

DOCSIS® Best Practices and Guidelines

Proactive Network Maintenance Using Pre-equalization

CM-GL-PNMP-V02-110623

RELEASED

Notice

This DOCSIS® guidelines document is the result of a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. for the benefit of the cable industry and its customers. This document may contain references to other documents not owned or controlled by CableLabs®. Use and understanding of this document may require access to such other documents. Designing, manufacturing, distributing, using, selling, or servicing products, or providing services, based on this document may require intellectual property licenses from third parties for technology referenced in this document.

Neither CableLabs nor any member company is responsible to any party for any liability of any nature whatsoever resulting from or arising out of use or reliance upon this document, or any document referenced herein. This document is furnished on an "AS IS" basis and neither CableLabs nor its members provides any representation or warranty, express or implied, regarding the accuracy, completeness, noninfringement, or fitness for a particular purpose of this document, or any document referenced herein.

© 2010-2011 Cable Television Laboratories, Inc.
All rights reserved.

Document Status Sheet

Document Control Number: CM-GL-PNMP-V02-110623				
Document Title: Proactive Network Maintenance Using Pre-equalization				
Revision History: V01 – Released 4/15/10				
V02 – Released 6/23/11				
Date: June 23, 2011				
Status:	Work in Progress	Draft	Released	Closed
Distribution Restrictions:	Author Only	CL/Member	CL/Member/Vendor	Public

Key to Document Status Codes

- Work in Progress** An incomplete document, designed to guide discussion and generate feedback that may include several alternative requirements for consideration.
- Draft** A document in specification format considered largely complete, but lacking review by Members and vendors. Drafts are susceptible to substantial change during the review process.
- Issued** A stable document, which has undergone rigorous member and vendor review and is suitable for product design and development, cross-vendor interoperability, and certification testing. Changes to Issued test plans are made through engineering change (EC) requests.
- Closed** A static document, reviewed, tested, validated, and closed to further engineering changes to the specification through CableLabs.

Trademarks

CableCARD™, CableHome®, CableLabs®, CableNET®, CableOffice™, CablePC™, DCAST™, DOCSIS®, DPoE™, EBIF™, eDOCSIS™, EuroDOCSIS™, EuroPacketCable™, Go2BroadbandSM, M-Card™, M-CMTS™, OCAP™, OpenCable™, PacketCable™, PCMM™, and tru2way® are marks of Cable Television Laboratories, Inc. All other marks are the property of their respective owners.

Contents

1	SCOPE	1
1.1	Introduction and Purpose.....	1
2	REFERENCES	2
2.1	Informative References.....	2
2.2	Reference Acquisition	2
3	TERMS AND DEFINITIONS	3
4	ABBREVIATIONS AND ACRONYMS	9
5	BACKGROUND	11
5.1	Reactive versus Proactive Network Maintenance	11
5.2	Linear Impairments.....	11
5.2.1	<i>Micro-reflection Types</i>	12
5.3	Pre-equalization Mechanism Enabled through DOCSIS Ranging	15
5.3.1	<i>Pre-equalization Enabling Messages</i>	17
5.3.2	<i>CM and CMTS Equalization Information</i>	19
5.4	Upstream Pre-equalization in DOCSIS 1.0, DOCSIS 1.1 and DOCSIS 2.0	21
5.4.1	<i>DOCSIS 1.1 Pre-equalization Considerations</i>	21
5.5	Limitations on Pre-equalization Compensation.....	22
5.6	DOCSIS Pre-equalization MIBs.....	23
5.6.1	<i>DOCSIS 2.0 and 3.0 Pre-equalization MIBs</i>	24
6	METHODOLOGY	25
6.1	General Approach and Processes	25
6.2	Format Verification, Normalization and Guidelines	26
6.2.1	<i>Four Nibble 2's Complement Pre-equalization Coefficient Representation</i>	27
6.2.2	<i>Three Nibble 2's Complement Pre-equalization Coefficient Representation</i>	27
6.2.3	<i>Universal Decoding</i>	27
6.3	Key Metrics	28
6.3.1	<i>Adaptive Equalizer Main Tap Energy</i>	28
6.3.2	<i>Main Tap Nominal Energy and Main Tap Nominal Amplitude</i>	28
6.3.3	<i>Pre-Main Tap Energy</i>	28
6.3.4	<i>Post-Main Tap Energy</i>	28
6.3.5	<i>Total Tap Energy</i>	28
6.3.6	<i>Main Tap Compression</i>	29
6.3.7	<i>Non-Main Tap to Total Energy Ratio (Distortion Metric)</i>	29
6.3.8	<i>Pre-Main Tap to Total Energy Ratio</i>	29
6.3.9	<i>Post-Main Tap to Total Energy Ratio</i>	30
6.3.10	<i>Pre-Post Energy Symmetry Ratio</i>	30
6.3.11	<i>Group Delay Distortion</i>	30
6.4	DOCSIS Pre-equalization Coefficient Data Collection.....	36
6.4.1	<i>SNMP Implementation and Performance Considerations</i>	36
6.4.2	<i>Data Collection Strategy</i>	37
6.5	Calibration Mechanisms	38
6.5.1	<i>CMTS-CM Short Reference Plant</i>	38
6.5.2	<i>Pre-equalization Calibration Approach</i>	40
6.6	Fault Localization.....	44
6.6.1	<i>Fault Localization Examples</i>	45
6.6.2	<i>Determining Micro-reflection Signatures</i>	45

6.6.3	<i>Determining Micro-reflection Boundaries Edges</i>	48
6.7	Severity Assessment	49
6.7.1	<i>Initial CM Selection for Analysis</i>	50
6.7.2	<i>Severity Analysis Strategy for Static or Single Data Point Scenario</i>	50
6.7.3	<i>Severity Analysis Strategy for Trending</i>	51
6.7.4	<i>Severity Analysis Strategy for Intermittent Issues</i>	52
6.8	Use Case Examples for Adaptive Equalizer Coefficients.....	53
6.8.1	<i>Use Cases</i>	54
6.8.2	<i>Use Case: Upstream Ingress or Noise Detection</i>	57
6.8.3	<i>Use Cases 3 -11</i>	61
6.9	Post-equalization	63
6.9.1	<i>Advantages of Pre-equalization</i>	63
6.9.2	<i>Disadvantages of Pre-equalization</i>	64
6.9.3	<i>Advantages of Post-equalization</i>	65
6.9.4	<i>Disadvantages of Post-equalization</i>	65
6.9.5	<i>Pre- and Post-equalization Measurements: Difference in Performance</i>	66
7	INTEGRATION OF DOCSIS NETWORK HEALTH METRICS WITH PRE-EQUALIZATION METRICS AND PROACTIVE NETWORK MAINTENANCE APPROACH	68
7.1	DOCSIS Robustness Against Impulse and Noise.....	68
7.1.1	<i>Impulse Noise Statistics</i>	70
7.1.2	<i>Impulse and Burst Robustness for the Different DOCSIS Modes of Operation</i>	71
8	CONCLUSION	77
APPENDIX I	TUTORIAL	78
I.1	Nonlinear Distortions	78
I.2	Linear Distortions	79
I.3	Frequency Response	79
I.4	Group Delay	79
I.5	Impedance Mismatch.....	84
I.6	Amplitude Ripple	87
I.7	Amplitude Tilt	91
I.8	Modulation Error Ratio	92
I.9	RxMER Measurement in a Digital Receiver	95
I.10	Adaptive Equalization	98
I.11	Adaptive Pre-equalization	110
I.12	Velocity of Propagation.....	111
APPENDIX II	SNMP COLLECTION METHODOLOGY	113
APPENDIX III	SNMP MIB REFERENCE	114
APPENDIX IV	MICRO-REFLECTION CALCULATOR	115
APPENDIX V	TWO TYPES OF ECHOES	116
APPENDIX VI	DOCSIS PRE-EQUALIZER COEFFICIENTS ANALYSIS - SOFTWARE SEQUENCE DIAGRAMS	123
VI.1	Software Sequence Diagram SD-PNM200	123
VI.2	Software Sequence Diagram SD-PNM201	123
APPENDIX VII	ACKNOWLEDGEMENTS	124

Figures

Figure 1 - Micro-reflection with multiple-transit echoes.....	13
Figure 2 - Micro-reflection with single impedance mismatch interface.....	14
Figure 3 - Composite Micro-reflection resulting from Type 1 and Type 2 Micro-reflections.....	15
Figure 4 - Upstream Equalizer Structure.....	16
Figure 5 - CM-CMTS Ranging Interaction Enabling Pre-equalization Process.....	17
Figure 6 - Range Response Message Format.....	18
Figure 7 - Range Request Message Format.....	19
Figure 8 - CM Pre-eq Coefficients Values and Frequency Response Scenarios.....	20
Figure 9 - CMTS CM Pre-equalization Coefficients Values and Frequency Response Scenarios.....	21
Figure 10 - Pre-equalization Compensation Capabilities under Short and Long Delay Micro-reflection Scenarios.....	23
Figure 11 - MIB Format for docsIfCmStatusEqualizationdata.....	24
Figure 12 - Proactive Network Maintenance Processes based on Pre-equalization.....	25
Figure 13 - Equalizer Structure HEXDECIMAL-to-DECIMAL Conversion.....	26
Figure 14 - Group Delay increase with increasing cascade depth.....	31
Figure 15 - Pre Main Tap Energy Increase with Cascade Depth (Fc=40.4 MHz, Ch. W=3.2 MHz, No Micro-reflections, first 12 taps shown).....	32
Figure 16 - Pre Main Tap Energy Increase with Cascade Depth (Fc=40.4 MHz, Ch. W=3.2 MHz, with 0.5 μ s Micro-reflection, first 12 taps shown).....	34
Figure 17 - Tap Energy for Different Cascade Depth Scenarios (Fc=14 MHz, Ch. W=3.2 MHz, with 0.5 μ s Micro-reflection, first 12 taps shown).....	35
Figure 18 - Short reference plant block diagram.....	40
Figure 19 - CM & CMTS Elements Contributing To US Distortion (In Orange).....	41
Figure 20 - Pre-equalizer frequency response with (0.5 μ s, -10 dBc) and without micro-reflection.....	43
Figure 21 - Calibrated pre-equalizer frequency response obtained from micro-reflection on (0.5 μ s, -10 dBc) and off scenarios.....	43
Figure 22 - Correlation of Topology with Distortion to Provide Fault Localization.....	45
Figure 23 - Observation of multiple CMs frequency response.....	46
Figure 24 - Identified Micro-reflection Patterns.....	47
Figure 25 - Clustering of Common Micro-reflection Signatures.....	47
Figure 26 - Common Micro-reflection Signature - Case A.....	48
Figure 27 - Common Micro-reflection Signature - Case B.....	49
Figure 28 - Severity Classification Mechanisms.....	50
Figure 29 - Micro-reflection amplitude data of same two CMs to highlight trending over time.....	51
Figure 30 - CM micro-reflection amplitude over time highlighting intermittent issues.....	53
Figure 31 - Amplitude vs. Frequency peak/valley of 10.55 dB with echo present in impulse response.....	57
Figure 32 - Entire upstream scan shows no similar signatures shared by other modems.....	57
Figure 33 - Distance calculation applied with customer address and mapping.....	57
Figure 34 - Single customer modem demonstrates the effects of ingress.....	60
Figure 35 - Multiple modems on the same upstream demonstrate the effects of ingress.....	60
Figure 36 - Effects of ingress when mapped.....	61
Figure 37 - Time domain representation of typical impulse.....	69
Figure 38 - Frequency and time capture of typical impulse.....	70

Figure 39 - Impulse duration characteristics71

Figure 40 - Impulse power level distribution characteristics.....71

Figure 41 - Burst/impulse noise lasting N symbol periods requires burst protection of N+1 symbols if the burst crosses symbol boundaries72

Figure 42 - Byte Interleaver Operation.....73

Figure 43 - Burst robustness (in red) and added latency (in blue) for a 64 byte packet using different FEC and interleaver block size settings.....74

Figure 44 - Burst robustness (in red) and added latency (in blue) for a 256 byte packet using different FEC and interleaver block size settings.....74

Figure 45 - Burst robustness (in red) and added latency (in blue) for a 1518 byte packet using different FEC and interleaver block size settings.....74

Figure 46 - Burst impact on S-CDMA subframe.....75

Figure 47 - Burst robustness in microseconds for 16-QAM, 3.2 MHz bandwidth channel with TCM off and diverse FEC and subframe settings (packet occupies at least one subframe)76

Figure 48 - Burst robustness in microseconds for 64-QAM, 6.4 MHz bandwidth channel with TCM off and diverse FEC and subframe settings (packet occupies at least one subframe)76

Figure I-1 - A filter's time delay-versus-frequency curve often has a bathtub shape79

Figure I-2 - Phase-versus-frequency for 100 feet of coax81

Figure I-3 - Time delay-versus-frequency for 100 feet of coax82

Figure I-4 - Complex frequency response in the return path.....83

Figure I-5 - Ideal transmission line model.....84

Figure I-6 - Real-world transmission line model.....85

Figure I-7 - Creation of reflections in a cable network's feeder plant86

Figure I-8 - Graphic representation of incident signal and second reflection86

Figure I-9 - Reflection example that will be used to illustrate the formation of amplitude ripple87

Figure I-10 - Phasor view of incident signal vector (long arrow) and reflection vector (short arrow)88

Figure I-11 - Reflection vector rotated 45 degrees from original position88

Figure I-12 - Reflection vector rotated 90 degrees from original position88

Figure I-13 - Reflection vector rotated 135 degrees from original position88

Figure I-14 - Reflection vector rotated 180 degrees from original position89

Figure I-15 - Reflection vector rotated 225 degrees from original position89

Figure I-16 - Reflection vector rotated 270 degrees from original position89

Figure I-17 - Reflection vector rotated 315 degrees from original position89

Figure I-18 - Reflection vector back at original position after rotating 360 degrees.....90

Figure I-19 - Amplitude-versus-phase plot of phasor view vector sum vectors.....90

Figure I-20 - Example of flat amplitude-versus-frequency response91

Figure I-21 - Example of tilted or sloped amplitude-versus-frequency response.....91

Figure I-22 - Example of upstream 64-QAM signal with substantial in-channel tilt91

Figure I-23 - Example of 64-QAM signal after adaptive pre-equalization eliminated most of the in-channel tilt.....92

Figure I-24 - A signal carried in the sloped portion of the widely spaced amplitude ripple will exhibit in-channel tilt92

Figure I-25 - 64-QAM signal with good (36.3 dB) MER92

Figure I-26 - 64-QAM signal with poor (23.2 dB) MER93

Figure I-27 - 16-QAM constellation showing target symbol, transmitted (or received) symbol, and modulation error vectors.....94

Figure I-28 - MER is the ratio of average symbol power to average error power.....	94
Figure I-29 - QAM receiver block diagram.....	96
Figure I-30 - Each vector has a real (in-phase or I) and imaginary (quadrature or Q) component	97
Figure I-31 - Unequalized 64-QAM constellation	99
Figure I-32 - Equalized 64-QAM constellation.....	100
Figure I-33 - Generic 4-tap adaptive equalizer.....	102
Figure I-34 - Amplitude-versus-time plot of an incident signal and micro-reflection.....	103
Figure I-35 - Amplitude-versus-frequency response.....	103
Figure I-36 - Phase-versus-frequency response.....	104
Figure I-37 - Required magnitude- and phase-versus-frequency response to cancel echo.....	104
Figure I-38 - Adaptive equalizer that will be used in the example in the text.....	105
Figure I-39 - Operation of the adaptive equalizer's first tap.....	106
Figure I-40 - Operation of the adaptive equalizer's second tap	106
Figure I-41 - Summing the outputs of the adaptive equalizer's first and second taps	107
Figure I-42 - Operation of the adaptive equalizer's third tap.....	107
Figure I-43 - Summing the output of the adaptive equalizer's third tap with the previously summed first and second taps.....	108
Figure I-44 - Operation of the adaptive equalizer's fourth tap.....	108
Figure I-45 - Final summing process provides an equalized output.....	109
Figure I-46 - Final amplitude and phase-versus-frequency response after adaptive equalization.....	109
Figure I-47 - Upstream pre-equalization is able to compensate for in-channel tilt	110
Figure IV-1 - Equivalent Reflection Coefficient Γ_E	115
Figure V-1 - A Multiple Recursion Echo	117
Figure V-2 - Wring Diagrams To Make Echoes In a Lab	118
Figure V-3 - A Single Recursion Echo Example.....	119
Figure V-4 - How A Multiple Recursion Echo Can be Canceled with Pre-distortion	120
Figure V-5 - Comparison of Signal Path Impulse Responses and Programming for Adaptive Equalizers.....	121
Figure VI-1 - Software Sequence Diagram SD-PNM200	123
Figure VI-2 - Software Sequence Diagram SD-PNM201	123

Tables

Table 1 - Maximum Delays Generated by Pre-equalization Filter Structures in DOCSIS 1.1 and 2.0.....	22
Table 2 - DOCSIS 2.0 and 3.0 Transmit Pre-equalization MIBs	24
Table 3 - Maximum amplitude and encoding formats for the 16 most popular 2.0 CMs in US	26
Table 4 - Band-Edge Operation Impact on Tap Energy (no Micro-reflections).....	32
Table 5 - Pre-equalization Metrics at Band-Edge (No Micro-reflections)	33
Table 6 - Band-Edge Operation Impact on Tap Energy (with 0.5 μ s Micro-reflection)	33
Table 7 - Pre-equalization Metrics at Band-Edge (with 0.5 μ s Micro-reflection).....	34
Table 8 - Micro-reflection Impairment on Pre and Post Main Tap Energy	35
Table 9 - Pre-equalization Metrics at Middle of Upstream Band (with 0.5 μ s Micro-reflection)	36
Table 10 - Low Rate – Once Daily – Rotating Eight Hour Time Shifts - Three day cycle.....	38
Table 11 - Low Rate for Three CM groups.....	38
Table 12 - Medium Rate – Once Every Four Hours - One day cycle.....	38
Table 13 - Medium Rate – Once Every Four Hours - One day cycle - Four Groups (All Times in EST)	38
Table 14 - Tested parts for Short Reference Plant.....	39
Table 15 - Pre-equalization Coefficients of upstream path with and without micro-reflection.....	42
Table 16 - Two CMs showing micro- reflection amplitude over 2 days.....	51
Table 17 - Micro-reflection amplitude of two CMs showing intermittent issue	52
Table I-1 - Frequency, wavelength, and phase relationships in 100 feet of coax.....	81
Table I-2 - Velocity of propagation versus transit delay	112

1 SCOPE

1.1 Introduction and Purpose

As cable networks evolve, and many diverse services such as telephony, data, video, business and advanced services (e.g., tele-medicine, remote education, home monitoring) are carried over them, the demand for maintaining a high level of reliability for services increases. To achieve such high reliability, operators have to fix problems before they have any impact on service.

Increasingly, intelligent end devices are deployed in cable networks, and termination devices and monitoring instruments are installed in headends (HEs) and hubs. The new devices being deployed by operators such as digital set-top boxes (STBs), multimedia terminal adapters (MTAs) and embedded MTAs, hybrid monitoring systems and even high end television sets are DOCSIS capable, resulting in DOCSIS ubiquity. A conservative scenario in a serving area assuming 60% penetration of STBs, 35% of cable modems (CMs) and 15% of eMTAs all enabled with DOCSIS clearly highlights the trend towards DOCSIS ubiquity.

As DOCSIS devices evolve and are equipped with elaborate monitoring tools, it becomes practical to use them for plant monitoring purposes. By using these devices as network probes, cable operators can collect device and network parameters. Combining the analysis of the data along with network topology and device location, it is possible to isolate the source of a problem. A proactive maintenance plan can be developed using this information.

This document describes guidelines and best practices for proactive network maintenance mechanisms that rely on DOCSIS upstream pre-equalization coefficients. The processes described here will help cable operators and industry vendors implement smart monitoring tools, improve maintenance practices, gain better insight in network problems, and enhance network reliability, among other things.

Even though the focus for the development of a proactive network maintenance strategy is through the use of pre-equalization coefficients, the intent is for this effort to expand in the future to include other plant metrics that could help identify and resolve plant issues.

The key outcome of this effort is the reduction of troubleshooting and problem resolution time thereby reducing operational costs. In addition, improvements in network reliability enable the introduction of business and advanced services that require SLAs (service level agreements) thereby generating new revenue. This mechanism adds the capability to detect and resolve problems before they impact customer service, which helps with churn reduction.

2 REFERENCES

2.1 Informative References

- [DOCSIS OSSIV2.0] Operations Support System Interface Specification, CM-SP-OSSIV2.0-C01-081104, November 4, 2008, Cable Television Laboratories, Inc.
- [DOCSIS OSSIV3.0] Operations Support System Interface Specification, CM-SP-OSSIV3.0-I14-110210, February 10, 2011, Cable Television Laboratories, Inc.
- [DOCSIS PHYv3.0] Physical Layer Specification, CM-SP-PHYv3.0-I09-101008, October 8, 2010, Cable Television Laboratories, Inc.
- [DOCSIS RFIv2.0] Radio Frequency Interface Specification, CM-SP-RFIv2.0-C02-090422, April 22, 2009, Cable Television Laboratories, Inc.

2.2 Reference Acquisition

- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027; Phone +1-303-661-9100; Fax +1-303-661-9199; <http://www.cablelabs.com>

3 TERMS AND DEFINITIONS

This document uses the following terms:

adaptive equalizer	A circuit in a QAM receiver that compensates for channel response impairments. In effect, the circuit creates a digital filter that has approximately the opposite complex frequency response of the channel through which the desired signal was transmitted.
adaptive equalizer tap	See <i>tap</i> .
adaptive pre-equalizer	A circuit in a DOCSIS 1.1 or newer cable modem that pre-equalizes or pre-distorts the transmitted upstream signal to compensate for channel response impairments. In effect, the circuit creates a digital filter that has approximately the opposite complex frequency response of the channel through which the desired signal is to be transmitted.
additive impairment	Noise which is added to the desired signal, and which is generally independent of the signal. Includes thermal noise, narrowband ingress, and impulse/burst noise.
amplitude ripple	Nonflat frequency response in which the amplitude-versus-frequency characteristic of the channel or operating spectrum has a sinusoidal or scalloped sinusoidal shape across a specified frequency range.
amplitude tilt	Nonflat frequency response in which the amplitude-versus-frequency characteristic of the channel or operating spectrum is sloped or tilted across a specified frequency range.
cable modem (CM)	A modulator-demodulator at subscriber locations intended for use in conveying data communications on a cable television system.
cable modem termination system (CMTS)	A device located at the cable television system headend or distribution hub, which provides complementary functionality to the cable modems to enable data connectivity to a wide-area network.
channel	A portion of the electromagnetic spectrum used to convey one or more RF signals between a transmitter and receiver.
characteristic impedance	In a transmission line such as coaxial cable, a constant that is the ratio of voltage E to current I in a traveling wave, expressed in ohms, and defined mathematically as $Z_c = (E/I)_{\text{traveling wave}}$. Coaxial cable characteristic impedance is further related to the diameters of the center conductor and inside surface of the shield, and the dielectric material's dielectric constant (see the tutorial in Appendix I.5, Impedance Mismatch).
coefficient	Complex number that establishes the gain of each tap in an adaptive equalizer.
common path distortion (CPD)	Clusters of second and third order distortion beats generated in a diode-like nonlinearity such as a corroded connector in the signal path common to downstream and upstream. The beats tend to be prevalent in the upstream spectrum. When the primary RF signals are digitally modulated signals instead of analog TV channels, the distortions are noise-like rather than clusters of discrete beats.
composite second order (CSO)	Clusters of second order distortion beats generated in cable network active devices that carry multiple RF signals. When the primary RF signals are digitally modulated signals instead of analog TV channels, the distortions are noise-like rather than clusters of discrete beats.

composite triple beat (CTB)	Clusters of third order distortion beats generated in cable network active devices that carry multiple RF signals. When the primary RF signals are digitally modulated signals instead of analog TV channels, the distortions are noise-like rather than clusters of discrete beats.
convolution	A process of combining two signals in which one of the signals is time-reversed and correlated with the second signal. The output of a filter is the convolution of its impulse response with the input signal.
correlation	A process of combining two signals in which the signals are multiplied sample-by-sample and summed; the process is repeated at each sample as one signal is slid in time past the other.
cross modulation (XMOD)	A form of television signal distortion where modulation from one or more television channels is imposed on another channel or channels.
decibel (dB)	Ratio of two power levels expressed mathematically as $dB = 10\log_{10}(P_1/P_2)$.
decibel millivolt (dBmV)	Unit of RF power expressed in terms of voltage, defined as decibels relative to 1 millivolt, where 1 millivolt equals 13.33 nanowatts in a 75 ohm impedance. Mathematically, $dBmV = 20\log_{10}(\text{value in mV}/1 \text{ mV})$.
decision feedback equalizer (DFE)	An adaptive equalizer, usually comprising a combination of feedforward and feedback filters, that uses previously detected symbols to suppress inter-symbol interference in the current symbol being detected.
discrete Fourier transform (DFT)	Part of the family of mathematical methods known as Fourier analysis, which defines the “decomposition” of signals into sinusoids. Forward DFT transforms from the time to the frequency domain, and inverse DFT transforms from the frequency to the time domain.
downstream	The direction of RF signal transmission from headend to subscriber. In North American cable networks, the downstream or forward spectrum occupies frequencies from 54 MHz to as high as 1002 MHz.
drop	Coaxial cable and related hardware that connects a residence or service location to a tap in the nearest coaxial feeder cable. Also called drop cable or subscriber drop.
embedded multimedia terminal adapter (eMTA)	A multimedia terminal adapter that has been combined with a cable modem (see <i>multimedia terminal adapter</i>).
equalizer tap	See <i>tap</i> .
fast Fourier transform (FFT)	An algorithm to compute the discrete Fourier transform (DFT), typically far more efficiently than methods such as correlation or solving simultaneous linear equations.
feeder	Outside plant “hardline” coaxial cables that are part of the coaxial distribution network. These coaxial cables are installed on utility poles or buried underground, are routed near the homes in the service area, and have taps installed that are used to provide connections to the subscribers’ premises.
feeder tap	See <i>tap</i> .
feedforward equalizer (FFE)	An adaptive equalizer that corrects the received waveform using samples of the waveform itself at successive time delays, not using information about the logical decisions made on the waveform.
fiber node	See <i>node</i> .

finite impulse response (FIR)	A type of filter or system in which the impulse response is finite in duration—that is, the impulse response settles to zero in a finite number of sample intervals. A FIR filter is usually implemented as a tapped delay line, with the tap outputs weighted and summed to produce each output sample.
forward	See <i>downstream</i> .
forward error correction (FEC)	A method of error detection and correction in which redundant information is sent with a data payload in order to allow the receiver to reconstruct the original data if an error occurs during transmission.
frequency response	A complex quantity describing the flatness of a channel or specified frequency range, and that has two components: amplitude (magnitude)-versus-frequency, and phase-versus-frequency.
group delay	The negative derivative of phase with respect to frequency, expressed mathematically as $GD = -(d\phi/d\omega)$ and expressed in units of time such as nanoseconds.
group delay variation (GDV) or group delay distortion	The difference in propagation time between one frequency and another. That is, some frequency components of the signal may arrive at the output before others, causing distortion of the received signal.
group delay ripple	Group delay variation which has a sinusoidal or scalloped sinusoidal shape across a specified frequency range.
headend	A central facility that is used for receiving, processing, and combining broadcast, narrowcast and other signals to be carried on a cable network. Somewhat analogous to a telephone company's central office. Location from which the DOCSIS cable plant fans out to subscribers.
hybrid fiber/coax (HFC)	A broadband bidirectional shared-media transmission system or network architecture using optical fibers between the headend and fiber nodes, and coaxial cable distribution from the fiber nodes to the subscriber locations.
impedance	The combined opposition to current in a circuit that contains both resistance and reactance, represented by the symbol Z and expressed in ohms. See also <i>characteristic impedance</i> .
impedance mismatch	Any variation in the uniformity of the nominal impedance of a transmission line or device connected to a transmission line, and which generates a reflected wave.
impulse noise	Noise that is bursty in nature, characterized by non-overlapping transient disturbances. May be repetitive. Generally of short duration—from about 1 microsecond to a few tens of microseconds—with a fast risetime and moderately fast falltime.
impulse response	The output of a filter when its input is excited by an impulse function.
impulse function	A sequence of samples consisting of a single 1, surrounded by all 0s. Also called Kronecker delta function.
incident wave	A traveling wave in a transmission line that is propagating from the source toward the load.
index of refraction	The ratio of the velocity of an electromagnetic wave—specifically what is known as a transverse electromagnetic (TEM) mode wave—in a vacuum to its velocity in a dielectric material, $v_{TEM}(\text{vacuum})/v_{TEM}(\text{dielectric})$.

infinite impulse response (IIR)	A type of filter or system in which the impulse response is infinite in duration—that is, the impulse response keeps going on forever. A IIR filter is usually implemented as a feedback mechanism, where both the input and the previous output are used to produce the next output.
least mean squares (LMS)	Search or adaptive algorithm used in adaptive transversal (tapped delay line) filters. The algorithm attempts to minimize the error energy that occurs between the output and the detected or desired signal.
linear distortion	Distortion that occurs when the overall response of the system (including transmitter, cable plant, and receiver) differs from the ideal or desired response. This class of distortions maintains a linear, or 1:1, signal-to-distortion relationship (increasing signal by 1 dB causes distortion to increase by 1 dB), and often occurs when amplitude-versus-frequency and/or phase-versus-frequency depart from ideal. Linear distortions include impairments such as micro-reflections, amplitude ripple/tilt, and group delay variation, and can be corrected by an adaptive equalizer.
media access control (MAC)	A sublayer of the OSI Model’s data link layer (Layer 2), which manages access to shared media such as the OSI Model’s physical layer (Layer 1).
media access control (MAC) address	The “built-in” hardware address of a device connected to a shared medium.
micro-reflection	A short time delay echo or reflection caused by an impedance mismatch. A micro-reflection’s time delay is typically in the range from less than a symbol period to several symbol periods.
modulation error ratio (MER)	The ratio of average symbol power to average error power. The higher the MER, the cleaner the received signal.
multimedia terminal adapter (MTA)	A device that provides an interface between analog telephones and an IP network.
node	An optical-to-electrical (RF) interface between a fiber optic cable and the coaxial cable distribution network. Also called fiber node.
noise	See <i>thermal noise</i> .
nonlinear distortion	A class of distortions caused by a combination of small signal nonlinearities in active devices and by signal compression that occurs as RF output levels reach the active device’s saturation point. Nonlinear distortions generally have a nonlinear signal-to-distortion amplitude relationship—for instance, 1:2, 1:3 or worse (increasing signal level by 1 dB causes distortion to increase by 2 dB, 3 dB, or more). The most common nonlinear distortions are even order distortions such as composite second order (CSO), and odd order distortions such as composite triple beat (CTB). Passive components can generate nonlinear distortions under certain circumstances.
pre-equalizer	See <i>adaptive pre-equalizer</i> .
QAM receiver	A circuit that receives, processes, and demodulates a QAM signal.
QAM signal	Analog RF signal that uses quadrature amplitude modulation to convey information.
quadrature amplitude modulation (QAM)	A modulation technique in which an analog signal’s amplitude and phase vary to convey information, such as digital data. The name “quadrature” indicates that amplitude and phase can be represented in rectangular coordinates as in-phase (I) and quadrature (Q) components of a signal.
quadrature	Two sine waves are in quadrature if their phases differ by 90 degrees, such as sine and cosine.

radio frequency (RF)	That portion of the electromagnetic spectrum from a few kilohertz to just below the frequency of infrared light.
reflected wave	A traveling wave in a transmission line, caused by an impedance mismatch, that is propagating from the point where the impedance mismatch exists back toward the incident wave's source.
reflection coefficient	Ratio of reflected voltage E^- to incident voltage E^+ , expressed mathematically as $\Gamma = E^-/E^+$ where Γ (uppercase Greek letter gamma) is the reflection coefficient.
return	See <i>upstream</i> .
return loss (R)	The ratio of incident power P_I to reflected power P_R , expressed mathematically as $R = 10\log_{10}(P_I/P_R)$, where R is return loss in decibels.
reverse	See <i>upstream</i> .
root mean square (RMS)	A statistical measure of the magnitude of a varying quantity such as current or voltage, where the RMS value of a set of instantaneous values over, say, one cycle of alternating current is equal to the square root of the mean value of the squares of the original values.
receive MER (RxMER)	The modulation error ratio at the receiver, at the point at which symbol decisions are made. A high RxMER results in a clean constellation plot, where each symbol point exhibits a tight cluster separated from the neighboring symbols. (See Appendix I.)
standing wave	A distribution of fields along a transmission line caused by the interaction of an incident and reflected wave, such that the peaks and troughs of the wave are stationary in location.
subscriber drop	See <i>drop</i> .
tap	(1) In the feeder portion of a coaxial cable distribution network, a passive device that comprises a combination of a directional coupler and splitter to "tap" off some of the feeder cable RF signal for connection to the subscriber drop. So-called self-terminating taps used at feeder ends-of-line are splitters only and do not usually contain a directional coupler. Also called a multitap. (2) The part of an adaptive equalizer where some of the main signal is "tapped" off, and which includes a delay element and multiplier. The gain of the multipliers are set by the equalizer's coefficients. (3) One term of the difference equation in a finite impulse response or a infinite impulse response filter. The difference equation of a FIR follows: $y(n) = b_0x(n) + b_1x(n-1) + b_2x(n-2) + \dots + b_Nx(n-N)$.
thermal noise	The fluctuating voltage across a resistance due to the random motion of free charge caused by thermal agitation. When the probability distribution of the voltage is Gaussian, the noise is called additive white Gaussian noise (AWGN).
upstream	The direction of RF signal transmission from subscriber to headend. Also called return or reverse. In most North American cable networks, the upstream spectrum occupies frequencies from 5 MHz to as high as 42 MHz.
vector	A quantity that expresses magnitude and direction (or phase), and is represented graphically using an arrow.
velocity factor (V)	The reciprocal of index of refraction, expressed in decimal format.
velocity of propagation (VP or VoP)	The speed at which an electromagnetic wave travels through a medium such as coaxial cable, expressed as a percentage of the free space value of the speed of light. For example, VP in a typical coaxial cable is about 85% to 87% of the speed of light. VP in a typical single mode optical fiber is about 67% to 69%.

voltage standing wave ratio (VSWR)

Ratio of a standing wave's maximum voltage E_{\max} to its minimum voltage E_{\min} along a transmission line, expressed mathematically as $VSWR = E_{\max}/E_{\min}$, or $VSWR = (1+|\Gamma|)/(1-|\Gamma|)$.

4 ABBREVIATIONS AND ACRONYMS

This document uses the following abbreviations:

AC	alternating current
AWG	American Wire Gauge
CATV	cable television
CM	cable modem
CMTS	cable modem termination system
CNR	carrier-to-noise ratio
CPD	common path distortion
CSO	composite second order
CTB	composite triple beat
dB	decibel
dBc	decibels relative to carrier
dBmV	decibel millivolt
DOCSIS®	Data-Over-Cable Service Interface Specification
E_s/N_0	energy per symbol-to-noise density ratio
FEC	forward error correction
FSE	fractionally spaced equalizer
GD	group delay
GDV	group delay variation (also called group delay distortion)
HEX	hexadecimal
HFC	hybrid/fiber coax
Hz	hertz
IEEE	Institute of Electrical and Electronics Engineers
ISI	inter-symbol interference
IP	Internet protocol
kHz	kilohertz
km	kilometer
LMS	least-mean-square
MAC	media access control
MTC	main tap compression
MTNE	main tap nominal energy
MER	modulation error ratio
MHz	megahertz
MIB	management information base

ML	maximum likelihood
MPEG	Moving Picture Experts Group
ms	millisecond
MSE	mean square error
Msym/sec	megasymbols per second
mV	millivolt
NMTER	non-main tap to total energy ratio
ns	nanosecond
NTSC	National Television System Committee
OID	object identifier
PostMTE	post-main tap energy
PPTSR	pre-post tap symmetry ratio
PreMTE	pre-main tap energy
QAM	quadrature amplitude modulation
R	return loss
RF	radio frequency
RNG-REQ	ranging request
RNG-RSP	ranging response
RxMER	receive modulation error ratio
SAW	surface acoustic wave
S-CDMA	synchronous code division multiple access
SID	service identifier
SNMP	simple network management protocol
SNR	signal-to-noise ratio
TDMA	time division multiple access
TEM	transverse electromagnetic
TLV	type length value
TV	television
VSWR	voltage standing wave ratio
XMOD	cross modulation
UDP/IP	user datagram protocol/Internet protocol
μs	microsecond
16-QAM	16-state quadrature amplitude modulation
64-QAM	64-state quadrature amplitude modulation
256-QAM	256-state quadrature amplitude modulation

5 BACKGROUND

5.1 Reactive versus Proactive Network Maintenance

In this document the definition of reactive network maintenance is a stringent one and in the era of multimedia, business and advanced services is perhaps one that the cable industry should follow. Reactive network maintenance consists of network maintenance practices that are initiated by metrics which show service performance has been impacted. Under this definition, it is not assumed the need of customer feedback/calls for network maintenance to be reactive. As long as conditions such as FEC statistics, power starvation, CPD, narrowband interferers, laser clipping, impulse noise and other service impacting symptoms are detected, the response to such is reactive. In the case of distortion that is completely corrected by pre-equalization there is no performance impact, therefore maintenance actions that arise from these impairment discoveries are considered to be proactive. If there are symptoms that combine performance impacting conditions with distortion detected through equalization, the maintenance that arises from it is still considered reactive.

Impairments that result in network maintenance classified as reactive would likely be given higher priority for resolution because they are already impacting performance.

In most scenarios, upstream pre-equalization mechanisms can completely compensate for certain problems in the network and no symptoms are detected in FEC statistics or through other metrics. The fact that equalization can fully compensate for network linear distortion can buy the operator time in resolving the issue before there is service impact, thus enabling a proactive network maintenance strategy.

5.2 Linear Impairments

In the upstream portion of the CATV network there are different types of impairments. These can be classified based on the impact these impairments have on the signal as linear as well as nonlinear impairments. In the case of a “linear” impairment, the impact on the signal will be given by a change in amplitude and phase of the original signal. In the case of a “nonlinear impairment”, the signal generates distortion components, including harmonics of the original signal or multiplies the original signal with other energy present in the return band. For example, in the linear distortion case, transmitting information across the upstream channel will result in an amplitude and phase deviation for a given frequency point. A micro-reflection, which is analogous to a wireless multipath signal, could result from the bouncing back and forth of a signal between two interfaces that have impedance mismatches, generating an amplitude and phase distortion of the signal as summation of a time-delayed signal copy—the reflection or echo—combined with the desired signal. A second example of a linear distortion occurs at the diplex filter rolloff that marks the upper edge of the upstream frequency spectrum around 42 MHz. At this rolloff frequency, the amplitude and the phase suffer considerable distortion. In particular, the phase distortion is noticeable prior to reaching the band-edge. This phase distortion is more easily shown when expressed as group delay which is defined as:

$$GroupDelay = -\frac{d\phi}{d\omega}$$

where ϕ is the phase in radians, ω is the frequency in radians per second, and group delay is in seconds. Group delay is ideally constant across the band of interest. Group delay variation across the band is known as group delay distortion. Additional discussion about linear impairments can be found in the tutorial material in Appendix I.

One thing to keep in mind is that even impairments that are considered “nonlinear” such as common path distortion (CPD) may have associated linear distortion elements. For instance, the corrosion of a center conductor that generates a mixer effect also results in an impedance mismatch that can generate a noticeable micro-reflection.

Examples of impairments that are nonlinear include the previously mentioned CPD, as well as composite second order (CSO) and composite triple beat (CTB) distortions, cross-modulation, and laser clipping. Ingress and impulse noise are considered additive impairments, although if an impairment such as impulse noise is high enough to cause laser clipping, it can be considered nonlinear in nature too.

5.2.1 Micro-reflection Types

The following sections describe multiple scenarios in which micro-reflections may be generated within the HFC network. Micro-reflection Example 1 describes a case where there are HFC components whose low return loss (R) values are contributing to the micro-reflection source. Micro-reflection Example 2 describes a case where an HFC component whose isolation performance and an impedance mismatch are contributing to the micro-reflection source. Lastly, micro-reflection Example 3 describes a case where the micro-reflection impairments represent a combination of the two previous cases. Micro-reflection sources are not limited to the examples presented here. For more information about micro-reflections, refer to the tutorial section of this document.

5.2.1.1 Micro-reflection Example 1

This type of commonly experienced micro-reflection is exhibited when the upstream signal encounters impedance mismatches somewhere in its upstream path to the CMTS, causing redirection of a fraction of the signal's energy back towards the CM. Once the redirected signal becomes incident on another impedance mismatch, it is then directed back toward the CMTS. Figure 1 illustrates an upstream signal, labeled Upstream Signal #2, on the cable originating between two reflection points, Γ_1 , and Γ_2 . However, the upstream signal may originate anywhere downstream from the first reflection point, Γ_1 , illustrated in the figure as Upstream Signal #1. The majority of the upstream signal passes through the impedance mismatch and continues toward the CMTS and is labeled Main Signal in the figure. A fraction of the upstream signal is reflected back towards the signal source at Γ_1 . This reflected signal encounters a second reflection point, Γ_2 which reflects a fraction of the reflected signal energy back into the direction of the original signal, represented as Reflected Signal in Figure 1.

This general case describes the signal passing through two reflectors, neither of which is the CM itself. One technique to determine whether the CM is one of the reflectors is to add a 3 dB inline attenuator at the CM port (upstream output). If the micro-reflection magnitude is lessened by 6 dB or more, because of the reflected signal passing through the attenuator twice, then the CM is indeed one of the reflectors. If there is no significant change in micro-reflection magnitude, then the pair of reflectors lies remote to the CM location. This technique of adding attenuation would work equally well to isolate reflections elsewhere in the network, but it is much more difficult to install attenuation between line passives.

The reflected signal from Γ_2 proceeds upstream and again encounters the reflection point Γ_1 . This causes yet another reflection back downstream, and the process repeats endlessly. (It is analogous to looking at one's reflection in a mirror, when there is another mirror behind. There will be an endless sequence of images, each one progressively smaller.) Each passage between Γ_1 and Γ_2 is called a transit. The main micro-reflection echo results from the triple transit, up/down/up. The next echo is from the 5th-order transit up/down/up/down/up, and so on. This type of response, which goes on forever, is called infinite impulse response (IIR). It consists of a geometric series of echoes, separated by equal delays (equal to twice the propagation delay between the points Γ_1 and Γ_2), with each echo value smaller than the last (by the same ratio $\Gamma_1 \Gamma_2$ or, dB difference $10 \log \Gamma_1 + 10 \log \Gamma_2$). The decaying micro-reflections may quickly be lost in the noise floor after two or three trips through the micro-reflection source.

Conversely, the adaptive equalizer response, which is approximately the inverse of the channel response, will have only a single tap after the main tap for this case.

Though the reflector examples in Figure 1 are feeder taps, it should be noted that many devices can produce similar results, including damaged cables or corroded splices, which are often causes of micro-reflections as well.

Note that the delay between echoes equals twice the propagation delay between the two reflection points Γ_1 and Γ_2 , so the distance between the two reflectors is known. However it does not relate how far along the plant these two

reflectors are, that is, the location of the impairment in the plant. To determine the location of the fault, additional information is necessary as described in Section 6.6, Fault Localization.

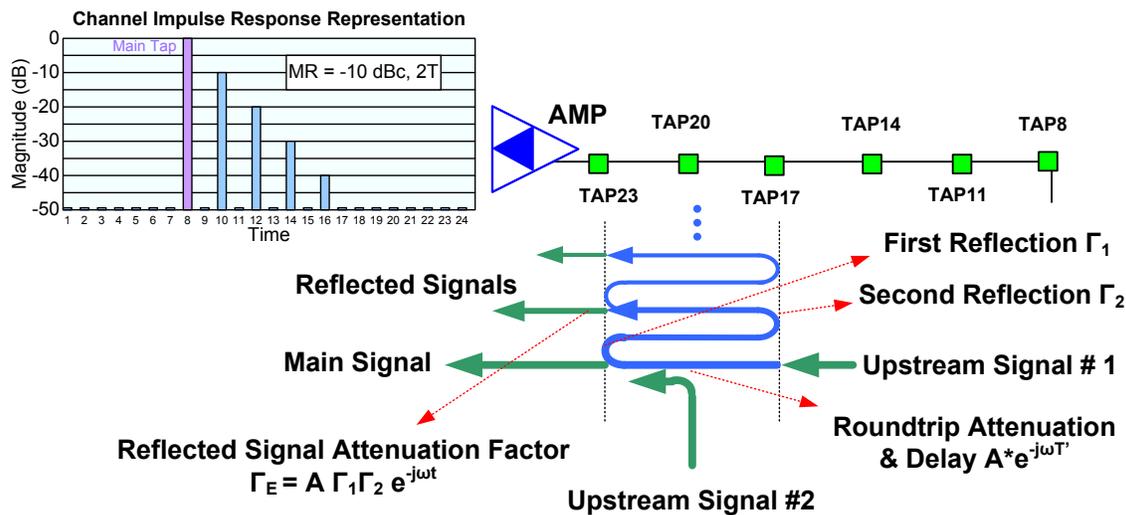


Figure 1 - Micro-reflection with multiple-transit echoes

5.2.1.2 Micro-reflection Example 2

The second type of micro-reflection may occur anywhere in the network, but its magnitude is most noticeable in CMs located off the highest value feeder taps. Figure 2 describes the upstream signal flow from a CM sending its upstream signal into the port of the 23 dB feeder tap. The intent of the plant design is that the primary path of upstream signal flow is toward the left of the page, toward the CMTS. Some signal energy may be reflected from the amplifier to the left of the 23 dB feeder tap, but this energy is usually insignificant because of good impedance matching of the amplifier.

The 23 dB feeder tap is a combination directional coupler and splitter meaning it directs the upstream signal in the upstream direction. The directivity of the feeder tap is not perfect. Some of the energy from the CM may be leaked towards the output connector. This is a function of the isolation of the feeder tap. Isolation is specified with all ports terminated, and cable industry field practice has often not adhered to terminating unused feeder tap ports. In addition, corrosion of the feeder tap can degrade its isolation performance.

If a 23 dB feeder tap has a tap-port-to-output-port isolation of only 38 dB, there would be two signal paths created, one with 23 dB attenuation, and one additional signal going the other direction with only 38 dB attenuation. If this additional signal meets a reflector downstream from the 23 dB feeder tap, the additional signal will be reflected back in the direction of the upstream signal and will be combined with the main signal at the 23 dB feeder tap where it originated. This micro-reflection is easily observed in cases where there is long un-tapped span of cable between the first feeder tap, the 23 dB feeder tap in this case, and low-value feeder taps with open terminations. All of these unterminated feeder tap ports create their own micro-reflections. This condition usually creates multiple micro-reflections, and since there was an open port at the 11 dB feeder tap and another at the 8 dB feeder tap, there would be two different cable lengths resulting in two different micro-reflection delay characteristics.

This case is important to mention since the reflector, Γ_1 , is not in the tree of the devices between the CM and the CMTS, but includes devices that are downstream from the CM feeder tap location. This phenomenon, as mentioned earlier, is more noticeable at the high value feeder tap, because as the feeder tap value decreases, the amplitude difference between the desired signal and the micro-reflection increases. If both the 23 dB feeder tap and the 11 dB feeder tap had 35 dB tap-port-to-output-port isolation, there would be 12 dB greater separation between desired and the micro-reflection signal. Additionally, the 23 dB feeder tap location is near the amplifier followed by a long

length of cable, where the 11 dB feeder tap is found near the end of the cable so the cable length for the 11 dB feeder tap is shorter and the micro-reflection delay characteristic is correspondingly less.

This type of micro-reflection does not exhibit an unending IIR response, since the reflected energy from Tap 8 passes through Tap 23 relatively unimpeded and continues upstream. This type of response, which stops after a single echo, is called finite impulse response (FIR). The signature will show a single main echo without trailing echoes.

Conversely, the adaptive equalizer response, which is approximately the inverse of the channel response, will have a sequence of smaller and smaller taps for this case. This is because the equalizer internally generates additional echoes as it cancels the original echo in the signal. As it generates an echo, it must use another tap to cancel the new echo. This process goes on until the end of the equalizer tapped delay line is reached. Any remaining echo energy is uncompensated after this point, and results in reduced RxMER.

To summarize, a multiple-transit echo scenario (Example 1) has an unending sequence of decaying echoes in the channel response, and the corresponding adaptive equalizer response has a single echo. A single-reflection scenario (Example 2) has a single echo in its channel response, and the corresponding adaptive equalizer response has a decaying sequence of echoes continuing to the end of the equalizer tapped delay line.

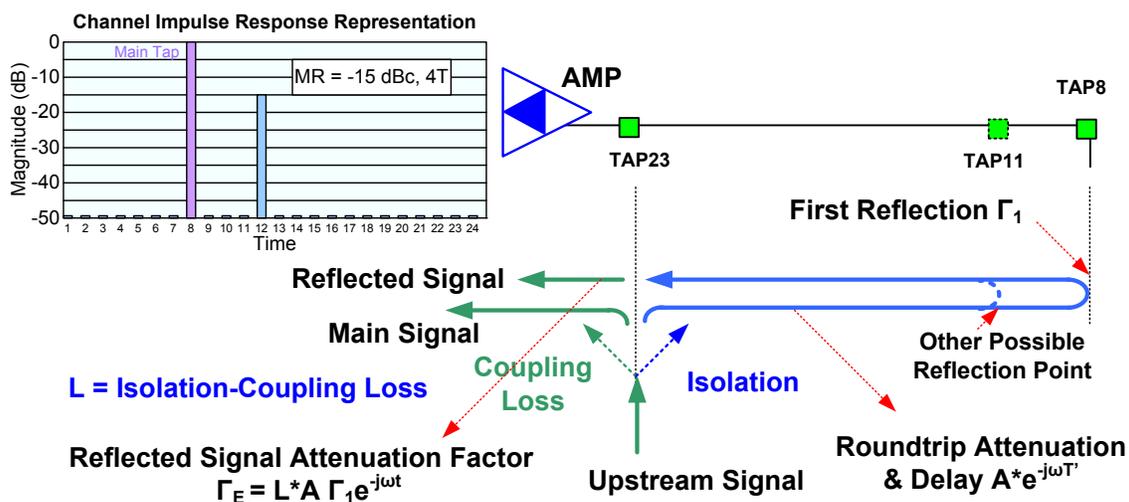


Figure 2 - Micro-reflection with single impedance mismatch interface

5.2.1.3 Micro-reflection Example 3

Figure 3 represents a case that is a superposition of the previous cases described in Micro-reflection Example 1 and Micro-reflection Example 2. A reflective point exists if the amplifier has poor return loss. A fraction of the desired signal is reflected off the output of the amplifier, Γ_1 , and propagates downstream to the open connection at the end-of-line or an open feeder tap port, Γ_2 , where it reverses direction back towards the amplifier and rejoins the original signal at the feeder tap location.

If the amplifier and feeder tap are co-located or are very close together, there will little difference in micro-reflection delay characteristic between the previously described Micro-reflection Example 1 and Micro-reflection Example 2. In fact, they may add or cancel each other out.

However, if the amplifier is a pole span or more away from the feeder tap, there may be two different distinct micro-reflections created for the single unterminated feeder tap. These micro-reflections are usually low in magnitude, 30 dB or more lower than the incident signal, but they can be numerous in a single amplifier-to-termination span.

It is conceivable that multiple micro-reflections may emanate from successive multiple passes through HFC components comprising the micro-reflection source. Cases in which only a single micro-reflection exists may be limited to laboratory simulation of the micro-reflection impairment and may not be practically encountered within the HFC network.

The cable industry has favored “capping” unused feeder tap ports rather than terminating the ports. Some operators will defend that position because it is claimed to reduce potential ingress sources. Indeed it is a tradeoff between craft integrity and good impedance matching practices, which causes more stress on the upstream adaptive equalizer.

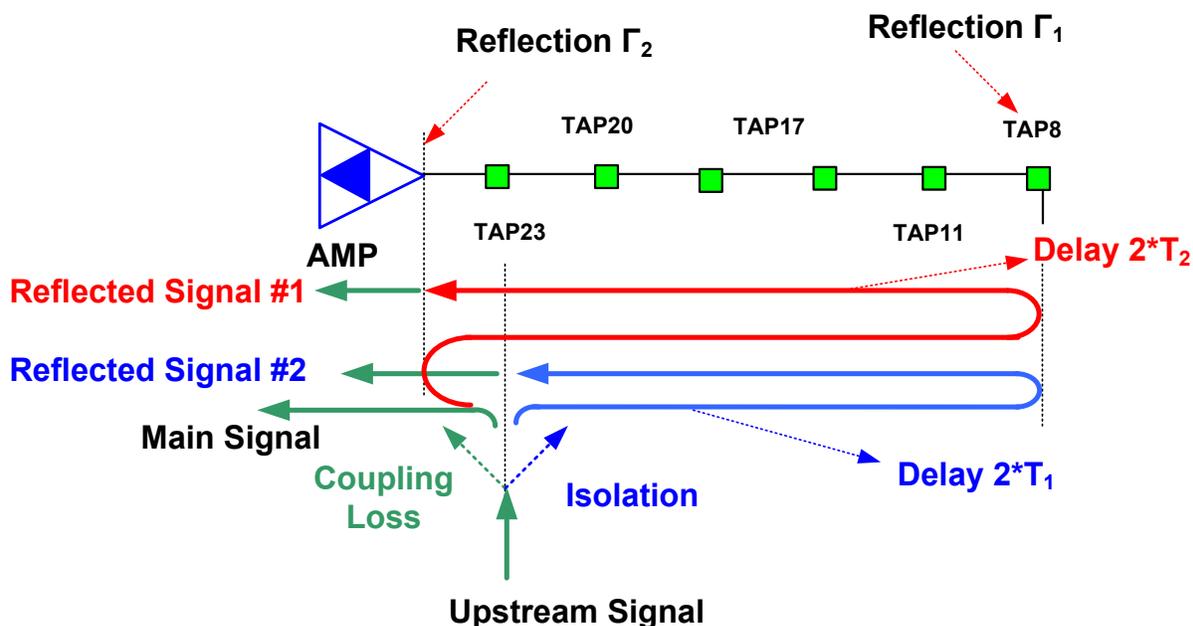


Figure 3 - Composite Micro-reflection resulting from Type 1 and Type 2 Micro-reflections

5.3 Pre-equalization Mechanism Enabled through DOCSIS Ranging

The upstream pre-equalization mechanism relies on the interactions of the DOCSIS ranging process in order to determine and adjust the CM pre-equalization coefficients. The intent is for the CM to use its coefficients to pre-distort the upstream signal such that the pre-distortion equals the approximate inverse of the upstream path distortion, so that as the pre-distorted upstream signal travels through the network it is corrected and arrives free of distortion at the upstream receiver at the CMTS.

The pre-equalization coefficients of the CM are the complex coefficients (F1 through F24) of the 24-tap linear transversal filter structure shown in Figure 4.

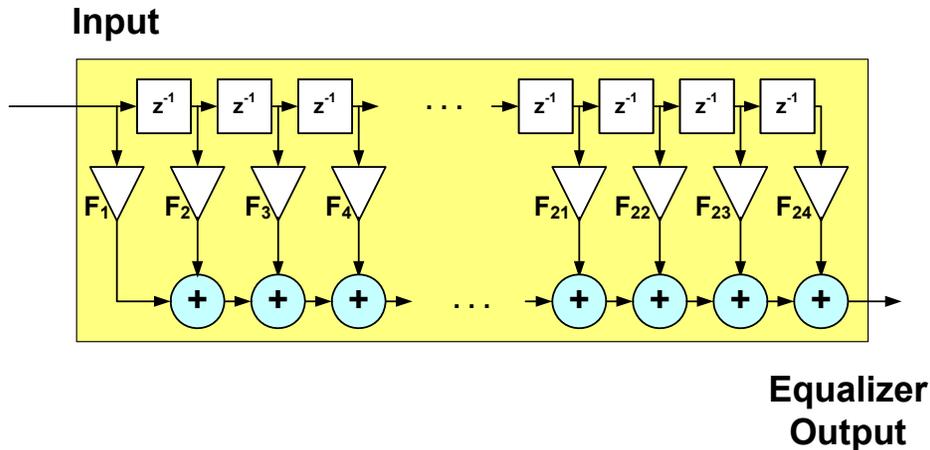


Figure 4 - Upstream Equalizer Structure

In this structure the blocks with z^{-1} label represents delay elements, each of which in the DOCSIS 2.0 pre-equalizer is the symbol period T (in DOCSIS 1.1 it can also represent delays equal to $T/2$ and $T/4$).

In the ranging process the CM sends a ranging request message (RNG-REQ) to the CMTS. The CMTS may use a known portion of this message, such as the preamble, as well as other known messages to determine the quality of the received signal, as well as to determine the adjustment the CM should make to its pre-equalization coefficients to better compensate the upstream distortion. In response to the RNG-REQ message, the CMTS sends a ranging response (RNG-RSP) message with a set of 24 coefficients and a parameter that indicates whether these coefficients are intended to result in a set or adjust operation by the CM. In the case of a set command, the CM will replace its existing coefficients with the ones sent by the CMTS. In the case of an adjust command, the CM convolves its coefficients with the ones sent by the CMTS to achieve the adjusted coefficients (Figure 5).

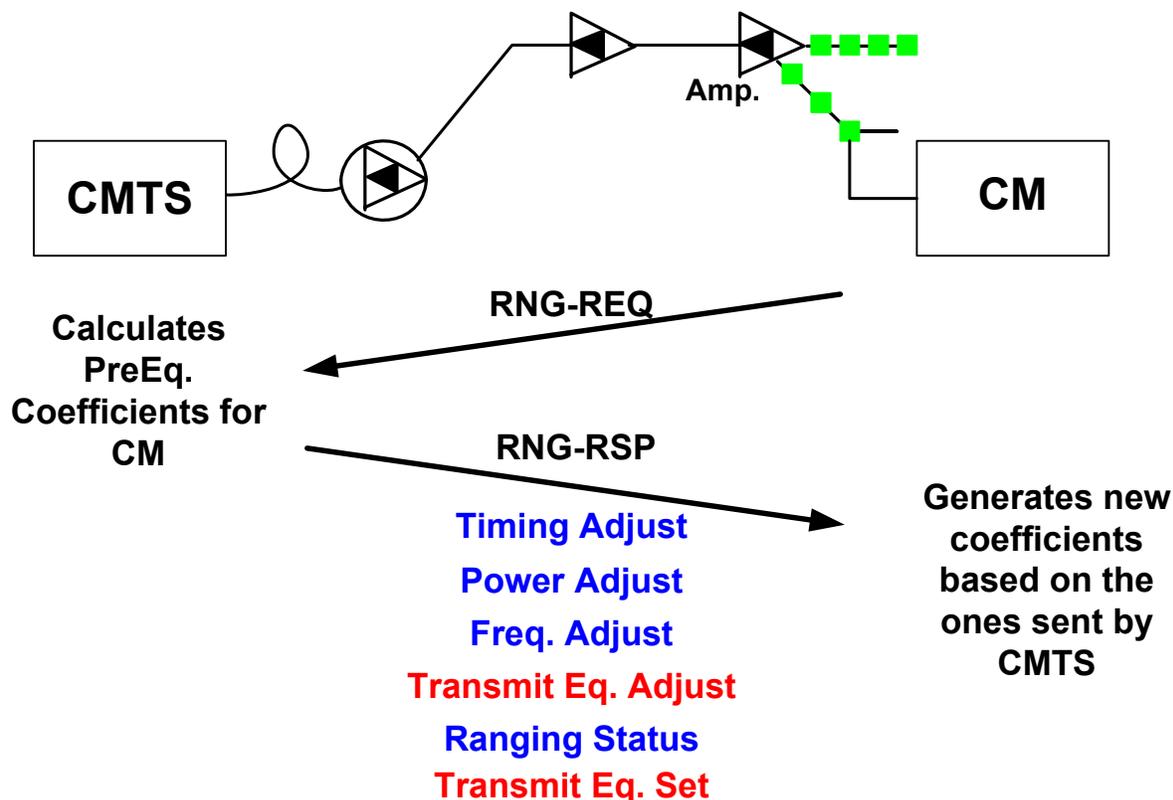


Figure 5 - CM-CMTS Ranging Interaction Enabling Pre-equalization Process

The CMTS may not be completely satisfied with the quality of the signal the CM is sending after the initial try. This is an iterative process which may take a few interactions before the coefficients are stable.

CMTS implementations use for the most part the transmit-equalization-adjust option to convey information. Only after the initial ranging request, one may see a CMTS send a transmit-equalization-set message to make sure that the CM initializes properly. In principle the CMTS could use this message when it needs to reset the coefficients.

A CMTS that is completely satisfied with the values of the pre-equalization coefficients sends an adjust message where all coefficients are zero except for the pre-equalizer's main tap coefficients, which has maximum or nominal value. This represents a Kronecker delta or impulse function, and any data set convolved with an impulse results in the original data set, which in this case is the CM pre-equalization coefficients, unchanged.

5.3.1 Pre-equalization Enabling Messages

As described previously, the two messages that are key in the ranging process are the range response (RNG-RSP) and range request (RNG-REQ) messages. The RNG-RSP message, which is generated by the CMTS in response to a RNG-REQ message, carries timing, frequency, power level, and equalization adjustment information as well as equalization set or load information and ranging status. This information is encoded following what is known as type-length-value (TLV) format. DOCSIS 1.1 pre-equalization coefficients are identified by type 04 and DOCSIS 2.0 or 3.0 by type 09. The RNG-RSP messages that the RNG-REQ messages correspond to are linked by the service ID or SID. SIDs identify upstream service flows. It may be that a CM has several SIDs. In that case a CM will get ranging information through each of the SIDs it has. For example, if a CM has a SID that is used for telephony service and one that is used for data service, there will be two parallel ranging processes within a single CM. In addition to the SID, the RNG-RSP message payload also carries the upstream channel ID. Figure 6 shows the structure of the RNG-RSP message.

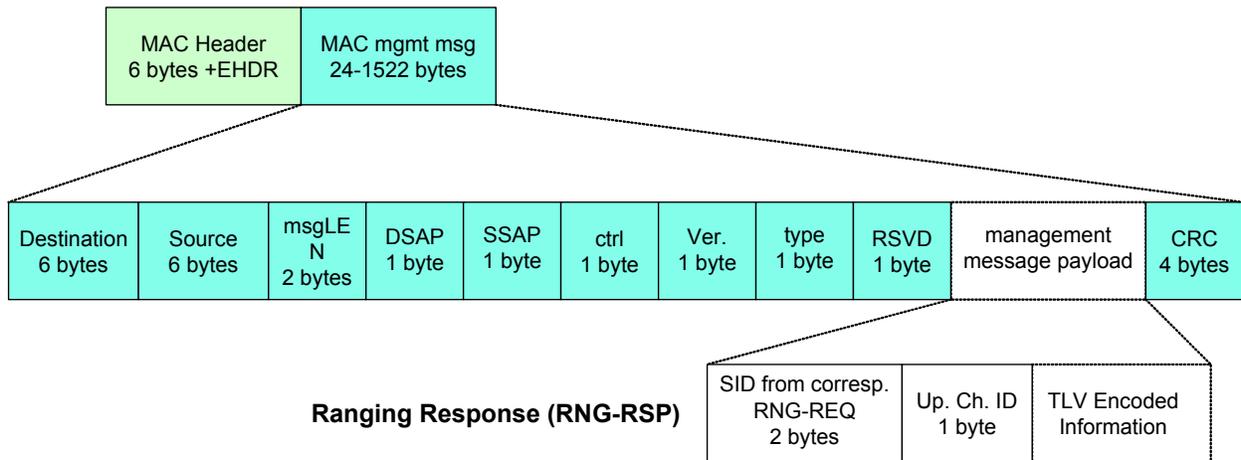


Figure 6 - Range Response Message Format

The RNG-REQ message is generated by the CM and sent to the CMTS. The RNG-REQ is used as the reference to determine whether the CM signal needs any adjustment. These adjustments could be in frequency, power level, timing offset, and distortion. Once the CMTS receives the RNG-REQ message it uses a known portion of this message as the reference of what the signal should look like. Typically that known portion of the message is the preamble. If the CM is not finished implementing the changes the CMTS is asking for, the CM includes in the RNG-REQ message a ranging status indicating whether or not the ranging changes are still pending. This is the “pending till complete” field in the RNG-REQ message payload. The RNG-REQ message also carries a downstream channel ID that associates the upstream being used with a downstream channel. Figure 7 shows the structure of the RNG-REQ message.

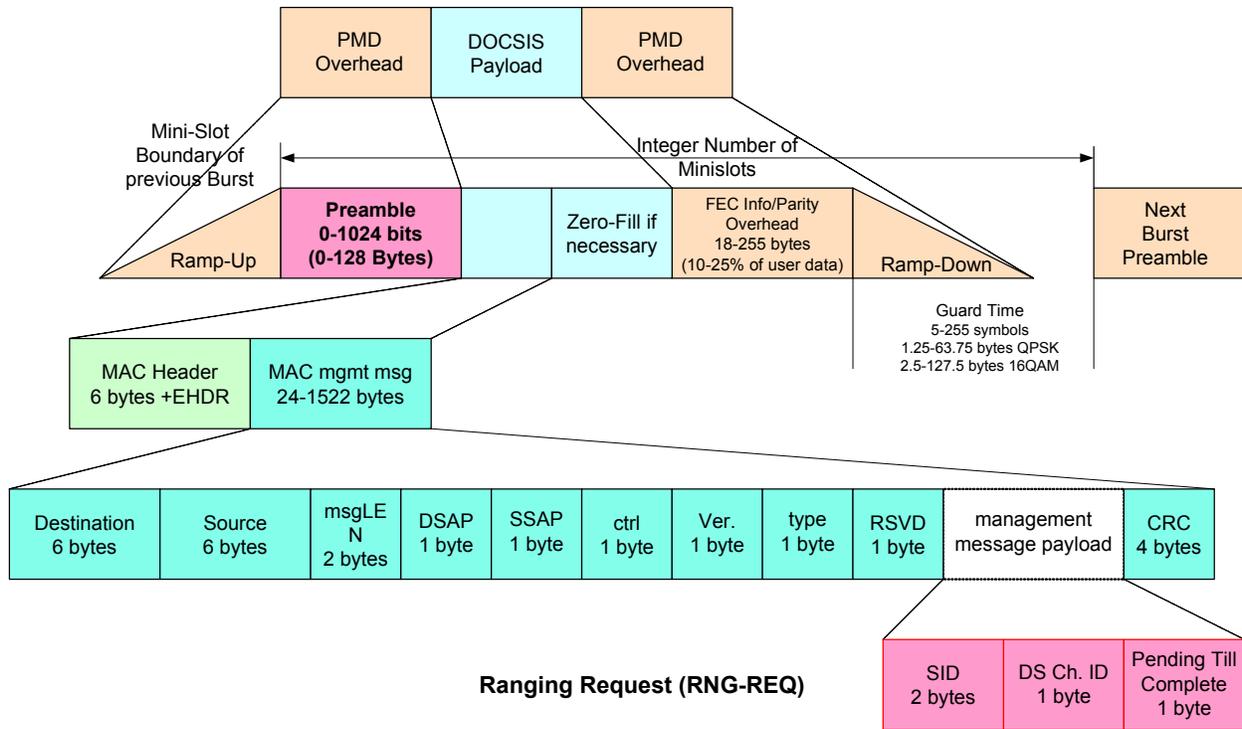


Figure 7 - Range Request Message Format

5.3.2 CM and CMTS Equalization Information

The pre-equalization coefficients are loaded into the adaptive pre-equalizer in the CM, which is used to compensate for upstream linear distortion(s). Hence the CM pre-equalization data indirectly describes the distortion in the plant for which it compensates. The pre-equalizer response is approximately the inverse or opposite response of the plant. The pre-equalization coefficients provide detailed characteristics of the channel distortion, although the coefficients do not directly indicate the level of micro-reflections. Assuming negligible group delay distortion and a single micro-reflection, a quick estimate of micro-reflection level can sometimes be obtained using the energy in the adaptive equalizer’s non-main taps. In general, an elaborate analysis is required to uniquely resolve micro-reflection level/delay signature characteristics. An upstream channel that exhibits no distortion has all the energy concentrated in the adaptive equalizer main tap while one that exhibits distortion also has energy in taps other than the main tap (Figure 8).

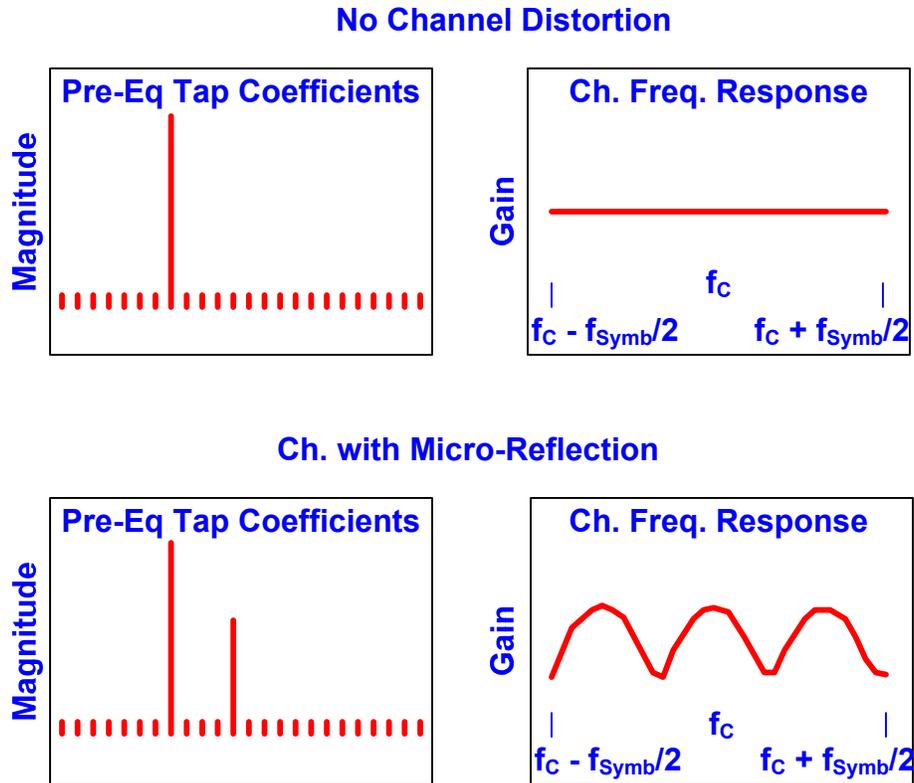


Figure 8 - CM Pre-eq Coefficients Values and Frequency Response Scenarios

The pre-equalization data which the CMTS continues to send to the CM indicates how successful a CM has been in compensating for the distortion by showing what is left to compensate to achieve ideal reception. Ideally and typically, the CM starts with no compensation and after a few ranging intervals, achieves a steady state where the CM compensates for all the distortion. At that point the CMTS pre-equalization data exhibits a flat response indicating that further compensation is not required (Figure 9).

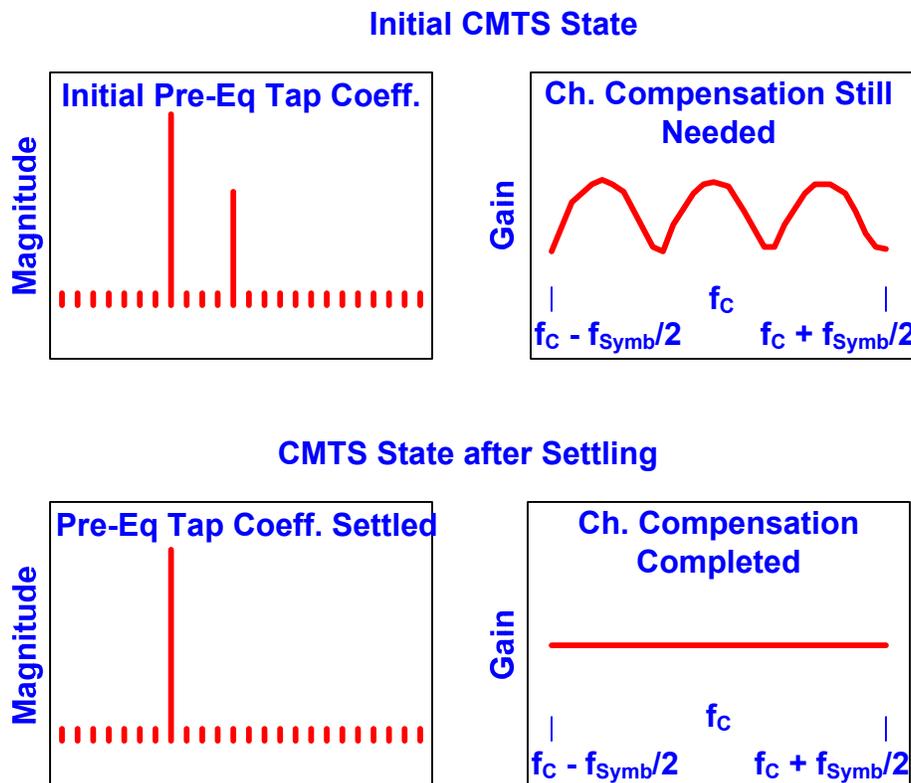


Figure 9 - CMTS CM Pre-equalization Coefficients Values and Frequency Response Scenarios

The upstream CM equalization data collected by the CMTS is analyzed to verify that any plant distortion has been compensated. There is the possibility of a distortion being so severe (e.g., a micro-reflection having a very long delay) that the pre-equalization process would not be able to fully compensate for it. These scenarios are rare in current HFC architectures, but if this does occur, one must be aware that an impairment identification process using only CM pre-equalization data will not yield accurate results.

5.4 Upstream Pre-equalization in DOCSIS 1.0, DOCSIS 1.1 and DOCSIS 2.0

Upstream pre-equalization in DOCSIS 1.0 was left as optional and the equalization process between CMTS and CM was not defined in sufficient detail. An unexpected result occurred when DOCSIS 1.1 and 2.0 were introduced with a well defined process. A few 1.0 CMs that implemented pre-equalization exhibited erratic behavior in the presence of downstream RNG-RSP messages that were generated by 1.1 or 2.0 CMTSs. For quite some time operators have not been motivated to turn pre-equalization on, in part because the demand for capacity and spectrum availability have not been significant enough to warrant the use of wider channels, higher order modulations, or frequencies near the edges of the upstream spectrum where linear distortion occurs.

Some 1.0 CMs exhibiting the problem have been successfully upgraded with firmware that corrects this issue. Unfortunately it has not been possible to correct this issue on all affected CMs. To support reliable use of upstream pre-equalization, operators have been replacing 1.0 CMs having known issues.

5.4.1 DOCSIS 1.1 Pre-equalization Considerations

The percentage of DOCSIS 1.1 CMs deployed is still significant enough not to take advantage of the pre-equalization compensation. Nevertheless, based on the percentage of 1.1 CM population and the rate at which 1.1

CM versions are decreasing with time, it is important to determine at what point the procedures described in this documentation will be worthwhile to implement.

5.5 Limitations on Pre-equalization Compensation

In a scenario of an upstream path that exhibits a micro-reflection, the maximum delay compensation that can be achieved using pre-equalization is limited by the amount of delay that can be generated within the pre-equalization filter structure shown in Figure 4. The maximum delay that can be generated is given by the delay between the adaptive equalizer's main tap and the last adaptive equalizer tap.

In DOCSIS 2.0 and 3.0, the delay or spacing between each adaptive equalizer tap location is equal to the symbol period, because it always has a parameter of adaptive equalizer taps/symbol equal to 1. Typical implementations in DOCSIS 2.0 and 3.0 have the main equalizer tap in the eighth position out of a 24-tap delay line. Therefore the maximum delay that can be generated in that filter structure is 16T (last tap position – main tap position) where T equals the symbol period.

In DOCSIS 1.1 the delay between different adaptive equalizer tap locations can be a fraction of a symbol period. That is, the number of equalizer taps/symbol parameter is allowed to be 1, 2 or 4, resulting respectively in delay differences between adaptive equalizer tap locations of T, T/2 and T/4. This option has not been implemented in a CMTS. Therefore, in DOCSIS 1.1 CMTS scenarios, the maximum delay that can be generated is equal to 4T (last tap position – main tap position). Table 1 shows the maximum delays that are generated in DOCSIS 1.1 and 2.0 or 3.0 filter structures at different symbol rates using the typical equalizer main tap configurations (position 4 for DOCSIS 1.1 and position 8 for DOCSIS 2.0 and 3.0).

Table 1 - Maximum Delays Generated by Pre-equalization Filter Structures in DOCSIS 1.1 and 2.0

Symbol Rate MHz	Symbol Period (T) μ sec	4* T μ sec	16 T μ sec
5.12	0.195		3.125
2.56	0.391	1.563	6.250
1.28	0.781	3.125	12.500

It cannot be assumed that a DOCSIS 2.0 filter structure can fully compensate for a micro-reflection with a delay of 3.125 microseconds. Typically energy in the neighboring equalizer taps help in the fine tuning of that compensation and a micro-reflection that pushes the delay to the limit of the equalizer won't have longer delay equalizer taps available to help in the representation of the exact value. This is especially true if the echo delay is not a multiple of the symbol period, since the equalizer taps are then not spaced at the exact intervals to efficiently cancel the echo, and more equalizer taps are needed to provide effective cancellation. This will impact more severely higher order modulation scenarios such as 64-QAM where the adjustment is more critical.

In addition, in the case of strong micro-reflections, the equalizer may have a decaying sequence of taps as described in Example 2. For proper cancellation of the echo, taps at 2 or 3 times the echo delay may be needed. This implies that the echo must be 2 or 3 times shorter than the equalizer length.

Here are some examples of a few micro-reflection scenarios in potential HFC plant configurations. The first scenario is a micro-reflection that occurs between an amplifier and a feeder tap that are separated by 75 feet (150 feet round trip distance) and that have a return loss of 6 dB on each reflection interface (interface where impedance mismatch occurred). The feeder cable between these interfaces has a diameter of 0.625" and an attenuation of 1.2 dB/300 feet. This is considered a strong and short micro-reflection. Figure 10 shows the level and delay of the third transit and its subsequent multiple transit echoes.

The second scenario is a micro-reflection that occurs between two amplifiers and there are no feeder taps in between. They are separated by 1200 feet (2400 feet round trip distance) and that have a return loss of 8 dB on each reflection interface (interface where impedance mismatch occurred). The feeder cable between these interfaces has a

diameter of 1.000" and an attenuation of 0.8 dB/300 feet. This is considered a strong and long-delay micro-reflection. Figure 10 shows the level and delay of the third transit echo (top large blue square). The fifth transit echo of this micro-reflection is too low in amplitude to be noticeable.

The third scenario is a micro-reflection that occurs between two amplifiers and with feeder taps in between. The amplifiers are separated by 1200 feet (2400 feet round trip distance) and each has a return loss of 8 dB on each reflection interface (interface where impedance mismatch occurred). The aggregate insertion loss (also called through loss) in the feeder taps is equal to 6 dB (12 dB round trip). The feeder cable between the interfaces has a diameter of 1.000" and an attenuation of 0.8 dB/300 feet. This is considered a mild and long-delay micro-reflection. Figure 10 shows the level and delay of the third transit echo. The fifth transit echo of this micro-reflection is too low in amplitude to be noticeable.

Figure 10 also indicates which scenarios can be compensated in the different DOCSIS configurations. The scenarios that lie to the left of the vertical line that corresponds to a given channel width /DOCSIS mode combination can be compensated, while the ones that lie to the right of the line cannot be properly compensated. It is also worth noting that in cases close to the vertical line, higher order modulation may not be possible.

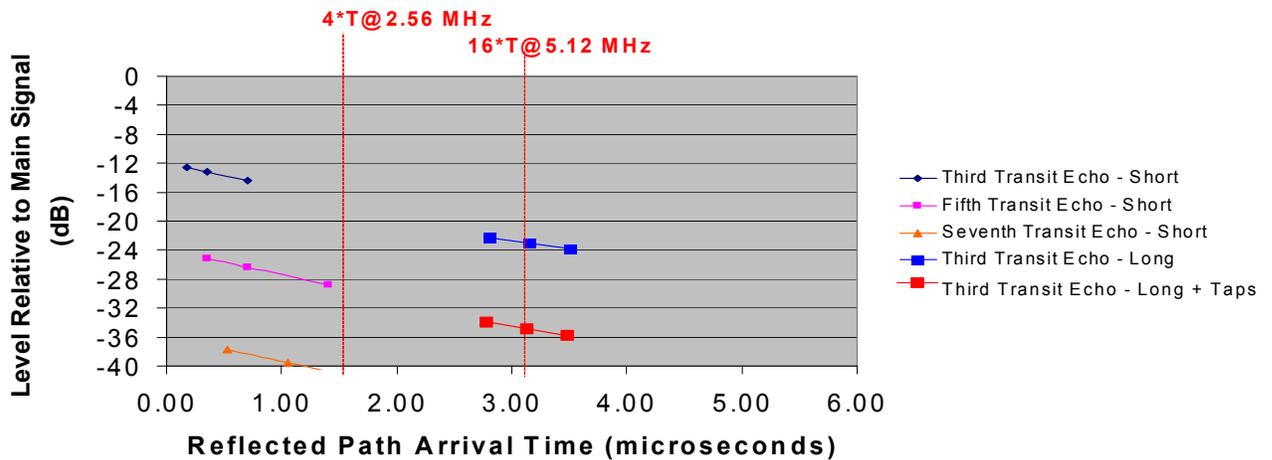


Figure 10 - Pre-equalization Compensation Capabilities under Short and Long Delay Micro-reflection Scenarios

The examples just discussed assumed 0.625" cable for the short time delay reflection and 1" cable for the long time delay reflection. The short time delay reflection scenario includes data points at 150', 300', and 600' round trip distances and the long time delay reflection includes data points at 2400', 2700', and 3000'.

5.6 DOCSIS Pre-equalization MIBs

DOCSIS pre-equalization coefficients indicate different things depending whether the CMTS or the CM is being queried. The information that is available through MIBs relate to what the CMTS and CM keep track of at the time the respective devices are being queried. Through the ranging interaction discussed in Section 5.3.1, the CMTS MIB (docsIfCmtsCmStatusEqualizationData) provides the adjustment necessary to update the CM coefficients and achieve upstream path distortion compensation. The CM MIB (docsIfCmStatusEqualizationData) indicates the current pre-distortion that is applied to the upstream signals.

The MIB format is as follows:

Bits	0	15	31
	Main Coeff.	Coeff/Symb	# Coeff
	F1real		F1imag
	F2real		F2imag
	⋮		
	F24real		F24imag

Figure 11 - MIB Format for docslfCmStatusEqualizationData

5.6.1 DOCSIS 2.0 and 3.0 Pre-equalization MIBs

Pre-equalization data is relevant in CM-CMTS channel combinations. In DOCSIS 2.0 the CM supports a single upstream channel, meaning the CM and CMTS reports a single pre-equalization data value. In DOCSIS 3.0 the pre-equalization data is measured for each of the upstream channels of the CM. To accommodate DOCSIS 3.0, the RFI management requirements were changed and DOCSIS 2.0 and 3.0 have separate MIBs for pre-equalization measurements. An additional per channel pre-equalization data measurement is also available and briefly discussed in this section to avoid confusion on usage of the appropriate information.

5.6.1.1 Per CM Pre-equalization

Table 2 presents semantically identical management objects for CM and CMTS for both DOCSIS 2.0 and 3.0.

Table 2 - DOCSIS 2.0 and 3.0 Transmit Pre-equalization MIBs

	DOCSIS 2.0 DOCS-IF-MIB [RFC4546]		DOCSIS 3.0 DOCS-IF3-MIB [DOCSISv3.0]	
	MIB Table	MIB Object	MIB Table	MIB Object
CM	docslfCmStatusTable Single upstream channel	docslfCmStatusEqualizationData	docslf3CmStatusUsTable Per configured upstream channel	docslf3CmStatusUsEqData
CMTS	docslfCmtsCmStatusTable Single instance per CM	docslfCmtsCmStatusEqualizationData	docslf3CmtsCmUsStatusTable Per CM, per configured upstream channel	docslf3CmtsCmUsStatusEqData

5.6.1.2 Per Interface Pre-equalization

For the original RFI MIB requirements there is a per interface pre-equalization data element that is common for DOCSIS 2.0 and 3.0, the docslfSigQEqualizationData from the docslfSignalQualityTable.

For the CM this data provides equalization information of the downstream receiver at the CM as defined in the DOCSIS DRFI specification. For the CMTS this object was intended to report some type of aggregated equalization value for the entire upstream channel. RFI MIB [RFC4546] clarifies the CMTS does not need to report a value other than an empty string.

Note that this equalization data is not relevant to the scope of this document.

6 METHODOLOGY

6.1 General Approach and Processes

The proactive network maintenance methodology that is based on pre-equalization coefficients can be described in terms of a few key general components.

The first general component is the data collection process. It comprises polling all CMs and CMTSs to obtain pre-equalization data from all configured upstream channels. The gathered data is verified for format integrity and is normalized to be useful for comparison. For scalability purposes, the data collection process is conducted using a more frequent polling cycle for the CMs that exhibited apparent distortion above a pre-determined level and a less frequent cycle for all CMs.

The second general process incorporates the initial distortion assessment that is conducted on all CMs that are monitored more frequently. This process uses the non-main tap to total energy (NMTER) ratio to discriminate which CMs should be examined in more detail and which should be left for evaluation in the next coarse monitoring cycle.

The third component in this approach conducts the detail analysis that includes the calibration process and the determination of the distortion signatures from frequency domain and time domain analysis. These signatures include group delay and micro-reflections. In case of multiple different micro-reflections, the signatures are obtained after a discrimination process.

The fourth component takes the distortion signatures and evaluates whether from a static perspective they should be classified as red which implies the need for immediate action, or as yellow which indicates the CM should be monitored more frequently and its distortion data be stored for observation over time. The information describing which CMs have to be examined more frequently is communicated to the data collection process. Green classification indicates that no action is necessary.

The fifth process takes the CM signatures and identifies within a fiber node's service area which micro-reflections are common to several CMs. The next process identifies by comparing historical data collected in the yellow classified CMs whether intermittent issues or trending issues are of concern and may require action.

The last process is the one that correlates the affected CM or CMs with the outside plant topology and uses that information to determine fault location. Figure 12 shows a diagram of the process just described.

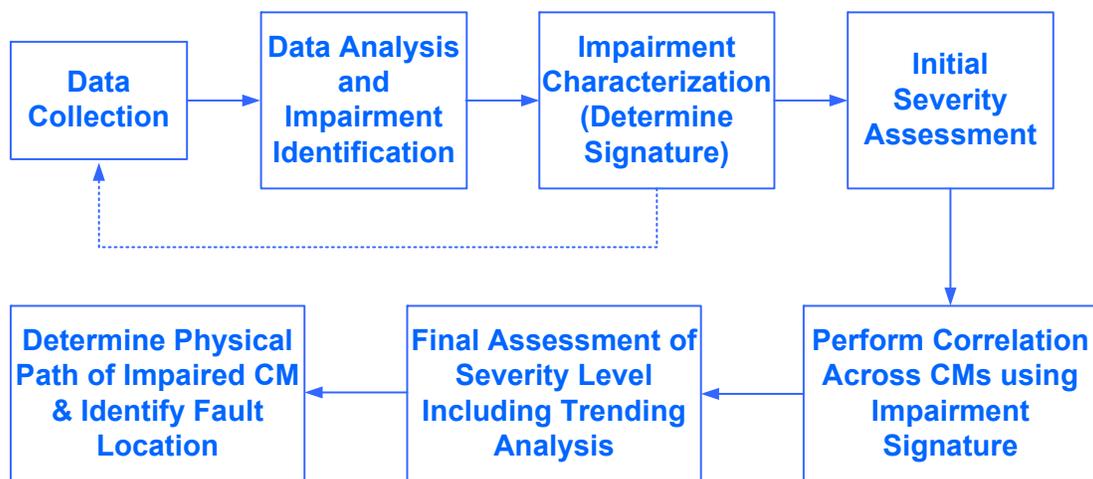


Figure 12 - Proactive Network Maintenance Processes based on Pre-equalization

6.2 Format Verification, Normalization and Guidelines

The structure of the pre-equalization information has been described in Section 5.6. How the values within this structure are interpreted depends on implementation. The first four byte-long elements in the header are to be interpreted in HEX mode. For example, the number of adaptive equalizer taps value of 18 in HEX is 24 in decimal (Figure 13). The rest of the equalizer structure defined in two byte increments containing the real and imaginary coefficients should be interpreted according to 2's complement over the entire two bytes or 4 nibbles describing the real or the imaginary coefficients. For example, the 2's complement of a 2 byte such as the fourth real coefficient is FF FC which in 2's complement is -4 (red circle).

MainT	T/Symb	#Taps	Rsvd	F1r	F1i	F2r	F2i	F3r	F3i	F4r	F4i
08	01	18	00	FF FF	00 02	FF FF	00 01	00 03	FF FF	FF FC	00 00
				-1	2	-1	1	3	-1	-4	0

F5r	F5i	F6r	F6i	F7r	F7i	F8r	F8i	F9r	F9i
00 0B	FF FF	FF EE	00 04	00 21	FF FB	07 FE	00 31	FF F3	FF E8
11	-1	-18	4	33	-5	2046	49	-13	-24

F10r	F10i	F11r	F11i	F12r	F12i	F13r	F13i	F14r	F14i
00 18	FF F0	FF F4	00 05	00 09	FF FB	FF FA	00 01	00 06	FF FD
24	-16	-12	5	9	-5	-6	1	6	-3

F15r	F15i	F16r	F16i	F17r	F17i	F18r	F18i	F19r	F19i
FF FD	00 00	00 05	FF FD	FF FD	00 01	00 01	00 01	FF FE	00 00
-3	0	5	-3	-3	1	1	1	-2	0

F20r	F20i	F21r	F21i	F22r	F22i	F23r	F23i	F24r	F24i
FF FF	00 00	FF FF	FF FF	00 00	FF FF	00 00	FF FF	FF FE	FF FD
-1	0	-1	-1	0	-1	0	-1	-2	-3

Figure 13 - Equalizer Structure HEXDECIMAL-to-DECIMAL Conversion

The representation of coefficients often differs among CM vendors. There are variations in maximum amplitude as well as variations in the way the coefficients are interpreted. Table 3 highlights the different interpretations that exist for the most popular CMs deployed.

Table 3 - Maximum amplitude and encoding formats for the 16 most popular 2.0 CMs in US

CM	Maximum Amplitude	Encoding Format
1	2047	4 Nibbles 2's Complement
2	2047	3 Nibbles 2's Complement
3	2047	4 Nibbles 2's Complement
4	511	4 Nibbles 2's Complement
5	511	4 Nibbles 2's Complement
6	511	4 Nibbles 2's Complement
7	2047	3 Nibbles 2's Complement
8	2047	4 Nibbles 2's Complement
9	2047	4 Nibbles 2's Complement
10	2047	3 Nibbles 2's Complement
11	1023	3 Nibbles 2's Complement

CM	Maximum Amplitude	Encoding Format
12	511	4 Nibbles 2's Complement
13	2047	4 Nibbles 2's Complement
14	2047	3 Nibbles 2's Complement
15	2047	4 Nibbles 2's Complement
16	2047	3 Nibbles 2's Complement

The CM vendors had two interpretations of how to decode the coefficients. One is the four nibble 2's complement interpretation and the other is the three nibble 2's complement interpretation. The four nibble 2's complement interpretation is the one assumed by the spec but there is a significant number of CMs deployed with the three nibble 2's complement interpretation. Regarding maximum amplitude, CMs have maximum amplitude equal to 2047, 1023 or 511. If the coefficients are normalized, the difference in CMs' maximum coefficient amplitude turns into a difference in granularity. The difference then becomes one of decoding interpretation of coefficients.

6.2.1 Four Nibble 2's Complement Pre-equalization Coefficient Representation

In this real or imaginary coefficient representation the entire four nibbles (two bytes) are used. This means that if the first bit is 0 the rest of the bits represent a positive integer binary number. If the first bit is 1, it is a negative 2's complement number. The actual value can be calculated by inverting the bits and adding 1, resulting in the negative of the number. Both positive and negative examples have been included below.

a) 000D = 0000 0000 0000 1101 = 13

b) FFFE = 1111 1111 1111 1110

after inverting it and adding 1 → 0000 0000 0000 0001 + 1 = 0000 0000 0000 0010 = 2

→ The number is - 2

6.2.2 Three Nibble 2's Complement Pre-equalization Coefficient Representation

In this real or imaginary coefficient representation the last three nibbles out of the two coefficient bytes are used for the calculation. In this interpretation the first nibble is 0 which could erroneously lead one to believe that all coefficients are positive. Only after eliminating the first nibble one can tell if a number is positive or negative. This means that if the fifth bit is 0 the rest of the bits represent a positive integer binary number. If the fifth bit is 1, it is a negative 2's complement number. The actual value can be calculated by inverting the bits of the three nibbles and adding 1. This is the negative of the number in question. Both positive and negative examples have been included below.

a) 000D = 0000 0000 0000 1101 → eliminate the first nibble → 0000 0000 1101 = 13

b) 0FFE = 0000 1111 1111 1110 → eliminate the first nibble → 1111 1111 1110

after inverting it and adding 1 → 0000 0000 0001 + 1 = 0000 0000 0010 = 2 → The number is - 2

6.2.3 Universal Decoding

Since in current implementations the maximum value that a coefficient can take is always less than or equal 2047, the first nibble is never used and can be removed to generate a universal decoder. After removing the first nibble, the decoding process would be identical to the third nibble process.

6.3 Key Metrics

The real and imaginary complex coefficients of a DOCSIS 2.0 upstream pre-equalizer defined as:

$$F1_R, F1_I, F2_R, F2_I, F3_R, F3_I, F4_R, F4_I, \dots, F23_R, F23_I, F24_R, F24_I,$$

and will be used to define several key metrics that follow.

6.3.1 Adaptive Equalizer Main Tap Energy

The adaptive equalizer main tap in DOCSIS 2.0 is typically in tap position eight although some CMTS implementations can have it in as low as the sixth position. In DOCSIS 1.1 the main tap is in the fourth position. The equalizer tap energy is given by the sum of the squares of the real and imaginary components of the coefficient.

The main tap energy (MTE), assuming it in eighth position, is defined as:

$$MTE = F8_R^2 + F8_I^2$$

6.3.2 Main Tap Nominal Energy and Main Tap Nominal Amplitude

The DOCSIS pre-equalization taps exhibit different nominal or maximum amplitudes depending on CM implementations. The maximum of amplitude implementations from CMs are 2047, 1023 or 511. This parameter is defined here as the main tap nominal amplitude (MTNA). The square of the nominal amplitude yields the nominal tap energy.

The main tap nominal energy (MTNE), assuming main tap is in the eighth position, is defined as:

$$MTNE = F8_{nominal_R}^2 + F8_{nominal_I}^2$$

6.3.3 Pre-Main Tap Energy

The summation of the energy in all equalizer taps prior to the main tap provides the pre-main tap energy (PreMTE).

The pre-main tap energy assuming a main tap in the eighth position is defined as:

$$PreMTE = F1_R^2 + F1_I^2 + F2_R^2 + F2_I^2 + F3_R^2 + F3_I^2 + \dots + F7_R^2 + F7_I^2$$

6.3.4 Post-Main Tap Energy

The summation of the energy in all equalizer taps after the main tap provides the post-main tap energy (PostMTE).

The post-main tap energy assuming a main tap in the eighth position is defined as:

$$PostMTE = F9_R^2 + F9_I^2 + F10_R^2 + F10_I^2 + \dots + F23_R^2 + F23_I^2 + F24_R^2 + F24_I^2$$

6.3.5 Total Tap Energy

The summation of the energy in all equalizer taps provides the total tap energy (TTE).

The total tap energy is defined as:

$$TTE = \text{PreMTE} + \text{MTE} + \text{PostMTE}$$

6.3.6 Main Tap Compression

Adaptive equalizer main tap compression (MTC) at the CM is a good indicator of the available margin for the continued reliance on the equalization compensation process. An MTC ratio greater than 2 dB may suggest that equalization compensation can no longer be successfully achieved. This metric is given by the ratio of the energy in all taps to the main tap energy.

The main tap compression expressed in dB is defined as:

$$MTC = 10\text{Log}\left(\frac{TTE}{MTE}\right)$$

Main tap compression at the CM translates to a less RF power level delivered to the CMTS. An MTC of 2 dB results in the CMTS receiving 2 dB less input power.

Main tap compression at the CMTS is not expected under normal operating conditions. Any level of main tap compression at the CMTS should raise an alarm.

6.3.7 Non-Main Tap to Total Energy Ratio (Distortion Metric)

The adaptive equalizer's non-main tap to total energy ratio is a useful "distortion metric" to determine the distortion level in the upstream path. This parameter can be used as an initial assessment tool to determine which CMs need to be examined further and more frequently. This distortion metric is given as the ratio of the aggregate energy that exists in all but the main tap to the energy in all of the adaptive equalizer's taps.

The non-main tap to total energy ratio expressed in dB is defined as:

$$NMTER = 10\text{Log}\left(\frac{\text{PreMTE} + \text{PostMTE}}{TTE}\right)$$

Notice that the main tap energy in the numerator is missing. The non-main tap energy ratio is also a good estimation of the MER assuming that the signal is not impacted by impairments that are not considered linear distortions, such as burst noise and nonlinear impairments.

Non-main tap to total energy ratio at the CMTS is a good indicator of the type of upstream performance the CM signals have based on the amount of linear distortion present. If a 27dB CNR is assumed for negligible errors with a 64-QAM signal, a NMTER target value of -27 dB can be assumed for comparable performance. If a 30 dB CNR is the threshold where correctable errors are beginning to appear, that would also correspond to a threshold of -30 dB NMTER when correctable errors begin to appear. This CNR to NMTER relationships are useful in determining thresholds from the NMTER values. An operator could assume an immediate action (red) NMTER threshold of -27dB for 64-QAM operation and a monitor more frequently (yellow) NMTER threshold of -30 dB.

6.3.8 Pre-Main Tap to Total Energy Ratio

The adaptive equalizer's pre-main tap to total energy ratio (PreMTTER) is a useful parameter along with the adaptive equalizer's pre-post tap symmetry to determine the group delay level in the upstream path. This distortion metric is the ratio of the pre-main tap energy to the energy in all taps.

The pre-main tap to total energy ratio expressed in dB is defined as:

$$PreMTTER = 10\text{Log}\left(\frac{PreMTE}{TTE}\right)$$

6.3.9 Post-Main Tap to Total Energy Ratio

The adaptive equalizer's post-main tap to total energy ratio (PostMTTER) is a useful parameter to assess micro-reflection impairment contribution. This distortion metric is the ratio of the post-main tap energy to the energy in all taps.

The post-main tap to total energy ratio expressed in dB is defined as:

$$PostMTTER = 10\text{Log}\left(\frac{PostMTE}{TTE}\right)$$

6.3.10 Pre-Post Energy Symmetry Ratio

The pre-post energy symmetry ratio (PPESR), along with pre-main tap to total energy ratio, is a useful parameter to indicate the presence of group delay in the upstream path. This distortion metric is the ratio of the post to pre-main tap energy ratios.

The pre-post energy symmetry ratio expressed in dB is defined as:

$$PPESR = 10\text{Log}\left(\frac{PreMTE}{PostMTE}\right)$$

For practical purposes, the pre-post energy symmetry may be approximated using only the two taps adjacent to the main tap, giving the pre-post tap symmetry ratio (PPTSR):

$$PPTSR = 10\text{Log}\left(\frac{F7_R^2 + F7_I^2}{F9_R^2 + F9_I^2}\right)$$

6.3.11 Group Delay Distortion

Group delay distortion is a type of linear distortion that is also corrected by the pre-equalization process. Figure 14 shows group delay increase with frequency for increasing number of cascaded actives. This configuration does not include the impact of chokes in the lower portion of the spectrum (typically 5-10 MHz) and it does not assume additional in-line equalizers.

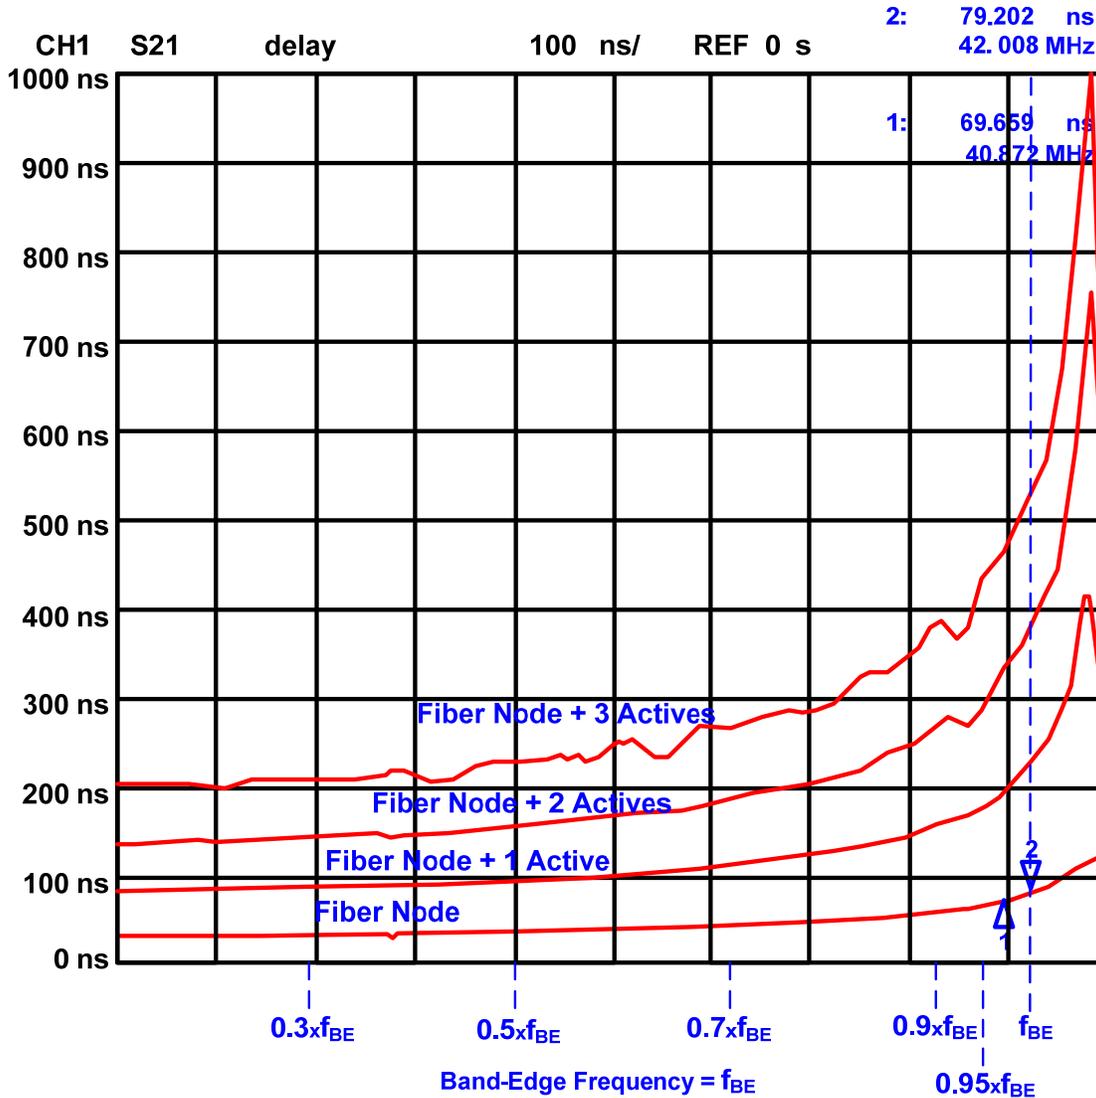


Figure 14 - Group Delay increase with increasing cascade depth

As observed in Figure 14, there is a variation in group delay of up to 300 ns in a channel that is located at the band-edge. The pre-tap equalization coefficient energy increases in the presence of group delay distortion. Additional group delay details are discussed in the tutorial section (Appendix I).

6.3.11.1 Group Delay Distortion at Band-Edge with No Micro-reflection

Table 4, Figure 15, and Table 5 relate to the scenario where the DOCSIS channel is at the band-edge and there is no micro-reflection present. Table 4 shows the impact of group delay distortion on tap energy when the upstream channel is at the band-edge and no micro-reflection is present.

Table 4 - Band-Edge Operation Impact on Tap Energy (no Micro-reflections)

Tap (Coeff)#	Pre-Main Tap Coefficients							Post-Main Tap Coefficients																	
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	
Cascade	N+5	-49.1	-43	-41.6	-35.4	-31.7	-25.9	-17.1	-0.11	-26.1	-43	-45.1	-53.2	-53.2	-60.2	-57.2	-57.2	-49.1	-53.2	-53.2	-50.7	-54.2	-53.2	-53.2	
	N+4	-51.2	-44.1	-40.7	-38.2	-32.6	-27.1	-18.2	-0.08	-31	-47.2	-40.9	-49.1	-47.9	-60.2	-54.2	-53.2	-57.2	-100	-53.2	-49.1	-57.2	-53.2	-50.2	-57.2
	N+3	-60.2	-47.2	-44.9	-37.7	-34.2	-28.5	-19.8	-0.06	-38.1	-34.9	-44.2	-60.2	-46.2	-57.2	-60.2	-54.2	-53.2	-60.2	-50.2	-53.2	-53.2	-47.2	-50.2	-100
	N+2	-57.2	-47.9	-44.9	-38.6	-36.4	-30.7	-22.3	-0.03	-34.2	-40.3	-46.2	-54.2	-50.7	-60.2	-60.2	-60.2	-45.6	-53.2	-51.2	-47.2	-46.1	-57.2	-57.2	-57.2
	N+1	-60.2	-54.2	-50.2	-44.6	-40.7	-34.2	-26	-0.01	-41.9	-49.1	-54.2	-60.2	-50.7	-51.2	-57.2	-60.2	-49.1	-50.7	-49.1	-57.2	-57.2	-51.2	-100	-49.1
	N+0	-60.2	-57.2	-60.2	-46.2	-44.6	-39.6	-30.2	-0.01	-47.2	-46.2	-46.2	-57.2	-54.2	-60.2	-53.2	-49.1	-57.2	-50.2	-60.2	-60.2	-57.2	-53.2	-53.2	-100

Table Note: No micro-reflections, 40.4 MHz center frequency, 3.2 MHz channel width

Figure 15 shows how when operating at the band-edge, the pre-main tap energy increases proportionally with increasing cascade depth. This is characteristic of the group delay distortion impact on tap energy. The effect of group delay distortion could be hidden just by looking at cascade depth in a plant topology map. In-line equalizers which may not be obvious in a plant topology diagram may contribute to distortion just as diplexers within amplifiers do. In severe distortion cases, assessment of pre-main tap energy could be used to determine whether certain CMs should be moved to lower distortion channels in the middle of the upstream band.

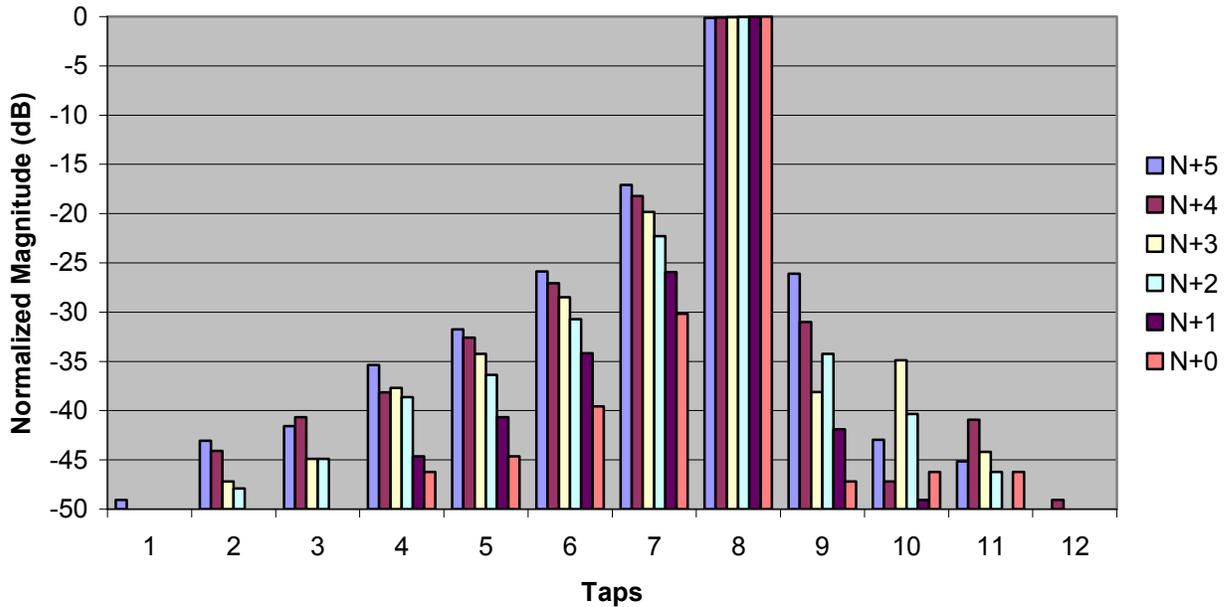


Figure 15 - Pre Main Tap Energy Increase with Cascade Depth (Fc=40.4 MHz, Ch. W=3.2 MHz, No Micro-reflections, first 12 taps shown)

Key metrics worth highlighting in Table 5 are the PreMTTER that increases with cascade depth while the PostMTTER shows low values as no micro-reflections are present. The PPESR show high positive values since in this scenario the dominant impairment is group delay distortion. The low values of MTC are indicative that the pre-equalization compensation is effective.

Table 5 - Pre-equalization Metrics at Band-Edge (No Micro-reflections)

Cas cade		MTC	NMTER	PreMTTER	PostMTTER	PPESR
N+5		0.11 dB	-15.9 dB	-16.3 dB	-25.9 dB	9.53 dB
N+4		0.08 dB	-17.3 dB	-17.5 dB	-30.1 dB	12.7 dB
N+3		0.06 dB	-18.9 dB	-19.1 dB	-32.2 dB	13.1 dB
N+2		0.03 dB	-21.1 dB	-21.4 dB	-32.2 dB	10.8 dB
N+1		0.01 dB	-24.9 dB	-25.1 dB	-38 dB	12.9 dB
N+0		0.01 dB	-29 dB	-29.4 dB	-39.4 dB	9.93 dB

6.3.11.2 Group Delay Distortion at Band-Edge with 0.5 μ s Micro-reflection

Table 6, Figure 16, and Table 7 relate to the scenario where the DOCSIS channel is at the band-edge but there is also a 0.5 μ s micro-reflection present. Table 6 shows the impact of group delay distortion on tap energy.

Table 6 - Band-Edge Operation Impact on Tap Energy (with 0.5 μ s Micro-reflection)

Tap (Coeff) #	Pre-Main Tap Coefficients								Post-Main Tap Coefficients															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24
N+5	-60.2	-47.9	-43.9	-37.5	-31.2	-24.8	-16.5	-0.48	-14.4	-14.5	-31.7	-24.7	-32.5	-32.9	-35.5	-37	-41.4	-43.9	-50.2	-57.2	-50.7	-57.2	-60.2	-49.1
N+4	-53.2	-57.2	-43.8	-38.5	-31.6	-25.8	-17.1	-0.46	-14.4	-14.3	-33.3	-25.1	-32.6	-33.2	-37.2	-38.7	-43.8	-46.4	-49	-46.9	-51.2	-53.2	-54.2	-60.2
N+3	-60.2	-50.7	-43.3	-37.9	-31.9	-26.7	-17.6	-0.46	-15.1	-13.6	-34.8	-24.6	-32.6	-32.2	-38.2	-37.1	-53.2	-47.2	-47.6	-50.2	-53.2	-47.2	-54.2	-54.2
N+2	-100	-50.9	-44.6	-41.3	-33.7	-28.2	-19.6	-0.38	-14.6	-14.9	-36.4	-25.7	-37.5	-34.7	-38.4	-42.5	-48.1	-52.1	-50.9	-59.2	-59.2	-50.1	-49.2	-66.2
N+1	-54.2	-50.2	-44.6	-39.3	-33.2	-27.5	-19.1	-0.43	-14.5	-14	-34.6	-24.7	-33.4	-32.7	-37.2	-37.2	-44.1	-47.9	-43	-53.2	-50.7	-44.2	-44.9	-53.2
N+0	-59.2	-52.2	-47.1	-40.8	-35.1	-29.7	-21.7	-0.4	-14	-14.5	-35.3	-25.5	-37.9	-34.9	-40.5	-40.6	-47.2	-52.1	-48.6	-63.2	-56.7	-46.2	-57.2	-47.1

Table Note: 0.5 microsecond delay micro-reflection, 40.4 MHz center frequency, 3.2 MHz channel width

Figure 16 shows how when operating at the band-edge, the pre-main tap energy increases with increasing cascade depth but not as noticeably as the scenario without micro-reflections depicted in Figure 15. This is due to some pre-main tap energy needed to compensate for the micro-reflection which adds to the pre-main tap energy that is caused by group delay distortion. This leaking of energy into the pre-main tap region is more prevalent with shorter micro-reflections that use the taps closer to the main tap for compensation than the longer micro-reflections that use the higher value taps.

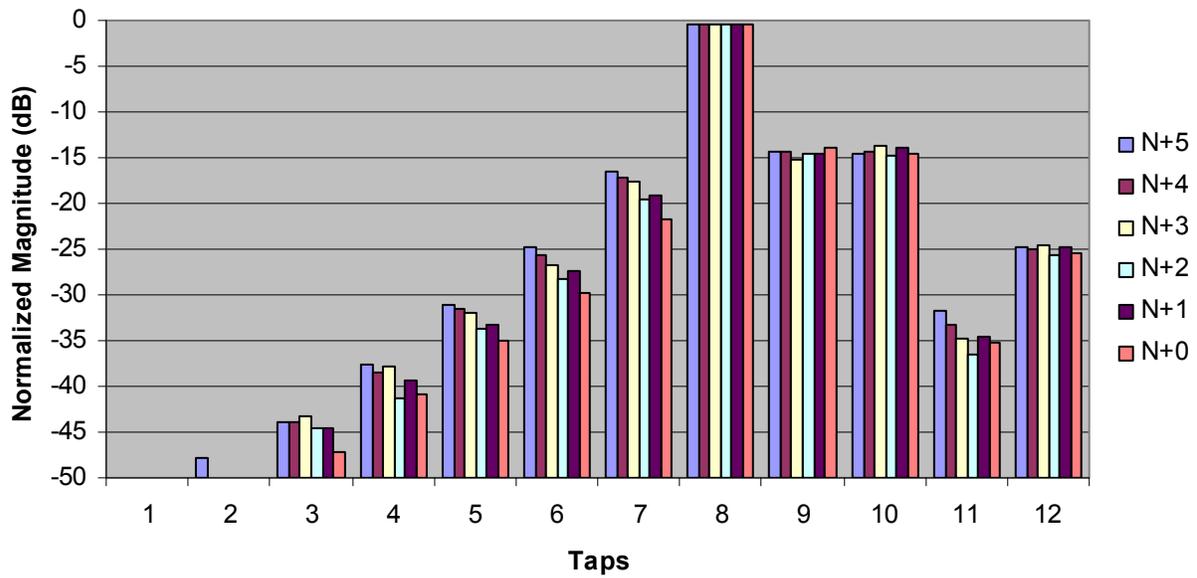


Figure 16 - Pre Main Tap Energy Increase with Cascade Depth
(Fc=40.4 MHz, Ch. W=3.2 MHz, with 0.5 μs Micro-reflection, first 12 taps shown)

Key metrics worth highlighting from Table 7 are the PreMTTER that increases with cascade depth. The PostMTTER shows a high value indicative of the micro-reflection present. The PPESR show negative values since in this scenario the dominant impairment is micro-reflection. The combined group delay distortion and micro-reflection by themselves is properly compensated through the pre-equalization process, although some increase in MTC begins to show.

Table 7 - Pre-equalization Metrics at Band-Edge (with 0.5 μs Micro-reflection)

		MTC	NMTER	PreMTTER	PostMTTER	PPESR
Cascade	N+5	0.48 dB	-9.84 dB	-15.8 dB	-11.1 dB	-4.65 dB
	N+4	0.46 dB	-9.94 dB	-16.4 dB	-11.1 dB	-5.36 dB
	N+3	0.46 dB	-10 dB	-16.9 dB	-11 dB	-5.95 dB
	N+2	0.38 dB	-10.8 dB	-18.9 dB	-11.5 dB	-7.4 dB
	N+1	0.43 dB	-10.2 dB	-18.3 dB	-10.9 dB	-7.35 dB
	N+0	0.4 dB	-10.6 dB	-20.8 dB	-11 dB	-9.83 dB

6.3.11.3 Distortion in Middle of Upstream Band with Micro-reflection

Table 8, Figure 17, and Table 9 relate to the scenario where the DOCSIS channel is operating in the middle of the upstream band with a 0.5 μs micro-reflection present. Table 8 shows the impact of distortion on tap energy.

Table 8 - Micro-reflection Impairment on Pre and Post Main Tap Energy

Tap (Coeff)#	Pre-Main Tap Coefficients							Post-Main Tap Coefficients																	
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	
Cascade	N+5	-60.2	-54.2	-46.2	-40.2	-33	-27.2	-18.1	-0.8	-11.9	-11.6	-22	-21.8	-27.7	-30.5	-31.8	-36.2	-39.4	-40.9	-41.1	-60.2	-44.6	-53.2	-54.2	-51.2
	N+4	-59.2	-60.2	-46.2	-41.1	-34.8	-29.2	-20.7	-0.8	-11.2	-11.8	-21.2	-21.9	-28.2	-31.8	-33.4	-39.4	-40.5	-45.5	-45.1	-56.2	-50.9	-51.2	-53.9	-47.2
	N+3	-59.2	-53.2	-47.6	-41.6	-35.1	-30.3	-22.3	-0.81	-10.9	-11.8	-21.1	-22.5	-29	-33.1	-35.3	-41.4	-42.1	-50.9	-48.4	-50.2	-63.2	-53.9	-57.2	-59.2
	N+2	-59.2	-55.1	-46.2	-41.7	-34.2	-28.5	-20.4	-0.8	-11.3	-11.8	-21	-22.1	-28.2	-32.2	-33.7	-39.2	-39.8	-44.6	-49	-46.7	-50.9	-48.6	-47.1	-47.6
	N+1	-100	-56.2	-45.4	-42	-34	-29.1	-20.6	-0.83	-11	-11.7	-21.6	-21.8	-28.6	-31.5	-33.4	-39.1	-40.2	-42.8	-57.2	-53.2	-56.2	-49	-63.2	-46.4
	N+0	-57.2	-60.2	-47.2	-41.6	-32.9	-27.8	-18.7	-0.79	-11.8	-11.6	-22.2	-21.7	-28.7	-31.2	-33.1	-39.6	-37.6	-47.9	-42.6	-50.2	-49.1	-60.2	-57.2	-57.2

Table Note: 0.5 microsecond delay micro-reflection, 14 MHz center frequency, 3.2 MHz channel width

Figure 17 shows how when operating in the middle of the band, there is no increase in pre-main tap energy with increasing cascade depth as the group delay is fairly flat. (See also Figure 14.) Figure 17 shows how post-main tap energy is used in compensating for the micro-reflection. Figure 17 also illustrates how a small amount of pre-main tap energy is used in compensating for the fractional delay 0.5 μs micro-reflection. When the micro-reflection delay doesn't coincide with a tap delay, neighboring taps are used for compensation. This explains why in Figure 16 the increase of pre-main tap energy with increasing cascade depth is not as noticeable as in the “group delay only” case depicted in Figure 15. The combined effect of group delay distortion and a 0.5 μs micro-reflection of Figure 16 can be approximated to a rough superposition of Figure 15 representing the “group delay only” case and Figure 17 representing the “micro-reflection only” case.

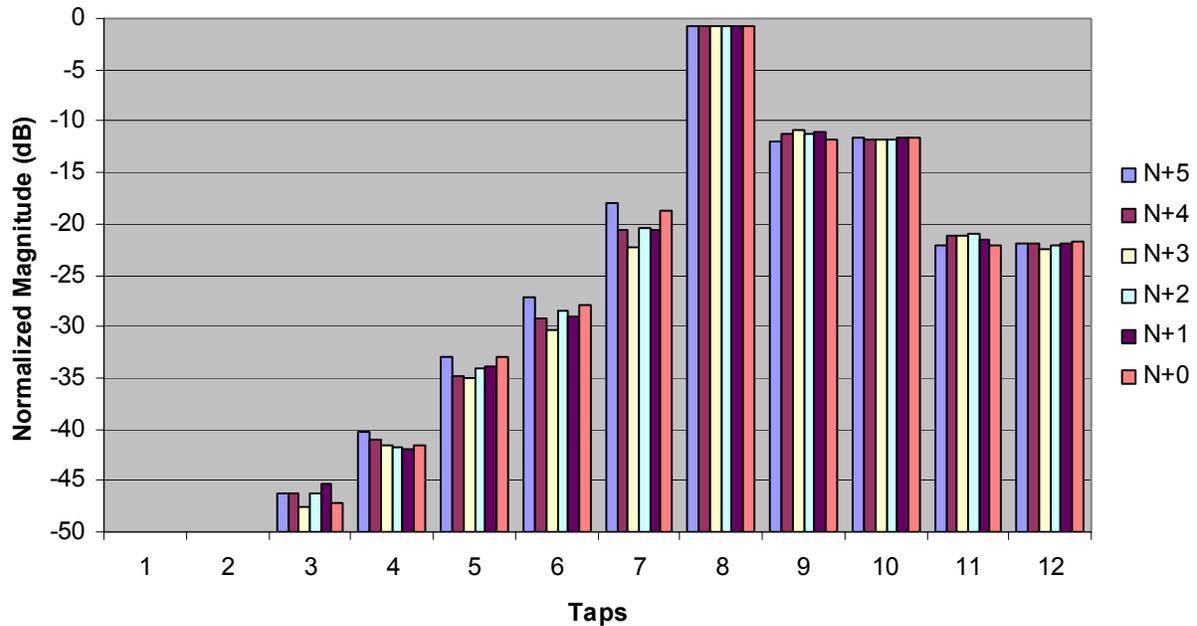


Figure 17 - Tap Energy for Different Cascade Depth Scenarios (Fc=14 MHz, Ch. W=3.2 MHz, with 0.5 μs Micro-reflection, first 12 taps shown)

Some key metrics from Table 9 are worth highlighting. The low PreMTTER value is indicative of negligible group delay distortion. The PostMTTER high value is indicative of the strong micro-reflection present. The PPSR show negative values since in this scenario the dominant impairment is the micro-reflection. The MTC value is higher than the previous two scenarios which indicates that the pre-equalizer may be beginning to lose its equalization compensation effectiveness.

Table 9 - Pre-equalization Metrics at Middle of Upstream Band (with 0.5 μs Micro-reflection)

Cascade		MTC	NMTER	PreMTTER	PostMTTER	PPESR
N+5		0.8 dB	-7.73 dB	-17.4 dB	-8.23 dB	-9.18 dB
N+4		0.8 dB	-7.72 dB	-19.9 dB	-7.99 dB	-11.9 dB
N+3		0.81 dB	-7.68 dB	-21.4 dB	-7.87 dB	-13.5 dB
N+2		0.8 dB	-7.72 dB	-19.6 dB	-8.01 dB	-11.6 dB
N+1		0.83 dB	-7.6 dB	-19.8 dB	-7.87 dB	-11.9 dB
N+0		0.79 dB	-7.78 dB	-18.1 dB	-8.21 dB	-9.84 dB

6.4 DOCSIS Pre-equalization Coefficient Data Collection

DOCSIS pre-equalization coefficient information is obtained by using simple network management protocol (SNMP). The SNMP management information base (MIB) describes information which is available in a standard manner across all implementations that comply with the DOCSIS MIB definitions.

At a minimum, there are three data elements required to complete the mathematical transformations described in this document. The three values are the DOCSIS equalization coefficients, upstream center channel frequency, and the RF bandwidth of the channel. The three data points are obtained from the cable modem and the CMTS. Having the cable modem coefficients provides a view of the inverted channel response prior to equalization. The CMTS offers a post-equalization view of the inverted channel response as it is received from the cable modem. This can be helpful when evaluating the performance of the modems’ upstream pre-equalizers. It can also help identify issues where the impairment is changing frequently.

In addition to the three required data elements, there are other metrics which can be used to help identify and localize areas of impairment. These other elements may already be available as a result of pre-existing monitoring systems. One should extend or reuse existing data as opposed to over-polling the modems or CMTS.

6.4.1 SNMP Implementation and Performance Considerations

Large SNMP implementations can represent unique challenges when scale is considered. First, assuming DOCSIS 2.0 pre-equalization coefficients and their storage as characters, the minimum over-the-wire network impact is 614 bytes per modem/CMTS pair. This is calculated using the PDU sizes of the minimum required data elements. In this case, both the modem and CMTS coefficients are 299 bytes while the frequency and RF bandwidth together are 16. Assuming binary representation the storage requirements are 100 bytes each (4 bytes header + 96 bytes real and imaginary coefficients data) for CMTS and CM and 4 bytes each for upstream frequency and channel width totaling a minimum of 208 bytes.

Example:

docsIfCmStatusEqualizationData (299 bytes character storage, 100 bytes binary storage)

```
08 01 18 00 00 04 ff fd ff fb ff fa ff fd ff fd 00 07 00 04 ff f8 00 00 00 17 ff ff ff d6 ff e8 07 f7 ff f9 ff 8a ff 94 ff f7
00 28 00 11 ff ec ff f7 00 19 00 06 ff f5 ff fc ff ff 00 0d ff fb 00 01 00 01 00 04 00 04 ff f6 00 07 00 07 ff fb 00 00
00 08 ff ff ff fe 00 00 00 04 ff fc ff ff 00 08 00 00
```

docsIfCmtsCmStatusEqualizationData (299 bytes character storage, 100 bytes binary storage)

```
08 01 18 00 ff c8 ff f0 ff f8 ff e8 ff c0 00 20 ff f8 00 38 ff c8 ff d8 00 18 00 18 00 38 ff f0 3f 20 00
00 00 08 00 40 ff d8 ff e8 00 38 ff d0 ff d8 ff b8 00 08 ff d8 00 40 00 40 ff e8 00 00 00 00 ff c0 00 48
00 20 00 20 00 58 00 10 ff f0 00 00 ff e8 ff f0 ff d8 00 50 00 00 ff d8 ff f0 00 18 ff e8
```

UpStreamFrequency (8 bytes character storage, 4 bytes binary storage)

29300000

UpStreamWidth (8 bytes character storage, 4 bytes binary storage)

6400000

A payload of 614 bytes with a million devices translates to nearly 600 megabytes. At 5 million devices, a single poll represents almost 3 gigabytes. This traffic would be distributed between both the CMTS and the cable modems. Careful consideration should be made regarding the impact on routers, firewalls, traffic monitors, the CMTS, and other network elements.

In large implementations, heavily threaded/distributed SNMP processing will be required to accommodate the large number of requests. Assuming 250 milliseconds round trip query latency, 1 million devices would require 70 hours of CPU time if executed serially. Given that most of the CPU time is spent in wait state, a horizontally scaled execution can be very effective. Specifically, “horizontal scale” can be achieved by adding additional CPUs or polling nodes. When this same load is distributed across 1000 threads of process, that time is reduced to sub five minutes.

If the network topology is well known at the time of polling, a geographically sensitive polling approach might be considered. Also, market specific maintenance windows may be among the factors to consider in polling cycle timing.

6.4.2 Data Collection Strategy

Initially, three polling rates are discussed to accommodate different levels of analysis. First, a low-rate cycle should be used to acquire a daily baseline of data used for coarse grained analysis. This coarse grained analysis will provide several metrics calculated from the raw coefficient data prior to Fourier transformation. Analysis of the coarse metric will narrow the scope of devices which qualify for medium-rate polling. This secondary polling cycle will be used to drive further equalization analysis which includes a Fourier transformation to the frequency domain. The intent of the medium-rate polling is to insulate the data processing layer from unnecessarily processing modems that are currently operating within acceptable thresholds. Finally, the high-rate poll cycle will be used to provide low latency visibility to a small subset of devices. Given the expense associated with polling, storage and analysis, this stage should be reserved for only the most critical or otherwise material-interest devices.

6.4.2.1 Low Rate (once daily, rotating across eight-hour time-shifts, adjustable)

Modem based pre-equalization coefficient data should be collected once per day. It's possible that certain plant problems might be specific to the time of day. These scenarios might include weather patterns, watering systems and other time-based anomalies. Assuming an eight-hour window to complete polling (including accommodation for maintenance), a three shift pattern should work well. This could be achieved by a single daily poll with a time offset of 32 hours. Likewise, the modem population could be divided by three and the load distributed throughout the day. In the latter approach, it would be important to rotate the three modem populations to correctly achieve the desired result (Table 10 and Table 11).

CMTS based equalization coefficient data should be collected as close in time to the modems as possible. Understanding that the SNMP process for gathering data from the CMTS will be decoupled from the modems is important. The most efficient way to obtain this data from the CMTS would be a bulk walk (SNMP) of the docsIfCmtsCmStatusEqualizationData table. This will return a large swath of data, while the modem collection threads are completely independent and non-synchronized with the CMTS process.

Using the main tap compression and non-main tap to total energy ratio formulas (see Key Metrics, Section 6.3), modems of interest are identified and promoted to the medium-rate polling cycle. These will remain under medium-rate scrutiny until the correct threshold is met for some predetermined time. Initially, 48 hours is recommended.

Table 10 - Low Rate – Once Daily – Rotating Eight Hour Time Shifts - Three day cycle

Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8	Day 9
6 AM	2 PM	10 PM	6 AM	2PM	10 PM	6 AM	2 PM	10 PM

Table 11 - Low Rate for Three CM groups

Groups	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8	Day 9
1	6 AM	2 PM	10 PM	6 AM	2PM	10 PM	6 AM	2 PM	10 PM
2	2 PM	10 PM	6 AM	2PM	10 PM	6 AM	2 PM	10 PM	6 AM
3	10 PM	6 AM	2PM	10 PM	6 AM	2 PM	10 PM	6 AM	2 PM

6.4.2.2 Medium – Rate (every four hours, adjustable)

The secondary level of analysis will require greater resources to gather, store, and analyze, so sizing should be done based on preliminary results of the initial low-rate poll. This poll will handle the subset devices of interest. Devices may be promoted/demoted across low, medium and high rate polling as described in the analysis section of this document. Table 12 and Table 13 show medium rate polling examples.

Table 12 - Medium Rate – Once Every Four Hours - One day cycle

Day 1	Day 2	Day 2	Day 2	Day 2				
6 AM	10 AM	2 PM	6 PM	10PM	2 AM	6 AM	10 AM	2 PM

Table 13 - Medium Rate – Once Every Four Hours - One day cycle - Four Groups (All Times in EST)

Groups	Day 1	Day 2	Day 2	Day 2	Day 2				
EST	6 AM	10 AM	2 PM	6 PM	10PM	2 AM	6 AM	10 AM	2 PM
CST	7 AM	11 AM	3 PM	7 PM	11PM	3 AM	7 AM	11 AM	3 PM
MT	8 AM	12 PM	4 PM	8 PM	12AM	4 AM	8 AM	12 PM	4 PM
PST	9 AM	1 PM	5 PM	9 PM	1 AM	5 AM	9 AM	1 PM	5 PM

6.4.2.3 High – Rate (ten minute, adjustable)

Only devices of the highest interest will be candidates for high-rate polling. This will most often be the result of manual intervention to facilitate real-time troubleshooting or field analysis. In general, there will be limits placed on high-rate entries which automatically expire in a designated period of time

6.5 Calibration Mechanisms

6.5.1 CMTS-CM Short Reference Plant

Improved accuracy may be achieved by subtracting the contributions of the CMTS and cable modem(s) from the end-to-end channel response, leaving just the channel response signature of the cable network. Determining the channel response distortion contributions of these two components can be done using a short reference plant—sometimes called a six foot plant—which isolates the CMTS and cable modem(s) from the cable network. Adaptive pre-equalization coefficients can then be used to characterize the response signature of the combined devices, much the same as is done to characterize end-to-end pre-equalization signatures in an operating cable network. Once the CMTS and modem response signatures are known, they can be subtracted from operational field measurement results. It is recommended that this characterization be performed for each make/model CMTS and cable modem combination in use.

The short reference plant illustrated in Figure 18 has been measured with a vector network analyzer, and verified to provide optimum performance using the parts listed in Table 14. Choose a diplex filter that has a frequency split well above the normal operating upper frequency limit of the return path. This will reduce the group delay contribution of the diplex filter. For example, a 5-42 MHz DOCSIS return should use at least the 65/85 MHz split diplex filter in the short reference plant.

Interconnecting cables should be as short as practicable, and where it is possible to eliminate cables altogether, one should use appropriate male-male or female-female F-type splices. In any case, avoid exceeding the maximum cable lengths shown in Figure 18.

Table 14 - Tested parts for Short Reference Plant

Vendor	Part Number	Description
Eagle Comtronics	EDPF-65/85	Diplex filter, 65 MHz/85 MHz split
Eagle Comtronics	EDPF-88/108	Diplex filter, 88 MHz/108 MHz split
Holland Electronics	GHS-8	8-way splitter
Holland Electronics	GHS-2FC LI	2-way splitter, low intermod
Holland Electronics	GHS-2	2-way splitter
Holland Electronics	FAM-* HR	Fixed attenuator pads (3, 6, 8, 10, 12, 16, 20 dB)
Holland Electronics	DCG-20SB	20 dB directional coupler
Holland Electronics	F-59TH	Precision 75-ohm terminator
Holland Electronics	G-F81F	F female-female splice
Holland Electronics	F-71	F male-male splice

Bill of Materials

- 22 ea. precision male pin F connectors to match coax
- 2 ea. adapters, female F-to-connector type used on CMTS (if needed)
- 3 ea. male-male F splice
- 1 ea. 30 dB fixed attenuator
- 1 ea. 20 dB fixed attenuators
- 1 ea. 10 dB fixed attenuator
- 11 ea. 6 dB fixed attenuators
- 2 ea. precision 75-ohm male F terminators
- 50~300 feet quad-shield coax
- 1 ea. diplex filter
- 1 ea. 8-way splitter
- 1 ea. 20 dB directional coupler
- 1 ea. 2-way splitter
- 8 ea. DOCSIS 2.0 or later cable modems

Note: DOCSIS specifies that a CMTS must support +50 to +61 dBmV downstream RF output (single channel). Typical CMTS downstream level is +55 to +58 dBmV as shown here.

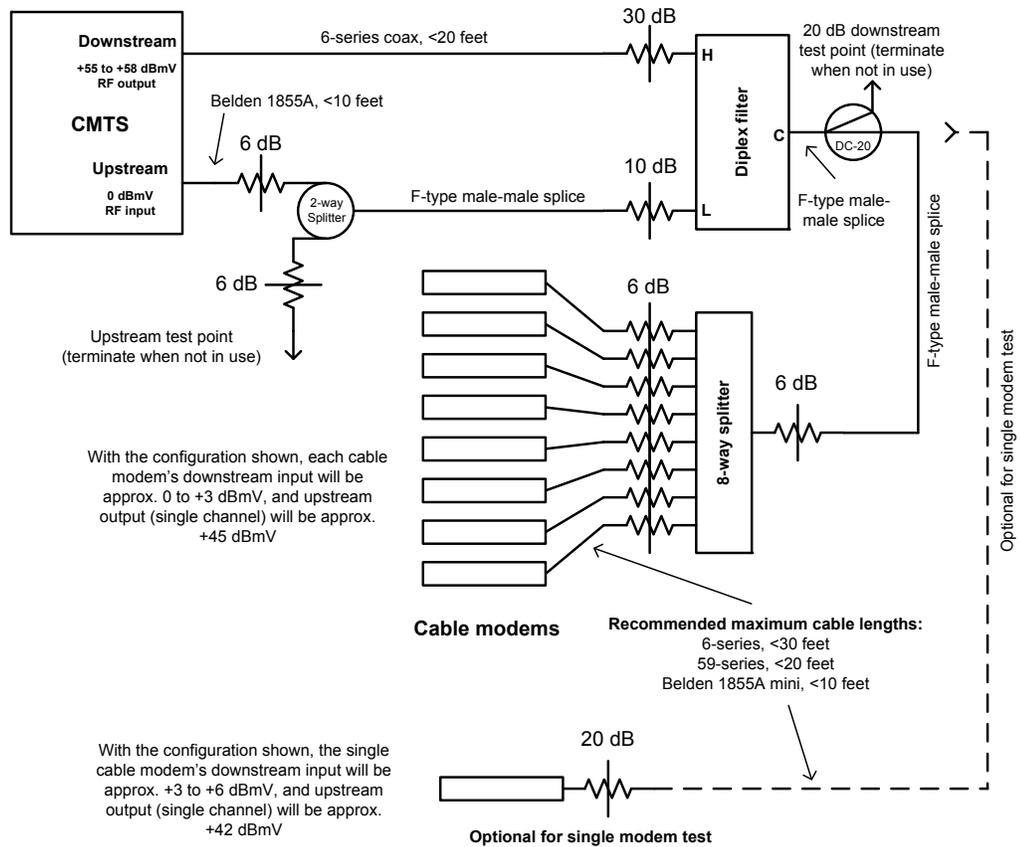


Figure 18 - Short reference plant block diagram

6.5.2 Pre-equalization Calibration Approach

The pre-equalization process between CMTS and CM compensates not only for the distortion that occurs in the plant but for any existing distortion that exists in the upstream path following the transmit equalizer within the CM up to the baseband receiver inside the CMTS. This includes filters, modulators, and amplifiers in the CM, the home network, combining network at the hub or headend as well as the front-end components in the CMTS (Figure 19).

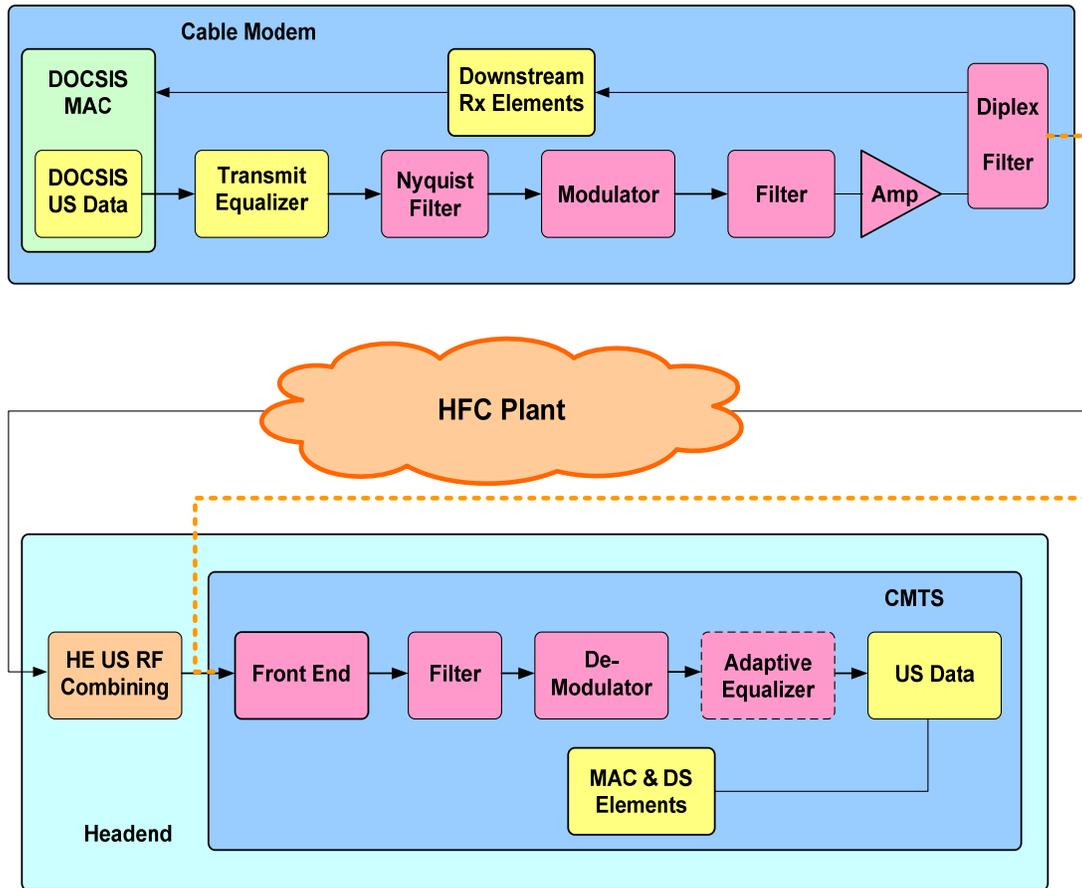


Figure 19 - CM & CMTS Elements Contributing To US Distortion (In Orange)

If there was a-priori knowledge of the distortion contributed by the CM and CMTS, it could be calibrated out in order to have a more accurate representation of the distortion in the field.

Generating a database with all the possible CM and CMTS model pairs is feasible given that there are a limited number of CM and CMTS implementations. A known short plant setup consisting of a CMTS-CM pair and a few fully characterized components enables determining almost exclusively the contribution of the CM and CMTS internal distortion. Section 6.5.1 describes an example of a known short plant where the internal characteristics of the CMTS and CM can be measured. The CMTS-CM internal distortion is obtained by gathering the pre-equalization coefficients of the CM after allowing a few maintenance intervals to elapse to achieve convergence of the coefficients.

Assuming that the real and imaginary values of the 24 CM pre-equalization coefficients for a particular CMTS-CM model pair measured on a short calibrated plant are given by:

$$F_{1CR}, F_{1CI}, F_{2CR}, F_{2CI}, F_{3CR}, F_{3CI}, F_{4CR}, F_{4CI}, \dots F_{23CR}, F_{23CI}, F_{24CR}, F_{24CI}$$

For every i^{th} real and imaginary coefficient, the resulting complex number is obtained,

$$F_{iCR} + j F_{iCI} = F_{iC}$$

Resulting in the following complex coefficients

$$F_{1C}, F_{2C}, F_{3C}, F_{4C}, F_{5C}, F_{6C}, F_{7C}, F_{8C}, \dots F_{21C}, F_{22C}, F_{23C}, F_{24C}$$

Similarly, assuming that the real and imaginary values of the 24 CM pre-equalization coefficients obtained in the field, matching the CMTS-CM model pair being analyzed, are given by:

F1R, F1I, F2R, F2I, F3R, F3I, F4R, F4I, . . . F23R, F23I, F24R, F24I

The resulting complex coefficients are

F1, F2, F3, F4, F5, F6, F7, F8, . . . F21, F22, F23, F24

Given 24 complex coefficients and assuming that the equalizer’s main tap is located at tap position number 8, a 32 element fast Fourier transform can be used to translate from the time domain into the frequency domain. Tap 8 would coincide with FFT input element 16 to preserve the relative position of the response in frequency.

The mapping of coefficients to FFT input parameters follows:

Fin1 through Fin8 = 0
 Fin9 = F1, Fin10 = F2, Fin11 = F3, . . . , Fin16= F8, . . . , Fin31 = F23, Fin32= F24

The mapping for the coefficients obtained using the short plant measurement results in the following:

FCin1 through FCin8 = 0
 FCin9 = F1C, FCin10 = F2C, . . . , FCin16= F8C, . . . , FCin31 = F23C, FCin32= F24C

After calculating the 32 point FFT the following frequency response values are obtained:

Fouti for i =1 to 32 and FCouti for i =1 to 32

For calibration, the magnitude of the field coefficients are divided by the magnitude of the short plant coefficients and the phase of the short plant coefficients are subtracted from the phase of the field coefficients.

$|F'outi| = |Fouti|/|FCouti|$ for i =1 to 32
 $\Phi(F'outi) = \phi(Fouti) - \phi(|FCouti)$ for i =1 to 32

From the calculated magnitude and phase values, the corrected F’out frequency response values are obtained

F’outi for i =1 to 32

Additional granularity in the frequency response representation can be obtained by inserting zeroes to a larger size FFT such as a 64, 128 or 256 FFT.

An example of the calibration process is illustrated next.

Figure 20 shows the distortion of a CMTS/CM pair in a short calibrated plant versus the same CMTS/CM pair measured in the field. The pre-equalization coefficients obtained from the CM MIBs are shown in Table 15.

Table 15 - Pre-equalization Coefficients of upstream path with and without micro-reflection

Field CM Eq Data MIB	08 01 18 00 FF FF 00 00 00 00 01 FF FE FF FD 00 03 00 04 FF FA FF FB 00 08 00 09 FF F0 FF EA 01 EE 00 00 FF EF FF EC 00 38 FF D8 00 55 FF D7 FF E2 00 03 00 1F FF E5 FF FC FF FA 00 01 FF FE 00 01 FF F7 FE 00 02 FF FF FF FD FF FF 00 00 00 00 00 00 FF FF 00 00 00 00 00 00
Short Plant CM Eq Data MIB	08 01 18 00 00 00 FF FF FF FF 00 01 00 00 FF FE FF FF 00 02 00 01 FF FA FF FE 00 0A FF FE FF E5 01 FE 00 00 FF FF FF E3 00 05 FF F5 FF FE 00 04 FF FF FF FD 00 00 00 02 FF FF FF FE FF FF 00 02 FF FF FF FF 00 00 00 01 00 00 FF FE 00 00 00 00 FF FD FF FF FF FF FF FF FF 00 00 00 00 01 FF FF 00 00

The first four bytes 08 01 18 00 provide the main tap position “08”, the number of taps per symbol “01”, the number of taps “18” (hex number for 24) and a reserved byte. The rest of the information are the real and imaginary coefficient data (2 bytes each) for the 24 taps.

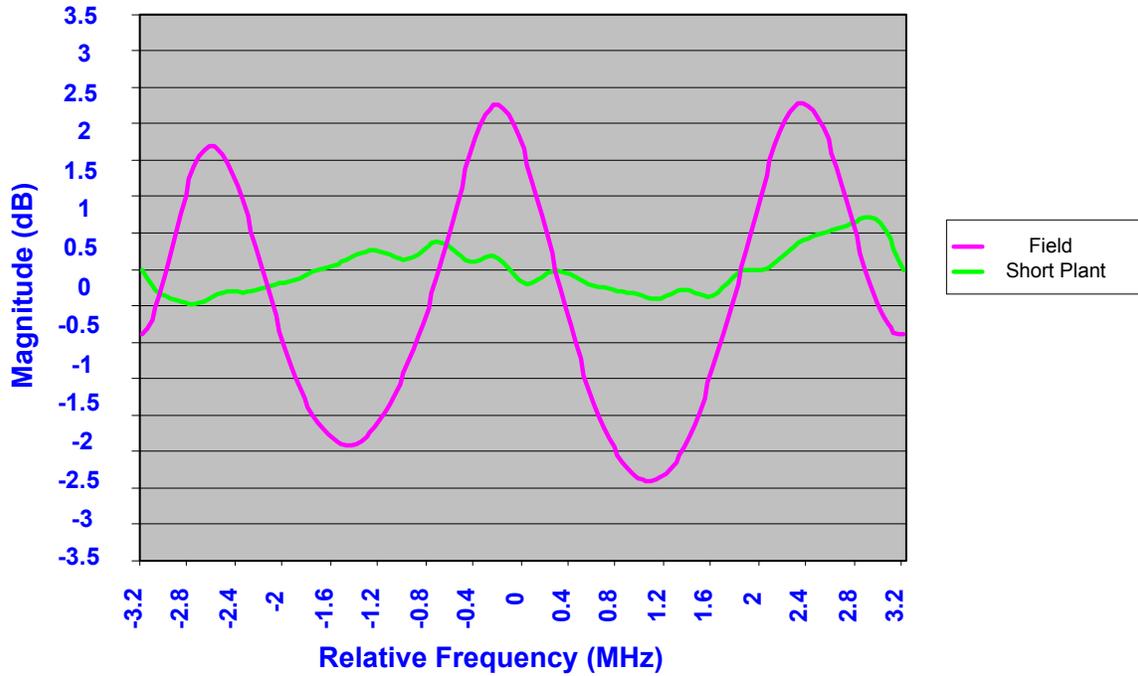


Figure 20 - Pre-equalizer frequency response with (0.5µs, -10 dBc) and without micro-reflection

The distortion in the short plant scenario is predominantly impacted by the CMTS receiver and CM transmitter. Therefore, to calibrate out the impact of the CMTS and CM, the frequency response of the short plant (micro-reflection off) is subtracted from the frequency response of the micro-reflection on scenario. This calibrated response is shown in blue in Figure 21.

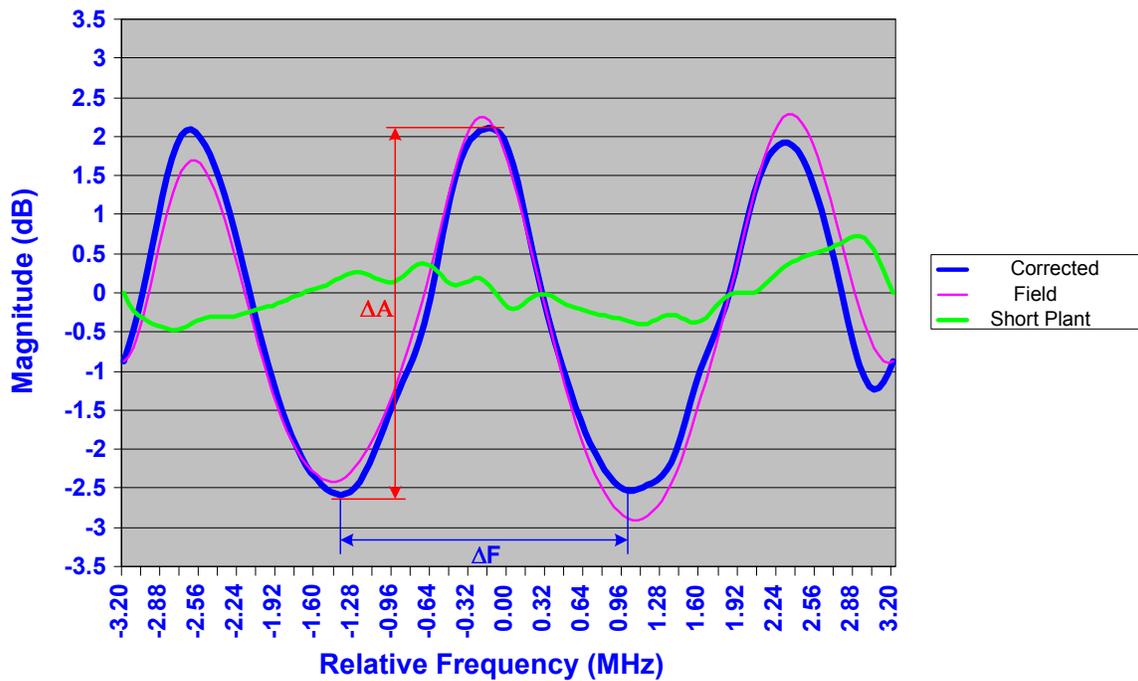


Figure 21 - Calibrated pre-equalizer frequency response obtained from micro-reflection on (0.5µs, -10 dBc) and off scenarios

Averaging ΔF and ΔA values is no longer necessary since the calibrated response is fairly even across the frequency range under observation.

6.6 Fault Localization

The process of fault detection and localization relies on monitoring the network for general plant-wide or neighborhood-localized problems as well as for specific end devices. In this process it is assumed that there is detailed knowledge of the node service area's topology. It is also assumed that distortion data (pre-equalizer coefficients and other applicable information) has been collected from the CMs and analyzed to determine the distortion signatures of the affected CM(s). Next, a process is described by which, through correlation of topology with distortion signatures, the location of faults can be determined.

In the example highlighted in Figure 22 is a group of CMs, identified in red, that exhibit the same unique distortion. The CMs in green are CMs that don't share that specific distortion.

It is assumed that to obtain the distortion signatures, an analysis and classification process of the impairments has already taken place.

If only information from one CM were available, the problem area could only be isolated to somewhere along the path between the CM and the fiber node (dashed line). The more interesting process is when the relationships of CMs that share specific impairments (as well as those that do not) to upstream paths are examined.

In order to estimate the impairment location, the common path shared by the end devices showing the specific impairment is found. This path containing the impairment is further constrained by excluding the path that is shared with the end devices that operate properly.

Knowledge of the micro-reflection signature also helps localize the problem. For example, in Figure 22 a triple transit reflection signature with delay matching the distance between known devices on the plant, such as distance between taps, length of drops, or distance between amplifiers, can point to the likely cause or narrow down the possible set of causes of the problem.

After analysis and path manipulation of the end devices showing impairments such as micro-reflections, a potential location of the problem is determined. These areas are shown in purple on Figure 22. This mechanism maps the devices that have the same unique micro-reflection attribute and pinpoints the portion of the network that exhibits the impairment.

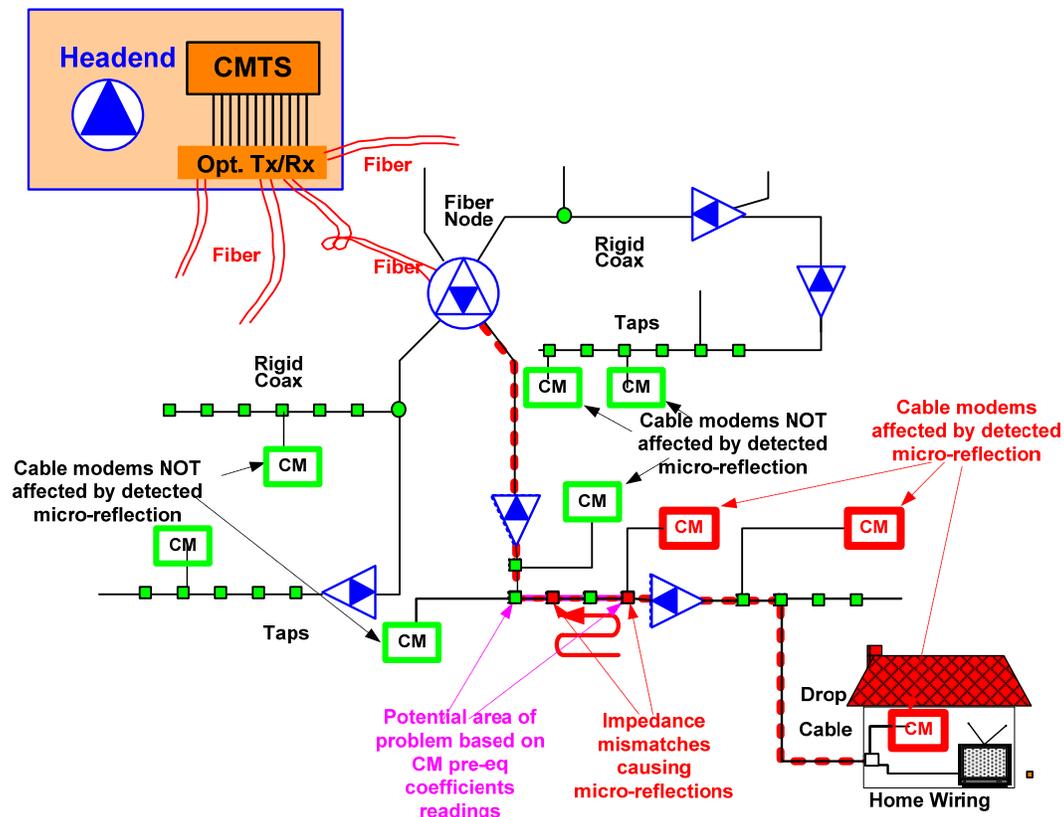


Figure 22 - Correlation of Topology with Distortion to Provide Fault Localization

6.6.1 Fault Localization Examples

This section details basic examples of fault location based on micro-reflection signatures. Data from a 6.4 MHz DOCSIS channel on a single fiber node was used. A few limitations of the data set are worth noting. First, the node was selected not for the diversity of problem scenarios but for the availability of the digital maps. Therefore, the micro-reflections found were not significant, and in most cases the scenarios detected corresponded to issues either near a line termination or impedance mismatches inside or near-the-home. Second, there was a relative low number of CMs in the 6.4 MHz channel in this node which led to few scenarios for evaluations as well as few “monitoring probes/CMs” for accurate determination of the impairment location. High penetration of CMs narrows the potential problem area where the leading edge of the micro-reflection can be located.

6.6.2 Determining Micro-reflection Signatures

A micro-reflection signature consists of a pair of characterizing elements, the relative level of the reflected signal compared to the main signal and its delay. The level provides an indication of the severity of the micro-reflection (see Section 6.7.4 for Severity Analysis Strategy for Intermittent Issues considerations). The delay represents the extra distance traversed by the reflected signal or echo and is represented in microseconds (see Section 5.2.1 for definitions of micro-reflections types).

Figure 23 shows the frequency response of several CMs sampled from a fiber node. For clarity, only a few CMs are included in the figure. However, in the field, if encompassing an entire fiber node, the technician will have to pay close attention to the chart to distinguish and extract the common micro-reflections (labeled as A and B in Figure 24) from all the micro-reflections present. Nonetheless, the intention of this section is to explore systematic mechanisms where human intervention can be minimized as a first pass at extracting the most relevant cases for

analysis and troubleshooting. In summary, the micro-reflection signature which provides a two-dimensional characterization from which the selection of common frequency responses as noted in Figure 24 is feasible.

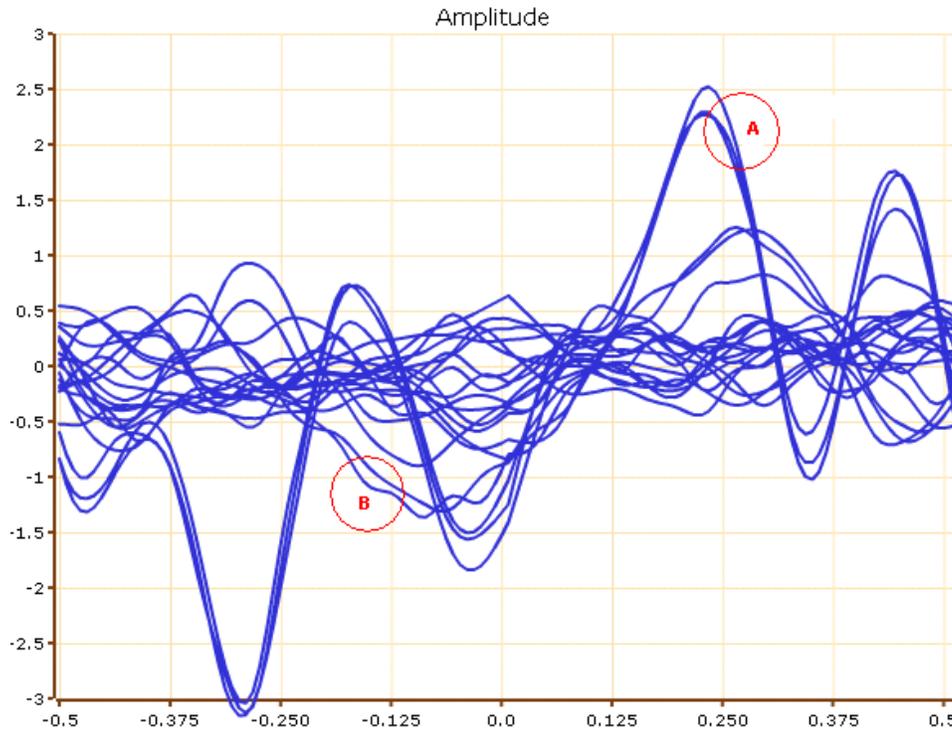


Figure 23 - Observation of multiple CMs frequency response

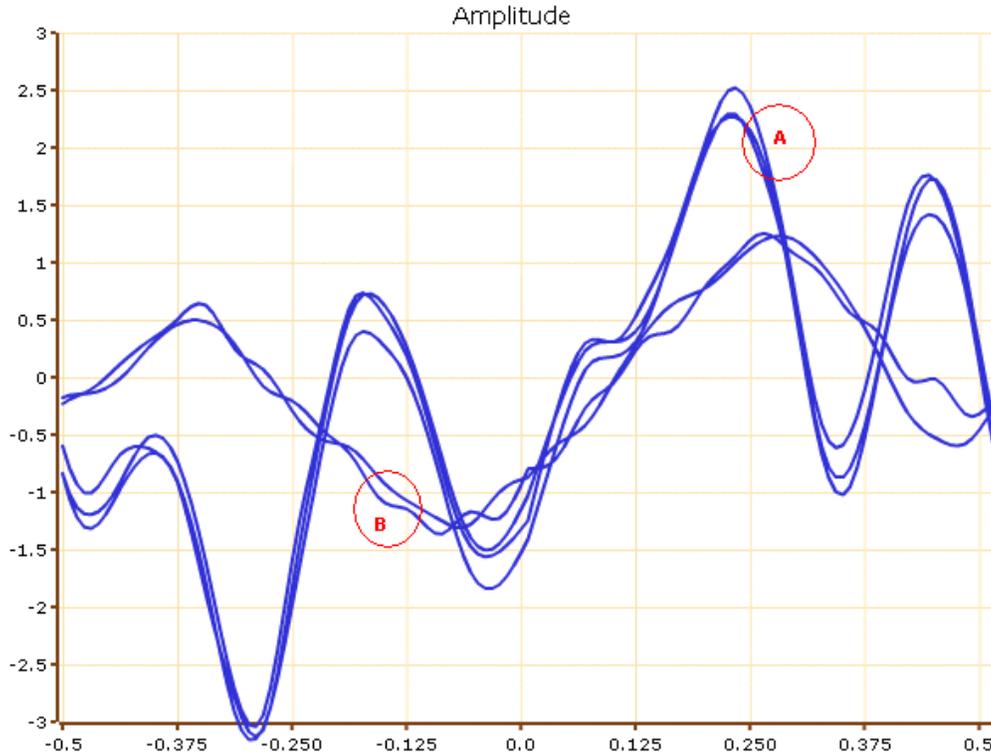


Figure 24 - Identified Micro-reflection Patterns

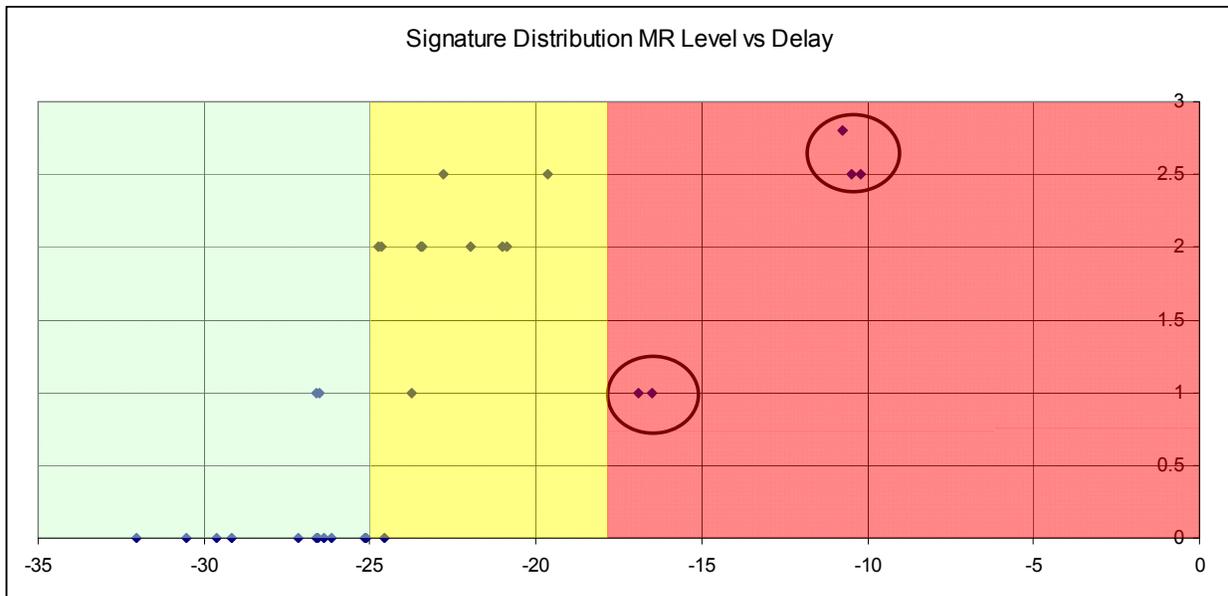


Figure 25 - Clustering of Common Micro-reflection Signatures

As seen in Figure 24, the amplitude ripple of the frequency response is sufficient to determine group the CMs with same micro-reflection. Some amplitude variability margin must be allowed. Note that in a few cases, the same ripple magnitude range could correspond to different micro-reflections. In such cases the estimation of echo delay can provide an unambiguous answer. For typical fiber node sizes the probability of uniqueness of a micro-reflection is

very high. In rare exceptions, the micro-reflection signature pair will not be sufficient to determine the CM clusters (e.g., a sub-T echo delay). In such cases, only the topology resolution will provide final resolution of such conditions. Figure 25 shows the way that common signatures of CMs are correlated. The delay is an approximation of the estimated maximum delay and it is expressed as $2T$ where T is the inverse of the symbol period. That means the distance between the two reflectors is calculated based on half the delay. Delay below the "no action required" threshold is not relevant.

6.6.3 Determining Micro-reflection Boundaries Edges

For the relevant micro-reflections cases A and B, the next steps consist of performing the localization within the HFC plant topology.

By considering the upstream direction and other CMs in the path that do not present the same signature, the trailing edges of the micro-reflection is determined as discussed at the beginning of this section.

Figure 26 shows micro-reflection case A. This micro-reflection can be considered as a composite of type 1 and 2. CM1 and CM2 are close to the Tap reflecting the signal and CM 3 has an additional transit for the echo relative to the distance to the reflecting tap. Hence the echo from CM3 ($2.8T$) is longer than CM1 and CM2.

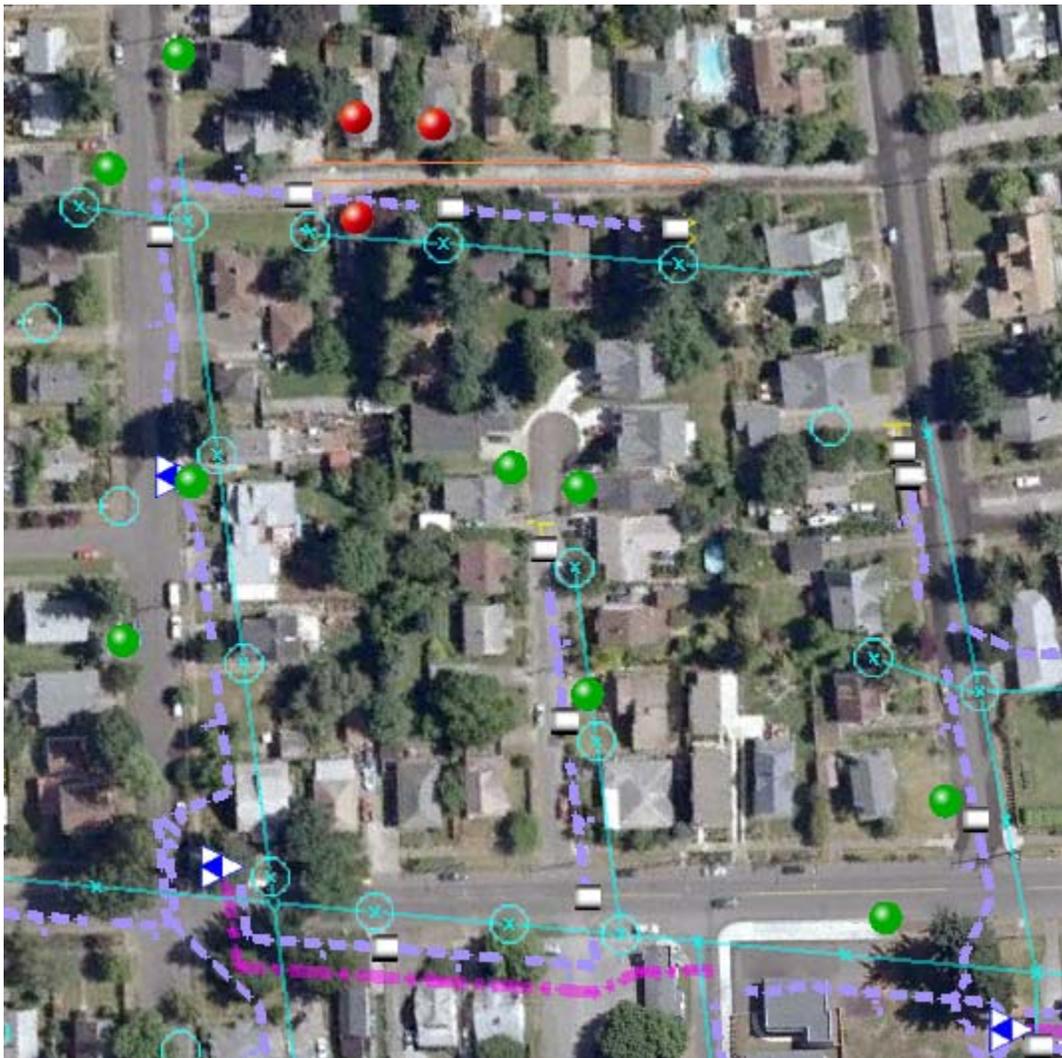


Figure 26 - Common Micro-reflection Signature - Case A

Figure 27 shows the case of a single sub-T echo. The two signatures come from the same household, H1, and note the neighbor CMs: CM4, 5, 6, and 7 do not register the same micro-reflection signature as the CMs in observation.

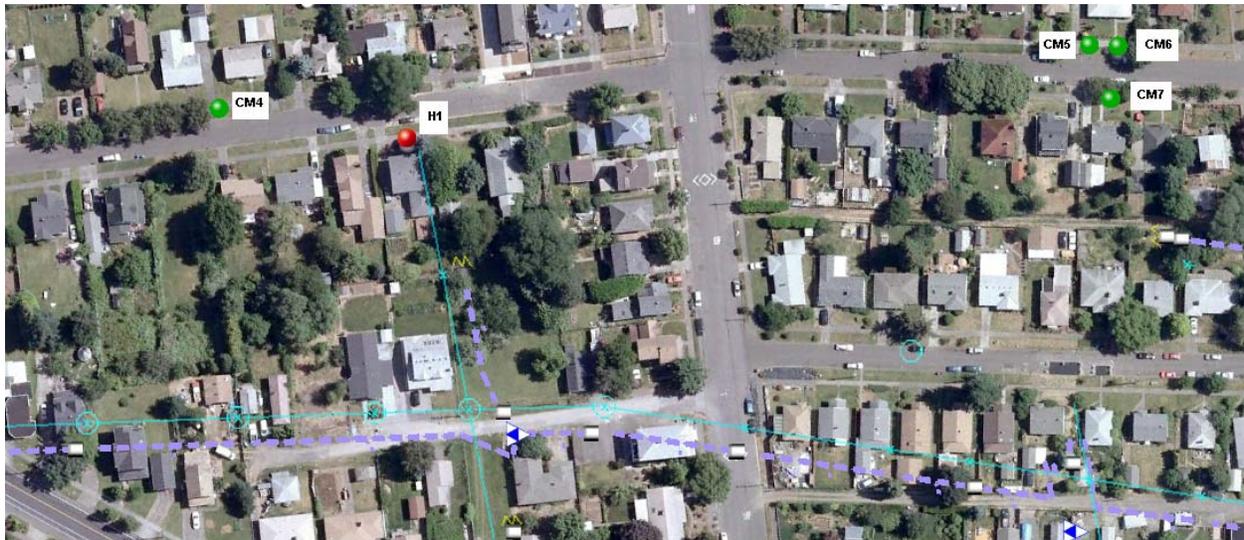


Figure 27 - Common Micro-reflection Signature - Case B

Note that limited data could lead to non-optimal location of the micro-reflection. For example, fewer CMs reporting pre-equalization data or few customers in the fiber node branch will reduce the possibility to accurately estimating the location of the micro-reflection.

6.7 Severity Assessment

Using the pre-equalization mechanism defined in DOCSIS is efficient, resulting in no performance degradation even under strong micro-reflections. Pre-equalization may help decrease the urgency for plant repair, but it should not be used to circumvent required plant maintenance. The purpose of the proactive network maintenance method is to listen to important network health metrics and take action before service is impacted.

Micro-reflection severity can be graded into three categories: immediate action required (Red), high monitoring frequency (Yellow), and no action required (Green).

The three severity assessment metrics that are considered for a single CM are defined as follows:

- The first severity metric is a static classification conducted solely based on the relative level of the micro-reflection amplitude.
- A second severity metric results from a trending analysis that is conducted so that network operators can identify an impairment degrading at a rate that will result in immediate action in the near future (e.g., one week). This method requires time stamping of the measurements gathered so that a history of the impairment is obtained.
- A third severity metric is generated from historical comparison of the micro-reflection levels of the high-frequency monitored CMs. This third metric measures intermittent fluctuations that would be considered significant, but which did not rise to the immediate action level. The measurements should be conducted over multiple days. The comparison of the measurements is done at the same time of the day.

Figure 28 illustrates the severity assessment metrics.

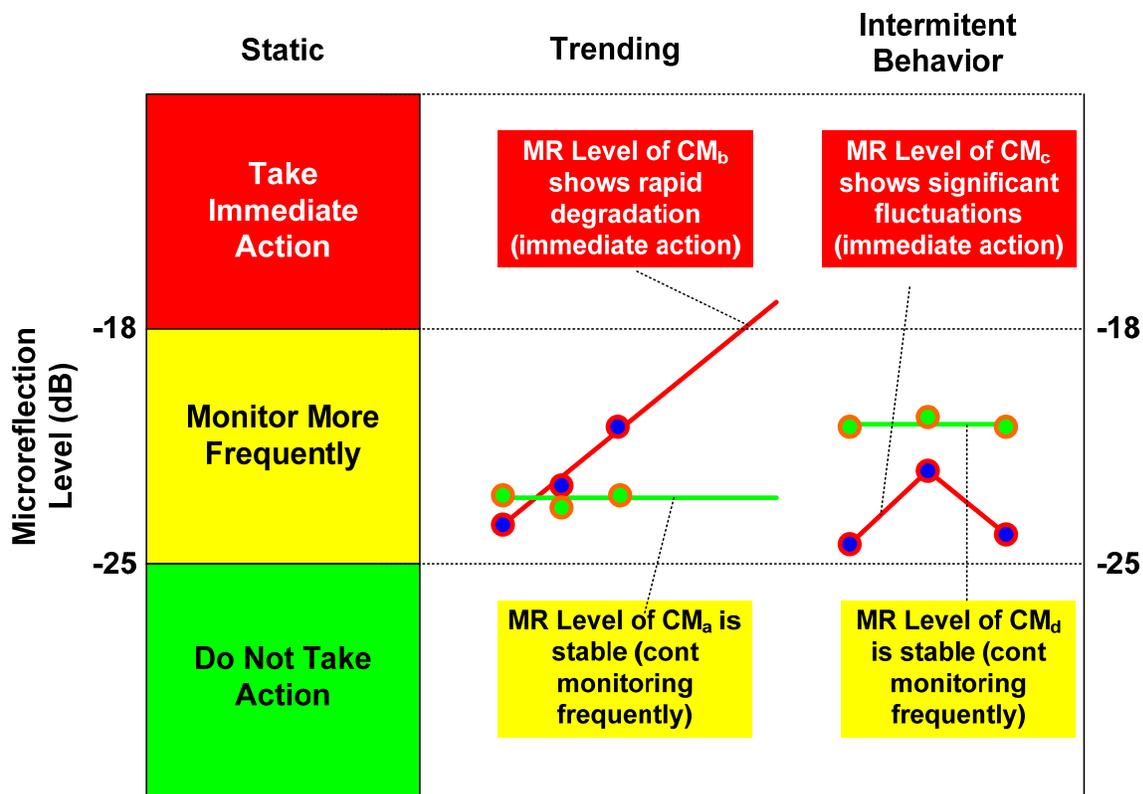


Figure 28 - Severity Classification Mechanisms

From Figure 28, a single micro-reflection level of -18 dBc results in an amplitude ripple of 2.2 dB and a -25 dBc micro-reflection results in an amplitude ripple of 1 dB. See Micro-reflection Calculator in Appendix IV for more information on converting ripple to relative micro-reflection levels.

6.7.1 Initial CM Selection for Analysis

It is assumed that the monitoring strategy combines a standard measurement interval (MIntSTD) to monitor all CMs in the network with a frequent measurement interval (MIntFREQ) for CMs that are deemed of interest for further analysis. This initial assessment of which CMs need to be examined further and more frequently can be done using the Distortion Metric, the ratio of the aggregate energy that exists in all but the main adaptive equalizer tap to energy in all adaptive equalizer taps, defined as NMTER in Key Metrics Section 6.3.

This distortion metric, expressed in dB, can quickly indicate which CMs have to go through detailed signature analysis. These CMs of interest will also be placed in the “monitor more frequently pool” to gather performance history and detect signature trending and intermittent behavior.

Consider a scenario in which the threshold value of the distortion metric is -25 dBc shown in Figure 28. This means that any CM that exhibits a NMTER > -25 dB will undergo detailed coefficient analysis through which one or more micro-reflection signatures will be determined. It is assumed for the analysis that follows that the CM pre-equalization coefficients have already been analyzed and manipulated, and distinct micro-reflection amplitude and delay signatures have been obtained.

6.7.2 Severity Analysis Strategy for Static or Single Data Point Scenario

In a static environment where no change of pre-equalization data is assumed, or when there is only one data point and not able to determine change, a simple set of fixed thresholds can be used to determine severity categories.

For example, a total distortion energy amplitude of -25 dBc or lower may fall into the category of “No Action Required,” while -18 dBc or greater may be considered in the category of “Immediate Action Required.” Two thresholds determine the three categories. In Figure 28, the values below -25 dBc belong to the “No Action Required” category or green severity. The values above -18 dBc belong to “Immediate Action Required” category or red severity. The values between -18 dBc and -25 dBc belong to the category “High Monitoring Frequency” Category or yellow severity.

As additional data points are collected, time dependent analysis can be conducted which is described in the next section. Since time dependent analysis is conducted for CMs that are already in the yellow category from a static perspective, the time dependent severity classifications only have to define a red classification criteria.

6.7.3 Severity Analysis Strategy for Trending

If, for example, an analysis resulted in a single micro-reflection with a signature of -22.2 dBc in amplitude and 0.5 μs delay, it would fall into the yellow severity or “monitor more frequently” category. At the same time, a second CM showing an initial signature of -19.2 dBc in amplitude and 1 μs delay also falls into the yellow severity or “monitor more frequently” category.

Assuming that the standard measurement interval (MInt_{STD}) is once a day and the frequent measurement interval (MInt_{FREQ}) is every four hours, the data from these two CMs’ micro-reflection amplitudes tracked in time is detailed in Table 16. The micro-reflection time delay is not used for comparison here and is not shown. It is expected though that the changes in delay will be negligible.

Table 16 - Two CMs showing micro-reflection amplitude over 2 days

8/19/09	T0_12	T0_16	T0_20	T1_0	T1_4	T1_8	T1_12	T1_16	T1_20	T2_0	T2_4	T2_8
CM#1	-22.2	-22.2	-20.4	-22.8	-21.8	-21	-22.5	-20.4	-22	-20.6	-20	-21
CM#2	-19.2	-19.5	-20	-20.3	-20	-19.5	-19.2	-19.5	-20	-20.3	-20	-19.5

The time measurement reference parameter T_{n_m} is coded as follows; the first number (n) indicates the number of days after the initial measurement date, while the second number (m) corresponds to the approximate time of day when measurement was taken. This means that measurement T0_16 was taken at 4:00 p.m. on the same day as the initial stored measurement, and measurement T2_12 was taken two days after the first measurement and at 12:00 p.m.

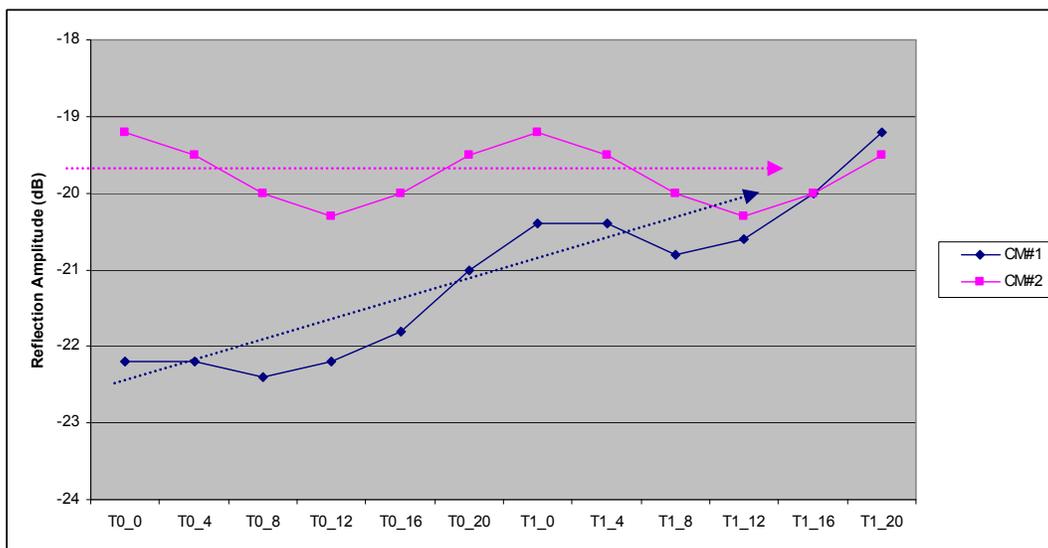


Figure 29 - Micro-reflection amplitude data of same two CMs to highlight trending over time

Figure 29 plots Table 16’s micro-reflection amplitude versus time, indicating that in less than three days, CM#1 will be going from yellow to red according to the static criteria while CM#2 stays in the yellow region. These trends should be conducted using data points measured at the same time of the day so that approximately equal temperature data points are obtained.

If CM#1 shows a delta of 1.8 dB between T0_0 (-22.2 dBc) and T1_0 (-20.4 dBc), at that rate of change it will reach -18 dBc in less than three days with respect to T1_0. The estimated level at T2_0 is -18.6 dBc and T3_0 would be -16.8 dBc, crossing into the red region.

If delta amplitude per day is D, the last amplitude measured was MR_{Last} and if MRLast + 3*D > Trending Red Threshold then it is classified as red for the trending criteria.

3D is a variable subject to operator adjustment and is intended as a recommendation.

In addition, another indicator of urgency is determined by the number of days that the operator has until the micro-reflection amplitude reaches the red region.

$$\text{Target\#ofDaysToRepair} = (-18 - \text{MR}_{\text{Last}})/D$$

6.7.4 Severity Analysis Strategy for Intermittent Issues

Another time-related behavior of adaptive equalizer coefficients operators may observe is the rapid change of micro-reflection amplitude. This intermittent behavior may be hiding a problem that is on the verge of causing service interruption and should be addressed promptly.

Table 17 shows two CMs that are monitored more frequently as they initially were categorized in the yellow category but one of the CMs shows a behavior that cannot be explained with daily temperature variations.

Table 17 - Micro-reflection amplitude of two CMs showing intermittent issue

8/19/09	T0_12	T0_16	T0_20	T1_0	T1_4	T1_8	T1_12	T1_16	T1_20	T2_0	T2_4	T2_8
CM#1	-22.2	-22.2	-20.4	-22.8	-21.8	-21	-22.5	-20.4	-22	-20.6	-20	-21
CM#2	-19.2	-19.5	-20	-20.3	-20	-19.5	-19.2	-19.5	-20	-20.3	-20	-19.5

Figure 30 plots Table 17’s micro-reflection amplitude versus time showing CM#1 with significant micro-reflection amplitude swings. Both CMs remained in the yellow region under the static severity criteria but CM#1 shows drastic changes in micro-reflection amplitude which could be a loose connector or something likely to break in the very near future.

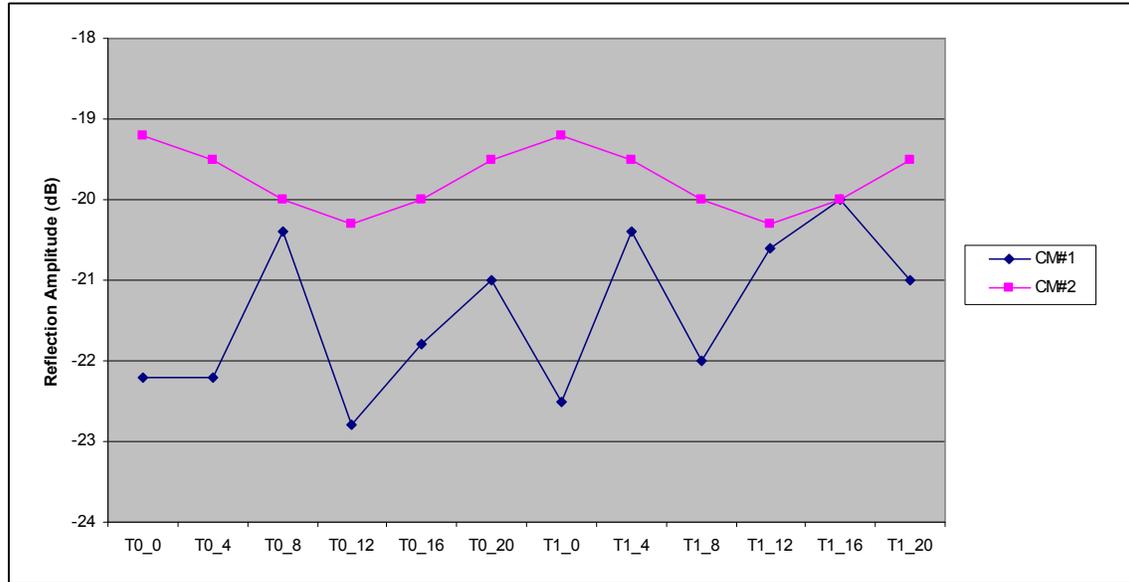


Figure 30 - CM micro-reflection amplitude over time highlighting intermittent issues

CM#2 shows gradual amplitude swings that could be attributed to daily temperature variations. CM#1 shows significant variations in micro-reflection amplitudes. A metric that could define the red classification criteria for intermittent behavior is proposed as follows.

Intermittent Red Threshold = $\text{Avg 4 hour } \Delta \text{ Ref. Amp} / (\text{Static Red Threshold} - \text{Avg. Refl. Amp.}) > 0.25$

- where the average four hour delta of the reflection amplitude represents the average swing in dB over the four hour interval which corresponds the monitoring granularity in time of the devices that were deemed to need high-frequency monitoring and historical data tracking. The average reflection amplitude is measured up to a fixed number of days (for instance, four days). If this metric approaches 1 it indicates that the reflection amplitude is varying with swings that approach the static threshold level. If this metric is smaller than 0.1, for example, one can say that the reflection amplitude swings occur at a safe margin below the static threshold.

The scenarios just discussed dealt with the quantification of the degree of impairment of a single CM. It is assumed that a cable operator may add weighting factors to indicate urgency in addressing these issues after the previous information is correlated with the number of CMs affected by a particular impairment. The analysis used to determine fault location can also be used to determine how many users will be impacted.

6.8 Use Case Examples for Adaptive Equalizer Coefficients

The following four listed states describe the condition of the plant and the performance of the measurement devices for different scenarios operators may encounter in the field.

State 1. Adaptive equalizer tool is working properly and plant is within acceptable limits.

State 2. Adaptive equalizer tool is working properly, and plant exhibits severe linear distortion.

Plant is stable

Plant is unstable (intermittent or trending)

State 3. Adaptive equalizer tool is working properly, but CMTS/CM is reacting badly

Transmission characteristics in channel have resulted in no solution for CM coefficients (e.g. a deep suck-out or echo is too long or too severe)

Upstream impulsive noise is causing wrong or unstable adaption.

CM needs to be replaced

State 4. Adaptive equalizer tool is not working correctly because of equipment design or configuration.

CMs are giving wrong MIB data, but working properly.

CMs are giving wrong/no MIB data and not working properly.

Wrong configuration of CMs or CMTS.

The CMTS is not configured for adaptive pre-equalization in the upstream

CM includes DOCSIS capable devices STBs and MTAs

Distortion red = static or trending or intermittent

6.8.1 Use Cases

Use cases are a processes to describe how tools will be used. Examples are provided in the following subsections.

6.8.1.1 Use Case 1: In or Near Home Problem Identification

Id: UC- PNM100

Description

Support User detects possible service problem by monitoring (proactive) or trouble call (reactive). Analysis tool is used to interrogate the customer's modem and CMTS for performance information. The performance data is analyzed for radio-frequency reflections to determine approximate problem location. Location information is used to instruct Support User or dispatch a repair technician to resolve the problem.

Level: User Goal

Primary Actor

Maintenance User or Support User

Supporting Actors

Customer, Repair Technician, Dispatch Personnel, Customer Modem, CMTS

Stakeholders and Interests

Customer – may be unaware of service trouble in proactive scenario.

Plant Maintenance – non-critical repair may be scheduled during routine maintenance.

Billing Personnel – billing system accounts may need to be reconciled.

Pre-Conditions

Support User must have specific knowledge of a possible service problem.

Customer's modem must be online.

CMTS and customer modem must both support Adaptive Equalization.

Analysis tool must be available to the Support User.

User must have sufficient access and knowledge of using the tool and interpreting the data.

Impairments to the customer's signal must be such that the modem can still respond correctly.

Post Conditions

Success end condition

Reactively, the problem was identified and repaired to the satisfaction of the customer. Proactively, the problem was identified and repaired before the customer perceived a problem. Customer's modem and entire node show improved performance.

Failure end condition

Modem fails to demonstrate any actionable performance metrics. No remediation of problem, customer may experience recurrence or deterioration of service.

Trigger

Proactive alert identifies possible service affecting problem OR reactive alert, customer calls with trouble.

Main Success Scenario

1. Customer calls with service trouble.
2. Support User initiates analysis tool, providing customer identification.
3. Customer identification is used to obtain the modem MAC address and CMTS for the customer.
4. MAC address is used to query both the modem and CMTS for performance information.
 - Reference SD-PNM200 for software sequence (Appendix VI).
5. Performance data is analyzed to obtain reflection information if available, used for distance calculations.
6. Support User is presented with customer's service address and an estimate of problem distance from home.
7. Support User evaluates distance information to determine "in or near home" problem, escalates to Dispatch Personnel.
8. Dispatch Personnel evaluates trouble call information and approves a repair call.
9. Repair Technician is provided with analysis and dispatched to the customer premise.
10. Repair Technician locates trouble and resolves issue.

Extensions

4a. in Step 4, the customer information is not located.

1. Trouble is escalated to billing department for account reconciliation.

4b. in Step 4, the modem or CMTS do not support equalization analysis.

1. Trouble resolution reverts to conventional process.

5a. in Step 5, analysis data shows no trouble present.

1. Customer account is noted for future reference
2. Proactive monitoring is escalated to a higher rate

5b. in Step 5, customer has multiple devices at premises.

1. All devices are analyzed for similar distortion characteristics
 - 1a. multiple devices share distortion - common fault is noted in analysis
 - 1b. Single device demonstrates distortion - scope of analysis is narrowed to “in home”

5c. in Step 5, adjacent homes may also be analyzed

1. All devices within a specified radius are analyzed for similar distortion characteristics
 - 1a. multiple customer devices share distortion - common fault is likely to be at or above the tap.
 - 1b. single customer devices demonstrate distortion - fault is more likely to be between the tap and something in the customer’s home.

Variations

- 1a. in Step 1, proactive monitoring may provide use case trigger.
- 1b. in Step 1, automations such as an interactive telephone system may provide use case trigger.
- 2a. in Step 2, the Support User role may be implemented as an automated telephone system.
- 3a. in Step 3, the customer identification may be any one of MAC, telephone number or account number.

Frequency

Use case is executed per each trouble call, day of install or subject to the frequency of proactive monitoring.

Assumptions

Users of the system must have basic knowledge of troubleshooting cable service problems.

Users must have required access and training of the analysis tools.

Special Requirements

All systems must support authentication and encryption pursuant to corporate security standards.

Issues

1. Distance calculations are preliminary, subject to market specific conditions and trial findings.
2. Distance estimates actually describe the distance between two impedance mismatches (reflectors). In the case of in or near home, the assumption is made that one of the reflections is within the subscriber’s home. This may not always be the case however the majority of the time, this holds correct. There will be some cases where an outside plant issue may be incorrectly characterized as “in or near home” problem.

To do

1. Validate use case in arid climates with less water ingress and corrosion.
2. Track distance estimates with actual problems found to identify opportunities of improvement in the distance calculations.

Visualizations

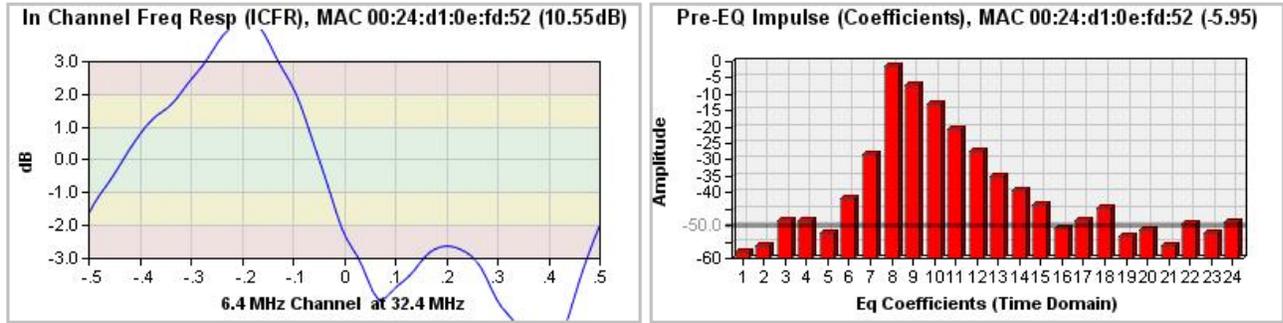


Figure 31 - Amplitude vs. Frequency peak/valley of 10.55 dB with echo present in impulse response

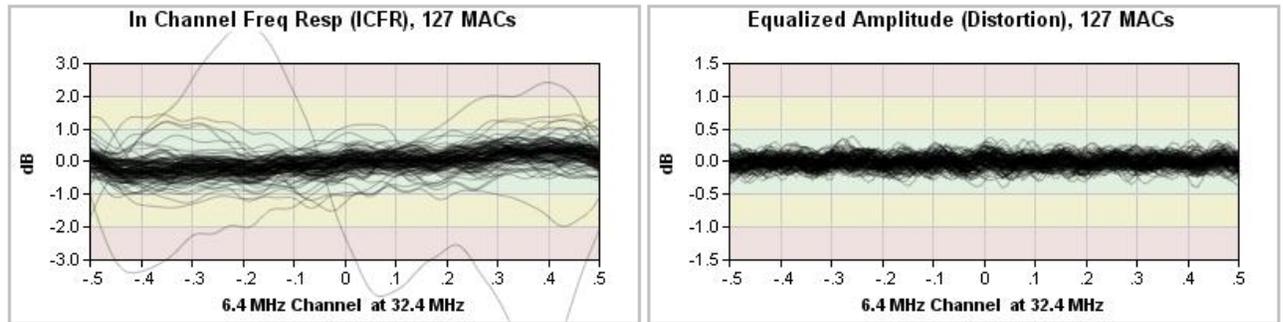


Figure 32 - Entire upstream scan shows no similar signatures shared by other modems

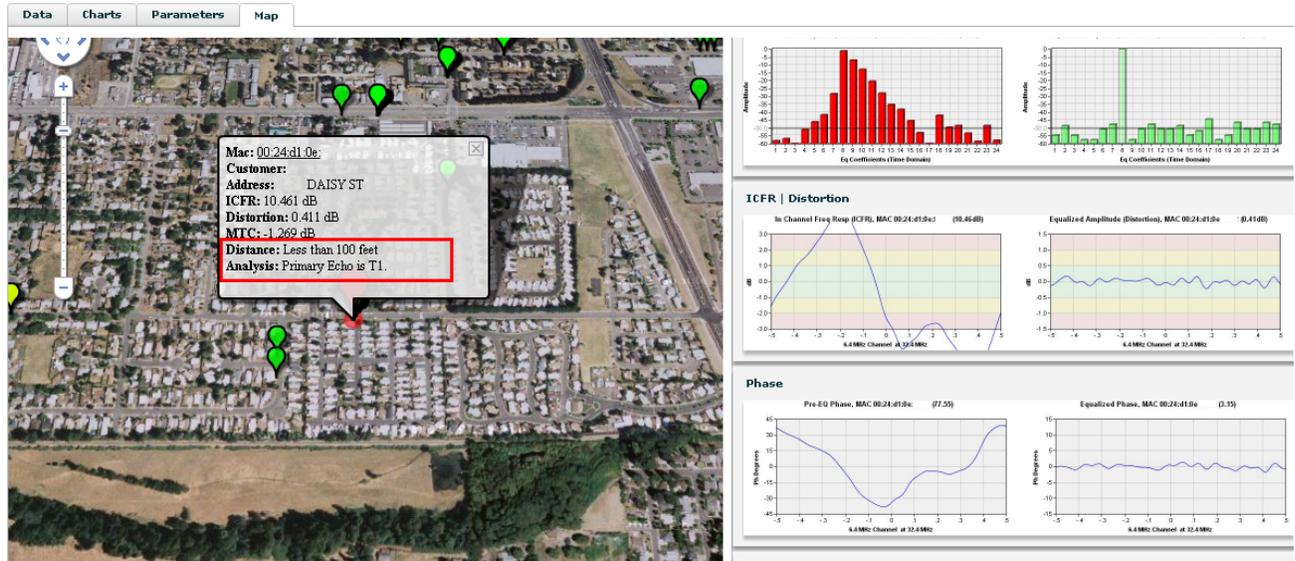


Figure 33 - Distance calculation applied with customer address and mapping

6.8.2 Use Case: Upstream Ingress or Noise Detection

Id: UC- PNM101

Description

Support User detects possible service problem by monitoring (proactive) or trouble call (reactive). Analysis tool is used to interrogate the customer's modem and CMTS for performance information. The performance data is analyzed for additive influence in the modem transmission.

Level: User Goal

Primary Actor

Maintenance User or Support User

Supporting Actors

Customer, Repair Technician, Dispatch Personnel, Customer Modem, CMTS

Stakeholders and Interests

Customer – may be unaware of service trouble in proactive scenario.

Plant Maintenance – non-critical repair may be scheduled during routine maintenance.

Billing Personnel – billing system accounts may need to be reconciled.

Pre-Conditions

Support User must have specific knowledge of a possible service problem.

CMTS and customer modem must both support Adaptive Equalization.

Analysis tool must be available to the Support User.

Support User must have sufficient access and knowledge of using the tool and interpreting the data.

Post Conditions

Success end condition

Reactively, the problem was identified and repaired to the satisfaction of the customer. Proactively, the problem was identified and repaired before the customer perceived a problem. Customer's modem and entire node show improved performance.

Failure end condition

Intermittent ingress such as impulse may no longer be present. Modem fails to demonstrate any actionable performance metrics. No remediation of problem, customer may experience recurrence or deterioration of service.

Trigger

Proactive alert identifies possible service affecting problem OR reactive alert, customer calls with trouble.

Main Success Scenario

1. Customer calls with service trouble.
2. Support User initiates analysis tool, providing customer identification.
3. Customer identification is used to obtain the modem MAC address and CMTS for the customer.

4. MAC address is used to query both the modem and CMTS for performance information.
 - Reference SD-PNM200 for software sequence. (Appendix VI)
5. Performance data is analyzed to obtain the distortion present at the CMTS after equalization.
6. Support User evaluates performance of all modems on the same node.
7. Support User determines that only the single customer demonstrates a noisy response.
8. Customer is asked to tighten connectors and remove extra cable or splitters.
9. Support User rescans device for performance information and perceives that the problem has been resolved.

Extensions

- 4a. in Step 4, the customer information is not located.
 2. Trouble is escalated to billing department for account reconciliation.
- 4b. in Step 4, the modem or CMTS do not support equalization analysis.
 2. Trouble resolution reverts to conventional process.
- 5a. in Step 5, analysis data shows no trouble present.
 3. Customer account is noted for future reference
 4. Proactive monitoring is escalated to a higher rate
- 6a. in Step 6, modems on same node may also be analyzed
 2. All devices sharing the common upstream interface are analyzed
 - 1a. multiple customer devices share ingress – one or more points of ingress need to be resolved.
 - a. Trouble is escalated to Dispatch for maintenance scheduling
 - 1b. single customer device demonstrates noisy response – likely impedance problem caused by loose, damaged or corroded connectors or cable.

Variations

- 1a. in Step 1, proactive monitoring may provide use case trigger.
- 1b. in Step 1, automations such as an interactive telephone system may provide use case trigger.
- 2a. in Step 2, the Support User role may be implemented as an automated telephone system.
- 3a. in Step 3, the customer identification may be any one of MAC, telephone number or account number.

Frequency

Use case is executed per each trouble call, day of install or subject to the frequency of proactive monitoring.

Assumptions

Users of the system must have basic knowledge of troubleshooting cable service problems.

Users must have required access and training of the analysis tools.

Special Requirements

All systems must support authentication and encryption pursuant to corporate security standards.

Issues

- Magnitude of equalized response ripples from the CMTS need to be correlated to concrete BER / MER values. Pending lab work.

To do

- Better correlate the perception of “noise” with unstable equalizer operation. This is reproducible with corrosion.

Visualizations

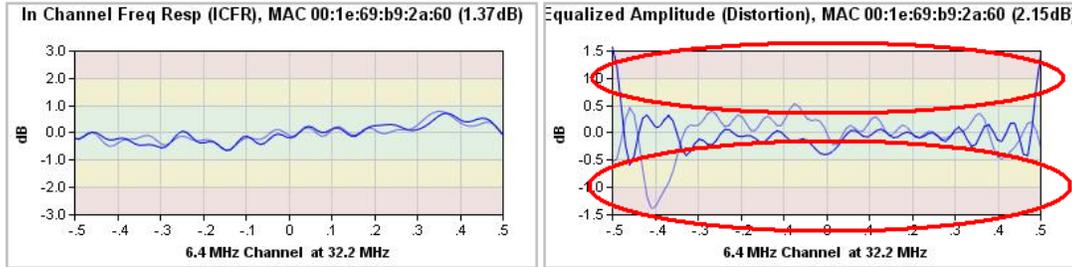


Figure 34 - Single customer modem demonstrates the effects of ingress

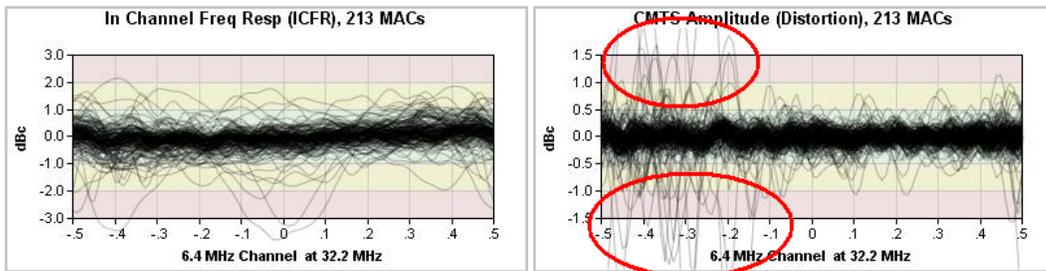


Figure 35 - Multiple modems on the same upstream demonstrate the effects of ingress

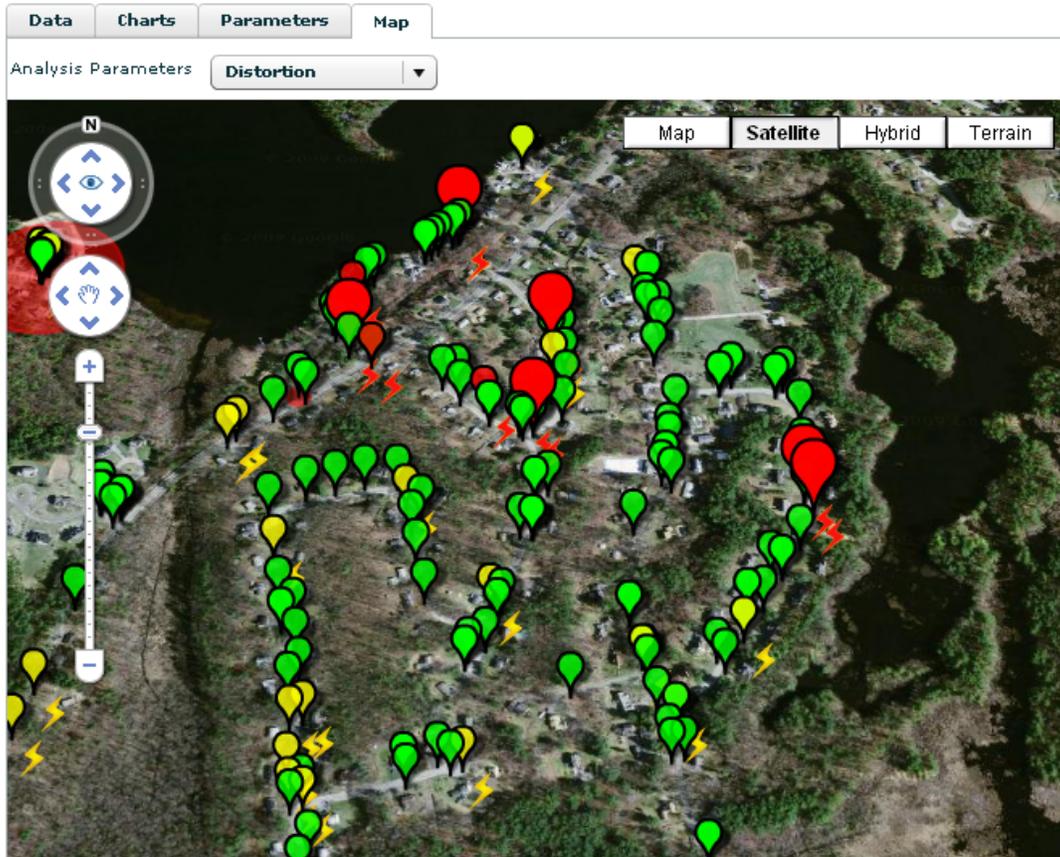


Figure 36 - Effects of ingress when mapped

6.8.3 Use Cases 3-11

3. Plant technician trying to solve a frequency response problem with the upstream or augment system upstream sweeping. State 2.

Actors: Technician, NOC

External Event: Distortion red, other indicators yellow/green or plant sweep shows problem/anomaly

Proactive/High distortion detected

Group of CMs

Expected Outcome: Distortion removed

4. Tech Supervisor: To deploy technicians. Verify technicians got the upstream frequency response issues resolved. State 2 becoming State 1.

Actors: Tech Supervisor, Technician

External Event: Close out trouble ticket

Reactive and Proactive

Group of CMs

Expected Outcome: Resolved Issue (non-main adaptive equalizer tap or coefficient energy reduced)

5. CMTS having issues due to equalization. Unable to compensate for channel. Micro-reflection too long or no inverse solution, e.g. notch in channel. Upstream burst noise interfering with the correct ranging of CMs' coefficients. State 3.

Actors: Technician, NOC, CSR, HE Tech

External Event: Distortion red, other indicators yellow or red

Reactive/High distortion detected

Group of CMs

Expected Outcome: Fix cause of impairment

6. QA for accessing a node health score. Goes along with upstream MER (on a per CMTS port or per CM basis). States 1 and 2.

Actors: Technician, NOC (QA)

External Event: Periodic scan and alarms - Distortion red, other indicators red/yellow/green

Reactive and Proactive/High distortion detected

Group of CMs

Expected Outcome: Distortion removed

7. Identification or location of faults. Combine location technique with GPS data or plant connectivity data. Accurate micro-reflection time delay is important. Condition 2. (Subset)
8. Qualifying an upstream for a wider channel RF bandwidth or higher order modulation. Qualification can also be done to verify a service level agreement (SLA). States 1 and 2.

Actors: Initially NOC, then technicians

External Event: Distortion red, other indicators yellow/green

Use higher order modulation and wide bandwidth channel as reference

Generally Proactive/High distortion detected

Group of CMs

Expected Outcome: Node qualified

9. CSR quick check. See if something abnormal is going on, and if the neighbors have the same problem. All Conditions

Actors: CSR, Technician

External Event: Customer complaint

Reactive/High distortion detected

1 device possible others

Expected Outcome: Trouble ticket or no distortion problem, service call avoided

10. Trend and Intermittent Issues Tracking. State 1 becoming State 2

Actors: NOC

External Event: Distortion turning red within predetermined period, significant delta distortion within short time period, other indicators yellow/green

Proactive/Change in distortion detected

Group of CMs

Expected Outcome: Problem identified and located, trouble ticket generated

11. Resolving flapping problems and unstable equalization solutions. State 3 or State 4.

Actors: NOC

External Event: CMTS not compensating or improperly compensating for distortion, high uncorrectable FEC statistics, high impulse noise scenario

Reactive/CMTS change in distortion detected

Group of CMs (if distortion generated), all CMs in node (if impulse generated)

Expected Outcome: Problem identified, located and resolved

6.9 Post-equalization

As its name indicates, post-equalization is the process of distortion compensation after the signal has been received. In the cable environment, post-equalization has been used mostly in the downstream direction where a CPE device always receives a continuous signal coming from one location, usually the headend. That fixed location signal source simplified downstream distortion compensation implementation using adaptive equalization at the receiver. In the upstream direction the upstream signals come from multiple sources. Transmission is bursty and one burst may have suffered a different distortion from the next because the upstream paths traversed may be different. A post-equalization approach would require the receiver to compensate for distortion on a per burst basis. Early in DOCSIS, it was deemed that a post-equalization approach would be a significant processing burden on the CMTS, which led to the implementation of the upstream pre-equalization approach covered in this document. Processing capabilities improvements have enabled the implementation of per burst equalization which is now common in CMTSs. The advantages and disadvantages of pre- versus post-equalization are discussed next.

6.9.1 Advantages of Pre-equalization

Reliable diagnostic transmission analysis is enabled by the regular complex coefficient updates the CMTS equalizer provides to individual cable modems when adaptive pre-equalization is enabled. The repeated convolution of

coefficients effectively generates a time averaged response of the upstream channel, which leads to the following advantages:

- Reasonably stable per modem equalized MER as reported by either the CMTS or SNMP Polling
- Ability to view all of the linear transmission path characteristics unique to the individual cable modem due to the complex coefficient impulse response reported via SNMP MIB query. From the SNMP MIB query one can also compute the effective unequalized MER in addition to amplitude and group delay characteristics, as well as enable analysis of micro-reflection conditions for any given cable modem.
- Superior equalized MER capability compared to that provided by post-equalization when the channel is at the upper edge of the return path spectrum. The superiority is also a function of the depth of the return path amplifier cascade
- Another advantage of pre-equalization over post-equalization is that throughput is somewhat higher when pre-equalization is used because there is no need to equalize the data transmission burst. This is related to the difference in the length of the data burst preamble.

Pre-equalization is beneficial to virtually all DOCSIS 2.0 and DOCSIS 3.0 services. In some cases there would appear to be some instability or inconsistency in the metrics being reported when using pre-equalization. Some causes for such behavior are discussed next.

6.9.2 Disadvantages of Pre-equalization

Pre-equalization disadvantages are worth mentioning although they can be considered minor compared to the advantages.

Memory

Operation of pre-equalization has memory associated with the function, and that memory lives in the convolution of the equalizer updates in the cable modem. As stated previously there are many benefits to having this function exist in the cable modem, but there are a few deficiencies as well.

For example, if there is suddenly a major impulse noise event or a change in ingress characteristics during a station maintenance transmission, when the CMTS is assessing the equalized impulse response update to send back to any given modem, there is the possibility that the station maintenance will be corrupted. That would cause the CMTS equalized impulse update to become corrupted, and the corrupted update would be transmitted to the cable modem. The cable modem will convolve the corrupted update, so now that particular cable modem's equalized impulse response will no longer be valid for a period of time.

One weakness of pre-equalization is that the CMTS must be assumed to be responsible as the final arbiter as to which equalized updates are sent and which updates are potentially corrupted, and compensate for those differences in one manner or another. There are indeed many options available to the CMTS to deal with equalizer corruptions, but that topic is beyond the scope of this document.

CM TX Level Increased

One unfortunate but direct consequence of the transmit pre-equalization process occurs when there is significant amplitude roll-off present on a given channel such as at the very end of the cascade, and on the highest available carrier frequency. A primary function of pre-equalization is to eliminate the post equalization noise enhancement that occurs when raising the signal power over a specific portion of the spectrum and thereby raising the flat noise floor of the receiver right along with it.

Equalizer noise enhancement elimination is the single biggest improvement that pre-equalization provides over post-equalization, especially when operating near the upper band edge. Transmit pre-equalization eliminates post-equalization noise enhancement by hitting the receiver with an ideally flat spectrum. This happens even as the equalization process compensates for other impairments such as group delay distortion.

The very act of updating the cable modem to increase the transmit power in the heavily attenuated or rolloff portion of the spectrum at the expense of lowering the power in the less attenuated portion of the spectrum eventually ends up with a perfectly flat response at the input to the CMTS receiver. The final amount the cable modem may be instructed to increase the transmit level over the original non pre-equalized transmit signal power is a function of how much amplitude roll off exists on the channel.

If there is sufficient headroom for the cable modem to increase its upstream transmit level, then there is no penalty for the significantly improved performance. If the cable modem does not have sufficient margin to increase its transmit level to the correct amount then the CMTS will only allow as much low input signal level to exist as was defined in the CMTS's operating configuration.

6.9.3 Advantages of Post-equalization

The following summarizes some of the benefits of post-equalization.

No Memory

Since post-equalization operation starts off every new data burst from scratch, there are no pre-stored conditions required. As such, there is no memory condition to become corrupted from any given error condition.

In other words, while any given post-equalized burst can be corrupted, all that is lost during that corruption period is that particular burst, and not, say, multiple bursts until a station maintenance can correct the corrupted update from the CMTS.

Because the post-equalization process possesses no memory in the operation, one might argue that the post-equalization mode of operation is potentially more stable, or the simplicity of the post-equalization process results in a more stable performance loop over all. This is particularly true in the below 20 MHz frequency range where both ingress and impulse noise are common.

Lower Cable Modem Transmit Levels

At the upper end of the return path spectrum there is generally no mechanism that will result in the post-equalization process requesting more cable modem transmit level, at least as a result of the equalization process itself. This is because the post-equalization process requires absolutely no cooperation on behalf of the cable modem. That is, there is no pre-equalization function being convolved in the cable modem for post-equalization to work.

A reasonably strong case can be argued for using post-equalization only when in TDMA or A-TDMA mode below 30 MHz, and definitely below 20 MHz.

6.9.4 Disadvantages of Post-equalization

The following summarizes several disadvantages of post-equalization.

- Lower throughput than pre-equalization operation due to longer data preamble – the throughput loss is mostly a small amount compared to FEC correction capabilities variations, but nonetheless it is measurable.
- Poorer equalized MER performance compared to pre-equalization because of the convergence time required, versus a data preamble length most operators would be willing to suffer with to attain better

equalized MER performance per burst. Also, given the possibility of multiple impairments being present during any given data burst, there can be a significant variation regarding the equalized MER estimate from burst to burst.

- Vastly poorer micro-reflection analysis capabilities in part because of convergence time and in part because of any impairment present during the data burst.
- Equalizer noise enhancement phenomena are unavoidable at the diplex filter band-edge, and the phenomena become increasingly worse as the amplifier cascade increases. In other words, in an N + 1 or perhaps even up to an N + 3 cascade depth, the equalizer noise enhancement penalty is perhaps a minor point that could be forgotten. However, as the cascade depth increases to N + 6 or higher, the equalizer noise enhancement penalty becomes more significant to the point that the post-equalization mode of operation total available or usable bandwidth on a given return path becomes much less than when using pre-equalization.
- Generally speaking, much poorer estimate for both amplitude roll-off and group delay distortion when attempting to extract the information from a single equalized data burst for any given cable modem.

6.9.5 Pre- and Post-equalization Measurements: Difference in Performance

Cascade Depth	US RX Lvl	IP Address	MAC Address
N + 2	5.0	150.31.230.15	0011.aeff.f9c2
N + 3	4.7	150.31.230.14	0011.aeff.f9dc
N + 4	4.3	150.31.230.13	0012.2503.52ba
N + 3	4.6	150.31.230.12	0012.c90b.366a
N + 2	5.1	150.31.230.11	0012.c90f.b61c
N + 1	5.0	150.31.230.10	0012.c9f6.e030

PRE-EQ – 36.5 MHz		CH BW	=	6.4 MHz
MAC Address	Cascade Depth	RXPwr(dBmV)	EQ-MER(dB)	
0011.aeff.f9c2	N + 2	5.3	35.1	
0011.aeff.f9dc	N + 3	4.9	35.1	
0012.2503.52ba	N + 4	4.6	34.8	
0012.c90b.366a	N + 3	4.9	34.8	
0012.c90f.b61c	N + 2	4.8	35.1	
0012.c9f6.e030	N + 1	5.0	35.3	

POST-EQ – 36.5 MHz		CH BW	=	6.4 MHz
MAC Address	Cascade Depth	RXPwr(dBmV)	EQ-MER(dB)	
0011.aeff.f9c2	N + 2	5.3	33.0	
0011.aeff.f9dc	N + 3	4.9	31.6	
0012.2503.52ba	N + 4	4.6	30.4	
0012.c90b.366a	N + 3	4.9	31.5	
0012.c90f.b61c	N + 2	4.8	32.7	
0012.c9f6.e030	N + 1	5.0	35.3	

PRE-EQ – 36.8 MHz		CH BW	=	6.4 MHz	
MAC Address	Cascade Depth	RXPwr(dBmV)	EQ-MER(dB)	CM TX Lvl	
0011.aeff.f9c2	N + 2	5.3	35.1	46.3	
0011.aeff.f9dc	N + 3	4.9	34.0	53.7	
0012.2503.52ba	N + 4	4.0	33.3	55.2	
0012.c90b.366a	N + 3	4.9	34.2	53.6	
0012.c90f.b61c	N + 2	4.8	35.1	44.2	
0012.c9f6.e030	N + 1	5.0	35.3	37.4	

MAC Address	Cascade Depth	POST-EQ – 36.8 MHz RXPwr(dBmV)	CH BW EQ-MER(dB)	=	6.4 MHz CM TX Lvl
0011.aeff.f9c2	N + 2	5.3	31.5		45.7
0011.aeff.f9dc	N + 3	4.9	30.4		39.9
0012.2503.52ba	N + 4	4.6	28.6		55.2
0012.c90b.366a	N + 3	4.9	30.0		52.4
0012.c90f.b61c	N + 2	4.8	31.4		43.7
0012.c9f6.e030	N + 1	5.0	33.6		37.2

7 INTEGRATION OF DOCSIS NETWORK HEALTH METRICS WITH PRE-EQUALIZATION METRICS AND PROACTIVE NETWORK MAINTENANCE APPROACH

This document thus far has focused on the pre-equalization metrics and processes that are leveraged in the implementation of a proactive network maintenance approach. The mechanisms discussed here can be further enhanced by taking advantage of additional intelligence obtained from other DOCSIS network health and configuration parameters, such as forward error correction statistics and configuration settings, operational modes and interleaver configurations, transmit and receive power level statistics, and the time dependencies of these parameters. In addition to these parameters and statistics, the upstream spectrum analysis capabilities of CMTS upstream receivers can be used to provide a detail view of the spectrum within a channel when no DOCSIS bursts are present and of the entire upstream spectrum. Downstream receive equalization parameters can be leveraged to troubleshoot impairments in a similar fashion as pre-equalization parameters were used in the upstream. The following sections explore these different mechanisms and how they can further enhance the proposed proactive network maintenance approach.

7.1 DOCSIS Robustness Against Impulse and Noise

Impulses in cable networks can have many origins, but most come from external sources and enter the cable plant via common ingress mechanisms. They can be generated when cable network defects such as damaged cables are in close proximity to sources including electric motors, appliances, thermostats, arc welders, defective light fixtures, power line noise (arcing, sparking), static from lightning, and so forth. Impulses are short in duration, typically less than 20 microseconds (μs), and can fully cover the upstream spectrum.

Impulse repetition rates can vary, but they are typically in the order of a few kilohertz (kHz). In the time domain, one characteristic of an impulse is its exponential decay. Figure 37 shows a time domain representation of an impulse. A narrow peak is an indicator of high-frequency components, while a longer duration impulse is indicative of lower-frequency components. These time domain impulse characteristics are consistent with the wideband spectral nature of impulses.

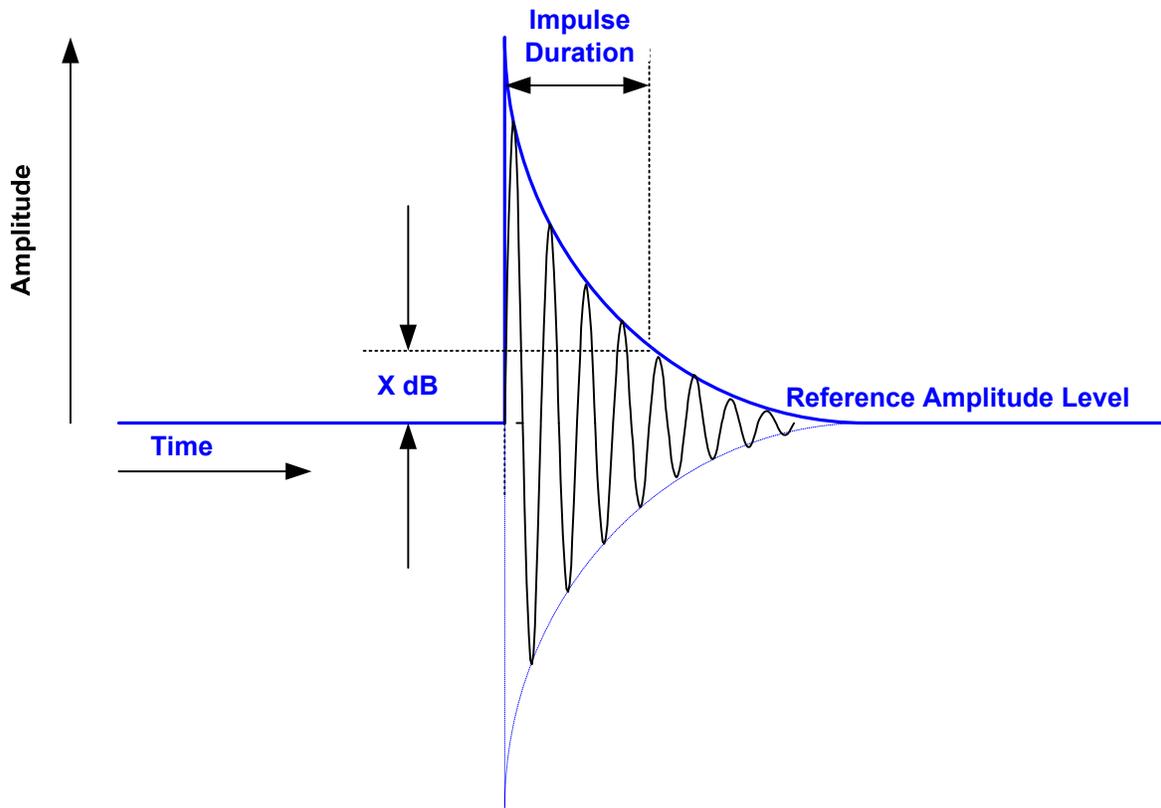


Figure 37 - Time domain representation of typical impulse

Figure 38 shows an impulse measured in the time domain (lower trace) and frequency domain (upper trace), confirming the wideband nature in frequency as well as its exponential decay in time. In this actual measurement, the characteristics of the upstream path modify the view of the impulse in time and frequency. A traversed micro-reflection, for example, would exhibit ripples in the frequency domain and a second peak in the time domain.

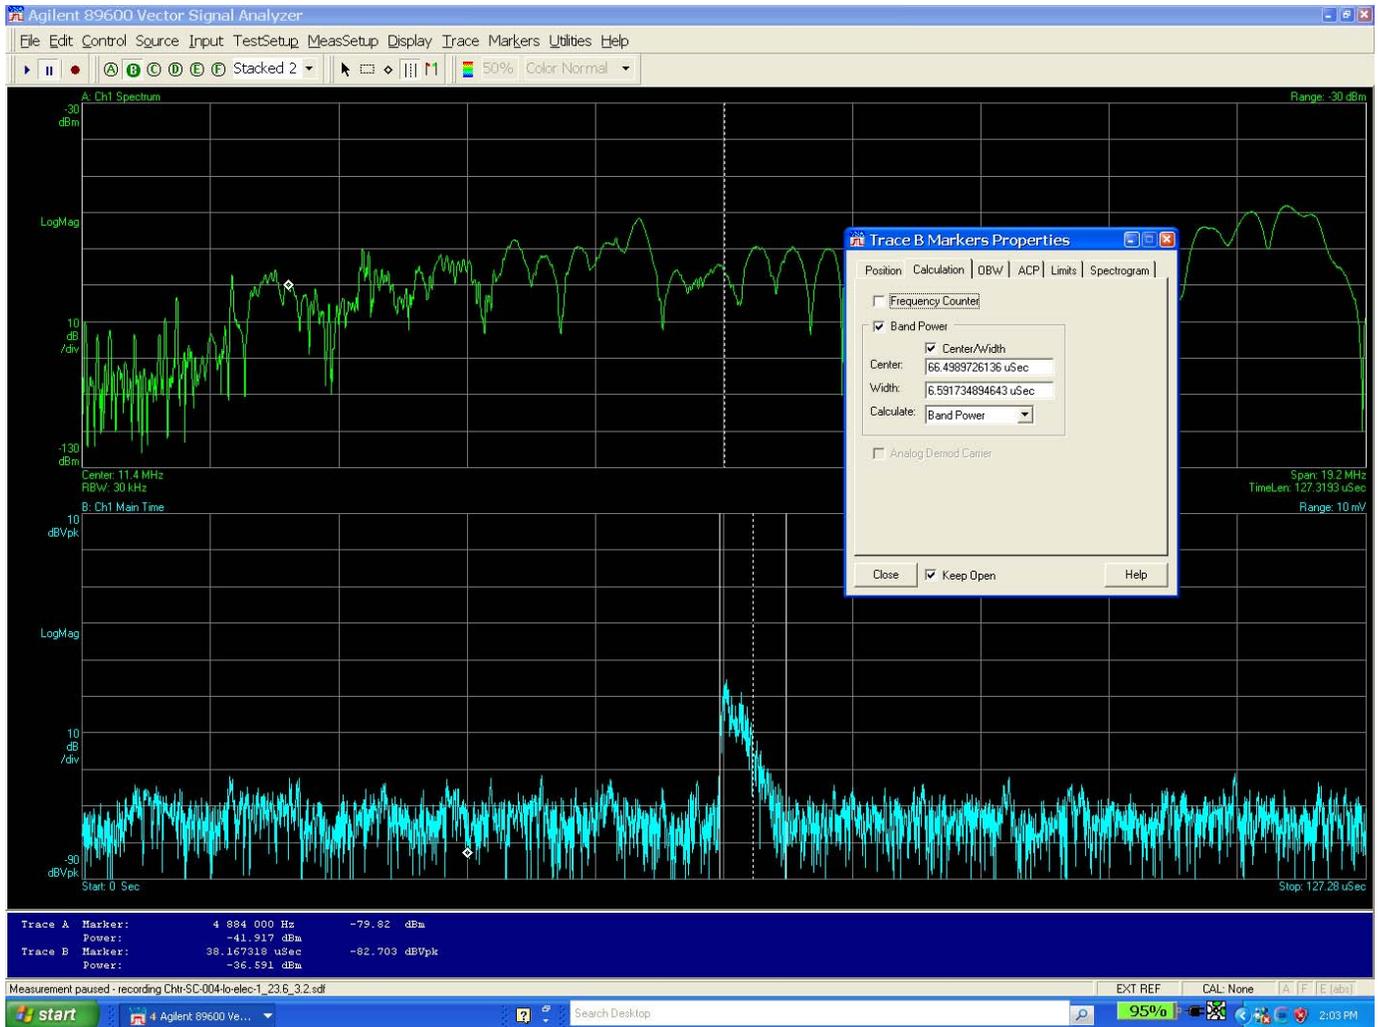


Figure 38 - Frequency and time capture of typical impulse

7.1.1 Impulse Noise Statistics

Key parameters describing impulse noise impairment to quantify its impact, as well as to assess how to better counteract or mitigate its effects, are impulse duration, impulse repetition rate, and impulse power level.

Figure 39 and Figure 40 illustrate the duration and power level statistics of impulses measured across 346 nodes. These statistics do not include upstream bursts generated by overdriving upstream lasers and/or reverse amplifiers. From a duration perspective, it can be observed that most impulses have a duration that lasts less than 10 microseconds. This can provide guidance regarding interleaver design and configuration. From a power perspective, a wide range of power levels hit the upstream receiver, although the majority of those are between -65 dBm and -55 dBm (-16.25 to -6.25 dBmV), which correspond to levels between 23 dB and 13 dB above the noise floor.

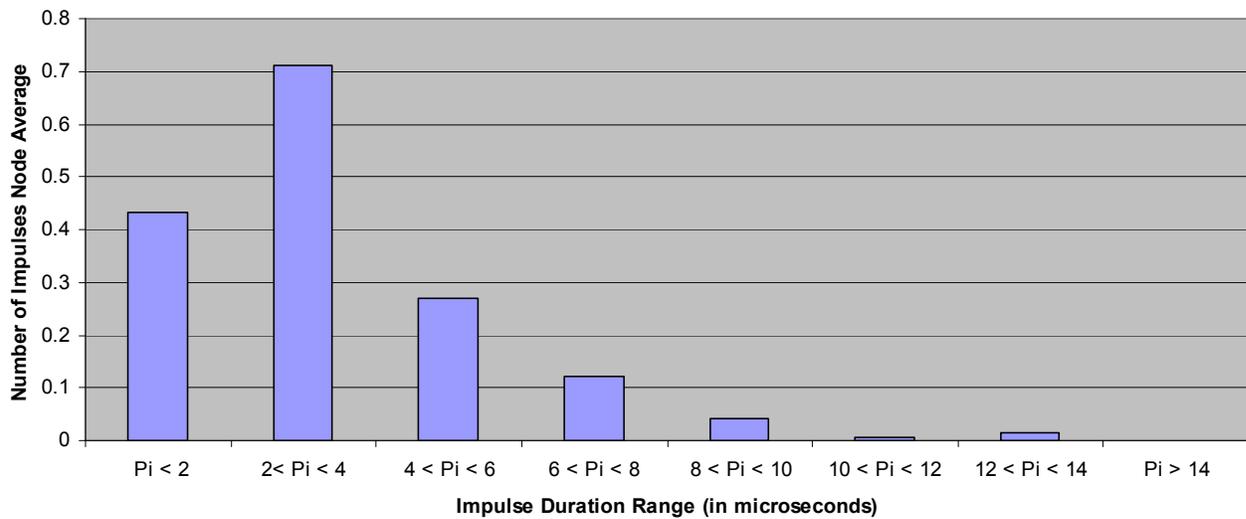


Figure 39 - Impulse duration characteristics

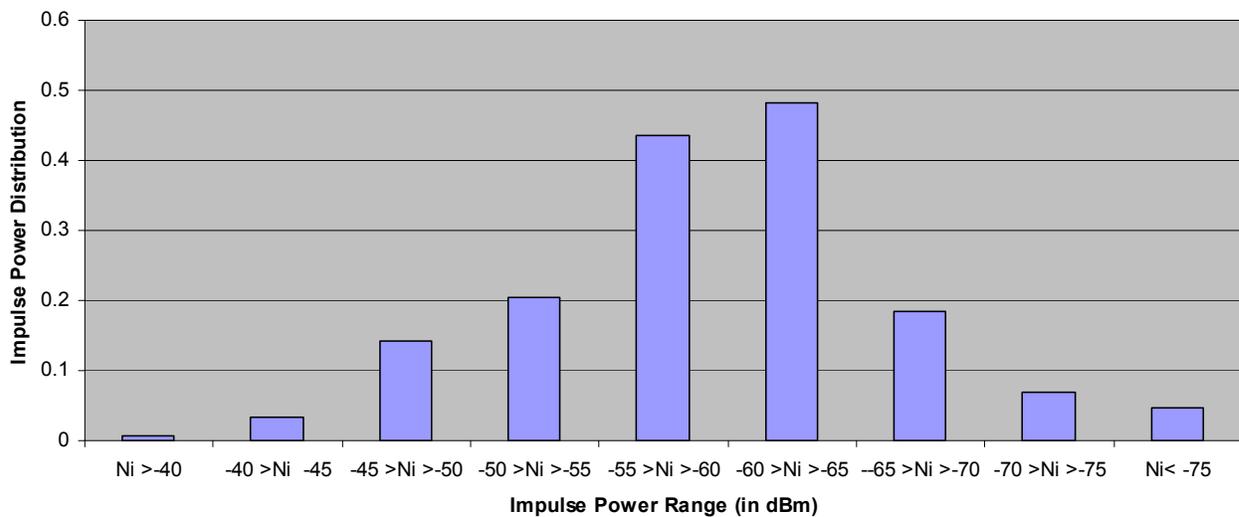


Figure 40 - Impulse power level distribution characteristics

By selecting the appropriate impulse/burst robustness configuration settings, a node that shows significant noise through standard spectrum analysis, observed by turning the test equipment's MAX-HOLD on, could in fact be operating without uncorrectable errors when leveraging the benefits of interleaving and FEC described next.

7.1.2 Impulse and Burst Robustness for the Different DOCSIS Modes of Operation

The upstream transmissions in DOCSIS can take place using either of two modes of operation. One is time division multiple access or TDMA - commonly called advanced or A-TDMA in DOCSIS 2.0 and later - and the other is synchronous code division multiple access or S-CDMA. In TDMA, the transmissions are sequential with only one CM transmitting at a time as directed by the CMTS scheduler. In S-CDMA, the information can be carried in up to 128 orthogonal codes that are transmitted simultaneously. This can result in transmissions coming from more than

one CM at a time. In S-CDMA the CMTS manages transmissions across two dimensions, in the time domain and in the code domain. There are many common physical layer parameters, such as channel width, excess bandwidth, modulation order, codeword length, and number of corrected bytes per codeword, as well as some mode-specific parameters like block size for TDMA, and modulation intervals per frame and number of codes per subframe for S-CDMA. These parameters will impact the burst robustness to different degrees.

7.1.2.1 TDMA Impulse and Burst Noise Robustness

When using TDMA, the impulse robustness that can be achieved without the use of interleaving will depend on the Reed-Solomon FEC configuration available in DOCSIS. The RS-FEC configuration in DOCSIS consists of codeword lengths = $k+2T$ that can be up to 255 bytes in length. The codeword payload length in bytes (not containing parity) is k and the codeword portion containing parity information is $2T$, where T is the number of bytes that can be corrected per codeword. In a RS-FEC long packet configuration scenario of $k = 220$ and $T = 16$ for 3.2 MHz channel bandwidth and 16-QAM, the 16 correctable bytes correspond to a burst protection of 32 symbols or 12.512 μs ($32 * 0.391 \mu\text{s}$). A short packet configuration scenario of $k = 78$ and $T = 8$ for 3.2 MHz channel bandwidth and 16-QAM, the burst protection is 6.256 μs ($16 * 0.391 \mu\text{s}$). When operating in 64-QAM and 6.4 MHz channel bandwidth, we have that 16 bytes represents 21 ($16 * 8/6$) symbols and 8 bytes represents 10 ($8 * 8/6$) symbols. The burst protection for the long packet scenario can be as high as 4.095 μs ($21 * 0.195 \mu\text{s}$) and as high as 1.95 μs ($10 * 0.195 \mu\text{s}$) for the short packet scenario. In case the actual burst crosses symbol boundaries, the burst protection capabilities are reduced by 1 symbol (Figure 41).

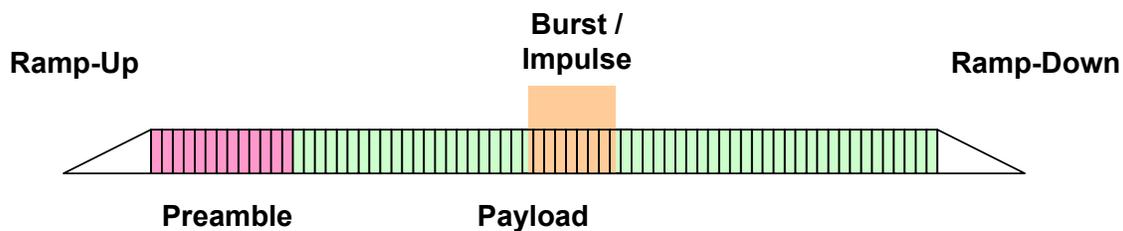


Figure 41 - Burst/impulse noise lasting N symbol periods requires burst protection of $N+1$ symbols if the burst crosses symbol boundaries

As can be seen from the impulse noise statistics presented previously, a significant percentage of impulses will not be compensated for in 64-QAM operation just through the use of FEC.

7.1.2.2 TDMA Burst Protection with Interleaving

One option in TDMA operation is to use byte interleaving in conjunction with FEC to improve robustness against burst or impulse noise. In TDMA, burst robustness is achieved through block interleaving. Upstream interleaving can be switched on and off. Interleaving multiplies the effectiveness of FEC by aggregating multiple codewords in a block and changing the order in which bytes are transmitted through a shuffling or interleaving process of the codeword bytes so that a burst in time does not predominantly impact a single codeword. Instead, the impact is spread among all the codewords contained in a block, thereby providing enhanced protection.

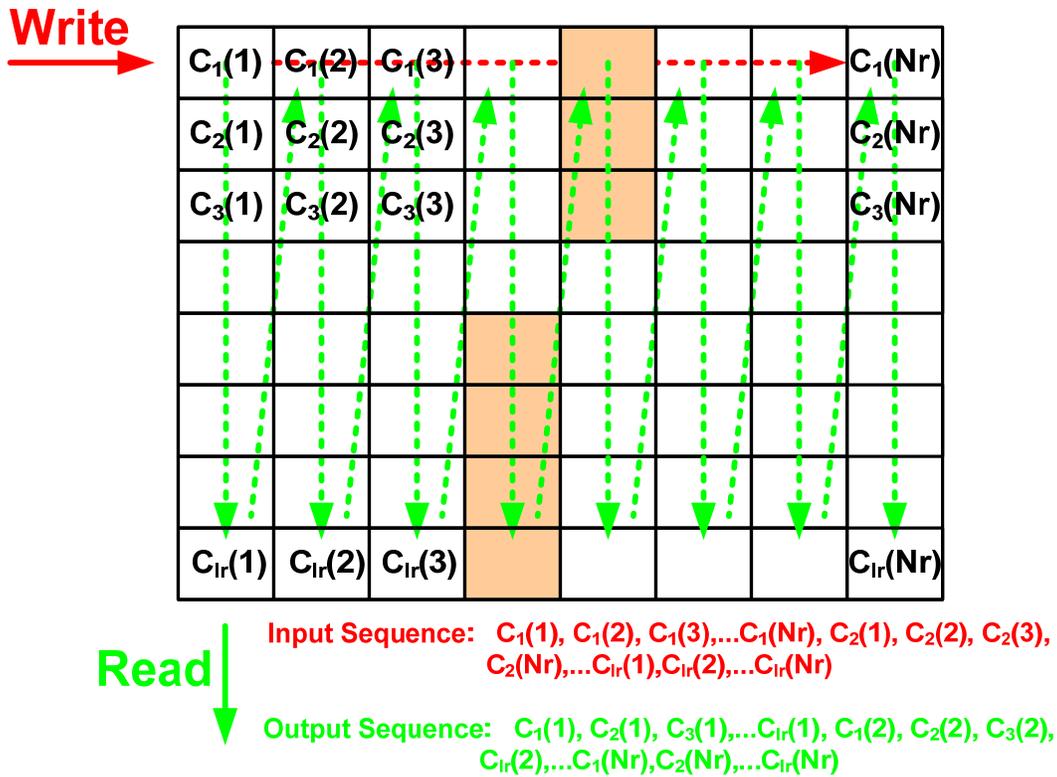


Figure 42 - Byte Interleaver Operation

Figure 42 shows the interleaving of codewords in a block so that the burst impact spreads across the different codewords in that block. The codewords are loaded row by row horizontally into the block but are read one column at a time (vertically), generating the interleaving of bytes. The output sequence shows how the burst spreads across different codewords.

Since the interleaving process holds a number of codewords in the interleaver block before reading, it generates an added latency. The interleaver block in DOCSIS can be up to 2048 bytes in length.

In TDMA, the use of large blocks provides more burst robustness at the expense of latency because it can contain more codewords. Unfortunately, when transmitting shorter packets, the block size is limited, resulting in limited burst robustness. The latency overhead has to be counted in both the interleaving process at the transmitter, as well as the de-interleaving process at the receiver.

In the case of operation using dynamic interleaving mode, the interleaver depth is chosen dynamically to achieve optimum burst robustness. It is advisable for cable operators to take advantage of DOCSIS optimizations such as shortened last codeword in the FEC configuration and dynamic interleaving mode when turning interleaving mode on. In the analysis that follows, it is assumed that these DOCSIS optimizations are always on.

Figure 43, Figure 44 and Figure 45 estimate the burst robustness that can be achieved with different FEC configurations for different packet sizes as well as the added latency incurred through block interleaving.

FEC Settings				Nr=k+2T	80	92	60	72	42	52	40	32	RS Codeword Length							
				k	68	68	48	48	32	32	32	16	CW-w/o-correc.							
				T	6	12	6	12	5	10	4	8	Correctable Errors /CW							
Bit Rate	Symbol Rate kspss	Symbol Period microsec.	Symbols /Byte	Upstream Burst Protection & Dynamic Interleaving and De-Interleaving Latency Overhead (in µs)																
				160	184	120	144	126	156	120	160	Int. Block Size (Bytes)	Int. Depth (#CW/Block)							
2.56	1280	0.781	4.00	38	588	75	738	38	588	75	738	47	625	94	813	38	588	125	938	QPSK
5.12	2560	0.391	4.00	19	294	38	369	19	294	38	369	23	313	47	406	19	294	63	469	
10.24	5120	0.195	4.00	9	147	19	184	9	147	19	184	12	156	23	203	9	147	31	234	
5.12	1280	0.781	2.00	19	294	38	369	19	294	38	369	23	313	47	406	19	294	63	469	16QAM
10.24	2560	0.391	2.00	9	147	19	184	9	147	19	184	12	156	23	203	9	147	31	234	
20.48	5120	0.195	2.00	5	73	9	92	5	73	9	92	6	78	12	102	5	73	16	117	
7.68	1280	0.781	1.33	13	196	25	246	13	196	25	246	16	208	31	271	13	196	42	313	64QAM
15.36	2560	0.391	1.33	6	98	13	123	6	98	13	123	8	104	16	135	6	98	21	156	
30.72	5120	0.195	1.33	3	49	6	61	3	49	6	61	4	52	8	68	3	49	10	78	

Figure 43 - Burst robustness (in red) and added latency (in blue) for a 64 byte packet using different FEC and interleaver block size settings

FEC Settings				Nr=k+2T	255	204	174	116	90	72	36	18	RS Codeword Length							
				k	223	176	150	100	78	64	32	16	CW-w/o-correc.							
				T	16	14	12	8	6	4	2	1	Correctable Errors /CW							
Bit Rate	Symbol Rate kspss	Symbol Period microsec.	Symbols /Byte	Upstream Burst Protection & Dynamic Interleaving and De-Interleaving Latency Overhead (in µs)																
				510	408	348	348	360	360	324	306	Int. Block Size (Bytes)	Int. Depth (#CW/Block)							
2.56	1280	0.781	4.00	100	2038	88	1988	75	1938	75	1938	75	1938	63	1888	56	1863	53	1850	QPSK
5.12	2560	0.391	4.00	50	1019	44	994	38	969	38	969	38	969	31	944	28	931	27	925	
10.24	5120	0.195	4.00	25	509	22	497	19	484	19	484	19	484	16	472	14	466	13	463	
5.12	1280	0.781	2.00	50	1019	44	994	38	969	38	969	38	969	31	944	28	931	27	925	16QAM
10.24	2560	0.391	2.00	25	509	22	497	19	484	19	484	19	484	16	472	14	466	13	463	
20.48	5120	0.195	2.00	13	255	11	248	9	242	9	242	9	242	8	236	7	233	7	231	
7.68	1280	0.781	1.33	33	679	29	663	25	646	25	646	25	646	21	629	19	621	18	617	64QAM
15.36	2560	0.391	1.33	17	340	15	331	13	323	13	323	13	323	10	315	9	310	9	308	
30.72	5120	0.195	1.33	8	170	7	166	6	161	6	161	6	161	5	157	5	155	4	154	

Figure 44 - Burst robustness (in red) and added latency (in blue) for a 256 byte packet using different FEC and interleaver block size settings

FEC Settings				Nr=k+2T	255	204	174	116	90	72	36	18	RS Codeword Length							
				k	223	176	150	100	78	64	32	16	CW-w/o-correc.							
				T	16	14	12	8	6	4	2	1	Correctable Errors /CW							
Bit Rate	Symbol Rate kspss	Symbol Period microsec.	Symbols /Byte	Upstream Burst Protection & Dynamic Interleaving and De-Interleaving Latency Overhead (in µs)																
				1785	1836	1914	1856	1800	1728	1728	1728	Int. Block Size (Bytes)	Int. Depth (#CW/Block)							
2.56	1280	0.781	4.00	350	10925	394	11100	413	11175	400	11125	375	11025	300	10725	300	10725	300	10725	QPSK
5.12	2560	0.391	4.00	175	5463	197	5550	206	5588	200	5563	188	5513	150	5363	150	5363	150	5363	
10.24	5120	0.195	4.00	88	2731	98	2775	103	2794	100	2781	94	2756	75	2681	75	2681	75	2681	
5.12	1280	0.781	2.00	175	5463	197	5550	206	5588	200	5563	188	5513	150	5363	150	5363	150	5363	16QAM
10.24	2560	0.391	2.00	88	2731	98	2775	103	2794	100	2781	94	2756	75	2681	75	2681	75	2681	
20.48	5120	0.195	2.00	44	1366	49	1388	52	1397	50	1391	47	1378	38	1341	38	1341	38	1341	
7.68	1280	0.781	1.33	117	3642	131	3700	138	3725	133	3708	125	3675	100	3575	100	3575	100	3575	64QAM
15.36	2560	0.391	1.33	58	1821	66	1850	69	1863	67	1854	63	1838	50	1788	50	1788	50	1788	
30.72	5120	0.195	1.33	29	910	33	925	34	931	33	927	31	919	25	894	25	894	25	894	

Figure 45 - Burst robustness (in red) and added latency (in blue) for a 1518 byte packet using different FEC and interleaver block size settings

Figure 43, Figure 44, and Figure 45 highlight the sensitivity of burst robustness to packet size that exists in TDMA even with interleaving turned on. Nevertheless, if a cable operator is willing to incur some additional efficiency overhead, enough robustness can be achieved for typical impulse durations that are observed in cable networks. The

impulse statistics in Figure 39 show that most impulses last less than 10 μs , which is an impulse duration robustness that is still achievable for shorter packets in wider channels and higher modulation orders.

7.1.2.3 Impulse and Burst Protection in S-CDMA Mode

In S-CDMA transmission, information is carried in time and code domains. Orthogonality among codes is required for simultaneous transmission of the 128 codes. Each symbol transmitted is spread across 128 codes. The time required to transmit one symbol per code across all 128 codes is the spreading interval. Transmissions in S-CDMA are carried over a frame that consists of K spreading intervals and up to 128 codes. In an S-CDMA frame, the information is grouped in minislots by defining codes per minislot, and in subframes by defining codes per subframe. Interleaving in an S-CDMA channel is constrained within subframe boundaries. A subframe is typically an integer number of Reed-Solomon codewords. The impulse protection capability in S-CDMA is measured in spreading intervals across the block defined by the subframe. The number of bytes that can be corrected by the number of codewords within a subframe has to be translated into symbols contained in the columns of a subframe that are impacted by the impulse. Figure 46 shows a frame divided in subframes and a burst impact over multiple spreading intervals. The latency overhead process in S-CDMA is somewhat similar to TDMA block interleaving, as the frame has to be assembled first before transmitting and has to be assembled at the receiver prior to decoding, thereby generating an added latency equal to twice the frame size.

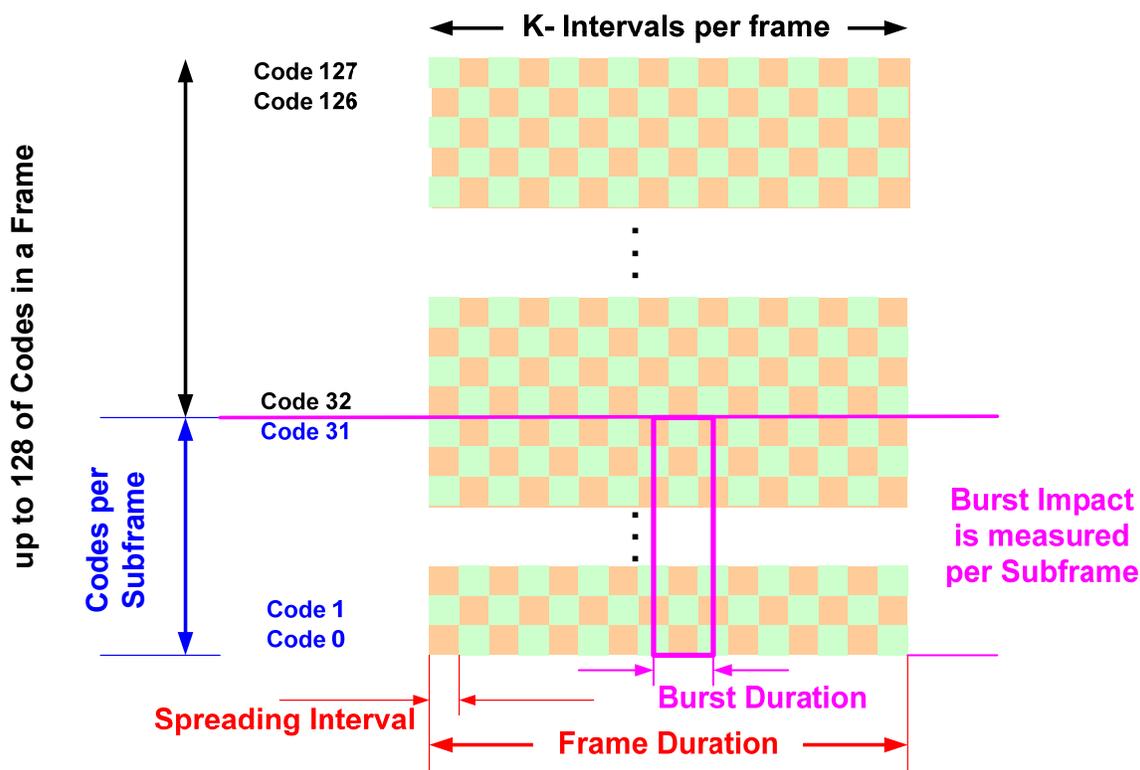


Figure 46 - Burst impact on S-CDMA subframe

Figure 47 and Figure 48 highlight the burst robustness capabilities that can be obtained in S-CDMA under different S-CDMA frame and FEC configuration settings.

Active Codes	126	126	126	126	126	126	126	126	126	126	126
Frame Length k	32	32	32	16	32	32	32	16	16	32	64
Codes per Subframe	7	7	7	14	14	28	28	56	28	14	28
FEC k	92	88	80	80	80	80	192	192	192	192	192
FEC T	10	12	16	16	16	16	16	16	16	16	16
Max # Symbols per Subframe	224	224	224	224	448	896	896	896	448	448	1792
# of FEC Symbols per codeword	224	224	224	224	224	224	448	448	448	448	448
Payload Symbol Efficiency	82%	79%	71%	71%	71%	71%	86%	86%	86%	86%	86%
# of codewords per subframe	1	1	1	1	2	4	2	2	1	1	4
Remaining FEC symbols in subframe	0	0	0	0	0	0	0	0	0	0	0
Corrected symbols by FEC	20	24	32	32	64	128	64	64	32	32	128
Intervals corrected in Subframe	2	3	4	2	4	4	2	1	1	2	4
Guaranteed burst robustness in μ s	50	100	150	50	150	150	50	0	0	50	150
S-CDMA Tx & Rx Frame Latency in μ s	3200	3200	3200	1600	3200	3200	3200	1600	1600	3200	6400

Figure 47 - Burst robustness in microseconds for 16-QAM, 3.2 MHz bandwidth channel with TCM off and diverse FEC and subframe settings (packet occupies at least one subframe)

Active Codes	126	126	126	126	126	126	126	126	126	126	126
Frame Length k	32	32	32	16	32	16	32	16	16	32	64
Codes per Subframe	7	7	7	14	14	28	28	56	28	14	28
FEC k	136	144	148	148	136	80	192	192	136	136	192
FEC T	16	12	10	10	16	16	16	16	16	16	16
Max # Symbols per Subframe	224	224	224	224	448	448	896	896	448	448	1792
# of FEC Symbols per codeword	224	224	224	224	224	149.3	298.7	298.7	224	224	298.7
Payload Symbol Efficiency	81%	86%	88%	88%	81%	71%	86%	86%	81%	81%	86%
# of codewords per subframe	1	1	1	1	2	3	3	3	2	2	6
Remaining FEC symbols in subframe	0	0	0	0	0	0	0	0	0	0	0
Corrected symbols by FEC	21.33	16	13.33	13.33	42.67	64	64	64	42.67	42.67	128
Intervals corrected in Subframe	3	2	1	0	3	2	2	1	1	3	4
Guaranteed burst robustness in μ s	50	25	0	0	50	25	25	0	0	50	75
S-CDMA Tx & Rx Frame Latency in μ s	1600	1600	1600	800	1600	800	1600	800	800	1600	3200

Figure 48 - Burst robustness in microseconds for 64-QAM, 6.4 MHz bandwidth channel with TCM off and diverse FEC and subframe settings (packet occupies at least one subframe)

As can be seen when comparing Figure 47 and Figure 48 with Figure 43, Figure 44, and Figure 45, the burst robustness that can be achieved with S-CDMA is significantly greater than what can be achieved with interleaving in TDMA. In S-CDMA though, one must be careful, because the FEC and subframe configuration should be such that at least two spreading interval columns in a subframe can be corrected to guarantee any robustness. Figure 48 includes some examples where only one spreading interval column is protected. This cannot guarantee burst or impulse protection and will only protect in cases where the impulse is shorter than a spreading interval and it does not cross spreading interval boundaries.

8 CONCLUSION

This best practices and guidelines report covers different aspects of network maintenance that rely on upstream pre-equalization to proactively maintain the CATV network. Even though the initial purpose of this effort was to implement a network maintenance strategy that is able to take corrective action before service is impacted, the outcome of this effort has also proven to be very powerful in implementing and optimizing numerous “reactive” maintenance tasks present in day to day operations. Reactive maintenance is defined here as maintenance triggered by any perceivable deterioration in service performance.

Linear distortion including two types of micro-reflections and group delay distortion have been discussed in detail, as well as the ranging process that CMTS and CM go through to sound the upstream channel and compute the equalization coefficients. From these parameters the CM information indicates the estimated distortion in the upstream path while the CMTS indicates the extent to which the distortion compensation of the upstream path of a CM has been completed.

Interpretation of the CM/CMTS MIB data has been provided including the implementation details of a universal decoder that translates all known CM MIB interpretations.

This report describes how to take advantage of pre-equalization analysis both in the time and frequency domains. Other topics covered here include scalable data collection approaches; distortion compensation capabilities of DOCSIS pre-equalization; pre-equalization calibration techniques, which included the details of a short reference plant; severity assessment for static, trending and intermittent distortion environments; definitions of key equalization metrics; and extrapolation of coefficients for neighboring channels.

A key outcome of this effort is the fault localization processes to pinpoint the problem location which leads to reduced mean time to repair and improved reliability. This is obtained through the correlation of the CATV network topology with the impairment-unique characteristics derived from the CM pre-equalization data.

This gathered knowledge has been leveraged to implement use cases describing guidelines to resolve proactive as well as reactive operational, engineering and maintenance issues.

Even though the emphasis here has been on leveraging pre-equalization, incorporation of additional metrics and the correlation among metrics will increase the efficacy in troubleshooting network problems.

Appendix I Tutorial

I.1 Nonlinear Distortions

Active devices such as amplifiers and optoelectronics are devices that do not operate perfectly linearly. For example, the RF signals at the output of a perfect amplifier would be identical to the input signals, but at a higher power level. The RF signals at the output of a real-world amplifier are increased to a higher power level and are almost identical to the input signals, but are distorted somewhat by nonlinearities in the amplifier. Nonlinear operation in an amplifier is caused by small-signal nonlinearities in the active device's semiconductors, and by signal compression that takes place as higher RF output levels reach the amplifier's saturation point. All of this means that a real-world amplifier's output includes both the amplified signals and distortions. As the amplifier's RF output level is increased to even higher amplitudes, the distortions get worse (think of a stereo that distorts the sound as the volume is turned up too high).

Nonlinear distortions of most interest to equipment manufacturers and cable operators include even order distortions such as composite second order (CSO), odd order distortions such as composite triple beat (CTB), and cross modulation (XMOD). One might think that as an amplifier's RF output level changes by, say, 1 decibel (dB), the amplitude of the distortions also would change by 1 dB. But they do not.

CSO distortion amplitudes change by 2 dB for every 1 dB change in the amplifier's output level, while CTB distortion amplitudes change 3 dB for every 1 dB change in the amplifier's output level.¹ Nonlinear distortions get their name in part because of the "nonlinear" 1:2 and 1:3 signal-to-distortion relationship, as well as from the fact that the distortions are a function of the active device's inherent nonlinear operation. As RF output levels approach or reach the active device's saturation threshold, the distortion amplitudes may begin to change at ratios other than the expected 1:2 or 1:3.

The tendency is to think of nonlinear distortions coming primarily from active devices whose output levels are too high, as just discussed. Common path distortion (CPD) is an interesting class of nonlinear distortion in that it's often generated in a diode-like junction somewhere in the transmission path that is common to both the forward and return, hence the name common path distortion. The culprit is typically corrosion of some sort, where the corrosion's oxide layer itself—which may be only a few molecules thick—behaves somewhat like a diode. As you know, diodes are used to make electronic mixer circuits. The presence of downstream signals at the diode-like corrosion-based "mixer" results in the generation of various second- and third-order distortions. Many of those distortions appear in the return spectrum. For instance, in an NTSC network, CPD's second order beat clusters may appear every 6 MHz: 6, 12, 18, 24, 30, 36 MHz and so on. Third order distortions may appear 1.25 MHz either side of the 6 MHz-spaced second order distortions. Sometimes only the second order distortions are present, sometimes only the third order distortions are present, sometimes both are present, and sometimes CPD manifests itself as an elevated noise floor.

In some circumstances the mechanism that generates CPD may also be a noticeable impedance mismatch that creates a micro-reflection.

Nonlinear distortions can be generated in a cable network's passive devices, separate from the mechanism that creates CPD. One example is passive device intermodulation, which occurs when excessive RF levels saturate the ferrite material in, say, drop splitter toroidal transformers (the ferrite material generally must have some residual magnetism present). Newer passive designs minimize passive device intermodulation.

¹ The *ratios* of desired signal amplitudes to CSO and CTB distortion amplitudes change by 1 dB and 2 dB respectively for each 1 dB change in amplifier output level.

I.2 Linear Distortions

Linear distortions include impairments such as *micro-reflections*, *amplitude ripple/tilt*, and *group delay*. Linear distortions comprise a class of impairments that from a general perspective maintains a linear signal-to-distortion amplitude relationship. For example, if an impedance mismatch creates a micro-reflection, the amplitude of the reflection relative to the incident signal remains constant as the incident signal's amplitude changes. That is, if the incident signal's amplitude changes 1 dB, the amplitude of the reflection also will change 1 dB. This is in contrast to nonlinear distortions, where a 1 dB signal change results in a distortion amplitude change of 2 or 3 dB.

Amplitude ripple/tilt, group delay, impedance mismatch, and micro-reflection are discussed in other sections of this tutorial.

I.3 Frequency Response

Cable industry technical personnel have for decades used broadband sweep equipment comprising a sweep transmitter and sweep receiver to align and maintain the outside plant's active devices. In its simplest configuration, a sweep transmitter transmits a test signal that has constant amplitude versus frequency. The transmitted signal's frequency is varied over a frequency range of interest—for instance, the 54-870 MHz downstream range. The sweep receiver recovers the transmitted sweep signal, and plots the sweep signal's amplitude-versus-frequency on the receiver display. The latter shows a graphic representation of the cable network's frequency response. It's important to note that frequency response is a complex quantity that has two components: amplitude (magnitude)-versus-frequency, and phase-versus-frequency. Broadband sweep equipment displays only amplitude-versus-frequency, the latter more generally known as frequency response in cable industry vernacular.

Ideally, amplitude-versus-frequency response should be flat, and phase-versus-frequency response should change in proportion to frequency.

What happens when amplitude-versus-frequency deviates from ideal? The result may be amplitude ripple (standing waves), amplitude tilt (slope), or some combination of the two. When phase-versus-frequency deviates from ideal, the result is group delay. Refer to the Amplitude Ripple, Amplitude Tilt, and Group Delay tutorial sections for more information.

I.4 Group Delay

Consider the 6 MHz spectrum occupied by an NTSC analog TV channel or digitally modulated signal, the 5-42 MHz upstream spectrum, or any specified bandwidth or passband as the equivalent of a bandpass filter. A signal takes a certain amount of time to pass through a filter. The finite time required for a signal to pass through a filter—or any device for that matter—is called delay. Absolute delay is the delay a signal experiences passing through the device at some reference frequency.

In many instances the delay through a filter varies with frequency—that is, the delay may be less in the center of a filter's passband, but greater near the filter's band edges. Indeed, if delay through a filter is plotted on a graph of frequency (x-axis)-versus-time delay (y-axis), the plot often has a parabola- or bathtub-like shape, as shown in Figure I-1.

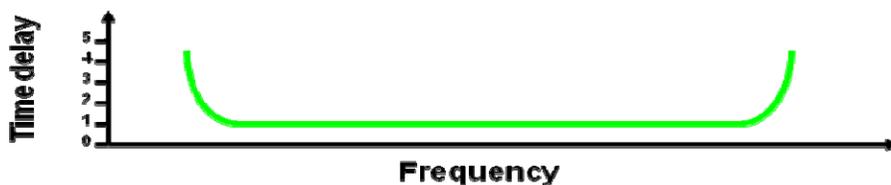


Figure I-1 - A filter's time delay-versus-frequency curve often has a bathtub shape

If propagation or transit time through a device is the same at all frequencies, phase is said to change proportionally with respect to frequency. If phase changes proportionally with frequency, an output signal will be identical to the input signal—except that it will have a time shift because of the uniform delay through the device. If propagation or transit time through a device is different at different frequencies, the result is delay shift or nonlinear phase shift. If phase does not change proportionally with frequency, the output signal will be distorted.

Delay distortion—also known as phase distortion—is usually expressed in units of time: milliseconds (ms), microseconds (μs) or nanoseconds (ns) relative to a reference frequency.

Phase distortion is related to phase delay, and is measured using a parameter called envelope delay distortion, or group delay distortion.

According to the *IEEE Standard Dictionary of Electrical and Electronics Terms*, group delay is “the [negative] derivative of radian phase with respect to radian frequency. It is equal to the phase delay for an ideal non-dispersive delay device, but may differ greatly in actual devices where there is a ripple in the phase versus frequency characteristic.”²

Group delay is expressed mathematically as

$$GD = -\frac{d\phi}{d\omega}$$

where GD is group delay in seconds, ϕ is phase in radians, and ω is frequency in radians per second.

Translation? If phase-versus-frequency response does not change in proportion to frequency, group delay exists. In a system, network, device or component with no group delay variation or group delay distortion, all frequencies are transmitted through the system, network, device or component in the same amount of time—that is, with equal time delay. If group delay distortion exists, signals at some frequencies travel faster than signals at other frequencies.

Group delay distortion exists if phase-versus-frequency deviates from ideal. But just what does that mean? As an example, using a 100 ft. piece of .500 feeder cable, hardline coax installed in cable networks for feeder applications has a velocity of propagation of around 87% (refer to the tutorial section on Velocity of Propagation for more information). The speed of light in free space or a vacuum is 299,792,458 meters per second, or 983,571,056.43 feet per second—1 foot in about 1.02 ns.

In coaxial cable with a velocity of propagation of 87%, electromagnetic signals travel at a velocity equal to 87% of the free space value of the speed of light. That works out to 260,819,438.46 meters per second, or 855,706,819.09 feet per second—1 foot in about 1.17 ns.

So, electromagnetic signals will travel 100 feet in a vacuum in 101.67 ns, and through a 100 ft. piece of coax in 116.86 ns.

Wavelength (λ) is the speed of propagation of an electromagnetic signal divided by its frequency (f) in hertz (Hz). It is further defined as the distance a wave travels through some medium in one period of an electromagnetic signal, or the distance over which a wave’s shape repeats. The period (T) of an electromagnetic signal (in seconds) = $1/f$, that is, the reciprocal of the electromagnetic signal’s frequency in Hz. In a vacuum, wavelength in feet (λ_{ft}) = $983,571,056.43/f_{Hz}$, which is the same as $\lambda_{ft} = 983.57/f_{MHz}$. In coaxial cable with 87% velocity of propagation, $\lambda_{ft} = 855,706,819.09/f_{Hz}$ or $\lambda_{ft} = 855.71/f_{MHz}$.

For instance, the period of a 10 MHz sine wave is $1/10,000,000 \text{ Hz} = 1 \times 10^{-7}$ second, or 0.1 μs . That means a 10 MHz signal takes 0.1 μs to complete one cycle, or 1 second to complete 10,000,000 cycles. In a vacuum an electromagnetic signal travels 98.36 ft. in 0.1 μs . This distance is one wavelength in a vacuum for a 10 MHz signal.

² The IEEE’s dictionary definition does not include the word “negative.”

In 87% velocity of propagation coax, the 10 MHz signal travels 85.57 ft. in 0.1 μ s. This distance is one wavelength in coax for a 10 MHz signal.

By calculating a given frequency's wavelength in feet, it can be said that a 100 ft. piece of .500 coax is equivalent to a certain number wavelengths at that frequency! From the previous example, it stands to reason that a 100 ft. piece of coax is equivalent to just over one wavelength at 10 MHz. That is, the 10 MHz signal's 85.57 ft. wavelength in coax is just shy of the 100 ft. overall length of the piece of coax.

Consider the wavelength in feet for several frequencies in a vacuum and in our 100 ft. piece of coax, using the previous formulas. Because of the cable's velocity of propagation, each frequency's wavelength in the cable will be a little less than it is in a vacuum. The number of wavelengths can also be figured out for each frequency in the 100 ft. piece of coax. Finally, knowing that one wavelength (cycle) of a sine wave equals 360 degrees of phase, the total number of degrees of phase the 100 ft. piece of cable represents at each frequency can be calculated. All of this is summarized in the following table.

Table I-1 - Frequency, wavelength, and phase relationships in 100 feet of coax

Frequency	λ_{ft} in a vacuum	λ_{ft} in coax	Number of λ in 100 ft of coax	Total phase in degrees in 100 ft of coax
1 MHz	983.57 feet	855.71 feet	0.12 λ	42.07°
5 MHz	196.71 feet	171.14 feet	0.58 λ	210.35°
10 MHz	98.36 feet	85.57 feet	1.17 λ	420.71°
30 MHz	32.97 feet	28.52 feet	3.51 λ	1262.12°
42 MHz	23.42 feet	20.37 feet	4.91 λ	1766.96°
50 MHz	19.67 feet	17.11 feet	5.84 λ	2103.53°
65 MHz	15.13 feet	13.16 feet	7.60 λ	2734.58°
100 MHz	9.84 feet	8.56 feet	11.69 λ	4207.05°

Next, plot the 100 ft. piece of cable's phase-versus-frequency on a graph. In the example in Figure I-2, the line is a sloped straight line—that is, phase varies proportionally with frequency.

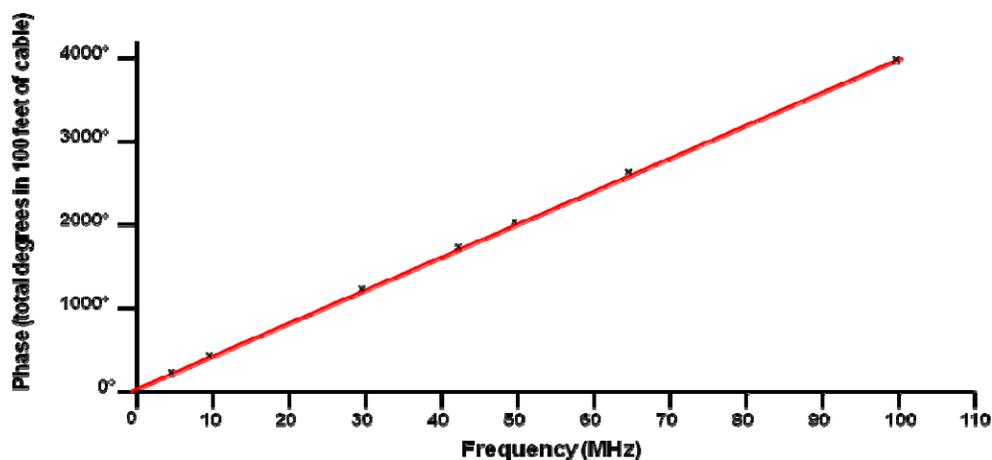


Figure I-2 - Phase-versus-frequency for 100 feet of coax

Finally, plot the time delay for each frequency through the 100 ft. piece of cable. The line in Figure I-3 is the negative derivative of radian phase with respect to radian frequency. It's a flat straight line; because the delay is the same at all frequencies, there is no group delay variation!

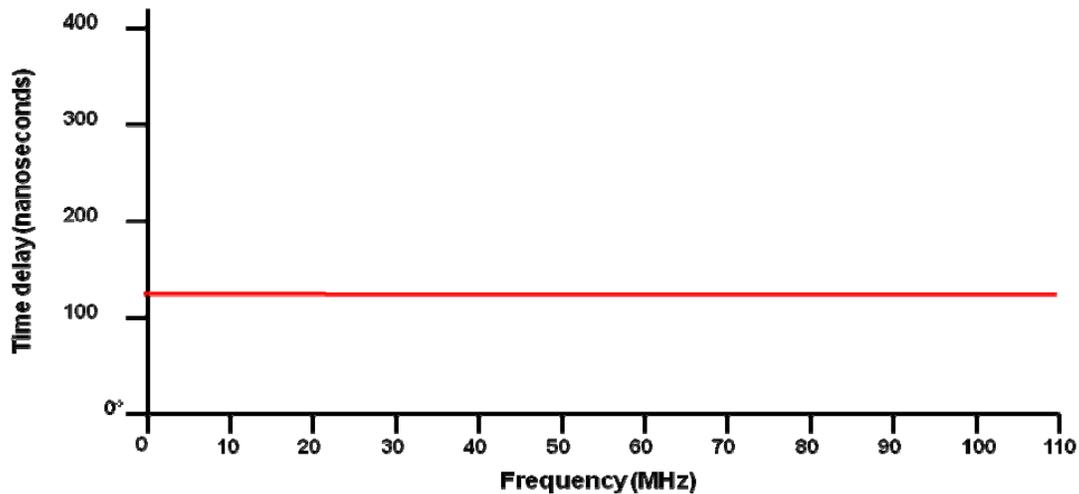


Figure I-3 - Time delay-versus-frequency for 100 feet of coax

Another way of looking at this is to say that the cable's velocity of propagation is the same at all frequencies! In other words, every frequency takes 116.86 ns to travel from one end of the 100 ft. piece of cable to the other end. But what happens if something in the signal path causes some frequencies to travel a little slower than other frequencies?

Look at the phase-versus-frequency trace (the sawtooth-shaped waveform) in Figure I-4. Where phase does not vary proportionally with frequency—that is, where the sloped line is not straight—group delay variation exists. The bottom bathtub-shaped trace in the figure is group delay, plotted as frequency in the horizontal axis versus time (100 ns/div) in the vertical axis. The group delay trace reveals that a signal at 40 MHz takes about 300 ns longer to travel through the network than a signal at 20 MHz.

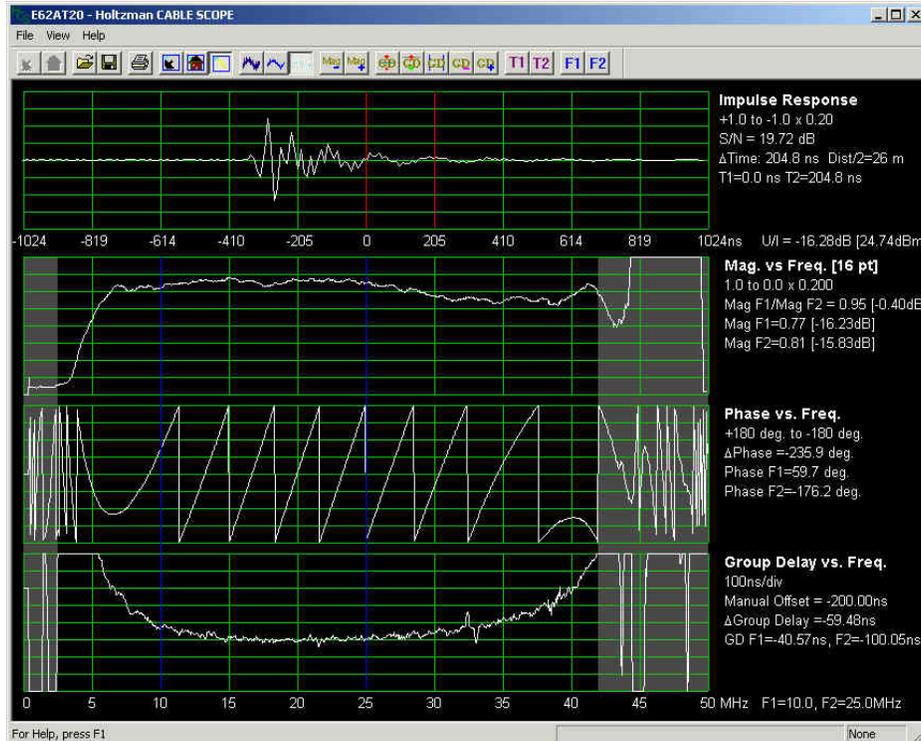


Figure I-4 - Complex frequency response in the return path

Common sources of group delay in a cable network include:

AC power coils/chokes (affects 5~10 MHz in the upstream)

Node and amplifier diplex filters (affect frequencies near the diplex filter cutoff region in the upstream and downstream, typically the upper end of the return spectrum, and the lower end of the downstream spectrum)

Band edges and rolloff areas

High-pass filters, data-only filters, step attenuators, feeder taps or in-line equalizers with filters

Group delay ripple caused by impedance mismatch-related micro-reflections and amplitude ripple

The Fix?

Use adaptive pre-equalization available in DOCSIS 1.1, 2.0, and 3.0 modems (generally not supported in DOCSIS 1.0 modems)

If possible, avoid frequencies where diplex filter-related group delay is common

Sweep the forward and reverse to ensure amplitude-versus-frequency response is flat (set the sweep equipment to the highest resolution available, and use resistive test points or probe seizure screws to see amplitude ripple)

Identify and repair impedance mismatches that cause micro-reflections

Use specialized test equipment to characterize and troubleshoot group delay (group delay can exist even when amplitude-versus-frequency response is flat)

I.5 Impedance Mismatch

All transmission lines and their components have a property known as characteristic impedance, expressed in ohms. Coaxial cable characteristic impedance is related to the diameters of the center conductor and inside surface of the shield, and the dielectric material's dielectric constant. Mathematically, that relationship is

$$Z_0 = \frac{138}{\sqrt{\epsilon_r}} \log\left(\frac{D}{d}\right)$$

where:

Z_0 = coaxial cable characteristic impedance

D = inside diameter of shield

d = outside diameter of center conductor

ϵ_r = dielectric constant

Consider 6-series coaxial cable, commonly used in cable subscriber drop applications. A typical nominal impedance specification for 6-series subscriber drop cable is 75 ohms \pm 3 ohms, with a nominal 85 percent velocity of propagation. Assume the center conductor is 18 AWG (0.040 inch diameter), and the inside diameter of the shield is 0.180 inch. The cable's 85 percent velocity of propagation translates to a dielectric constant of 1.3841 (see the tutorial section on Velocity of Propagation). Using the previous formula, the calculated impedance Z_0 of the 6-series cable in this example is 76.6 ohms, well within the specified 75 ohms \pm 3 ohms range.

Transmission line fundamentals typically are based upon a simplified model that includes—as shown in Figure I-5—a signal source S , a lossless transmission line T , and a terminating impedance comprising some sort of load or termination L . The assumption is that all three components have the same characteristic impedance Z , such that all power in the incident wave transmitted from the source is absorbed by the load.

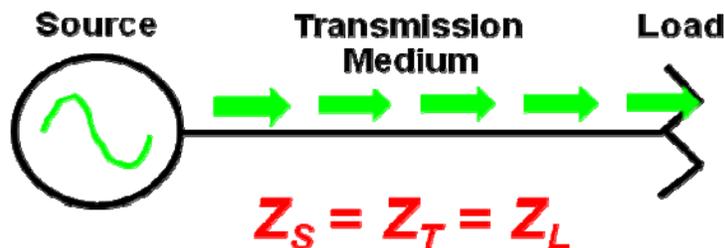


Figure I-5 - Ideal transmission line model

In the real world the signal source, transmission line, and load rarely have the same impedance at all frequencies. As well, the transmission line has attenuation which reduces the RF power of the signal(s) passing through it. When the impedance of, say, the load L is different from that of the transmission medium T (see Figure I-6), it is said that an impedance mismatch exists. An impedance mismatch causes some or all of the incident wave to be reflected back toward the source³. The ratio of reflected to incident voltages is known as the reflection coefficient Γ , described mathematically as $\Gamma = E^-/E^+$, where E^- is the reflected wave's voltage, and E^+ is the incident wave's voltage⁴.

³ For an overview of the mechanics of how an impedance mismatch—specifically an open or short circuit—creates a reflection, see “Impedance Mismatches and Reflections” in the December 2005 issue of *Communications Technology* magazine. <http://www.cable360.net/ct/operations/testing/15296.html>

⁴ Reflection coefficients are ratios of phasors and as such are complex quantities. Mathematically, $\Gamma = |\Gamma|\angle\theta$, where θ is the angle of Γ , and is the angle by which the reflected voltage leads the incident voltage. The magnitude of Γ can vary from 0 to 1, where 0 indicates all power is absorbed by the load (no reflection), and 1 indicates 100% reflection from an open circuit, short circuit, or pure reactance.

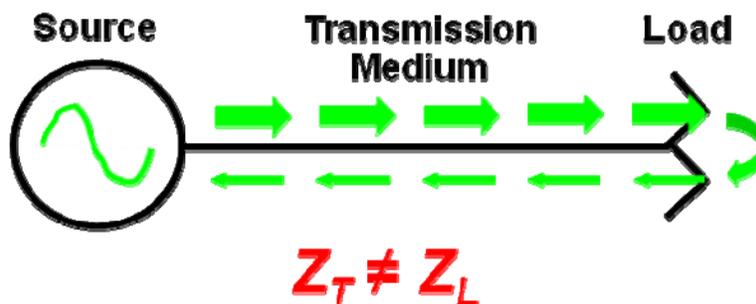


Figure I-6 - Real-world transmission line model

The reflected wave interacts with the incident wave to produce standing waves, also known as amplitude ripple. Impedance mismatches exist everywhere in our cable networks. Indeed, while the nominal impedance of a coaxial cable network and its components is said to be 75 ohms, every connector, amplifier, node, splitter, coupler, power inserter, feeder tap, terminator, and even the cable itself represent an impedance mismatch of some sort. The question is just how severe are those impedance mismatches? There are several ways to characterize the severity of an impedance mismatch. One is known as voltage standing wave ratio (VSWR), defined as the ratio of the standing wave's maximum voltage E_{\max} to its minimum voltage E_{\min} along the transmission line:

$$\text{VSWR} = E_{\max}/E_{\min}$$

VSWR also is related to reflection coefficient Γ :

$$\text{VSWR} = (1+|\Gamma|)/(1-|\Gamma|)$$

Another way to characterize the severity of an impedance mismatch is return loss (R), which is the ratio in decibels of incident power P_I to reflected power P_R :

$$R = 10\log_{10}(P_I/P_R)$$

Consider a scenario in which the incident power is 20 watts, and an impedance mismatch causes the reflected power to be 5 watts. The return loss is $10\log_{10}(20 \text{ watts}/5 \text{ watts}) = 6.02 \text{ dB}$. The relationship between VSWR and R is given by $\text{VSWR} = (10^{R/20} + 1)/(10^{R/20} - 1)$ and $R_{(\text{dB})} = 20\log_{10}[(\text{VSWR}+1)/(\text{VSWR}-1)]$. In this example, 6.02 dB return loss equals a VSWR of 3.0:1. Note that return loss is always a positive number, since it is the number of decibels between the amount of power in the incident and reflected signals. Some references (and test equipment!) show return loss as a negative number, but this is incorrect.

The worst case impedance mismatch is a short circuit, open circuit, or pure reactance, all of which reflect 100% of the incident wave back toward the source. In these three cases $R = 0 \text{ dB}$. When a perfect impedance match exists—that is, when $Z_T = Z_L$ —all of the incident wave is absorbed by the load and R is infinite (∞). While there are exceptions, most real-world impedance mismatches cause return loss to be somewhere between 0 dB and ∞ , since impedance mismatches are rarely pure short or open circuits, and a truly perfect impedance match with infinite return loss is more of a mathematical construct. From a practical perspective, typical return loss values of components in cable networks vary from a few dB to 30 dB or more.

When multiple impedance mismatches exist in a transmission medium—one example might be two water-damaged feeder taps in the outside plant separated by a span of feeder cable—multiple reflections occur! Figure I-7 shows an example in which a +31 dBmV incident signal travels right to left from the upstream output of the first water-damaged feeder tap (right side of figure) through a 100 ft. span of cable with 1 dB of loss. The first reflection occurs when the now +30 dBmV incident signal reaches an impedance mismatch—the upstream input of a second water-damaged feeder tap (left side of figure)—that has 7 dB return loss. A +23 dBmV reflection is created (+30 dBmV incident signal – 7 dB return loss = +23 dBmV reflection), which travels back toward the first feeder tap that also has 7 dB return loss. The now +22 dBmV reflection is re-reflected at a level of +15 dBmV at the first feeder tap (+22 dBmV reflection – 7 dB return loss = +15 dBmV second reflection), which travels back toward the second feeder tap and arrives there at +14 dBmV. And so on.

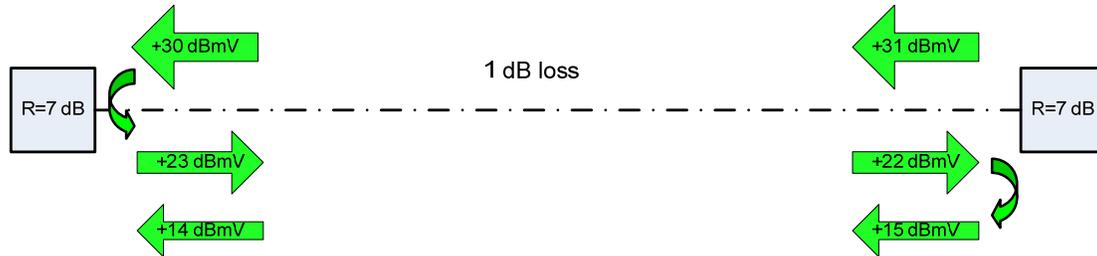


Figure I-7 - Creation of reflections in a cable network's feeder plant

The scenario in Figure I-7 can be represented graphically as shown in Figure I-8.

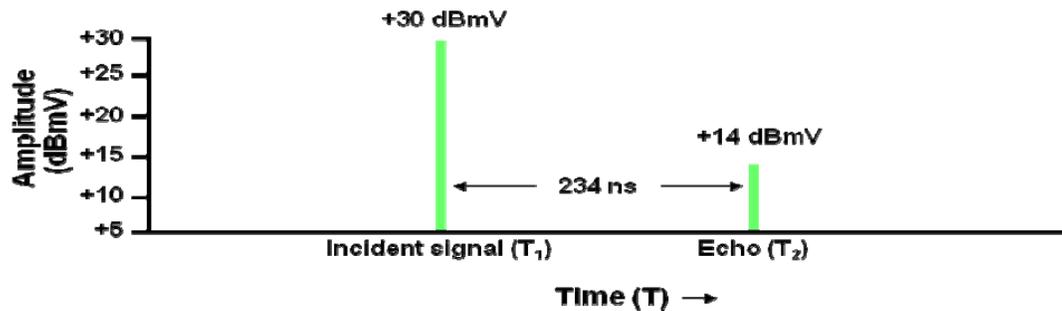


Figure I-8 - Graphic representation of incident signal and second reflection

The tall vertical line marked +30 dBmV represents the incident signal at the upstream input to the feeder tap on the left side of Figure I-7. The second, somewhat shorter vertical line marked +14 dBmV represents the second reflection, also at the input to the first feeder tap on the left side of Figure I-7. The horizontal separation between the two vertical lines represents the time delay between the +30 dBmV incident signal and the +14 dBmV reflection. Assuming the 100 ft. span of cable has a velocity of propagation of 87%, RF travels through each foot of coax in about 1.17 nanosecond. Since the roundtrip distance for the reflection is 100 ft. + 100 ft. = 200 ft., the second reflection's time delay is 200 ft. x 1.17 nanosecond = 234 nanoseconds.

Reflection or Micro-reflection?

When an impedance mismatch causes some or all of the incident wave to be reflected back toward the source, the reflected wave is called simply a reflection or an echo. In the world of high-speed data communications over cable networks, the term micro-reflection is frequently used. A micro-reflection is still a reflection or echo. "Micro-reflection" generally denotes a reflection with a short time delay—that is, a close-in reflection whose time delay relative to the incident signal ranges from less than a symbol period to perhaps several symbol periods.

As noted earlier in this tutorial, reflections (and micro-reflections) are caused by impedance mismatches. Some of the more typical causes of impedance mismatches and the resulting micro-reflections in cable networks include:

- Damaged or missing end-of-line terminators
- Damaged or missing chassis terminators on directional coupler, splitter, or multiple-output amplifier unused ports
- Loose center conductor seizure screws
- Unused feeder tap ports not terminated; this is especially critical on low value feeder taps, but *all* unused feeder tap ports should be terminated with 75-ohm terminations (locking terminators *without* resistors or stingers do not terminate the feeder tap port)

- Poor isolation in splitters, feeder taps, and directional couplers
- Unused customer premises splitter and directional coupler ports not terminated
- Use of so-called self-terminating feeder taps at ends-of-line; these are the equivalent of splitters, and do not properly terminate the feeder cable unless *all* feeder tap ports are terminated
- Kinked or damaged cable (includes cracked cable, which causes a reflection *and* ingress)
- Defective or damaged actives or passives (water-damaged, water-filled, cold solder joint, corrosion, loose circuit board screws, etc.)
- Cable-ready TVs and VCRs connected directly to the drop (return loss on most cable-ready devices is poor)
- Some traps and filters have been found to have poor return loss in the upstream, especially those used for data-only service

I.6 Amplitude Ripple

As noted in the Impedance Mismatch tutorial section, a reflected wave interacts with an incident wave to produce standing waves (amplitude ripple). The double reflection example from the same tutorial section is shown in Figure I-9, and will be used to illustrate how amplitude ripple forms.

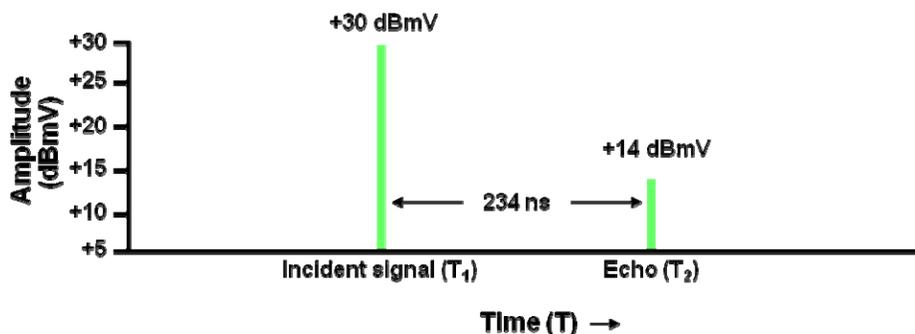


Figure I-9 - Reflection example that will be used to illustrate the formation of amplitude ripple

First, convert the incident and reflected signals from dBmV to millivolts (mV) using the formula $mV = 10^{(dBmV/20)}$. This results in $+30 \text{ dBmV} = 31.62 \text{ mV}$ and $+14 \text{ dBmV} = 5.01 \text{ mV}$. Next look at phasor views of the incident and reflected signals, with each represented as a vector⁵. The length of each vector corresponds to the signal magnitude in millivolts, and the direction vectors point is phase. In this example make the incident signal's vector (long arrow) stationary while the reflection's vector (short arrow) rotates counterclockwise around the end of the incident signal's vector. The vector sum of the two vectors will be plotted on a graph from which the amplitude ripple is derived. In Figure I-10 the incident and reflected signals are in phase, represented by the two arrows placed end-to-end in series. The vector sum in this case is simply $31.62 \text{ mV} + 5.01 \text{ mV} = 36.63 \text{ mV}$.

⁵ A vector is a quantity that has magnitude and direction, and is represented graphically with an arrow. One example of the use of vectors is the weather forecast on local TV news. Wind vectors (arrows) are overlaid on a map, and those vectors represent wind—the length of the arrows correspond to the wind's velocity, and the direction the arrows point show the direction the wind is blowing. In the case of an RF signal, a vector's length represents the magnitude or amplitude of the RF signal, and the direction the vector points represents the phase of the RF signal compared to an agreed reference signal.

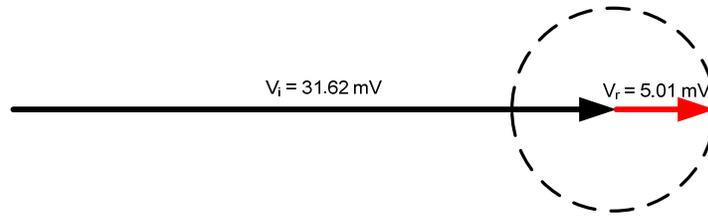


Figure I-10 - Phasor view of incident signal vector (long arrow) and reflection vector (short arrow)

Figure I-11 shows the reflection’s vector rotated counterclockwise 45 degrees from its original position. In reality the reflection vector rotates continuously around the end of the incident signal vector; it does not move in increments of 45 degrees. The latter is for illustrative purposes. The vector sum of the incident and reflection vectors is 35.34 mV (fine dashed arrow), which is calculated using basic geometry.

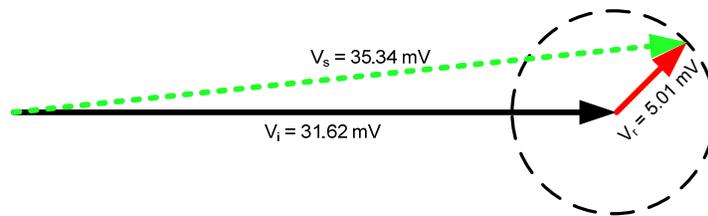


Figure I-11 - Reflection vector rotated 45 degrees from original position

The next figure shows the reflection vector has rotated counterclockwise 90 degrees from its original position. The vector sum of the incident and reflection vectors is 32.01 mV (fine dashed arrow).

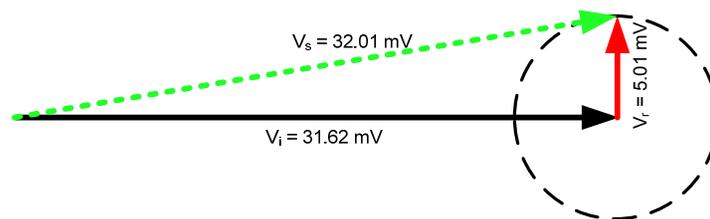


Figure I-12 - Reflection vector rotated 90 degrees from original position

Figure I-13 shows the reflection vector rotated counterclockwise 135 degrees from its original position. The vector sum of the incident and reflection vectors is 28.3 mV (fine dashed arrow).

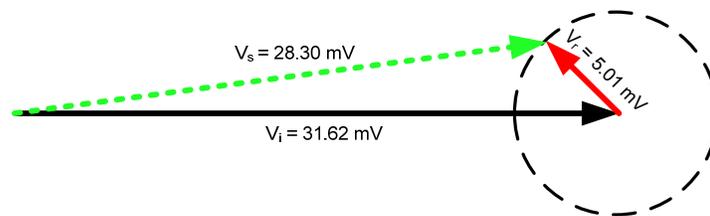


Figure I-13 - Reflection vector rotated 135 degrees from original position

Figure I-14 shows the reflection vector rotated counter clockwise 180 degrees from its original position. Since the incident and reflection vectors are out of phase, the vector sum is $31.62 \text{ mV} - 5.01 \text{ mV} = 26.61 \text{ mV}$.

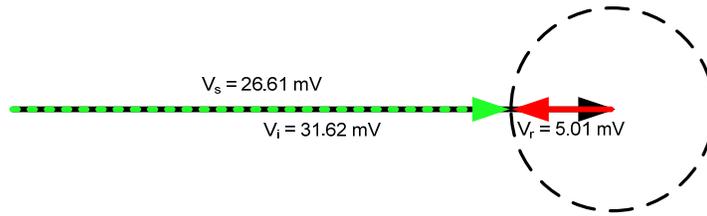


Figure I-14 - Reflection vector rotated 180 degrees from original position

The following figures show the reflection vector’s phase in 45 degree increments, until it reaches 360 degrees (0 degrees), from which the counterclockwise rotation continues.

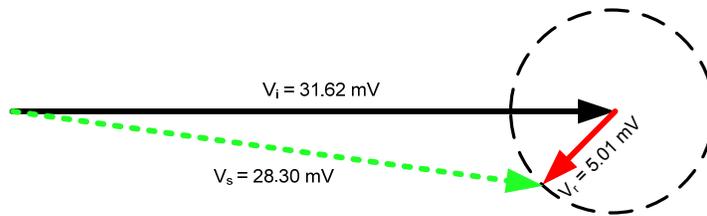


Figure I-15 - Reflection vector rotated 225 degrees from original position

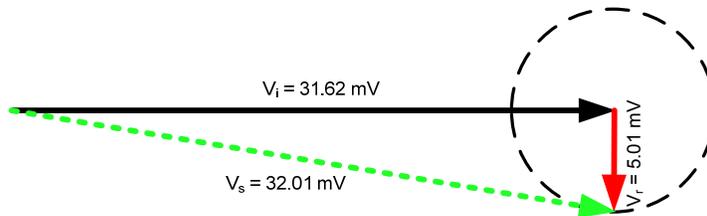


Figure I-16 - Reflection vector rotated 270 degrees from original position

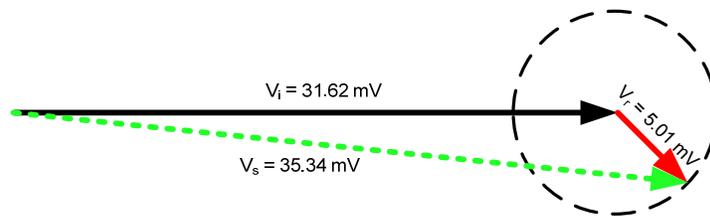


Figure I-17 - Reflection vector rotated 315 degrees from original position

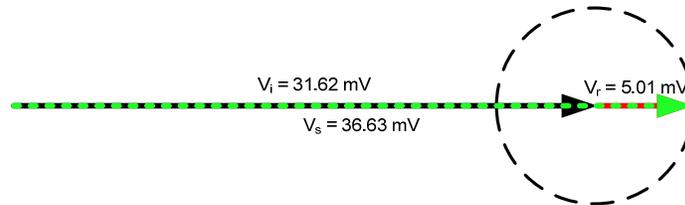


Figure I-18 - Reflection vector back at original position after rotating 360 degrees

The next step is to plot the vector sum vectors—the fine-dashed arrows—on a graph of amplitude-versus-phase (time), as shown in Figure I-19. Note that “connecting the dots” along the tops of the plotted vector sum vectors traces out the amplitude ripple caused by the reflection!

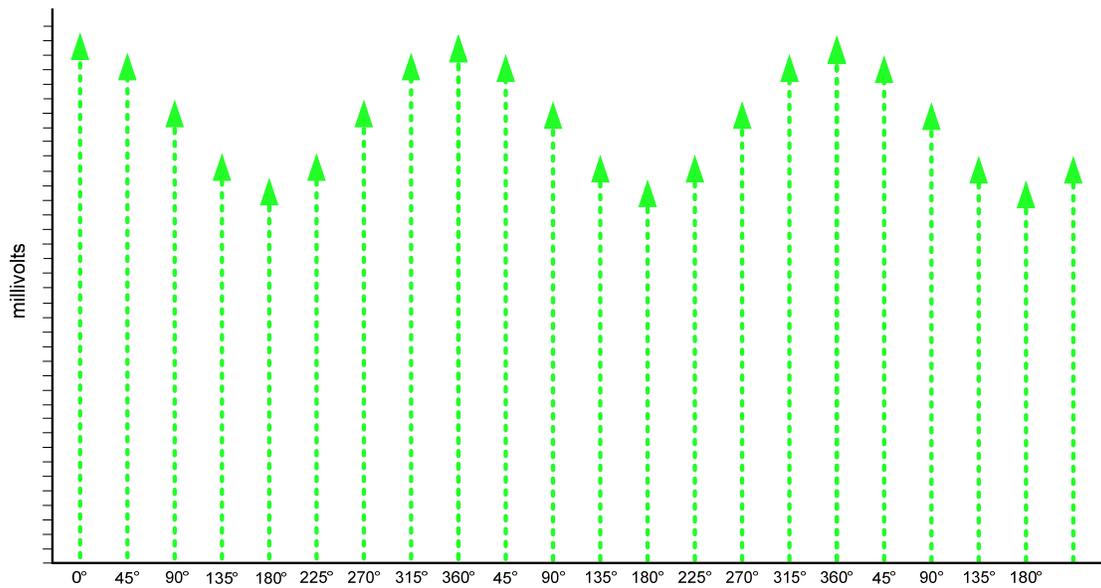


Figure I-19 - Amplitude-versus-phase plot of phasor view vector sum vectors

Now that the magnitude of the amplitude ripple has been plotted, other details of interest can be derived from the response. The peak-to-valley amplitude ripple in decibels is $20\log_{10}(E_p/E_N)$ —where E_p is the ripple’s peak voltage and E_N is the ripple’s null voltage—or, for the example in Figure I-19, $20\log_{10}(36.63 \text{ mV}/26.61 \text{ mV}) = 2.78 \text{ dB}$. By already knowing the incident signal amplitude is +30 dBmV and the reflection’s amplitude is +14 dBmV (16 dB difference), the relative amplitude difference between the incident and reflection signals from the amplitude ripple plot can be validated. That is done using the formula $20\log_{10}[(E_p - E_N)/(E_p + E_N)]$, or $20\log_{10}[(36.63 - 26.61)/(36.63 + 26.61)] = -16 \text{ dBc}$.

The frequency of the amplitude ripple example being discussed here is the reciprocal of the reflection’s 234 ns time delay: $1/0.000000234 \text{ second} = 4,273,504 \text{ Hz}$ or 4.27 MHz. This is equal to the frequency spacing between adjacent ripple peaks or nulls. The distance between the two impedance mismatches in feet is $492 \cdot (V/F_{\text{MHz}})$, where V is the coaxial cable’s velocity factor (velocity of propagation expressed in decimal form rather than percentage), and F_{MHz} is the frequency spacing in MHz between adjacent peaks or nulls. Assuming 0.87 velocity factor (87% velocity of propagation) and approximately 4.27 MHz frequency spacing between adjacent peaks or nulls, the distance is $492 \cdot (0.87/4.27) \approx 100 \text{ feet}$.

I.7 Amplitude Tilt

One manifestation of nonflat amplitude-versus-frequency response is amplitude tilt. That is, the amplitude-versus-frequency characteristic of the channel or operating spectrum is sloped or tilted across a specified frequency range. Figure I-20 illustrates flat amplitude-versus-frequency response, in which each frequency has the same amplitude—indicated by the horizontal dashed line.

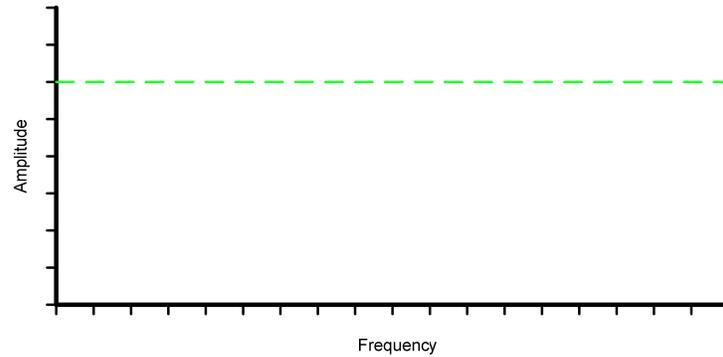


Figure I-20 - Example of flat amplitude-versus-frequency response

Figure I-21 shows amplitude tilt, indicated by the sloping dashed line.

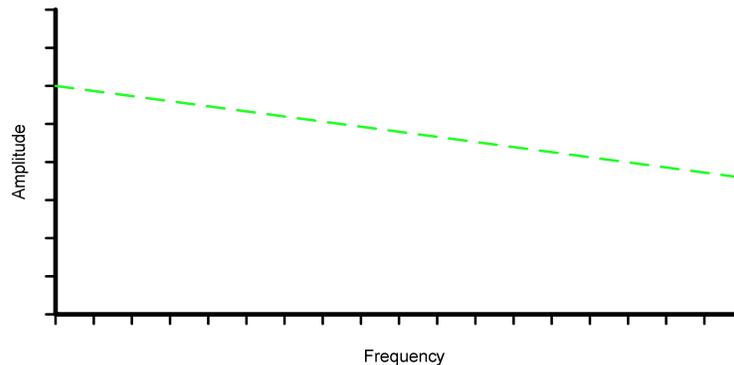


Figure I-21 - Example of tilted or sloped amplitude-versus-frequency response

Figure I-22 and Figure I-23 show two upstream 6.4 MHz 64-QAM signals, one with significant in-channel tilt and one that is mostly flat.

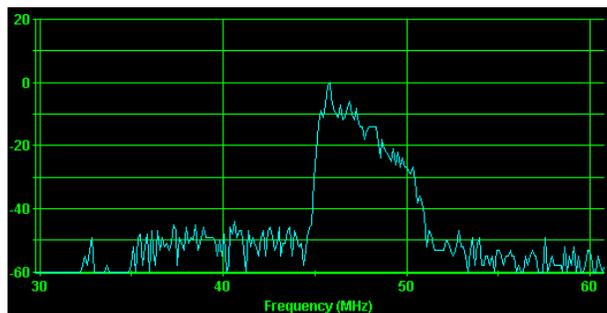


Figure I-22 - Example of upstream 64-QAM signal with substantial in-channel tilt

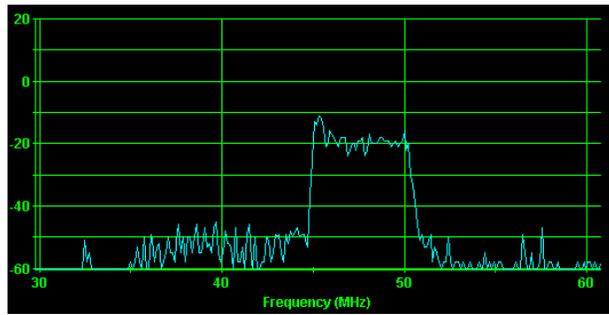


Figure I-23 - Example of 64-QAM signal after adaptive pre-equalization eliminated most of the in-channel tilt

Among the causes of amplitude tilt are defective active or passive devices, as well as active device misalignment. Note that many cable network distribution actives are often set up with an intentionally sloped output in part to compensate for the attenuation-versus-frequency characteristic of coaxial cable following the active device.

Other causes of amplitude tilt include operation near band edges or rolloff areas of the spectrum (or a filter), where amplitude-versus-frequency may not be flat. As well, short time delay reflections will result in widely spaced amplitude ripple, which may tilt the amplitude-versus-frequency response of signals in the sloped portion of the ripple. Figure I-24 shows an example of the latter scenario, where amplitude tilt is present between the two vertical dashed lines.

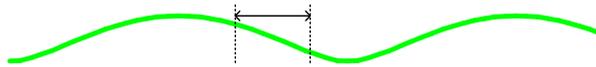


Figure I-24 - A signal carried in the sloped portion of the widely spaced amplitude ripple will exhibit in-channel tilt

I.8 Modulation Error Ratio

Modulation error ratio (MER) is a parameter reported by most QAM (quadrature amplitude modulation) analyzers when measuring the performance of digitally modulated signals. QAM analyzer screen shots showing a 64-QAM signal with good (36.3 dB) and poor (23.2 dB) MER can be seen in Figure I-25 and Figure I-26. Digital set-tops and cable modems report signal-to-noise ratio (SNR)—actually MER—for the downstream QAM signal to which the device is tuned, and cable modem termination systems (CMTSs) report upstream SNR (MER). But what is MER?

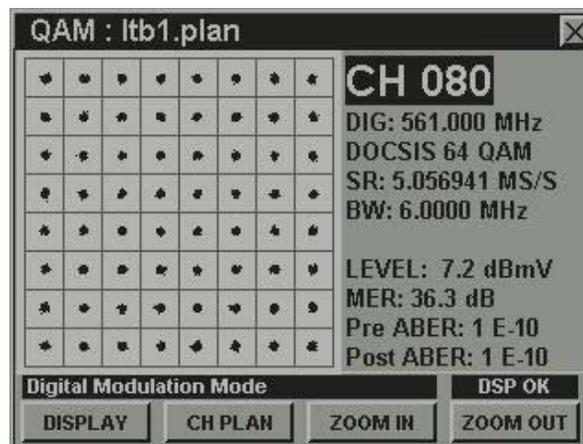


Figure I-25 - 64-QAM signal with good (36.3 dB) MER

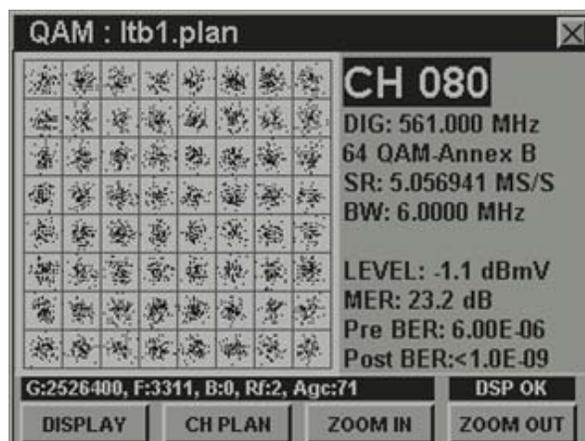


Figure I-26 - 64-QAM signal with poor (23.2 dB) MER

One analogy that might help to clarify the concept of MER is target shooting. A typical target used at the range comprises a set of concentric circles printed on a piece of paper. The center of the target is called the bull's-eye, which carries the highest point value. The further away from the bull's-eye, the lower the assigned points. Ideally, one would always hit the bull's-eye and get the maximum possible score. In the real world, this seldom happens. Instead, one or two shots might hit at or near the bull's-eye, and most of the rest hit somewhere in the circles surrounding the center of the target. For a person who is a decent shot, a round of target shooting usually results in a fairly uniform "fuzzy cloud" of holes in and around the bull's-eye. The smaller the diameter of this cloud and the closer it is to the bull's-eye, the higher the score.

A variety of factors affects how close to the bull's-eye the shots land. Some of those factors include the quality and accuracy of the firearm, type of ammunition used, weather conditions if outdoors, ambient lighting, distance to the target, steadiness of aim, breathing control, and so on.

Now visualize the constellation display on a QAM analyzer (refer to Figure I-25 and Figure I-26). Each symbol landing on the constellation can be thought of as a target of sorts. For instance, a 64-QAM constellation has 64 targets arranged in an eight-by-eight square-shaped grid. Ideally, when the 64 symbols are transmitted, they should land exactly on their respective targets' "bull's-eyes" in the constellation display. In reality, the symbols form a fuzzy cloud at and around the constellation's target centers. Measuring MER is, in effect, measuring the fuzziness of those clouds. The smaller the fuzzy clouds, the higher the MER. Like a high score in target shooting, the higher the MER, the better.

Modulation error ratio is digital complex baseband SNR—in fact, in the data world, the terms "SNR" and "MER" are often used interchangeably, adding to the confusion that sometimes exists about SNR, especially considering that in the telecommunications world, the terms "carrier-to-noise ratio (CNR)" and "SNR" are often used interchangeably.

Why use MER to characterize a data signal? Modulation error ratio is a direct measure of modulation quality and has linkage to bit error rate. Modulation error ratio is normally expressed in decibels, so it is a measurement that is familiar to cable engineers and technicians. It's a useful metric with which to gauge the end-to-end health of a network, although by itself, MER provides little or no insight about the type of impairments that exist.

Figure I-27 illustrates a 16-QAM constellation. A perfect, unimpaired 16-QAM digitally modulated signal would have all of its symbols land at exactly the same 16 points on the constellation over time. Real-world impairments cause most of the symbol landing points to be spread out somewhat from the ideal symbol landing points. If a constellation diagram is used to plot the landing points of a given symbol over time, as previously mentioned the resulting display forms a small "cloud" of symbol landing points rather than a single point.

Figure I-27 shows the vector for a target symbol – the ideal symbol to be transmitted. Because of one or more impairments, the transmitted symbol vector (or received symbol vector) is a little different than ideal. Modulation error is the vector difference between the ideal target symbol vector and the transmitted symbol vector. That is,

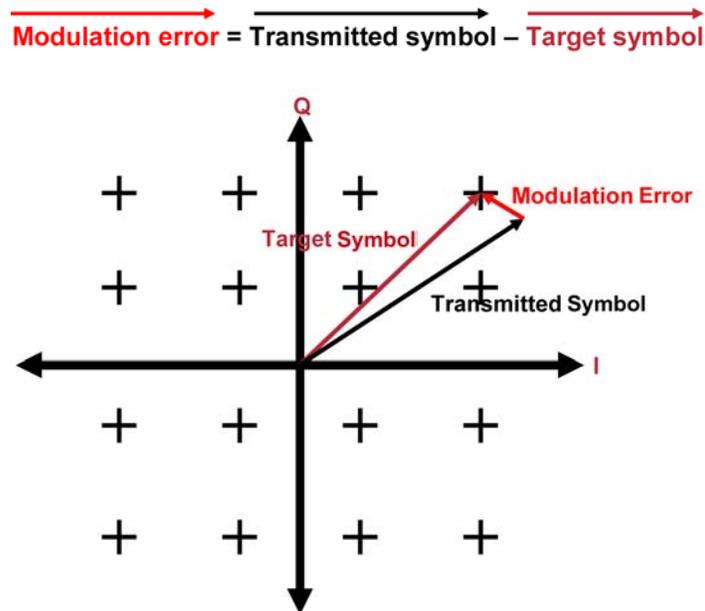


Figure I-27 - 16-QAM constellation showing target symbol, transmitted (or received) symbol, and modulation error vectors

Modulation error ratio is the ratio, in decibels, of average symbol power to average error power: $MER_{(dB)} = 10\log(\text{average symbol power}/\text{average error power})$. From this, you can see that the fuzzier the symbol cloud—in other words, the greater the average error power—the lower the MER. See Figure I-28.

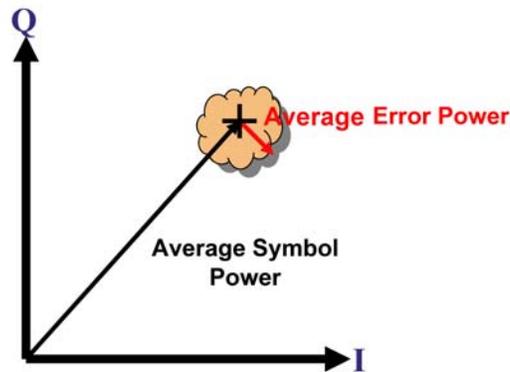


Figure I-28 - MER is the ratio of average symbol power to average error power

Mathematically, a more precise definition of MER is:

$$MER = 10 \log_{10} \left[\frac{\sum_{j=1}^N (I_j^2 + Q_j^2)}{\sum_{j=1}^N (\delta I_j^2 + \delta Q_j^2)} \right]$$

where I and Q are the real (in-phase) and imaginary (quadrature) parts of each sampled ideal target symbol vector, and δI and δQ are the real (in-phase) and imaginary (quadrature) parts of each modulation error vector. This definition assumes that a long enough sample is taken so that all the constellation symbols are equally likely to occur. Refer to the RxMER Measurement in Appendix I.9 tutorial section for more information on how MER is computed.

MER is affected by pretty much everything in a digitally modulated signal's transmission path: transmitted phase noise; carrier-to-noise ratio; nonlinear distortions (composite triple beat, composite second order, cross modulation, common path distortion); linear distortions (micro-reflections, amplitude tilt/ripple, group delay); in-channel ingress; laser clipping; data collisions; and even suboptimal modulation profiles. Some of these can be controlled fairly well, but no matter what is done, a digitally modulated signal is going to be impaired as it makes its way through a cable network. The worse the impairments, the fuzzier the constellation landings. The fuzzier the constellation landings, the lower the MER.

As such, the constellation's symbol landings will never be perfectly small points. They will always be spread out at least a little, the extent of which is described by MER. By itself, a low measured MER value doesn't determine what caused it to be low in the first place, only that it is low. Poor carrier-to-noise ratio? In-channel ingress? Group delay? Micro-reflection? Hard to say, until additional diagnostics are done with test equipment such as a QAM analyzer.

I.9 RxMER Measurement in a Digital Receiver ⁶

This tutorial discusses how a digital receiver is implemented, and how RxMER is measured. Figure I-29 is a generalized block diagram of a digital QAM receiver. The receiver may reside in the CMTS, in which case it receives time-division multiple access (TDMA) or synchronous code division multiple access (S-CDMA) upstream bursts; or it may reside in a cable modem or set-top box (STB), in which case it receives a continuous stream of downstream digital data. The RF signal from the cable plant enters at the left of the diagram, and is processed by analog and digital front-end components that perform tuning, automatic gain control, channel selection, analog-to-digital conversion, and related functions. The square-root Nyquist filter has a response "matched" to the symbol or S-CDMA chip (a "chip" is a bit in the pseudorandom spreading code used in S-CDMA). An adaptive equalizer compensates for channel response effects, including group delay variation, amplitude tilt or ripple, and micro-reflections. An ingress canceller is normally included in a CMTS burst receiver to remove in-channel narrowband interference. Acquisition and tracking loops provide estimates of frequency, phase, and symbol timing, allowing the receiver to lock to the incoming signal. In the CMTS burst receiver, preamble symbols are used as a reference to aid in the acquisition and tracking of each upstream burst. In the case of S-CDMA, the chips are despread. The received QAM symbol, or "soft decision," is passed to the slicer, which selects the nearest ideal symbol or "hard decision," from the QAM constellation. The decisions are passed to the Trellis decoder, descrambler, deinterleaver, Reed-Solomon (RS) FEC decoder and MPEG deframer, and on to the MAC layer, which assembles and outputs received packets to the user.

⁶ Excerpted from the paper "Digital Transmission: Carrier-to-Noise, Signal-to-Noise, and Modulation Error Ratio," by Bruce Currihan (Broadcom) and Ron Hranac (Cisco).

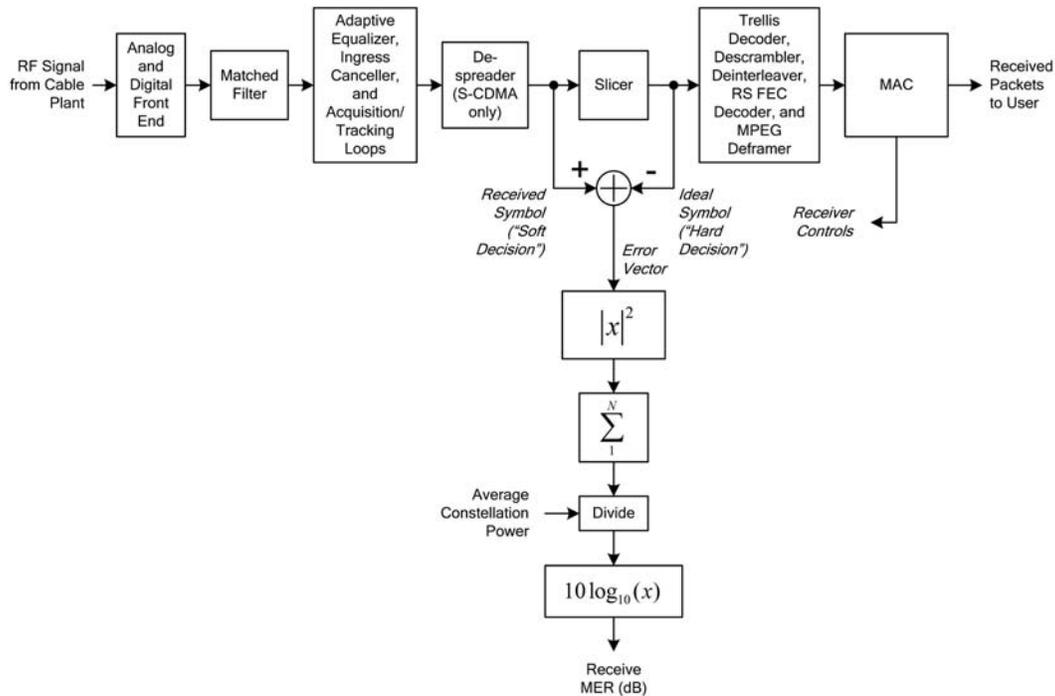


Figure I-29 - QAM receiver block diagram

What Is Inside the Blocks in a Digital QAM Receiver?

Analog and digital front end: Analog and digital front-end components perform tuning, automatic gain control, channel selection, analog-to-digital conversion, and related functions. Their purpose is to preprocess the signal so that the individual QAM RF channels are available for further digital processing.

Matched filter: The square-root Nyquist filter has a response matched to the symbol or S-CDMA chip. An identical filter is located in the transmitter; this “matched-filter” arrangement gives optimal receive SNR in white noise. The cascade of the transmit and receive square-root filters gives a response with the “Nyquist property.” This property, expressed in the time domain, results in ideally zero ISI, even when symbols are transmitted so close together in time that their responses significantly overlap.

Adaptive equalizer: This element compensates for channel effects, including group delay variation, amplitude slope or tilt, and micro-reflections. It adapts its filter coefficients to dynamically varying channel responses so as to maximize the receive MER. In effect, an adaptive equalizer creates a digital filter with the opposite response of the impaired channel.

Ingress canceller: An ingress canceller is normally included in a CMTS burst receiver to remove narrowband interference (including CB, ham and shortwave radio signals, etc.). It operates by dynamically detecting and measuring the interference, and adapting its coefficients to cancel it.

Acquisition and tracking loops: Tracking loops provide estimates of frequency, phase, and symbol timing, allowing the receiver to lock to the incoming signal. Acquisition refers to the initialization and pull-in process that occurs when the receiver is first powered on or changes channels.

Despreader: (S-CDMA upstream only) Despreading consists of multiplying the composite received signal by a given code sequence, and summing over all 128 chips in the code. There are 128 despreaders, one for each code. The output of the despreader is a soft symbol decision.

Slicer: The slicer selects the nearest ideal symbol, known as a hard decision, from the QAM constellation.

Trellis decoder: (Downstream and some S-CDMA upstream modes) The trellis decoder uses the Viterbi algorithm to choose the most likely sequence of symbols and thereby reject noise.

Descrambler: The descrambler adds a pseudorandom bit sequence to the received data bits, reversing the scrambling operation performed at the transmitter. The purpose of scrambling is to randomize the transmitted data in order to provide an even distribution of QAM symbols across the constellation.

Deinterleaver: The deinterleaver pseudo-randomly reorders groups of received bits, reversing the interleaving operation performed at the transmitter. The purpose of deinterleaving is to break up long bursts of noise so that the errored bits can be corrected by the Reed-Solomon decoder.

Reed-Solomon (RS) FEC decoder: This device processes groups of bits (7- or 8-bit symbols) arranged in codeword blocks, in terms of an algebraic code using Galois field arithmetic. By processing the received code words, which include redundant parity symbols, receive symbol errors can be found and corrected, up to one corrected RS symbol for each two redundant RS parity symbols.

MPEG deframer: The downstream DOCSIS signal is grouped into 188-byte MPEG transport packets, permitting the multiplexing of video and data over the common physical layer. The MPEG deframer removes the MPEG transport overhead to recover the bytes that are delivered to the MAC layer.

MAC: The media access control (MAC) layer controls the physical (PHY) layer and is the source and sink of PHY data. The MAC layer processes data frames delineated by DOCSIS headers. In the upstream, the MAC layer governs how cable modems share the channel through a request or grant mechanism.

The input and output of the slicer are complex numbers or vectors, each represented by two components: magnitude and phase, or equivalently, real (in-phase or “I”) and imaginary (quadrature or “Q”) parts, as shown in Figure I-30. In an ideal zero-noise, zero-ISI condition, the soft decision would lie exactly on one of the constellation points, and the magnitude of the error between them would be zero. In a real-world receiver, subtracting the hard-decision vector from the soft-decision vector gives the error or noise vector at each symbol time. The implicit assumption is that a low symbol error rate exists – that is, very few decisions are incorrect, ensuring that the “decision-directed” error vector from the nearest symbol nearly always equals the true error vector from the correct reference symbol.

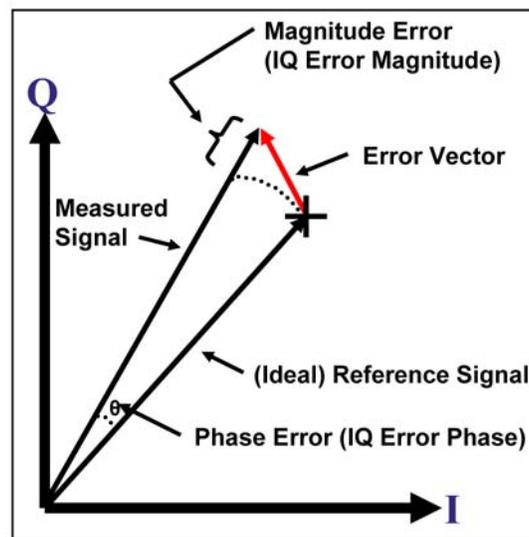


Figure I-30 - Each vector has a real (in-phase or I) and imaginary (quadrature or Q) component

For RxMER, the concern is with the average power of the error vector, which is computed by taking the complex magnitude-squared of the error vector and accumulating or averaging it over a given number of symbols N . This process gives the error vector power (or noise power) at the slicer. To obtain the ratio of signal to noise, the average signal power (a known constant for each constellation, such as 64-QAM or 256-QAM) is divided by the average error vector power. Then take the logarithm to convert to decibels, giving RxMER in dB. To summarize: RxMER is simply the ratio of average symbol power to average slicer error power, expressed in dB.

I.10 Adaptive Equalization

Digital communications systems are designed to transmit high-speed data through band-limited channels—for example, 6 MHz bandwidth downstream channels or 3.2 MHz bandwidth upstream channels. In the vast majority of cases a digitally modulated signal must be transmitted through a channel that is susceptible to various distortions. The presence of distortion in the channel results in something called inter-symbol interference (ISI), which may cause data transmission errors. One way to compensate for or reduce ISI is to incorporate an equalizer in the receiver or transmitter. If the channel characteristics are known and do not change over time, fixed-value equalizers are one possible solution. However, in a typical cable network the signal path between the headend and each cable modem (or digital set-top) is unique, so a one-size-fits-all fixed-value equalizer is impractical. Furthermore, distortions causing ISI often change over time, so the equalizer must somehow be adjustable to compensate for changes in channel conditions. Adaptive equalizers are most often used for this purpose.

All QAM receivers—whether in a digital set-top, cable modem, QAM analyzer, or even the upstream receiver of a cable modem termination system (CMTS)—incorporate a circuit known as an adaptive equalizer. Adaptive equalization is used in data transmission to compensate for certain in-channel impairments that degrade the quality of digital signals. Most discussions about adaptive equalization are full of complex mathematics and difficult-to-understand concepts. This tutorial provides a high-level overview of what adaptive equalization is and how it works, without the gnarly equations and Ph.D.-level techno-speak. Included is a step-by-step example of how a simple four-tap adaptive equalizer can be used to reduce in-channel impairments caused by a severe micro-reflection.

Consider first the concept of equalization from the perspective of a cable TV distribution network. It is understood that in a given length of coaxial cable higher frequencies are attenuated more than lower frequencies. For instance, if all downstream signals in the 54-870 MHz spectrum have the same amplitude at the output of an amplifier, the overall frequency response—technically speaking, amplitude (or magnitude)-versus-frequency response—is flat. For the following example, assume there is no slope at an amplifier's output (and the amplifier has no internal slope or tilt), and our equal-amplitude signals leave the output of the first amplifier and travel through 1,000 feet of 0.500 inch diameter coax to the second amplifier.

Since 0.500 inch diameter distribution cable's attenuation is about 0.5 dB/100 ft. at 54 MHz and 2.3 dB/100 ft. at 870 MHz, our hypothetical 1,000 ft. span of coax has a total of 5 dB of attenuation at 54 MHz and 23 dB of attenuation at 870 MHz. The 54-870 MHz spectrum will be tilted a bunch at the second amplifier's input! The goal is to see a flat amplitude-versus-frequency response across the spectrum, so it is necessary to install a fixed-value plug-in equalizer at the second amplifier. The equalizer is a small passive circuit that has the opposite amplitude-versus-frequency response of the 1,000 feet of coaxial cable preceding the amplifier. The equalizer is in effect a broadband filter that “cancels” the tilted response from 54 to 870 MHz, resulting in a flat amplitude-versus-frequency spectrum at the second amplifier's internal gain stages.

Adaptive equalization performs a function similar to that of a cable amplifier's fixed-value plug-in equalizer. Rather than equalizing the entire 54-870 MHz downstream or 5-42 MHz upstream RF spectrum, it deals with just a single channel. Adaptive means the equalizer can change its characteristics as channel conditions change.

An adaptive equalizer is a digital circuit that compensates for a digitally modulated signal's in-channel complex frequency response impairments. The cable industry has long used the term frequency response to describe amplitude (or magnitude)-versus-frequency—that is, what is seen on the display of test equipment used to sweep

outside plant. True frequency response is a complex entity that has two components: amplitude-versus-frequency, and phase-versus-frequency. An adaptive equalizer can compensate for in-channel amplitude- and phase-versus-frequency impairments.

The adaptive equalizer uses sophisticated algorithms to derive coefficients for an equalizer solution “on the fly”—in effect, creating a digital filter with essentially the opposite complex frequency response of the impaired channel. Because the adaptive equalizer’s complex frequency response is in theory a mirror image of the impaired channel’s complex frequency response, the adaptive equalizer cancels out most or all of the degraded in-channel frequency response that is affecting the digital signal—within the limits of the adaptive equalizer’s capabilities, of course.

It’s important to note that at high signal-to-noise ratio (E_s/N_0), the adaptive equalizer will synthesize the opposite response of the channel. At lower SNR doing so would cause noise enhancement, so a compromise solution is derived.

Ideal adaptive equalizer coefficients yield maximum modulation error ratio (MER) by minimizing total impairments including inter-symbol interference (ISI), within the limits of the equalizer’s capabilities (number of adaptive equalizer taps, etc.).

If the in-channel impairment suddenly changes or goes away, the adaptive equalizer will distort the signal, at least until new equalizer coefficients for the current channel conditions are derived and the equalizer’s operation updated. This adaptation process is very fast, typically completed in milliseconds.

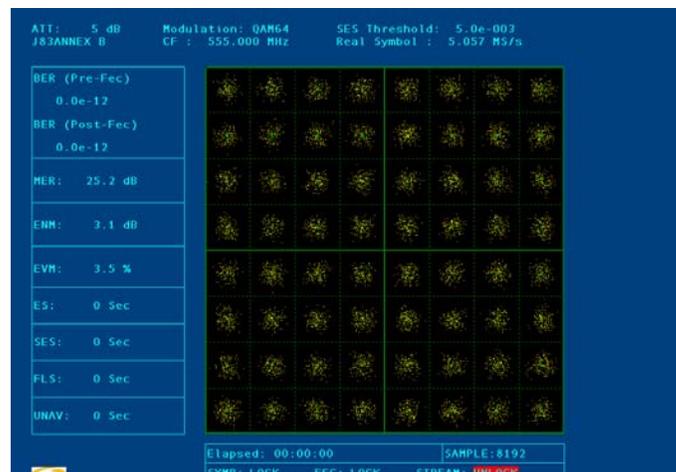


Figure I-31 - Unequalized 64-QAM constellation

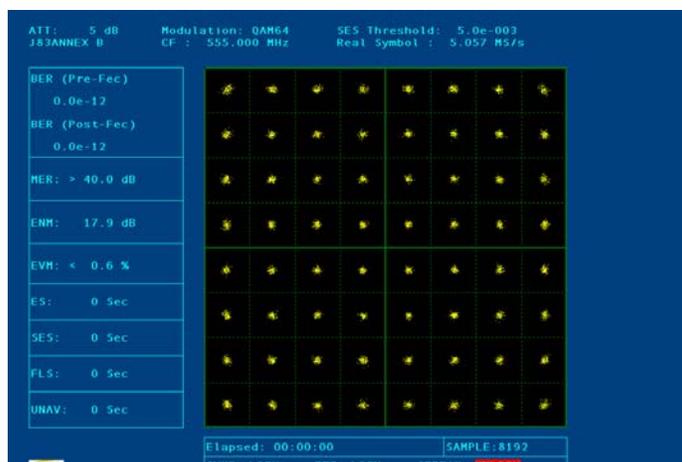


Figure I-32 - Equalized 64-QAM constellation

The test equipment screen shot in Figure I-31 shows an unequalized 64-QAM constellation. Note that the constellation's 64 symbol landings are “fuzzy,” caused by degraded modulation error ratio. Indeed, the unequalized MER is about 25 dB, only a few dB above the lower E_s/N_0 threshold for 64-QAM.⁷

The screen shot in Figure I-32 shows the same 64-QAM signal after adaptive equalization. The 64 symbol landings are small and almost dot-like. The equalized MER is >40 dB. The adaptive equalizer has effectively compensated for in-channel response impairments, but the equalized MER and constellation do not provide a way to determine how close to the failure threshold the signal really is.

By evaluating both unequalized and equalized MER, one can determine available MER headroom, and also see how well adaptive equalization is able to improve the signal.

In order for an adaptive equalizer's algorithms to begin the process of deriving coefficients that will be used to create a “filter” whose complex frequency response is opposite of the channel's complex frequency response, the equalizer starts with an adaptation source. The adaptation source can be a transmitted training sequence, or the signal itself.

Transmitted training sequence: In conventional zero-forcing or minimum mean square error (MSE) equalizers, a known training sequence is transmitted to the receiver for the purpose of initially adjusting adaptive equalizer coefficients. In the DOCSIS® upstream pre-equalization process (discussed in the Adaptive Pre-equalization tutorial section), data transmissions from all cable modems include a preamble at the beginning of each burst. The preamble is used as a training signal for the cable modem termination system's (CMTS's) adaptive equalizer.

The signal itself: Adaptive equalizers that do not rely upon transmitted training sequences for the initial adjustment of the coefficients are called self-recovering or blind adaptive equalizers. The adaptive equalizer in the downstream receiver of a DOCSIS cable modem (or a digital set-top) is a blind adaptive equalizer. In the case of cable modems, DOCSIS does not specify a training sequence in the downstream signal.

Several types of adaptive equalization methods are used in data transmission. For example, one method is based on maximum-likelihood (ML) sequence detection; another uses a linear filter with adjustable coefficients. Decision-feedback equalization—comprising a combination of feedforward and feedback filters—uses previously detected symbols to suppress inter-symbol interference in the current symbol being detected. A commonly used adaptive equalization type is a combination of decision-feedback and feedforward equalizers.

⁷ The lower E_s/N_0 threshold is in effect the unequalized MER failure threshold for the modulation format in use.

Powerful algorithms are used to automatically adjust adaptive equalizer coefficients to achieve optimum performance, and to rapidly adapt to changing channel characteristics. General criteria for defining optimum performance include minimizing peak distortion at the equalizer output, or minimizing the MSE at the equalizer output. In other words, the algorithm adjusts equalizer coefficients “on the fly” to converge on a solution that best reduces, say, MSE.

Zero-forcing and least-mean-square (LMS) are among the algorithms used in adaptive equalizers. Fractionally spaced equalizers (FSE) may use a LMS algorithm or a tap-leakage algorithm.

If faster convergence is desired in an equalizer, more complex algorithms generally have to be used. Examples include a recursive least-squares (Kalman) algorithm. Blind equalizers may use stochastic gradient algorithms such as Godard, Sato, Benveniste-Goursat, or stop-and-go.

Digital Communications, 4th Edition, by John G. Proakis (McGraw-Hill, ©2001, ISBN 0-07-232111-3), includes in-depth explanations of the adaptive equalizer and algorithm types mentioned here.

An important parameter in an adaptive equalizer is its span, defined mathematically as **(number of equalizer taps – 1) x equalizer tap spacing**. This particular definition assumes that the equalizer’s first tap is the main tap. Adaptive equalizer tap spacing is the amount of time delay per equalizer tap. An adaptive equalizer’s span is directly related to the maximum amount of time delay in a micro-reflection that can be compensated for.

The adaptive equalizer tap spacing of a DOCSIS 2.0 cable modem’s upstream pre-equalizer taps is called T-spaced, or symbol-spaced. Symbol-spaced means the time delay per adaptive equalizer tap is equal to the symbol period, which is the reciprocal of the symbol rate—that is, $1/T$. For example, if the upstream digitally modulated signal’s symbol rate is 5.12 megasymbols per second (Msym/sec), the adaptive equalizer tap spacing is 0.1953125 microsecond (μs) per tap: $1/5,120,000 = 1.953125 \times 10^{-7}$ second, or 0.1953125 μs .

Since DOCSIS 2.0 specifies 24-tap T-spaced pre-equalization in the upstream, the maximum span of an adaptive equalizer for a 5.12 Msym/sec signal is $(24 - 1) \times 0.1953125 \mu\text{s} = 4.49 \mu\text{s}$. Another way to calculate this value is $(24 - 1)/5.12 = 4.49 \mu\text{s}$.

Currently available CMTS burst receivers, which incorporate DOCSIS 2.0 and later 24-tap adaptive equalization, support main tap positions of 2 through 10. Adaptive equalizer tap position 8 generally is the default setting. How does that affect the maximum usable equalizer span when it comes to dealing with micro-reflections? Assuming that adaptive equalizer tap #8 is the main tap, that results in a usable span of $(24 - 8) \times 0.1953125 \mu\text{s} = 3.13 \mu\text{s}$. If the adaptive equalizer main tap is changed to #10, the usable span becomes $(24-10) \times 0.1953125 = 2.73 \mu\text{s}$, and if the adaptive equalizer main tap is changed to #2 the usable span becomes $(24-2) \times 0.1953125 = 4.3 \mu\text{s}$.

T-spaced equalizers are the most commonly used. As discussed previously, T-spaced means the equalizer taps are spaced at the reciprocal of the symbol rate.

A fractionally spaced equalizer is based on sampling the incoming signal at least as fast as the *Nyquist rate*. (From Wikipedia: “The Nyquist rate is the minimum sampling rate required to avoid aliasing when sampling a continuous signal. In other words, the Nyquist rate is the minimum sampling rate required to allow unambiguous reconstruction of a band limited continuous signal from its samples. If the input signal is real and band limited, the Nyquist rate is simply twice the highest frequency contained within the signal.”)

A $\frac{1}{2}T$ -spaced (also written as “ $T/2$ -spaced”) equalizer is used in many applications requiring a FSE. Other applications may use $\frac{1}{4}T$ -spaced ($T/4$ -spaced) equalizers, and so forth.

Why use a FSE? They often perform better than T-spaced equalizers in the presence of symbol clock timing errors, because FSEs are less sensitive to timing phase. Despite their better performance, FSEs are not as common as T-spaced equalizers because of computational complexity and convergence performance.

DOCSIS supports T-spaced and fractionally spaced equalizers under the following conditions for upstream pre-equalization, defined in the DOCSIS 2.0 Radio Frequency Interface Specification: “*There are two modes of operation for the pre-equalizer of a CM: DOCSIS 1.1 mode, and DOCSIS 2.0 mode. In DOCSIS 1.1 mode, the CM MUST support a (T)-spaced equalizer structure with 8 taps; the pre-equalizer MAY have 1, 2 or 4 samples per symbol, with a tap length longer than 8 symbols. In DOCSIS 1.1 mode, for backwards compatibility, the CMTS MAY support fractionally spaced equalizer format (T/2 and T/4). In DOCSIS 2.0 mode, the pre-equalizer MUST support a symbol (T)-spaced equalizer structure with 24 taps.*”

If the channel response contains a reflection (caused by an impedance mismatch) that is further out in delay than the span of the adaptive equalizer, the adaptive equalizer cannot compensate for that reflection.

For example, if there is a significant amount of ISI from a SAW filter’s triple transit, and this ISI is equivalent to a reflection at a large delay that is beyond the span of the equalizer’s taps, then the equalizer will still do its best on the other impairments. But it won’t be able to cancel the triple transit reflections since they are beyond the limits of the adaptive equalizer’s capabilities—in this case, the equalizer span.

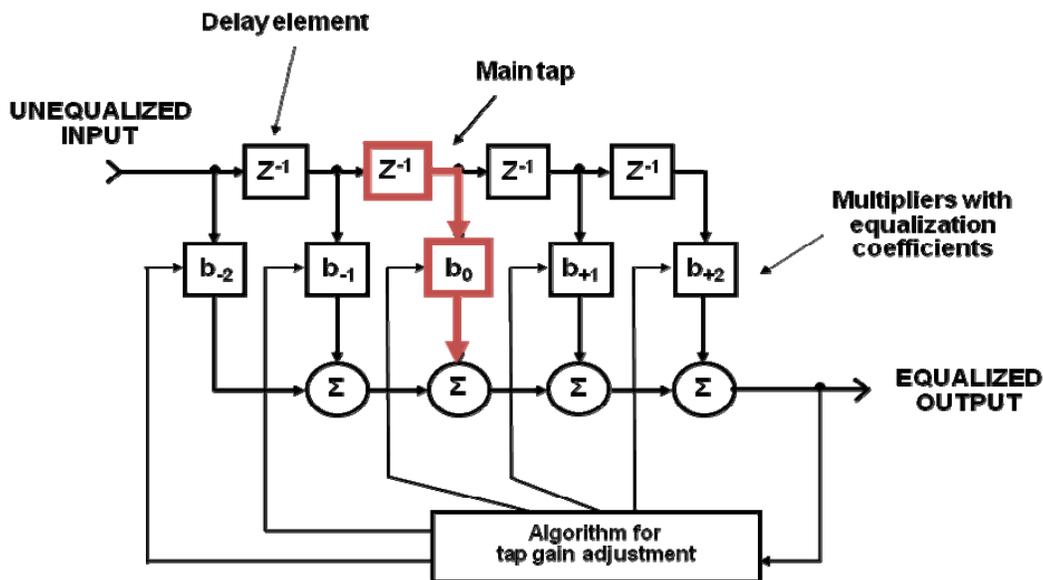


Figure I-33 - Generic 4-tap adaptive equalizer

Figure I-33 illustrates a generic adaptive equalizer. The top row with boxes labeled Z^{-1} can be thought of as a tapped delay line. Each box marked Z^{-1} is a *delay element*, with the amount of time delay per “box” equal to the reciprocal of the symbol rate in a T-spaced equalizer. A delay element is often called a tap, but an adaptive equalizer tap also can be considered the combination of a delay element, the point where some of the signal is “tapped” off, and a multiplier. The boxes labeled b_{-2} , b_{-1} , b_0 , etc., are multipliers with coefficients that set the gain for each adaptive equalizer tap. The algorithm adjusts the equalization coefficients that set the gain for each multiplier. The circles marked Σ are summing or combining circuits.

One adaptive equalizer tap is called the main tap (the second Z^{-1} delay element to multiplier b_0 and highlighted in red in the figure). The main tap has a gain of 1, and passes the input signal at its original amplitude. Other adaptive equalizer taps represent either the “past” or “future” relative to the main tap, and vary the amplitudes of the respective signals passing through them as required.

While an adaptive equalizer is a digital circuit, one can think of its operation as functionally similar to an analog circuit that combines different amplitudes and phases of an input waveform to achieve a desired output waveform.

To better understand how a simple adaptive equalizer works, consider a scenario in which an impedance mismatch results in a severe micro-reflection. Assume that the incident signal has an amplitude of 1, and the echo (micro-reflection) caused by the impedance mismatch has an amplitude of -0.5 and is offset in time by $1\ \mu\text{s}$. The minus sign indicates that the echo has the opposite phase of the incident signal. Figure I-34 illustrates the incident signal and echo graphically in terms of amplitude-versus-time.

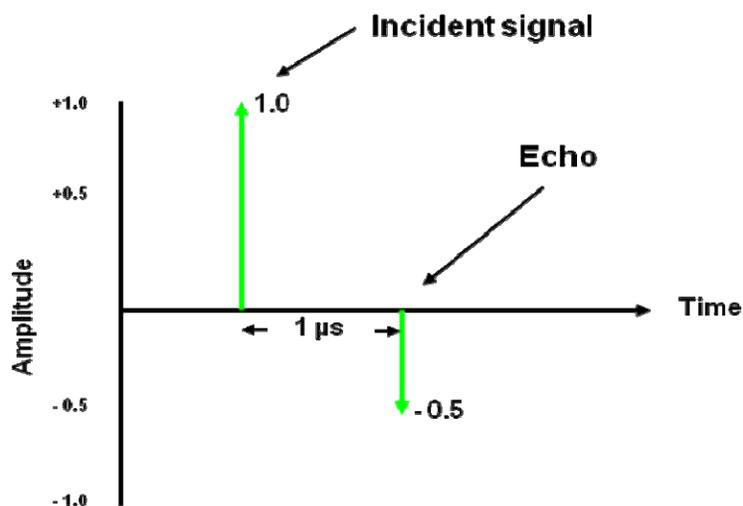


Figure I-34 - Amplitude-versus-time plot of an incident signal and micro-reflection

Figure I-35 shows amplitude (or magnitude)-versus-frequency response caused by the $1\ \mu\text{s}$ micro-reflection from Figure I-34. When the incident and micro-reflection vectors are in phase, they add to create a vector whose sum is 1.5; when the two vectors are out of phase, the resulting vector sum is 0.5. This makes the scalloped sine wave's peak-to-peak amplitude vary from 0.5 to 1.5, which, in decibels, is $20\log(1.5/0.5) = 9.54\ \text{dB}$. The frequency spacing between successive peaks is $1/1\ \mu\text{s}$, or 1 MHz.

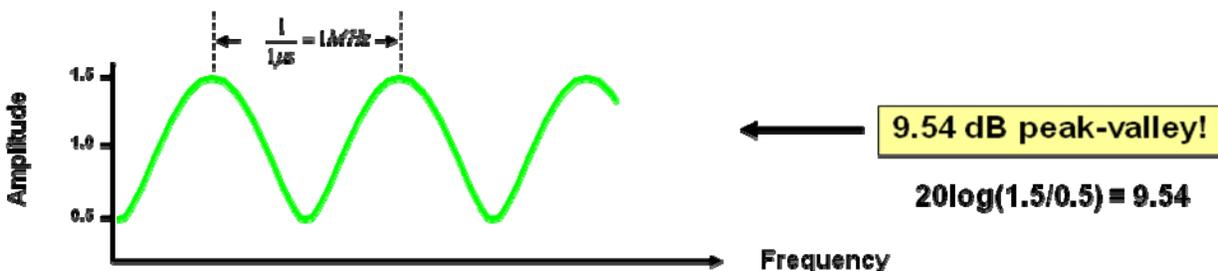


Figure I-35 - Amplitude-versus-frequency response

Figure I-36 shows the resulting phase-versus-frequency response caused by the micro-reflection in Figure I-34. The maximum phase excursion occurs when the vector sum vector is tangent to the circle on a phasor diagram representation of the incident and reflection vectors. That forms a right-triangle with *side c*—the hypotenuse—equal to 1.0 (incident signal vector length); *side b*—the opposite side—equal to 0.5 (reflection vector length), and *side a*—the

adjacent side—equal to 0.886 (vector sum length). Trigonometry can be used to calculate the peak phase excursion, which equals the right-triangle’s angle B: $\sin B = \text{opp/hyp}$, or $\text{angle } B = \text{arcsine}(\text{opp/hyp})$: $\text{arcsine}(0.5/1) = 30^\circ$.

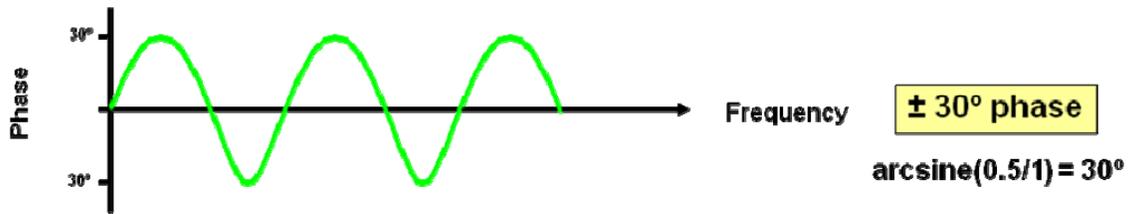


Figure I-36 - Phase-versus-frequency response

In order to cancel the echo and get a flat response, an adaptive equalizer with the opposite amplitude (magnitude)-versus-frequency and phase-versus-frequency response is needed. Figure I-37 shows the required amplitude and phase response to cancel the -0.5 amplitude 1 μs echo.

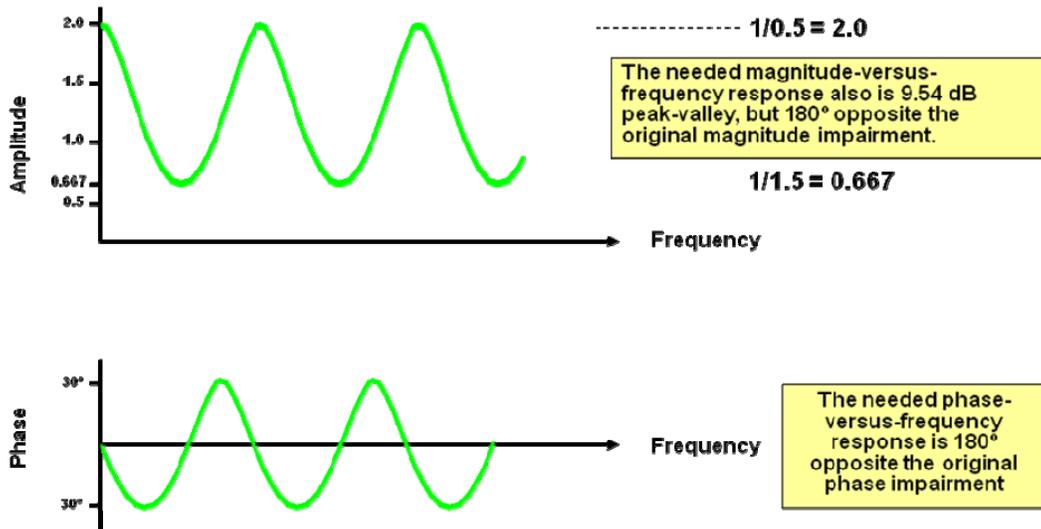


Figure I-37 - Required magnitude- and phase-versus-frequency response to cancel echo

Why is the peak-to-peak linear magnitude of the needed opposite amplitude-versus-frequency response 0.667 to 2.0 rather than the original 0.5 to 1.5?

A time-invariant system such as a filter can be characterized by its impulse response $h(t)$ or by its frequency response $H(f)$, which comprise what is known as a Fourier transform pair. Flat response occurs when $H(f)$ multiplied by $1/H(f)$ equals 1.00. If $H(f)$ is 0.5, then $1/H(f)$ is 2.0; likewise, if $H(f)$ is 1.5, then $1/H(f)$ is 0.667. In the example in the figure, $0.5 \times 2 = 1.00$ and $1.5 \times 0.667 = 1.00$.

For this example, use a simple 4-tap adaptive equalizer to compensate for the impaired frequency response caused by the -0.5 amplitude 1 μs echo. The first adaptive equalizer tap will be the main tap, so b_0 will have a gain of 1. Note that there is no delay element feeding the first equalizer tap.

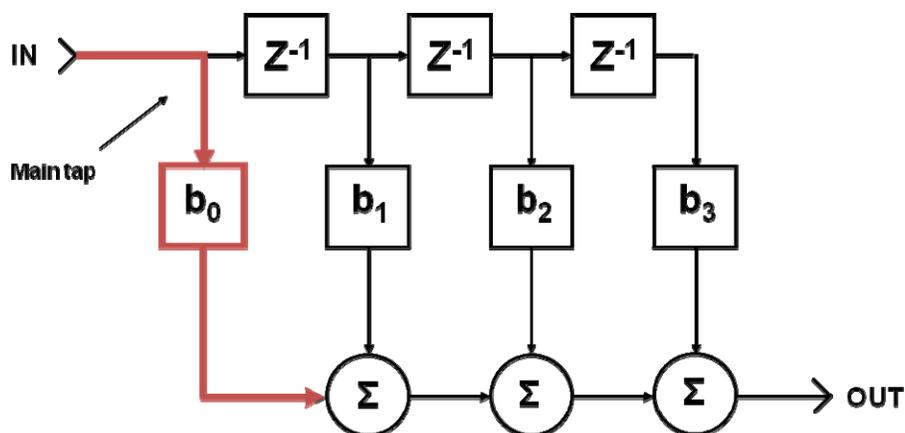


Figure I-38 - Adaptive equalizer that will be used in the example in the text

Assume the algorithm has converged on a solution that derived the following coefficients: $b_0 = 1.0$, $b_1 = +0.5$, $b_2 = +0.25$, and $b_3 = +0.125$, and each delay element Z^{-1} equals $1 \mu\text{s}$.⁸

An analog equivalent of what's going on here is something like the following: Each of the Z^{-1} delay elements is equivalent to about 1,000 feet of "lossless" coaxial cable (1 foot of cable has about 1 nanosecond of delay, so 1,000 feet of coax delays the signal by approximately 1,000 nanoseconds or $1 \mu\text{s}$). Each multiplier is equivalent to a variable attenuator. Attenuator b_0 is adjusted for no attenuation (unity gain), b_1 is adjusted to attenuate the signal by a factor of half, b_2 by a factor of one-fourth, and so on. Each of the Σ summing circuits can be thought of as backwards two-way splitters functioning as combiners, except that these, too, are assumed to be "lossless" for this example. Keep in mind that this analog equivalent is just that—an equivalent. A real-world adaptive equalizer is actually a digital circuit.

Figure I-39 shows the operation of the adaptive equalizer's first tap, which in this example also is the main tap. The unequalized input signal (1.0 amplitude incident signal and -0.5 amplitude $1 \mu\text{s}$ echo) does not pass through a delay element, but does pass through multiplier b_0 . Since the b_0 's coefficient is 1, the output of b_0 is identical to the input. The unaltered incident signal and its echo are routed to one input of the first summing circuit.

⁸ These equalizer coefficients have been simplified to illustrate this example. In practice, equalizer coefficients are complex coefficients, comprising real and imaginary components.

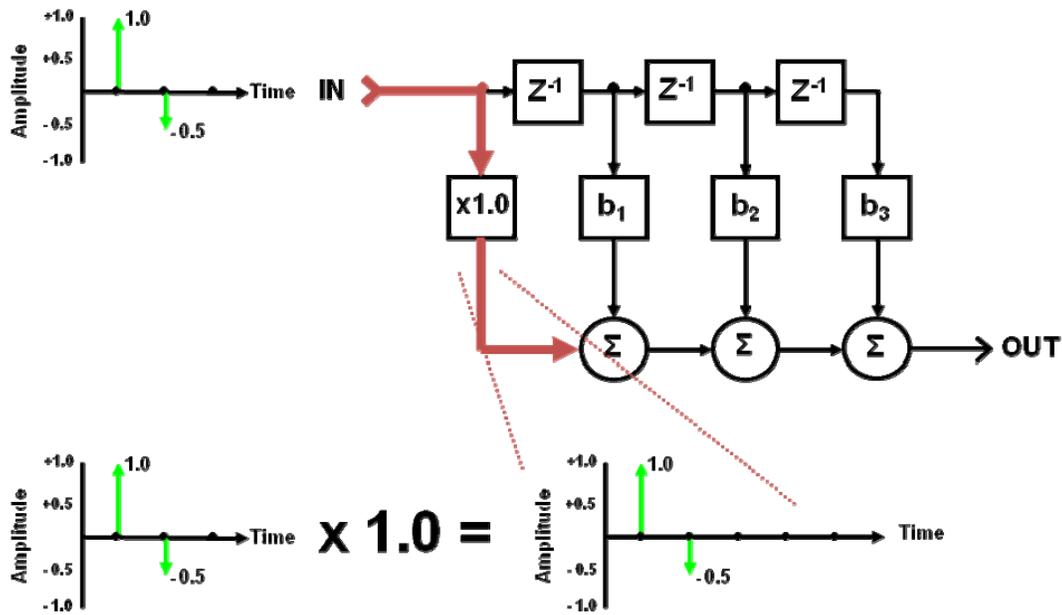


Figure I-39 - Operation of the adaptive equalizer's first tap

Figure I-40 shows the second tap's operation. The unequalized input signal passes through the first delay element Z^{-1} , which delays the incident signal and echo by $1 \mu s$. The delayed incident signal and echo next pass through multiplier b_1 , which has a coefficient of $+0.5$. The delayed incident signal and echo are multiplied by $+0.5$, which decreases the amplitude of the incident signal from 1.0 to 0.5 , and the echo from -0.5 to -0.25 . The delayed and attenuated incident signal and echo are routed to the second input of the first summing circuit.

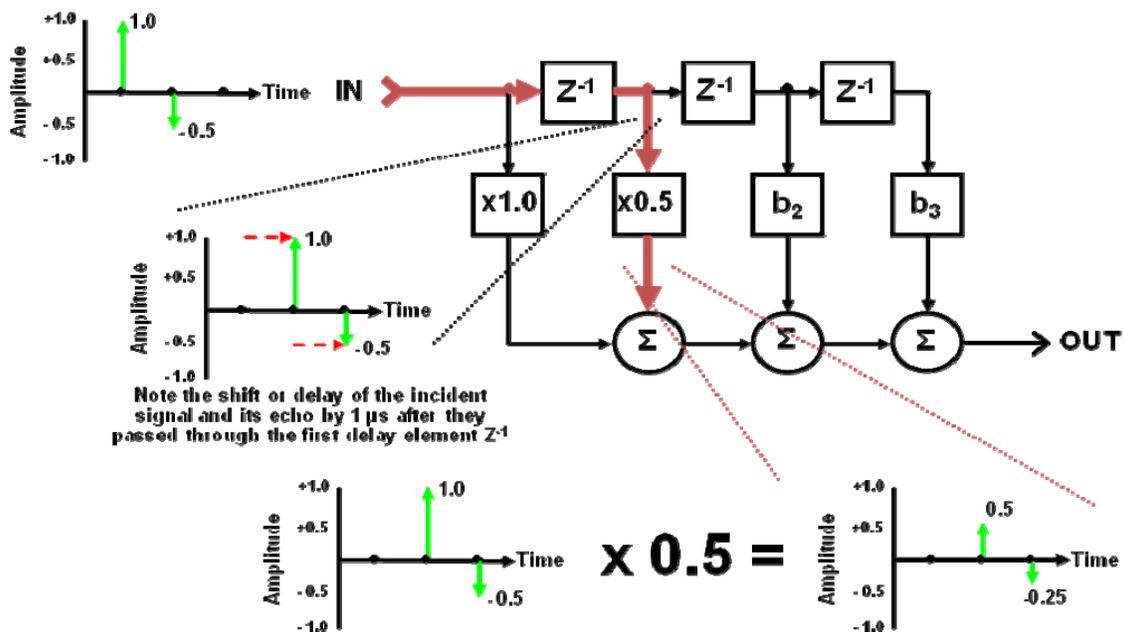


Figure I-40 - Operation of the adaptive equalizer's second tap

The first summing circuit combines the original undelayed, unattenuated incident signal and its echo with the delayed and attenuated incident signal and echo from the second tap. The output of the first summing circuit, which is connected to an input of the second summing circuit, has the original incident signal at amplitude 1.0, and a residual -0.25 amplitude $2 \mu\text{s}$ echo. The original -0.5 amplitude $1 \mu\text{s}$ echo was cancelled ($-0.5 + 0.5 = 0$). This is illustrated in Figure I-41.

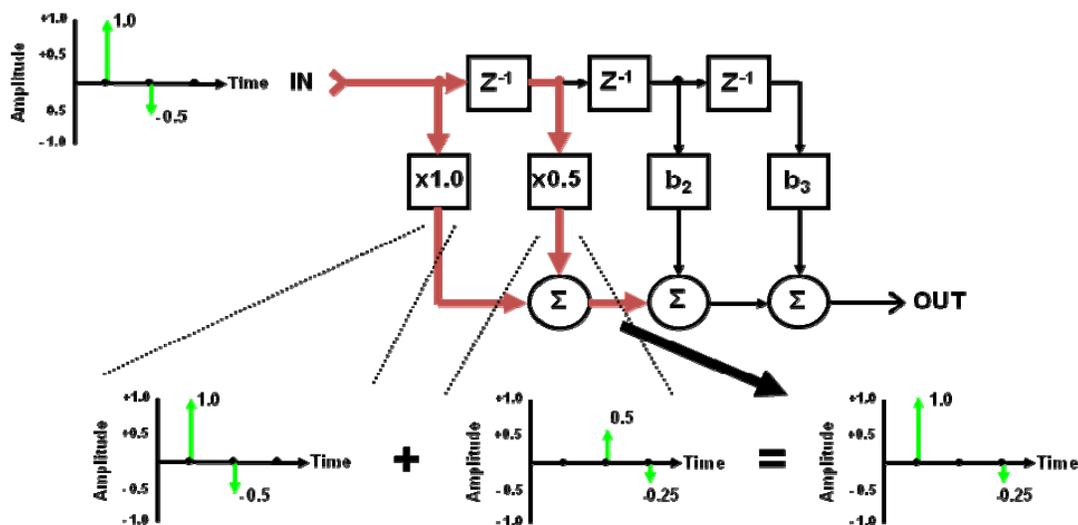


Figure I-41 - Summing the outputs of the adaptive equalizer's first and second taps

Figure I-42 shows the operation of the adaptive equalizer's third tap. The unequalized input signal passes through the first and second delay elements, which together delay the incident signal and echo by a total of $2 \mu\text{s}$. The delayed incident signal and echo next pass through multiplier b_2 , which has a coefficient of $+0.25$. The twice-delayed incident signal and echo are multiplied by $+0.25$, which decreases the amplitude of the incident signal from 1.0 to 0.25, and the echo from -0.5 to -0.125 . The delayed and attenuated incident signal and echo are routed to the second input of the second summing circuit.

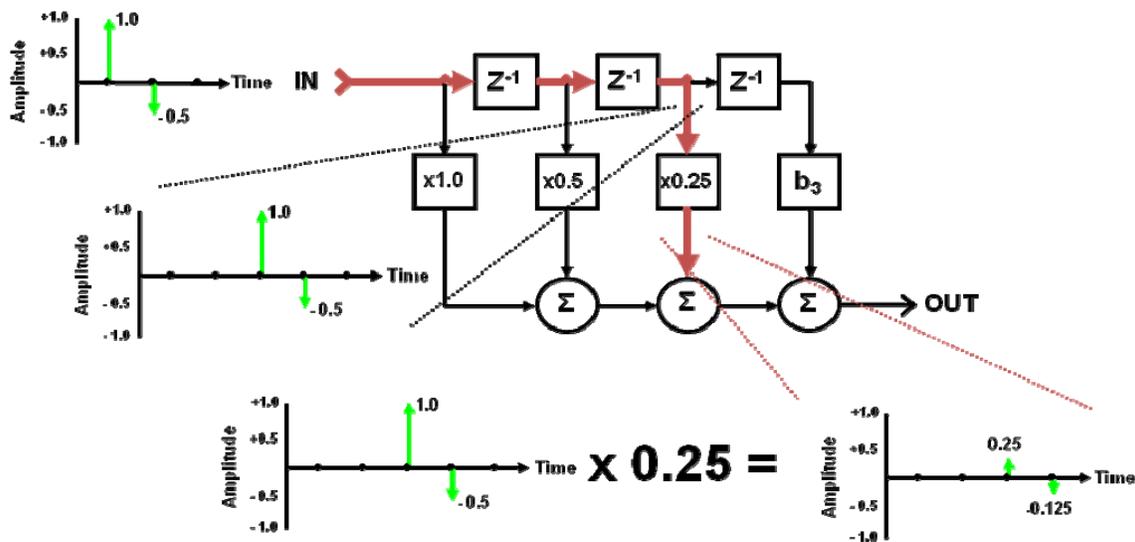


Figure I-42 - Operation of the adaptive equalizer's third tap

Referring to Figure I-43, the combined signal from the first summing circuit (1.0 amplitude incident signal and -0.25 amplitude 2 μ s echo) and the delayed and attenuated signal from the adaptive equalizer's third tap (0.25 amplitude incident signal and -0.125 amplitude echo) are combined in the second summing circuit. The output of the second summing circuit, which is connected to an input of the third summing circuit, has the original incident signal at amplitude 1.0, and a residual -0.125 amplitude 3 μ s echo. The -0.25 amplitude 2 μ s echo was cancelled (-0.25 + 0.25 = 0).

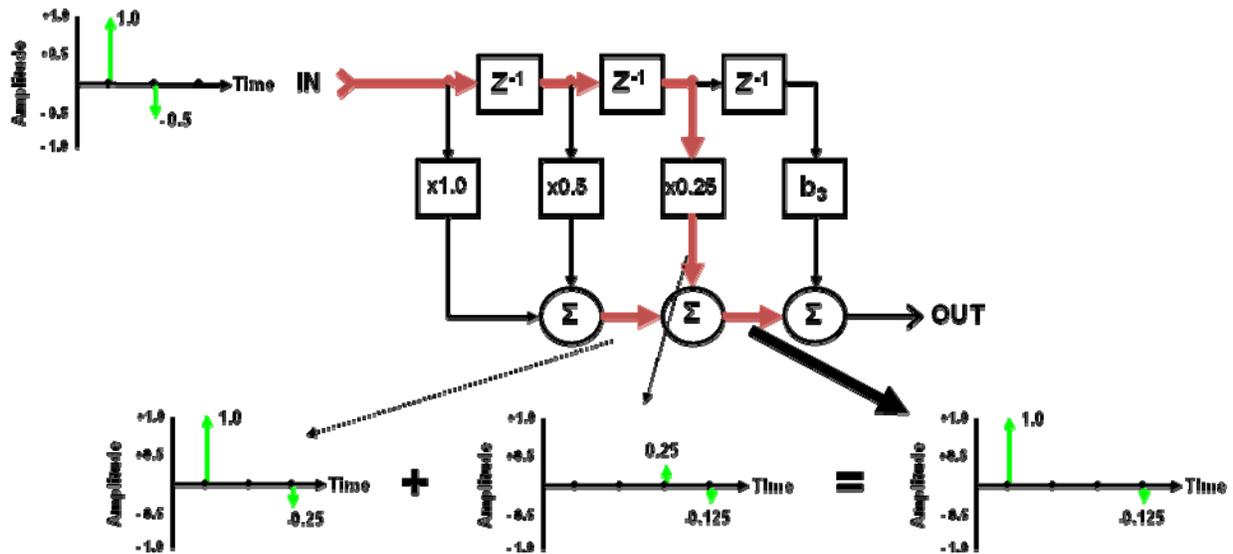


Figure I-43 - Summing the output of the adaptive equalizer's third tap with the previously summed first and second taps

Figure I-44 shows the operation of the fourth tap. The unequalized input signal passes through the first, second and third delay elements, which together delay the incident signal and echo by a total of 3 μ s. The delayed incident signal and echo next pass through multiplier b_3 , which has a coefficient of +0.125. The delayed incident signal and echo are multiplied by +0.125, which decreases the amplitude of the incident signal from 1.0 to 0.125, and the echo from -0.5 to -0.0625. The delayed and attenuated incident signal and echo are routed to the second input of the third summing circuit.

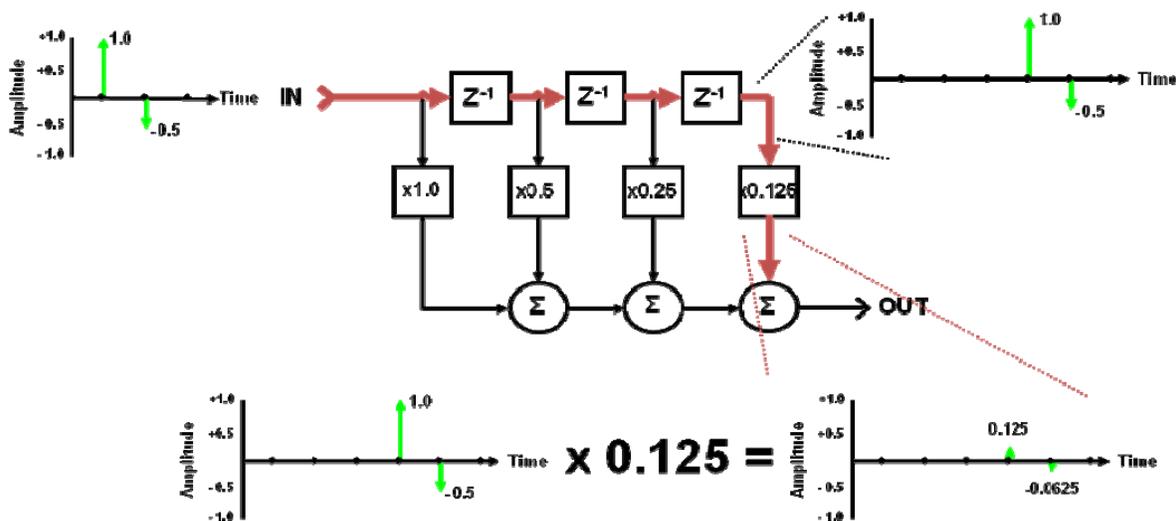


Figure I-44 - Operation of the adaptive equalizer's fourth tap

The combined signal from the second summing circuit (1.0 amplitude incident signal and -0.125 amplitude 3 μ s echo) and the delayed and attenuated signal from the fourth tap (0.125 amplitude incident signal and -0.0625 amplitude echo) are combined in the third summing circuit. The output of the third summing circuit is the equalized output signal, which has the original incident signal at amplitude 1.0, and a residual -0.0625 amplitude 4 μ s echo. The -0.125 amplitude 3 μ s echo was cancelled (-0.125 + 0.125 = 0). See Figure I-45.

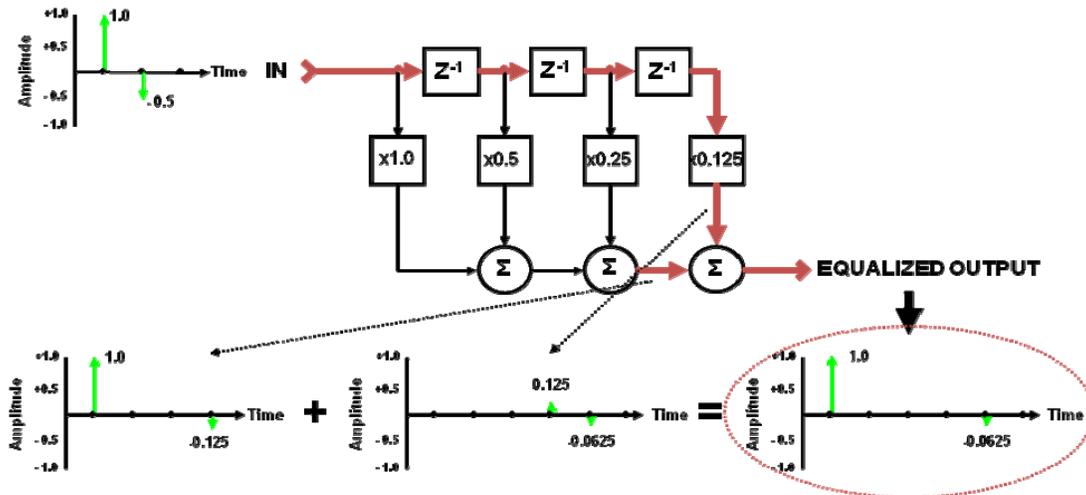


Figure I-45 - Final summing process provides an equalized output

The original -0.5 amplitude echo was reduced to an amplitude of -0.0625, or an 18 dB improvement: $20\log(-0.5/-0.0625) = 18.06$ dB. The echo also was shifted in time (delayed) to 4 μ s from the original 1 μ s.

Figure I-46 shows the resulting amplitude-versus-frequency and phase-versus-frequency response after the 4-tap equalizer has compensated for the original -0.5 amplitude 1 μ s echo. The residual echo is -0.0625 amplitude at 4 μ s, which results in the response shown in the figure. The amplitude (or magnitude)-versus-frequency response is now 1.87 dB, compared to the original 9.54 dB. The phase-versus-frequency response now is ± 3.58 degrees, compared to the original ± 30 degrees. Note that the amplitude and phase ripple, which was 1 MHz before, is now 250 kHz.

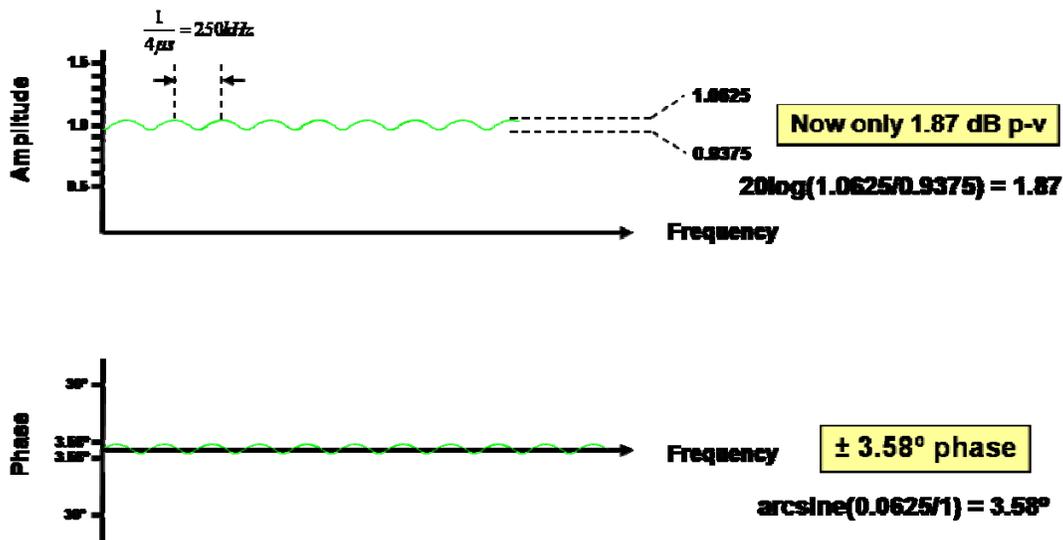


Figure I-46 - Final amplitude and phase-versus-frequency response after adaptive equalization

I.11 Adaptive Pre-equalization

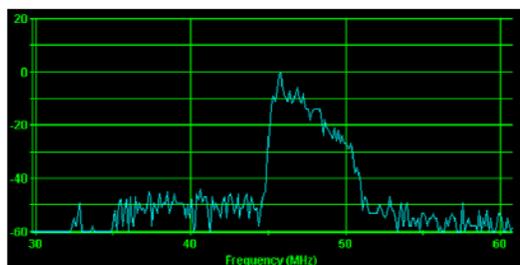
As noted in the Adaptive Equalization tutorial section, all QAM receivers—whether in a digital set-top, cable modem, QAM analyzer, or even the upstream receiver of a cable modem termination system (CMTS)—incorporate a circuit known as an adaptive equalizer. The purpose of an adaptive equalizer is to compensate for channel response impairments, by creating what amounts to a digital filter with the opposite response of the channel through which the QAM signal was transmitted. It’s called “adaptive” because the equalizer can change its characteristics on the fly, should channel response conditions change for some reason.

A cable modem or digital set-top uses a blind adaptive equalizer in the device’s downstream QAM receiver. DOCSIS 1.1, 2.0 and later cable modems are capable of equalizing—or more accurately, pre-equalizing—the upstream signal *prior* to transmission to compensate for channel response impairments. DOCSIS 1.1 supports 8-tap upstream pre-equalization, and DOCSIS 2.0 supports 24-tap upstream pre-equalization.⁹

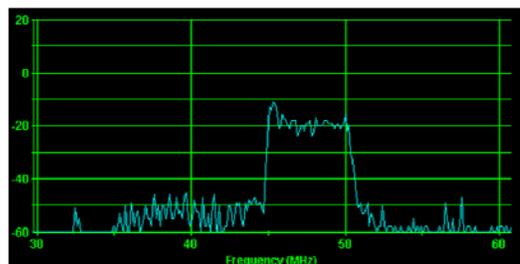
Since a cable modem is essentially a dumb box, it has no way of knowing what the upstream channel response is! So, ranging and station maintenance bursts transmitted by the modem are evaluated by the CMTS. The CMTS upstream receiver has its own adaptive equalizer, from which it derives equalizer coefficients that are transmitted to each modem. The modems, in turn, use those equalizer coefficients to configure their internal adaptive pre-equalizers to have the opposite response of the upstream channel. The “pre-distorted” signal is then transmitted by the modem, and in theory is received at the CMTS unimpaired.

You might ask why even bother with pre-equalization? Why not let the CMTS equalize each received signal? Well, the CMTS can do that (in fact, it does), but pre-equalization lets the modems do most of the adaptive equalization process’s heavy lifting.

Here’s an example. The upstream QAM signal from a given modem has several dB of in-channel tilt because of some problem in the outside plant. The CMTS “sees” that tilt and gives that modem the appropriate coefficients so that the modem can pre-equalize or tilt its transmitted QAM signal approximately the same amount in the OPPOSITE direction. When that pre-equalized QAM signal is received by the CMTS, it should now be flat. Refer to Figure I-47, which shows an upstream 6.4 MHz bandwidth 64-QAM signal as received at the CMTS. Note the substantial in-channel tilt before pre-equalization, and the same signal—now nearly flat—after pre-equalization.



Before adaptive pre-equalization:
Substantial in-channel tilt caused correctable FEC errors to increment at a rate of about 7000 errored codewords per second (232 bytes per codeword). The CMTS’s reported upstream MER (SNR) was 23 dB.



After adaptive pre-equalization:
DOCSIS 2.0’s 24-tap adaptive equalization—actually pre-equalization in the modem—was able to compensate for nearly all of the in-channel tilt. The result: No correctable or uncorrectable FEC errors and the CMTS’s reported upstream MER (SNR) increased to ~36 dB.

Figure I-47 - Upstream pre-equalization is able to compensate for in-channel tilt

⁹ “8-tap” and “24-tap” refer to the number of adaptive equalizer taps used for upstream pre-equalization in DOCSIS implementations.

I.12 Velocity of Propagation

Electromagnetic waves travels through a vacuum at the speed of light, 299,792,458 meters per second (983,571,056 feet per second). This is known as the free-space value of the speed of light. Since radio waves are part of the electromagnetic spectrum, they travel at the speed of light, too. But what happens when radio waves travel through coaxial cable? Their velocity is somewhat slower than it is in a vacuum! Indeed, radio waves travel through trunk and feeder-type coaxial cable at approximately 87% of the free space value of the speed of light, or 260,819,438 meters per second (855,706,819 feet per second). Looking at this another way, radio waves travel 1 foot in a vacuum in 1.02 nanosecond (ns), and through 1 foot of coaxial cable in 1.17 ns.

Why do radio waves travel more slowly through coaxial cable than they do through a vacuum? First, the conductors used in coaxial cable are not perfect. The loss in those conductors slows down the waves slightly, but the effect is almost negligible at frequencies used in cable networks. Of more importance is the effect of the dielectric material which separates the coaxial cable's center conductor and shield. Indeed, the presence of a dielectric other than a vacuum or air reduces the velocity of an electromagnetic wave, often by 10% to 20% or more.

The ratio of the velocity of an electromagnetic wave—specifically what is known as a transverse electromagnetic (TEM) mode wave—in a vacuum to its velocity in a dielectric material, $v_{\text{TEM}}(\text{vacuum})/v_{\text{TEM}}(\text{dielectric})$, equals what is called index of refraction.¹⁰ Velocity factor is the reciprocal of index of refraction.

The dielectric's magnetic permeability (represented by the symbol μ and expressed in henrys/meter) and electric permittivity (represented by the symbol ϵ and expressed in farads/meter) are two key properties that determine the velocity of electromagnetic waves in coaxial cable. The ratio of $\epsilon(\text{dielectric})/\epsilon(\text{vacuum})$ is the dielectric constant ϵ_r . The velocity of an electromagnetic wave in a dielectric is equal to the ratio of its velocity in a vacuum to the square root of the dielectric constant: $v_{\text{TEM}}(\text{dielectric}) = v_{\text{TEM}}(\text{vacuum})/\sqrt{\epsilon_r}$. A little number crunching with the latter equation and the ratio that defines index of refraction is $1/\sqrt{\epsilon_r} = \text{velocity factor}$.

The dielectric constant ϵ_r of the coaxial cable in the example in the first paragraph is approximately 1.32, so the velocity factor is $1/\sqrt{1.32} = 0.87$. Velocity of propagation is velocity factor expressed as a percentage. So, a velocity factor of 0.87 is 87% velocity of propagation, which means that radios waves travel through the coaxial cable at 87% the free-space value of the speed of light.

One application for understanding velocity of propagation is calculating propagation or transit delay in a cable network—that is, the time it takes for electromagnetic waves to travel from one point to another.

The Data-Over-Cable Service Interface Specification (DOCSIS®) *Radio Frequency Interface Specification* includes assumed downstream and upstream RF channel transmission characteristics for cable networks. Among those assumed characteristics is the previously mentioned transit delay. For instance, the transit delay from the headend to the most distant customer is assumed to be less than or equal to 0.800 millisecond (ms). Note that 0.800 ms (800 microseconds) is a one-way specification. The same assumed transit delay also applies in the upstream direction.

The approximate downstream or upstream transit delay can be calculated if one knows the length of the optical fiber link between the headend or hub and node, as well as the length of distribution network coaxial cable from the node to the most distant customer. The calculation is done using the reciprocal of the fiber's index of refraction—its velocity factor, which is then converted to velocity of propagation—and the velocity of propagation of the coaxial cable. The approximate index of refraction for single mode optical fiber at 1310 nm is 1.46, making its velocity factor 0.68 and its velocity of propagation 68%. In other words, light propagates through the optical fiber at a velocity that is 68% of the speed of light in a vacuum. A typical velocity of propagation for commonly used hardline distribution-type coaxial cables is the previously discussed 87%.

¹⁰ The ratio $v_{\text{TEM}}(\text{vacuum})/v_{\text{TEM}}(\text{dielectric})$ is called the index of refraction because the difference in the velocity of electromagnetic waves in a vacuum and some other medium results in refraction at the interface between the two media.

Since the speed of light in a vacuum is 299,792,458 meters per second, 68% of that value is 203,858,871 meters per second. Thus, light will propagate through 203,858,871 meters of optical fiber in one second. For coaxial cable with a velocity of propagation of 87%, RF will propagate through 260,819,438 meters of coax in one second.

Example: Assume a cable system with an optical fiber link from headend to node that is 30 kilometers (km) long. The coaxial cable distribution network connected to the node has a coax run that extends an additional 2 km beyond the node. What is the approximate transit delay from the headend to the most distant customer, excluding delay through active and passive devices?

Solution: Light propagates through 30 km (30,000 meters) of optical fiber in 1.47×10^{-4} second (30,000 meters/203,858,871.44 meters per second = 0.00014716 second). RF propagates through 2 km (2,000 meters) of coax in 7.6681×10^{-6} second (2,000 meters/260,819,438.46 meters per second = 0.0000076681 second). Combining these two numbers yields 1.55×10^{-4} second, or 0.155 ms. This is well within the DOCSIS one-way transit delay specification of 0.800 ms.

Table I-2 summarizes transit delay in ns-per-foot and ns-per-meter for several values of velocity of propagation. The velocity factor of a vacuum is 1.0 and its velocity of propagation is 100%, because electromagnetic signals travel at the free-space value of the speed of light.

The dielectric constant of dry air at a pressure of one atmosphere and a temperature of 23° C is 1.00068, so the velocity factor is $1/\sqrt{1.00068} = 0.999660173$ and the velocity of propagation is 99.966%. The values for a vacuum and air are usually considered to be the same in all but the most critical applications because of the negligible difference between them.

Table I-2 - Velocity of propagation versus transit delay

Velocity of Propagation	ns/foot	ns/meter	Velocity of Propagation	ns/foot	ns/meter
100%	1.02	3.34	81%	1.26	4.12
99%	1.03	3.37	80%	1.27	4.17
98%	1.04	3.40	79%	1.29	4.22
97%	1.05	3.44	78%	1.30	4.28
96%	1.06	3.47	77%	1.32	4.33
95%	1.07	3.51	76%	1.34	4.39
94%	1.08	3.55	75%	1.36	4.45
93%	1.09	3.59	74%	1.37	4.51
92%	1.11	3.63	73%	1.39	4.57
91%	1.12	3.67	72%	1.41	4.63
90%	1.13	3.71	71%	1.43	4.70
89%	1.14	3.75	70%	1.45	4.77
88%	1.16	3.79	69%	1.47	4.83
87%	1.17	3.83	68%	1.50	4.91
86%	1.18	3.88	67%	1.52	4.98
85%	1.20	3.92	66%	1.54	5.05
84%	1.21	3.97	65%	1.56	5.13
83%	1.22	4.02	64%	1.59	5.21
82%	1.24	4.07	63%	1.61	5.29

Typical published velocities of propagation for modern foam dielectric coaxial cables used by the cable industry are 84-85% for drop-type cables and 87-88% for hardline trunk and feeder cables. Published values for disc-and-air dielectric designs are as high as 93%. As previously noted, the typical velocity of propagation for single mode optical fiber at 1310 nm is about 68%.

Appendix II SNMP Collection Methodology

The fundamental requirements of the SNMP protocol are:

- network access (UDP/IP, port 161)

- community string

- IP address of each cable modem and/or CMTS

- Correctly formatted object identifiers (OIDs), including proper indexing

The initial CMTS step may be excluded if only the modem data is being considered for analysis. Otherwise, starting with the MAC address of the cable modem and the IP address of the CMTS, the pointer of the modem index must be resolved. This is achieved by passing the decimal formatted MAC address to the CMTS by way of the docsIfCmtsCmPtr MIB.

Example:

MAC = AA:BB:CC:00:11:22, converted to decimal would be 170.187.204.0.17.34

The SNMP command formation might look like this:

```
snmpget -v2c -c $COMMUNITY $CMTS_IP 1.3.6.1.2.1.10.127.1.3.7.1.2. 170.187.204.0.17.34
```

The resulting value would be the modem index required to obtain the IP address, coefficients, and other upstream information from the CMTS. Subsequent queries to the docsIfCmtsCmStatusTable will require this modem index. To obtain the modem IP and equalization data from the CMTS, a single SNMP query may be made by requesting the docsIfCmtsCmStatusIpAddress and docsIfCmtsCmStatusEqualizationData. Both of these OIDs require the previously obtained modem index. Example:

```
snmpget -v2c -c $COMMUNITY $CMTS_IP 1.3.6.1.2.1.10.127.1.3.3.1.3.$INDEX  
1.3.6.1.2.1.10.127.1.3.3.1.8.$INDEX
```

The resulting values will be the IP address of the cable modem and also the upstream pre-equalization coefficients as viewed by the CMTS.

Having the IP address of the modem, the three required values may be obtained by a query to UpStreamFrequency, UpStreamWidth and docsIfCmStatusEqualizationData. An example of this command formation would be:

```
snmpget -v2c -c $COMMUNITY $MODEM_IP 1.3.6.1.2.1.10.127.1.1.2.1.2.4 1.3.6.1.2.1.10.127.1.1.2.1.3.4  
1.3.6.1.2.1.10.127.1.2.2.1.17.2
```

Appendix III SNMP MIB Reference

REQUIRED FROM MODEM:

- 1.3.6.1.2.1.10.127.1.1.2.1.2.4 UpStreamFrequency
- 1.3.6.1.2.1.10.127.1.1.2.1.3.4 UpStreamWidth
- 1.3.6.1.2.1.10.127.1.2.2.1.17.2 docsIfCmStatusEqualizationData

OPTIONAL FROM MODEM:

- 1.3.6.1.2.1.1.1 sysDescr
- 1.3.6.1.2.1.1.2 sysObjectID
- 1.3.6.1.2.1.1.3 sysUpTime
- 1.3.6.1.2.1.10.127.1.1.4.1.2.3 docsIfSigQUnerrored
- 1.3.6.1.2.1.10.127.1.1.4.1.3.3 docsIfSigQCorrecteds
- 1.3.6.1.2.1.10.127.1.1.4.1.4.3 docsIfSigQUncorrectables
- 1.3.6.1.2.1.10.127.1.2.2.1.12.2 docsIfCmStatusT3Timeouts
- 1.3.6.1.2.1.10.127.1.2.2.1.13.2 docsIfCmStatusT4Timeouts
- 1.3.6.1.2.1.10.127.1.2.2.1.4.2 docsIfCmStatusResets
- 1.3.6.1.2.1.10.127.1.2.2.1.5.2 docsIfCmStatusLostSynchs
- 1.3.6.1.2.1.10.127.1.1.4.1.5.3 docsIfSigQSignalNoise
- 1.3.6.1.2.1.10.127.1.1.5.0 docsIfDocsisBaseCapability
- 1.3.6.1.2.1.10.127.1.1.1.1.2.3 docsIfDownChannelFrequency
- 1.3.6.1.2.1.10.127.1.2.2.1.3.2 docsIfCmStatusTxPower
- 1.3.6.1.2.1.10.127.1.1.1.1.3.3 DownStreamWidth
- 1.3.6.1.2.1.10.127.1.1.1.1.4.3 docsIfDownChannelModulation
- 1.3.6.1.2.1.10.127.1.1.2.1.4.4 docsIfUpChannelModulationProfile
- 1.3.6.1.2.1.10.127.1.1.1.1.6.3 docsIfDownChannelPower
- 1.3.6.1.2.1.10.127.1.1.2.1.6.4 docsIfUpChannelTxTimingOffset

REQUIRED FROM CMTS:

- 1.3.6.1.2.1.10.127.1.3.7.1.2 docsIfCmtsCmPtr
- 1.3.6.1.2.1.10.127.1.3.3.1.3 docsIfCmtsCmStatusIpAddress
- 1.3.6.1.2.1.10.127.1.3.3.1.8 docsIfCmtsCmStatusEqualizationData

OPTIONAL FROM CMTS:

- 1.3.6.1.2.1.10.127.1.3.3.1.6 docsIfCmtsCmStatusRxPower
- 1.3.6.1.2.1.10.127.1.3.3.1.9 docsIfCmtsCmStatusValue
- 1.3.6.1.2.1.10.127.1.3.3.1.13 docsIfCmtsCmStatusSignalNoise
- 1.3.6.1.2.1.10.127.1.3.3.1.7 docsIfCmtsCmStatusTimingOffset

OPTIONAL FROM CMTS US:

- 1.3.6.1.2.1.2.2.1.2 ifDescr
- 1.3.6.1.2.1.10.127.1.1.2.1.19 docsIfUpChannelPreEqEnable

OPTIONAL FROM CMTS US (vendor specific):

- 1.3.6.1.4.1.4998.1.1.15.3.2.1.7 Arris specific upstream padding
- 1.3.6.1.4.1.9.9.116.1.4.1.1.6 Cisco specific upstream padding (divide by 10)
- 1.3.6.1.4.1.4981.2.1.2.1.1 Motorola/RDN specific upstream padding (divide by 10)

Appendix IV Micro-reflection Calculator

A simple micro-reflection calculator to derive the relation between the amplitude ripple and the reflected wave level is shown here.

For simplification purposes it is assumed that the interference pattern consists of the incident signal combined with a single reflected wave. For this, take advantage of the voltage standing wave ratio (VSWR) definition and others already in place.

The equivalent reflection coefficient is defined by

$$\Gamma_E = \frac{Z_E - Z_0}{Z_E + Z_0}$$

Where Z_0 is the characteristic impedance of the transmission line, in this case 75 ohms, and Z_E is the equivalent impedance of the reflected wave that adds to and interferes with the incident signal.

It is considered an equivalent impedance because it contains the impact of the reflection against two mismatched interfaces as well as the round trip attenuation and delay between these two interfaces (See Figure IV-1).

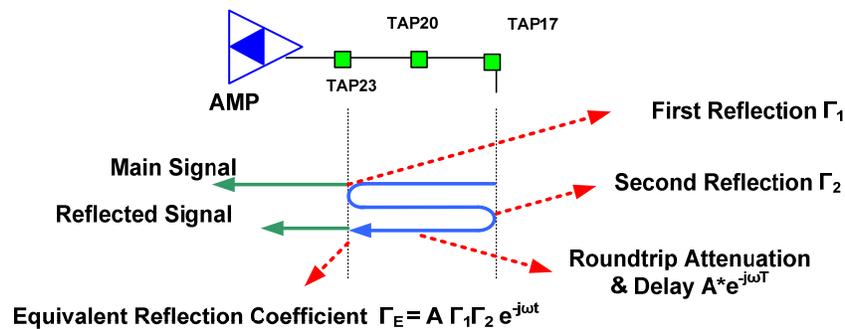


Figure IV-1 - Equivalent Reflection Coefficient Γ_E

The VSWR, which is the ratio between E_{\max} and E_{\min} , reveals the ripple magnitude. The VSWR in terms of the reflection coefficient is given by:

$$VSWR = \frac{1 + \Gamma_E}{1 - \Gamma_E}$$

Since VSWR is a ratio of voltages, to obtain the equivalent peak-to-valley power ratio or ripple in dB,

$$\text{Ripple (dB)} = 10 \log_{10}(VSWR^2)$$

Alternatively the process can be reversed by solving for Γ_E in the previous equation to obtain:

$$\Gamma_E = \frac{VSWR - 1}{VSWR + 1}$$

The equivalent reflection coefficient Γ_E provides the micro-reflection level in dB in the following expression

$$\text{Micro-reflection Level (dB)} = 10 \log_{10}(\Gamma_E^2)$$

Appendix V Two Types of Echoes

Portions of Appendix V have been previously published in a February 2010 *CED* magazine online article. An enhanced version of that article is included here with permission of the authors.

Upstream Cable Echoes Come In Two Flavors

Thomas H Williams, ARRIS Group

Alberto Campos, CableLabs

Bruce Currivan, Broadcom

Charles Moore, Motorola

Background

CableLabs has established a “DOCSIS Proactive Network Maintenance” working group with a goal of utilizing the data in DOCSIS-based communications systems to improve plant operations. The group’s first task is to produce a recommended practices document that enables cable operators to mine the pre-distortion coefficients from the cable modem’s (CM) upstream adaptive equalizer.

Cable lines typically have many small echoes, so-called micro-reflections, which will disrupt digital transmissions if not canceled. This is particularly true for high-speed upstream signals having a higher order modulation, such as 64-QAM, or having wider bandwidth, such as 6.4 MHz. The system chosen by DOCSIS uses pre-distortion (or pre-equalization), where a burst transmission is distorted prior to transmission, and arrives at a cable modem termination system (CMTS) receiver with the plant’s distortion canceled. The idea is that by reading the CM’s predistortion coefficients using a network management system, technicians can tell what plant impairments a CM is compensating for, and then compute what may be wrong with the cable plant. By reading the data from many CMs you can localize the problems using maps or connectivity data.

The process of programming the predistortion coefficients in the CM is handled during a periodic ranging process, which is controlled by the CMTS.

The group’s name, ‘DOCSIS Proactive Network Maintenance’, is somewhat limiting since the techniques being developed show a reactive ability to speed time-to-repair. The utility of the technique is that it can reduce expensive truck-rolls, either reactively or proactively.

Two Types of Echoes

As a result of data presented in the DOCSIS Proactive Network Maintenance face-to-face meetings and from data gathered in experiments, it appears that upstream cable plant has two distinctly different types of echoes with different mathematical properties. The first type is called multiple recursion, and is created by two or more impedance discontinuities inside the cable. This type of echo has a characteristic impulse response of a main signal followed by several echoes, with each succeeding echo having smaller amplitude. (An impulse response is nothing more than a time plot of how the signal path would respond to a narrow voltage spike.) The second type is called single recursion, where echoes that are created find two different routes upstream. This type of echo has a characteristic impulse response of a main signal followed by a single echo. The good news is that the existing upstream adaptive equalizer corrects both automatically. Also, if the echoes are not strong, only the first recursion is significant in both cases. For example, the second recursion of a -20 dBc echo is -40 dBc, which is generally below the noise threshold.

Multiple Recursion Echoes

In this echo situation, an upstream transmission gets bounced back and forth between two impedance discontinuities. The multiple recursion echo may be created in a situation such as is shown in the top of Figure V-1. An upstream signal travels from right to left through an amplifier, a span of cable with taps, and another amplifier. There are two impedance mismatches, labeled reflection points, which cause a portion of the upstream signal to reflect back and forth until the reflections eventually die out. The bottom of Figure V-1 shows the impulse response of the upstream channel. Note that there is a main signal, followed by multiple recursions, each caused by a re-reflection. In this example, the echo is very strong, so there are many significant recursions.

The multiple recursion scenario also has been observed in a drop cable, where a filter on one end has bad upstream return loss, and a house on the other end also has bad return loss due to un-terminated splitters. A CM's transmission from the house picks up multiple recursions.

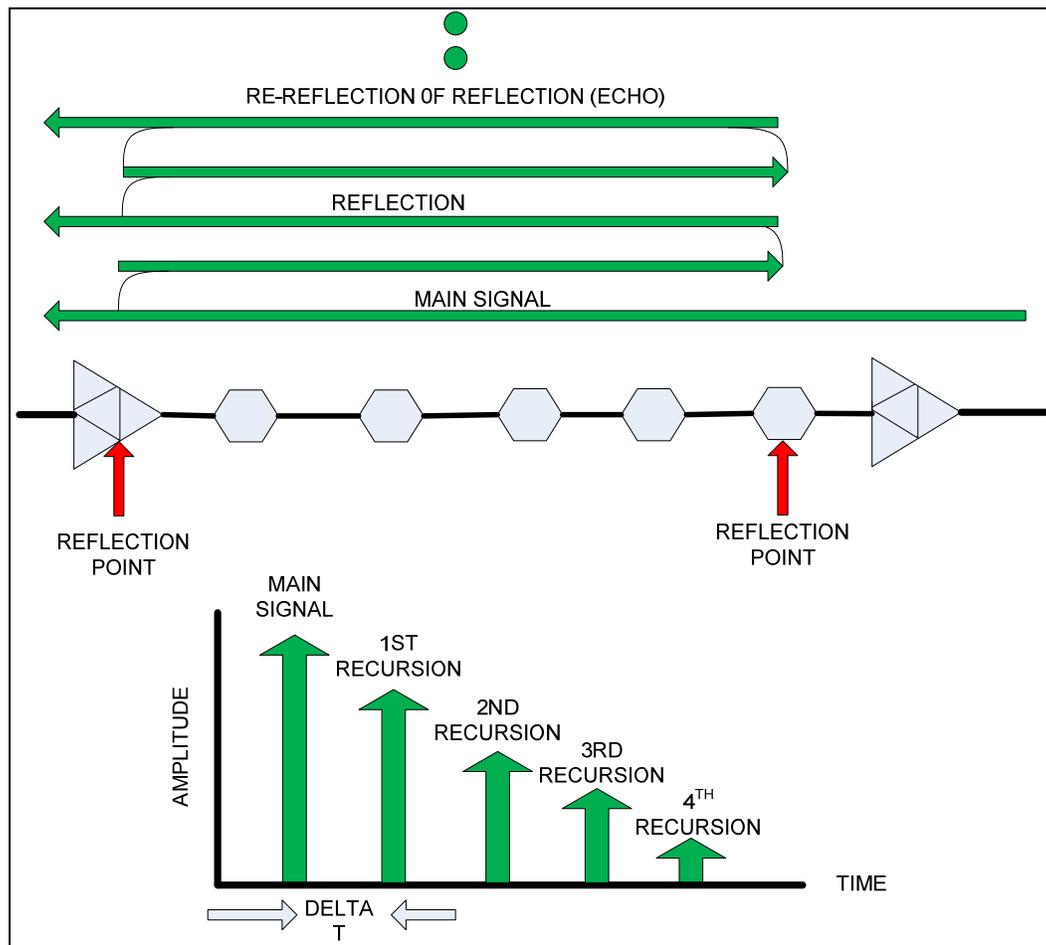


Figure V-1 - A Multiple Recursion Echo

Multiple Recursion Lab Echoes

The multiple recursion echo situation can be created in a lab environment with the diagram shown in the top of Figure V-2. A pair of directional couplers is chosen with tap values selected to adjust the strength of the echo. A cable length is chosen to give the desired echo delay. The tap leg of each directional coupler is left open to create a reflection. Assuming two-way splitters are used with 3.5 dB of insertion loss on each leg, and cable loss is 2 dB, 1st recursion echo strength is computed as:

Main path dB - echo path dB = (3.5 + 2.0 + 3.5) – (3.5 + 2.0 + 3.5 + 3.5 + 2.0 + 3.5 + 3.5 + 2.0 + 3.5) = 9.0 - 25.0 = 16 dB

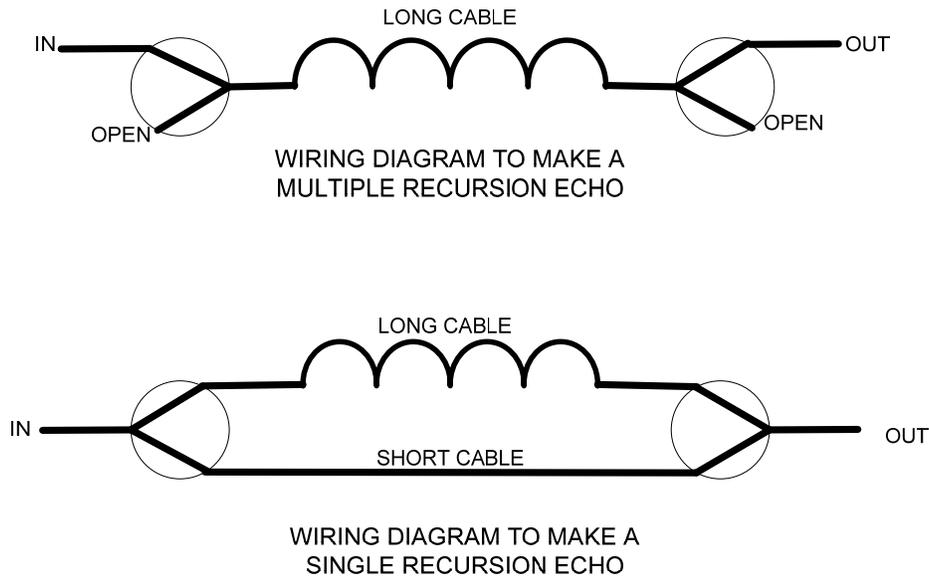


Figure V-2 - Wring Diagrams To Make Echoes In a Lab

Single Recursion Echo

In this echo situation, the transmission finds two different paths upstream. One scenario for this echo situation is illustrated in the top of Figure V-3. A signal is being transmitted from a house into a 28 dB tap. Unfortunately this tap has a port-to-output isolation of only 35 dB. The CM’s signal travels upstream attenuated by the tap’s value (28 dB), but the signal also travels downstream attenuated by only 35 dB. When the signal reaches the end of the line it encounters a two-way splitter which is not terminated. The lack of a termination causes the signal to reflect back upstream where it rejoins, after some time delay, the main signal. Because the input return loss of the amplifier and 28 dB tap is excellent, there is no further recursion of the echo. The resulting impulse response is shown the bottom of Figure V-3.

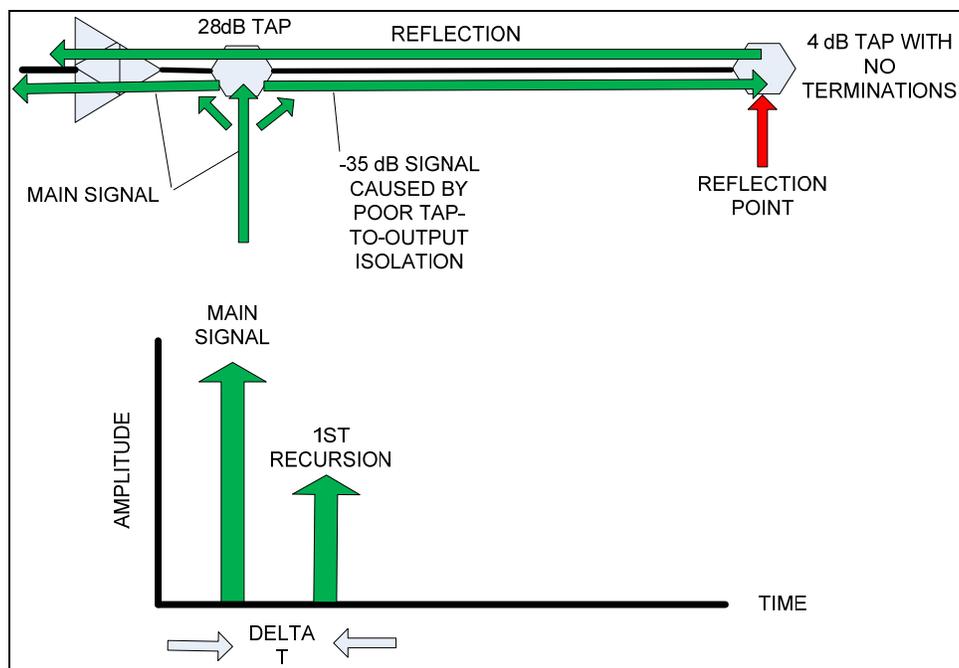


Figure V-3 - A Single Recursion Echo Example

Creating a Single Recursion Echo

A single recursion echo can be created using the wiring diagram in the bottom of Figure V-2. A signal is split by the first splitter on the left. One leg of the left splitter connects to a short piece of cable and the other connects to a long piece of cable. The ends of the two cables are combined to make an output signal in the right combiner. An attenuator (not shown) may be put in line with the long piece of cable to further attenuate the echo to a desired amplitude.

Echo Cancellation with Single and Multiple Recursion Echoes

It is well known by digital signal processing engineers that a single recursion echo can be canceled by an adaptive equalizer with enough taps to handle all of the significant recursions created in the equalizer, and that several recursions may be needed to cancel a strong echo. Multiple recursions are necessary to cancel a single echo because the adaptive equalizer sums an echo-corrupted received signal with a delayed echo-corrupted copy. The equalizer cancels the first recursion of the echo, but the echo in the delayed signal is not canceled. So yet another recursion is needed to take out the echo in the signal the equalizer used for cancellation, and so forth, repeating endlessly.

It is suspected (and confirmed by preliminary field and lab data) that under the proper conditions, a multiple recursion echo can be canceled by a pre-distorted signal with only one recursion. How this works is explained in Figure V-4. This diagram is a snapshot in time. At the illustrated point in time, a main signal's impulse already has propagated upstream past point A. A reflection created at REFLECTION POINT A has propagated backwards and is now at point C. The CM has already transmitted an inverse first recursion, shown at point D. When the signals propagating from point C and point D meet at REFLECTION POINT B they sum together and cancel. End of echo, period. No more recursions.

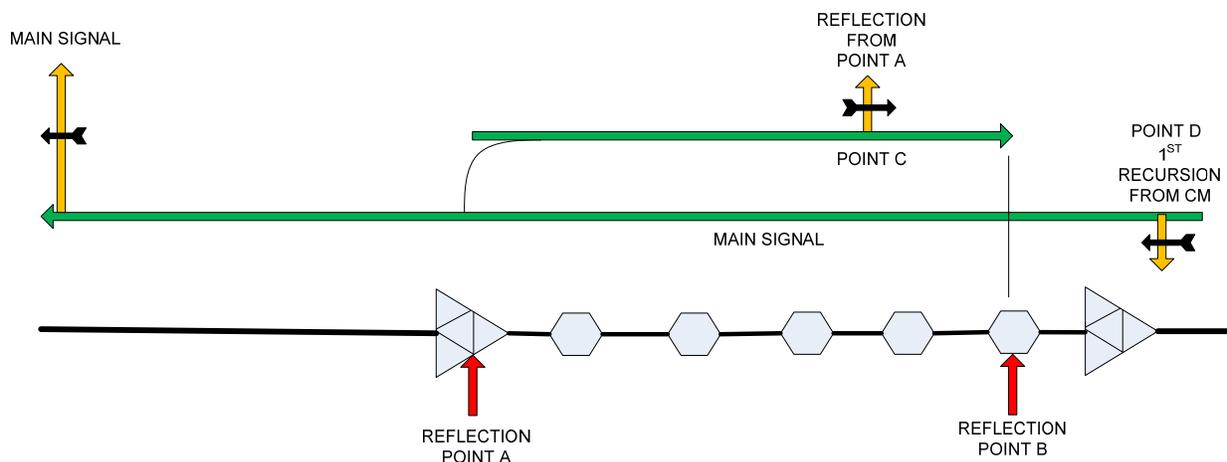


Figure V-4 - How A Multiple Recursion Echo Can be Canceled with Pre-distortion

The Math:

A multiple recursion echo may be modeled as:

$$1 + a + a^2 + a^3 + a^4 + \dots$$

where 1.0 is the main signal amplitude and ‘a’ is the first echo’s amplitude in linear terms. That is, for a -3 dBc echo a = 0.707. The ‘a²’, ‘a³’, ‘a⁴’ components and so on represent respectively the amplitude of the recurring second, third, fourth and higher order echoes. In a baseband channel ‘a’ is real, but in an RF channel ‘a’ may be complex. To be more precise,

$$a = Ae^{-j2\pi fT}$$

where A is the amplitude of the echo, f is carrier frequency (MHz) and T is the delay of the echo (µsec). However, the equations are easier to present by simply using ‘a’.

So its solution, meaning the inverse distortion that has to be applied for upstream path distortion compensation, can be computed as:

$$\frac{1}{1 + a + a^2 + a^3 + a^4 - \dots} = 1 - a$$

which shows that an infinitely recursive echo can be canceled by a single recursion. So you should be able to cancel a multiple recursion echo with an adaptive equalizer with only two taps – one for the main signal and one for the echo. To be precise, this example applies to an ideal case where the echo delay T equals a multiple of the symbol period, which for a 5.12 Msym/sec DOCSIS upstream symbol rate is Ts = 195 ns. So, the pre-equalizer can exactly cancel the echo with a single tap if the single echo has delay T = 195 ns, 390 ns, or 585 ns, etc. If the echo lies between these multiples, the equalizer will activate additional taps to provide interpolation. In that case it will be more difficult to see the pattern of a single main recursion.

A single recursion echo can be modeled by 1 + a where 1.0 is the main signal amplitude and ‘a’ is the echo’s amplitude in linear terms.

The equalizer solution is for a single recursion echo is:

$$\frac{1}{1+a} = 1 - a + a^2 - a^3 + a^4 - \dots$$

So the result is what was expected...infinitely recursive. The first term '1' represents the main tap of the equalizer. The second term '-a' acts to cancel the echo by subtracting it. However, in doing so, it causes another, smaller echo in the response. This smaller echo requires the third term to cancel it. This produces another, yet smaller echo, which requires the fourth term to cancel it, and so on until reaching the end of the equalizer delay line. After that, any remaining echo energy (hopefully very small) is not canceled, and shows up as reduced RxMER (received modulation error ratio), essentially a noise floor, in the receiver.

If 'a' is a relatively big number such as 0.707 (-3 dBc), and the delay T of the echo is several symbol periods, the taps might run out in the pre-equalizer before getting an accurate solution. DOCSIS 2.0 and later pre-equalizers have 24 taps, with 7 taps normally assigned ahead of the main tap, leaving 16 taps to cancel the echo. For a small value of 'a' such as 0.1 (-20 dBc), with a short echo delay, 16 taps are normally sufficient to cancel the echo with minimal residual energy.

Note that the recursions in the previous equation have alternating signs. To check for this effect, examine the real and imaginary parts of the equalizer taps. If the response is alternating in sign, it is a hint that this type of solution may be present.

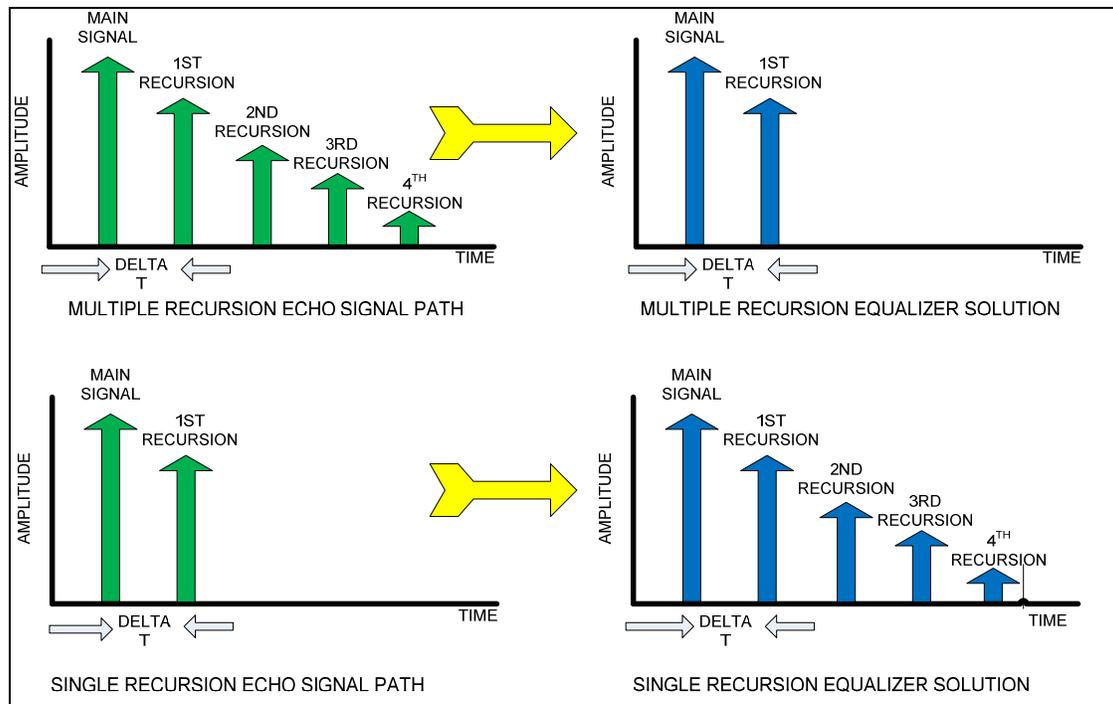


Figure V-5 - Comparison of Signal Path Impulse Responses and Programming for Adaptive Equalizers

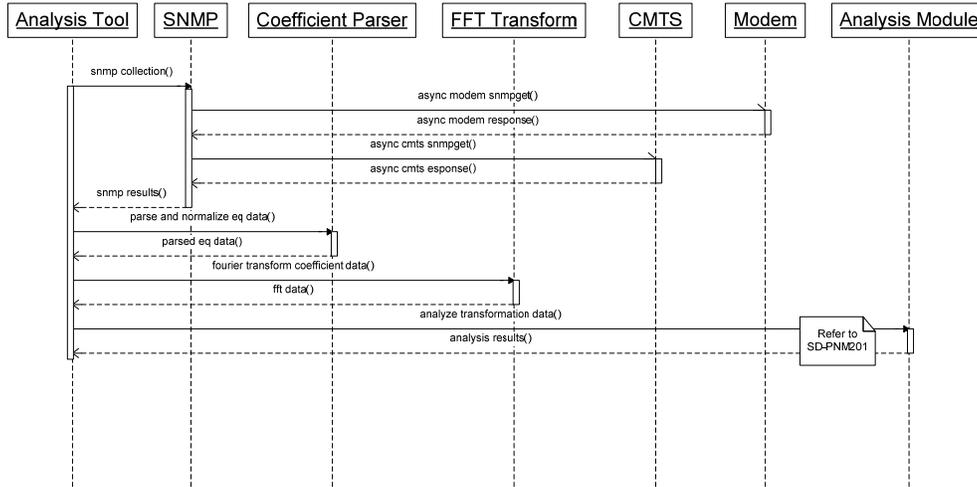
Figure V-5 summarizes both types of echoes and the resulting programming that could be found in the adaptive equalizer to cancel the echoes.

Conclusion

Upstream cable plant can have two distinctly different types of echoes, but if the echoes are relatively weak, the differences may not be significant. If the echoes are strong the differences can be exploited to help a cable technician diagnose and fix the cause of the strong echo.

Appendix VI DOCSIS Pre-equalizer Coefficients Analysis - Software Sequence Diagrams

VI.1 Software Sequence Diagram SD-PNM200



Note: Analysis module functionality is covered in SD-PNM201

Figure VI-1 - Software Sequence Diagram SD-PNM200

VI.2 Software Sequence Diagram SD-PNM201

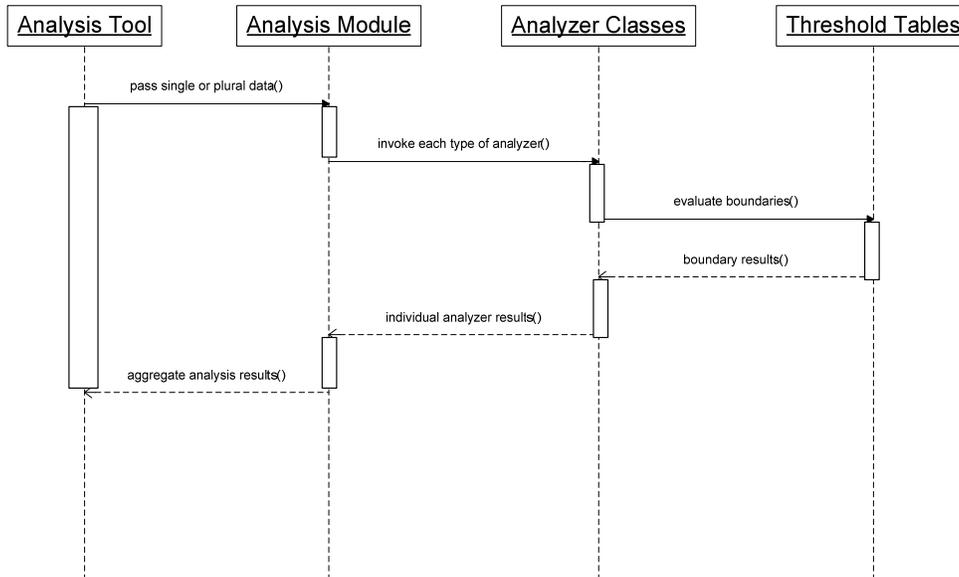


Figure VI-2 - Software Sequence Diagram SD-PNM201

Appendix VII Acknowledgements

On behalf of the cable industry and our member companies, CableLabs would like to thank the numerous individuals who contributed to the development of this document. In particular, we want to extend our sincere appreciation and gratitude to the following Proactive Network Maintenance Working Group members.

Contributor	Company Affiliation
Pradeep Anand	Time Warner Cable
Antonio Bonici	Videotron
L. Alberto Campos	Cable Television Laboratories, Inc.
Eduardo Cardona	Cable Television Laboratories, Inc.
Phillip Chang	Comcast Corp.
Bruce Currivan	Broadcom Corporation
Kirk Erichsen	Time Warner Cable
Jeff Finkelstein	Cox Communications
Alex Garcia	Videotron
Mark Geiger	Cox Communications
Jerry Green	Sunrise Telecom
Roger Fish	Broadcom Corporation
Ron Hranac	Cisco Systems
Bill Kostka	Cable Television Laboratories, Inc.
Jim Liu	Rogers
Ken Martin	Buckeye CableSystem
James Medlock	Ponderosa Systems
Walter Miller	JDSU
Charles Moore	Motorola, Inc.
Jack Moran	Motorola, Inc.
Christopher Reyes	Cox Communications
Rob Richards	VeEX, Inc.
Jason Sailor	Cox Communications
Steve Saunders	Sentosa Technologies
Rohini Telukutla	Arris International
Robert Thompson	Motorola, Inc.
Dustin Tracy	Cable Television Laboratories, Inc.
Huy Tran	Cox Communications
Michael Truong	Rogers
Tom Williams	Cable Television Laboratories, Inc.
Larry Wolcott	Comcast Corp.

We would particularly like to thank Robert Thompson of Motorola and Ron Hranac of Cisco for their commitment and leadership as the Proactive Network Maintenance Working Group editors, Alberto Campos of CableLabs for all his efforts as Working Group Lead, Eduardo Cardona of CableLabs for leading the reference implementation development effort, and all other participants who provided input to the working group.

Our sincere appreciation goes to operator members who supported this effort providing feedback and facilitating field data; in particular Nick Segura from Charter, Chris Nelson from Time Warner Cable, Wil Colon from Comcast, Joe Jensen from Buckeye, Mike Giobbi from Armstrong, Bob Cruickshank from Cablevision and George Hart from Rogers. Special thanks to Todd Szuter and Scott Johnston from Comcast for field verification of impairment scenarios and to Anupama Purohit from CableLabs for laboratory verification of micro-reflection and group delay scenarios.
