to 10 dB inclusive of the test fixture (refer to Cl. 92.8.3.5 and Annex 92A). A longer channel less than 10 dB is more common, and hence a suitable alternative when a short channel between 1.2 to 1.6 dB is not available.
  - **C2M (IEEE Std. 802.3-2018 Annex 109B and 83E):** At the end of channel including the test fixture with total channel loss between TX pin and the test fixture less than 10 dB at 12.89G.
- The programmable taps C(-1), C(0), and C(+1) are referred to as the pre, main, and post-cursor taps, respectively
- The programmable tap coefficient ranges are implementation-specific and can take positive or negative values.
- A tap value of '0' for any tap indicates 'off' state for that tap.
- The equalizer, when set to {C(-1)=0, C(0)=maximum, C(+1)=0}, provides no equalization.
- The following ratio is a measure of pre-emphasis or pre-cursor equalization:

$$\frac{C(0)+(C1)-(C-1)}{C(-1)+C(0)+C(+1)}$$

- The following ratio is a measure of de-emphasis or post-cursor equalization.

$$\frac{C(0)-(C1)+(C-1)}{C(-1)+C(0)+C(+1)}$$

- The voltage shift associated with an increment or decrement of any tap value is constrained by step size parameter defined in Cl. 92.8.3.5.4.
- Repeated increments or decrements result in voltage shifts that are monotonic.

Appendix B Background on Transmitter Equalization

Channel equalization is a task performed jointly, not necessarily equally, by the local transmitter and the far-end receiver. Link robustness is guaranteed by the link partner receiver meeting required PMD BER threshold for each PMD type as specified in the standard. Link modes that do not employ link training require manual transmitter equalization adjustment to ensure the link partner receiver is able to meet the required BER target. Some amount of equalization from the transmitter is always preferred over none, with few exceptions. The following discussions compare common methods for characterizing transmitter-based equalization.

Method 1: Oscilloscope Based Waveform Measurement

Requirements:

- Signal access point (SFP/QSFP) and test fixture (HCB).
- Real-time or sampling scope with built-in CDR.
- Scope and test fixture calibration along with measurement automation.

Measured Parameters:

- Eye height and width, jitter, rise/fall time.

Advantages:

- Objective standards based true measurement of transmitter performance.
- Measurement is independent of link partner receiver capability.

Limitations:

- Not suitable for chip-to-chip links in production environment.
- Restricted to PRBS patterns only.

Method 2: Receiver Eye Measurement

Requirements:

- On-die FW based eye measurement capability and automation.

Measured Parameters:

- Eye height and width.

Advantages:

- Measurement is pattern independent: PRBS or scrambled idles and data.
- Measurement is suitable for chip-to-chip links.
- Mission mode “end-end” channel coverage.

Limitations:

- Transmitter equalization sufficiency is inferred.
- Dependent on receiver equalization capability hence tuning necessary for each link partner.
- On-die FW based eye measurement not always available (example: switch link partners).
- Subjective, not based on standards, since receiver eye height has no specification.
- Pessimistic deviation from true eye opening due to “bath-tub” extrapolation math.

Method 3: Receiver BER Measurement

Requirements:

- On-die pattern checker and automation.
- Transmitter error injection capability preferred but not necessary.

Measured Parameters:

- Bit error ratio.

Advantages:

- Measurement suitable for chip-to-chip links.
- Measurement is standards based.
- Mission mode "end-end" channel coverage.

Limitations:

- Transmitter equalization sufficiency is inferred.
- Dependent on receiver equalization capability hence tuning necessary for each link partner.
- Restricted to PRBS patterns only.
- Pattern checker not always available (example: switch link partners).
- Slow method owing to BER gating time requirement: > 2.5 minutes (25G), 5 minutes (10G).

The method of choice to characterize transmitter equalization quality depends on factors such as:

- Link type viz., chip-to-chip (C2C) or chip-to-module (C2M).
- Transmitter compliance testing and qualification against standards specifications.
- Testing robustness of transmitter configuration to cover temperature and channel loss limits.
- Decision on FEC use in mission mode (example: Consortium mode 25G link without FEC).
- Instrumentation and automation availability to perform the time-consuming task of tuning.

Regardless of the method adopted, the overarching issue in transmitter equalization tuning is the vast search space that is a consequence of the ranges for the three coefficients defining the 3-tap FIR filter.

This is easily seen with an example 3-tap FIR filter which has the following coefficient widths:

- Pre-cursor: 3 bits (8 settings)
- Main cursor: 6 bits (64 settings)
- Post- cursor: 4 bits (16 settings).

A brute-force search for optimal coefficient requires 8192 individual measurement cycles, which is prohibitive. A heuristic method to reduce the search space is described in Appendix C, "Heuristic Method for Equalization Tuning".

Appendix C Heuristic Method for Equalization Tuning

C.1 Optimal Starting Point for Tuning

Search space reduction is central to transmitter equalization tuning and is best achieved through the selection of an optimal starting point for the search. This document recommends the use of Rpre and Rpst equalization settings, applied to the transmit function upon entering the "INITIALIZE" state of the link training state machine, as the starting point of tuning exercise regardless of channel loss.

Even though Rpre and Rpst are settings associated with link training, the basis for utilizing Rpre and Rpst as search starting point is that its intent is to provide some transmitter assistance to an untrained receiver, adapting its elements (VGA, CTLE, DFE etc.) in the early phase of link training. The values Rpre=1.29+/-10% and Rpst=2.57+/-10%, stipulated in IEEE Std. 802.3-2018 Cl. 72 (10G) and 110 (25G), provide 2.3 dB of pre-cursor and 8.2 dB of post-cursor equalization respectively.

Although the clauses indicated are specific to link modes that use training, the 3-tap FIR filter based programmable transmit equalization feature applies to all untrained modes. For details on 3-tap FIR filter characteristics, refer to Appendix A, "Programmable Transmit Equalization Feature".

Rpre and Rpst are ratios, and hence independent of coefficient tap range and waveform amplitude achieved by a specific PHY transmit function implementation. Multi-rate transmitters have fast edge rates (rise/fall time) optimized for the highest rate, which must be considered when adjusting transmitter equalization for lower rates supported. The use of Rpre and Rpst to set the direction of transmitter equalization adjustment for different channel loss profiles is illustrated in Figure 3.

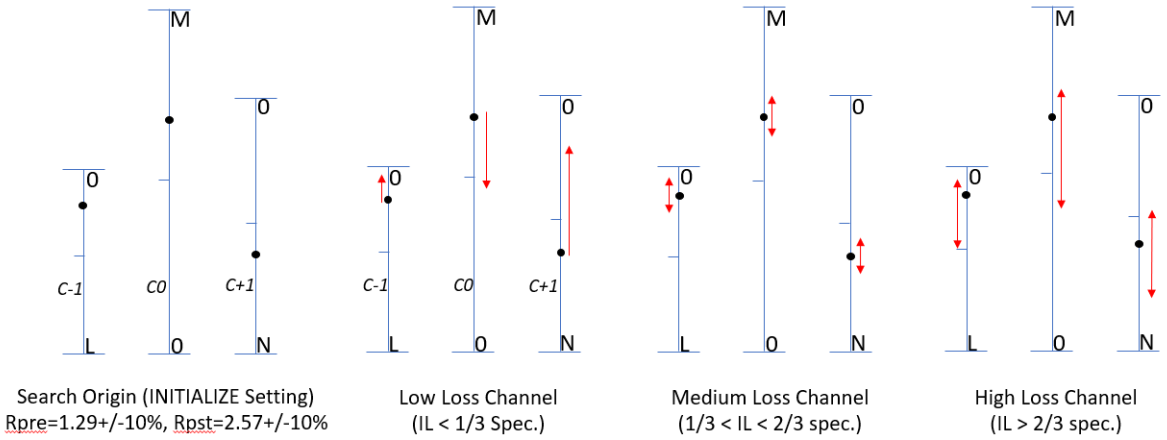

Figure 3. Recommended Tuning Strategy with INITIALIZE Setting as Search Starting Point

The coefficient range limits L and N are negative integers, while M is a positive integer. The following are general observations guiding the tuning directions depicted in Figure 3:

- **Low Loss Channels:** The main impairment is reflection, which is best mitigated by reducing incident energy (amplitude). This is particularly important for channels that terminate in SFP/QSFP cages. Also, since channel loss is low, there is less need for equalization through pre-cursor and post-cursor taps. Thus, the recommended tuning direction for all three taps is toward reducing their individual contributions.

- **Medium Loss Channels:** The impairment from this channel is a combination of loss and reflection, which is best mitigated by some amount of equalization and amplitude adjustment. These channels typically offer receivers the least amount of stress, which translates to a wide range of transmit equalization settings suitable for a receiver to meet its target BER threshold. The recommendation here is for small adjustments around the INITIALIZE setting if required at all.
- **High Loss Channels:** The main impairment is loss, which is mitigated by boost from pre-cursor and post-cursor equalization. Impact of reflections are auto-regulated due to attenuation from high channel loss. These channels offer maximum stress to receivers, which results in a small set of transmit equalization settings suitable for a receiver to meet its target BER threshold. The recommendation here is for boosting equalization while optimizing amplitude.

C.2 Common Transmitter Equalization Rules and Considerations

Transmitter 3-tap FIR filter coefficient widths or ranges are implementation specific, but certain rules that apply for adjustment are common to all implementations. For parameter definitions, refer to Appendix A, "Programmable Transmit Equalization Feature".

Some common and useful rules to remember are:

- **Coefficient Signs:** Transmitter coefficient sign can be positive or negative, but in general, pre-cursor (C-1) and post-cursor (C+1) take negative sign, while the main cursor is always positive. The meaning of the signs is best understood by considering the summation of currents occurring at the output of the 3-tap FIR structure. A positive sign for a cursor implies current being sourced into the summation node from that cursor, while negative sign drains current. It is useful to visualize the main cursor as the path that sources (add) current into the summation node, and the pre-cursor and post-cursor as the paths that drain (subtract) current out of the summation node of the 3-tap FIR filter, illustrated in Figure 4 by red arrows signifying direction of current flow. The higher the current at the summation node, the higher the signal amplitude at the input of the receiver that is connected at the other end of the transmission line channel. Increasing contributions from pre-cursor and post-cursor results in amplitude reduction, making it a critical trade-off for handling high loss channels.

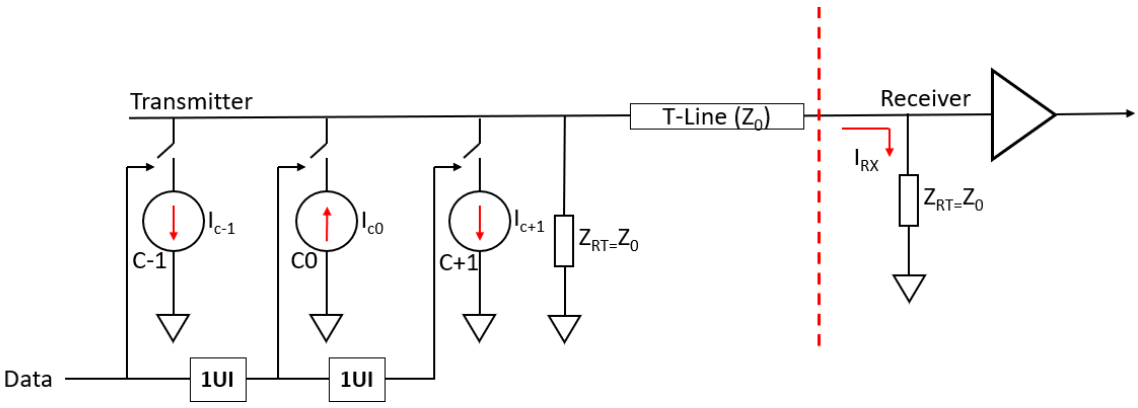

Single-ended circuit representation
PRESET Setting (max. Amplitude): $I_{RX} = I_0 = I_{C0}$ (max)
All other transmitter settings: $I_{RX} = I_{C0} - I_{C-1} - I_{C+1}$

Figure 4. Visualizing Single-Ended Current Flow and Summation in a 3-tap FIR Filter

- **Transmitter Frequency Response:** The frequency response of the 3-tap FIR filter is variable where both low-pass and high-pass responses can be realized. A good practice to interpret a given set of coefficients is to look at the coefficient sign. Alternating signs realize high-pass response, while same signs lead to low-pass. Apart from the case of an extremely short, low loss channel a high-pass response is always desired, as it compensates for low-pass (attenuating) characteristic of the channel.
- **Over Equalization:** Most transmitters operate by constraining the sum of the taps, which must be strictly adhered to avoid signal inversion. Signal inversion can be interpreted as switching from intended high-pass to low-pass response, thereby adding stress to the signal reaching the input of the receiver.
- **Inter-symbol Interference (ISI):** Inter-symbol interference (ISI) results from the combined effect of the random nature of the bit stream and low-pass frequency characteristics of the channel. The impact of ISI is best understood in time-domain by looking at the impulse response at various points along the channel and most importantly at the decision node (slicer) in the receiver where pre-cursor and post-cursor ISI components can be quantified with reference to the receiver sampling point. In general, post-cursor ISI is the dominant component at the receiver, but it is worth noting that the pre-cursor ISI component is elevated and non-negligible for 25 Gbps links. Thus, it is recommended to have some amount of equalization from the pre-cursor taps when dealing with 25 Gbps links, particularly for medium and high loss channels, to arrest the dispersion of impulse response seen by the receiver. This requirement is automatically satisfied by using INITIALIZE setting as a starting point.

Appendix D Glossary and Acronyms

Table 9. Definition of Terms

Term	Definition
ASIC	Application Specific Integrated Circuit
BER	Bit Error Rate
C2C	Chip-to-Chip
C2M	Chip-to-Module
CDR	Clock and Data Recovery
CLI	Command Line Interface
CTLE	Continuous-Time Linear Equalizer
DFE	Decision Feedback Equalizer
EEPROM	Electrically Erasable Programmable Memory
FEC	Forward Error Correction
FFE	Feed-Forward Equalizer
FIR	Finite Impulse Response
FW	Firmware
HCB	Host Compliance Board
IL	Insertion Loss
I^2C	Inter-Integrated Circuit
KR	K-Backplane, R-Scrambled
MDIO	Medium Dependent Input/Output
NVM	Non-Volatile Memory
PCT	Printed Circuit Trace
PMD	Physical Media Dependent
PRBS	Pseudo Random Binary Sequence
Rpre/Rpst	Equalization Ratio for Pre-cursor and Post-cursor
SFP/QSFP	Small Form-factor Pluggable (Q-Quad)
Tx/Rx	Transmit/Receive
VGA	Variable Gain Amplifier
25G AUI/LAUI/CAUI	25G Attachment Unit Interface (L-50, C-100)

LEGAL