

Data Over Cable Service Interface Specification Proactive Network Maintenance

Primer for PNM Best Practices in HFC Networks (DOCSIS® 3.1)

CM-GL-PNM-3.1-V02-210114

RELEASED

Notice

This DOCSIS Proactive Network Maintenance guideline 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. You may download, copy, distribute, and reference the documents herein only for the purpose of developing products or services in accordance with such documents, and educational use. Except as granted by CableLabs® in a separate written license agreement, no license is granted to modify the documents herein (except via the Engineering Change process), or to use, copy, modify or distribute the documents for any other purpose.

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. To the extent this document contains or refers to documents of third parties, you agree to abide by the terms of any licenses associated with such third-party documents, including open source licenses, if any.

© Cable Television Laboratories, Inc. 2020-2021

DISCLAIMER

This document is furnished on an "AS IS" basis and neither CableLabs nor its members provide 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. Any use or reliance on the information or opinion in this document is at the risk of the user, and CableLabs shall not be liable for any damage or injury incurred by any person arising out of the completeness, accuracy, infringement, or utility of any information or opinion contained in the document.

CableLabs reserves the right to revise this document for any reason including, but not limited to, changes in laws, regulations, or standards promulgated by various entities, technology advances, or changes in equipment design, manufacturing techniques, or operating procedures described, or referred to, herein.

This document is not to be construed to suggest that any company modify or change any of its products or procedures, nor does this document represent a commitment by CableLabs or any of its members to purchase any product whether or not it meets the characteristics described in the document. Unless granted in a separate written agreement from CableLabs, nothing contained herein shall be construed to confer any license or right to any intellectual property. This document is not to be construed as an endorsement of any product or company or as the adoption or promulgation of any guidelines, standards, or recommendations.

Document Status Sheet

Document Control Number:	CM-GL-PNM-3.1-V02-210114			
Document Title:	Primer for PNM Best Practices in HFC Networks (DOCSIS® 3.1)			
Revision History:	V01 – Released 05/06/2020 V02 – Released 01/14/2021			
Date:	January 14, 2021			
Status:	Work in Progress	Draft	Released	Closed
Distribution Restrictions:	Author Only	GL/Member	GL/ Member/ Vendor	Public

Trademarks

CableLabs® is a registered trademark of Cable Television Laboratories, Inc. Other CableLabs marks are listed at <http://www.cablelabs.com/specs/certification/trademarks>. All other marks are the property of their respective owners.

Contents

1	SCOPE	13
1.1	Introduction and Purpose.....	13
1.2	PNM Background and History.....	13
1.3	Proactive Network Maintenance Informed by Cable Plant Measurements.....	16
1.3.1	<i>Plant Problems and Their Causes</i>	16
1.3.2	<i>Measurable Impact, Impact of Measurements</i>	17
2	REFERENCES	18
2.1	Informative References.....	18
2.2	Reference Acquisition.....	19
3	TERMS AND DEFINITIONS	20
4	ABBREVIATIONS	23
5	DOWNSTREAM PNM MEASUREMENTS	26
5.1	Downstream Transmit and Received Power.....	26
5.1.1	<i>Overview</i>	26
5.1.2	<i>Background</i>	26
5.1.3	<i>OFDM Power at the CMTS</i>	26
5.1.4	<i>Setting OFDM Channel Power at the CMTS</i>	27
5.1.5	<i>RF Input to the Cable Modem</i>	27
5.1.6	<i>Downstream Power Measurement on a Spectrum Analyzer</i>	27
5.1.7	<i>Downstream Power Measurement on Other Instruments</i>	28
5.1.8	<i>Impact of Pilots on OFDM Power</i>	29
5.1.9	<i>How to Implement</i>	29
5.1.10	<i>Conclusions and Future Work</i>	29
5.2	Downstream Wideband Spectrum Analysis.....	30
5.2.1	<i>Overview</i>	30
5.2.2	<i>Description of Measurement</i>	30
5.2.3	<i>Capabilities and Limitations</i>	30
5.2.4	<i>Averaging</i>	32
5.2.5	<i>Channel Power Accuracy and Calibration</i>	33
5.2.6	<i>Spectral Impairment Detection</i>	35
5.2.7	<i>Background Theory</i>	54
5.2.8	<i>How to Implement</i>	54
5.2.9	<i>Conclusion and Suggestions for Future Work</i>	57
5.2.10	<i>Detailed Discussion of Wideband Spectrum Analysis</i>	58
5.3	Downstream Receive Modulation Error Ratio (RxMER).....	62
5.3.1	<i>Overview</i>	62
5.3.2	<i>Description of Measurement</i>	64
5.3.3	<i>Background Theory</i>	65
5.3.4	<i>Data Presentation and Interpretation</i>	66
5.3.5	<i>Summary Metrics</i>	73
5.3.6	<i>How to Implement</i>	75
5.3.7	<i>Conclusion</i>	75
5.4	Downstream FEC Statistics.....	76
5.4.1	<i>Overview</i>	76
5.4.2	<i>Description of Measurement</i>	76
5.4.3	<i>Background Theory</i>	77
5.4.4	<i>Content Sections</i>	77
5.4.5	<i>Why High Correctables Are Acceptable in Downstream OFDM</i>	77
5.4.6	<i>Conclusion and Suggestions for Future Work</i>	81
5.5	Downstream Channel Estimate Coefficients.....	82

5.5.1	Overview.....	82
5.5.2	Description of Measurement.....	82
5.5.3	Background Theory.....	82
5.5.4	Relationship of Channel Estimate Coefficients to RxMER per Subcarrier.....	83
5.5.5	Examples.....	83
5.5.6	Tilt and Ripple Summary Metrics for Magnitude and Group Delay.....	85
5.5.7	Equations for Summary Metrics.....	86
5.5.8	Conclusion and Suggestions for Future Work.....	89
5.6	Downstream NPR Measurement.....	89
5.6.1	Overview.....	89
5.6.2	Description of Measurement.....	89
5.6.3	Background Theory.....	90
5.6.4	Measurement Description.....	91
5.6.5	Conclusion and Suggestions for Future Work.....	94
5.7	Downstream Constellation Display.....	95
5.7.1	Overview.....	95
5.7.2	Description of Measurement.....	95
5.7.3	Background Theory.....	95
5.7.4	Impaired Constellation Pattern Identification.....	96
5.7.5	Conclusion and Suggestions for Future Work.....	96
5.8	Downstream Symbol Capture.....	97
5.8.1	Overview.....	97
5.8.2	Description of Measurement.....	97
5.8.3	Summary of Uses of Downstream Symbol Capture.....	102
5.8.4	Conclusion and Suggestions for Future Work.....	102
5.9	Downstream Histogram.....	102
5.9.1	Overview.....	102
5.9.2	Description of Measurement.....	102
5.9.3	Background Theory.....	103
5.9.4	Examples of Histograms.....	103
5.9.5	Conclusion and Suggestions for Future Work.....	110
5.10	Physical Layer Link Channel (PLC) and Next Codeword Pointer (NCP).....	110
5.10.1	PLC Overview.....	110
5.10.2	A Closer Look.....	110
5.10.3	Useful PLC Parameters.....	111
5.10.4	NCP Overview.....	112
5.10.5	Useful NCP Parameters.....	112
5.10.6	How to Implement.....	112
5.10.7	Conclusion.....	112
5.11	MAC CRC Failures.....	112
5.12	SNR Margin for Candidate Profile.....	113
6	UPSTREAM PNM MEASUREMENTS.....	115
6.1	Upstream Power Levels.....	115
6.1.1	Overview.....	115
6.1.2	Description of Measurement.....	115
6.1.3	How to Implement.....	116
6.2	Upstream Pre-Equalization Coefficients.....	117
6.2.1	Overview.....	117
6.2.2	Description of Measurement.....	119
6.2.3	Background Theory.....	120
6.2.4	How to Implement.....	120
6.2.5	Conclusion and Suggestions for Future Work.....	120
6.3	Upstream OFDMA RxMER per Subcarrier.....	121
6.3.1	Description of Measurement.....	121
6.3.2	Conclusion and Suggestions for Future Work.....	123

6.4	Upstream Triggered Spectrum Capture	124
6.4.1	<i>Overview</i>	124
6.4.2	<i>How to Implement</i>	125
6.4.3	<i>Conclusion and Suggestions for Future Work</i>	125
6.5	Upstream Impulse Noise Statistics	125
6.5.1	<i>Overview</i>	125
6.5.2	<i>Description of Measurement</i>	125
6.5.3	<i>Conclusion and Suggestions for Future Work</i>	126
6.6	Upstream FEC Statistics	126
6.6.1	<i>Overview</i>	126
6.6.2	<i>Description of Measurement</i>	126
6.6.3	<i>Background Theory</i>	126
6.6.4	<i>Rolling FEC Summaries</i>	127
6.6.5	<i>Conclusion and Suggestions for Future Work</i>	127
6.7	Upstream Spectrum Capture at the CM	127
6.7.1	<i>Overview</i>	127
6.7.2	<i>Description of Measurement</i>	128
6.7.3	<i>Conclusion and Suggestions for Future Work</i>	128
6.8	Upstream Histogram	128
6.8.1	<i>Overview</i>	128
6.8.2	<i>Example</i>	128
6.8.3	<i>Conclusion and Suggestions for Future Work</i>	129
6.9	Upstream Capture for Active and Quiet Probes	129
6.9.1	<i>Overview</i>	129
6.9.2	<i>Description of Measurement</i>	129
6.9.3	<i>Background Theory</i>	130
6.9.4	<i>Conclusion and Suggestions for Future Work</i>	130
7	EXAMPLES OF USE CASES	131
7.1	Use Cases Versus Measurements	131
7.2	LTE and Other Interference	131
7.2.1	<i>Overview</i>	131
7.2.2	<i>Signal Leakage</i>	132
7.2.3	<i>Ingress and Direct Pickup</i>	132
7.2.4	<i>How to Implement</i>	134
7.2.5	<i>Conclusion and Suggestions for Future Work</i>	134
7.3	Fault Localization Based on Downstream RxMER	134
7.3.1	<i>Fault Localization Examples</i>	134
8	PNM DATA ISSUES	137
8.1	Data Collection Methods	137
8.1.1	<i>Operational Considerations</i>	138
8.1.2	<i>Security and Reliability Considerations</i>	139
8.1.3	<i>Deployment Considerations</i>	139
8.1.4	<i>IPDR Bulk Data collection</i>	139
8.2	Measurement Calibration	139
9	PNM PRACTICES	140
9.1	Measures of Access Network Health	140
9.1.1	<i>Quality of downstream plant (for each OFDM signal)</i>	140
9.1.2	<i>Quality of downstream service with respect to downstream plant</i>	142
9.1.3	<i>Consideration of FEC for Upstream health</i>	142
9.1.4	<i>About Averages</i>	143
9.1.5	<i>Sampling Frequency for Intermittent Problems</i>	143
9.1.6	<i>Stability</i>	144
9.1.7	<i>Impairments</i>	144

9.1.8	Conclusions.....	144
9.2	Network Maintenance Decision Process.....	145
9.3	FEC, MER and Decoding in the Downstream.....	146
APPENDIX I ERROR CORRECTION MARGIN SOLUTION ANALYSIS.....		151
I.1	Introduction	151
I.2	Error Correction Margin Test Plan	151
I.2.1	Test Set A	151
I.2.2	Test Set B	152
I.2.3	Test Set C.....	152
I.2.4	Figures for Test Configurations	152
I.3	Relating FEC to RxMER in OFDM for avoiding uncorrectable codeword errors.....	153
I.3.1	AWGN flat noise only	153
I.3.2	LTE Ingress Results with Correctable and Uncorrectable Codeword Errors (Whale Plots).....	158
I.3.3	Testing with raised noise floor, and simulated LTE ingress together.....	163
I.4	Conclusions	166
APPENDIX II FUTURE TECHNOLOGY		167
II.1	Technology Evolution	167
II.1.1	Technology Milestones	167
II.2	What's Next.....	169
II.2.1	Full Duplex DOCSIS Technology.....	169
II.2.2	FDX Amplifier	169
II.2.3	Extended Spectrum DOCSIS Technology	169
II.2.4	10G Initiative.....	170
II.2.5	Remote PHY Tap.....	170
II.2.6	Wireless PNM.....	170
II.2.7	Optical PNM.....	170
II.2.8	Coherent Optics.....	170
II.2.9	Fiber to the Home (FTTH).....	171
II.2.10	Conclusion and Suggestions for Future Work	171
APPENDIX III FURTHER READING		172
III.1	CableLabs DOCSIS 3.1 Specifications.....	172
III.2	CableLabs Remote PHY Specifications	172
III.3	SCTE•ISBE Cable-Tec Expo Papers	172
III.4	Articles.....	173
III.5	Training Videos and Programs	174
III.6	Other Useful References.....	174
APPENDIX IV ACKNOWLEDGEMENTS		175

Figures

Figure 1 - The Cable Plant Degradation Model that PNM Addresses.....	16
Figure 2 - Example of Same PSD Used for OFDM Signal and SC-QAM Signals.....	27
Figure 3 - Downstream Channel Lineup with a 96 MHz-Wide OFDM Signal at the Upper End of the Spectrum.....	28
Figure 4 - Power Measurements for a 96 MHz-Wide OFDM Channel as 16 Occupied Channels in 6-MHz Segments and for the Lower and Upper Adjacent SC-QAM Channel at Left and Right Edges	29
Figure 5 - Spectrum Analyzer Display of an OFDM Signal Showing the Boosted Continuous Pilots	29
Figure 6 - Digital Spectrum Analyzer Block Diagram	30
Figure 7 - Typical Downstream Wideband Spectrum Analysis of RF Spectrum	31
Figure 8 - Downstream Wideband Spectrum Analysis of RF Spectrum Including Interference	31

Figure 9 - Dynamic Range Limitations of Downstream Wideband Spectrum Analysis	32
Figure 10 - Upstream Spectrum and Diplexer Filter Rejection	32
Figure 11 - Maximum On-Chip Averaging Achieved with Different Implementations.....	33
Figure 12 - Power Level Comparison of Field Gear and Wideband Spectrum Analysis after Calibration	35
Figure 13 - Examples of Typical RF Spectral Impairments	35
Figure 14 - Example of a Simplified Three-Input Combiner.....	37
Figure 15 - DOCSIS Single Channel Spectral Mask	37
Figure 16 - Effect of Channel Misalignment on Other Channels	38
Figure 17 - Spectral Input with Channels Set to the Same Power Levels.....	38
Figure 18 - Spectral Input with Channels Set with Misaligned Power Levels.....	38
Figure 19 - Spectral Input Misalignment Caused by Plant Impairments (Not by Power Level)	38
Figure 20 - Adjacent Misalignment in SC-QAM Channels in a CPE.....	39
Figure 21 - Adjacent Misalignment in Multiple CPEs with the Same Node and/or Transmitter	39
Figure 22 - Filter Types: (a) Band-Pass Filter, (b) Band-Stop Filter, (c) Low-Pass Filter, (d) High-Pass Filter	41
Figure 23 - Example Data for Band-Stop and High-Pass Filters	41
Figure 24 - Typical Spectrum Notch	42
Figure 25 - Wide Spectrum Notch, Typical Depth	43
Figure 26 - Shallow Spectrum Notch, Typical Width	43
Figure 27 - Suckout as Demonstrated by PNM	43
Figure 28 - Expanded View of Suckout Notch and Impact on Several Digital Channels.....	44
Figure 29 - Examples of Ingress.....	47
Figure 30 - Example of a Significant Ingress Event with Several Types of Ingress.....	47
Figure 31 - Example of Resonant Peaking Affecting Adjacent Channels	48
Figure 32 - Example of Resonant Peaking with High Change in Amplitude	49
Figure 33 - Example of Resonant Peaking that Exceeds the Measurement Window	49
Figure 34 - Additional Alignment Issues in the Spectrum from Figure 33.....	49
Figure 35 - Examples of Spectral Roll Off at the High End of the Spectrum.....	50
Figure 36 - Example of Standing Waves	51
Figure 37 - Examples of Tilt: Negative Only, Positive Only, and Both Negative and Positive	52
Figure 38 - Examples of Leveling	53
Figure 39 - Example of Leveling Impairment: Increased Signal Level at CPE.....	53
Figure 40 - Example of Possible Leveling Across Several CPEs.....	54
Figure 41 - Block Diagram of a Digital Spectrum Analyzer in a Cable Modem.....	54
Figure 42 - Typical Window Functions.....	58
Figure 43 - Ripples Indicating an Echo Tunnel; No Phase Data Available.....	59
Figure 44 - Impulse Associated with Frequency Domain Ripple Among Comb Teeth (at Intervals of 166.67 ns) ...	60
Figure 45 - Diagram Showing Detection of a Single Reflection	60
Figure 46 - Digital Cable Signal Captured by Rapidly Retuning an SDR; Standing Wave Indicates a Reflected Signal Is Present	61
Figure 47 - Processed Signal Showing a Single Reflection, with Harmonics Caused by Roll Off of the 6-MHz Haystacks at Band Edges.....	61
Figure 48 - Cablelabs Engineer Making a TDR Measurement in the Field (the SDR Is in His Backpack)	61
Figure 49 - List of RxMER-per-Subcarrier Data from a DOCSIS 3.1 Cable Modem.....	63
Figure 50 - Plot of RxMER-per-Subcarrier Data from a Field Meter with a DOCSIS 3.1 Cable Modem; Downstream OFDM Signal Is 96 MHz Wide and Average RxMER per Subcarrier Is 41.2 dB	64

Figure 51 - Using Pilot Symbols to Compute RxMER.....	64
Figure 52 - Computation of RxMER for a Given Subcarrier	65
Figure 53 - Plot of RxMER per Subcarrier Showing Impact of Simulated LTE Ingress on Affected Subcarriers.....	66
Figure 54 - Inverted Plot of RxMER per Subcarrier with LTE Ingress.....	67
Figure 55 - Graph of OFDM RxMER per Subcarrier Showing Ingress Interference at About 890 MHz.....	67
Figure 56 - Plot of Actual RxMER per Subcarrier Showing Examples of Various LTE Interference and an Exclusion Band at About 760–763 MHz.....	68
Figure 57 - Spectrum Plot with Ripple, Case 1, Noise Added Before Micro-Reflection	69
Figure 58 - Channel Estimate Plot with Ripple, Case 1, Noise Added Before Micro-Reflection	69
Figure 59 - RxMER-per-Subcarrier Plot with Little to No Ripple, Case 1, Noise Added Before Micro-Reflection ..	70
Figure 60 - Spectrum Plot with Ripple, Case 2, Noise Added After Micro-Reflection.....	70
Figure 61 - Channel Estimate Plot with Ripple, Case 2, Noise Added After Micro-Reflection.....	71
Figure 62 - RxMER-per-Subcarrier Plot with Ripple, Case 2, Noise Added After Micro-Reflection	71
Figure 63 - Example of Edge Roll Off in RxMER per Subcarrier.....	73
Figure 64 - RxMER Percentile Comparisons	75
Figure 65 - RxMER 2 nd Percentile Versus RxMER Standard Deviation for a Population of Cable Modems.....	75
Figure 66 - Screen Capture from Field Meter Showing Performance of a DOCSIS 3.1 OFDM Signal.....	77
Figure 67 - Comparison of Uncorrectable Codeword Percentage for SC-QAM and OFDM, Measured	79
Figure 68 - Measured Performance of SC-QAM Downstream 256-QAM.....	80
Figure 69 - Measured Performance of OFDM 256-QAM.....	81
Figure 70 - Measured Performance of OFDM 4K-QAM	81
Figure 71 - Use of Downstream Channel Estimate Coefficients	83
Figure 72 - Channel Estimate Coefficients for 70-dB Notch at Edge of OFDM Channel.....	84
Figure 73 - Channel Estimate Coefficients for 500-ns Echo at -10 dB.....	84
Figure 74 - Amplitude Response of Channel Estimate Coefficients for 500-ns Echo at -10 dB	84
Figure 75 - Residual Group Delay of Channel Estimate Coefficients for 500-ns Echo at -10 dB.....	85
Figure 76 - Channel Estimate Expressed as Time-Domain Taps for 500-ns Echo at -10 dB	85
Figure 77 - Frequency Response of Channel Estimate Coefficients, with Best-Fit Line in Red	87
Figure 78 - Group Delay Response of Channel Estimate Coefficients, with Best-Fit Line in Red	88
Figure 79 - Use of Exclusion Band to Support NPR Measurement in OFDM Signal	89
Figure 80 - Use of Exclusion Band to See Ingress Interference Under an OFDM Signal.....	89
Figure 81 - Band-Limited Noise Added to the Input of a Device or System Under Test (Band of Interest Is 5 MHz to 42 MHz)	90
Figure 82 - Notch Filter Added to Test Setup to Create a Notch in Band-Limited Noise Prior to the Device or System Under Test.....	90
Figure 83 - NPR (Depth of the Notch in Decibels) Measured by Spectrum Analyzer at Output of the Device or System Under Test.....	91
Figure 84 - Example of NPR Graph	91
Figure 85 - Exclusion Band Width Including (B) Taper Regions and (C) NPR Measurement Region.....	92
Figure 86 - OFDM Signal, 192 MHz, Before Addition of Exclusion Band (Note PLC Near Left Edge of Signal)....	93
Figure 87 - OFDM Signal from Figure 86 with 4 MHz-Wide Exclusion Band	93
Figure 88 - Close-up View of Exclusion Band in Figure 87; Combined Width of Each Taper Region (0.975 MHz) and Exclusion Band (4 MHz) Is Sufficient to Accommodate an NPR Measurement	94
Figure 89 - Example Exclusion Band in which Combined Width of Each Taper Region (1.875 MHz) and Exclusion Band (4 MHz) Is Not Adequate for a Reliable NPR Measurement.....	94
Figure 90 - Downstream OFDM Constellation	96

Figure 91 - Nonlinear SC-QAM 16-QAM Constellation Showing Evidence of Compression or Clipping	96
Figure 92 - Constellation (Frequency-Domain QAM Values) of Downstream OFDM Symbol Captured at the Transmitter	98
Figure 93 - Real Part of Frequency-Domain QAM Values of Transmitted OFDM Symbol.....	98
Figure 94 - Real Part of Time-Domain Samples of OFDM Symbol Received at CM.....	99
Figure 95 - Soft Decisions (Time-Domain Samples) of OFDM Symbol Received at CM.....	99
Figure 96 - FFT of Soft Decisions (Resulting Frequency-Domain Samples) of OFDM Symbol Received at CM...	100
Figure 97 - Frequency Response of Downstream Channel Computed from a Single Captured Downstream OFDM Symbol.....	101
Figure 98 - Computed Phase Response of Downstream Channel.....	101
Figure 99 - Computed Unwrapped Phase Response of Downstream Channel	101
Figure 100 - Normal OFDM Signal with Gaussian-Shaped Histogram; (a) Time Domain, (b) Histogram	104
Figure 101 - Laser Clipping in Signal with One Tail of Histogram Replaced with Spike; (a) Time Domain, (b) Histogram	104
Figure 102 - Histogram of Actual Upstream Clipping Event	105
Figure 103 - Histogram of Undistorted OFDM and SC-QAM Downstream Signal.....	105
Figure 104 - Log Histogram of Undistorted Downstream Signal.....	106
Figure 105 - Delta from Gaussian (in Decibels) for Histogram of Undistorted Downstream Signal	106
Figure 106 - Histogram of Downstream RFoG Signal with Severe Laser Distortion (190 MHz, 16-QAM, 3,800 Subcarriers).....	107
Figure 107 - Log Histogram of Downstream RFoG Signal with Laser Distortion.....	107
Figure 108 - Delta from Gaussian (in Decibels) for Histogram of Downstream RFoG Signal with Severe Laser Distortion.....	107
Figure 109 - Spectrum Analyzer Display of Downstream Signal with Low Laser Distortion, OFDM RxMER = 37 dB	108
Figure 110 - Spectrum Analyzer Display of Downstream Signal with High Laser Distortion, OFDM RxMER = 23 dB	108
Figure 111 - Log Histogram of Downstream Signal with Low Distortion, OFDM RxMER = 37 dB.....	109
Figure 112 - Delta from Gaussian of Downstream Signal with Low Distortion, OFDM RxMER = 37 dB.....	109
Figure 113 - Log Histogram of Downstream Signal with High Distortion, OFDM RxMER = 23 dB.....	109
Figure 114 - Delta from Gaussian of Downstream Signal with High Distortion, OFDM RxMER = 23 dB	110
Figure 115 - Example Spectrum Analyzer Display Showing the PLC Band Within a 192 MHz-Wide OFDM Signal	111
Figure 116 - Example Spectrum Analyzer Display Showing PLC Placement in the PLC Band, with a Close-Up of the PLC Band (Yellow), PLC (Red), and Boosted Continuous Pilots in a Unique Pilot Pattern.....	111
Figure 117 - Examples of Frequency Response for Upstream OFDMA Signals	119
Figure 118 - Upstream RxMER-per-Subcarrier Data Plotted by Frequency for a Single CM at Four Different Times	123
Figure 119 - Upstream RxMER-per-Subcarrier Data Plotted in 3 Dimensions by Frequency for a Single CM at Four Different Times	123
Figure 120 - Upstream Spectrum Capture of Multiple Upstream SC-QAM Channels.....	124
Figure 121 - Upstream Spectrum Capture of Quiet Time.....	125
Figure 122 - CM Upstream Capture with Diplex Rejection	127
Figure 123 - CM Upstream Capture with Diplex Rejection and Noise	128
Figure 124 - Histogram of Actual Upstream Clipping Event.....	129
Figure 125 - Use Cases Addressable by PNM Measurements.....	131

Figure 126 - Cable Network Signal Leakage in the 700 MHz LTE Spectrum, Overlapping LTE Band 13 (Verizon) Tower-to-UE Downlink Signals	132
Figure 127 - Full-Band Capture Spectrum from a Modem Showing Several Types of Ingress	133
Figure 128 - OFDM RxMER-per-Subcarrier Plot Showing the Impact of LTE and Other Types of Ingress	133
Figure 129 - "Ingress Under the Carrier" Display from a Field Meter	134
Figure 130 - Example of Downstream RxMER Data Showing Ripple Location	135
Figure 131 - Example of Downstream RxMER Data Showing LTE Ingress Location	136
Figure 132 - Data Collection Sequence Diagram	137
Figure 133 - Example of Downstream RxMER per Subcarrier	138
Figure 134 - TFTP Upload Diagram.....	138
Figure 135 - Test Setup for Test Sets A and B	152
Figure 136 - Test Setup for Test Set C	152
Figure 137 - CM and VSA Input Signal Power -10 dBmV	153
Figure 138 - CM and VSA Input Signal Power -6 dBmV	154
Figure 139 - CM and VSA Input Signal Power 0 dBmV	155
Figure 140 - CM and VSA Input Signal Power +10 dBmV	156
Figure 141 - Whale plots (Correctable and Uncorrectable Codeword Errors) of 0 dBmV Power Level, 4k-QAM, with Increasing Noise to Reduce Average RxMER	157
Figure 142 - Whale Plots of 0 dBmV Power Level, 4k-QAM, 96 MHz-wide Channel, 8 MHz-wide LTE Ingress, with Increasing LTE level to Reduce the RxMER in the LTE Band. MER (dB) in these Plots is a Notch MER (dB) Value	159
Figure 143 - Whale Plots of 0 dBmV Power Level, 4k-QAM, 192 MHz-wide Channel, 8 MHz-wide LTE Ingress, with Increasing LTE level to Reduce the Average RxMER (x-axis) in the LTE Band, Reporting RxMER over the LTE Band	160
Figure 144 - Whale Plots of 0 dBmV Power Level, 256-QAM, 192 MHz-wide Channel, 8 MHz-wide LTE Ingress, with Increasing LTE Level to Reduce the RxMER in the LTE Band, Reporting Average RxMER over the LTE Band on the x-axis	161
Figure 145 - Whale Plots of 0 dBmV Power Level, 4k-QAM, 192 MHz-wide Channel, 8 MHz-wide LTE Ingress, with Increasing LTE Level to Reduce the RxMER in the LTE Band, this Time Reporting Average RxMER over the Whole Channel	162
Figure 146 - Whale Plots of 0 dBmV Power Level, 256-QAM, 192 MHz-wide Channel, 8 MHz-wide LTE Ingress, with Increasing LTE Level to Reduce the RxMER in the LTE Band, this Time Reporting Average RxMER over the Whole Channel	162
Figure 147 - Individual CMs in each Plot, Showing Various Noise Floors, and Increasing the LTE Level to Lower RxMER Values in the LTE Frequency Range; Correctable Codeword Errors and Uncorrectable Codeword Errors Proportions versus RxMER in the LTE Range.....	164
Figure 148 - Individual CMs in each Plot.....	165
Figure 149 - Example of Subsplit Band Plan Used in North America and Other Regions.....	168
Figure 150 - Frequency Spectrum Usage in FDX DOCSIS Systems (Bottom) Compared to Traditional Subsplit, Midsplit, or High-Split Band Plans (Top)	169
Figure 151 - Extended Spectrum DOCSIS Frequency Usage With and Without FDX DOCSIS Technology.....	169

Tables

Table 1 - Example of Downstream DOCSIS Channel Center Frequency and Receive Power.....	34
Table 2 - Example of Wideband Spectrum Channel Center Frequency and Receive Power.....	34
Table 3 - CmSpectrumAnalysisCtrlCmd Object Attributes	55

Table 4 - Spectrum Analysis File Format.....	57
Table 5 - Recommended OFDM RxMER for Each Modulation Order.....	72
Table 6 - Reliability Threshold for Cable Modem RxMER at 2 nd Percentile.....	74
Table 7 - Taper Region Widths Based on Roll-Off Configurations (Nrp).....	92
Table 8 - CmDsOfdmRequiredQamMer Object from [CM-OSSIV3.1]	141
Table 9 - CMTS Minimum CNR Performance in AWGN Channel from [PHYv3.1].....	141
Table 10 - Calculate Sample Frequency	143
Table 11 - CmDsOfdmRequiredQamMer Object from [CM-OSSIV3.1]	149
Table 12 - CM Minimum CNR Performance in AWGN Channel from [PHYv3.1]	149

1 SCOPE

1.1 Introduction and Purpose

As cable networks evolve, and as many diverse services such as telephony, data, video, business, and advanced services (e.g., tele-medicine, remote education, home monitoring) are carried over those networks, the demand for maintaining a high level of reliability and availability for services increases. To achieve this high reliability and availability, operators have found it advisable to fix problems before they have any impact on service. DOCSIS technology affords operators the ability to be proactive and correct network impairments before service is impacted significantly.

Some commonly tracked cable modem (CM) and cable modem termination system (CMTS) metrics have been updated and included in the DOCSIS 3.1 specifications. These measurements include wideband spectrum analysis, noise power ratio, channel estimate coefficients, constellation display, receive modulation error ratio (RxMER) per subcarrier, forward error correction (FEC) statistics, histogram plots, received power, equalizer coefficients, impulse noise statistics, active and quiet probes, and channel power. These metrics are good indicators of the existence of problems and, in many cases, can help reveal the cause of those problems or at least suggest causes to investigate.

Increasingly, intelligent end devices are being deployed in cable networks, and termination devices and monitoring instruments are being installed in headends (HEs) and hubs. Devices 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 often compatible with DOCSIS technology, making that technology ubiquitous.

Embedded monitoring and Proactive Network Maintenance (PNM) capabilities, introduced in the DOCSIS 3.1 specifications, allow cable networks to be maintained proactively to meet high-quality service needs. As DOCSIS devices evolve and become equipped with elaborate monitoring tools, it becomes practical to use them to monitor cable plants. By using these devices as network probes, cable operators can collect device and network parameters, and by combining the analysis of the data with network topology and device location, it is possible to isolate the source of a problem. This information can then be used to develop a proactive maintenance plan.

This document provides guidelines and best practices for proactive network maintenance mechanisms that exploit DOCSIS 3.1 measurements. The processes described herein will help cable operators and industry vendors implement smart monitoring tools, improve maintenance practices, gain better insight in network problems, and enhance network reliability. The intent is to update this document with the latest methods as these capabilities improve and become more available.

This effort has two key outcomes:

1. reducing the time required to troubleshoot and resolve problems, thereby reducing operational costs, and
2. enabling early warning of impending significant problems before they impact service, thereby enabling operations personnel to correct problems before service is impacted.

In addition, improvements in network reliability enable the introduction of business and advanced services that require SLAs (service level agreements), which can generate new revenue. These PNM mechanisms add the capability to detect and resolve problems before they impact customer service, which helps with churn reduction.

This guideline for PNM focuses on DOCSIS 3.1 technology, but it occasionally references the previous version, which focused on DOCSIS 3.0 techniques [PNMP-3.0]. This version differs significantly from the previous version because DOCSIS 3.1 technology introduced many new capabilities, including orthogonal frequency-division multiple/access (OFDM/A). The previous version stands on its own for DOCSIS 3.0 techniques and serves as reference for this and any future versions of the PNMP guidelines document.

1.2 PNM Background and History

The traditional DOCSIS service indicators that operators relied upon for years were CM status, upstream transmit level, upstream receive level, upstream signal-to-noise ratio (SNR) (RxMER), upstream codeword error rate, downstream receive level, downstream SNR (RxMER), and downstream codeword error rate. These indicators, though valuable, could not answer the questions of what exactly is wrong or what is the root cause of poor service. When PNM added upstream equalization and downstream full-band capture data, the data were able to speak much

more clearly as to root causes. Now, with DOCSIS 3.1 technology, the PNM tools are being expanded further to give greater precision. This document addresses the interpretation skill needed to fully utilize those tools.

Most radio frequency (RF) communication systems, including DOCSIS systems, use a variety of techniques to adjust and compensate for linear distortion and variations in timing, frequency, and receive power level. The industry has known for quite some time that useful plant health information can be derived from the parameters that describe the compensation and adjustments that take place. This network monitoring advantage has been enhanced by the ubiquity of DOCSIS devices operating as network probes and sensors across the entire HFC (hybrid fiber-coax) footprint.

Among the most important device parameters that provide insight into the characteristics of the DOCSIS upstream transmission channel are the pre-equalization coefficients. Pre-equalization techniques are used to compensate for linear distortion in the upstream channel, such as micro-reflections, amplitude ripple, tilt, and group delay distortion. In most cases, pre-equalization can completely compensate for linear distortion problems without having the customer perceive an impact on performance. If the operator can monitor and detect a change in the equalizer coefficients, then the operator often has time to fix the problem before any service degradation has taken place, thereby facilitating a proactive network maintenance strategy.

Although pre-equalization has been mandatory in DOCSIS systems since the DOCSIS 1.1 specifications, data on pre-equalization have only been exploited in PNM since around 2005, midway through the deployment of DOCSIS 2.0-compliant devices. DOCSIS systems have a variety of channel-width and modulation-order configuration options. Operators' early use of narrow bandwidth channels (predominantly 1.6 MHz) and low-order modulations (e.g., QPSK) often did not require the use of pre-equalization. Moreover, some operators were reluctant to turn on pre-equalization because certain DOCSIS 1.0 CMs had misbehaved when pre-equalization was turned on. However, the need to turn on pre-equalization became apparent in scenarios that needed increased upstream peak rates and transport robustness, which resulted in the migration to 3.2 MHz and 6.4 MHz bandwidth channels and the use of higher order modulations such as 64-QAM. These wider bandwidths and higher modulation orders require powerful pre-equalization.

Pre-equalization in DOCSIS 1.0 specifications was optional and was left unspecified. In DOCSIS 1.1 specifications, pre-equalization not only was defined accurately but also was mandatory for channels up to 3.2 MHz. The mandatory nature of pre-equalization in DOCSIS 1.1 specifications and the proper identification and isolation of misbehaving CMs enabled cable operators to turn on pre-equalization.

A couple of years later, after the DOCSIS 2.0 specifications were released, pre-equalization MIBs (management information bases) became available. The MIBs made it possible to use pre-equalization information for network management purposes.

Earlier work by Holtzman, Inc., Motorola, and CableLabs had highlighted the value of understanding linear distortions, additive impairments, and their impact on service performance. In 2005, work began on equalization coefficient decoding and normalization with a study by Cablelabs and Charter Communications in Estes Park, Colorado. They collected pre-equalization data from multiple nodes, and though only a portion of the data was easily readable, subsets of that data exhibited an apparent correlation. Cable operators could now begin to make sense of distortion signatures and how they relate to problems in the field.

The use of DOCSIS pre-equalization coefficients with plant topology information to pinpoint problems in the network were described in a series of CableLabs internal reports in 2006 and at the SCTE Cable-Tec Expo 2008 in Philadelphia. This proposed approach relied on the following steps:

1. derive frequency response signatures from pre-equalization coefficients,
2. identify and group linear distortions from the frequency response data, and
3. correlate impacted CMs with topology information to locate the cause of the significant problems detected in step 2.

In 2009, CableLabs sponsored the creation of a Proactive Network Maintenance (PNM) working group to leverage information obtained from DOCSIS devices and to troubleshoot the plant. This group comprised participants from cable operators, CableLabs, and silicon, CM, CMTS, and instrumentation vendors. Tasks of the PNM working group included the development of techniques to relate pre-equalization signatures to problems in the field. Participation in the PNM working group by the North American cable operator Comcast not only facilitated the use of large amounts of field data to understand pre-equalization coefficient information, but also provided valuable field information on specific associations between impairments and distortion signatures. One key output of the PNM working group was

the best practices and guidelines document "Proactive Network Maintenance Using Pre-Equalization" (first released version of [PNMP-3.0]), published in 2010.

The original PNM working group output also led operators to implement PNM-based tools. The first implementation of an operator-developed pre-equalization analysis tool was the Scout Flux tool from Comcast, implemented and released in 2009. Comcast's technical workforce played an integral role in determining a variety of plant impairment scenarios and their signatures. Tools from other cable operators followed, including Charter's Node Slayer PNM tool.

The PNM working group initially focused on upstream pre-equalization, but other topics were also discussed. One such topic was DOCSIS 3.0 downstream equalization, which still suffers from lack of compliance and discrepancies in MIB interpretation. Upstream spectrum analysis was also discussed; DOCSIS 3.0 specifications required it for a single channel, but it was supported in a proprietary fashion by numerous CMTS vendors across the full upstream spectrum. Cisco Systems lead efforts to assess the impact that LTE ingress has on performance. The introduction in the DOCSIS 3.0 specifications of multiple bonded channels resulted in high sampling rate receiver implementations. Industry leaders use this feature to leverage CPE spectrum capture, or full-band capture (FBC)¹, and Comcast introduced it into their network management systems shortly after its availability.

The DOCSIS 3.1 specifications introduced new modulation formats: orthogonal frequency-division multiplexing (OFDM) in the downstream and orthogonal frequency-division multiple access (OFDMA) in the upstream. New PNM tools have been developed to exploit the information in these new signals. For example, instead of a single receive modulation error ratio (RxMER) measurement for a downstream channel or a small number of RxMER measurements over a bonded channel group, there are now up to 4,000 or 8,000 RxMER values, one for each OFDM subcarrier. These measurements provide a detailed plot of the signal-to-noise ratio (SNR) of the channel (or, in the inverse, the noise underlying the channel) as a function of frequency across the whole band, with 25 kHz or 50 kHz resolution.²

The development of FBC gave operators significant operational advantages and resulted in operational cost savings. It enabled operators to have spectrum analysis capabilities anywhere modern DOCSIS 3.0/3.1 CMs are deployed, allowing them to take remote spectrum captures at the customer premises without having to carry an expensive spectrum analyzer or requiring access into a subscriber's home. As deployment of DOCSIS 3.0/3.1 CMs increased, these spectrum analysis probes gathered critical mass. The PNM system can correlate spectrum signatures from neighboring CMs to troubleshoot and locate problems. The width of the spectrum (full downstream band) has allowed operators to detect very short micro-reflections not visible when a single channel is analyzed through equalization. FBC also allows operators to verify channel level alignment and detect ingress. Correlation of CM signatures before and after actives enables the detection of nonlinear problems at amplifiers. Automated spectrum signature analysis and impairment detection are being implemented to scale the analysis to the millions of DOCSIS 3.0/3.1 CMs deployed with this functionality.

An extended capability of FBC is upstream spectrum capture at the CM. Although the lower frequency upstream signals are attenuated at the downstream receiver port by the diplex filter, in many cases, enough energy passes through to allow for the detection of impulse noise at the customer premises. This is a very promising technique that could be enhanced through CM design, including, for example, diplex filter bypass.

Leveraging field data collected at Comcast, Armstrong Cable, and Suddenlink Communications, CableLabs demonstrated the correlation between equalization coefficient variability and RxMER variation with impulse noise. This finding opened the door to solving ingress localization problems for certain types of noise. Ingress localization is one of the remaining challenges to conquer in HFC plant troubleshooting.

The PNM working group has been crucial to the incorporation of PNM tools into the DOCSIS 3.1 specifications. Under the leadership of Broadcom, hooks were incorporated in the DOCSIS 3.1 specifications that allow systems to emulate tools such as spectrum analyzers, vector network analyzers, and vector signal analyzers. These specifications provide a set of nearly 20 PNM measurements. As of this writing, the CM-based measurements have been implemented, but the specified CMTS measurements have generally not been implemented, with the exception of upstream triggered spectrum analysis. The remaining PNM measurements will be the subject of a future update of this document when they become available.

¹ Generically, the industry refers to CPE spectrum capture as full-band capture (FBC). MaxLinear calls it Full Spectrum Capture™.

² Multiple OFDM channels expand the observable spectrum in SNR measurements accordingly.

With the deployment of significant numbers of DOCSIS 3.1 CMs with the confirmed capability to support the specified PNM measures, operators have numerous advantages at their disposal. This document outlines updated methods for making use of these advantages for PNM.

It has been eight years since the release of the initial PNM publication [PNMP-3.0]. This document will incorporate the relevant PNM topics from that publication and update them for use with DOCSIS 3.1 systems.

1.3 Proactive Network Maintenance Informed by Cable Plant Measurements

Plant issues can be organized into just a few categories: impedance mismatches, leaks, noise, and corrosion. These categories can overlap as leaks can corrode and be seen as impedance mismatches. Therefore, better organization of plant issues is needed. An initial thought is to separate these issues according to impact on signal or transmission, and therefore service, versus impact on plant condition.

The primary goal of PNM is to find and fix problems before customers are impacted by them. Reliable indications of plant issues that can lead to future service impact are valuable, but current service impact is more reliable. Fortunately, DOCSIS systems do an excellent job of separating signal impairments from service impact.

A closely associated capability from PNM is to increase DOCSIS 3.1 capacity through enabling higher modulation constellation orders. A CM has a limited amount of bandwidth that it can use to receive and transmit. Plant issues that reduce a CM's ability to communicate with the CMTS in either direction result in a loss of service quality. If those plant issues are severe enough, service is impacted. DOCSIS technology, however, adjusts to many of these problems by cancelling echoes, adjusting power levels, correcting errors, and adjusting profiles, for example. When delivering full-quality service, a CM uses all provisioned bandwidth to the highest modulation order possible. Any reduction in bandwidth use leads to a loss of quality, and severe loss of quality will be noticed by the customer. Note, however, that a less-than-perfect plant can still provide full-quality service and bandwidth, so there is opportunity to be very proactive.

From an analog perspective, loss of service quality, particularly any loss of transmission or signal quality, comes from a few types of problems. Setting aside design problems and sabotage, quality can be reduced by human error, shock damage, and degradation. These causes manifest in the physical plant as bad splices or connectors, exposed shields, broken conductors, etc. These physical plant problems are represented in measurements as increased noise, decreased signal, or both. The signal or noise issues can be across the spectrum, be limited to some frequencies, or appear as a repeating pattern in the frequency or time domains.

PNM methods were developed to find and fix these plant issues. These conditions are rare, so it is far more efficient to begin with remote data collection. Because the goal is to find and fix problems before customers experience service impact, data should be collected with minimal service interruption. Operators currently collect data broadly and remotely at a low net cost, then separate anomalies for focused investigation using PNM tools.

The rest of this section discusses the causal chain from root cause to service impact, as depicted in Figure 1.




Figure 1 - The Cable Plant Degradation Model that PNM Addresses

1.3.1 Plant Problems and Their Causes

Generally, plant problems are caused by design, human error, sabotage, shock damage, and degradation. Causes can be made more actionable by defining them specifically—for example, degradation can be defined as water intrusion, corrosion, cracking, or loosening—but the added context may be particular to a specific plant problem or location.

1.3.2 Measurable Impact, Impact of Measurements

When looking at the impact of plant problems on transmission or service quality, not all of them will have a measurable impact, at least initially. For those that do have a measurable impact, the measurement might not be severe enough to manifest in the transmission or service quality to a degree noticeable by the customer.

The following sections describe problems that do have measurable impacts.

1.3.2.1 Reactive Measurements

Reactive measurements of cable service that come too late to be proactive but may still be informative include T3 and T4 timeouts and sysUptime (system uptime). Customers are highly likely to notice issues reflected by these measurements.

1.3.2.2 Short-Term Proactive Measurements

Proactive measurements that are reliable indicators of signal transmission issues include

- power levels,
- uncorrectable errors,
- low RxMER in one or more subcarriers,
- spectrum that deviates from ideal,
- pre-equalizer coefficients that have high energy off of the main tap,
- channel estimation coefficients that vary widely by frequency, and
- a modulation profile indicating subcarriers that are not transmitting at their expectation or potential.

Measurements in this group are reliable because they clearly indicate a loss of transmission signal quality, which puts service and, if severe enough, customer experience at risk. It is therefore important to address this class of measurements before they become more severe, which can occur when new problems manifest in the plant, when conditions change such that margins are pushed further (heat, cold, power limitations, etc.), or when existing problems degrade further. Used across time, locations, or even across measurements, these measurements can be used to help locate impairments (first tap showing an echo cavity) or identify causes (random ripple patterns indicating water in the cable).

1.3.2.3 Troubleshooting Measurements

Some PNM measurements are expected to be most useful for troubleshooting or performing periodic maintenance. For example, consider downstream symbol capture, upstream triggered spectrum capture, and in-channel frequency response characterization. Measurements in this category are useful for locating impairments or identifying causes but are not typically automated for monitoring plant quality.

1.3.2.4 Other Proactive Measurements

The remaining PNM measurements or tests indicate other plant or design issues related to loss of quality but may be minor or long-term proactive indicators. In some cases, several of these issues may contribute to reductions in transmission or service quality. However, many of the measurements in this category have yet to be fully explored for their utility in PNM or reactive repair or operations maintenance. Measurements in this category include histogram plots and constellation display.

2 REFERENCES

2.1 Informative References

- [47 C.F.R. §76.605] 47 C.F.R. §76.605 (b)(4)(i), Multichannel Video and Cable Television Service—Technical Standards, Code of Federal Regulations.
- [CCAP-OSSiv3.1] DOCSIS CCAP Operations Support System Interface Specification, CM-SP-CCAP-OSSiv3.1-I17-191204, December 4, 2019, Cable Television Laboratories, Inc.
- [CM-OSSiv3.1] DOCSIS Cable Modem Operations Support System Interface Specification, CM-SP-CM-OSSiv3.1-I16-190917, September 17, 2019, Cable Television Laboratories, Inc.
- [CTA-542-D] CTA-542-D, "Cable Television Channel Identification Plan," Consumer Technology Association, 2013 (previously identified by CEA-542-D when the CTA was known as the Consumer Electronics Association).
- [Currivan and Wolcott] "Leveraging DOCSIS 3.1 PNM Measurements for Plant Performance Optimization," SCTE•ISBE Cable-Tec Expo, 2015—B. Currivan, L. Wolcott.
- [DOCS-IF3-MIB] CableLabs DOCSIS DOC-IF3-MIB SNMP MIB Module, DOCS-IF3-MIB, <http://www.cablelabs.com/MIBs/DOCSIS/>
- [DOCS-IF31-MIB] CableLabs DOCSIS DOC-IF31-MIB SNMP MIB Module, DOCS-IF31-MIB, <http://www.cablelabs.com/MIBs/DOCSIS/>
- [DOCS-PNM-MIB] CableLabs DOCSIS DOC-PNM-MIB SNMP MIB Module, DOCS-PNM-MIB, <http://www.cablelabs.com/MIBs/DOCSIS/>
- [Ferreira, Harb, et al] "Convolutional Neural Networks for Proactive Network Management," Jude Ferreira, Maher Harb, Karthik Subramanya, Bryan Santangelo, Dan Rice, Cable-Tec Expo 2020.
- [Hranac and Currivan] "Digital Transmission: Carrier-to-Noise Ratio, Signal-to-Noise Ratio, and Modulation Error Ratio," white paper, 2006, Cisco—R. Hranac, B. Currivan.
- [Hranac and others] "Characterizing Network Problems Using DOCSIS 3.1 OFDM RxMER per Subcarrier Data," SCTE•ISBE 2020 Fall Technical Forum—R. Hranac, J. Medlock, B. Currivan, R. Fish, T. Kolze, J. Rupe, T. Williams, L. Wolcott.
- [Hranac, et al.] "Full Band Capture Revisited," SCTE•ISBE 2020 Fall Technical Forum—Ron Hranac, Chad Campbell, Roger Fish, Tom Kolze, Even Kristoffersen, James Medlock, Jason Rupe, Paul Schauer, Aleksander Soeberg, Tom Williams, Larry Wolcott.
- [ITU-T J.83] ITU-T Recommendation J.83 (12/2007), Digital Multi-Programme Systems for Television Sound and Data Services for Cable Distribution.
- [MULPIv3.1] DOCSIS MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I20-200407, April 7, 2020, Cable Television Laboratories, Inc.
- [PHYv3.0] DOCSIS Physical Layer Specification, CM-SP-PHYv3.0-C01-171207, December 7, 2017, Cable Television Laboratories, Inc.
- [PHYv3.1] DOCSIS Physical Layer Specification, CM-SP-PHYv3.1-I17-190917, September 17, 2019, Cable Television Laboratories, Inc.
- [PHYv4.0] DOCSIS Physical Layer Specification, CM-SP-PHYv4.0-I02-200429, April 29, 2020, Cable Television Laboratories, Inc.
- [PMI PNM] Profile Management Informed Proactive Network Maintenance, SCTE•ISBE Cable-Tec Expo 2020 Fall Technical Forum.
- [PNMP-3.0] DOCSIS PNM Best Practices: HFC Networks (DOCSIS 3.0), CM-GL-PNMP-V03-160725, July 25, 2016, Cable Television Laboratories, Inc.
- [Prodan, et al] R. Prodan, L. Montreuil, A Kliger, B. Shen, "Demystifying the DOCSIS 3.1 PHY," SCTE•ISBE Cable-Tec EXPO, September 22-25, 2014.
- [R-UEPI] DOCSIS Remote Upstream External PHY Interface Specification, CM-SP-R-UEPI-I13-201207, December 07, 2020, Cable Televisions Laboratories, Inc.
- [Rupe, Zhu] Jason Rupe, Jingji Zhu, "Kickstarting proactive network maintenance with the proactive operations platform and example application", 2019 SCTE-ISBE and NCTA.
- [SCTE 119] ANSI/SCTE 119, Measurement Procedure for Noise Power Ratio, 2018.
- [SCTE 257] ANSI/SCTE 257, Techniques for Accurate Measurement of Cable TV Network Downstream RF Signal Levels, 2019.
- [SCTE Measurement] SCTE Measurement Recommended Practices for Cable Systems, Fourth Edition, 2012.
- [Sundaresen, Zhu, et al] Karthik S, Jay Z., Mayank M., James L., "Practical Lessons from D3.1 Deployments and a Profile Management Application (PMA)," SCTE Cable-Tec EXPO, September 30 – October 3.

[Zhu, Sundaresan,
Rupe]

J. Zhu, K. Sundaresan and J. Rupe, "Proactive Network Maintenance using Fast, Accurate Anomaly Localization and Classification on 1-D Data Series," 2020 IEEE International Conference on Prognostics and Health Management (ICPHM), Detroit, MI, USA, 2020, pp. 1-11, doi: 10.1109/ICPHM49022.2020.9187045, <https://ieeexplore.ieee.org/document/9187045>.

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>
- International Telecommunication Union (ITU), Place des Nations, CH-1211, Geneva 20, Switzerland; Phone: +41-22-730-51-11; Fax: +41-22-733-7256; <http://www.itu.int>
- SCTE•ISBE - Society of Cable Telecommunications Engineers Inc., 140 Philips Road, Exton, PA 19341; Phone: +1-610-363-6888 / 800-542-5040; Fax: +1-610-363-5898; <http://www.scte.org/>

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 that is added to the desired signal and is generally independent of the signal. Includes thermal noise, narrowband ingress, and impulse/burst noise.
Amplitude Ripple	Nonflat frequency response in which the amplitude-vs-frequency characteristic of the channel or operating spectrum has a sinusoidal or scalloped sinusoidal shape 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 to the dielectric material's dielectric constant.
Coefficient	Complex number that establishes the gain of each tap in an adaptive equalizer.
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 television 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 television 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.
Decibel (dB)	Ratio of two power levels expressed mathematically as $\text{dB} = 10\log_{10}(P_1/P_2)$.
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, $\text{dBmV} = 20\log_{10}(\text{value in mV}/1 \text{ mV})$.
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 1,218 MHz.
Drop	Coaxial cable and related hardware that connect 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	A multimedia terminal adapter that has been combined with a cable modem.
Equalizer Tap	See <i>tap</i> .
Fast Fourier Transform (FFT)	An algorithm to compute the discrete Fourier transform, 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> .
Fiber Node	See <i>node</i> .
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. It has two components: amplitude (magnitude)-vs-frequency and phase-vs-frequency.
Full-Band Capture (FBC)	CPE-based spectrum analyzer-like functionality in which time domain samples are captured and Fourier transformed to produce a spectral display.
Group Delay	The negative derivative of phase with respect to frequency, expressed mathematically as $-(d\phi/d\omega)$ in units of time such as nanoseconds.
Group Delay Variation 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 to receive, process, and combine 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 that uses 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 that 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 a Kronecker delta function.
Incident Wave	A traveling wave in a transmission line that is propagating from the source to the load.
Index of Refraction	The ratio of the velocity of an electromagnetic wave—specifically what is known as a transverse electromagnetic mode wave—in a vacuum to its velocity in a dielectric material, expressed as $v_{TEM}(\text{vacuum})/v_{TEM}(\text{dielectric})$.
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-vs-frequency and/or phase-vs-frequency departs 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 typically ranges 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 RxMER, 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, such as 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, and odd order distortions such as composite triple beat. 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^+ , represented by gamma Γ and expressed mathematically as E^-/E^+ .
Return	See <i>upstream</i> .
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, for example, one cycle of alternating current is equal to the square root of the mean value of the squares of the original values.
Receive Modulation Error Ratio (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, in which each symbol point exhibits a tight cluster separated from the neighboring symbols.
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.
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 that includes a delay element and multiplier. The gain of the multipliers is set by the equalizer's coefficients. (3) One term of the difference equation in a finite impulse response or an infinite impulse response filter. The difference equation of a finite impulse response 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 caused by 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 42 MHz.
Vector	A quantity that expresses magnitude and direction (or phase) and is represented graphically by an arrow.
Velocity Factor (V)	The reciprocal of index of refraction, expressed in decimal format.
Velocity Of Propagation	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, velocity of propagation is about 85% to 87% of the speed of light in a typical coaxial cable and about 67% to 69% in a typical single-mode optical fiber.

4 ABBREVIATIONS

This document uses the following abbreviations.

µs	microsecond
ADC	analog-to-digital converter
AGC	automatic gain control
AWGN	additive white Gaussian noise
BCH	Bose, Ray-Chaudhuri, Hocquenghem
BER	bit error ratio
BPSK	binary phase-shift keying
CATV	community antenna television
CCF	Common Collection Framework
CM	cable modem
CMTS	cable modem termination system
CNR	carrier-to-noise ratio
CPE	customer premises equipment
CRC	cyclic redundancy check
CSO	composite second order
CTA	cable termination assembly
CTB	composite triple beat
CW	continuous wave
DAA	distributed access architecture
dB	decibel
dBc	decibels relative to carrier
dBmV	decibel millivolt
DFT	discrete Fourier transform
DOCSIS	Data-Over-Cable Service Interface Specifications
DPD	downstream profile descriptor
DRW	dynamic range window
DS	downstream
DSP	digital signal processing
DWDM	dense wavelength division multiplexing
EIA	Electronic Industries Alliance
E_s/N_0	energy per symbol-to-noise density ratio
FBC	full-band capture
FEC	forward error correction
FFT	fast Fourier transform
FM	frequency modulation
FMA	Flexible MAC Architecture
FR	fine ranging
FTTH	fiber to the home
Gb	gigabit
GHz	gigahertz
GPIB	General Purpose Interface Bus
GSM	Global System for Mobile Communications
HE	headend
HFC	hybrid fiber-coax
HSD	high-speed data

HTTPS	Hypertext Transfer Protocol Secure
Hz	hertz
I	in-phase
ICFR	in-channel frequency response
IEEE	Institute of Electrical and Electronics Engineers
IFFT	inverse fast Fourier transform
IP	Internet Protocol
IPDR	IP detail record
IR	initial ranging
ITU	International Telecommunications Union
kHz	kilohertz
LC	inductor capacitor or resonant (circuit)
LDPC	low-density parity check
LPF	low-pass filter
LSB	least significant bit
LTE	Long-Term Evolution
MAC	media access control
MCNS	Multimedia Cable Network System Partners Limited
MER	modulation error ratio
MHz	megahertz
MIB	management information base
ms	millisecond
MSB	most significant bit
Msym/s	megasymbol per second
MTA	multimedia terminal adapter
mV	millivolt
NCP	next codeword pointer
NDA	nondisclosure agreement
NPR	noise power ratio
ns	nanosecond
OCD	OFDM channel descriptor
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
OPT	OFDM Downstream Profile Test
P-IE	Probe Information Element
P-MAP	probe MAP, probe bandwidth allocation map
PDU	protocol data unit
PLC	PHY link channel
PNM	Proactive Network Maintenance
pCore	physical MAC core
PON	passive optical network
PSD	power spectral density
Q	quadrature
QAM	quadrature amplitude modulation
QPSK	quadrature phase-shift keying
R-PHY	Remote PHY
RBW	resolution bandwidth
REQ	request
RF	radio frequency

RFoG	radio frequency over glass
RMC	remote MAC core
RMD	remote MACPHY device
RMS	root mean square
RPD	R-PHY device
RSP	response
RxMER	receive modulation error ratio
SDR	software-defined radio
SID	spectral impairment detector
SLA	service level agreement
SNMP	Simple Network Management Protocol
SNR	signal-to-noise ratio
SoC	system on a chip
STB	set-top box
TCM	Trellis-coded modulation
TCS	transmit channel set
TDR	time domain reflectometer
TEM	transverse electromagnetic
TFTP	Trivial File Transfer Protocol
TLV	type length value
TR MB	Trigger Message Block
TV	television
UDP	User Datagram Protocol
UE	user equipment
UEPI	Upstream External PHY Interface
UHF	ultra high frequency
UTC	Coordinated Universal Time
vCore	virtual MAC core
VHF	very high frequency
VM	virtual machine
VOD	video on demand
VSA	vector signal analyzer

5 DOWNSTREAM PNM MEASUREMENTS

5.1 Downstream Transmit and Received Power

5.1.1 Overview

This section describes DOCSIS 3.1 OFDM downstream power measurement for CMTS transmit and cable modem receive, as well as how the measurement differs compared to legacy 6 MHz- or 8 MHz-wide SC-QAM signals.

5.1.2 Background

[SCTE Measurement] defines digital signal power (also called digital channel power) as

"...the power level as measured by a power meter which uses a thermocouple as a transducer. That is, the measurement is the average power in the signal, integrated over the actual occupied bandwidth of that signal. The signal power is normally expressed in decibels with respect to one millivolt rms in a 75Ω system (dBmV). Thus, the measurement reported is the rms value of the sinusoid that would produce the same heating in a 75Ω resistor as does the actual signal." (p. 17)

From the above definition, digital channel power for downstream SC-QAM signals carried in cable networks is the average power in the signal's 6-MHz occupied bandwidth (8 MHz in Europe).

This definition does not apply to a DOCSIS 3.1 downstream OFDM signal because the signal's occupied bandwidth can vary from 24 MHz to 192 MHz. [PHYv3.1] defines OFDM power as the power per CTA channel, which provides RF signal level information comparable to SC-QAM digital channel power, i.e., the average power per 6 MHz. [CTA-542-D] defines the frequencies for the 6 MHz-wide channels used in cable television networks.

5.1.3 OFDM Power at the CMTS

The first part of this section is taken from [PHYv3.1], Section 7.5.9.1, "CMTS Output Electrical Requirements."

For OFDM, all modulated subcarriers in an OFDM channel are set to the same average power (except pilots which are boosted by 6 dB). For purposes of spurious emissions requirements, the "commanded transmit power per channel" for an equivalent legacy DOCSIS channel is computed as follows:

- CMTS power is configured by power per CTA channel and number of occupied CTA channels for each OFDM channel.
- For each OFDM channel, the total power is Power per CTA channel + 10log₁₀(Number of occupied CTA channels) for that OFDM channel.
- CMTS calculates power for data subcarrier and pilots (using total number of non-zero valued (non-excluded) subcarriers).
- CMTS calculates power in 6 MHz containing PLC.
- For the spurious emissions requirements, power calculated for the 6 MHz containing the PLC is the commanded average power of an equivalent DOCSIS legacy channel for that OFDM channel.

A CMTS MUST output an OFDM RF modulated signal with the characteristics defined in Table 41 - RF Output Electrical Requirements, Table 42 - CMTS Output Power, and Table 43 - CMTS Output Out-of-Band Noise and Spurious Emissions Requirements. Legacy DOCSIS RF modulated signal characteristics are provided in Section 6.2.22.

The CMTS configuration for RF output level will be equivalent for an OFDM signal and an SC-QAM (ITU-T J.83 Annex B) signal since both are based on 6 MHz occupied bandwidths.

The required power per channel for N_{eq}' channels combined onto a single RF port is

$$\text{Power} = 60 - \text{ceil} [3.6 * \log_2(N^*)] \text{ dBmV}$$

Note the use of the ceiling (ceil) function and base 2 logarithm in the above formula.³

For more information about CMTS output power, see [PHYv3.1], Section 7.5.9.1.1, "Power per Channel for CMTS," Table 42 in particular.

5.1.4 Setting OFDM Channel Power at the CMTS

In most cases, OFDM channel power is set to the same power spectral density (PSD) as SC-QAM signals.⁴ In other words, the heights of the OFDM and SC-QAM "haystacks" as seen on a spectrum analyzer are the same (Figure 2).




Figure 2 - Example of Same PSD Used for OFDM Signal and SC-QAM Signals

5.1.5 RF Input to the Cable Modem

The following RF signal level parameters for the cable modem input are from [PHYv3.1].

- Total input power less than 40 dBmV, 54 MHz to 1.794 GHz (negligible input power outside this frequency range).
- Level range equal to -9 dBmV to +21 dBmV (in 24-MHz occupied bandwidth), which is equivalent PSD to -15 dBmV to +15 dBmV per 6 MHz-wide SC-QAM signal.

5.1.6 Downstream Power Measurement on a Spectrum Analyzer

When looking at the OFDM signal on a conventional spectrum analyzer with a typical resolution bandwidth (RBW) filter setting of 300 kHz, a correction factor needs to be calculated to get the power in a 6-MHz-equivalent bandwidth and, if desired, for total power.

For example, assume a spectrum analyzer power cursor reads 35 dBmV. This value is the power in the analyzer's 300-kHz RBW. The correction factor used to find the power in the 6-MHz bandwidth is $10 \cdot \log_{10}(6/0.3)$, or 13 dB. The correction factor is added to the power level from the analyzer to find the actual power equivalent for 6 MHz.

$$35 \text{ dBmV} + 13 \text{ dB} = 48 \text{ dBmV}$$

If the OFDM signal is flat, the total power of a 192 MHz-wide OFDM signal is found by this equation.

$$38 \text{ dBmV} + 10 \cdot \log_{10}(192/6) = 53 \text{ dBmV}$$

³ The ceiling function is a mathematical function that returns the lowest valued integer that is greater than or equal to a given value. For example, $\text{ceil}[32.1] = 33$, but $\text{ceil}[32] = 32$.

⁴ Some operators transmit the downstream OFDM signal at a higher PSD than SC-QAM signals to increase headroom for higher modulation orders such as 4096-QAM. Caution is urged because of the potential to increase downstream total power enough to cause laser clipping and/or amplifier (and possibly CPE) compression. Extensive lab and field testing beforehand is recommended.

This value can be useful information for laser loading concerns. The shape of the RBW filter can affect power measurements at the edge of the OFDM band or in the presence of exclusion bands, and the calculation of total power will be affected by a non-flat (e.g., tilted) OFDM signal.

In some cases, it might be preferable, or easier, to measure OFDM power with a dedicated DOCSIS 3.1 test instrument or a DOCSIS 3.1 cable modem (or modem-equipped field meter). Some test instruments are also capable of measuring total power within a specified bandwidth, such as that of an entire OFDM signal. Be aware of the absolute accuracy limitations of the device being used for the power measurement.

NOTE: What about OFDM power measurements in other than a 6-MHz equivalent? In areas that use Euro-DOCSIS and 8 MHz-wide SC-QAM signals, confusion about OFDM channel power can occur. Although [PHYv3.1] defines OFDM power as the power per CTA channel (i.e., power per 6 MHz), some operators might prefer the equivalent power per 8 MHz. OFDM channel power per 6 MHz reported by a DOCSIS 3.1 device can be converted to an 8-MHz equivalent by taking into account a correction factor added to the 6 MHz power equal to $10 \cdot \log_{10}(8 \text{ MHz}/6 \text{ MHz})$, or 1.25 dB. In the above example of 38 dBmV (power per 6 MHz), the addition of the correction factor would give a channel factor of 39.25 dBmV. The total power would, of course, be the same as before, taking in account that the number of 8 MHz-wide channels occupied by a 192 MHz-wide OFDM signal is 24: $39.25 \text{ dBmV} + 10 \cdot \log_{10}(24) = 53 \text{ dBmV}$.

5.1.7 Downstream Power Measurement on Other Instruments

In addition to a conventional spectrum analyzer, other types of test equipment can be used to measure downstream OFDM power. At the high end are laboratory-grade vector signal analyzers and DOCSIS analyzers. Additionally, several test equipment vendors have developed DOCSIS 3.1 field meters and similar instruments, most of which include an embedded DOCSIS 3.1 cable modem to support a variety of OFDM measurements (channel power, RxMER per subcarrier, codeword errors, etc.). Here, too, be aware of the absolute accuracy limitations of the device being used for the power measurement.

As mentioned, OFDM channel bandwidth (occupied bandwidth) can vary from a minimum of 24 MHz (minimum modulated spectrum of 22 MHz) to a maximum of 192 MHz (maximum encompassed spectrum of 190 MHz). In the channel lineup of Figure 3, a 96 MHz-wide OFDM signal includes a 1-MHz guard band on each end of the channel, so the encompassed spectrum is 94 MHz with 1,880 subcarriers at 50 kHz per subcarrier. The screen shots in Figure 3 and Figure 4 are from a field meter equipped with a DOCSIS 3.1 modem.




Figure 3 - Downstream Channel Lineup with a 96 MHz-Wide OFDM Signal at the Upper End of the Spectrum

In the case of the OFDM signal shown in Figure 3, the number of occupied CTA channels is 96 MHz/6 MHz, or 16. Figure 4 shows the OFDM power for that OFDM signal, with the power per 6 MHz represented by multiple adjacent horizontal lines whose heights represents the average power in each 6 MHz-wide CTA channel slot.




Figure 4 - Power Measurements for a 96 MHz-Wide OFDM Channel as 16 Occupied Channels in 6-MHz Segments and for the Lower and Upper Adjacent SC-QAM Channel at Left and Right Edges

As defined in [PHYv3.1], the OFDM signal level at the CM input should be between -15 dBmV and +15 dBmV (power per 6 MHz). The example in Figure 4 meets those requirements with the following measurements: average, -3.3 dBmV; minimum, -5.3 dBmV; and maximum: -2.7 dBmV.

An important point when looking at OFDM power as shown in Figure 4 is that the lower and upper edges of the OFDM signal might not fully occupy their respective CTA channel slots, even though those slots are considered "occupied" by the OFDM signal. If one considers each 6-MHz CTA channel occupied by the OFDM signal as a "bin" containing power data, the edge bins could have less power than the other bins between the edges. In other words, the 6-MHz slots on each end of the OFDM signal could show lower power than the fully occupied 6-MHz slots.

5.1.8 Impact of Pilots on OFDM Power

Downstream OFDM continuous and scattered pilots are boosted by 6 dB, per [PHYv3.1]. Figure 5 shows an example of a 192 MHz-wide downstream OFDM signal, in which the boosted continuous pilots are visible. There is some question as to whether pilot boosting affects OFDM channel power, but the presence of continuous and scattered pilots in an OFDM signal affects the signal's total power by less than 1 dB.




Figure 5 - Spectrum Analyzer Display of an OFDM Signal Showing the Boosted Continuous Pilots

5.1.9 How to Implement

Methods for polling these data elements from cable modems are given in [CM-OSSIV3.1]. Additional details are available from CableLabs in the Common Collection Framework.

5.1.10 Conclusions and Future Work

Downstream transmit and received power levels are basic, important measurements for a stable service. However, differences among measurement devices and conditions must be carefully considered when interpreting these measurements. For more information, see [SCTE 257]. Also, different CMTS and CM vendors should be compared to note any differences in specification interpretation.

5.2 Downstream Wideband Spectrum Analysis

5.2.1 Overview

The PNM Best Practices document for DOCSIS 3.0 technology [PNMP-3.0] included full-band capture (FBC), which is possibly one of the most useful features available in PNM. It allows operators to quickly capture and analyze the downstream RF spectrum from the location and perspective of the cable modem. This powerful tool allows remote RF power measurement without deploying a technician or requiring access to the cable modem's physical location. Section 7 of [PNMP-3.0] describes FBC in detail, so this document will focus on new information specific to DOCSIS 3.1 technology and additional improvements made since [PNMP-3.0].

5.2.2 Description of Measurement

Downstream wideband spectrum analysis in a CM measures the RF at the receiver front-end without relying on the demodulators (Figure 6). This measurement is taken by sampling RF in the time domain with an analog-to-digital converter (ADC) and then converting to the frequency domain using FFT. The resulting measurement is stored as logarithmic magnitude for each frequency bin of the given center frequency and the resolution bandwidth specified in the MIB configuration parameters listed in Table 58 of [CM-OSSiv3.1].




Figure 6 - Digital Spectrum Analyzer Block Diagram

5.2.3 Capabilities and Limitations

Although the ability to capture and report wideband RF spectrum is a powerful tool for cable operators to analyze their networks, there are limitations compared to traditional field instruments. One significant limitation of FBC is that the magnitude-only measurement lacks important information about the phase component of the signal. Without knowing the phase component, analysis is limited to power measurement, which precludes operators from taking signal quality measurements such as RxMER. Because of this, some impairments such as low-level ingress cannot be detected because the interference level does not exceed the peak-to-average power ratio of the signal. In cases when the interference is significant, this technique works only in extreme interference levels (Figure 8). Other limitations are based on FBC's capture speed and dynamic range, which makes it unsuitable for certain troubleshooting use cases. FBC has a relatively slow capture rate of several samples per second, so certain types of interference such as fast, bursty noise can be difficult or impossible to detect. Additionally, with the wideband front-end, dynamic range is typically between 30 dB and 40 dB, depending on implementation. This range generally works well for signal power measurement, but it does have limitations when the capture window power variation exceeds the dynamic range. In the example shown in Figure 7, the dynamic range limits can be seen at 800 MHz and 880 MHz as artificial power steps in the noise floor. This effect is caused by the front-end gain controller automatically adjusting to maximize the signal fidelity (Figure 9) at the expense of noise floor accuracy. Note that because full-band spectrum capture is a compilation of many individual frequency captures over time, discontinuities can occur that are not necessarily due to network impairments.

The example spectrum shown in Figure 7 includes 6-MHz SC-QAM, analog pilots, and OFDM signals. The capture in Figure 8 is from the same node but shows that the location has a shielding integrity problem. Note the FM, VHF, UHF, LTE, and GSM interference.

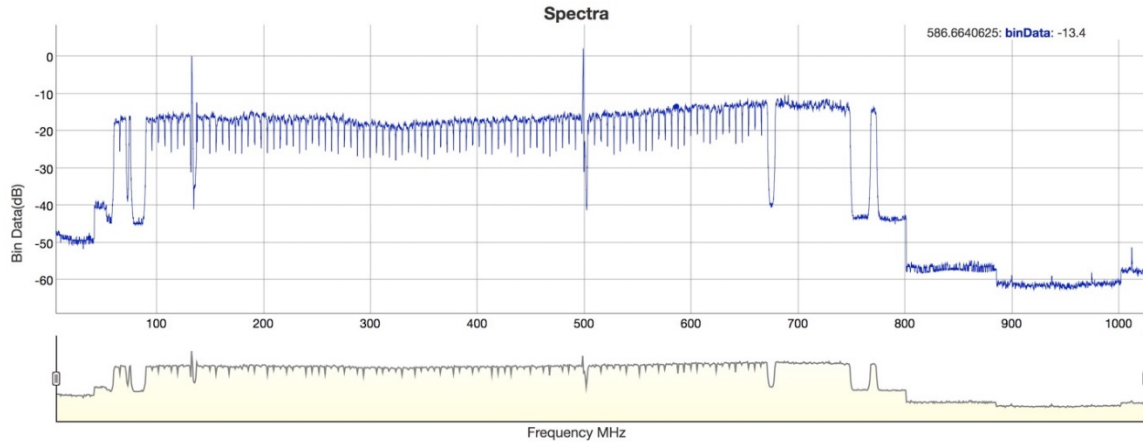


Figure 7 - Typical Downstream Wideband Spectrum Analysis of RF Spectrum

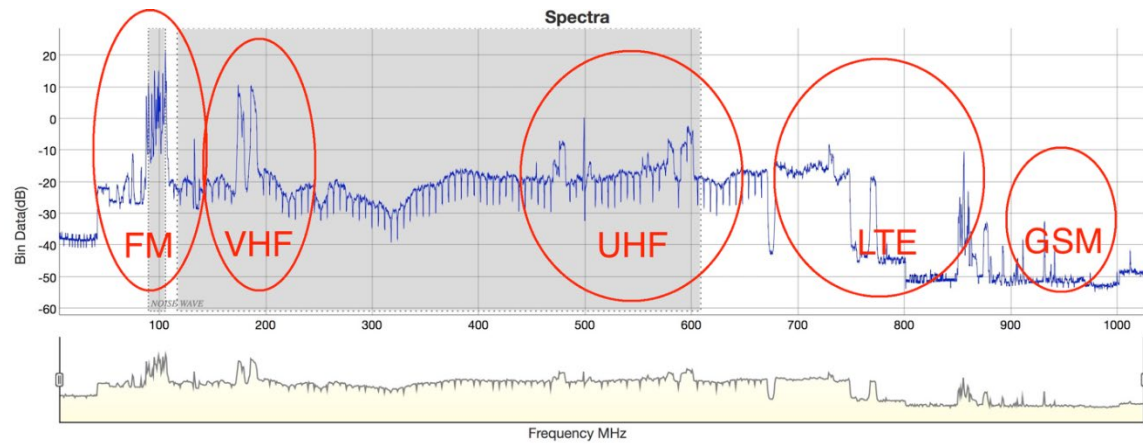


Figure 8 - Downstream Wideband Spectrum Analysis of RF Spectrum Including Interference

In the example in Figure 9, an artificial power step, the result of the dynamic range limitation of this device, is visible in the noise floor at 800 MHz. The actual noise floor is much lower, 900 MHz, but it cannot be acquired because of the presence of signal power that exceeds the dynamic range within the capture window.




Figure 9 - Dynamic Range Limitations of Downstream Wideband Spectrum Analysis

Another limitation of wideband spectrum capture is that it is constrained to the downstream spectrum available to the receiver. This constraint is a function of the filtering designed to protect the CM receiver from upstream burst energy. However, as seen in Figure 10, the upstream band is still available to the FFT and does produce a spectrum capture that is subject to the diplex filter's rejection. Noise can sometimes be observed through, around, and near the diplexer's rejection area. This technique can be used to detect upstream noise by inference rather than direct observation.

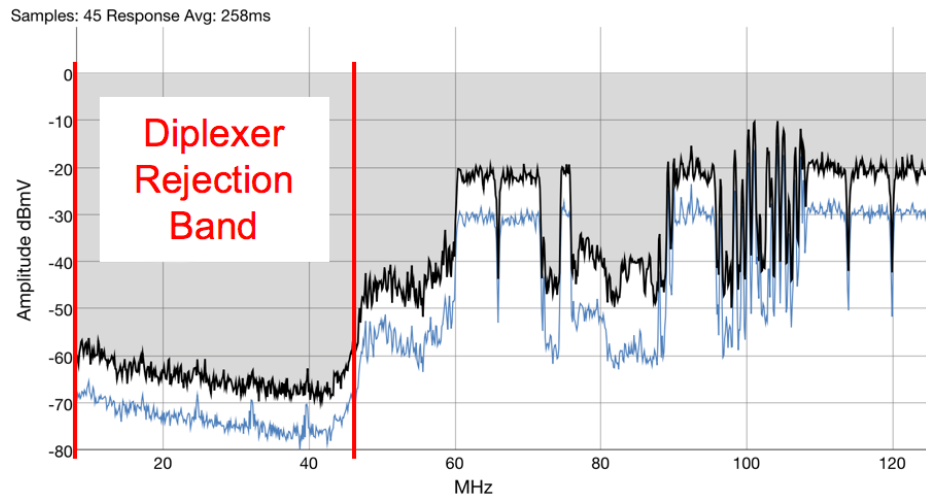


Figure 10 - Upstream Spectrum and Diplexer Filter Rejection

5.2.4 Averaging

As with most spectrum analyzers, averaging typically improves the accuracy of power measurement. However, there are many ways to implement averaging and different things to consider, such as capture speed and number of samples. In the case of PNM, the MIB specifies an averaging function for the wideband spectrum capture engine, going so far as to describe a uniform technique, the "leaky integrator" method, though the results when implemented in the field are less than uniform. Figure 11 illustrates this point by comparing the results of two DOCSIS SoC (system on a chip) manufacturers captured on the same RF plant. Both samples used the maximum averaging allowed for their respective implementations.




Figure 11 - Maximum On-Chip Averaging Achieved with Different Implementations

Given the differences in implementations, accommodations must be made in resolution bandwidth and the use of external averaging to improve detection consistency. External software applications are sometimes required to average multiple samples.

For power to be averaged correctly, the amplitude bins need to be converted from log decibels to linear power. The same frequency, resolution, and windowing parameters need to be used when collecting samples. Then the samples are summed and divided as linear power and converted back to log decibels. This operation must be done bin by bin by using common bin center frequencies. An example using Java code is provided.

```
double linearPower = 0;

// convert 3 dB bins (same center frequency) to linear power and add them
linearPower += Math.pow(10.0, (0.1 * dBbin1));
linearPower += Math.pow(10.0, (0.1 * dBbin2));
linearPower += Math.pow(10.0, (0.1 * dBbin3));

// divide by the number of samples to average
linearPower /= 3;

// convert back to log from linear
double totalLogPower = 10.0 * Math.log10(linearPower);
```

5.2.5 Channel Power Accuracy and Calibration

Due to variations in the component chain of the RF tuners, power calibration needs to be used to maximize the power accuracy of information reported by the DOCSIS MIBs. In most contemporary DOCSIS tuner implementations, the power levels are measured against a known source, and a calibration (power offset) value is stored in an internal table for future reference. Because the wideband spectrum measurement is taken ahead of the tuner and demodulator, it is not subject to equalization or power correction with the internal calibration tables, so the power measurements for a given spectrum bandwidth will generally not match the total channel power measured by the tuner and reported in the channel power MIBs.

The potential for mismatched information reported between the DOCSIS power and PNM wideband spectrum capture MIBs can be a problem for the people and systems responsible for interpreting the information. One technique for overcoming the reporting discrepancy is to compare the channel power measured through wideband capture with the power measured by the CM tuner. The power differential measured within each of the tuner channels can be averaged

to produce a calibration offset value. This technique essentially reveals the internal calibration offset, which can then be applied to the power measurements obtained by the wideband power measurement. Depending on the component chain variation, this calibration can vary significantly, so using it is strongly advised to avoid confusion.

Calibration begins by using three DOCS-IF MIBs to acquire the calibrated channel receiver power from the DOCSIS downstream tuners: docsIfDownChannelFrequency, docsIfDownChannelPower, and docsIfDownChannelWidth. In the case of a relatively flat RF spectrum, a simple strategy can be to create an average of all downstream receive power values. All averaging needs to be done in linear power, so the log decibel millivolt (dBmV) values need to be converted to linear, averaged, and then converted back to log decibel millivolt. In the example shown in Table 1, the average receive power across all channels is -0.1 dBmV.

The next step is to instruct the spectrum capture MIB to measure the same channel center frequency and RF bandwidth by setting FirstSegmentCenterFrequency, LastSegmentCenterFrequency, and SegmentFrequencySpan to produce the desired channelization. In the example shown in Table 1, these values would be 603000000, 645000000, and 6000000, respectively. After enabling the PNM MIB, the resulting BinAmplitudeFileData will return the TotalSegmentPower measurement using the wideband capture MIB. These values can be averaged in a manner similar to the previous step by using the DOCSIS tuner power measurement. The example shown in Table 2 yields an average channel power of 3.04 dBmV across the same eight channels of RF spectrum. Note that the values in the NumberOfAverages parameter of the calibration routine should match the values used for regular power measurement to minimize inconsistencies. In general, greater averaging results in more accurate (useful) power measurements of wideband signals.

Finally, the average channel powers are subtracted to produce a single offset value that can be applied to all of the channels measured using the wideband spectrum analysis engine. In this example, 3.04 dBmV minus 0.1 dB produces an offset value of 2.94 dBmV, which now can be applied to channels of similar spectrum bandwidth. Likewise, the 2.94 dBmV can be converted to different equivalent power bandwidths and used to calibrate any variation of resolution bandwidth power measurements.

Table 1 - Example of Downstream DOCSIS Channel Center Frequency and Receive Power

Center Frequency (MHz)	603	609	615	621	627	633	639	645
Channel Receive Power (dBmV)	0.1	-0.2	0	0.2	0.1	0	-0.2	-0.1
Channel Width (MHz)	6	6	6	6	6	6	6	6

Table 2 - Example of Wideband Spectrum Channel Center Frequency and Receive Power

Center Frequency (MHz)	603	609	615	621	627	633	639	645
Channel Receive Power (dBmV)	3.3	3.0	3.0	3.1	3.0	3.2	2.9	2.8
Channel Width (MHz)	6	6	6	6	6	6	6	6

This method of calibration has been demonstrated at scale to produce significant improvements in power accuracy when compared to commercial-grade field measurement tools (Figure 12). However, the achievable performance of these tools may differ depending on total input power, temperature, and filter design, among other things. If maximum power level accuracy is a requirement, then specific makes and models of cable modems should be characterized for a full understanding of the limitations under these different conditions.

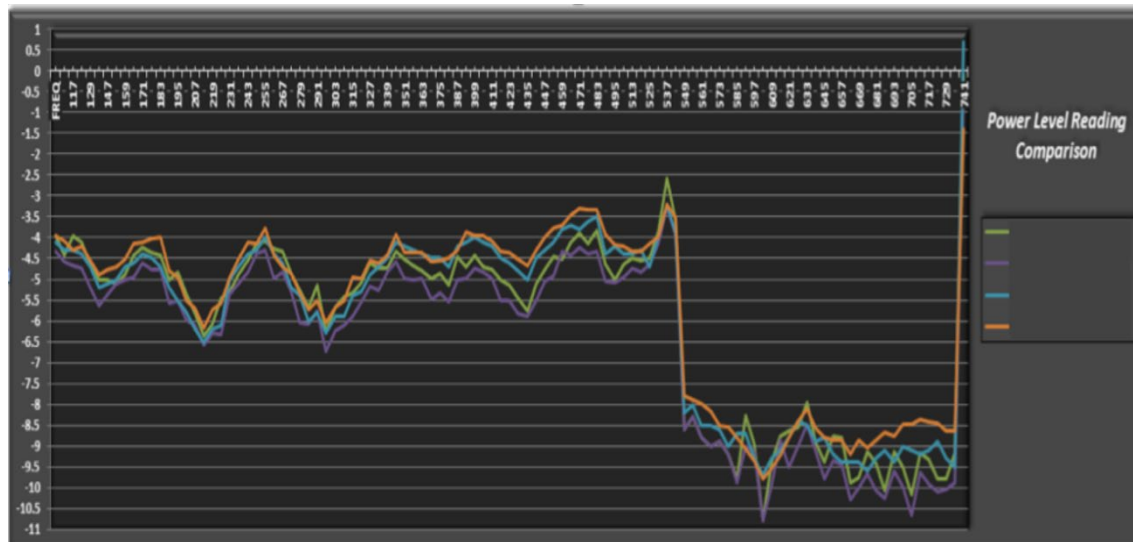


Figure 12 - Power Level Comparison of Field Gear and Wideband Spectrum Analysis after Calibration

(Figure Note: Individual product names have been removed.)

5.2.6 Spectral Impairment Detection

The ultimate goal of wideband spectral analysis is to proactively detect network impairments prior to impacting the performance of the tuners and demodulators. As stated earlier, having the full RF spectrum provides a larger context for troubleshooting beyond the relatively narrow bandwidth observed by the downstream tuners and signal demodulators. For example, many common impairments such as standing waves (see Figure 13) cannot be measured within the bandwidth of a single 6-MHz SC-QAM channel. However, having the full spectrum bandwidth of the wideband analyzer improves the ability to perceive these types of problems. In the example of a standing wave, knowing the peak-to-peak or null-to-null frequency periodicity of the wave is a useful tool for understanding fault distances of impedance cavities.

As an adjunct project for the CableLabs PNM working group, some of the members have collaborated on software libraries that automatically detect common RF impairments. This software, known as the spectral impairment detector (SID), was written in the Java programming language. It is available to NDA members. Many of the following impairments and functions are included, with a brief description of each.

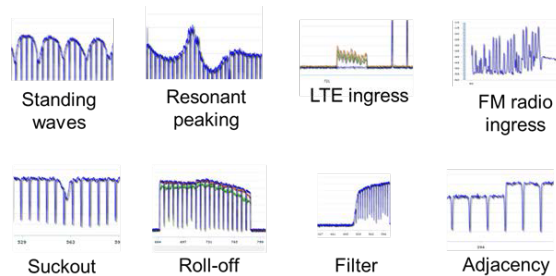


Figure 13 - Examples of Typical RF Spectral Impairments

5.2.6.1 SID Channelization

The SID detection algorithms are highly dependent on analysis of certain channel properties such as tilt, power, and frequency. To facilitate this analysis, an important function of the library is to channelize the FBC amplitude bin using one of several available methods. The channelization functions are contained in the PNMSpectrumChannelFactory and PNMSpectrumTuner software objects, which notably do not actually tune the RF spectrum. Because FBC lacks

a sufficiently fast capture speed (symbol rate) and is limited to log magnitude bins, true tuning is impossible. The library does, however, provide several tuner-like capabilities that are important to the subsequent function of the SID software.

Predefined channel maps and blind scanning are the two primary methods used by SID to convert amplitude versus frequency bins to usable channel information. The first method is the most common because most operators and users of the software will generally have this information available. The minimum information required includes the center frequencies of known channels within the RF lineup. It is also possible to provide discrete programming information for each of the RF channels, which is helpful when troubleshooting customer-reported issues. The `PNMSpectrumTuner` object may also be invoked with a standard U.S. or European channel plan, which contains a list of well-known channel center frequencies and widths. Either method results in a list of channel objects, which provides more useful information beyond the simple amplitude bins.

5.2.6.2 *SID Channel Measurements*

Channelizing the amplitude versus frequency bins makes the channel information more meaningful for analysis. The tuner library contains several software types within the channelization, including `PNMDocsisChannel`, `PNMSpectrumCATVChannel`, and `PNMSpectrumChannel`, which expose more useful attributes for use in analysis. Some of these attributes are listed here.


- **Start, center, and end frequency**—the respective frequencies of the channel in hertz
- **Name**—a user-friendly name for the RF channel that may be displayed in user interfaces
- **Width**—the channel width specified in hertz
- **CTA**—(formerly EIA) the industry standard identifier for a given channel and frequency that refers to a 6-MHz channel plan (Europe uses an 8-MHz plan)
- **ICFR**—in-channel frequency response measured as the minimum and maximum within the channel alpha
- **Tilt start and end levels**—the power levels of the leading and trailing bins within the channel alpha
- **Alpha median dB**—the median power value calculated from the median of the channel's alpha bins
- **Ripples**—calculated as a number of incidents in which a bin power transition has occurred within a given threshold
- **Tilt**—calculated as the difference between the tilt start and tilt end values; results may be positive or negative
- **MER/SNR**—known only for DOCSIS channels when calibrated using the tuner MIB values
- **Bin pointers**—useful indices that point to the respective channel locations within the main bin array

5.2.6.3 *Adjacency and Channel Alignment*

Adjacency misalignment is an RF impairment recognizable as a delta in channel power between adjacent channels or groups of channels. It may be observed multiple times in the spectrum. Adjacency misalignment issues can be attributable to RF combining or source issues at the headend or hub or in a node utilizing a broadcast narrowcast overlay or a split band configuration when the two spectrums are not combined correctly.

When adjacency misalignment is observed, the lower power channels may indicate poor performance, seen as a poor modulation error ratio (MER) when the delta between channels is large. This condition can manifest as lost packets, video tiling, freezing, or in very extreme cases, black screens at a customer's home. Because adjacency misalignment is introduced very early in the downstream signal path, it has the potential to impact a significant number of customers. The following will further describe how an adjacency misalignment can occur.

The downstream spectrum received at the CPE is produced by combining multiple non-overlapping regions of spectrum for different services. In a highly simplified example, shown in Figure 14, a three-input combiner is used to combine three types of services: video, high-speed data (HSD) plus voice, and video on demand (VOD).



[The services can be placed anywhere in the allocated spectrum.]

Figure 14 - Example of a Simplified Three-Input Combiner

The combiner takes the three service bandwidths and frequency multiplexes the channels from 96 MHz to 870 MHz for transmission in a single cable, which then feeds an amplifier or an optical transmitter. Adjacent channel misalignment can be introduced if the channels are not set correctly or if the combiner losses are not accounted for properly.

In the ideal scenario, every channel is set at the same power level, resulting in a flat spectrum. If, however, the combined channels arrive at the headend optical transmitter at differing channel power levels, the result will be a misaligned RF spectrum affecting all customers downstream of that transmitter.

As a point of reference, the Federal Communications Commission [47 C.F.R. §76.605] defines the maximum signal level of an adjacent channel as being within 3 dB of any visual carrier within 6 MHz.

For good engineering practices, this rule should apply for adjacent digital channels as well. From the headend, the digital channel levels should be within 1 dB if they are using the same QAM data rate.

Digital Channel Mask

Each digital channel has a spectral mask that defines how much energy is confined within the channel bandwidth and the allowed noise and distortion that can impact adjacent and non-adjacent channels. Figure 15 shows the maximum amplitude limits of the noise and distortions that can be found in adjacent and non-adjacent channels relative to the channel maximum power, referenced as 0 dBc.




Figure 15 - DOCSIS Single Channel Spectral Mask

A misaligned channel has an impact on the adjacent and non-adjacent spectrum. As shown in Figure 16, the red channel is 10 dB above where it should be, as are the channel limits as defined by the channel's spectral mask. The misalignment has a negative impact on the other channels, which may be seen at the CPE as poor RxMER.




Figure 16 - Effect of Channel Misalignment on Other Channels

Figure 17 shows an ideal spectral input to the transmitter, flat across the spectrum with each channel utilizing QAM signals with the same modulation rate, typically 256-QAM. Each channel is set at the same power level. Figure 18, however, shows channels that are not aligned properly prior to insertion into the transmitter. Once this alignment is introduced into the downstream path, it cannot be changed from a channel power relationship.




Figure 17 - Spectral Input with Channels Set to the Same Power Levels




Figure 18 - Spectral Input with Channels Set with Misaligned Power Levels

Other plant impairments can also cause a change in adjacent power levels. These impairments will impact the in-channel flatness as compared to a channel power level change. Figure 19 shows three adjacencies to measure: second to third channels, fourth to fifth channels, and fifth to sixth channels.




Figure 19 - Spectral Input Misalignment Caused by Plant Impairments (Not by Power Level)

Non-adjacent channel misalignment should be analyzed with the engineering performance parameters outlined in the previous section. A single misaligned channel will have an impact on the other analog or digital channels in the spectrum.

Figure 20 shows data for a single CPE. It has a misalignment between adjacent SC-QAM channels of about 8 dB. Each division is 5 dB. Depending on the level relationship to the control carriers, the channels at the higher power level may have a significant impact on downstream performance.



Figure 20 - Adjacent Misalignment in SC-QAM Channels in a CPE

The level difference of the misalignment is determined by first measuring the power levels of the lower CTA-L# channel number and the higher CTA-H# channel number. The delta is then recorded by subtracting the CTA-L# power level from the CTA-H# power level.

In Figure 21, data extracted from PNM show multiple CPEs connected to the same node and/or transmitter. The channel alignment issue will be found at either the headend or the node. As shown in the figure, the lower channels may be part of a broadcast narrowcast overlay that is not set correctly.




Figure 21 - Adjacent Misalignment in Multiple CPEs with the Same Node and/or Transmitter

The goal of this measurement is to identify adjacency misalignment issues within an operator's network. It is assumed that the impact to the customer experience will become more noticeable as adjacency misalignment increases in amplitude, as evidenced by customer contacts into operator's care and support channels. When adjacency misalignments are corrected, the customer experience should improve and be quantifiable by a decrease in customer contacts.

By implementing the logic to comprehensively measure adjacency misalignments, operators can measure the prevalence of the issue and its true impact to operations.

5.2.6.3.1 Adjacency and Channel Alignment Thresholds

The default adjacency detection logic in SID is contained in the AdjacentIdentifier code within the ident package of the software library. The following values are used within the analyzer and may be adjusted to accommodate differences in RF environments or operational standards. The code comments denoted by "//" are used to describe each value.

// MIN_DB_DELTA is the absolute minimum Db delta between two channels that can trigger adjacency. This is one of the very first checks, followed by some additional methods to verify this is adjacency vs. other impairment types.

static final double MIN_DB_DELTA = 2.9;

*// MAX_DB_DELTA is the absolute max Db delta between two channels that can trigger adjacency
// This check is trying to avoid false positives on vacant channels, where a drop in power on adjacent flat channels is simply vacant spectrum.*

static final double MAX_DB_DELTA = 14;

*// MAX_ICFR is the absolute max Db delta between two channels that can trigger adjacency.
// This check is trying to avoid false positives on short narrow spikes of power. Channels with power spikes should be thrown out (for adjacency detection)*

static final double MAX_ICFR = 4;

// MAX_TILT is the max tilt an individual channel can have before the channel data is thrown out as ineligible for adjacency detection.

// This tilt check is done on both the current channel and the previous channel being used to measure Db delta

static final double MAX_TILT = 1.75;

// HIGH_TILT_INDIVIDUAL_CHANNEL is the max tilt a series of individual channels following the current channel being used to measure adjacency

// can have before the data is thrown out as ineligible for adjacency.

static final double HIGH_TILT_INDIVIDUAL_CHANNEL = 2.5;

// CHANNEL_FLAT_CHECK is the number of channels used to check for individual channel high tilt (HIGH_TILT_INDIVIDUAL_CHANNEL)

static final int CHANNEL_FLAT_CHECK = 3;

// HIGH_TILT_INDIVIDUAL_CHANNEL is the max tilt a cumulative series of channels following the current channel being used to measure adjacency

// can have before the data is thrown out as ineligible for adjacency.

static final double HIGH_TILT_CUMULATIVE = 4.0;

// MIN_FREQUENCY is the minimum frequency we'll trigger adjacency on.

// This check is used to avoid a list of low frequency behavior that can trigger false positives on adjacency.

static final double MIN_FREQUENCY = 110000000;

// MAX_FREQUENCY is the maximum frequency we'll trigger adjacency on.

// This check is used to avoid the rolloff area at high frequencies

static final double MAX_FREQUENCY = 700000000;

// HIGH_CHANNEL_TILT_COUNT is the max number of channels in the entire spectrum that can have high tilt before we throw this device out as

// ineligible for adjacency detection.

static final int HIGH_CHANNEL_TILT_COUNT = 10;

// HIGH_CHANNEL_DELTA_COUNT is the maximum number of channels that can have a significant delta detected between adjacent digital channels before

// we throw this device out as ineligible for adjacency detection. If we have a lot of choppiness in the power levels across the spectrum, there

// is likely something else going on - not simple adjacency

static final int HIGH_CHANNEL_DELTA_COUNT = 9;

5.2.6.4 Filters and Missing Channels

Filtering is an RF impediment that may affect multiple channels at a time. Filters have been and continue to be used to remove unwanted signals in the RF path based on the customer's needs. When viewed at the CPE or other test equipment, the filter is characterized by the lack of channels in a specified bandwidth. Four types of filters are found in the field (Figure 22):

- (a) band-pass filter—passes frequencies within a specified range and rejects frequencies outside that range;
- (b) band-stop filter—rejects frequencies within a specified range and passes frequencies outside that range;
- (c) low-pass filter—passes frequencies below a specified frequency and rejects frequencies above;
- (d) high-pass filter—passes frequencies above a specified frequency and rejects frequencies below.

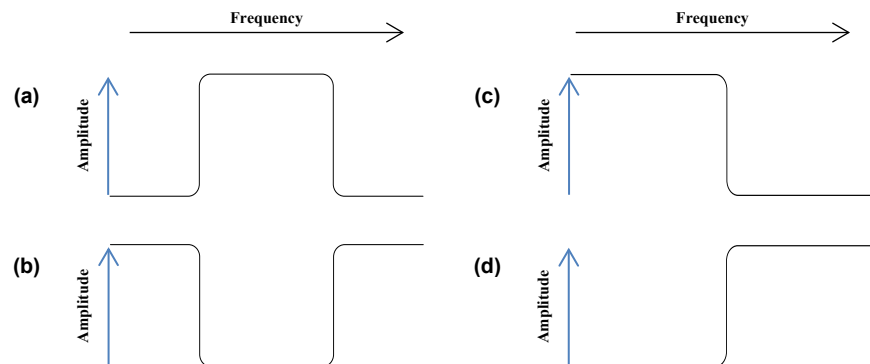


Figure 22 - Filter Types: (a) Band-Pass Filter, (b) Band-Stop Filter, (c) Low-Pass Filter, (d) High-Pass Filter

Figure 23 shows example data for two kinds of filters.

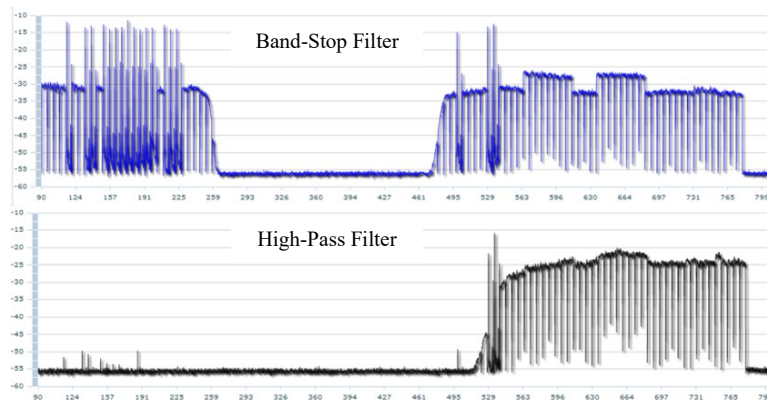


Figure 23 - Example Data for Band-Stop and High-Pass Filters

5.2.6.4.1 Filters and Missing Channels Thresholds

The default adjacency detection logic in SID is contained in the MissingChannelsIdentifier code within the ident package of the software library. The following values are used within the analyzer and may be adjusted to accommodate differences in RF environments or operational standards. The code comments denoted by "//" are used to describe each value.

```
// MIN_DIGITAL_CHANNELS is the minimum number of active channel - any less than this and we mark a
// missing channel impairment
private int MIN_DIGITAL_CHANNELS = 68;
```

// Any channel power over ICFR_THRESHOLD is marked as an active channel (examples - analog or other)

private double ICFR_THRESHOLD = 18.0;

// MIN_DB_THRESHOLD is used to identify devices that have low power across the entire spectrum - the entire spectrum is trashed and we're not sure it's a filter

private double MIN_DB_THRESHOLD = -35;

// a simple way to filter out old 550 plant that isn't using the higher spectrum

private long OLD_PLANT_MAX_FREQUENCY = 575000000;

5.2.6.5 Suckout

A suckout is an RF impairment that often spans multiple CTA channels. It is characterized by a concave notch with sinusoidal boundaries with attenuation in amplitude-power in the frequency domain, often observed over several megahertz. Typically, suckouts are caused by mechanical or grounding issues in active or passive network elements such as seizures, connectors, lids, or fittings. They can be attributable to multiple mismatches evenly spaced through the network. Each mismatch adds to the width and depth of the notch at the frequency of the suckout. An example is the repetitive impedance discontinuities created by the so-called "mold spike" in some disc-style dielectric cables. At a simpler level, a suckout is the result of an impedance mismatch. In structure, it resembles an LC circuit that is tuned to a specific frequency for the purpose of filtering an unwanted signal from the spectrum.

Impact on the customer depends on the depth and width of the notch. The notch can happen anywhere along the RF distribution path, making it one of the more complex issues to troubleshoot. The FBC tool can assist in determining where in the network the issue may originate and if the issue is periodic. Examining the spectrum from several CPEs associated with the node can determine if the issue is affecting multiple CPEs or a single CPE.

Anecdotally, operators have seen the suckout issue as one of the most prevalent RF impairments affecting spectral performance; however, no quantification has been made to date on the prevalence and scale of the impairments. From a customer point of view, a suckout may or may not have a significant impact on their service experience, depending on the services subscribed to and the location of the suckout in the spectrum.

Figure 24 - Typical Spectrum Notch, Figure 25 - Wide Spectrum Notch, Typical Depth, and Figure 26 - Shallow Spectrum Notch, Typical Width show various examples of suckouts. Each of these types of suckout will be seen in the network, but the slopes may not be as straight and the curves not as sharp. In each figure, each vertical line represents a 6-MHz bandwidth, and each horizontal line represents 5 db.




Figure 24 - Typical Spectrum Notch




Figure 25 - Wide Spectrum Notch, Typical Depth




Figure 26 - Shallow Spectrum Notch, Typical Width

Figure 27 and Figure 28 provide additional examples of suckouts demonstrated by PNM. The slopes are more jagged, but note especially the difference in the depth of each notch and the width of the effected spectrum. In each figure, each vertical line represents a 6-MHz bandwidth. Apart from the suckout notch, the spectrum does not appear to be very affected. It could not be confirmed if this is a notch filter signature.




Figure 27 - Suckout as Demonstrated by PNM

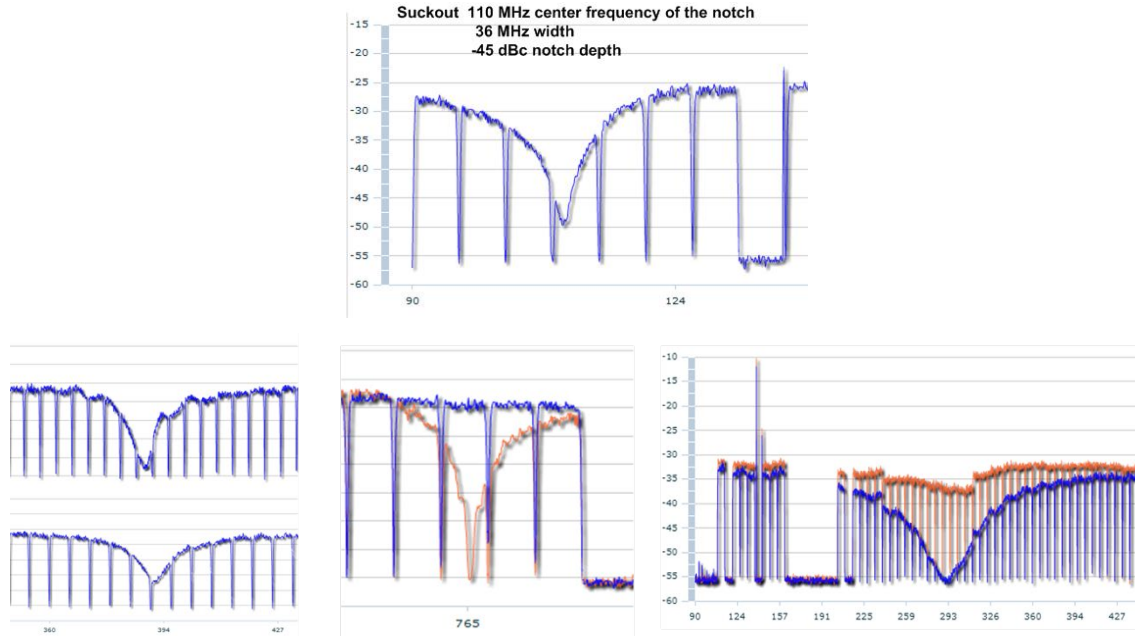


Figure 28 - Expanded View of Suckout Notch and Impact on Several Digital Channels

5.2.6.5.1 Suckout Thresholds

The default adjacency detection logic in SID is contained in the SuckoutIdentifier code within the ident package of the software library. The following values are used within the analyzer and may be adjusted to accommodate differences in RF environments or operational standards. The code comments denoted by "//" are used to describe each value.

// MIN_DB_DELTA is the minimum Db delta used for suckout detection

static final double MIN_DB_DELTA = 5.5;

// WAVE_DETECT_DB is the minimum Db delta used for wave detection

static final double WAVE_DETECT_DB = 4.5;

// MIN_FREQUENCY is the absolute minimum frequency we will trigger a suckout impairment on.

// The minimum is set to avoid a complicated list of things that happen in the lower frequency range.

static final double MIN_FREQUENCY = 110000000;

// HIGH_TILT_INDIVIDUAL_CHANNEL is used to detect substantial individual channel tilt in surrounding channels - tilt being a good indicator the impairment

// is a suckout/ wave and not adjacency

static final double HIGH_TILT_INDIVIDUAL_CHANNEL = 1.9;

// HIGH_TILT_INDIVIDUAL_CHANNEL is used to detect substantial cumulative tilt in surrounding channels - tilt being a good indicator the impairment

static final double HIGH_TILT_CUMULATIVE = 4.0;

// suckouts and waves do not typically have a drastic drop in Db between 2 adjacent channels. If we see a Db drop greater than this

// value, ignore it as a contribution to a suckout

double INDIVIDUAL_DIFF_MAX = 10.0;

// MAX_FREQUENCY is used to avoid looking for impairments in the rolloff area - where it's challenging to

```
differentiate normal
// rolloff from power suckout or wave affects
static final double MAX_FREQUENCY = 700000000;

// CHANNEL_TILT_CHECK is used as the number of adjacent channels to look for high tilt
static final int CHANNEL_TILT_CHECK = 4;

// CHANNEL_WINDOW is the window of channels to average Db values for - for delta calculations to identify a
potential suckout
static final int CHANNEL_WINDOW = 5;

// CHANNEL_COUNT is used as an indicator to differentiate suckout from adjacency. This is the number of
consecutive channels
// that have a declining power level. Adjacency should be an immediate power delta between 2 channels. If the
power continues to decline
// over multiple channels, that's a good indicator this is a suckout and not adjacency. There are suckouts that can
affect just a single channel
// and those must be found using tilt or other methods.
static final int CHANNEL_COUNT = 3;

// Don't try to detect a suckout that crosses over a large gap in vacant spectrum with no channels
static final int MAX_CHANNEL_GAP = 3;

// The sum of both the individual channel tilt and the channel icfr values cannot exceed this value to be eligible as
part of a suckout detection
// this is primarily used to filter out non digital channels returned in the digital channel list
int TILT_ICFR_SUM_MAX = 12;
```

5.2.6.6 Ingress

Off-air ingress is an RF impairment that occurs when signals external to the cable network find a means of getting into the plant and propagating with the intended signals. These signals are part of the same frequency domain that the FCC has allocated to other services and that operators use within the coaxial network. Ingress is an unwanted interaction between electronic systems, off-air and enclosed. The coaxial network is designed and built to minimize the effects of the off-air signals on the intended signals.

Off-air ingress is due to several causes, such as inadequate shielding at the receiver, damaged cable, node or amplifier RF shielding issues, loose ground connections, or poor coax connections, and can be from a number of sources. It often affects a relatively consistent band of spectrum, most of which are allocated and licensed by the FCC. The most common types of off-air ingress are described below.

- **FM Ingress**—FM ingress is the result of FM radio signals entering the plant from FM radio broadcasting sources. It is easily recognizable by the pattern it creates. It affects downstream frequencies between approximately 88 and 108 MHz because the FCC allocates a 200-kHz channel bandwidth for this service, and it often exhibits spikes in spectral power at odd intervals of tenths of a megahertz (e.g., 90.1, 90.3, 90.5 MHz) because the center frequency is located at half the channel bandwidth of the FM channel. Frequencies of concern differ throughout the world according to the FM frequencies used.
- **Cellular Ingress**—Cellular ingress is the result of cellular devices and base station transmissions leaking into the cable plant. Most often, cellular interference exists at 700 MHz and above.
- **VHF Ingress**—VHF ingress covers any ingress not classified as FM ingress in the range 90 to 300 MHz. Interference in these bands can arise from numerous sources, including broadcast television, radio location, amateur radio, land mobile, marine radio, aeronautical navigation and communications, and military communications.

- **UHF Ingress**—UHF ingress covers any ingress not classified as cellular ingress in the range 300 MHz to 3 GHz. Interference in these bands can arise from numerous sources, including broadcast television, amateur radio, land mobile, mobile satellite, aircraft navigation and communications, and military communications.

As shown above, there are numerous sources of ingress into a cable system, and accordingly, their spectral profiles can be quite complicated to characterize, identify, and fix. Ingress introduced early in the downstream path has a much higher probability of having negative customer impact, depending on the frequency, bandwidth, and amplitude of the ingress signal. If the plant fault is significant, the ingress can severely impair the downstream signal to the CPE.

Ingress can also sit close to the noise floor under the CTA channels. Such cases may be seen in the guard bands between channels as an observable increase in the noise or a sharp spike in the captured spectrum. Measuring wide band ingress under the CTA channels is difficult and highly dependent on the resolution at which the spectral response is captured. More study is required to determine if the 30-kHz resolution bandwidth will make capturing ingress in the guard band feasible.

In some systems, the power of the ingress signals creates interference persistent enough to leave certain CTAs impaired or without services. No signal is transported in these channels, based on historical data, and they are labeled as "poison channels." When a local solution is found and these channels demonstrate a clear spectral response, these channels can be loaded, resulting in the addition of useable spectrum for additional services.

Ingress can be sporadic in nature and is frequency independent. An ingress signal at 100 MHz does not mean that a 700 MHz signal will be observed; the opposite is also true. If the conditions in the plant allow an ingress signal at UHF channel 21, 513.25 MHz, then the spectrum around 100 MHz and 700 MHz may not show ingress as an issue. Alternatively, as can be seen in the examples of ingress in Figure 29, ingress may be seen throughout the spectrum. The second illustration in the figure shows the potential of using the guard bands to evaluate potential ingress issues. Figure 30 shows a significant ingress event with multiple issues, the ingress in the FM band being the most evident. It also shows periodicity issues that may clear up when the source of the ingress is found.

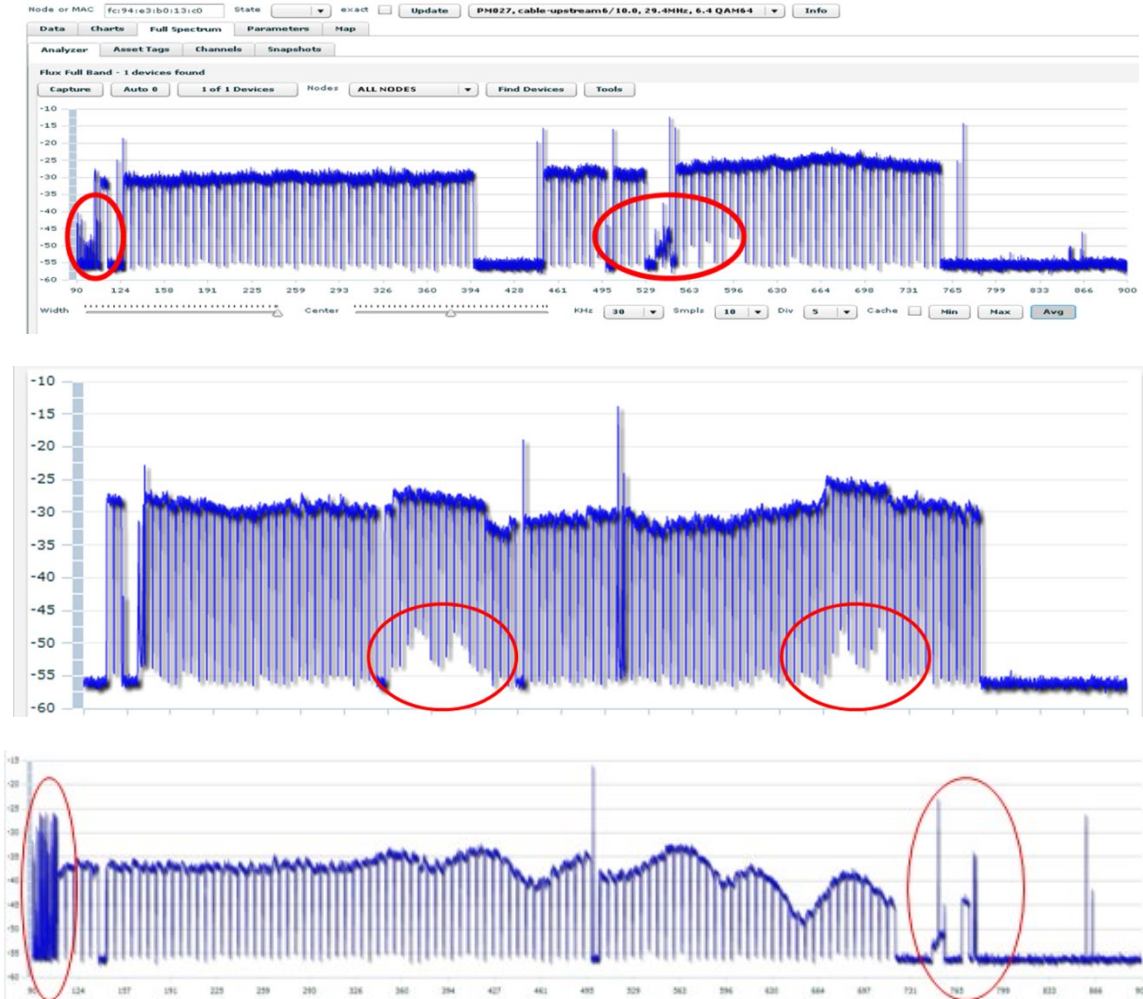


Figure 29 - Examples of Ingress




Figure 30 - Example of a Significant Ingress Event with Several Types of Ingress

5.2.6.6.1 Noise Thresholds

The default adjacency detection logic in SID is contained in the NoiseIdentifier code within the ident package of the software library. The following values are used within the analyzer and may be adjusted to accommodate differences in RF environments or operational standards. The code comments denoted by "//" are used to describe each value.

```
public static final double ANALOG_DB_THRESHOLD = -20;
```

5.2.6.7 Resonant Peaking

Resonant peaking is an RF spectrum impairment that can affect a series of RF or CTA channels. It is easily recognized when viewing the full spectrum and observing the frequency response. Resonant peaking is observed as an upward change in amplitude from a linear or flat spectral frequency response that peaks and then reduces in amplitude.

More specifically, it is characterized as an upward perturbation of the spectrum with convex and sinusoidal boundaries. It can be caused by one or more faulty network elements along the delivery path. Typically, the issue resides in an active device such as a node or amplifier. If multiple amplifiers have a spectrum signature with a peak at the same frequency, they become additive, causing the peak to increase in amplitude and the associated slopes to expand in bandwidth. A faulty passive element in the plant can exhibit what appears to be a peak.

It is important to note that resonant effects can be sporadic in nature. They may be visible for a while at a particular frequency but later disappear and reappear. As such, they can be complicated to diagnose. The peaking can also be temperature dependent in that temperature variability can change the effectiveness of the grounding or shielding or the quality of the connection.

The location in the delivery path where the resonant peak is introduced will determine the overall impact. The peak's relationship to the control carriers can impact the performance of the other channels being distributed.

- If the peak happens at the amplifier gain control carrier, the amplifier or series of amplifiers will react on a decibel-for-decibel basis to impact adjacent channels or the full spectrum.
- If the peak happens away from the control carrier, the impact will be observed in the CSO, CTB, and RxMER performance.

This impairment can manifest at the CPE as poor RxMER, poor codeword performance, and packet loss, which could result in slow data transfer speeds, tiling, or freezing, depending on which channels are affected. The following figures show examples of the effects of resonant peaking. In Figure 31, it is seen at 716 MHz and has an impact on four channels on either side of the resonant peak.




Figure 31 - Example of Resonant Peaking Affecting Adjacent Channels

Observe the channels impacted by the resonant peak in Figure 32. If the change in amplitude across the channel is greater than 3 dB, the channel may start to show impairments. In this example, the change from normal to the peak is approximately 17 dB, resulting in a performance impact that may be seen across the spectrum.




Figure 32 - Example of Resonant Peaking with High Change in Amplitude

In the example in Figure 33, the upper resonant peak is too high in amplitude for the measurement window, more than -15 dBmV/30 kHz, giving the spectrum a flat-top appearance. The delta between the top trace and the lower traces is approximately 18 dB. Observing the traces shows that a common device is causing the resonant peaking. As shown in Figure 34, this spectrum has additional alignment issues beyond the resonant peaking that need to be corrected.




Figure 33 - Example of Resonant Peaking that Exceeds the Measurement Window




Figure 34 - Additional Alignment Issues in the Spectrum from Figure 33

5.2.6.8 Roll Off

Spectral roll off is characterized by a gradual, non-linear decrease in amplitude and power over the frequency band, in frequency order. It is most often found towards the high end of the frequency range, closer to 1 GHz. There are numerous reasons why roll off can appear in the spectrum.

1. Devices in the network have a pass band through which frequencies will pass with a minimal loss of signal strength. Typical pass bands used to design components in the past are 450, 550, 650, 750, and 870 MHz. In some cases, cable networks transmit data at frequencies slightly above these pass bands. If components designed for these pass bands are in the network, they may be identified by roll off of the spectrum above their designed pass bands as signals are added. Components with the same pass band characteristics will be shown as additive and will produce a sharper roll off. In these cases, the signal strength at high frequencies degrades at a much

more rapid pace than lower frequencies, resulting in spectral roll off than can affect the customer depending on how far the customer's device is from the point of transmission.

2. Individual elements along the transmission path, such as amplifiers and less commonly defective cable, can produce roll-off characteristics, especially when a signal must propagate through numerous amplifiers that are not configured correctly for the network. For example, equalizers in these nodes or amplifiers may have an upper limit of 750 MHz though the plant is trying to pass 870 MHz.
3. Network designs that have an equalizer placed mid-span between the amplifier and the CPE can have the same roll-off issue if they were designed for a lower frequency than what is trying to be passed.
4. Some networks may use older coaxial cable that was not designed for the extended frequencies in use today. As an example, the coax may have typical losses up to 650 MHz, but the greater losses would be anticipated at higher frequencies. There are other contributors to additional loss in a coaxial cable such as water migration damage caused by improper installation or a plant fault at a connector. Loss may appear as roll off or excess loss depending on where in the network it occurs.
5. In more limited and extreme cases, what appears to be roll off may be something else. A large suckout at the high end of the spectrum with only one side of the suckout visible appears to be a roll off upon a quick visual inspection. Similarly, a standing wave can sometimes be misidentified as a suckout or a roll off.

Spectral roll off may cause tiling or freezing of video channels. Depending on the slope of each individual channel in the roll off and the power level at the receiver, the quality of service may be very poor or nonexistent.

Anecdotally, field technicians have observed roll off as one of the more prevalent RF impediments on the plant. However, to date, no actual quantification of the prevalence of this impairment has been made. To that point, as the roll off decreases the channel power of each channel in the roll-off region, the likelihood of a poor customer experience, and eventually a corresponding negative customer service touch point, are correlated. This outcome, however, needs to be tested and confirmed with data.

In both examples shown in Figure 35, the roll off appears to start at approximately 750 MHz. The end of the roll off is not shown because of the lack of additional carriers above 780 MHz. The slope per channel above 760 MHz may be impacting the performance of those channels.




Figure 35 - Examples of Spectral Roll Off at the High End of the Spectrum

5.2.6.9 Standing Waves

Standing waves indicate RF impairments that will impact the entire spectrum. The plant impairments are most often caused by impedance mismatches somewhere between the node and the modem. In the downstream, when seen by a cable modem, energy reflects off of the second impairment, back to the first, and then some of that energy continues downstream, but delayed, causing an impact to the signal that appears as a standing wave in the spectrum data. When observed in an amplitude-vs-frequency display, the standing wave will appear as a periodic amplitude change or a sine wave.

The height of the peaks of the standing wave will vary. The greater the delta in megahertz between the peaks, the shorter the distance is to the anomaly. The user can examine the PNM spectrum display to determine the distance between mismatches; the number of megahertz between standing waves can be calculated by using one of these equations:

$$\text{Distance (feet)} = 492 \times Vp/f$$

$$\text{Distance (meters)} = 150 \times Vp/f$$

where

- Vp = velocity of propagation in decimal form (velocity factor) in the network cable, typically 87% (0.87), and
- f = frequency of the periodicity, standing wave frequency expressed in megahertz, estimated as the frequency difference between adjacent peaks or troughs.

A mismatch can be caused by any active or passive network element (e.g., amplifier, splitter) that connects to the coaxial cable that eventually leads into the customer's home. Operators can utilize data extracts such as those shown in Figure 36, to measure the exact prevalence and impact of this RF impairment. Using the data from multiple CPEs associated with a node can help determine if the standing wave event is isolated or affecting several customers.




Figure 36 - Example of Standing Waves

5.2.6.10 Tilt

Spectral tilt is an RF characteristic that may or may not be an impediment depending on several factors. Tilt is defined as the amplitude difference between signals at specified frequencies, typically the lowest to highest. In the network, signals at a lower frequency will have less loss than signals at a higher frequency.

Tilt is introduced into the spectrum at the node for an optical transport or at the headend for a coax distribution network. The purpose of introducing tilt is to compensate for the anticipated signal losses in the network. In a coax distribution network, tilt also adds the performance benefit of making a flat spectrum possible. When the tilt compensation is calculated correctly, the delta between signals at a termination device will be minimized and each signal will have the same amplitude as the termination device (CPE).

It is possible that the positive tilt introduced into the network will be seen as a negative tilt at the CPE farthest from the node or amplifier (see Figure 37), but a negative tilt should not cause issues if it is operating within the system design parameters.

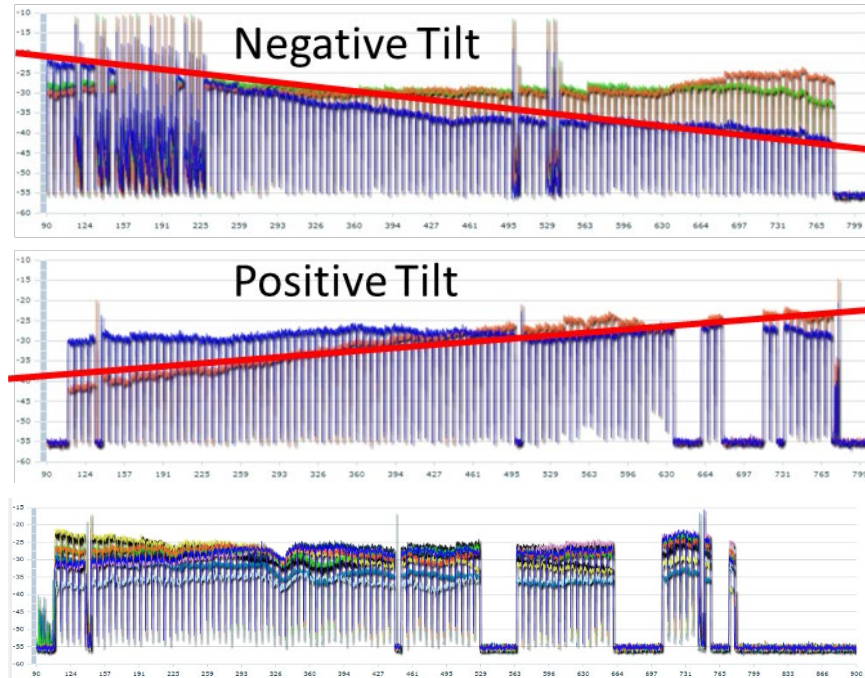


Figure 37 - Examples of Tilt: Negative Only, Positive Only, and Both Negative and Positive

It can be hypothesized that an operator is not likely to see a significant number of complaints from customers because of tilt alone, but that hypothesis needs to be tested and validated. Older CPE may have more problems with tilt than newer CPE, especially if the tilt is excessive and the total power across the spectrum is high.

Based on current DOCSIS specifications, North America does not have a maximum tilt specified for the input to the CPE. Instead, a dynamic range window is defined. The DOCSIS specifications do have a European tilt specified at 12 dB, either negative or positive.

5.2.6.10.1 Tilt Thresholds

The default adjacency detection logic in SID is contained in the TiltIdentifier code within the ident package of the software library. The following values are used within the analyzer and may be adjusted to accommodate differences in RF environments or operational standards. The code comments denoted by "//" are used to describe each value.

```
// MAX_SPECTRAL_TILT regardless of sign
private int MAX_SPECTRAL_TILT = 15;
```

5.2.6.11 Leveling

Leveling is characterized as a linear attenuation or increase of power across the downstream spectrum when compared to other devices in the related network. It is a comparative test to assist in determining alignment or balance issues within a node area. Leveling (see Figure 38) can be caused by one or more attenuation components in the path to the CPE, such as multiple splitters, wrong value directional couplers, added attenuators, or amplifiers set up incorrectly and affecting a series of devices. Because of the change in amplitude, the channel power may be attenuated to a level to cause freezing, tiling, lost packets, and voice issues in a customer's home.

Because tilt can occur naturally in the downstream path, it can be difficult to determine when leveling is an issue. The following questions should be asked when examining the issue.

- What is the minimum and maximum channel power in the spectrum, and does it fit within the input definitions for the CPE? For a DOCSIS network, that range is -15 to +15 dBmV.
- How does this CPE compare to another proximal CPE?

- How does this spectrum compare to the spectrum for a CPE related to a different leg of the associated node or an adjacent amplifier?

As can be seen by the questions, the detection method will apply general mathematical rules of attenuation based on the length of the cable or will compare the CPE to neighboring devices in the node area or cascade. This comparison is best made as an analytical exercise outside of the core Spectra channelization code and is better identified through a series of measurements of the general spectrum characteristics associated with CPEs fed by a common device. Based on past experience and anecdotal evidence, this characteristic is observed often in the plant and is not likely to affect a significant number of customers or drive much care-related activity.




Figure 38 - Examples of Leveling

One impairment caused by leveling is too much signal level received at the CPE. In the example in Figure 39, the top line shows a level to the CPE with power averaging about 10 dB higher than other CPE associated with the node.




Figure 39 - Example of Leveling Impairment: Increased Signal Level at CPE

The behavior shown in the example in Figure 40 can be attributed to leveling affecting multiple CPEs. Additional study would be required to verify if an amplifier has been set up incorrectly on the CPEs experiencing the negative tilt. Leveling is suspected because it is at the low end of the spectrum. Also, the two negative tilt lines are lower in level by about 5 dB than the average levels with a positive tilt. If a DOCSIS channel is placed close to 750 MHz on the black line (the lowest line at 750 MHz), the level would be too low after correcting for the bandwidth measurement RBW to channel bandwidth.




Figure 40 - Example of Possible Leveling Across Several CPEs

5.2.7 Background Theory

Figure 41 shows a block diagram of a digital spectrum analyzer that may reside in a cable modem or CMTS. The input signal enters at the left of the diagram; this signal is the full upstream or downstream band of the cable plant. An analog front end amplifies the signal and provides RF gain control. A high-speed analog-to-digital converter (ADC) provides digital samples of the signal. A digital tuner, consisting of a digital oscillator and a low-pass filter, selects the desired analysis band around a specified center frequency. The signal from the selected band is applied to the fast Fourier transform (FFT), which multiplies the signal by the discrete Fourier transform (DFT) matrix. Each bin of the FFT output comprises a complex value consisting of two numbers, real (I) and imaginary (Q), correlating the input signal with the particular frequency corresponding to a single row of the DFT matrix. Typically, a spectrum analyzer is only concerned with the magnitude, not the phase, of the FFT output. Therefore, the power (magnitude squared) of each bin is computed for each bin as $I^2 + Q^2$. If spectrum smoothing is to be applied, the previously described process is repeated with a fresh set of data from the same band, and the power values from several captures are averaged at each bin location. The smoothed bins are converted to decibels by taking $10 \cdot \log_{10}$ of each bin power value. These decibel values, one for each frequency bin, are then displayed as the spectrum of the input signal.




Figure 41 - Block Diagram of a Digital Spectrum Analyzer in a Cable Modem

Note that the tuner shown in Figure 41 is not necessary if the entire band can be processed as a single analysis band. However, if the band is being analyzed in segments, then the tuner is used to step through a sequence of analysis segments, and the individual spectrum segments are spliced together to produce the overall wideband spectrum.

5.2.8 How to Implement

Although well understood to be a PNM function, the downstream wideband spectrum analysis MIB is not defined in the [DOCS-PNM-MIB]. Instead, it is provided in [CM-OSSIV3.1], Section D.2.4, "CM Spectrum Analysis Objects." The spectrum analyzer control parameters, defined by Table 3, are used to instruct the capture engine. Most of the implementations have functional default values, so simply enabling the spectrum analyzer should result in usable information for testing the basic functions of TFTP and SNMP. In all cases, the measurements are configured and enabled by using SNMP with the following MIB objects.

Table 3 - CmSpectrumAnalysisCtrlCmd Object Attributes

Attribute Name	Type	Access	Type Constraints	Units	Default
Enable	Boolean	R/W			False
InactivityTimeout	UnsignedInt	R/W	0..86400	seconds	300
FirstSegmentCenterFrequency	UnsignedInt	R/W		Hz	93000000
LastSegmentCenterFrequency	UnsignedInt	R/W		Hz	993000000
SegmentFrequencySpan	UnsignedInt	R/W	1000000..900000000	Hz	7500000
NumBinsPerSegment	UnsignedShort	R/W	2..2048	bins-per-segment	256
EquivalentNoiseBandwidth	UnsignedShort	R/W	50..500	hundredths of bin spacing	150
WindowFunction	Enum	R/W	other(0), hann(1), blackmanHarris(2), rectangular(3), hamming(4), flatTop(5), gaussian(6), chebyshev(7)		
NumberOfAverages	UnsignedShort	R/W	1..1000		1
FileEnable	Boolean	R/W			False
MeasStatus	MeasStatusType	R/O			
FileName	AdminString	R/W	SIZE(1..255)		

5.2.8.1 FirstSegmentCenterFrequency

The FirstSegmentCenterFrequency attribute controls the center frequency of the first segment for the spectrum analysis measurement. The frequency bins for this segment lie symmetrically to the left and right of this center frequency. If the number of bins in a segment is odd, the segment center frequency lies directly on the center bin. If the number of bins in a segment is even, the segment center frequency lies halfway between two bins. Changing the value of this object may result in changes to the CmSpectrumAnalysisMeas object, as described in the description field for the object. Note that if this object is set to an invalid value, the device may return an error of inconsistentValue, or it may adjust the value of the object to the closest valid value.

5.2.8.2 LastSegmentCenterFrequency

The LastSegmentCenterFrequency attribute controls the center frequency of the last segment of the spectrum analysis measurement. The frequency bins for this segment lie symmetrically to the left and right of this center frequency. If the number of bins in a segment is odd, the segment center frequency lies directly on the center bin. If the number of bins in a segment is even, the segment center frequency lies halfway between two bins. The value of this object is typically equal to the value of the FirstSegmentCenterFrequency object plus an integer number of segment spans as determined by SegmentFrequencySpan. Changing the value of this object may result in changes to the CmSpectrumAnalysisMeas object, as described in the description field for the object. Note that if this attribute is set to an invalid value, the device may return an error of inconsistentValue, or it may adjust the value of the attribute to the closest valid value.

5.2.8.3 SegmentFrequencySpan

The SegmentFrequencySpan attribute controls the frequency span of each segment of the CmSpectrumAnalysisMeas object. If the attribute is set to a value of 0, then a default span will be chosen based on the hardware capabilities of the device. Segments are contiguous from the FirstSegmentCenterFrequency to the LastSegmentCenterFrequency, and the center frequency for each successive segment is incremented by the value of the SegmentFrequencySpan. The number of segments is equal to

$$\left[\frac{\text{LastSegmentCenterFrequency} - \text{FirstSegmentCenterFrequency}}{\text{SegmentFrequencySpan}} \right] + 1.$$

A segment is equivalent to an instance in the CmSpectrumAnalysisMeas object. The chosen SegmentFrequencySpan value affects the number of instances in the CmSpectrumAnalysisMeas object. A more granular value for SegmentFrequencySpan may adversely affect the amount of time needed to query the instances and possibly increase the acquisition time. Changing the value of this object may result in changes to the CmSpectrumAnalysisMeas object, as described in the description field for the object. Note that if this attribute is set to an invalid value, the device may return an error of inconsistentValue, or it may adjust the value of the attribute to the closest valid value.

5.2.8.4 NumBinsPerSegment

The NumBinsPerSegment attribute controls the number of bins collected by the measurement performed for each segment (instance) of the CmSpectrumAnalysisMeas object. Note that if this attribute is set to an invalid value, the device may return an error of inconsistentValue, or it may adjust the value of the attribute to the closest valid value.

5.2.8.5 EquivalentNoiseBandwidth

The EquivalentNoiseBandwidth attribute allows the user to request an equivalent noise bandwidth for the resolution bandwidth filter used in the spectrum analysis. It corresponds to the spectral width of the window function used when performing a discrete Fourier transform (DFT) for the analysis. The window function that corresponds to a value written to this attribute may be obtained by reading the value of the WindowFunction attribute. If an unsupported value is requested, the device may return an error of inconsistentValue, or it may choose the closest valid value to the one requested. If the closest value is chosen, then a subsequent read of this attribute will return the actual value in use.

5.2.8.6 WindowFunction

The WindowFunction attribute controls or indicates the windowing function that will be used when performing the DFT for the analysis. The WindowFunction and EquivalentNoiseBandwidth attributes are related. When a particular WindowFunction is selected, the EquivalentNoiseBandwidth attribute will report its value for that function. Alternatively, if an EquivalentNoiseBandwidth value is chosen and a WindowFunction representing that value is defined in the CM, then that value will be reported in the WindowFunction object, or a value of "other" will be reported. Use of "modern" windowing functions not yet defined will likely be reported as "other". Note that some devices may not support all window functions. An error will be returned if an attempt is made to set the attribute to an unsupported window function or if writing of the WindowFunction object is not supported by an implementation.

5.2.8.7 NumberOfAverages

The NumberOfAverages attribute controls the number of averages that will be performed on spectral bins. The average is computed by using the "leaky integrator" method, where reported bin value is equal to

$$[\alpha \times \text{accumulated bin values}] + [(1 - \alpha) \times \text{current bin value}].$$

Alpha is 1 minus the reciprocal of the number of averages. For example, if $N = 25$, then $\alpha = 0.96$. A value of 1 indicates no averaging. Rewriting the number of averages will restart the averaging process. If there are no accumulated values, then the accumulators are made equal to the first measured bin amplitudes. The number of averages will be set by writing the NumberOfAverages attribute. If an attempt is made to set the attribute to an unsupported number of averages, an error of inconsistentValue will be returned.

5.2.8.8 FileEnable

The FileEnable attribute, when set to "true", causes the CM to begin a Spectrum Analysis measurement with the parameters defined by the CmSpectrumAnalysisCtrlCmd set of attributes. In order to set FileEnable to "true", the Enable must already be set to "true". When the measurement is completed successfully, a file is generated and made available for transfer, and the MeasStatus attribute is set to "sampleReady". The file contains one complete snapshot of the spectrum data. Setting this object to a value of "false" instructs the CM to stop the measurement.

5.2.8.9 MeasStatus

The MeasStatus attribute is used to determine the status of the measurement. The PNM server will query this value to determine when the file is ready for transfer.

5.2.8.10 Filename

The Filename attribute is the name of the file at the CM that contains the spectrum analysis data and that is to be downloaded by the PNM server.

For the most part, the only change for wideband spectrum analysis in the DOCSIS 3.1 specifications is TFTP output. In addition to using the SNMP interface to obtain test output, TFTP can be used to transfer larger amounts of information through files. As shown in Table 3, use of TFTP is achieved by setting a value for the FileName attribute in CmSpectrumAnalysisCtrlCmd and by setting the FileEnable attribute to "true". When the measurement is completed successfully, a file is generated and made available for transfer, and the MeasStatus attribute is set to "sampleReady." The resulting file will contain one complete snapshot of the spectrum data, including all of the information used during test operation and acquired as a result. Table 4 describes the structure of this file, including byte offsets for each of the discrete values reported by the test. See [CM-OSSIV3.1], Section D.2.4, "CM Spectrum Analysis Objects," for the MIB details for implementation.

Table 4 - Spectrum Analysis File Format

Element	Size
Size File type (value = 504E4D09)	4 bytes
Major Version (value = 1)	1 byte
Minor Version	1 byte
Capture Time	4 bytes
Channel ID	1 byte
CM MAC Address	6 bytes
FirstSegmentCenterFrequency	4 bytes
LastSegmentCenterFrequency	4 bytes
SegmentFrequencySpan	4 bytes
NumBinsPerSegment	2 bytes
EquivalentNoiseBandWidth	2 bytes
WindowFunction	2 bytes
Length (in bytes) of SpectrumAnalysis Data	4 bytes
SpectrumAnalysisData	BinAmplitudeFileData

5.2.9 Conclusion and Suggestions for Future Work

As previously stated, wideband spectrum analysis is an eminently useful feature of the PNM tool stack. The use of spectrum analyzers is commonplace in troubleshooting RF performance, so very little additional training is required for users that adopt it. Because virtually all DOCSIS 3.0 and 3.1 devices are compatible, many operators are able to claim a majority of their cable modem population. These factors create a compelling opportunity for operators to participate in PNM with a minimum investment and potentially significant improvement to their operations.

One of the aforementioned limitations of wideband spectrum analysis is that it is currently subject to the diplex filter's rejection area. Although certain inferences can be made, this limitation precludes the spectrum analyzer from observing upstream noise, which would be useful for locating sources of upstream noise ingress. Some implementation options would facilitate upstream spectrum capture and analysis using the wideband capture engine, such as additional ADC and FFT engines, bypass switches, and software timing with the upstream burst transmitters. By using wideband spectrum analysis to detect and locate upstream noise sources, cable operators could see major improvements in the long-standing problem of upstream noise detection, location, and mitigation.

Another area of opportunity for wideband spectrum analysis is the inclusion of in-phase and quadrature (I and Q) measurements rather than log magnitude (dB). In summing the signal vectors ($I^2 + Q^2$) to achieve log magnitude, important information about the signal is lost. As a result, more sophisticated signal analysis methods such as RxMER and other digital signal processing techniques are difficult, inaccurate, or impossible. On the other hand, IQ

measurements would allow operators to perform RxMER measurements on non-DOCSIS channels such as legacy SC-QAM used for video. It would also be possible to use IFFT for time-domain analysis of the acquired signals rather than relying on minimum phase assumptions to synthesize the phase.

In addition to containing information on the enumerated impairment types, the SID software library is currently being updated to include a TDR-like functionality. The TDR-like functionality, described in [PNMP-3.0], Section 7.3, enables technicians to view the full spectrum response in the time domain, which easily translates to distances that they are accustomed to viewing in time domain reflectometer tools.

An anomaly detection capability based on machine learning methods is available as well, and you can learn more about that in a published IEEE paper entitled "Proactive Network Maintenance using Fast, Accurate Anomaly Localization and Classification on 1-D Data Series", [Zhu, Sundaresan, Rupe].

Finally, many contributors to this document also published a paper that expands on full band capture entitled "Full Band Capture Revisited," [Hranac, et al.].

5.2.10 Detailed Discussion of Wideband Spectrum Analysis

5.2.10.1 Windowing

Recall that the DFT matrix consists of rows of sines and cosines, with each row containing a whole number of cycles. If the input signal is a sine wave with a frequency exactly equal to one of the rows of the DFT matrix, it will correlate perfectly with that row and have zero correlation with the other rows. However, if the frequency of the input signal falls somewhere between two DFT rows, or "off bin," which is actually most likely, the signal will correlate slightly with all the DFT rows. This scenario will cause what is called "spectral leakage" wherein an off-bin continuous wave (CW) signal, instead of producing a single spike in the spectrum, produces a large number of spikes.

To solve the spectral leakage problem, a data-tapering window is often used. A window is a sequence with gradual reduction at the edges that is multiplied by the input signal before the signal is multiplied by the DFT matrix. Its purpose is to taper the ends of the input signal vector, providing a smooth transition to zero at the two ends. Tapering reduces spectral leakage and causes a CW signal to produce a compact spectral spike, which is typical when using an analog spectrum analyzer. However, the spectral spike for a CW signal with windowing is slightly wider than it would be without windowing, implying that windowing slightly degrades the resolution of the spectrum measurement. Figure 42 shows some typical window functions. The Hanning window is a popular window function. It has a raised-cosine shape that reaches 0 at both ends and rises smoothly to 1 in the center. The resolution bandwidth of a Hanning window is 1.5 times the FFT bin spacing, an example of the reduction in resolution due to windowing.




Figure 42 - Typical Window Functions

5.2.10.2 Method to Find a Time Response from an IFFT when Phase Data Are Not Available

FBC is used by cable modems and set-top boxes to provide magnitude-only spectral data about RF path conditions in a remote location, such as a home. In some cases, the downstream channels being monitored are digital channels, such as 64- or 256-QAM. In other cases, the signals are analog signals or noise and ingress.

The following method applies to blocks of QAM signals and is used to identify the existence of an echo tunnel that causes ripple in the frequency response.

1. Pick a block of averaged (smoothed) contiguous digital signals, as many as possible. For example, each 7.5-MHz block of frequency domain data may have 256 spectral components, and multiple blocks are pasted together to make a wide spectral response.
2. Extract samples from the lower band edge of the lowest QAM signal to the upper band edge of the highest QAM signal, and convert the values into linear values. Use these values as I (in-phase) components.
3. Use zeroes for all Q (quadrature) values.
4. If necessary, zero-pad the values to fill out a 2ⁿ IFFT transform, such as 16,384 or 4,096.
5. Optionally, a window should be applied to the data.
6. A frequency region with another signal, such as an analog RF carrier, or vacant band can be filled in with a straight line connecting the channel just above the vacant band to the channel just below the vacant band.
7. Perform an IFFT to put the data into the time domain.
8. Transformed data will be symmetrical because quadrature values were not provided. The image can be discarded.
9. A DC term will be present. Comb teeth will be present every 166.67 ns because of the notch between 6-MHz channels.
10. If there is an echo in the frequency response, there will be a ripple in the frequency domain (Figure 43). The ripple will linearly transform to an impulse located among the comb teeth (Figure 44). If the echo is an exact multiple of 166.67 ns, it cannot be observed as easily. The delay between the main impulse and echo is the round-trip time of an echo tunnel formed by two impedance mismatches, corrected for propagation velocity of cable. Because the shape of the teeth on the comb is known, the teeth can be removed by subtraction.

This method is valuable because the wide bandwidth of the multiple QAM signals makes for exceedingly accurate time resolution, allowing the cable operator to repair a buried cable by making a hole rather than digging a trench.

Other methods to remove effect of the notches between carriers are to interpolate over the notches or to equalize the magnitude response, but equalization cannot reach zero because of negligible energy in the notch.

This method could also work with analog spectrum analyzers. For example, GPIB or other interface technology supported by the analyzer could be used to extract the magnitude data. Note that this method cannot be used to detect group delay problems because no phase information is available.

The code for this method is available to NDA members in the CableLabs Spectrum Impairment Detector (SID) in the PNM repository.




Figure 43 - Ripples Indicating an Echo Tunnel; No Phase Data Available

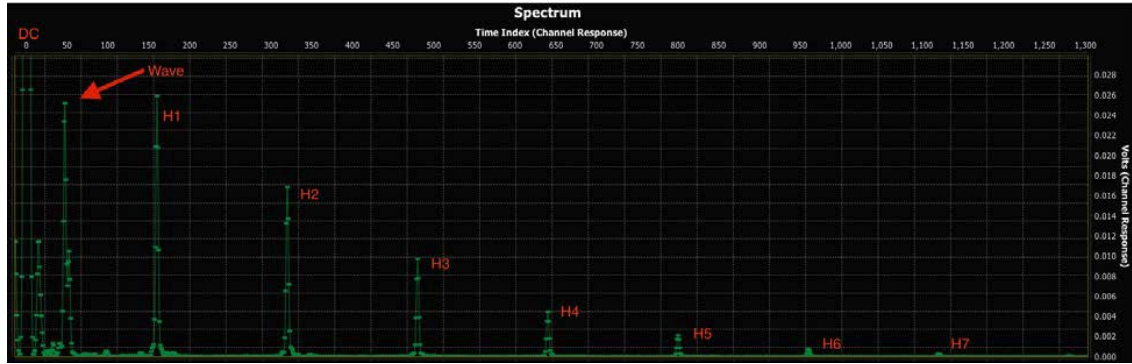


Figure 44 - Impulse Associated with Frequency Domain Ripple Among Comb Teeth (at Intervals of 166.67 ns)

5.2.10.3 CableLabs Time Domain Reflectometer (TDR)

As a cable plant ages, plant damage such as corrosion, animal chews, or stress fractures can cause impedance mismatches, which degrade signal quality. Two impedance mismatches form an echo tunnel that can be observed in a hub site for upstream signals or at a cable modem for downstream signals. A single reflection is reabsorbed and cannot be observed at either end, although signals are degraded (Figure 45). However, if a high impedance probe is connected to the center conductor of the coax, ripples in the downstream digital frequency response can reveal a problem. CableLabs has shown how a downstream signal with standing waves can be captured with a software-defined radio (SDR) and signal processing to form an approximate equivalent to a time domain reflectometer (TDR). This TDR, because of its passive nature, is not service disrupting and, because of the wide bandwidth capture, makes extremely accurate distance measurements.




Figure 45 - Diagram Showing Detection of a Single Reflection

CableLabs uses an inverse fast Fourier transform (IFFT) to convert the magnitude frequency response (Figure 46) into a time response (Figure 47). The method uses a math trick, using zero degrees (falsely) for phase value, allowing the inverse transform to be performed. The utility of the idea is enhanced by many sources of magnitude data, including FBC cable modem chips, analog spectrum analyzers, SDRs (Figure 48), field meters that are already deployed, and random noise generators.




Figure 46 - Digital Cable Signal Captured by Rapidly Retuning an SDR; Standing Wave Indicates a Reflected Signal Is Present




Figure 47 - Processed Signal Showing a Single Reflection, with Harmonics Caused by Roll Off of the 6-MHz Haystacks at Band Edges



Figure 48 - Cablelabs Engineer Making a TDR Measurement in the Field (the SDR Is in His Backpack)

The ideal location for a high-impedance test probe is on the output of an amplifier or fiber node, and the test point needs to show signals in both directions. Based on field testing, connecting a test probe in the middle of the echo cavity will make it difficult to find the impedance mismatches.

5.3 Downstream Receive Modulation Error Ratio (RxMER)

5.3.1 Overview

For many years, cable operators have relied on the RxMER of downstream SC-QAM signals to determine the health of those signals and the network. A QAM receiver computes a single value of RxMER for an SC-QAM signal, and that value is often compared to a company-defined threshold or to a value determined by good engineering practice.

RxMER is more useful with a DOCSIS 3.1 OFDM signal because it can comprise thousands of subcarriers, each of which is a narrow-bandwidth QAM signal. It would be unwieldy and impractical, however, to try to manage a list of RxMER values for thousands of subcarriers. Consider the tabular list in Figure 49, from an operational DOCSIS 3.1 cable modem. The far-left column is the subcarrier number in hex. The values to the right are RxMER values in increments of 0.25 dB, also in hex (two digits represent RxMER for a subcarrier). The zeros at the start and end are nulled for the taper regions. In practice, a list like this would have to be converted from hex to RxMER values (in decibels) to be useful. The original list includes 8,192 RxMER values (including the taper regions); several thousand values were removed to create this figure.

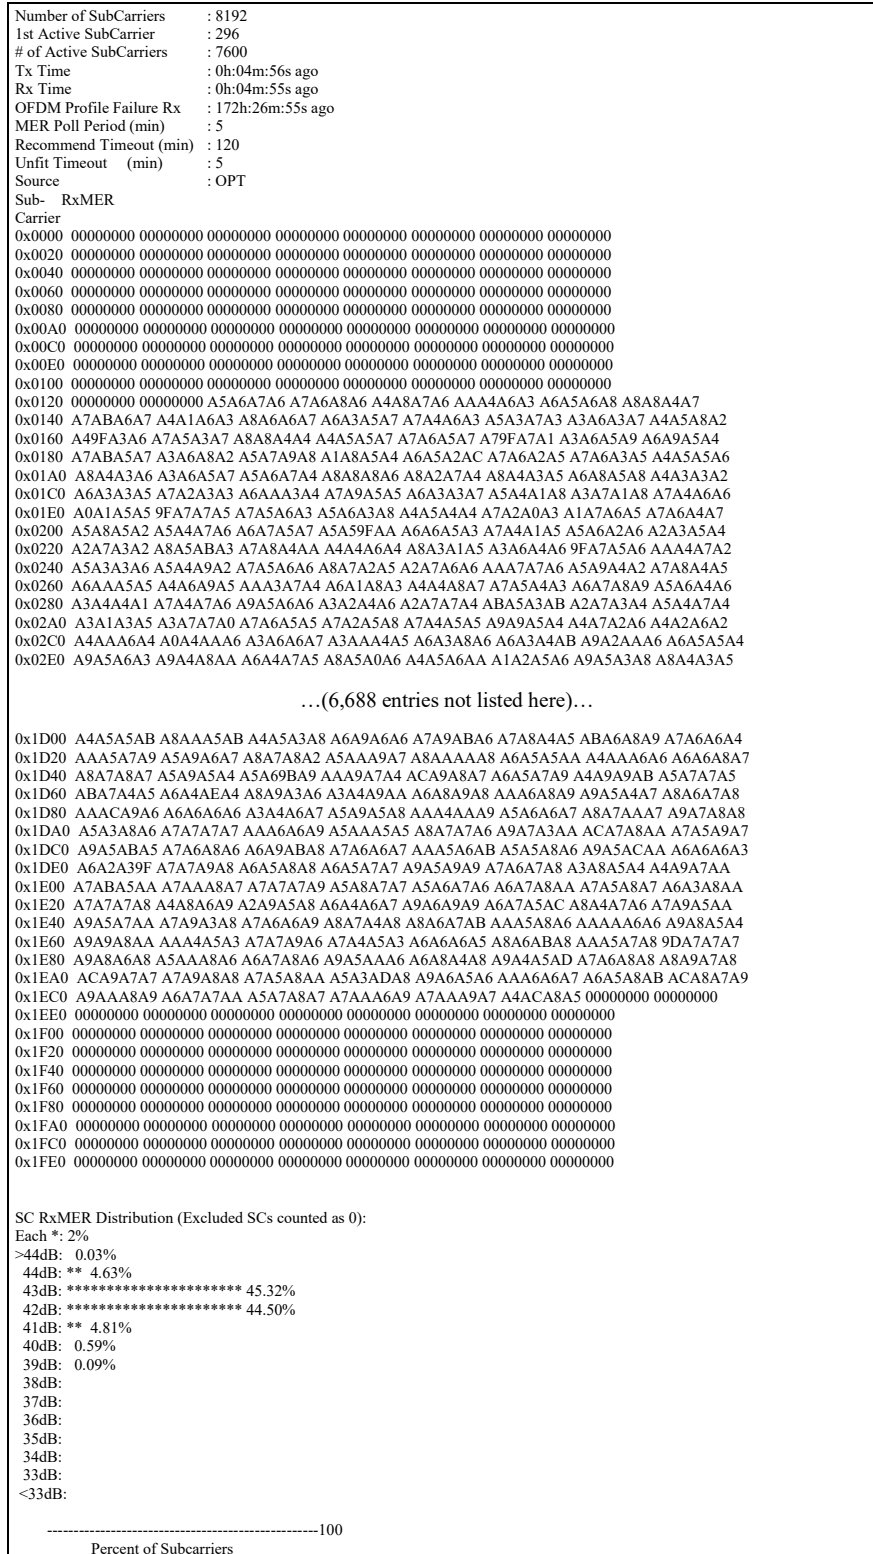


Figure 49 - List of RxCARRIER-per-Subcarrier Data from a DOCSIS 3.1 Cable Modem

Instead, it is much more convenient to plot the per-subcarrier RxMER on a graph (Figure 50), showing frequency or subcarrier numbers in the horizontal axis and RxMER in decibels in the vertical axis.




Figure 50 - Plot of RxMER-per-Subcarrier Data from a Field Meter with a DOCSIS 3.1 Cable Modem; Downstream OFDM Signal Is 96 MHz Wide and Average RxMER per Subcarrier Is 41.2 dB

Although the reported RxMER value for an SC-QAM signal does not indicate why the value is what it is, a plot of a downstream OFDM signal's RxMER per subcarrier, with its 25-kHz or 50-kHz frequency resolution, can provide important clues to the problems affecting an OFDM signal.

5.3.2 Description of Measurement

One reason the RxMER measurement is useful under a wide range of conditions is the way it is measured. It does not rely on data subcarriers, which are subject to symbol errors and cannot provide a reliable measurement of RxMER at low SNR or CNR. Rather, the cable modem measures the RxMER with pilots and PLC preamble symbols, which have known values regardless of SNR. The scattered pilots move across the OFDM band, repeating their scan every 128 OFDM symbols. They are boosted by 6 dB for robustness and visit all data subcarrier frequency locations throughout the OFDM signal. When a scattered pilot falls on a continuous pilot or PLC preamble symbol location, those values become known. Thus, the RxMER of all active subcarriers across the whole OFDM band is periodically measured over time. The cable modem indicates any subcarrier locations (such as exclusion bands) that it cannot measure, or "holes," in the measurement data.

For this measurement, RxMER is defined as the ratio of the average power of the ideal QAM constellation (numerator) to the average error vector power (denominator), as shown in Figure 51. The error vector is the difference between the equalized received pilot or preamble value and the known correct pilot value or preamble value. With this definition, because the numerator is the power of the QAM constellation rather than the boosted pilot power, the RxMER measurement yields the true QAM RxMER even though the pilots are boosted by 6 dB relative to the data subcarriers. That is, the design takes the pilot boost into account and no further compensation of the measurement is necessary to remove the effect of the pilot boosting.

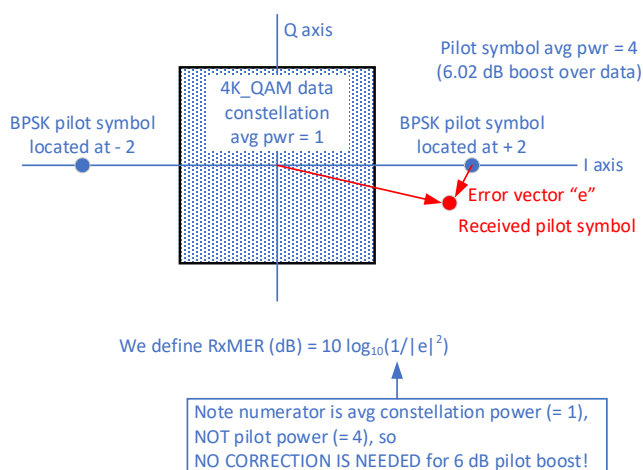


Figure 51 - Using Pilot Symbols to Compute RxMER

To make this definition clear, the following test case is defined. For an ideal AWGN channel, an OFDM signal containing a mix of QAM constellations, with data-subcarrier carrier-to-noise ratio (CNR) equal to 35 dB on the QAM subcarriers, will yield an RxMER measurement of nominally 35 dB averaged over all subcarrier locations.

RxMER may be more clearly defined in mathematical notation as shown in Figure 52, which shows an ideal transmit and receive model, with no intent to imply a particular implementation. The figure uses the following variables:

- p = scattered pilot (or PLC preamble) symbol before the transmit IFFT,
- n_1 = noise added before the channel,
- H = channel coefficient for a given subcarrier frequency (incorporating, for example, tilt and ripple of cable plant),
- n_2 = noise added after the channel, and
- y = unequalized received symbol after the receive FFT.

The receiver computes the estimate of H (H_{est}) and then the equalized received symbol (r) with the equation $r = y/H_{est}$. In the ideal case of perfect channel estimation, H_{est} equals H , so the receiver divides by H to effectively invert the channel and undo the channel effect. By using the known modulation value of the pilot or preamble symbol (p), the receiver computes the equalized error vector (e) as $e = r - p$. All of these quantities are complex scalars for a given subcarrier.

To compute RxMER, the receiver first computes E_{dB} with the equation

$$E_{dB} = 10 * \log_{10}(E)$$

where E equals the time average of $|e|^2$ over many visits (may be tens or hundreds) of the scattered pilot to the given subcarrier frequency (or PLC preamble symbol as applicable). Let S_{dB} equal the average power of an ideal QAM data subcarrier constellation (not including pilots) expressed in decibels. In DOCSIS 3.1 systems, all QAM constellations are defined to have the same average power. The modem then reports

$$RxMER_{dB} = S_{dB} - E_{dB}$$

and provides measurements of RxMER for all active subcarrier locations for the given OFDM downstream signal.

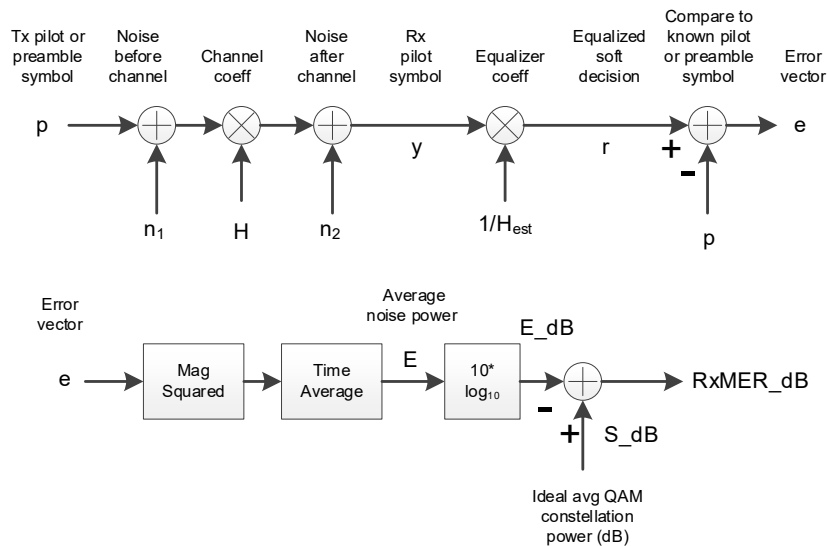


Figure 52 - Computation of RxMER for a Given Subcarrier

5.3.3 Background Theory

RxMER theory is discussed in [Hranac and Currivan]. RxMER for a given subcarrier is defined as the ratio of average QAM constellation symbol power to average constellation error power, expressed in decibels. Standard measures of

RxMER look at the demodulated complex baseband constellation symbols and measure their quality. The RxMER measurement gives the near "bottom line" status of the communications link because it is the demodulated symbols that produce correct bits, or bit errors, at the receiver output after processing by the FEC decoder.

In the case of DOCSIS 3.1 OFDM signals, RxMER per subcarrier is measured with pilots and PLC preamble symbols instead of data subcarriers to provide greater fidelity in the measurement. The pilots have the following advantages.

- Pilots have known values, so they are not subject to symbol errors. In a normal (non-pilot) QAM data symbol, if the noise is large enough to cause the soft decision to cross the boundary into the decision region of a neighboring symbol, the error vector will be measured incorrectly, resulting in an RxMER that is limited in range and cannot go below a certain value depending on the QAM constellation level. By using pilots, the measurement range is increased significantly so that subcarriers with very low RxMER values can be measured accurately. This is important because in an OFDM symbol, although the average RxMER may be within specification to allow error-free data transmission, individual subcarriers may have very low RxMER due to narrowband ingress or LTE transmissions.

On the other hand, when using pilots to measure RxMER, the following disadvantages should be considered.

- The level (6-dB boost) and constellation (BPSK) of pilots differ from that of data subcarriers, so one may question whether they provide a true representation of the RxMER of the actual QAM data subcarriers. This objection is believed to not have a practical significance, and for all intents and purposes, the pilot RxMER should very closely match the true RxMER of the data symbols (as explained in Section 5.3.2).

5.3.4 Data Presentation and Interpretation

5.3.4.1 Plot of Equalized and Unequalized Noise and Ingress Floor

Figure 53 shows a simulated example of RxMER per subcarrier for a 96 MHz-wide OFDM signal in which ingress from an idealized 10 MHz-wide LTE signal causes a reduction in RxMER of approximately 10 dB on the affected subcarriers.




Figure 53 - Plot of RxMER per Subcarrier Showing Impact of Simulated LTE Ingress on Affected Subcarriers

An inverted plot of RxMER versus subcarrier frequency (that is, -RxMER per subcarrier) shows the underlying noise (including ingress) in the channel. Figure 54 shows the same simulated OFDM signal with the RxMER plot inverted.




Figure 54 - Inverted Plot of RxMER per Subcarrier with LTE Ingress

The noise that is thus plotted is the equalized noise floor. That is, the signal and noise have both passed through the receive equalizer in the CM. The equalizer removes any linear channel effects such as ripple or slope from the signal, but it affects the noise floor as well. To visualize the unequalized noise—the noise as originally received from the channel before equalization—the RxMER and channel estimate can be combined as follows:

$$\text{unequalized_noise} = \text{RxMER} + \text{channel_estimate}$$

where

unequalized_noise = unequalized noise, ingress, etc., in decibels for each subcarrier,

RxMER = RxMER measurement in decibels for each subcarrier, and

channel_estimate = channel estimate measurement in decibels for each subcarrier.

Both the equalized noise floor and the unequalized noise floor are of value in PNM. The equalized noise floor is the bottom-line noise that the receiver is dealing with at the slicer, and the unequalized noise floor shows what is happening in the downstream channel.

Figure 55 and Figure 56 show examples of real in-channel ingress.




Figure 55 - Graph of OFDM RxMER per Subcarrier Showing Ingress Interference at About 890 MHz




Figure 56 - Plot of Actual RxMER per Subcarrier Showing Examples of Various LTE Interference and an Exclusion Band at About 760–763 MHz

5.3.4.2 Why an RxMER Plot May Not Show Channel Ripple

Reflections in the cable plant, such as from broken or corroded connectors, can cause shaping such as ripple in the channel response. This shaping will be seen in the plot of channel estimate coefficients and in the spectrum plot, but it may or may not appear in the RxMER-per-subcarrier plot. If most of the underlying noise floor responsible for the RxMER value is added before the channel shaping is introduced, the shaping will not appear in the RxMER plot. However, if most or all of the noise floor is added after the channel shaping is introduced, such as if it were the result of the modem's noise figure, the shaping will appear in the RxMER plot.

To understand why this occurs, take the simple example of a single micro-reflection in the channel causing periodic ripples in the frequency response. The receive equalizer in the cable modem or test set receiver contains an adaptive equalizer that applies the inverse micro-reflection response to remove the channel effect.

Refer to the RxMER block diagram in Figure 52. The figure shows a single OFDM subcarrier (for example, a scattered pilot) being sent from the transmitter through the channel and processed by the receiver. (OFDM allows one subcarrier to be considered at a time because each subcarrier is narrow enough to not experience frequency-selective effects.) The ideal case is assumed to be when $H_{est} = H$, so the receiver divides by H to undo the channel effect.

The following equations are used to calculate the soft decision, r , and the error vector, e .

$$\begin{aligned} r &= (H(p + n_1) + n_2)/H \\ &= p + n_1 + n_2/H \\ e &= r - p \\ &= n_1 + n_2/H \end{aligned}$$

The last equation shows that the error vector, which is used to compute the RxMER, includes the pre-reflection noise, n_1 , unmodified by the channel ripple, but that the post-reflection noise, n_2 , has been modified by the channel ripple.

The following details three cases of noise relative to the micro-reflection: before, after, and both before and after.

Case 1 assumes AWGN with value SNR_1 added before the micro-reflection. The following equations apply.

$$\begin{aligned} n_2 &= 0 \\ e &= n_1 \\ RxMER &= S/\langle |e|^2 \rangle = S/\langle |n_1|^2 \rangle = SNR_1 \end{aligned}$$

The equation shows that the original flat SNR due to noise n_1 appears as unmodified by the channel response, that is, there is no ripple. Figure 57, Figure 58, and Figure 59 show plots of the spectrum, channel estimate, and RxMER per subcarrier, respectively, for a laboratory test of this case. Though the spectrum and channel response plots show the channel ripple, the RxMER plot has little or no noticeable ripple. The spectrum plot also shows the pilots, CW lines with 6-dB boost above the data subcarriers.

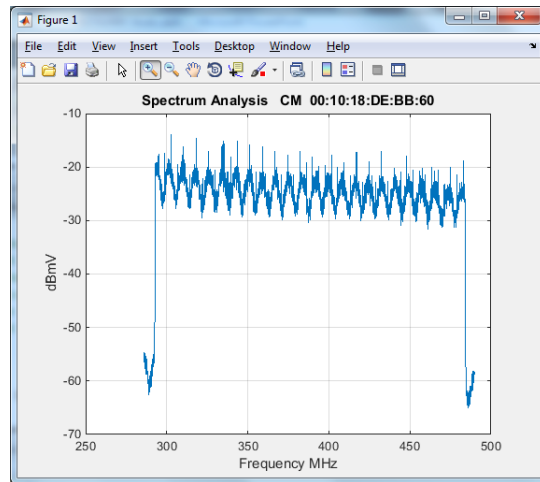


Figure 57 - Spectrum Plot with Ripple, Case 1, Noise Added Before Micro-Reflection

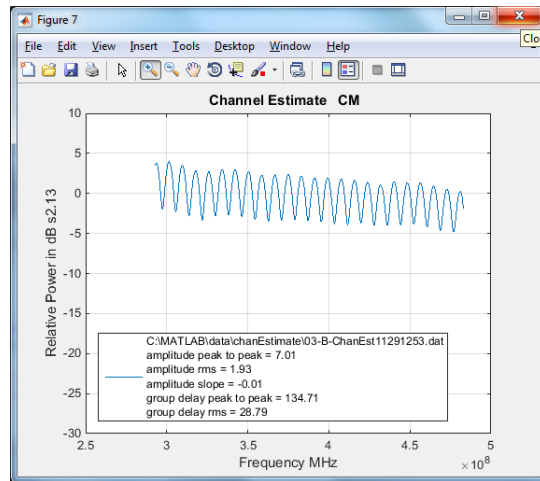


Figure 58 - Channel Estimate Plot with Ripple, Case 1, Noise Added Before Micro-Reflection

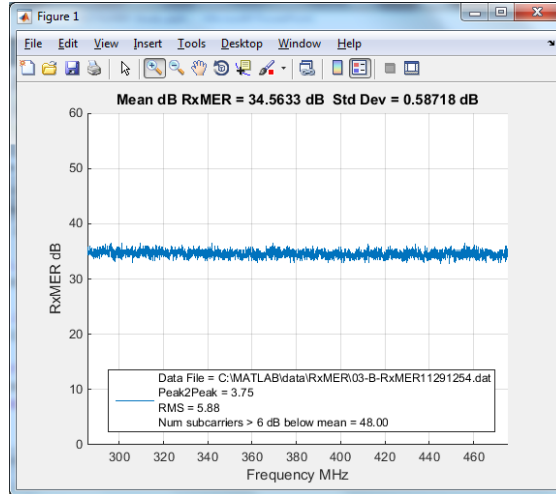


Figure 59 - RxMER-per-Subcarrier Plot with Little to No Ripple, Case 1, Noise Added Before Micro-Reflection

Case 2 assumes AWGN with value SNR_2 added *after* the micro-reflection. The following equations apply.

$$n_1 = 0$$

$$e = n_2/H$$

$$RxMER = S/\langle |e|^2 \rangle = |H|^2 S/\langle |n_2|^2 \rangle = |H|^2 \times SNR_2$$

The equation shows that the original SNR_2 due to noise n_2 has been modified by the squared magnitude of the channel response H . In other words, the RxMER will show the same ripple as the spectrum display and channel response plot, that is, there is a ripple. Figure 60, Figure 61, and Figure 62 show plots of the spectrum, channel estimate, and RxMER per subcarrier, for a laboratory test of this case. Note that all three plots for this case show the same uninverted ripple.

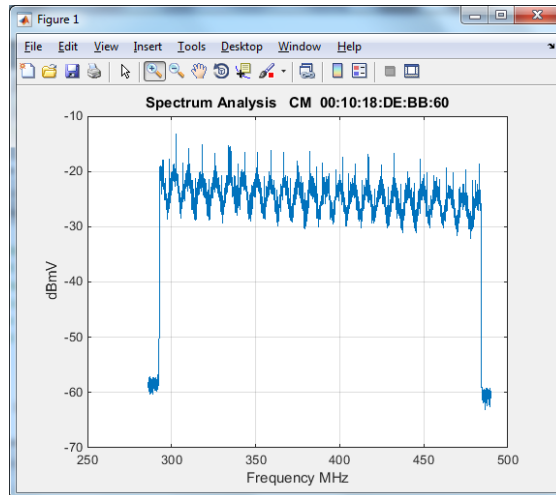


Figure 60 - Spectrum Plot with Ripple, Case 2, Noise Added After Micro-Reflection

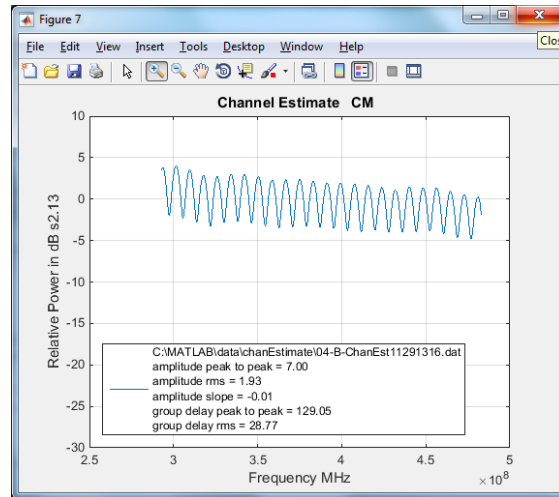


Figure 61 - Channel Estimate Plot with Ripple, Case 2, Noise Added After Micro-Reflection

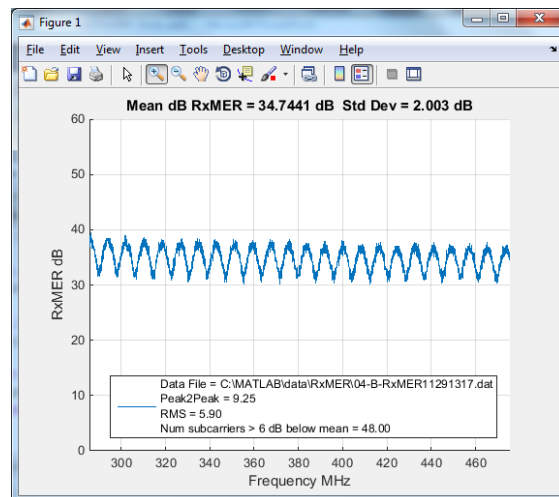


Figure 62 - RxMER-per-Subcarrier Plot with Ripple, Case 2, Noise Added After Micro-Reflection

Case 3 assumes noise is added both before and after the micro-reflection: AWGN with value SNR_1 is added before the micro-reflection and SNR_2 is added after the micro-reflection. In this intermediate case, the RxMER will show some ripple, depending on the relative values of SNR_1 and SNR_2 . If the receive signal level is low, the ripple may be hidden by the underlying noise floor.

The appearance of ripple in the RxMER-per-subcarrier plot can be used as a troubleshooting tool to augment amplitude ripple observed in a frequency domain (spectrum capture) display. The ripple indicates the presence of one or more micro-reflections. If multiple modems exhibit the same amplitude ripple signature, the PNM application can group modems with the same or similar response signatures to help locate the source of the problem. If one modem exhibits ripple and others (including those connected to the same tap) do not, the problem likely is in the subscriber drop. As is the case with amplitude ripple in the frequency domain, ripple in an RxMER-per-subcarrier plot can help determine the echo cavity length. The following formulas can be used to calculate the echo cavity length in feet or meters:

$$D_{feet} = 492 \times (VF/F_{MHz})$$

$$D_{meters} = 150 \times (VF/F_{MHz})$$

where

D_{feet} is the length of the echo cavity in feet,

D_{meters} is the length of the echo cavity in meters,

VF is the coaxial cable's velocity factor (velocity of propagation in decimal format), and

F_{MHz} is the amplitude ripple spacing (the frequency spacing between adjacent ripple peaks or adjacent ripple nulls) in megahertz.

For example, Figure 62 shows ripple in an RxMER per subcarrier plot; the ripple spacing is approximately 8 MHz. Assuming the cable's velocity factor is 0.87, the length of the echo cavity is 53.5 ft.

$$D_{feet} = 492 \times (0.87/8) = 53.5$$

5.3.4.3 Use of RxMER to Estimate Link Margin

The RxMER measurement can be used to determine how far from threshold the link is operating. Refer to Figure 50, which shows an average RxMER per subcarrier of 41.2 dB across the 96 MHz-wide OFDM signal. For example, if the modulation profile were 4096-QAM, one could infer about 6 dB of margin above the known ~35 dB RxMER threshold for that modulation order.

The values in Table 5 include approximately 5 to 6 dB of margin above the RxMER thresholds for the respective modulation orders. The data in the table reflect good engineering practice for operational RxMER values, but some cable operators choose to reduce the margin to a lower value, such as 3 dB. Though the respective modulation orders will generally work at lower margins, there is a risk that the lower margins could reduce reliability.

Table 5 - Recommended OFDM RxMER for Each Modulation Order

Modulation Order	RxMER
256-QAM	≥ 29~30 dB
512-QAM	≥ 31~33 dB
1024-QAM	≥ 34~36 dB
2048-QAM	≥ 37~39 dB
4096-QAM	≥ 40~42 dB

5.3.4.3.1 Computation of Average RxMER over All Subcarriers

The RxMER values in Table 5 are the average RxMER over all subcarriers in the OFDM symbol. Various methods have been proposed for this average: over decibel values, over noise values, over inverse noise values, and over BER.

5.3.4.4 RxMER Roll Off at Ends of OFDM Signal

Under some circumstances, the RxMER-per-subcarrier plot can exhibit roll off at one or both ends of the OFDM signal (see Figure 63). That roll off indicates a degradation of RxMER on the subcarriers closest to edges of the OFDM signal. The visibility of that RxMER edge roll off is related to the OFDM signal's taper region configuration, cyclic prefix setting, the presence of energy in the adjacent channel slot (e.g., SC-QAM signal), the total power at the modem input, and even things such as the make/model modem and silicon.




Figure 63 - Example of Edge Roll Off in RxMER per Subcarrier

Table 7 includes roll-off period samples (N_{rp}) versus taper region width in megahertz. Operators are cautioned to correctly configure the roll-off period setting when the channel(s) adjacent to the edge(s) of the OFDM signal will be occupied. It is important to understand that the N_{rp} setting impacts the OFDM signal's spectral edge sharpness and consequently the OFDM signal's transmitted fidelity. The latter can help prevent the OFDM signal from interfering with signals in the adjacent channel(s). The N_{rp} setting does not necessarily prevent signals (or other spectral energy) in the adjacent channel(s) from interfering with the OFDM signal.

5.3.5 Summary Metrics

The RxMER measurement provides high resolution (25 kHz or 50 kHz) of received SNR vs frequency, with up to 4,096 or even 8,192 measurements. However, the sheer number of data points can be unwieldy (see the tabular data in Figure 49), especially if the PNM server is rapidly polling many cable modems. To reduce the amount of data, the following summary metrics have been defined in [CM-OSSIV3.1].

5.3.5.1 Percentile

The Percentile attribute specifies the percentile (such as 2nd or 5th percentile) of all active subcarriers in an OFDM signal at which ThresholdRxMerValue occurs. For example, the 2nd percentile means that 98% of the subcarriers have an RxMER higher than ThresholdRxMerValue.

5.3.5.2 Mean

The Mean attribute is the mean of the decibel values of the RxMER measurements of all active subcarriers. The mean is computed directly on the decibel values as follows.

$$\text{Mean} = \text{sum of (RxMER decibel values)} / \text{number of RxMER values}$$

5.3.5.3 StdDev

The StdDev attribute is the standard deviation of the decibel values of the RxMER measurements of all active subcarriers. Good engineering practice suggests the standard deviation should be less than 2 dB, but 1 dB or less is better. The standard deviation is computed directly on the decibel values as follows.

$$\text{StdDev} = \text{sqrt}(\text{sum of (RxMER decibel values} - \text{RxMER}_{\text{mean}})^2 / \text{number of RxMER values})$$

5.3.5.4 ThresholdRxMerValue

The ThresholdRxMerValue attribute is the RxMER value corresponding to the specified percentile value. The modem sorts the subcarriers in ascending order of RxMER, resulting in a post-sorting subcarrier index ranging from 1 to the number of active subcarriers. If the percentile value corresponds to a non-integer post-sorting subcarrier index, the post-sorting index is truncated (floor function is applied); that is, the post-sorting index selected is the greatest integer less than or equal to the corresponding percentile value. For example, if there are 3,677 active subcarriers and the 2nd percentile is specified, the CM computes $\text{floor}(3,677 \times 0.02) = 73$. That is, the RxMER value of the 73rd subcarrier in the sorted list is associated with the 2nd percentile.

5.3.5.5 *ThresholdRxMerHighestFreq*

The *ThresholdRxMerHighestFreq* attribute is the frequency in hertz of the highest frequency subcarrier having an RxMER equal to *ThresholdRxMerValue*.

5.3.5.6 *Application of Summary Metrics*

The RxMER per subcarrier provides meaningful information regarding network health. When coupled with the active modulation profile per subcarrier, it becomes possible to determine the amount of operating margin for each subcarrier.

The DOCSIS 3.1 OSSI MIB providing percentile RxMER is also a good measure of network health. For example, showing that 2% of subcarriers have dropped below the threshold for the modulation profile of interest implies the channel is no longer in full capacity service and must drop down to a lower profile. When the RxMER falls below the acceptable threshold for a given modulation profile, the CM uses a lower order modulation, which results in lower capacity. Depending on channel utilization, the customer may experience degraded connection speed, particularly top-tier speed customers.

Practical application of the RxMER summary metrics is most easily explained by way of an example. The LDPC FEC algorithm shows that some number of bit errors can be successfully corrected.

The Percentile summary metric can be used to show how poorly the worst subcarriers within the OFDM channel are performing. If this metric is configured for the lowest 2% of subcarriers, a channel can be monitored for when becomes at risk of no longer delivering a reliable capability for a given modulation order.

As shown in Table 6, for reliable operation at 1024-QAM with RxMER equal across all subcarriers, empirical evidence shows that the RxMER must be above 30 dB. If the *ThresholdRxMerValue*, the RxMER for the lowest 2% of subcarriers, falls below 29 dB, uncorrectable errors begin to appear, which means 1024-QAM can no longer be supported by the relevant channel. The number of bit errors across the poor-quality subcarriers exceeds what the LDPC FEC can correct. Instead of looking at the RxMER of every individual subcarrier, the 2% threshold RxMER value can be monitored to confirm whether or not a specific modulation profile can be supported by the channel.

Table 6 - Reliability Threshold for Cable Modem RxMER at 2nd Percentile

DS Modulation	Metric	Pass	Marginal Pass	Fail
4096-QAM	2 nd Percentile RxMER	> 38	35 to 38	< 35
2048-QAM		> 35	32 to 35	< 32
1024-QAM		> 32	29 to 32	< 29
512-QAM		> 29	26 to 29	< 26
256-QAM		> 26	24 to 26	< 24

The *StdDev* value provides additional insight into the distribution of the RxMER values across all subcarriers in the OFDM channel. If the *StdDev* is a small number, the channel has a uniform RxMER across all subcarriers and can be described as well behaved. There is nothing odd in the frequency response, and no discrete ingress is corrupting a subset of the subcarriers. If the *StdDev* is a large number, a frequency response anomaly or significant ingress is likely. In either case, further investigation is warranted to find the cause and make a repair in the cable plant.

Rather than scanning FBC frequency response plots for signs of trouble, a simple filter on *ThresholdRxMerValue* and *StdDev* can be used to quickly identify nodes that require maintenance. The RxMER standard deviation of the subcarriers can be used to prioritize corrective/maintenance activities. For example, for two OFDM MAC domains that have the same failed RxMER 2nd percentile level, the MAC domain with the smaller RxMER standard deviation should be prioritized over the domain with the larger RxMER standard deviation. A smaller RxMER standard deviation results in more subcarriers closer to the failed 2nd percentile RxMER level (which therefore have overall poorer subcarrier performance) than a larger RxMER standard deviation.

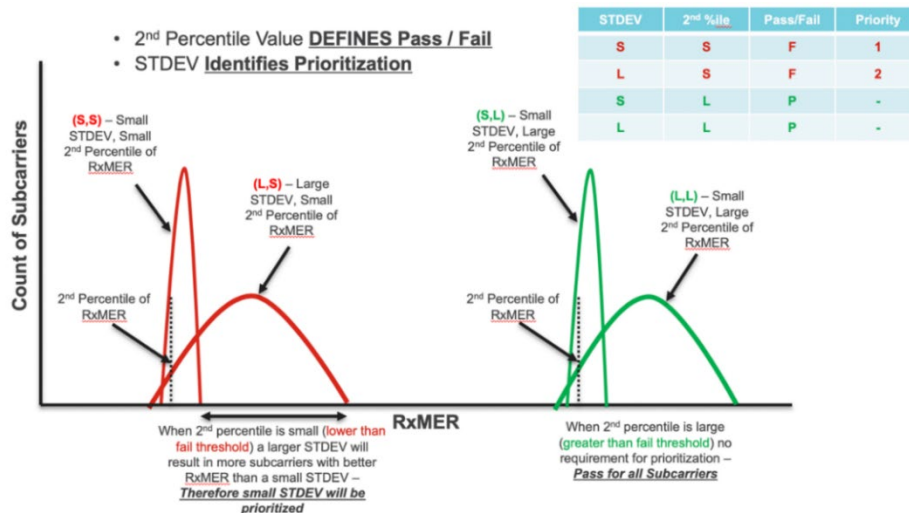


Figure 64 - RxMER Percentile Comparisons




Figure 65 - RxMER 2nd Percentile Versus RxMER Standard Deviation for a Population of Cable Modems

5.3.6 How to Implement

Methods for polling these data elements from cable modems are given in [CM-OSSIV3.1]. Additional details are available from CableLabs in the Common Collection Framework.

5.3.7 Conclusion

RxMER per subcarrier is one of the most useful downstream PNM measurements. It is a robust measurement with a wide dynamic range because of its use of pilots. It gives the health of each subcarrier, shows the underlying noise and ingress floor, provides information about the possible cause of problems, and gives the overall performance of the downstream channel.

Additional work in this area that the cable industry should undertake includes the following topics.

- Methods that use percentile values provided with downstream RxMER measurements
- How to calculate other metrics in addition to RxMER average

- Further study of skewness, a third moment of the distribution, as another possible summary metric that may help detect LTE and similar interference
- How and whether to present and summarize results in groups of 6 MHz, at a minimum, similar to DS power
- How to correlate RxMER vs time with FEC error history

5.4 Downstream FEC Statistics

5.4.1 Overview

This section describes the downstream OFDM FEC statistics and compares them to the downstream SC-QAM FEC statistics ([ITU-T J.83], Annex B) which are familiar in the industry.

5.4.2 Description of Measurement

The purpose of the downstream FEC statistics are to monitor downstream link quality via FEC and related statistics. The presence of bit errors reported through the FEC statistics can be caused by a number of factors:

- sweep transmitter interference,
- amplifier overload,
- laser clipping,
- impulsive noise or ingress interference,
- loose or intermittent connections,
- poor combiner or splitter isolation, and
- low signal-to-noise ratio.

Statistics are taken on FEC codeword error events, taking into account both the inner low-density parity check (LDPC) coding and the outer Bose, Ray-Chaudhuri, Hocquenghem (BCH) coding. Statistics are provided on each OFDM channel and for each profile being received by the CM. For example, if the CM is receiving four downstream profiles, there will be four sets of FEC counters plus a set of counters for the transition profile used for OFDM Downstream Profile Test (OPT). For profiles 1–4, statistics for data codewords include all codewords. For profile 5 (transition profile), statistics for data codewords include all codewords if codeword tagging is disabled; if codeword tagging is enabled, statistics include only codewords marked with T bit = 1 in the next codeword pointer (NCP). Similar statistics are taken on the NCP and PHY link channel (PLC).

For each profile, the modem will report

- uncorrectable codewords (number of codewords that failed BCH decoding),
- correctable codewords (number of codewords that failed initial parity check in LDPC decoding but passed BCH decoding), and
- total number of FEC codewords.

For the NCP, the modem will report

- NCP CRC failures (number of NCP fields that failed CRC check), and
- total number of NCP fields.

For the PLC, the modem will report

- unreliable PLC codewords (number of PLC codewords that failed LDPC decoding), and
- total number of PLC codewords.

For each OFDM channel, the modem will report FEC summaries for each profile over short and long time periods.

- Short-term summary statistics: number of uncorrectable and correctable codewords and total number each second for a rolling 10-minute period (600 values)
- Long-term summary statistics: number of uncorrectable and correctable codewords and total number in one-minute intervals for a rolling 24-hour period (1,440 values)

5.4.3 Background Theory

Unlike the Reed Solomon FEC (RS-FEC) used with SC-QAM (which is also included in DOCSIS 3.1 systems), OFDM in DOCSIS 3.1 technology introduced a new concatenated scheme using BCH (outer coding) and LDPC (inner coding) FEC (BCH-LDPC FEC). When transmitting, the CMTS first encodes the data payload using BCH and then encodes that using LDPC. When receiving, the cable modem first decodes the LDPC and then the BCH.

5.4.4 Content Sections

FEC statistics are maintained on a per-profile basis and provide insight into the performance of all active profiles. The modem maintains both short-term (last 10 minutes) and long-term (last 24 hours) statistics providing insight into profile performance over time. LDPC is a more robust error correction algorithm than Reed Solomon, so it is common for OFDM to have a high number of correctable codewords with a margin greater than 3 dB in SNR, unlike SC-QAM. If the number of uncorrectable codewords is high, then the profile is not performing well, resulting in retries and packet loss.

Unlike previous RS-FEC statistics, the presence of correctable codewords is acceptable with OFDM FEC and, in fact, it can be common to have close to 100% correctable codewords. In OFDM, the presence of uncorrectable codewords as high as 8 dB above the FEC threshold (10^{-6} codeword error probability) is the primary indicator of problems and that the modulation profile is not able to operate well given the network conditions.

Analyzing and understanding the per-profile FEC statistics is important to ensure both stability of the OFDM channel and that the highest achievable performance is being attained. The PLC and NCP operate at lower modulation profiles to ensure their robustness. It is critical that these profiles report no uncorrectable codewords because they are the foundation for the operation of the OFDM channel. Profile A (boot profile) likewise should have few or no uncorrectable codewords to ensure a usable base profile is available. The remaining profiles may show increasing numbers of uncorrectable codewords, and their presence will affect throughput and are indications of cable plant impairments that are limiting the ability of the OFDM channel to operate at its highest performance.

An example of OFDM FEC statistics over a short time period with a healthy plant is shown in Figure 66.

PROFILE	LOCKED	CWE (Corr)	CWE (Uncorr)
A	YES	1.5e-2	0.0
B	YES	3.4e-1	0.0
NCP	YES	0.0	0.0
PLC	YES	0.0	0.0

Figure 66 - Screen Capture from Field Meter Showing Performance of a DOCSIS 3.1 OFDM Signal

Parameters displayed include the ratio of correctable codewords and uncorrectable codewords each to the total number of codewords, as well as profiles, NPC, and PLC lock. Note that the ratio of correctable codewords is relatively high at 0.34, but the uncorrectable codewords are zero, indicating healthy performance.

5.4.5 Why High Correctables Are Acceptable in Downstream OFDM

PNM statistics for downstream correctable codewords in OFDM versus SC-QAM show a significant difference in their behavior when operating in healthy links. DOCSIS 3.0 technology, as well as earlier versions, uses SC-QAM in the downstream, whereas DOCSIS 3.1 technology uses SC-QAM and introduces OFDM in the downstream. More importantly, FEC for the two different downstream modulations is completely different. The SC-QAM FEC correctable codeword statistics (usually shown as "pre-FEC" bit errors in most field test instruments) for healthy links show relatively small number of errors (per total number of codewords). In contrast, for OFDM, it is typical that healthy links have a ratio of correctable codewords (for data) compared to the total number of data codewords over

an interval of time of nearly 100%; in other words, even with high link margin, with OFDM there can be nearly 100% of pre-decoding (input or decoder) erred codewords, while the decoder corrects almost all of them (fewer than 10^{-6} failing to decode). This difference has raised questions in operations using OFDM. This section explains the reason for this difference and shows why the high number of correctable codewords for data in OFDM downstream (even near 100%) is not a reason for concern and is, in fact, expected in most cases.

The key difference in the FEC statistics reporting between SC-QAM and OFDM data is that the codeword statistics are over the entire decoding process in OFDM FEC but only for the outer code in the SC-QAM FEC.

5.4.5.1 Differences in Modulation Between DOCSIS 3.0 and DOCSIS 3.1 Downstream

DOCSIS 3.0 SC-QAM downstream modulation is typically 256-QAM (8 total bits per constellation symbol), but 64-QAM (6 bits per constellation symbol) is sometimes used. DOCSIS 3.1 downstream uses SC-QAM and OFDM. OFDM uses many different constellation densities for downstream, from BPSK to 4096-QAM or 4K-QAM (1 to 12 total bits per constellation symbol), and different constellation densities can be applied in a downstream profile for data. Different modulation density choices are specified for different portions of DOCSIS 3.1 downstream transmissions (NCP, PLC, and data). Although these differences between DOCSIS 3.0 and DOCSIS 3.1 downstream modulation are significant, they have (essentially) no contribution to the differences in the FEC characteristics, which have been observed between the SC-QAM and OFDM downstream signals.

5.4.5.2 Differences in FEC Between SC-QAM and OFDM Downstream

Both SC-QAM and OFDM (for data) use FEC, or "concatenated coding," but similarity ends there. This section addresses SC-QAM FEC as described in annex B of [ITU-T J.83], the North American standard for the cable industry. This section also addresses the OFDM FEC for data codewords but not for DOCSIS 3.1 coding classes NCP and PLC because they do not report "correctable codewords" (see [PHYv3.1], Section 9.3.7, "Downstream FEC Statistics").

In a concatenated code, the data are coded in a first encoding, the outer code, and the result is encoded in a second encoding, the inner code. In OFDM, the outer code is a BCH block code, and the inner code is an LDPC code (see [PHYv3.1], Section 7.5.4.2, "FEC Encoding"). In SC-QAM, the outer code is a Reed-Solomon block code, and the inner code is a convolutional code, also known as a Trellis code or Trellis-coded modulation (TCM). Convolutional codes do not consist of distinct codewords. The use of a convolutional code as the inner code of SC-QAM FEC, as opposed to a block code, is the underlying reason for the difference in reporting results for downstream correctable codewords between SC-QAM and OFDM.

5.4.5.3 Differences in FEC Correctable Codewords Between SC-QAM and OFDM Downstream

As mentioned, both SC-QAM and OFDM downstream data use concatenated coding, but that is the only similarity. At the receiver in a concatenated coded link, the inner code is decoded first, and the result of that decoding is passed to the outer code for the final decoding.

The CM provides the following metrics on OFDM data codewords.

- Uncorrectables: number of codewords that failed final, outer BCH decoding
- Correctables: number of codewords that failed pre-decoding LDPC syndrome check (a parity check to determine whether the LDPC block is error-free) and passed BCH decoding
- Total number of FEC codewords

The LDPC decoder performs several iterations of its parity or "syndrome" check during the decoding process. At the beginning of the first iteration, it performs the syndrome check on the raw data from the slicer. Because error correction has not yet been applied to these data, the "pre-decoding LDPC syndrome check" is almost guaranteed to contain at least one bit error somewhere in the 16,200-bit block unless the link margin is high enough (well over 6 dB) for the raw slicer data to be error-free over a whole block. The block length of the OFDM LDPC code is 16,200 bits (in comparison, the block length of the SC-QAM RS code is 896 bits). A larger block will have a proportionally greater chance of having at least one bit in error. Thus, the large LDPC block length enhances the probability that a given block will have an error before decoding.

The total codeword count for OFDM is the number of BCH codewords in a counting interval. The uncorrectable codeword count is the number of failed BCH codeword decodings that the receiver determined in the counting interval;


that is, those codewords failed the final, outer BCH code. The correctable codeword count is the number of LDPC codewords that were determined to contain errors prior to inner code decoding and to be correct after outer code (BCH) decoding. The correctable codeword count in OFDM is essentially the number of codewords that contained an error after the receiver slicer, without benefit of any decoding, but then were corrected by the combination of the inner and outer decoding.

By contrast, the inner code for SC-QAM is a convolutional code and, as such, does not contain distinct codewords. At the receiver, therefore, in reporting correctable codewords and total codewords, the SC-QAM receiver counts only the Reed-Solomon codewords of the outer code, i.e., the only codewords that exist. These outer code, Reed-Solomon codewords only exist in the receiver after the benefit of the inner Trellis decoding. Therefore, in SC-QAM, FEC and the inner code already provide a significant coding benefit prior to the counting of the outer code's block codewords. This difference is the key difference in the correctable codeword performance between the SC-QAM and OFDM data FEC statistics.

The previous results are further explained or interpreted as follows. In SC-QAM counting, the correctable codewords are those Reed-Solomon codewords that contain errors after the inner code decoding and are corrected by using the outer code (the Reed-Solomon decoding). In OFDM counting, the correctable codewords are the codewords that had errors after slicing but prior to any decoding and then are corrected by using both the inner and outer code decodings. The key reason for the difference between the correctable codeword statistics in OFDM and SC-QAM is that the corrected codewords in OFDM have some slicer errors but are correct after completing both inner and outer decoding, whereas the corrected codewords in SC-QAM have errors after the inner decoding but are then correct after the outer decoding. There is substantial coding gain already included in the process of determining a codeword that has "input" errors in SC-QAM "prior to decoding" because the counting is only counting the performance of the outer code, whereas in OFDM the counting process is based on the entire coding scheme.

5.4.5.4 Differences in Probability of Failed Decoding Between SC-QAM and OFDM

As mentioned, the key difference in FEC statistics reporting between SC-QAM and OFDM data is that the codeword statistics are over the entire decoding in OFDM and only for the outer code in SC-QAM. Additionally, FEC for OFDM is more powerful than FEC for SC-QAM because of advances in coding theory and practice over two decades. Performance of the overall coding schemes for the two standards is best measured by the "fail to decode" probabilities. As a basis of comparison, a link that is ideal with an AWGN assumption is generally also appropriate. The overall performance, in terms of failure to decode a codeword versus E_s/N_0 , differs for the two standards because of the advances in FEC for the DOCSIS 3.1 OFDM and the different constellation densities available in the two standards. Figure 67 shows an example of measured performance for failure to decode versus RxMER, which is a measure of E_s/N_0 at the receiver slicer, for SC-QAM 256-QAM, OFDM 256-QAM, and OFDM 4K-QAM.



Series 1 = DOCSIS downstream SC-QAM, 256-QAM
 Series 2 = DOCSIS OFDM, 256-QAM
 Series 3 = DOCSIS OFDM, 4K-QAM

Figure 67 - Comparison of Uncorrectable Codeword Percentage for SC-QAM and OFDM, Measured

5.4.5.5 FEC Statistics Reported in DOCSIS 3.0 and DOCSIS 3.1 Systems

From the OFDM PNM statistics that the CM is required to provide—(a) uncorrectables, (b) correctables, and (c) total number of FEC codewords—several key link performance values can be calculated. The number of codewords containing an error prior to attempting any decoding is the sum of (a) and (b). The number of correct codewords after decoding can be computed as the sum of correctable codewords (b) and codewords that did not have any errors after slicing and prior to inner decoding, which is obtained by subtracting the sum of (b) and (a) from (c). Thus, the number of correct codewords after decoding (those that had no errors prior to decoding and those that were corrected by the decoding) is equal to (c) – (a).

In SC-QAM, the downstream FEC statistics are not required by [PHYv3.0]; rather, they are reported as MIBs per [CM-OSSIV3.1]. The three values reported are (1) the number of codewords received without errors, (2) the number of codewords received with correctable errors, and (3) the total number of codewords. Note that the codewords being counted in each of the categories are after the decoding of the inner convolutional code and are the codewords passed into the outer code (Reed-Solomon) decoding. Also note that the number of codewords received without errors (1) is not part of OFDM PNM but is calculable from the three OFDM PNM reported statistics. The number of uncorrectable codewords, required to be reported in OFDM PNM, is not required to be directly reported by the SC-QAM MIBs, but it can be computed from the required reported values by subtracting the sum of (1) and (2) from (3). The number of codewords with errors prior to Reed-Solomon decoding is computed by subtracting (1) from (3).


Thus, the following four data points can be obtained in both sC-QAM and OFDM, either through directly reported information or by calculating from available data:

- uncorrectable codewords,
- codewords without errors prior to decoding (of the outer code, for SC-QAM),
- correctable codewords, and
- total codewords.

Recall that the counts pertain to only the outer code for SC-QAM and to the entire coding scheme for OFDM.

5.4.5.6 Performance for SC-QAM and OFDM Correctable Codewords Statistics

Performance for SC-QAM uncorrectable codewords is based on the performance of the outer code. Figure 68 shows the 256-QAM signal with AWGN at the receiver slicer relative to RxMER or E_s/N_0 . The curves indicate (a) the theoretical probability of error for a single 256-QAM symbol at the slicer, (b) the measured probability of codeword error into the Reed-Solomon decoder (after the Trellis decoding) versus RxMER, and (c) the uncorrectable codeword error rate (after Reed-Solomon decoding).




Series 1 = Theoretical symbol error rate of 256-QAM

Series 2 = Downstream probability of codeword error into outer decoder, with 256-QAM

Series 3 = Downstream probability of codeword error after outer decoder, with 256-QAM (uncorrectable codeword probability)

Figure 68 - Measured Performance of SC-QAM Downstream 256-QAM

Figure 69 shows measured performance for OFDM with 256-QAM versus RxMER or E_s/N_0 , particularly the probability of erred codeword after the slicer and prior to any decoding and the probability of failed decoding. Figure 70 shows the same OFDM FEC statistics performance with 4K-QAM.




Series 1 = Theoretical symbol error rate of 256-QAM

Series 2 = Downstream probability of codeword error into decoder, with 256-QAM

Series 3 = Downstream probability of codeword error after outer decoder, with 256-QAM, (uncorrectable codeword probability)

Figure 69 - Measured Performance of OFDM 256-QAM



Series 1 = Theoretical symbol error rate of 4K-QAM

Series 2 = Downstream probability of codeword error into decoder, with 4K-QAM

Series 3 = Downstream probability of codeword error after outer decoder, with 4K-QAM (uncorrectable codeword probability)

Figure 70 - Measured Performance of OFDM 4K-QAM

5.4.6 Conclusion and Suggestions for Future Work

The FEC statistics reporting for SC-QAM and OFDM is well understood. Additional measured performance curves can show expected fluctuations in healthy links. Also, measured performance curves in the presence of other impairments besides AWGN (if available) would be useful as examples to guide troubleshooting.

SC-QAM FEC statistics provide correctable counts that can indicate poor service health and plant impairments but OFDM FEC statistics do not, so correctable counts are not a reliable indication of poor service health for OFDM. Therefore, a solution must be developed that replaces this well-known but lost method and that allows technicians to conduct PNM in a familiar way.

The following items are identified for future work.

- Theory of LDPC, such as the history and the definition of a syndrome
- Practical implementation of LDPC decoders, such as iterations
- Comparison to SC-QAM FEC
- Field examples and screen shots
- A waterfall plot of RxMER (plot intensity) vs frequency (x) and time (y)

5.5 Downstream Channel Estimate Coefficients

5.5.1 Overview

This section describes the origin and use of the CM's downstream channel estimate coefficients for an OFDM channel.

5.5.2 Description of Measurement

The downstream channel estimate coefficient measurement is used by the CM to report its estimate of the downstream channel response. The downstream channel estimate coefficients can be used to compute ICFR, group delay, and phase response for the downstream.

The channel estimate data returned by the CM consist of a single complex value per subcarrier for a specified OFDM downstream channel. The channel response coefficients are expressed as 16-bit two's complement numbers using 2.13 format. The CM samples are scaled such that the average power of the samples is approximately 1, in order to avoid excessive clipping and quantization noise.

Summary metrics (slope, ripple, and mean) available via SNMP are defined in order to avoid having to send all coefficients with each request.

5.5.3 Background Theory

The OFDM channel estimate coefficients are analogous to adaptive equalizer tap data from an SC-QAM CM receiver, with the following differences.

- The OFDM channel estimate coefficients are expressed in the frequency domain: one coefficient per subcarrier, with spacing of 25 kHz or 50 kHz. By contrast, SC-QAM equalizer coefficients are in the time domain: one coefficient per equalizer tap, with typical tap spacing of one symbol period or about 190 ns.
- The OFDM case has many more coefficients: up to 3,800 or 7,600 coefficients covering up to 190 kHz of bandwidth. By contrast, a typical North American SC-QAM adaptive equalizer has 16 to 68 equalizer coefficients for a given channel covering 5.36 MHz.
- The OFDM channel estimate coefficients directly represent the frequency response of the downstream channel. By contrast, SC-QAM equalizer coefficients are intended to provide the inverse frequency response of the downstream channel, that is, to undo the effect of the channel.

The block diagram in Figure 71 shows how OFDM downstream channel estimate coefficients are used. The transmitter sends an OFDM symbol, consisting of up to 3,800 or 7,600 subcarriers, through the downstream channel to the CM receiver. Each subcarrier can be considered independently of the others, so Figure 71 shows only one subcarrier.

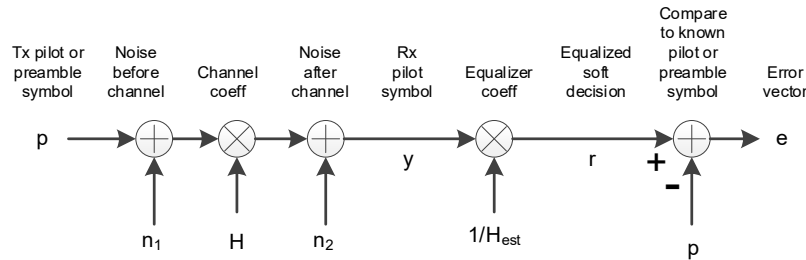


Figure 71 - Use of Downstream Channel Estimate Coefficients

The transmitter impresses a modulation value p onto the subcarrier and transmits it into the channel. Additive white Gaussian noise (AWGN) n_1 is added before the channel. The channel is represented by a multiplication by a number H , which can represent ripple, tilt, suckout, or another linear channel response. The value of H will depend on the frequency of the subcarrier. A second AWGN n_2 is added after the micro-reflection. The receiver applies its adaptive equalizer, which consists of division by the channel estimate coefficient. If the channel estimate were perfect, it would equal the actual channel value H , so the equalizer would perform a division by the channel value H as shown in the diagram. In reality, the CM cannot perfectly estimate the channel, so a value approximately equal to H is used. The result after equalization is the soft decision r . For a data subcarrier, r goes to the slicer and on to the forward error correction (FEC) decoder. For a pilot, the soft decision r is used to estimate the channel H as well as other receiver parameters such as timing and frequency errors.

For the number of subcarriers N , as large as 3,800 or 7,600, the process in the block diagram is repeated for each subcarrier, that is, from $k=1$ to N . The result is N values of the channel coefficient $H(k)$, one for each subcarrier k . The CM makes these coefficients available to the PNM server for analysis.

Thus, the downstream channel estimate coefficients report an estimate of the downstream channel response. The cable modem reports its downstream channel estimate for any requested OFDM channel.

5.5.4 Relationship of Channel Estimate Coefficients to RxMER per Subcarrier

The response of the downstream channel (including tilt, ripple, suckouts, etc.) may be seen in both the channel estimate coefficients and RxMER per subcarrier displays. See Section 6.3 for a discussion of how the channel response is observable by these two metrics and what the differences say about the channel.

5.5.5 Examples

The examples in this section were created in a controlled lab environment. Case 1 is a 70-dB notch at the edge of an OFDM channel, and Case 2 is a 500-ns echo at -10 dB.

Figure 72 and Figure 73 are plots of the magnitude of the channel estimate coefficients for Case 1 and Case 2, respectively. The plot for Case 2 shows a ripple across the band caused by the micro-reflection.




Figure 72 - Channel Estimate Coefficients for 70-dB Notch at Edge of OFDM Channel

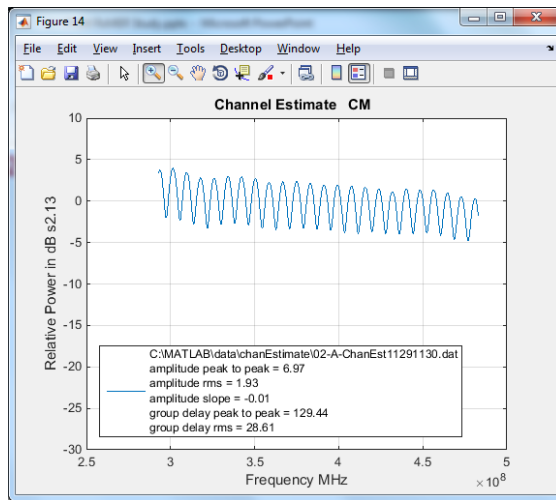


Figure 73 - Channel Estimate Coefficients for 500-ns Echo at -10 dB

Figure 74 and Figure 75 show the amplitude response and residual group delay, respectively of the channel estimate coefficients for Case 2.

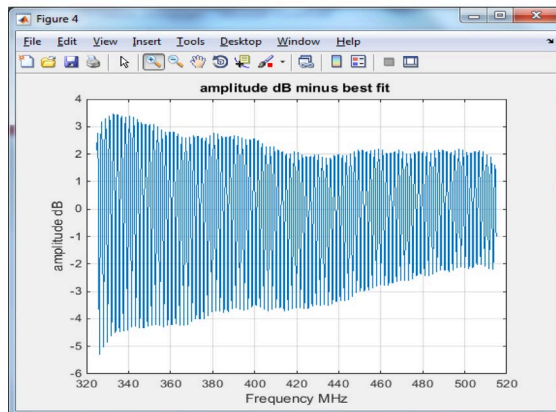


Figure 74 - Amplitude Response of Channel Estimate Coefficients for 500-ns Echo at -10 dB

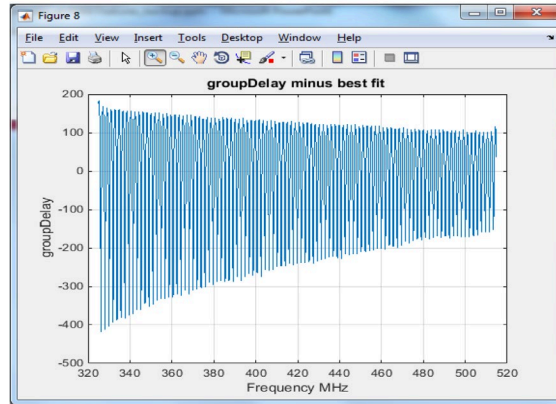


Figure 75 - Residual Group Delay of Channel Estimate Coefficients for 500-ns Echo at -10 dB

Figure 76 shows the channel estimate for Case 2 expressed as time-domain taps. The taps were produced by taking the magnitude of the IFFT of the channel estimate coefficients. To produce taps equivalent to SC-QAM equalizer taps, the channel estimate coefficients would be inverted before taking the IFFT; this inversion was not done in this example.

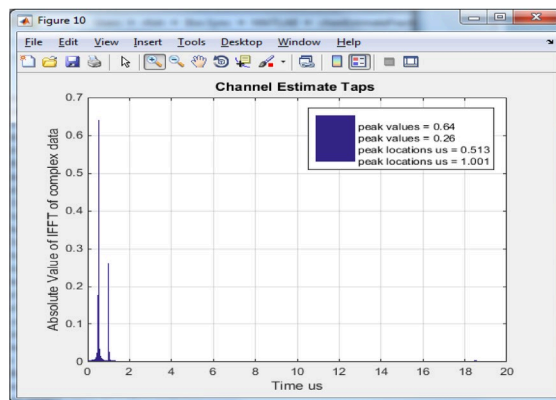


Figure 76 - Channel Estimate Expressed as Time-Domain Taps for 500-ns Echo at -10 dB

The channel response is measured from pilots, and the correction for elevated pilots is already included. The channel response is not a spectral plot, so it does not always agree with the spectral plot.

5.5.6 Tilt and Ripple Summary Metrics for Magnitude and Group Delay

This section describes the summary metrics of slope (tilt) and ripple for magnitude and group delay. Summary metrics are defined in order to avoid having to send all 4K/8K channel estimate coefficients on every query.

5.5.6.1 Tilt

The tilt attribute represents the tilt or slope, in decibels per megahertz, in the magnitude of the equalizer coefficients. The tilt is calculated as the slope of a linear least-squares fit of the frequency domain data. This attribute represents the tilt across the entire OFDM channel.

5.5.6.2 RMS Ripple

The RMS ripple attribute represents the value of the RMS (root mean square) ripple in the magnitude of the equalizer coefficients. After the tilt is subtracted from the frequency domain data, the ripple is calculated from the residual data. This attribute represents the ripple across the entire OFDM channel.

5.5.6.3 Peak-to-Peak Ripple

The peak-to-peak ripple attribute represents the value of the peak-to-peak ripple in the magnitude of the equalizer coefficients. After the tilt is subtracted from the frequency domain data, the ripple is calculated from the residual data. This attribute represents the ripple across the entire OFDM channel.

5.5.6.3.1 GroupDelayRipplePkToPk

The GroupDelayRipplePkToPk attribute represents the peak-to-peak Group Delay Ripple expressed in units of 0.001 ns. After the slope component calculated for the GroupDelaySlope is subtracted from the frequency domain data, the peak-to-peak ripple is calculated from the resulting data. This attribute represents the group delay variation across the entire OFDM channel.

5.5.6.3.2 GroupDelayRippleRMS

The GroupDelayRippleRMS attribute represents the RMS value of the Group Delay Ripple expressed in units of 0.001 ns. After the slope component calculated for the GroupDelaySlope is subtracted from the frequency domain data, the RMS ripple is calculated from the resulting data. This attribute represents the group delay variation across the entire OFDM channel. This attribute is not stored in the data file.

5.5.6.3.3 GroupDelaySlope

The GroupDelaySlope attribute represents the slope in 0.001 ns/MHz in the group delay of the equalizer coefficients. This attribute represents the slope across the entire OFDM channel.

5.5.6.3.4 GroupDelayMean

The GroupDelayMean attribute represents the mean of the group delay expressed in units of 0.001 ns.

5.5.7 Equations for Summary Metrics

This section describes the mathematical formulas used to compute the summary metrics of tilt, ripple, and group delay. For illustrations, the descriptions use example lab data from a DOCSIS 3.1 OFDM downstream receiver with a channel start frequency of 543.6 MHz and a bandwidth of 94 MHz. The channel was nominally AWGN with high SNR, although there is some slope and ripple resulting from the test setup.

5.5.7.1 Best-Fit Equations

A best-fit line to a set of data may be computed using the following equations taken from SCTE best practices ([SCTE Measurement], section 6.5, "Nominal Relative Carrier Power Levels and Carrier Level Variations," and Appendix). The equations are general and apply to both the magnitude and group delay responses.

$$m = \frac{\sum_{i=1}^N (x_i - \bar{x})(y_i - \bar{y})}{\sum_{i=1}^N (x_i - \bar{x})^2}$$

$$b = \bar{y} - m\bar{x}$$

where

$$\bar{x} = \frac{1}{N} \sum_{i=1}^N x_i,$$

$$\bar{y} = \frac{1}{N} \sum_{i=1}^N y_i, \text{ and}$$

\bar{x} and \bar{y} are the means of x and y, respectively.

The best-fit (f_i) line is calculated using the following equation.

$$y_{fi} = mx_i + b$$

The residual after removing the fit line (r_i) is calculated using the following equation.

$$y_{ri} = y_i - y_{fi}$$

The RMS ripple (R_{rms}) and peak-to-peak ripple (R_{pp}) are computed based on the residual values, as follows.

$$R_{rms} = \sqrt{\frac{1}{N} \sum_{i=1}^N y_{ri}^2}$$

$$R_{pp} = \max_{1 \leq i \leq N} y_{ri} - \min_{1 \leq i \leq N} y_{ri}$$

5.5.7.2 Magnitude Summary Metrics

A set of N complex channel estimate coefficients are provided in which the coefficient for each subcarrier (C) is calculated by the following equation.

$$C_i = I_i + jQ_i$$

For the magnitude case, the inputs to the best-fit equations are x , the list of N subcarrier frequencies in megahertz, and y , the list of N channel estimate coefficient magnitudes in decibels dB for each subcarrier, as follows.

$$y_i = 10 \log_{10}(I_i^2 + Q_i^2)$$

Figure 77 shows the log-magnitude-vs-frequency response of the channel estimate coefficients for the example lab data read out from the receiver.




Figure 77 - Frequency Response of Channel Estimate Coefficients, with Best-Fit Line in Red

The calculations are performed by using the above best-fit equations directly on the decibel values of y . The summary metrics are the slope m in decibels per megahertz, the RMS ripple R_{rms} in decibels, and the peak-to-peak ripple R_{pp} in decibels. The metrics for the above example data are as follows.

$$m = -0.0128 \text{ dB/MHz}$$

$$R_{rms} = 0.2658 \text{ dB}$$

$$R_{pp} = 1.2379 \text{ dB}$$

5.5.7.3 Group Delay Summary Metrics

Figure 78 shows the group-delay-vs-frequency response for the above example channel.




Figure 78 - Group Delay Response of Channel Estimate Coefficients, with Best-Fit Line in Red

Given the complex channel estimate coefficients

$$C_i = I_i + jQ_i$$

the group delay is computed as follows. First, the angle in radians of each channel coefficient is computed with the following equation.⁵

$$\phi_i = \arg C_i$$

This phase is unwrapped, i.e., it is not bounded over a single interval of length 2π , which prevents jumps at interval boundaries that could cause spikes in the group delay. Group delay is generally defined as follows.

$$\tau(\omega) = -\frac{d\phi(\omega)}{d\omega}$$

In this case, however, there is no continuous function of radian frequency ω because the channel estimate coefficients are defined at discrete frequencies corresponding to each active subcarrier. Group delay can be approximated by using phase and frequency differences from subcarrier to subcarrier in place of the derivative. This gives the following formula for group delay in nanoseconds, where Δf is the subcarrier spacing in hertz.

$$\tau_i = -\frac{\phi_i - \phi_{i-1}}{2\pi \cdot \Delta f} \cdot 10^9$$

Note that because this process uses differences between subcarriers, the value of N is decreased by 1, i.e., there is one less value of group delay than there are active subcarriers. The mean value of the group delay is dependent on receiver implementation and may be removed.

Having computed the group delay, the above best-fit equations are applied. The inputs to the best-fit equations are N , the number of active subcarriers minus 1; x , the list of frequencies in megahertz; and $y = \tau$, the list of group delay values in nanoseconds. The outputs are the slope m in nanoseconds per megahertz, the RMS ripple R_{rms} in nanoseconds, and the peak-to-peak ripple R_{pp} in nanoseconds. The summary group delay metrics for the example data are as follows.

⁵ The mathematical function “arg” returns the argument of the complex number, C in this case, providing a real polar angle, the angle between the positive real axis and the line connecting the point C to the origin. The function is also known as atan2.

$$m = -0.0012 \text{ ns/MHz}$$

$$R_{rms} = 4.0300 \text{ ns}$$

$$R_{pp} = 25.4990 \text{ ns}$$

5.5.8 Conclusion and Suggestions for Future Work

Field use of this measure is pending. Additional data from lab and field testing are needed to address this measure further. This document will be updated as use of this measure develops.

- Why this measure is useful and how it can be used
- How to set limit values for alarms based on summary metrics
- The format of the data and how to create graphs of the coefficients
- Examples, including a vector signal analyzer (VSA) capture with a notch
- How to use this data element as a TDR-like function, with wide spectrum results in finer time resolution
- The effect of exclusion bands on the estimated channel response and how to account for missing data

5.6 Downstream NPR Measurement

5.6.1 Overview

This section describes how to perform the equivalent of a noise power ratio (NPR) measurement by using an exclusion band within a DOCSIS 3.1 OFDM signal.

5.6.2 Description of Measurement

The purpose of downstream NPR measurement is to view the noise, interference, and intermodulation products that underly a portion of the OFDM signal. As part of its normal operation or in an out-of-service test, the CMTS or the operator can define an exclusion band of zero-valued (excluded) subcarriers that forms a spectral notch in the downstream OFDM signal for all profiles of a given downstream channel. The cable modem provides its normal spectral capture measurements, or symbol capture, which permit analysis of the notch depth.

In addition to using an exclusion band to identify an NPR measurement to characterize noise and intermodulation products caused by network alignment issues or defective equipment (as shown in Figure 79), another possible use case for an exclusion band is to observe LTE or similar ingress or direct pickup interference occurring within an OFDM band (see Figure 80).




Figure 79 - Use of Exclusion Band to Support NPR Measurement in OFDM Signal




Figure 80 - Use of Exclusion Band to See Ingress Interference Under an OFDM Signal

Because the introduction and removal of a notch (exclusion band) affects all profiles, causing possible link downtime, this measurement is intended for infrequent testing and maintenance. The width of the notch needs to accommodate the taper regions, intrinsic characteristics of an OFDM signal, on both sides of the notch in addition to the spectral band intended to observe the noise, interference, and intermodulation products underlying the OFDM signal.

5.6.3 Background Theory

The noise power ratio is a means of characterizing the linearity and noise contribution of an optical fiber link, an amplifier cascade, or individual active devices.

First, additive white Gaussian noise is applied to the input of the device or system under test (Figure 81). The noise is band-limited to the frequency spectrum of interest in the test; the band of interest in this example is 5 MHz to 42 MHz.




Figure 81 - Band-Limited Noise Added to the Input of a Device or System Under Test (Band of Interest Is 5 MHz to 42 MHz)

Next, a notch filter is added to the test setup to create a notch in the wideband noise (Figure 82).




Figure 82 - Notch Filter Added to Test Setup to Create a Notch in Band-Limited Noise Prior to the Device or System Under Test

At the output of the device or system under test, a spectrum analyzer measures the depth of the notch. The depth is reduced from its original value by the presence of intermodulation products or noise within the notch (Figure 83).




Figure 83 - NPR (Depth of the Notch in Decibels) Measured by Spectrum Analyzer at Output of the Device or System Under Test

The NPR is measured over a range of input levels to the device or system under test, and the results are plotted on a graph similar to Figure 84.




Figure 84 - Example of NPR Graph

The same type of measurement can be done using a downstream OFDM signal as the band-limited "noise." An exclusion band within the OFDM signal functions as the notch, and the depth of the notch can provide data similar to a conventional NPR test.

5.6.4 Measurement Description

When configuring an exclusion band within an OFDM signal to conduct an NPR measurement, the width of the taper regions on the edges of the exclusion band needs to be considered. The best place to measure the notch (exclusion band) depth is at or near its center, in order to avoid the effects of the taper regions.

First, determine the desired width of the notch measurement. The following guidance is given in [SCTE 119]:

The width of the notch at the depth of interest must be at least as wide as the resolution BW setting on the analyzer. For instance, to use the recommended RBW of 100 kHz, the notch should be at least 100 kHz wide at the depth of interest.

Practically speaking, a notch width of 1 to 2 MHz will be sufficient to accommodate most common spectrum analyzer resolution bandwidth (RBW) settings (e.g., 100 kHz, 300 kHz, etc.), but a wider width can be used.

Next, determine the OFDM taper region widths. The taper region width depends on the OFDM signal source's configured roll-off period samples (N_{rp}) (Table 7).

Table 7 - Taper Region Widths Based on Roll-Off Configurations (N_{rp})

FFT	Roll-Off Period Samples (N _{rp})	Taper Region (MHz)
4K	64	3.575
	128	1.845
	192	1.325
	256	0.975
8K	64	3.3375
	128	1.7125
	192	1.1625
	256	0.9875*

* The taper region of 0.9875 MHz is in accordance with the requirement for a minimum taper region of 1 MHz minus half subcarrier spacing. Achieving up to approximately 0.5 dB impact to the noise power in the adjacent spurious emissions integration region would allow a taper region of 0.8625 MHz, if the specification did not mandate the minimum taper region to be larger than this.

Table taken from [PHYv3.1], Appendix V, "CMTS Proposed Configuration Parameters (Informative)."

The following formula can be used to calculate the necessary minimum width of the exclusion band being used for an NPR measurement. Refer to Figure 85.

$$A = 2B + C$$

where

- A* is the exclusion band width in MHz,
- B* is the taper region width in MHz, and
- C* is the desired notch measurement width in MHz.




Figure 85 - Exclusion Band Width Including (B) Taper Regions and (C) NPR Measurement Region.

For example, a 50-kHz subcarrier spacing (4K FFT) and a configured N_{rp} of 256 results in a taper region width of 0.975 MHz. If the desired notch width for the NPR measurement is 2 MHz, the exclusion band's width should be at least 3.95 MHz.

$$\begin{aligned}
 A &= (2 \times 0.975 \text{ MHz}) + 2 \text{ MHz} \\
 A &= 1.95 + 2 \\
 A &= 3.95 \text{ MHz}
 \end{aligned}$$

The following figures show examples of exclusion band details.




Figure 86 - OFDM Signal, 192 MHz, Before Addition of Exclusion Band (Note PLC Near Left Edge of Signal)




Figure 87 - OFDM Signal from Figure 86 with 4 MHz-Wide Exclusion Band




Figure 88 - Close-up View of Exclusion Band in Figure 87; Combined Width of Each Taper Region (0.975 MHz) and Exclusion Band (4 MHz) Is Sufficient to Accommodate an NPR Measurement




Figure 89 - Example Exclusion Band in which Combined Width of Each Taper Region (1.875 MHz) and Exclusion Band (4 MHz) Is Not Adequate for a Reliable NPR Measurement

After a suitable exclusion band has been configured, the depth of the notch can be measured with a spectrum analyzer. The measurement can be done to yield an "as-is" NPR, or the amplitude of the OFDM signal (or all signals in the spectrum) can be varied through the device or network under test while measuring the depth of the notch at each signal level.

If the exclusion band is configured to look for LTE or other interference under the OFDM signal, the exclusion band's width can be configured to a wider value, such as 10 MHz. A wider exclusion band, however, will affect overall throughput more than a narrower width because more subcarriers are set to zero-value.

Whether for an NPR measurement or to look for interference under the OFDM signal, configuring an exclusion band as described here is intrusive if done on an in-service basis. As mentioned earlier, this procedure is intended for infrequent testing and maintenance.

5.6.5 Conclusion and Suggestions for Future Work

Noise power ratio is a means of characterizing the linearity or nonlinearity and noise contribution of an optical fiber link, an amplifier cascade, or individual active devices. This application of NPR uses an OFDM signal instead of the

wideband noise typically used in an NPR measurement and an exclusion band inside of the OFDM signal, representing a notch in the noise. During NPR testing with an OFDM signal, the depth of the notch (the exclusion band) is measured to obtain an equivalent NPR value. This measurement indicates which frequencies are impaired and at what levels, which can indicate the plant quality and help locate problems in the plant that should be resolved.

Pending future work by the industry, improvements to this section may include the following.

- Manual measurement, which is intrusive, and how to set it up
- The effect of roll-off length or taper region width on the sharpness of the notch or exclusion band edges; current guidance in [PHYv3.1], Appendix V, "CMTS Proposed Configuration Parameters", says to make a notch wide enough, with a minimum exclusion band of 3 MHz.
- Example captures showing a notch filling in with noise, for reference
- A demonstration of how to set up an exclusion band or zone by port, across all subcarriers, and whether removing the notch may cause a reboot of CMs
- How to create a new profile with zero subcarriers
- How to change the signal level and look for a change in the notch depth
- MIBs and other information for developers, operators, etc., to configure an exclusion band within the OFDM signal and use a CM's full band capture (or symbol capture) for analysis of the notch depth

5.7 Downstream Constellation Display

5.7.1 Overview

Constellation displays are useful for detecting a variety of impairments including noise, phase noise, and nonlinearity. Impairments can appear as a spreading of constellation symbol points, apparent rotation of the constellation display, non-square constellation shape (when IQ imbalance exists), rounding of the corners (symbol points appear to be moving inward), and so forth.

5.7.2 Description of Measurement

The purpose of the PNM measurement for downstream constellation display is to provide data on symbol spreading patterns, which may assist with identifying dominant impairments. It does so by capturing 8,192 symbols' worth of data at the CM and placing the data in memory for retrieval by the CmDsConstDispMeas MIB object. These data are processed by the PNM server to provide dominant impairment notifications, alarms, or trends to the operator.

5.7.3 Background Theory

Figure 90 shows a typical constellation plot for a mixed-modulation OFDM profile with high symbol density. The square fuzzy box is the constellation for the high-modulation orders used in the profile; the 16-QAM PLC (brighter yellow dots) and the boosted BPSK pilots are apparent. When dealing with the higher DOCSIS 3.1 modulation orders (such as 2K-QAM or 4K-QAM), a single constellation cannot simultaneously display all of the symbol points in a usable manner (see Figure 90). Therefore, constellation displays must be limited to lower modulation orders or to subsets of a larger modulation order ("zooming in" to part of the display). Large constellation subsets can be thought of as a zoom-in view.




Figure 90 - Downstream OFDM Constellation

5.7.4 Impaired Constellation Pattern Identification

For constellations with high symbol density, it is useful to zoom-in to symbol subsets of the constellation, which are most likely to be affected by impairments. The zoom-in display is limited to a subset of symbols, such as a 16-symbol cluster associated with the outermost corner of the larger constellation.

Impairment pattern matching can identify dominant impairments affecting the constellation, like nonlinear compression of a SC-QAM (Figure 91). A set of known dominant impairment patterns derived for OFDM signals can be used to facilitate likely impairment matches and ultimately diagnose dominant impairments of new constellation samples.




Figure 91 - Nonlinear SC-QAM 16-QAM Constellation Showing Evidence of Compression or Clipping

5.7.5 Conclusion and Suggestions for Future Work

A method for making large constellation displays useful for dominant impairment diagnostics has been proposed, i.e., by limiting it to 16-QAM or large constellation subsets.

In some cases, the constellation display can be supplemental information to RxMER per subcarrier because of their relationship (see Section 5.3).

Future work to be done with this measurement includes the following.

- Lab measurements and a thorough literature search are needed to define impairment matching algorithms based on representative symbol spreading data associated with common downstream impairments including linear distortion, nonlinearity, and phase noise
- How to create zoom-in views of various types in software plots of the data

- How to construct a constellation of pilots
- How to compute a constellation by using downstream symbol capture
- How to color-code each constellation point to visualize symbol errors and how to use downstream symbol capture to do so
- How to use downstream symbol capture to index a constellation point-to-subcarrier frequency
- How to gather data from a subset of frequencies

5.8 Downstream Symbol Capture

The purpose of downstream symbol capture is to provide partial functionality of a network analyzer to measure the linear and nonlinear response of the cable plant. The requirements are given in [PHYv3.1], Section 9.3.1, "Downstream Symbol Capture;" [MULPIv3.1], Section 6.5.5, "Trigger Message Block;" [CM-OSSIV3.1], Section D.2.5, "CmSymbolCapture;" and [CCAP-OSSIV3.1], Section 7.3.4.1, "DsOfdmSymbolCapture."

As of this writing, this measurement has been implemented in CMs but only in reference CMTS platforms. This section will be updated when field data are available.

Downstream symbol capture is potentially the most comprehensive of all downstream PNM measurements and provides capture of the same symbol at the CM and CMTS.

5.8.1 Overview

The basic principle of downstream symbol capture is to capture an OFDM symbol at the CMTS and the exact same symbol at one or more CMs. Because the OFDM downstream does not have a probe defined, as does the OFDMA upstream, an ordinary OFDM downstream data symbol is used. The modulation values of the designated OFDM symbol are captured at the transmitter so that it becomes a known signal serving the same function as a probe symbol. In practice, several (e.g., 10) symbols may be captured to allow the results to be averaged.

At the CMTS, the transmitted frequency-domain modulation values of one full OFDM symbol before the IFFT are captured and made available for analysis. These values include the I and Q modulation values of all subcarriers in the active bandwidth of the OFDM channel, including data subcarriers, pilots, PLC preamble symbols, and excluded subcarriers. The number of samples captured depends on the OFDM channel width, i.e., it is equal to bandwidth in kilohertz divided by subcarrier spacing in kilohertz. For example, for a 50-kHz subcarrier spacing in a 192-MHz channel with an active bandwidth of 190 MHz, 3,800 samples will be captured; for a 25-kHz subcarrier spacing in a 192-MHz channel with an active bandwidth of 190 MHz, 7,600 samples will be captured; and for a 25-kHz subcarrier spacing in a 24-MHz channel with an active bandwidth of 22 MHz, 880 samples will be captured.

At the CM, the received I and Q time-domain samples of one full OFDM symbol before the FFT, not including the guard interval, are captured and made available for analysis. This capture will result in a number of data points equal to the FFT length in use, which are approximately time-aligned for receiver FFT processing. The number of samples is equal to sampling rate in kilohertz divided by subcarrier spacing in kilohertz. For example, for a 50-kHz subcarrier spacing in a 192-MHz channel with a 204.8-MHz sampling rate, 4,096 samples will be captured; for a 25-kHz subcarrier spacing in a 192-MHz channel with a 204.8-MHz sampling rate, 8,192 samples will be captured. The capture includes a bit indicating if receiver windowing effects are present in the data.

The CMTS MAC provides signaling via the PLC Trigger Message to ensure that the same symbol is captured at the CMTS and CM. The Trigger Message Block (TR MB) provides a mechanism for synchronizing an event at the CMTS and one or more CMs. The CMTS inserts a TR MB into the PLC and indicating an action to be performed at a specific time aligned with the PLC frame. When any given CM detects the TR MB, it performs a capture at the same time relative to specified times aligned with the PLC frame received at the CM. The details of the TR MB are given in [MULPIv3.1], Section 6.5.5, "Trigger Message Block."

5.8.2 Description of Measurement

The CMTS captures the known frequency-domain modulation values of the downstream symbol selected for analysis, and the CM captures the received time-domain samples of the same symbol after it has passed through the

cable channel. As part of its analysis, the PNM software may compute the FFT of the CM capture so that the data from the transmitter and receiver will both be in the frequency domain for further processing.

5.8.2.1 Example Data Capture

The figures in this section show an example of downstream symbol capture.

Figure 92 shows the constellation (frequency-domain QAM values) of the designated downstream OFDM symbol captured at the transmitter. Because these values are known, the symbol may be used as a probe to measure the channel. The constellation shown includes 16-QAM data along with BPSK pilots.




Figure 92 - Constellation (Frequency-Domain QAM Values) of Downstream OFDM Symbol Captured at the Transmitter

Figure 93 shows the real part of the frequency-domain QAM values of the transmitted OFDM symbol. The x-axis represents frequency, i.e., the subcarrier index. Of the 4,096 QAM symbols, the zero-valued guard symbols appear at the beginning and end of the OFDM symbol.




Figure 93 - Real Part of Frequency-Domain QAM Values of Transmitted OFDM Symbol

Figure 94 shows the real part of the time-domain samples of the designated OFDM symbol received at the CM. There are 4,096 time-domain samples that correspond to this OFDM symbol.




Figure 94 - Real Part of Time-Domain Samples of OFDM Symbol Received at CM

Figure 95 shows the complex time-domain samples of the designated OFDM symbol received at the CM. The x-axis represents the real part (I), and the y-axis represents the imaginary part (Q) of the soft decisions. There are 4,096 points plotted. The distribution is close to a joint Gaussian (normal) distribution, similar to the distribution of thermal noise.




Figure 95 - Soft Decisions (Time-Domain Samples) of OFDM Symbol Received at CM

Figure 96 shows the FFT of the above time-domain samples of the OFDM symbol received at the CM. The FFT output has 4,096 complex frequency-domain samples. These samples correspond to the original transmit constellation, except they have not been synchronized in time and frequency or equalized, so the constellation is distorted and rotating.

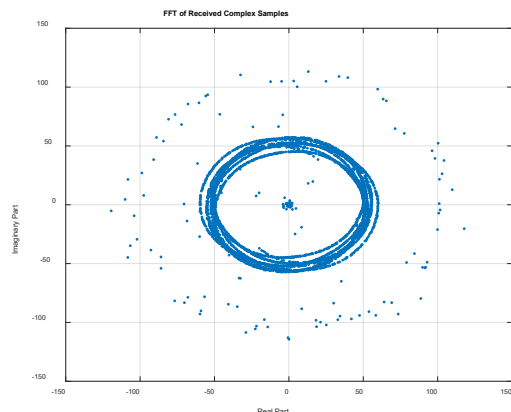


Figure 96 - FFT of Soft Decisions (Resulting Frequency-Domain Samples) of OFDM Symbol Received at CM

5.8.2.2 Constellation Plot Using Known Data Points

The downstream symbol capture measurement can be used to plot the received constellation while taking advantage of the knowledge of the correct symbols. Further processing by the PNM server (not shown here) could remove the time and frequency offsets from the data in Figure 96 and perform channel equalization, thereby recovering the receive constellation. Because each point in the constellation is known at the transmitter and receiver, the receive constellation could be annotated with different colors for different classifications of points. For example, points corresponding to low-frequency subcarriers could be plotted in green, high-frequency subcarriers in red, etc. Symbol errors, i.e., points that land outside the decision boundary around that point on the constellation, could also be highlighted in a color. The color assigned to each constellation point could provide a color map of the length of the error vector from the ideal constellation point. Many creative options are available for constellation displays, allowing the operator to visualize various impairments, which would be analogous to the constellation analysis performed in typical test equipment used on SC-QAM in many current applications.

5.8.2.3 Channel Response Computation

The downstream symbol capture measurement is extremely versatile and does not necessarily require demodulation (time/frequency synchronization and equalization) of the received OFDM symbol. The approach treats the downstream channel as a black box; a known time domain input, $x(t)$, is put into the channel at the CMTS, and the time domain output, $y(t)$, is captured at the CM. This approach provides a probe of the channel, which allows solving for the channel response. For convenience, the CMTS captures the frequency-domain symbol, $X(f)$, rather than the time-domain signal, $x(t)$, because the frequency-domain signal is more readily available and, if desired, the PNM server can easily connect to convert between the time and frequency domains by using the FFT and IFFT. The FFT of the CM signal, $y(t)$, is $Y(f)$. The CMTS signal is already in the frequency domain, so there is no need to take its FFT.

To compute the channel response, PNM processing divides the FFT of the CM signal, $Y(f)$, point-by-point by the CMTS signal, $X(f)$, yielding the complex frequency response, $H(f)$, of the downstream channel. This process can be repeated with multiple OFDM symbol captures, and 10 or more instances of $H(f)$ may be averaged. Because some subcarriers in a single OFDM symbol may have small QAM symbols (inner constellation points) and some have large QAM symbols (outer constellation points), it has been found that weighted averaging improves the final result. If there is a phase ramp caused by timing offset, it can be removed by post-processing. The goal is to plot the magnitude, phase, and group delay of $H(f)$.

Figure 97 shows the magnitude (in decibels) of the frequency response of the downstream cable channel computed from a single captured downstream OFDM symbol, i.e., without averaging.




Figure 97 - Frequency Response of Downstream Channel Computed from a Single Captured Downstream OFDM Symbol

Figure 98 shows the phase of the frequency response of the channel.




Figure 98 - Computed Phase Response of Downstream Channel

Figure 99 shows the unwrapped phase of the frequency response of the channel.




Figure 99 - Computed Unwrapped Phase Response of Downstream Channel

The group delay of the channel may also be computed.

5.8.3 Summary of Uses of Downstream Symbol Capture

The following is a partial list of the capabilities of the downstream symbol capture measurement.

- Synchronize capture over CMTS and multiple CMs, providing information for ingress location
- Compute enhanced constellation plots showing frequency/subcarrier dependence and QAM decision boundary errors, allowing impairments to be visualized
- Measure downstream channel response (amplitude, phase, group delay) even under low SNR conditions
- Compute histogram of downstream signal with rapid response
- Analyze nonlinear distortion in downstream channel
- Compute channel noise under signal

5.8.4 Conclusion and Suggestions for Future Work

Downstream symbol capture provides the possibility of comparing a symbol sent by the CMTS to that symbol as received by the CM; this capability covers a multitude of use cases. Many of its results complement and extend other downstream measurements. As the measurement begins to be implemented, this section can be updated with field results.

Work do be done with symbol capture, when feasible, is to discover which impairments it identifies best and if effective methods to improve troubleshooting and repair can be developed. Once the methods are deployed in the field, initial data collection and analysis should be conducted to develop theories for further testing and then to develop new methods around what is proven.

5.9 Downstream Histogram

5.9.1 Overview

This section describes the downstream histogram measurement.

5.9.2 Description of Measurement

The purpose of the downstream histogram is to provide a measurement of nonlinear effects in the channel such as amplifier compression and laser clipping. For example, ideal one-sided laser clipping would cause one tail of the histogram to be truncated and replaced with a spike. The CM captures a histogram of time domain samples after the wideband front end of the receiver (full downstream band). It is a two-sided histogram, encompassing far-negative and far-positive values. The histogram has either 256 equally spaced bins with even symmetry about the origin or 255 equally spaced bins with odd symmetry about the origin. These bins typically correspond to the 8 MSBs of the wideband ADC in the CM receiver. The histogram provides a dwell count and hit count. The dwell count is a 32-bit unsigned integer that gives the number of samples observed while counting hits for a given bin; it typically will have the same value for all bins. The hit count, or bin count, is a vector of 32-bit unsigned integers that gives the number of samples that fell in any given bin during the measurement period. The CM reports the dwell count and hit count for each bin. When the histogram measurement is enabled, the CM computes a histogram with a dwell of at least 10 million samples in 30 seconds or less. With this many samples, the histogram can reliably measure a probability density per bin as low as 10^{-6} with at least 10 hits in each bin. The CM will continue accumulating histogram samples until it is restarted, disabled, approaches its 32-bit overflow value in one of the bins, or times out. The CM also provides a UTC timestamp in the header of the PNM file to indicate the start time of the capture.

The ADC in a typical CM samples the entire downstream band, so the histogram is not limited to a single OFDM channel and can therefore show the distortion effects of the downstream laser.

Because the histogram is captured after the CM front end (specifically, after the automatic gain control (AGC) of the CM front end), the histogram is influenced by the CM's front end as well as by the plant and the transmitter. Also, because the histogram is built over a long time frame (e.g., 30 seconds, but it can be longer or shorter), dynamic plant

influences may cause a single histogram to capture a mix of conditions. Further, there are no standardized requirements on the CM front end or its AGC, and AGCs can display hysteresis of some sort (e.g., "intentionally" changing the loading target of the ADC). Thus, under dynamic plant conditions, the AGC may react during histogram collection and influence a small portion of the histogram collection while in transient. More significantly, the set point of the ADC may change during histogram collection, causing significant portions of the histogram to be collected with different ADC loading targets.

In many millions of histograms collected and reviewed, a very small percentage were found to exhibit some seemingly anomalous results, and it has to be considered that these results could be influenced by intended CM front end (and AGC) behavior or by unintended CM front end behavior in a small percentage of cases. Most of the millions of downstream histograms reviewed to-date show benign, common histogram results. Several types of pathological histograms were observed, rarely, but often enough to be noteworthy. The cause of these pathological results has not been understood or explained, nor can it be determined if the if the cause of any particular pathological histogram was the CM front end, the plant, or an interaction of the two.

The histograms in this section were generated by simulation; they are idealized, involving computer models of plant impairments. They likely do not represent realistic models of CM front end behavior (filtering and AGC, especially), so they do not necessarily illustrate histogram results expected from CMs operational in cable plants. The idealized nature of computer-simulated histograms, especially with idealized, modeled impairments, is not to be confused with cable modems' actual histogram results.

5.9.3 Background Theory

There are two straightforward ways to describe how a histogram works: from the standpoint of practical implementation and from the standpoint of probability theory.

5.9.3.1 Description Based on Practical Implementation

In its most basic implementation, the histogram is simply a count of how many times each ADC code occurred during the measurement. Take a measurement run for 10 million samples (it should take only a few seconds given the high sample rate in the CM). For every sample, the 8 MSBs are examined and the other bits are ignored (for example, if the ADC has 10 bits, the lower 2 bits are ignored). There are 255 or 256 "hit count" registers, each corresponding to one of the ADC codes. For example, the ADC codes may have 256 twos-complement values from -128 to +127 or, in some cases, 255 twos-complement values from -127 to +127. A register is assigned to each of these codes. On each sample, the ADC outputs a code, and the register corresponding to that code is incremented. When this process has been applied 10 million times, codes with high probability will have occurred often and thus will have high hit counts, while codes that occurred rarely will have low hit counts.

5.9.3.2 Description Based on Probability Theory

The histogram (when hit counts are normalized so the sum over all bins equals 1) is an empirical probability distribution of a discrete random variable. When graphed, the x-axis contains the possible values of the variable (in this case, the 255 or 256 values of the 8 MSBs of the ADC), and the y-axis gives the hit count, or number of times each value occurred during the measurement. In a CM, the input RF signal is a continuous quantity sampled by the ADC. Therefore, the histogram is an approximation to the probability density function of the RF input signal.

5.9.4 Examples of Histograms

As described above, a histogram is a plot of the values of a signal on the x axis and the number of times that value occurred on the y axis. When properly normalized, it represents the probability density function of the signal. Each y-axis value represents the total number of occurrences of the x-axis value (or range of values) over the time during which the histogram data were taken. A non-bell-curve appearance of the histogram enables the visual identification of nonlinearities in the plant such as amplifier compression and laser clipping.

5.9.4.1 Example 1: Ideal Histogram

To show how the histogram illustrates the effect of laser clipping, Figure 100 contains the time domain and histogram plots of a computer-simulated normal OFDM signal with no clipping. The histogram is approximately Gaussian (normal or bell-shaped) because of the summation of a large number of independent subcarriers in an OFDM

symbol. The figure shows an ideal Gaussian distribution (dotted red line) for reference. In contrast, Figure 101 shows the same plots for an OFDM signal after idealized laser clipping. The signal is normalized to standard deviation of 1. It is an ideal simulation for the purposes of illustration, in which the clipped laser causes all values of +2 or greater to be replaced with the value +2. Laser clipping replaces one tail of the histogram with a spike at the value $x = 2$.




Figure 100 - Normal OFDM Signal with Gaussian-Shaped Histogram;
(a) Time Domain, (b) Histogram




Figure 101 - Laser Clipping in Signal with One Tail of Histogram Replaced with Spike;
(a) Time Domain, (b) Histogram

5.9.4.2 Example 2: Upstream Soft Clipping

The one-sided hard clipping in Figure 101 starkly illustrates the effect of ideal laser clipping. In reality, the result is more subtle because of softer, more gradual clipping or other effects such as bandlimiting. Figure 102 shows the histogram of a signal that has passed through an actual upstream cable plant with laser clipping. The upper (rightmost) tail of the distribution, though not completely cut off, has been reduced noticeably below the reference Gaussian distribution. The missing samples from the tail are evident in the fairly pronounced spike on the right side of the distribution at around 1.25 standard deviations on the x axis. Although the effect is not as pronounced as in the idealized computer simulation, this histogram could be analyzed by the appropriate algorithm to show a strong indication of laser clipping, especially if a longer set of data is averaged.




Figure 102 - Histogram of Actual Upstream Clipping Event

5.9.4.3 Example 3: Undistorted Downstream

For the example shown in Figure 103, an OFDM downstream was observed under conditions where no laser distortion was present. The histogram closely matches the Gaussian reference (red dotted line), defined as the Gaussian (normal) distribution with the same mean and variance as the measured histogram, but some ADC code irregularities are noticeable around the peak. The tails closely match the Gaussian reference, at least to the extent it is possible to see in this plot.




Figure 103 - Histogram of Undistorted OFDM and SC-QAM Downstream Signal

To magnify the behavior in the tails, the logarithm base 10 of the bin count is displayed in Figure 104. It shows some slight departure from an ideal Gaussian distribution along the tails.




Figure 104 - Log Histogram of Undistorted Downstream Signal

It is useful to plot the ratio, in decibels, between the ideal Gaussian distribution and the signal histogram. The delta from Gaussian may be defined by the following formula.

$$\text{gaussian_delta_dB} = 10 * \log_{10}(\text{gaussian_distribution} / \text{signal_histogram})$$

The signal histogram has been normalized to a discrete probability distribution, i.e., the sum of its bin counts = 1. (The use of decibels defined as $10 * \log_{10}$ of the histogram ratio is arbitrary and is selected because of its familiarity from signal measurements.) This measure is plotted in Figure 105. In this benign case, the ratio remains below 7 dB even in the tails, where the data are less robust because fewer counts occurred. Points where the signal histogram is zero (values where no counts occurred) are omitted from the plots.




Figure 105 - Delta from Gaussian (in Decibels) for Histogram of Undistorted Downstream Signal

Example 4: Downstream RFoG System with Laser Distortion

For this example, an OFDM downstream was transmitted through a radio frequency over glass (RFoG) system assembled in a laboratory under extreme conditions (intentionally constructed) in which the laser was distorting. A single OFDM channel was operating in the downstream. The distortion was severe enough to barely support 16-QAM being received. Figure 106 shows the resulting histogram. The distortion causes the histogram to depart noticeably from the Gaussian reference; in particular, the peak is lower, the tails are suppressed, and there is a slight asymmetrical bulge on the right.




Figure 106 - Histogram of Downstream RFoG Signal with Severe Laser Distortion (190 MHz, 16-QAM, 3,800 Subcarriers)

The behavior in the tails can be examined in a log histogram (Figure 107). It shows a significant departure from the ideal Gaussian distribution, especially in the tails.




Figure 107 - Log Histogram of Downstream RFoG Signal with Laser Distortion

As in the benign case, the delta between the signal histogram and the ideal Gaussian distribution was computed (Figure 108). The plot shows a significant difference from Gaussian, more than 45 dB.




Figure 108 - Delta from Gaussian (in Decibels) for Histogram of Downstream RFoG Signal with Severe Laser Distortion

5.9.4.4 Example 5: Downstream with Varying Laser Distortion

For this example, a combined OFDM and multiple-SC-QAM downstream was transmitted through a downstream system while the amount of input power to the laser, and hence the distortion, was varied by summing in an extra OFDM channel with adjustable power. The OFDM RxMER was reduced from 37 dB to 23 dB as the distortion increased. Figure 109 shows the spectrum analyzer display for the benign case with low laser distortion. The downstream contains an OFDM channel with an RxMER of 37 dB and a bank of SC-QAM channels. Figure 110 shows the spectrum analyzer display for the case with high laser distortion. An additional OFDM channel is inserted at a lower frequency than the other channels, resulting in an RxMER of 23 dB in the OFDM channel of interest.



Figure 109 - Spectrum Analyzer Display of Downstream Signal with Low Laser Distortion, OFDM RxMER = 37 dB



Figure 110 - Spectrum Analyzer Display of Downstream Signal with High Laser Distortion, OFDM RxMER = 23 dB

Figure 111 and Figure 112 show the log histogram and Gaussian delta of the histogram, respectively, for the benign case with an RxMER of 37 dB. This case has a delta from an ideal Gaussian distribution of about 5 dB max. In contrast, Figure 113 and Figure 114 show the log histogram and Gaussian delta, respectively, for the case with an

RxMER of 23 dB. The higher laser distortion causes a delta of about 10 dB max. The difference in the histogram delta between the benign and distorted cases is about 5 dB.




Figure 111 - Log Histogram of Downstream Signal with Low Distortion, OFDM RxMER = 37 dB




Figure 112 - Delta from Gaussian of Downstream Signal with Low Distortion, OFDM RxMER = 37 dB




Figure 113 - Log Histogram of Downstream Signal with High Distortion, OFDM RxMER = 23 dB




Figure 114 - Delta from Gaussian of Downstream Signal with High Distortion, OFDM RxMER = 23 dB

5.9.5 Conclusion and Suggestions for Future Work

The histogram is a useful tool for detecting and measuring compression (including laser clipping) in the downstream cable plant, but it does not necessarily indicate where the impairments occur.

A possible future effort to consider would be to run the histogram data through various statistical tests to gauge their capabilities at detecting plant issues—including a Shapiro-Wilks, Kolmogorov-Smirnov, Chi-squared, or Pearson test—and a Q-Q plot to look for the difference between the measured plant histogram and a Gaussian distribution.

The possibility of a more accurate reference distribution for OFDM signals, rather than an ideal Gaussian distribution, may be investigated. Use of more OFDM signals may lead to larger tails, but field studies are required to determine whether this is true and under what conditions.

5.10 Physical Layer Link Channel (PLC) and Next Codeword Pointer (NCP)

5.10.1 PLC Overview

The PLC is a critical part of the downstream OFDM signal that conveys physical layer parameters⁶ from the CMTS to the cable modem. Specifically, the OFDM channel descriptor (OCD) and the downstream profile descriptor (DPD) for Profile 0 are transmitted periodically (at least once per every 250 ms) on the PLC. DOCSIS 3.1 operation by the modem is not possible if the cable modem cannot receive and demodulate the PLC.

5.10.2 A Closer Look

Figure 115 shows a spectrum analyzer display of a 192 MHz-wide downstream OFDM signal. The figure highlights the PLC band, a 6 MHz-wide contiguous portion of the OFDM signal within which the PLC is centered. This 6-MHz encompassed spectrum cannot have any excluded subcarriers or exclusion bands.

⁶ The PLC conveys a variety of information required by modems to come online: timestamp, energy management, DPD and OCD, and trigger messages for synchronizing events between a CMTS and modems. See [MULPIv3.1] for details on these messages.




Figure 115 - Example Spectrum Analyzer Display Showing the PLC Band Within a 192 MHz-Wide OFDM Signal

The continuous pilots in the PLC band have a unique pattern that tells the modem where to find the PLC (see Figure 116). The PLC band's pilot pattern is not user-adjustable; it is a fixed pattern defined in [PHYv3.1]. Likewise, the PLC's placement in the center of the PLC band is not user-adjustable.




Figure 116 - Example Spectrum Analyzer Display Showing PLC Placement in the PLC Band, with a Close-Up of the PLC Band (Yellow), PLC (Red), and Boosted Continuous Pilots in a Unique Pilot Pattern

The PLC is 400 kHz wide and comprises eight subcarriers (50-kHz spacing) or 16 subcarriers (25-kHz spacing), each using 16-QAM to convey physical layer parameters to the cable modems.

NOTE: The cable operator chooses where in the OFDM signal to place the PLC band (and PLC centered within that band). Ideally, the PLC band should be in a known clean part of the OFDM signal that is not susceptible to ingress, direct pickup, or other interference. Unfortunately, some operators leave the PLC band on the CMTS's default configured frequency, perhaps not realizing that the frequency is user-configurable. From [PHYv3.1]: "The 6 MHz encompassed spectrum containing the PLC at its center may be anywhere provided it contains 6 MHz of spectrum without any excluded subcarriers." If the default frequency happens to overlap an over-the-air service such as an LTE provider's downlink, interference to the PLC can occur.

For more information about the PLC, see [PHYv3.1], Section 7.5.13, "Physical Layer Link Channel (PLC)."

5.10.3 Useful PLC Parameters

The cable modem and DOCSIS 3.1 modem-equipped test equipment have to be able to locate, receive, and demodulate the PLC for the PHY parameters to be extracted. If the modem cannot receive and demodulate the PLC, DOCSIS 3.1 operation will not be possible. Good engineering practice suggests the following.

- PLC signal level: greater than or equal to -15 dBmV. The PLC's signal level will usually be similar to the OFDM channel power (-15 dBmV to +15 dBmV).

- PLC RxMER: greater than or equal to 15 dB. The PLC's RxMER likely will be similar to the OFDM signal's average RxMER per subcarrier.
- Frequency: Placement of the PLC band and PLC should be such that they do not overlap known sources of ingress, direct pickup, or other interference.
- Codeword errors: If the modem (or test instrument) reports any correctable codeword errors in the PLC, the errors can be ignored. However, there should not be any unreliable codeword errors in the PLC.

5.10.4 NCP Overview

NCPs are subcarriers containing messages that point modems to the start of a codeword. NCPs also contain information on things such as data profiles, zero bit-loaded subcarriers, NCP type, and T bits. NCPs direct traffic, in a way, and they are necessary because an OFDM symbol can contain codewords for several profiles; codewords can continue from one symbol to the next; and profiles can have different, and variable, QAM modulation orders.

The NCP has its own bit loading profile separate from that used for data subcarriers. It can use QPSK, 16-QAM, or 64-QAM.

More information about NCPs can be found in [PHYv3.1], Sections 7.5.14 and 8.3.4. OFDM requires NCPs to map the codewords within and across individual subcarriers. Because of the critical nature of their payload, they are modulated in low order to help improve their robustness. It is important that the NCP not contain uncorrectable codewords.

5.10.5 Useful NCP Parameters

The cable modem and DOCSIS 3.1 modem-equipped test equipment have to be able to properly receive the NCPs in order for the modem to know where codewords are located. Good engineering practice suggests the following.

- Codeword errors: If the modem (or test instrument) reports any correctable codeword errors in the NCP, the errors can be ignored. However, there should not be any CRC-24-D errors in the NCP field. If the NCP CRC field indicates an error in the NCP field, then the CM is required to reject all NCP data message blocks in the NCP field of the current symbol.

5.10.6 How to Implement

[CM-OSSiv3.1] states that one can obtain the total number of NCP fields received by the CM, and the total number of these fields that failed the CRC check, as part of the DsOfdmChannel object.

5.10.7 Conclusion

The PLC is an important part of the downstream OFDM signal that carries PHY parameters from the CMTS to the modems. If the PLC cannot be reliably demodulated, then DOCSIS 3.1 operation will not be possible. Operators need to ensure that the CMTS is configured to place the PLC band and PLC on a frequency within the OFDM signal that is not susceptible to interference. Parameters such as PLC signal level, RxMER, and codeword errors are useful metrics for PNM. Additionally, NCPs are critical for proper modem operation, telling modems where codewords begin. Parameters such as NCP codeword errors are useful metrics for PNM.

5.11 MAC CRC Failures

As discussed below, the primary indicator of error performance on downstream OFDM channels should be the downstream FEC codeword statistics, not the MAC CRC failure count. See Section 5.4 for more information.

[PHYv3.1] requires that the CM provide the following counts on MAC frames addressed to the CM for each downstream profile excluding the transition profile.

- MAC frame failures: number of frames that failed MAC CRC check
- Total number of MAC frames

This requirement clearly applies to OFDM channels because it appears in the downstream PNM section and mentions downstream profiles specifically. Note that this CRC is completely different from the PHY-MAC convergence CRC for NCP message blocks in an OFDM symbol.

In the docs-if31-mib pointed to by [CM-OSSIV3.1], the attribute docsIf31CmDsOfdmProfileStatsInFrameCrcFailures under DocsIf31CmDsOfdmProfileStatsEntry is a counter for "the number of MAC frames measured on this profile that failed the MAC frame CRC check."

DOCSIS 3.1 CMs discard bits of uncorrectable codewords at the downstream PHY layer and do not pass them to the MAC layer. (The DOCSIS 3.1 specifications do not mention if an uncorrectable codeword in the downstream results in the bits of that codeword being discarded or passed to the next higher layer.) Therefore, when a CM receives uncorrectable codewords on a profile, the docsIf31CmDsOfdmProfileStatsInFrameCrcFailures attribute does not increment. Though CMs will support this count, it will usually read 0 or a very small number because of very unlikely random bit errors in internal handling or the very unlikely event of the PHY FEC "correcting" to a wrong codeword, thus generating a "correctable codeword" that actually contains errors. For this reason, codeword statistics are the best indicator of error performance for the particular profile received by the DOCSIS 3.1 CM on an OFDM channel.

There is interest in determining the impact of FEC codeword errors on a particular CM or users behind a CM. However, even if a downstream physical layer implementation passed all the bits of an uncorrectable codeword to the MAC and MAC layer CRC check for the impacted DOCSIS frames failed, there is no reliable way to determine whether the packet was intended for the particular CM or user behind the CM because the MAC layer CRC check on the DOCSIS MAC frame also covers the destination MAC address of the frame. Also, an uncorrectable codeword may have errored bits in the two-byte LDPC codeword header, and the codeword header contains a pointer to where the first MAC frame begins in the codeword to aid in MAC frame alignment. Thus, an errored codeword header may result in MAC frame misalignment and, as a result, CRC errors if the bits of uncorrectable codewords are passed to the MAC. So, regardless of whether a CM PHY layer implementation passed the bits of uncorrectable codewords, there is no reliable way to use CM counts to associate uncorrectable FEC codeword errors with errored or undelivered frames to a particular CM or users behind a CM.

From the CMTS side, the CMTS can count frames and octets per downstream profile (ProfileOutFrames and ProfileOutOctets, respectively). A comparison can be made between the CMTS frame and octet counts per profile and the CM frame counts statistics (ProfileInFrames, ProfileInOctets) for that profile. However, as noted above, at the CM, this count would be for all packets and octets on the profile at the MAC layer and not for packets and octets specifically destined to the CM or a user behind the CM.

A CMTS also can count packets per service flow. If the service flow is not bonded, then the only frames transmitted for that service flow will be on a single profile on a single channel. In this case, the frames transmitted on the flow would correspond to the frames received by the CM on that single profile on the single channel. Thus, a comparison can be made between the CMTS packet count for the service flow and the CM frame count statistics for that profile (ProfileInFrames). Note that because CMs are not required to classify frames in the downstream to identify the service flow the frames were sent on, a count for received PDUs at the CM for a particular service flow will likely read 0.

In almost all cases, MAC CRC failures are not useful in identifying a problematic or failing DOCSIS 3.1 OFDM channel. On a downstream OFDM channel, the MAC CRC frame failure count is expected to read 0 or close to 0.

5.12 SNR Margin for Candidate Profile

Section 9.3.6.1 and Appendix VI of [PHYv3.1] discuss an estimate of the SNR margin available on a downstream OFDM channel with respect to a candidate OFDM modulation profile. The CMTS can send test data to the CM to measure the performance of a candidate profile. The OPT-REQ message (described in [MULPIV3.1]) is used by the CMTS to cause a CM to report various aspects of an OFDM downstream signal, including the CM's ability to receive a candidate downstream profile. The CM sends an OPT-RSP message to acknowledge the request; if the request was to start a test, then it sends another OPT-RSP message to report the results.

The OPT-REQ message from the CMTS includes a TLV that specifies the statistics requested from the CM. The following are statistics related to RxMER and SNR that can be requested from the CM.

- RxMER Statistics per Subcarrier
- RxMER per Subcarrier Threshold Comparison for Candidate Profile

- SNR Margin for Candidate Profile

In the OPT-REQ message from the CMTS, the following parameters are communicated to the CM.

- RxMER vs Bit-Loading Target (threshold RxMER value used by the CM to calculate the SNR margin for the candidate profile under test and reported in the associated OPT-RSP message)
- RxMER Margin (used by the CM to adjust the RxMER threshold value used to calculate the "Number of subcarriers whose RxMER is RxMER Margin below the RxMER Target" reported in the OPT-RSP message)

In the OPT-RSP message, under the category of RxMER and SNR Margin Data, the following data can be reported.

- RxMER per Subcarrier
- RxMER per Subcarrier Threshold Comparison Result
- Number of subcarriers whose RxMER is RxMER Margin below the RxMER Target—
RxMER < (RxMER vs Bit-loading Target – RxMER Margin)
- SNR Margin
- Average RxMER

If a CM receives an OPT-REQ message with the requested statistic of the SNR Margin for Candidate Profile, the CM is required to include the SNR Margin in the OPT-RSP message.

For SNR Margin, the CM is required to implement an algorithm to estimate this parameter on the downstream data channel for a candidate profile. Appendix VI of [PHYv3.1] suggests an algorithm to compute the SNR Margin for the candidate profile, but it is merely an example without any normative language. Vendors are encouraged to implement this algorithm, but alternative, vendor-specific algorithms are not precluded.

The text of Appendix VI is reprinted here with square bracket comments added for clarity.

The CM measures the RxMER value for each data subcarrier as specified in Section 9.3.6 [PHYv3.1]. From these measurements it calculates the average RxMER over all data subcarriers [Average RxMER], MER1. The CM accepts as an input the required RxMER delivering a defined threshold of CER = 1e-5 under ideal AWGN conditions for each bit loading [RxMER vs Bit-Loading Target]. The CM computes the difference of the measured RxMER values from the required RxMER values. The CM computes the required average RxMER, denoted MER2, over all data subcarriers for the candidate profile. The averaging computations for MER1 and MER2 use values in the log (dB) domain. The SNR margin is defined as MER1 - MER2, where all quantities are in dB. As an example, if the CM measures MER1 = 33 dB, and the candidate profile requires MER2 = 30 dB, the CM reports an SNR margin of 3 dB. In addition, the CM reports the number of subcarriers whose RxMER is at least x dB below the defined threshold for the bit loading of the given subcarrier [Number of subcarriers whose RxMER is RxMER Margin below the RxMER Target], where x is a configurable parameter with default value = 1.

6 UPSTREAM PNM MEASUREMENTS

6.1 Upstream Power Levels

6.1.1 Overview

This section describes DOCSIS 3.1 upstream power levels for CMTS or R-PHY device (RPD) receive and CM transmit and how they differ from legacy SC-QAM DOCSIS power levels.

6.1.2 Description of Measurement

6.1.2.1 CM Transmit

The requirements for transmit power are from [PHYv3.1], Section 7.4.13, "Cable Modem Transmitter Output Requirements." The following list of CM transmitter output signal characteristics is based on [PHYv3.1], Table 16.

1. The CM MUST be capable of transmitting a total average output power of 65 dBmV.
2. The CM MAY be capable of transmitting a total average output power greater than 65 dBmV.

Assume that a hypothetical modem transmits a 95 MHz-wide OFDMA signal at a maximum total average power of +65 dBmV, the digital channel power for the full 95 MHz-wide signal. That hypothetical modem's reported power in a 1.6-MHz bandwidth would be calculated as follows.

$$65 \text{ dBmV} - 10\log[\text{ceiling}(95/1.6)] = 47.22 \text{ dBmV}$$

NOTE: OFDMA transmit power should always be reported in a 1.6-MHz bandwidth, but SC-QAM reported power depends on the combination of equipment.

1. If both the CM and CMTS are DOCSIS 3.1 devices, SC-QAM transmit power is reported in a 1.6-MHz bandwidth.
2. Of the CM and the CMTS, if one is a DOCSIS 3.1 device and the other is a DOCSIS 3.0 device, then the SC-QAM transmit power is reported in the occupied bandwidth (e.g., 6.4 MHz).

DOCSIS 3.1 CMs report their transmit level based the on 1.6-MHz bandwidth equivalent. When a DOCSIS 3.1 CM replaces a DOCSIS 3.0 CM and transmits SC-QAMs, that transmit level is reported as 1.6-MHz equivalent and not in the actual channel width used, which could be 6.4 MHz. In this example, the transmit level would be higher than the reported power by 6 dB, calculated by $10\log(6.4/1.6)$.

6.1.2.2 Maximum and Minimum Transmit Power

[PHYv3.1] states a maximum total transmit power of at least 65 dBmV and at least maximum total transmit power spectral density of 53.2 dBmV/1.6 MHz; vendors are allowed to provide higher maximum power.

Usable bandwidth could subtract guardbands and any exclusion zones. For example, if configuring a 48 MHz-wide upstream OFDMA signal, the usable bandwidth could be 47 MHz if there are 0.5-MHz guardbands on each end. The following formula can be applied.

$$65 \text{ dBmV} - 10\log[\text{ceiling}(47/1.6)] = 50.23 \text{ dBmV}$$

If monitoring the CMTS for cable modem information, this might appear with a `show cable modem "mac address" verbose` type of output.

At maximum transmit power, a minimum of modulated spectrum is allowed because of the limit on power spectral density. If transmitting at maximum power (e.g., 65 dBmV), the ratio of modulated spectrum (in megahertz) divided by 1.6 cannot be lower than 15. The maximum power spectral density required in this case is calculated as follows.

$$65 \text{ dBmV} - 10\log(15) = 53.2 \text{ dBmV/1.6 MHz}$$

Smaller modulated spectrum is allowed but cannot achieve the maximum total transmit power because of the power spectral density limitation.

6.1.2.3 Modem Ranging

Initial ranging (IR) frequency can be set in the OFDMA config; it is one-third from the bottom bandedge by default. A higher frequency could be best to range on as it would represent the worst case for cable attenuation, but doing so could affect the TCS selection (upstream bonding group) and the 12-dB dynamic range window (DRW) decision.

Fine ranging (FR) frequency is equivalent to `init(r2)` or station maintenance as known in DOCSIS 3.0 systems. Probes would be responsible for determining the frequency response of the channel because they cover all frequencies in the channel. Measuring the frequency response of the probes allows the CMTS to provide the pre-equalization coefficients to the modem in ranging-response (RNG-RSP) messages. Variations in transmit power across the channel because of pre-equalization are not subject to limitations of the 12-dB DRW.

NOTE: DOCSIS 3.1 CMs must register upstream `mtc-mode` for DOCSIS 3.1 technology to work on the downstream. The CM may get `reject(na)` and keep cycling if upstream bonding (`mtc-mode`) is deactivated or the CM has upstream power level issues that keep it from selecting an upstream bonding group. The code should still allow a single-channel bonding group, which would suffice.

6.1.2.4 CMTS Receive

The CMTS upstream port is looking for 6.4-MHz equivalent receive and is typically set for 0 dBmV. The CM's actual transmit level is based on a 1.6 MHz-wide channel equivalent during ranging. When looking at a spectrum analyzer at the CMTS, a 6.4 MHz-wide SC-QAM channel and a 6.4 MHz-wide OFDMA channel would "appear" to have the same PSD.

6.1.3 How to Implement

Operators are accustomed to seeing CM transmit power levels on SC-QAM channels as total power rather than PSD in 1.6 MHz. Because OFDMA channels can exist over various bandwidths and can be much larger than SC-QAM channels, DOCSIS 3.1 technology describes OFDMA channels by their power spectral density (decibel millivolts per 1.6 MHz); this is a departure from describing SC-QAM channels by channel power. Two MIBs are used to report the CM transmit (Tx) power for a DOCSIS 3.1 CM operating on a DOCSIS 3.1 CMTS.

The MIBs associated with upstream power levels are as follows ([DOCS-IF31-MIB]).

```
The docsIf31CmUsOfdmaChanTxPower
    represents the operational transmit power for
    the associated OFDMA upstream channel. The CM reports its Target
    Power, P1.6r_n as described in [PHYv3.1]. Valid values for this
    object are 68 to (213 + (4*(Pmax - 65 dBmV))), since 68 quarter dBmV
    represents the lowest Tx power value 17 dBmV and 213 represents the
    nearest quarter dBmV to the highest Tx power value 53.2 dBmV."
```

```
The docsIf31CmUsScQamChanTxPsd
    represents P1.6r_n, the power spectral density
    in 1.6 MHz, for the associated SC-QAM upstream channel."
```

This value is used by the CM in the RNG-REQ message and by the CMTS in the Commanded Power TLV of the RNG-RSP message.

The following MIB associated with upstream power levels in DOCSIS 3.0 also applies to DOCSIS 3.1 and is found in ([DOCS-IF3-MIB]).

```
The docsIf3CmStatusUsTxPower represents
    the operational CM transmit
    power for this SC-QAM upstream channel.
    In order for this attribute to provide consistent information
    under all circumstances, a 3.1 CM will report the average total
    power for the SC-QAM channel the same as was done for
    DOCSIS 3.0, regardless of whether it is operating with a 3.1 or
    a 3.0 CMTS. The value that is reported was referred to as Pr in
    the DOCSIS 3.0 PHY Spec."
```

6.2 Upstream Pre-Equalization Coefficients

6.2.1 Overview

Upstream pre-equalization allows cable modems to adapt their transmission to maximize the fidelity at the receiver. To achieve this, the CMTS measures the received signal and calculates the amplitude and phase corrections that are needed. This correction information is represented as a set of complex coefficients that are transmitted to the cable modem in the downstream channel. When the cable modem receives the coefficients, it applies them to the transmitter to pre-distort the signal by using a waveform that is the inverse of the channel response. This process runs periodically as directed by the operator. The adaptive pre-equalizer has served a very important role in PNM and continues to be valuable in OFDMA. This section discusses the differences between DOCSIS 3.1 pre-equalization and its predecessors. The CMTS can instruct the CM to send a signal either with or without equalization (if it was given equalization levels to use before). In the former case, the CMTS tells the CM the changes (deltas) to implement from the previous equalization used. In the latter case, the CMTS tells the CM what equalization settings to use.

One of most obvious differences between DOCSIS 3.1 and DOCSIS 3.0 adaptive pre-equalization is the wider bandwidth of DOCSIS 3.1 systems. Channel width for DOCSIS 3.0 systems was fixed at either 1.6, 3.2 MHz or 6.4 MHz and used 24 T-spaced time domain equalizer taps, whereas DOCSIS 3.1 OFDMA has a channel width of up to 96 MHz and uses per subcarrier frequency domain coefficients. There is greater time resolution when wider channels are implemented, which provides more precise distance calculations than narrower channels. The time resolution is calculated as the reciprocal of the total equalizer bandwidth. For example, a 96 MHz-wide OFDMA channel has 1,920 coefficients at 50-kHz spacing. In the time domain, each point represents 1/96 MHz, or 10.416 ns. This scenario can be compared to previous versions of 24-tap equalization coefficients of a 6.4 MHz-wide SC-QAM in which the signal's symbol rate is 5.12 Msym/s, which gives a time resolution of 195 ns. Because DOCSIS 2.0 and 3.0 technology specifies 24-tap T-spaced pre-equalization in the upstream, the maximum span of an adaptive equalizer for a 5.12-Msym/s signal is calculated by one of these methods.

$$(24 - 1) \times 0.1953125 \mu\text{s} = 4.49 \mu\text{s}$$

$$(24 - 1)/5.12 \text{ MHz} = 4.49 \mu\text{s}$$

This value compares to the maximum span of the 96 MHz-wide OFDMA channel with 50-kHz spacing, shown below.

$$1,920 \times 0.0104166 \mu\text{s} = 20 \mu\text{s}$$

Another significant difference between DOCSIS 3.0 and DOCSIS 3.1 PNM capabilities is the addition of summary metrics. Anticipating the additional complexity of extracting, decoding, and calculating a string of 1,920 coefficient values, the specification [PHYv3.1] provides summary values that are already computed. These summary values are discussed in the next section. In DOCSIS 3.1 systems, OFDMA pre-equalization occurs in the frequency domain, whereas it is in the time domain for DOCSIS 3.0 systems.

Figure 117 shows a frequency response plot of an upstream OFDMA signal.

The pre-equalizer coefficients at the CM are the values in use by the CM for upstream pre-equalization (PNMUsPreEq). The CMTS coefficients are the last set of pre-equalization updates that were sent to the CM by the CMTS (PNMUsPreEqLastUpdate). Both sets of coefficients are stored in the CM. In ideal plant conditions, the adjustments are minor. In typical plant conditions, some adjustments are needed, but the adjustments that the CM makes, reflected in the Pre-Eq coefficients, will fully reflect the condition of the plant (noise, in practice, will always prevent “perfection”) as long as the channel impulse response is negligible beyond the duration of the cyclic prefix. Thus, the pre-equalization in typical plant conditions provides that the signal presented to the burst receiver is “nearly optimal” when the cyclic prefix is maintained the same as used with the Probe. However, there can be some less-common cases where the CM cannot fully adjust and accept the coefficients given to it by the CMTS.

The pre-equalization coefficients are developed in the CMTS in response to a Probe transmitted by the CM. A Probe contains all subcarriers of an OFDMA channel which are not excluded. The CMTS can instruct the CM to transmit a Probe without using pre-equalization coefficients, which is analogous to a “reset” operation and in which case the RNG-RSP sent to the CM contains Pre-EQ coefficients defined to be used in a Set operation meaning that the coefficients are applied directly. In ongoing, stable operation, it is expected that the CMTS will instruct the CM to transmit the Probe using its current pre-equalization coefficients and any changes if needed will be sent in an RNG-RSP with coefficients defined to be used as an Adjust.

The pre-equalization coefficients sent by the CM to the PNM server are complex numbers using s2.13 format and expressed as 32 bits with the most significant 16 bits for the I component while the least significant 16 bits are for the Q value; however, the CMTS uses s1.14 format to send the updates to the CM. To update the pre-equalization coefficients, the CM performs a complex-multiply of the existing CM pre-equalization coefficients by their corresponding updates, to yield the new pre-equalization coefficients. As such, an update pre-equalization coefficient which is “ $1 + j0$ ” (where “ j ” is square-root of negative 1), results in “no change” to the corresponding existing pre-equalization coefficient in the CM.

Normally the updates from the CMTS would be minimal (e.g., near unity magnitude) indicating a stable plant; large adjustments (magnitude of an update value which is much larger or much smaller than unity) would indicate instability in which the CMTS (working with the CM) was unable to achieve acceptable equalization. Pre-equalization is conducted periodically, and not much variation should appear under normal operational conditions. Thus, this measurement is a valuable tool in support of PNM goals.

Ideally, the pre-equalizer at the CM will provide an optimal response to the upstream burst receiver, for that CM’s transmissions. In this case, the receiver will not need any adaptive equalization. The pre-equalization may deviate from “nearly optimal” for the channel, because of a) noise in and received by the receiver when creating the pre-equalization, b) the channel’s impulse response is non-negligible beyond the duration of the cyclic prefix, or c) a change in the channel impulse response since the most recent pre-equalization. In the event of a) or c), there will be some residual equalization needed at the burst receiver of the CMTS for best reception. Today, there is no way to get information about any additional equalization done by the burst receiver. A sufficient approach for estimating the channel in almost all cases is to invert the CM transmitter pre-equalization which resulted from the most recent update from the CMTS.

Because the pre-equalization typically should be able to take care of the entire needed plant response, and thus the burst receiver should have no additional work to do to equalize the received response, any indication that the burst receiver had additional work to do, i.e., updates which deviate significantly from unity magnitude, can be an indication of needed maintenance work.

When the CM receives the pre-equalization coefficient updates from the CMTS, typically the CM responds (via PreEqCoAdjStatus) with a “success” once it fully applies the requested adjustments to its pre-eq coefficients. If the CMTS coefficients cannot be adopted fully by the CM, the response will not be a “success.” For example, if there is too much tilt or ripple in the channel forcing the CM to clip the coefficients before applying them, then the status response is “clipped.” Other reasons that the CM needs to modify the coefficients results in a status of “other.” If the CM can’t apply the adjustments at all, it sends a status of “rejected.” Any of these last three status results will indicate that the CM response is not able to fully adjust to the plant conditions to provide an optimal signal to the CMTS. In this case, if there is an adaptive equalizer in the CMTS burst receiver, there will be residual work for the CMTS burst receiver to do.⁷ Further, simply knowing that the CM is not able to apply the adjustments fully or at all is sufficient for determining that maintenance should be conducted, which is our main purpose.

Therefore, Pre-equalization values are of use in a few ways.

- Large deviations from 0 dB in the pre-equalization values can indicate the plant needs large adjustments in equalization so the CM energy may be stressed to compensate for these plant conditions.
- Large differences between the current and last pre-equalization values can indicate that the plant conditions are rapidly changing. If several CMs show this condition, or the condition appears repeatedly over time, or both, then this large difference can indicate unstable plant that needs attention through maintenance. Without attention, there may be changes in the plant that get worse and cause poor performance due to the equalizers not being able to compensate.

⁷ In the event that the CM pre-equalization is not implementing the updates it was provided by the CMTS, or the plant has changed, if the burst receiver employs an adaptive equalizer, the plant channel response could be theoretically characterized by using both a) the CM’s pre-equalization coefficients (which are obtainable via PNM) and b) the CMTS receiver adaptive pre-equalization if implemented (which is not obtainable via PNM). In this case, the inverse of the composite of a) the transmit pre-equalization, and b) the receive adaptive equalization, could theoretically provide a more accurate estimate of the channel impulse response. However, the latter information is not obtainable today.

- If a CM is not able to accept and implement the pre-equalizer settings it was given, this is an indication that there is a problem isolated to a single CM. If CMs in a group are often not able to implement the pre-equalization fully, it can indicate that plant conditions are poor and need attention.

Pre-equalization can be considered for use in these ways to determine the urgency of maintenance, proactive or otherwise.




Figure 117 - Examples of Frequency Response for Upstream OFDMA Signals

6.2.2 Description of Measurement

The upstream adaptive pre-equalization coefficients are extracted from the cable modem MIB, which also includes summary values to simplify the process. The PNM MIB enumerates these summary values as follows.

6.2.2.1 Amplitude Ripple Peak-to-Peak

The amplitude ripple peak-to-peak measurement provides the magnitude of ripples in the measured spectrum. The delta value is useful for determining the fault distance of an echo cavity. Echo cavity distances (distance within the cavity) can be calculated as follows.

In feet: $492 \times (\text{velocity factor} / \text{ripple spacing in megahertz})$

In meters: $150 \times (\text{velocity factor} / \text{ripple spacing in megahertz})$

6.2.2.2 Amplitude Ripple RMS

The ripple RMS provides the ripple amplitude as the square root of the mean squares. This value allows calculation of the magnitude of the reflection. It is important to note that the amplitude is attenuated at longer distances as a function of loss on the transmission line. When distance and cable type are known, the amplitude can be adjusted to account for the loss. This technique can help evaluate fault magnitude in a consistent manner.

6.2.2.3 Amplitude Slope

Amplitude slope is calculated as the linear least squares fit of the frequency domain data. This value, also referred to as the in-channel tilt, represents the slope across the entire OFDMA channel. It is useful for setting up fixed equalization within the system amplifiers, for example. It is common to remove the tilt from the channel before calculating ripples.

6.2.2.4 Group Delay Ripple Peak-to-Peak

Similar to the amplitude ripple peak-to-peak measurement, the group delay can be calculated within the measured spectrum. Typically, in the case of reflective cavities, there is a symmetry between amplitude and group delay ripple. However, certain impairments may look like echo cavity ripples though they are caused by something else. Comparing the amplitude and group delay ripple peak-to-peak values can help evaluate these conditions.

6.2.2.5 Group Delay Ripple RMS

The group delay ripple RMS measurement is the square root of the mean squares of the ripple in the group delay. This attribute expresses the RMS value of the group delay ripple in units of 0.001 ns. Group delay ripple RMS can be used with amplitude ripple RMS to evaluate for symmetries associated with echo cavity-induced standing waves.

6.2.2.6 Amplitude Mean

This value represents the mean of the magnitude of the equalizer coefficients.

6.2.2.7 Group Delay Slope

Group delay slope is calculated as the linear least squares fit of the frequency domain data (see [CM-OSSiv3.1], Section D.4, "Slope and Ripple Algorithms"). This value represents the slope across the entire OFDMA channel. Significant group delay is common when a large number of diplex filters are cascaded. Group delay slope will indicate when a channel is located near the band edge of cascaded filters. These types of issues can cause latent ranging times and occasionally ranging failure when the equalizer is over-stressed.

6.2.2.8 Group Delay Mean

This attribute represents the mean of the group delay in nanoseconds.

6.2.3 Background Theory

For upstream adaptive pre-equalization theory of operation, refer to [PNMP-3.0], Section 6, "Methodology for PNM Using Upstream Equalization;" Section I.10, "Adaptive Equalization;" and Appendix VI, "DOCSIS Pre-equalizer Coefficients Analysis - Software Sequence Diagrams." See [PHYv3.1] for more information in OFDMA.

6.2.4 How to Implement

One can access the CM MIBs for the most recent coefficient updates received by the CM and the coefficients currently being applied by the CM for upstream transmissions.

6.2.5 Conclusion and Suggestions for Future Work

Measurement has not been field tested. Additional information will be provided in a future version of this document.

6.3 Upstream OFDMA RxMER per Subcarrier

6.3.1 Description of Measurement

Upstream data for RxMER per subcarrier on OFDMA are collected as follows.

1. Find the OFDMA Ifindex Containing the CM

```
'docsIf31CmtsCmRegStatusUsProfileIucList': ".1.3.6.1.4.1.4491.2.1.28.1.3.1.3"
```

The above MIB (found in [DOCS-IF31-MIB] Core MIB) shows encoded data where the Ifindex can be decoded. The encoding format of the MIB is as follows.

```
docsIf31CmtsCmRegStatusUsProfileIucList OBJECT-TYPE
    SYNTAX      OCTET STRING (SIZE (0 | 6..72))
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This attribute is a variable length series of hexadecimal
        octets where each series entry consists of the following fields
        (encoded in the following order):

        - The ifIndex (4 octets) of the OFDMA channel where the
        Profile IUCs are assigned.

        - The number or count of Data IUCs (1 octet with valid values
        of 1-2) assigned to this CM on that channel.

        - The list of Data IUCs (1 octet each with valid values of
        5, 6, 9-13) assigned to this CM on that channel.
```

The CCAP encodes each OFDMA channel in a CM's TCS as a separate n-octet entry in the UsProfileIucList.

Examples: a CM with a single OFDMA channel (ifIndex 36) and two assigned Data IUCs (5 and 6) would have a ProfileIdList value of 0x00000024020506. A CM with a 2 OFDMA channel bonding group each with one assigned Data IUC (IUC 5 on channel with ifIndex 34 and IUC 13 on channel with ifIndex 35) would have a UsProfileIucList value of 0x00000022010500000023010D).

Note: the CCAP MUST NOT include transitional IUCs or test IUCs in the UsProfileIucList.

Note that octet string lengths greater than 18 are optional.

This object is applicable to DOCSIS 3.1 modems but not to prior versions of DOCSIS modems. If the CM is a pre-DOCSIS 3.1 modem, the CMTS returns a zero length octet string."

```
::= { docsIf31CmtsCmRegStatusEntry 3 }
```

Correlating the above information to the following MIB walk (on Core) can generate the Ifindex → CM MAC address mapping (if the Core implements the above encoding correctly).

```
'docsIf3CmtsCmRegStatusMacAddr': ".1.3.6.1.4.1.4491.2.1.20.1.3.1.2",
```

2. Trigger Upstream OFDMA RxMER Test

```
us_collection_mibs = {
    "docsPnmCmtsUsOfdmaRxMerEnable": ".1.3.6.1.4.1.4491.2.1.27.1.3.7.1.1",
    "docsPnmCmtsUsOfdmaRxMerCmMac": ".1.3.6.1.4.1.4491.2.1.27.1.3.7.1.2",
    "docsPnmCmtsUsOfdmaRxMerNumAvg": ".1.3.6.1.4.1.4491.2.1.27.1.3.7.1.4",
    "docsPnmCmtsUsOfdmaRxMerFileName": ".1.3.6.1.4.1.4491.2.1.27.1.3.7.1.6",
    "docsPnmCmtsUsOfdmaRxMerPreEq": ".1.3.6.1.4.1.4491.2.1.27.1.3.7.1.3"
}
```

Set example:

```
# set pre eq
snmpset -v2c -c private 10.xx.xx.xx .1.3.6.1.4.1.4491.2.1.27.1.3.7.1.3.<ofdma_ifindex> i 2

# set num averages
snmpset -v2c -c private 10.xx.xx.xx .1.3.6.1.4.1.4491.2.1.27.1.3.7.1.4.<ofdma_ifindex> u 1

# set file name
snmpset -v2c -c private 10.xx.xx.xx .1.3.6.1.4.1.4491.2.1.27.1.3.7.1.6.<ofdma_ifindex> s <file_name>

# set CM MAC
snmpset -v2c -c private 10.xx.xx.xx .1.3.6.1.4.1.4491.2.1.27.1.3.7.1.2.<ofdma_ifindex> x FFFFFFFFFF

# enable
snmpset -v2c -c private 10.xx.xx.xx .1.3.6.1.4.1.4491.2.1.27.1.3.7.1.1.<ofdma_ifindex> i 1
```

3. Bulk Transfer MIBs

```
tftp_mibs = {
  "docsPnmBulkCtl": ".1.3.6.1.4.1.4491.2.1.27.1.1.1.0",
  "docsPnmBulkDestIpAddrType": ".1.3.6.1.4.1.4491.2.1.27.1.1.1.1.0",
  "docsPnmBulkDestIpAddr": ".1.3.6.1.4.1.4491.2.1.27.1.1.1.2.0",
  "docsPnmBulkDestPath": ".1.3.6.1.4.1.4491.2.1.27.1.1.1.3.0",
  "docsPnmCmCtlTest": ".1.3.6.1.4.1.4491.2.1.27.1.2.1.1.0",
  "docsPnmCmCtlTestDuration": ".1.3.6.1.4.1.4491.2.1.27.1.2.1.2.0",
  "docsPnmCmCtlStatus": ".1.3.6.1.4.1.4491.2.1.27.1.2.1.3.0",
  "docsPnmBulkFileUploadStatus": ".1.3.6.1.4.1.4491.2.1.27.1.1.2.1.4.0",
  "docsPnmBulkFileName": ".1.3.6.1.4.1.4491.2.1.27.1.1.2.1.2.0",
  "docsPnmCmtsUsOfdmaRxMerMeasStatus": ".1.3.6.1.4.1.4491.2.1.27.1.3.7.1.5"
}
```

The MIBs docsPnmBulkDestIpAddrType, docsPnmBulkDestIpAddr, and docsPnmBulkDestPath above need to be set to finish the file upload config.

Figure 118 and Figure 119 are screenshots captured in the lab.

Because OFDMA is only now beginning to be deployed, current use of upstream RxMER-per-subcarrier data is limited to profile management. It is used with active subcarrier, excluded subcarrier, and unused subcarrier information to calculate modulation orders for each OFDMA mini-slot as OFDMA interval usage codes.




Figure 118 - Upstream RxMER-per-Subcarrier Data Plotted by Frequency for a Single CM at Four Different Times




Figure 119 - Upstream RxMER-per-Subcarrier Data Plotted in 3 Dimensions by Frequency for a Single CM at Four Different Times

6.3.2 Conclusion and Suggestions for Future Work

This measure has not been well developed as it is not generally available according to specification methods. This section uses an implementation-specific method to gather the data so that operators can see a proof of concept from which to build. This section will be updated as the method becomes more frequently available in the field.

6.4 Upstream Triggered Spectrum Capture

6.4.1 Overview

Upstream spectrum analysis is one of the most important tools cable operators rely on to troubleshoot RF interference in the return band. But given the return funnel⁸ effect in cable networks, locating ingress sources continues to be inference-based, which is manual and tedious. Upstream triggered spectrum capture helps improve and automate this process.

The upstream triggered spectrum analysis measurement provides a wideband spectrum analyzer function by using the CMTS burst receiver, which has knowledge of the scheduler. This is important because it can be triggered to examine desired upstream transmissions as well as underlying noise/interference during a quiet period when no devices are scheduled to transmit. It can also be triggered to capture on specific timestamps or time slots granted to a specific MAC address.

In addition to having scheduler information available, the CMTS also has demodulators that run continuously on active upstream channels. These demodulators can provide additional information for each of the channels to the spectrum analyzer, such as RxMER and FEC statistics, that is useful for troubleshooting. These data can be used to help automate and assist manual troubleshooting, such as by changing color indicators from red to green when certain thresholds are exceeded.

In Figure 120 and Figure 121, the spectrum traces of the respective channels are colored green and red according to their channel performance based on RxMER and FEC.

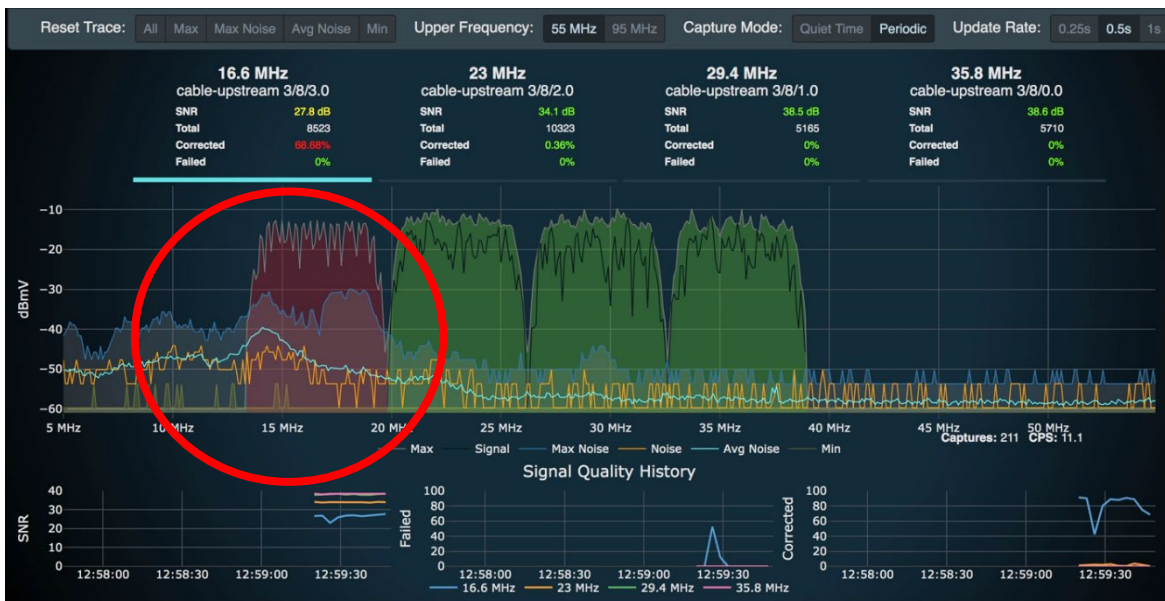


Figure 120 - Upstream Spectrum Capture of Multiple Upstream SC-QAM Channels

⁸ Also referred to as noise funneling. Downstream signal transmission in a cable network is point to multipoint (headend or hub site to multiple subscribers); upstream transmission is the reverse. Upstream interference such as noise or ingress can enter the network at multiple locations and "funnel" back to the headend or hub, hence the reverse funnel effect.

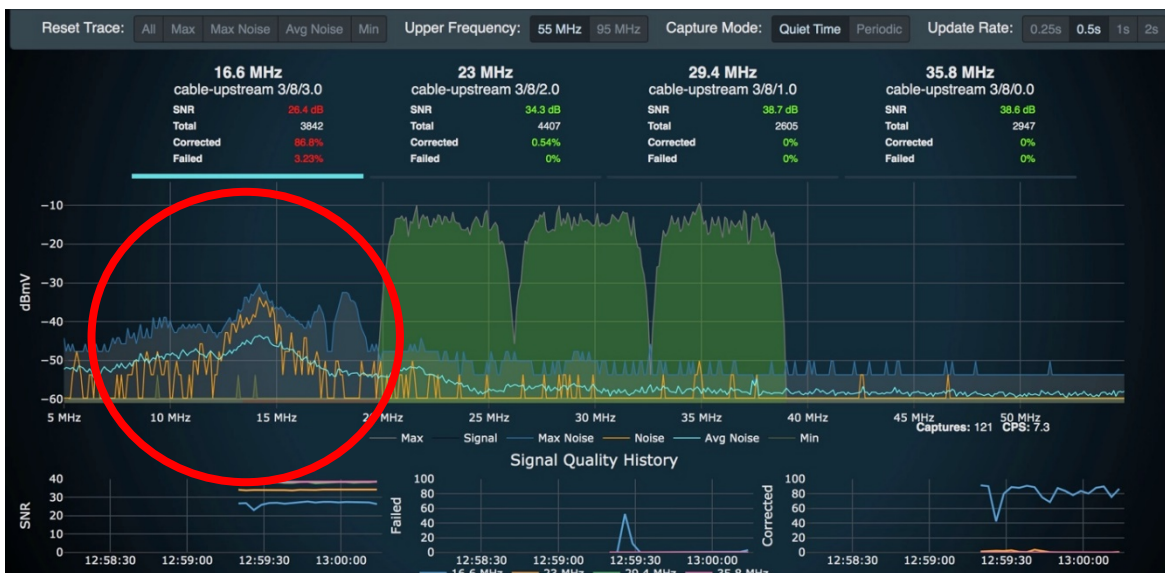


Figure 121 - Upstream Spectrum Capture of Quiet Time

In Figure 121, the first SC-QAM signal does not appear because the capture occurred precisely when there was no scheduled upstream transmission. This allows noise power to be measured and the noise floor to be seen without the presence of the SC-QAM signal on that frequency.

6.4.2 How to Implement

Vendor proprietary implementations and standardized CableLabs MIBs allow operators to take advantage of this feature. Consult equipment manufacturers for the former. For the latter, see [CM-OSSIV3.1] and [CCAP-OSSIV3.1].

6.4.3 Conclusion and Suggestions for Future Work

Triggered spectrum capture is in use in CMTS line cards and in remote PHY equipment. Test equipment companies have been integrating this capability into their products for some time, especially in distributed access architectures where there is no RF returning to the headend or hub.

However, some CMTS manufacturers may not be following the specification to provide this feature in favor of proprietary approaches. Therefore, significant effort or proprietary tools are required to benefit from upstream spectrum capture methods. In general, there is some movement toward the methods outlined in the specifications.

6.5 Upstream Impulse Noise Statistics

6.5.1 Overview

This section describes the operation of the measurement of upstream impulse noise statistics.

6.5.2 Description of Measurement

The upstream impulse noise statistics measurement gathers data about burst and impulse noise occurring in a selected narrow band as defined in [CCAP-OSSIV3.1]. To perform the measurement, a bandpass filter in the CMTS is positioned in an unoccupied upstream band. A threshold is set, and energy exceeding the threshold triggers the beginning of event measurement. Once an event has started, energy in the band is continuously monitored, and energy falling below the threshold causes the measurement to stop and the statistics of the event to be recorded. The CMTS may also allow the threshold to be set to zero, in which case only the average power in the band will be measured. The measurement is timestamped using the 32-bit DOCSIS 3.0 timestamp (bits 9-40 of the 64-bit extended timestamp, where bit 0 is the LSB), which provides a resolution of 98 ns and a range of 7 minutes. The CMTS provides the capability to capture the following statistics in a selected upstream band up to 5.12 MHz wide:

- timestamp of event,

- duration of event, and
- average power of event.

The CMTS provides a time history buffer of these three measurements for up to 1,024 events.

6.5.3 Conclusion and Suggestions for Future Work

This measurement has not been field tested, other than proprietary implementations. Additional information will be provided in a future version of this document.

6.6 Upstream FEC Statistics

6.6.1 Overview

This section describes the upstream FEC statistics.

6.6.2 Description of Measurement

The following text is taken from [PHYv3.1], Section 9.4.5, "Upstream FEC Statistics."

Upstream FEC statistics provide monitoring upstream link quality via FEC and related statistics. Statistics are taken on codeword error events. An LDPC codeword that fails post-decoding syndrome check will be labeled "unreliable", but the data portion of the codeword may not contain bit errors; hence the "unreliable codeword" count will tend to be pessimistic. All codewords, whether full-length or shortened, are included in the measurements. The codeword (or frame) and error counters are long (64-bit) integers, so that rollover is not an issue.

The CMTS reports the following FEC statistics for any specified single upstream channel.

- Pre-FEC Error-Free Codewords: Number of codewords that passed pre-decoding syndrome check.
- Unreliable Codewords: Number of codewords that failed post-decoding syndrome check.
- Corrected Codewords: Number of codewords that failed pre-decoding syndrome check, but passed post-decoding syndrome check.
- MAC CRC failures: Number of frames that failed MAC CRC check.
- Total number of FEC codewords.
- Total number of MAC frames.
- Start and stop time of analysis period, or time that snapshot of counters was taken.
- SID corresponding to upstream user being measured.

The CMTS reports the following FEC summaries over a rolling 10 minute period for any single upstream channel:

- Total number of seconds.
- Number of errored seconds (seconds during which at least one unreliable codeword occurred).
- Count of codeword errors (unreliable codewords) in each 1-second interval (600 values over 10 minutes).
- Start and stop time of summary period.

6.6.3 Background Theory

At the CMTS, codewords are first assessed to determine if they are error free. If error free, they are counted as pre-FEC error-free codewords. If they contain errors, the error correction of LDPC is applied to fix them. If a codeword can be fixed, it is counted as a corrected codeword. If it cannot be fixed, it is counted as an unreliable codeword.

Independent of whether errors are fixed or not, codewords are stripped of their redundant bits and reassembled into their original MAC frames. Once they are back in MAC frames, the CMTS uses the CRC to confirm if the MAC frame is error free. If the MAC frame contains errors, it is counted as a MAC CRC failure. It is important to note that, from the customer's perspective, the remaining data will be perfect if the CMTS's MAC frames are error free.

For each CM, the CMTS keeps counters for every type of codeword and MAC frame error listed above as well as for both the total number of codewords the CMTS received and the total number of MAC frames the CMTS received.

6.6.4 Rolling FEC Summaries

With the 10-minute rolling record of errors in one-second intervals, the time characteristics of the errors may identify the possible source of the interference. For example, if correctable errors occur in every sample interval, then the possible source may be related to electrical arcing from utility power. If uncorrectable errors occur in several consecutive seconds, then the possible source may be another data transmission. If long periods of errors are separated by long periods of no errors, then the possible source may be a two-way radio transmission.

Examples are not yet available and will be provided in a future version of this document.

6.6.5 Conclusion and Suggestions for Future Work

Measurement has not been fielded. Additional information will be provided in a future version of this document.

6.7 Upstream Spectrum Capture at the CM

This section describes the use of full-band capture (FBC) in the upstream frequency range to isolate noise sources.

6.7.1 Overview

Traditional upstream capture techniques are obtained at the burst receiver in the CMTS or RPD, which is subject to return noise funneling. This measurement is useful to manually troubleshoot noise, but it is particularly difficult to locate noise sources. The difficulty is a side effect of HFC designs that use diplex filters and directional coupling to facilitate upstream communication in lower frequency RF spectrum.

Cable modem upstream capture allows operators to sample the upstream RF spectrum at the cable modem. This capability is important because the upstream RF spectrum is also subject to directional coupling such as port-to-port isolation in the splitters and taps. The source can be quickly identified and located if the noise is entering at the connection of the cable modem, ground block, drop, or tap connections.

Figure 122 shows a typical spectrum capture of the upstream without any discernable impairments in the return band. In contrast, Figure 123 shows observable noise near and around the diplex cutoff frequency and interference within the downstream SC-QAM, including FM radio. This result is a good indication that noise is likely present in the return band, which the diplex filter is attenuating.




Figure 122 - CM Upstream Capture with Diplex Rejection




Figure 123 - CM Upstream Capture with Diplex Rejection and Noise

6.7.2 Description of Measurement

See Section 5.2 for descriptions of full-band capture.

6.7.3 Conclusion and Suggestions for Future Work

Although the upstream frequency range is not precluded from FBC in the DOCSIS suite of specifications, it has limited usefulness in the presence of a diplex filter. There are cable modem reference designs that allow vendors to implement the spectrum capture in a manner that bypasses the diplex filter without risk of leakage to the diplex filter circuit. However, this capability is optional and may have limited value until wider adoption is achieved.

6.8 Upstream Histogram

6.8.1 Overview

This section describes the upstream histogram measurement. See Section 5.9 on downstream histograms for a description of the measurement; the specifics are the same except as described in that section.

6.8.2 Example

As noted in Section 5.9, one-sided hard clipping starkly illustrates the effect of ideal laser clipping. In reality, the result is more subtle because of softer, more gradual clipping or other effects such as bandlimiting. Figure 124 shows the histogram of a signal that has passed through an actual upstream cable plant with laser clipping. The upper (rightmost) tail of the distribution, though not completely cut off, has been reduced noticeably below the reference Gaussian distribution. The missing samples from the tail are evident in the fairly pronounced spike on the right side of the distribution at around 1.25 standard deviations on the x axis. Although the effect is not as pronounced as in an ideal simulation, this histogram could be analyzed by the appropriate algorithm to show a strong indication of laser clipping, especially if a longer set of data is averaged.




Figure 124 - Histogram of Actual Upstream Clipping Event

6.8.3 Conclusion and Suggestions for Future Work

The upstream histogram measurement has not yet been developed in the field such that it adheres to specifications. Additional information will be provided in a future version of this document.

6.9 Upstream Capture for Active and Quiet Probes

6.9.1 Overview

This section describes the operation of the upstream capture measurement for active and quiet probes.

6.9.2 Description of Measurement

The purpose of upstream capture is to measure upstream cable plant response and to view the underlying noise floor. It does so by capturing at least one OFDMA symbol during a scheduled active or quiet probe (the quiet probe captures the symbol period but no OFDMA symbol).

The PNM server selects an active CM to analyze by specifying its MAC address or requests a quiet probe measurement. The CMTS selects a specified transmitting CM or quiet period when no CMs are transmitting for the capture. The CMTS sets up the capture as described in [MULPIv3.1], selecting either an active SID corresponding to the specified MAC address or the idle SID and defining an active or quiet probe. The active probe symbol for this capture normally includes all non-excluded subcarriers across the upstream OFDMA channel and has pre-equalization off. The quiet probe symbol normally has all subcarriers off; that is, there are no transmissions in the given upstream OFDMA channel during the quiet probe time. The CMTS captures samples of a full OFDMA symbol including the guard interval. The CMTS begins the capture with the first symbol of the specified probe. The sample rate is the FFT sample rate (102.4 megasamples per second).

To ensure the input to the channel is known, the CMTS reports the list of excluded subcarriers, cyclic prefix length, and transmit window roll-off period in order to fully define the transmitted waveform. The CMTS also reports the index of the starting sample used by the receiver for its FFT, which helps specify the correct timing of the captured samples. The CMTS reports the timestamp corresponding to the beginning of the probe, which provides information as to when the measurement was made, for comparison between measurements and to correlate with plant conditions at different times. When the P-MAPs for the OFDMA upstream being analyzed are being sent in an OFDM downstream, the timestamp reported is the 64-bit extended timestamp. When there are OFDMA upstream channels but no OFDM downstream channels, the reported timestamp is the 32-bit DOCSIS 3.0 timestamp (bits 9-40 of the extended timestamp, where bit 0 is the LSB). For an active probe, the CMTS also reports the contents of the Probe Information Element (P-IE) message, which gives detailed information describing that probe.

This is a time domain capture normally reported in the time domain to preserve the cyclic prefix.

6.9.3 Background Theory

An active probe provides the partial functionality of a network analyzer because the input is known and the output is captured, permitting full characterization of the linear and nonlinear response of the upstream cable plant in the OFDMA band being measured. A quiet probe provides an opportunity to view the underlying noise and ingress while no traffic is being transmitted in the OFDMA band being measured.

6.9.4 Conclusion and Suggestions for Future Work

At the time of publication, the measurement has not been field tested. Work is ongoing to implement the measurement and provide test results. This section will be updated as this measurement and its use cases develop.

7 EXAMPLES OF USE CASES

7.1 Use Cases Versus Measurements

Figure 125 was updated from [Currivan and Wolcott] and provides an example for a number of use cases. Future versions of this document will contain additional details about use cases.

THIS PNM MEASUREMENT ADDRESSES THESE USE CASES														and these repair use cases					
	Local micro-reflections	Spectrum notches, ripple, rolloff, filters, adjacency	Group delay distortion	Non-flat noise under carrier	Laser misalignment, amplifier compression	Digital CFO/CSO/CTB	Raised thermal noise floor	Ingress, spurious, FM, LTE	QAM carrier performance issues	HI or low received power	Disturbance interference: gpon lights, motor, switching PS	Disturbance/noise transmitted by CM	Automatic data gathering to find impairments	Manual data gathering to find impairments	Verification of impairment or impact	Remote isolation (local ton) of impairment	Remote troubleshooting service (pro active or reactive)	Field isolation (local ton) of impairment	Field troubleshooting service (pro active or reactive)	Verifying repair of service
DS Symbol Capture	x	x	x	x	x	x	x	x	x	x	x	x	x	x						
DS Spectrum Analysis in CM	x				x	x	x	x	x					x	x	x				
Noise Power Ratio					x			x												
Channel Estimate Coefficients	x		x			x								x	x					
Constellation Display					x			x	x	x	x				x					
ModMER Per Subcarrier	x			x			x	x	x	x				x	x	x				
Histogram					x			x							x					
FEC Errored Seconds and Minutes							x		x					x	x					
CM Received Power																				
US Spectrum Analysis at CMTS	x		x	x	x	x	x	x	x	x	x	x								
US Received Power at CMTS																				

Figure 125 - Use Cases Addressable by PNM Measurements

7.2 LTE and Other Interference

Coaxial cables and other components used in the distribution and subscriber drop portions of a cable network provide a shielded transmission medium that is independent of the over-the-air environment. Over-the-air radio frequencies are allocated to various services by government agencies, but cable operators largely have the ability to use frequencies within their closed networks as they see fit.

As long as the shielding integrity of a cable network is maintained, frequencies inside of the network’s cables and components can be used for different purposes than what they are used for in the over-the-air environment. Should the shielding integrity be compromised for any reason, over-the-air signals can leak into the cable network and interfere with signals and services carried by that network, and signals inside the cable network can leak out into the over-the-air environment and potentially cause interference.

7.2.1 Overview

Long-Term Evolution (LTE), or 4G, is a standard that describes high-speed wireless data telecommunications, typically between user equipment (UE) such as smart phones, tablets, and other mobile devices and fixed tower and similar sites. LTE operating frequencies range from 450 MHz to more than 3 GHz, depending on country, region, and LTE band. Other frequencies in the over-the-air environment are used for services such as broadcast radio and television, two-way radio, and aircraft navigation. Of particular interest to cable operators is RF signal transmission on over-the-air frequencies that overlap frequencies used in cable networks.

Signal leakage, or egress, occurs when RF signals inside the coaxial cables and other cable network components can leak out into the over-the-air environment and interfere with licensed services such as LTE. Ingress occurs when over-the-air signals leak into a cable network and interfere with the cable network’s signals. (A variation of ingress

is called direct pickup interference.) In either scenario, the cable operator, not the provider of licensed over-the-air service, is responsible for resolving the interference.

The ability to identify, troubleshoot, and locate potential sources of ingress and direct pickup (and perhaps signal leakage) is a valuable part of PNM.

7.2.2 Signal Leakage

Cable network signal leakage is regulated by the respective federal government in many countries. If signal leakage causes harmful interference to over-the-air signals, the cable operator can be subject to fines and/or temporary suspension of the ability to use the affected frequencies until the harmful interference is eliminated. Figure 126 shows an example of signal leakage from a cable network, with SC-QAM signals overlapping LTE signals.



Figure 126 - Cable Network Signal Leakage in the 700 MHz LTE Spectrum, Overlapping LTE Band 13 (Verizon) Tower-to-UE Downlink Signals

Although PNM tools cannot directly show signal leakage or provide a measurement of the leakage field strength, it can detect the presence of ingress, as discussed in the next section. If a cable network has ingress problems, it likely has signal leakage problems.

7.2.3 Ingress and Direct Pickup

Ingress is the unwanted entrance of over-the-air signals into a cable network because of degraded shielding effectiveness of the network's coaxial cables and other components. Direct pickup, similar to ingress, occurs when the over-the-air signal(s) directly enters a susceptible set-top box, cable modem, television set, test instrument, or other device directly without any cables or other external devices being physically connected. If the susceptible device's outer case or cover is inadequately shielded, then the internal wiring, printed circuit board traces, and/or components can directly receive interfering over-the-air signals.

PNM tools can detect the presence of ingress interference and direct pickup interference in a few ways. One is through full-band capture; the example in Figure 127 shows ingress from several sources, including LTE.




Figure 127 - Full-Band Capture Spectrum from a Modem Showing Several Types of Ingress

A second way to use a PNM tool to detect the presence of ingress or direct pickup interference is to monitor the RxMER on several frequencies. SC-QAM signals affected by ingress that might or might not be readily visible on a spectrum display likely will have degraded RxMER compared to SC-QAM signals not affected by ingress or direct pickup interference. For DOCSIS 3.1 downstream OFDM signals, a plot of RxMER per subcarrier can show the effect that LTE or other ingress has on subcarriers that overlap the OFDM signal. See Figure 128. In this example, ingress interference is visible where the RxMER drops abruptly, in the vicinity of 740 MHz, 760 MHz, and 870 MHz. It also shows evidence of possible narrowband interference at other frequencies.




Figure 128 - OFDM RxMER-per-Subcarrier Plot Showing the Impact of LTE and Other Types of Ingress

Another method that can be used to detect the presence of ingress or other interference under digital signals, described here for informational purposes, is based on deriving a spectrum from the signals' error vectors. The resulting error vector spectrum shows noise and ingress beneath the signal without having to turn off the signal. Some cable industry test equipment vendors have incorporated variations of this technique in their products. Figure 129 is an example screen capture from a DOCSIS 3.1 cable modem-equipped Viavi field meter showing the noise floor (black trace) under an OFDM signal and the lower and upper adjacent SC-QAM signals (blue trace).




Figure 129 - "Ingress Under the Carrier" Display from a Field Meter

Detection and identification of ingress and direct pickup on a per-CPE basis can be incorporated in a PNM tool or application to help cable operators locate where the ingress and/or direct pickup interference is entering the network.

7.2.4 How to Implement

Several ingress and egress detection methods for downstream and a few for upstream are described in other sections of this document, including the sections on RxMER per subcarrier (Section 5.3) and spectrum analysis (Section 5.2).

7.2.5 Conclusion and Suggestions for Future Work

The ability to troubleshoot downstream ingress and similar interference is an important part of PNM. At least two of the techniques described in this section—full-band capture and OFDM RxMER-per-subcarrier data—lend themselves to being a useful component of a PNM tool. Downstream ingress and direct pickup interference can degrade cable network signal quality. If ingress and/or direct pickup exists, it is likely that signal leakage is also present, which can potentially cause interference to over-the-air services. The ability to identify and locate potential sources of ingress can enable outside plant personnel to quickly fix problems.

Future work in this area might include incorporating error vector spectrum techniques in PNM tools to further support identification and troubleshooting of ingress and direct pickup interference.

7.3 Fault Localization Based on Downstream RxMER

The PNM Best Practices guide for DOCSIS 3.0 systems [PNMP-3.0] outlined a fault localization process based on correlating pre-equalization or full-band capture data from cable modems with their network topology to identify the point along their shared path where a fault or anomaly indicator changes. Understanding where this change occurs gives an indication as to the location or area of the cable plant that must be contributing to the fault. Based on field observations subsequently validated in test lab environments, downstream RxMER-per-subcarrier data can be used in a similar way to help locate faults.

The vicinity of the impairment can be identified by isolating end devices sharing a common path that show a ripple in the RxMER per subcarrier caused by a micro-reflection or ingress interference resulting from a break in the integrity of the cable plant. Comparing the reported RxMER per subcarrier of devices on the same shared path can show a boundary between where a problem exists and where it does not. The fault, therefore, must be between the last device showing the fault and the first device not showing the fault.

7.3.1 Fault Localization Examples

This section details basic examples of locating faults based on the reported downstream RxMER per subcarrier. Additional details for these examples can be found in [Hranac and others].

7.3.1.1 Downstream RxMER Ripple

Figure 130 shows a location where all of the cable modems connected to the amplifier (A) show a significant ripple in a plot of their downstream RxMER-per-subcarrier data. This ripple, however, is not seen on any cable modem connected to any other output of the upstream amplifier (B). This analysis would therefore indicate that the cause of the impairment is therefore located on the trunk cable (C) or connectors connecting amplifiers (A) and (B).

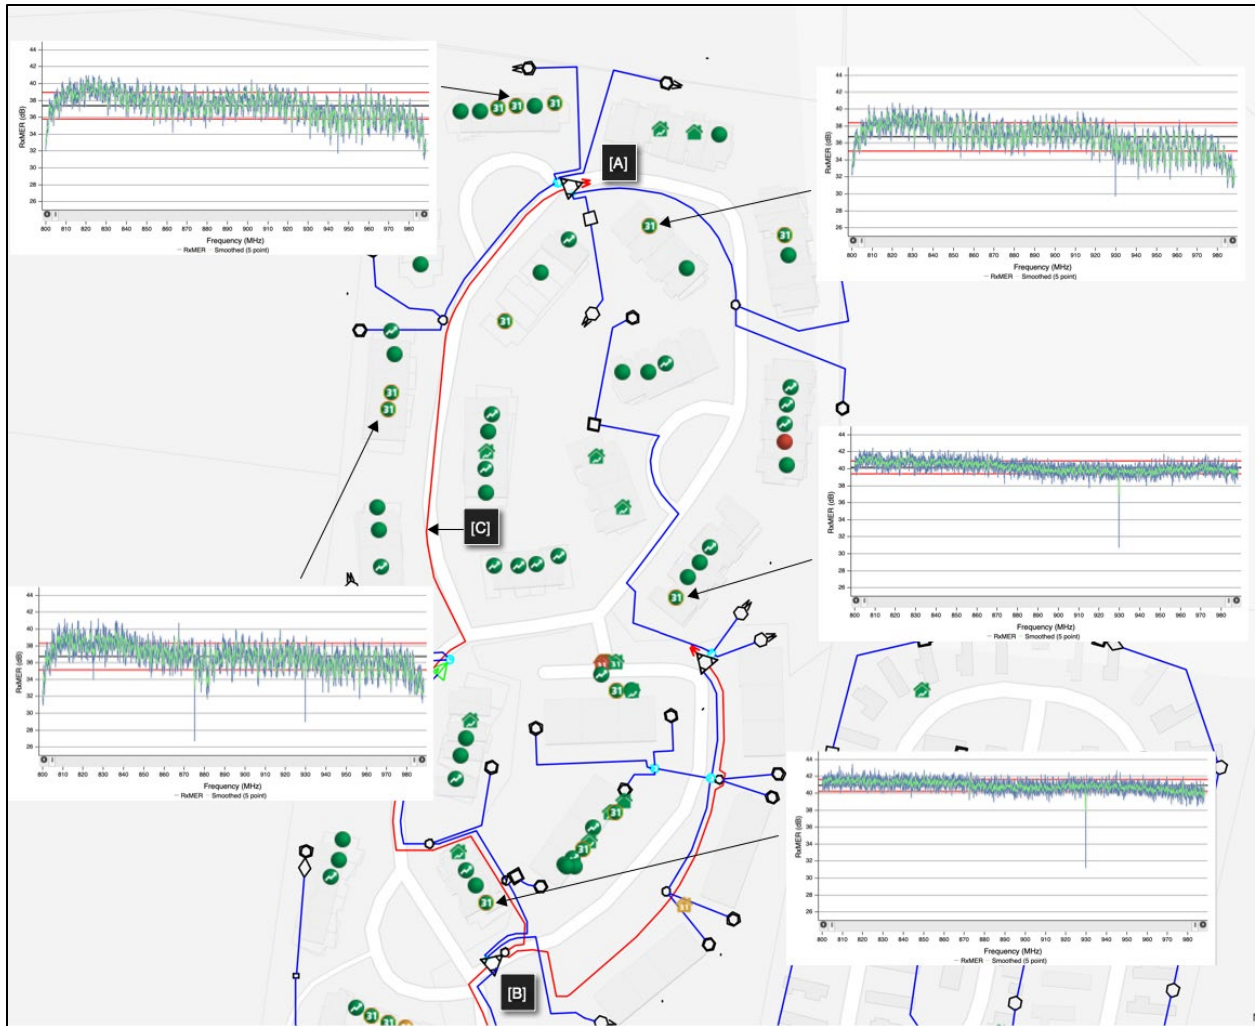


Figure 130 - Example of Downstream RxMER Data Showing Ripple Location

Map data ©2019 Google; ©Akleza, Inc.

7.3.1.2 LTE Ingress

Figure 131 is an example of downstream RxMER-per-subcarrier data showing interference affecting a group of cable modems all downstream of amplifier (A). Data from the cable modems connected on other outputs of the upstream amplifier (B) do not show the same interference. The results of comparing the isolated group of impacted cable modems and the group not showing the interference implies that the interference is related to a fault in the trunk cable feeding amplifier (A), the connection into amplifier (A), or the amplifier itself.

A closer look at the impacted frequencies shows an overlap with the LTE Band 5 downlink transmission spectrum. A survey of the area shows an antenna tower (C) close to amplifier (A), suggesting a possible integrity issue at or around the amplifier is allowing the LTE signal to interfere with the DOCSIS 3.1 OFDM channel.

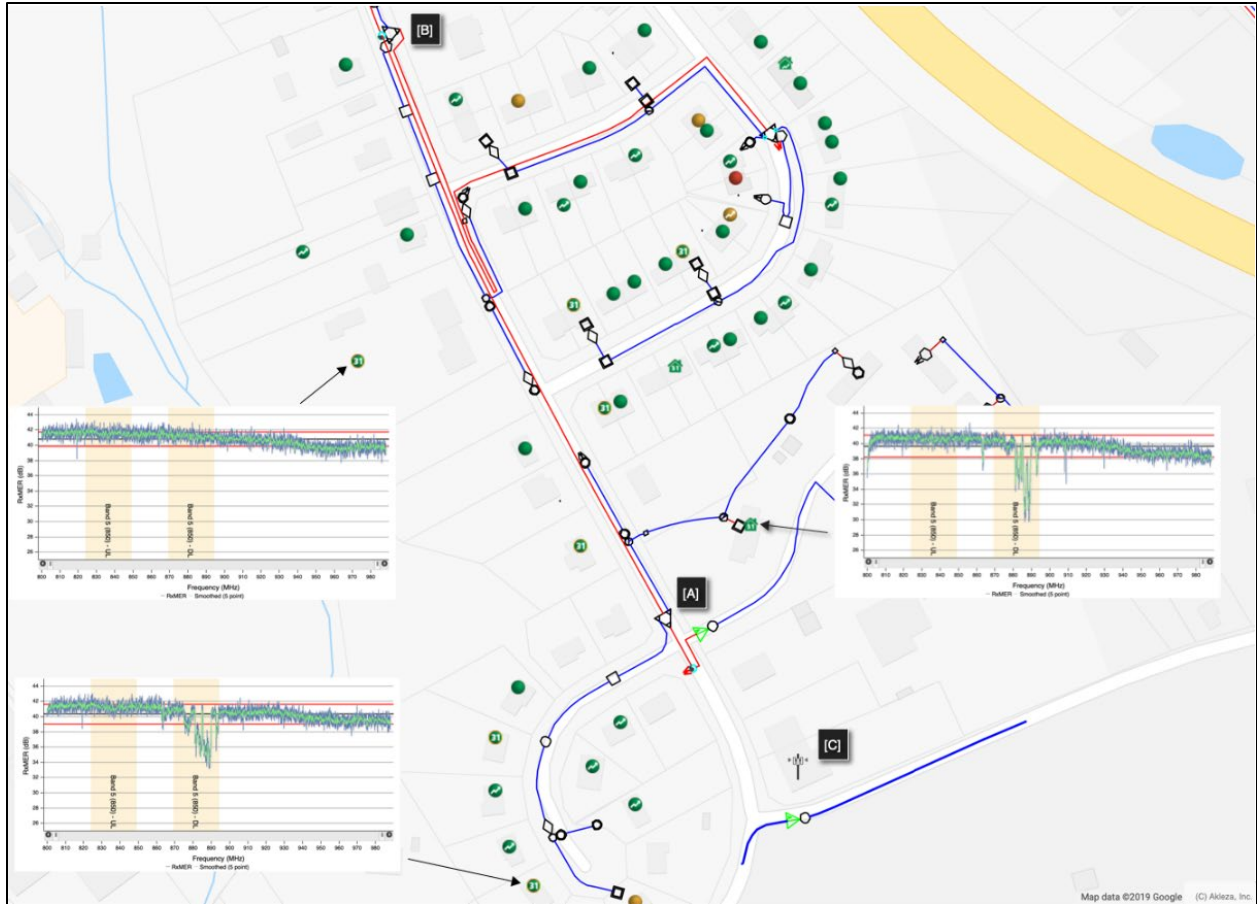


Figure 131 - Example of Downstream RxMER Data Showing LTE Ingress Location

Map data ©2019 Google; ©Akleza, Inc.

8 PNM DATA ISSUES

8.1 Data Collection Methods

In DOCSIS 3.0 systems, pre-equalization coefficient and full-band capture (FBC) data are obtained by using SNMP to request information from CMs and CCAPs. The requested information is described in [PNMP-3.0], and the SNMP MIB objects are defined in the [DOCS-IF3-MIB].

DOCSIS 3.1 technology introduced a new set of requirements on the CM and CMTS related to the collection and reporting of PNM data (see [PHYv3.1], Section 9, "Proactive Network Maintenance"). DOCSIS 3.1 technology also introduced a new Bulk-Data Transfer mechanism that uses TFTP to upload PNM data to a destination TFTP server. Configuration of the PNM test execution, file storage, and TFTP destination uses SNMP as defined in [CM-OSSIV3.1] and [CCAP-OSSIV3.1]. Figure 132 illustrates the steps and communication paths involved in collecting PNM data from a CM or CCAP device.

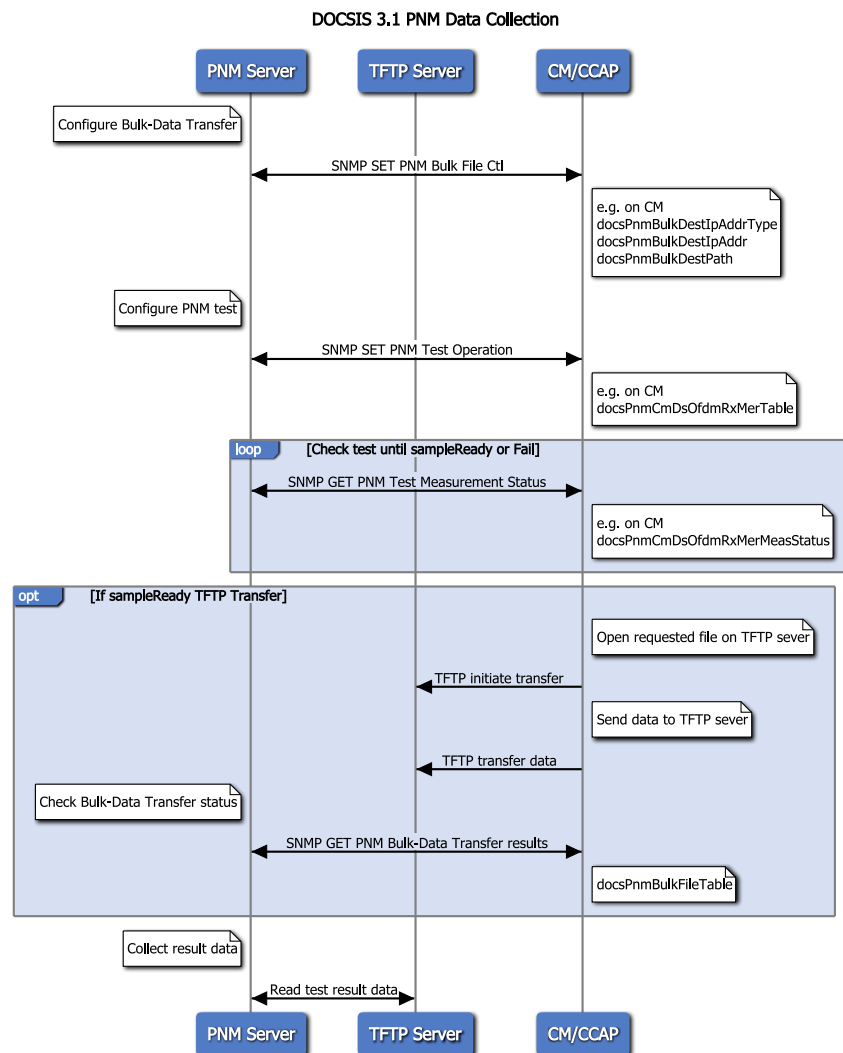


Figure 132 - Data Collection Sequence Diagram

The result of a successful PNM test and file upload is a PNM test result file that a PNM application can read and analyze. By default, each PNM test result file name contains the PNM test, a device identifier such as the CM MAC address, and a timestamp. The PNM test file is a binary data file containing a header and a data section as defined in [CM-OSSIV3.1] and [CCAP-OSSIV3.1].

For example, a successful downstream RxMER test against a CM with the MAC address 480033DDF724 and executed on Friday, March 22, 2019, 8:30:57 AM GMT would generate a PNM test result file with the name PNMDsMer_480033DDF724_1553243457. This file is fully configurable. By parsing the data contained in the file, a chart similar to that shown in Figure 133 can be plotted. This figure shows the RxMER per subcarrier across a DOCSIS 3.1 OFDM channel.




Figure 133 - Example of Downstream RxMER per Subcarrier

8.1.1 Operational Considerations

The use of TFTP by the DOCSIS 3.1 PNM Bulk-File Transfer mechanism may require some network configuration changes if any firewalls are between the CM or CCAP and the destination TFTP server. TFTP is a UDP-based protocol that uses ephemeral ports to transfer data. Figure 134 shows the upload sequence in which Ports X and Y are ephemeral port numbers.

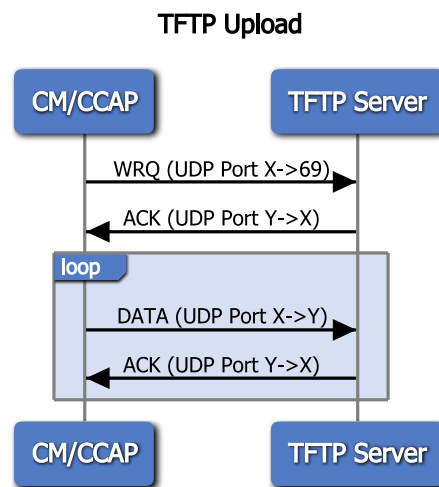


Figure 134 - TFTP Upload Diagram

For the TFTP upload process to complete, any network firewalls between the CM or CCAP and the TFTP server must be configured to allow UDP traffic from the ephemeral port range between the CM or CCAP and TFTP server to pass. Though it may be possible to use source and destination addresses to configure an explicit rule for CCAP devices, it is not practical for CM subnets. As such, rules allowing UDP Port 69 and all UDP traffic in the ephemeral port range destined for the TFTP server are required to support the PNM Bulk-File Transfer mechanism.

8.1.2 Security and Reliability Considerations

TFTP and SNMP are currently not specified to be encrypted. However, one can add application layer security if desired. For example, HTTPS proxy can be added in front of data collectors like the Common Collection Framework (CCF), which allows encrypted posts and gets.

UDP, which is low priority, is the assumed transport mechanism. As such, a capacity limit is reached if there is too much traffic, and the UDP packets can be dropped, causing a loss of PNM information. Further, there is a general assumption that if a test fails or if information is corrupted in the transfer process, then a retry of the test is warranted. The application layer (in this case the PNM application) is assumed to manage the information reliability.

In addition, CMs can be limited in their processing of PNM information. If a PNM request takes 5 minutes to complete and is only initiated when requested, then the CM will not likely respond to other requests until the first is completed. Consider this when requesting data from CMs.

8.1.3 Deployment Considerations

The network needs to have a TFTP server to receive and hold the TFTP files that are responses to the data requests.

Also, firewall rules on the network must allow the TFTP packets to travel from the CM to the server and to any applications that need the data.

Similarly, SNMP requires firewall rules to be configured to allow the traffic in both directions.

Virtual machines (VMs) are commonly used to deploy PNM data collectors such as CCF. The VM should have enough cores and be configured properly for the deployment size. It is common to have one PNM data collector instance for each CMTS core, but different models are possible and should be considered based on expected use.

8.1.4 IPDR Bulk Data collection

Although the current DOCSIS specifications that cover PNM require SNMP and TFTP mechanisms for data delivery, some cases require upstream data from the CMTS be obtained via IPDR.

Spectrum data are also specified for SNMP by DOCSIS 3.0 technology, but billing data is in IPDR only. IPDR may require a command-line interface to the CMTS.

8.2 Measurement Calibration

Discussion of this topic in a future version of this document is being considered. For now, it is acknowledged that calibration is an important factor when comparing measurements from different sources.

9 PNM PRACTICES

9.1 Measures of Access Network Health

This section provides a starting point for measuring access network health, which can be used as a performance measure for network operations, and potentially as a measure for selecting maintenance to prioritize.

There are two related concepts considered in this section: the quality of downstream (and upstream) plant, and the quality of downstream (and upstream) service with respect to the plant condition. This section first focuses on downstream measures, but also mentions briefly how to address this question for the upstream too.

General guidelines in summary are as follows.

- Quality of downstream plant – Start with RxMER per subcarrier, translate to bit loading, and use RxMER stats to assess when spectral efficiency is a concern. This is currently accessed on the CM via the CmDsOfdmRxMer MIB. But to access the actual bit loading, the table CmDsOfdmModProfile is used. A new measurement, it provides modulation profile by subcarrier. CM firmware is being updated as of this publication date, so this capability may become more common soon. There should be considerations for input signal level and FEC stats as well. Too low or high of a signal level can indicate problems, and very high correctable codewords or any uncorrectable codewords can indicate vulnerability to smaller impairments. To determine outliers, the values need to be compared to expectations, which should come from design guidelines, and may be specific for each operator.
- Quality of downstream service with respect to downstream plant – Use RxMER margin for the given profile for the CM, then sum the margin over all subcarriers. This is currently managed with the CmDsOfdmMerMargin MIB. CmDsOfdmModProfile data may be of use for this measurement too. Considerations should be given to the service provided to the customer too. More specifically, a low latency service may not tolerate retries as much as typical internet service.
- Averages over any period of time consisting of at least 10 samples, using statistical average or exponentially weighted moving average, for a stable measure. This would be managed outside the CM by the cable operator retaining the measures above over time and then assessing the data over that period to determine patterns and impacts over the period. Operators typically sample PNM statistics each hour, though some sample less frequently still. This is mostly a factor of the maturity and capability of the data collection systems available to the operators.
- Averages over any population or random sample from a population of CMs using statistical average to get estimates for the plant section. Samples should include at least 50% of the population for populations under 100, and at least the maximum of 20% or 50 CMs for larger populations (assumed, should be tested, should be sufficient). This item would be managed outside the CM by the cable operator retaining the measures above over time and then assessing the data over that period to determine patterns and impacts over the period. These measures should generally be collected over a given node.

9.1.1 Quality of downstream plant (for each OFDM signal)

9.1.1.1 RxMER per subcarrier approach

1. Retrieve RxMER per subcarrier data (say once per hour or every two hours) using the CmDsOfdmRxMer MIB.
2. Calculate the bit loading appropriate for each subcarrier, applying any operating margin desired⁹. Use the CmDsOfdmRequiredQamMer object to determine the RxMER target values required in each subcarrier, then determine the appropriate bit loading by comparing the RxMER values to the target values, per subcarrier. If docsPnmCmDsOfdmMerMarTable is not available, the default values from the specification are provided below. (Note that, when available, this can be done (or confirmed) using CmDsOfdmModProfile table.)
3. Translate into maximum symbol size if and as desired.
4. Use any of these three as the measure of downstream plant quality.

⁹ See section I.3 of this technical report which discusses FEC and RxMER per subcarrier test results and guidelines for operators for more information on this subject.

The bit loading recommendations from the current OSSI specification [CM-OSSIV3.1], and the similar table of CMTS Minimum CNR performance in AWGN channel from , are shown in the tables below, though operators may choose to adjust these recommendations to their own needs.

Table 8 - CmDsOfdmRequiredQamMer Object from [CM-OSSIV3.1]

Attribute Name	Type	Access	Type Constraints	Units	Default Value
ReqMerQam16	UnsignedInt	R/W		quarterDb	60
ReqMerQam64	UnsignedInt	R/W		quarterDb	84
ReqMerQam128	UnsignedInt	R/W		quarterDb	96
ReqMerQam256	UnsignedInt	R/W		quarterDb	108
ReqMerQam512	UnsignedInt	R/W		quarterDb	122
ReqMerQam1024	UnsignedInt	R/W		quarterDb	136
ReqMerQam2048	UnsignedInt	R/W		quarterDb	148
ReqMerQam4096	UnsignedInt	R/W		quarterDb	164
ReqMerQam8192	UnsignedInt	R/W		quarterDb	184
ReqMerQam16384	UnsignedInt	R/W		quarterDb	208

Table Note: the table above reports integer values of ¼ dB values, so translation to dB values may be needed in some cases. For example, ReqMerQAM128 reports a value of 96 in ¼ dB, which is equivalent to 96/4 =24 dB.

Table 9 - CMTS Minimum CNR Performance in AWGN Channel from [PHYv3.1]

Constellation	CNR ^{1,2} (dB)	Set Point (dBmV/6.4 MHz)	Offset
QPSK	11.0	-4 dBmV	0 dB
8-QAM	14.0	-4 dBmV	0 dB
16-QAM	17.0	-4 dBmV	0 dB
32-QAM	20.0	-4 dBmV	0 dB
64-QAM	23.0	-4 dBmV	0 dB
128-QAM	26.0	0 dBmV	0 dB
256-QAM	29.0	0 dBmV	0 dB
512-QAM	32.5	0 dBmV	0 dB
1024-QAM	35.5	0 dBmV	0 dB
2048-QAM	39.0	7 dBmV	0 dB
4096-QAM	43.0	10 dBmV	0 dB

Table Notes:

Note 1. CNR is defined here as the ratio of average signal power in occupied bandwidth to the average noise power in the occupied bandwidth given by the noise power spectral density integrated over the same occupied bandwidth.

Note 2. Channel CNR is adjusted to the required level by measuring the source inband noise including phase noise component and adding the required delta noise from an external AWGN generator.

Note 3. The channel CNR requirements are for OFDMA channels with non-booster pilots. For operation with booster pilots, which is optional at the CMTS, the CNR requirements are increased by

- 1 dB for channels with 50 kHz subcarrier spacing, and
- 0.5 dB for channels with 25 kHz subcarrier spacing. For example, the CNR requirement for QPSK with booster pilots is 12.0 dB with 50 kHz subcarrier spacing and 11.5 dB with 25 kHz subcarrier spacing.

9.1.1.2 RxMER stats approach

1. Poll RxMER stats at a suitable rate using CmDsOfdmRxMer MIB.
2. Identify any outlying values of 2nd percentile or 5th percentile statistics, and poll RxMER per subcarrier, then use the RxMER per subcarrier approach above to find their quality measure.

3. For those RxMER stats with n^{th} percentile values not far from the mean, apply a uniform bit loading to estimate the bit loading possible on the whole channel.
4. Translate into a maximum symbol size as desired.
5. Use any of the three results as measures of downstream plant quality, recognizing that those based on RxMER stats are approximations.

The nodes with outlier 2nd percentile or 5th percentile values can be quickly identified for closer examination. The standard deviation can also provide further insight into where to focus maintenance.

9.1.2 Quality of downstream service with respect to downstream plant

This section provides only a view of current practices with a view of general service health. Much more work should be done for consideration of specific services and use cases, including impact on latency, real time use cases, and other specific use case considerations. The general guidance in this section serves as a base and starting point for more development by operators who choose to differentiate. But what is provided here is informative toward maintenance decisions, which is an ultimate goal of making service health actionable.

Certainly, timeouts and sub-optimal profile assignment are key indicators of service health problems. With the ease of access to PNM measurements like profile, RxMER per subcarrier, and FEC stats, more can be done after timeout issues are discovered.

9.1.2.1 RxMER margin and profiles for downstream health

For downstream service quality, the profile in use is important. The profile in use can be compared to the service offered, and judgement about the suitability of the profile for providing the level of service may be rendered.

1. Select the desired profile, preferably the current or most likely used profile for the given CM.
2. Poll RxMER margin at a suitable interval of time using CmDsOfdmMerMargin MIB, providing the desired threshold offset value (ThreshldOffset).
3. Use either the number of subcarriers below threshold as the measure of quality, or repeat with a more suitable profile if NumSubcarriersBelowThreshld is greater than 0.

A possible replacement for network health indication for correctable error when using OFDM channels is the number of iterations it takes to correct all of the errored bits in an LDPC codeword. However, this is not available today. Our studies comparing FEC statistics to RxMER, which is used to assign downstream profiles, reveals the complex relationship between these measurements. In Section 9.3 of this document, those test results and guidelines for the use of FEC stats for maintenance decisions are provided.

9.1.3 Consideration of FEC for Upstream health

RxMER statistics are not easily obtainable in the upstream today. But when available, they can be used in a similar way as described for the downstream. FEC statistics are not generally available in the upstream.

A possible replacement for network health indication for correctable error when using OFDM or OFDMA channels is the number of iterations it takes to correct all of the errored bits in an LDPC codeword.

Low Density Parity Coding (LDPC) is used for the inner Forward Error Correction (FEC) of both DOCSIS OFDM and OFDMA channels. LDPC differs from traditional Reed-Solomon (RS) Encoding in that it uses much larger code words spread across all of the channel.

Traditionally FEC errors, both corrected and uncorrected, have been used as an indication of network health. With the introduction of DOCSIS 3.1 OFDMA and OFDMA channels, a more aggressive (powerful) FEC scheme is used, specifically LDPC. With LDPC coding it is normal to have corrected errors on a healthy channel. It can even be argued that if there are no FEC errors that the channel is being underutilized. Seeing uncorrectable errors is still an issue, as it indicates packet errors and loss.¹⁰

¹⁰ Further CRC check at Ethernet MAC level is required to determine which packets, if any, are errored.

LDPC decoders are iterative, if a codeword is not successfully decoded on first pass, a second pass is performed. The number of decoding passes continues to increase until either the codeword is successfully decode, or the codeword is declared uncorrectable. The number of passes is implementation dependent.

Currently the only way to observe the number of iterations is in the trailer of UEPI packets sent from RPDs. These trailers include a count of the average number of iterations for each LDPC codeword size. This is defined in Table 8 of the [R-UEPI]specification. It is possible to capture UEPI packets and monitor these counts on a per channel basis. The UEPI interface is used for Remote PHY connections to RPDs.

9.1.4 About Averages

For a single CM sampled over a time period, make the samples as evenly spaced as possible over the time period, but also account for any known cyclical issues such as temperature, time of day, weather, etc. For example, for a single day quality estimate, sample at least 10 times over the day, which may be every two hours for convenience, for a total of 12 samples. For a monthly sample, collect over random parts of the day for every other day of the month.

For a group of CMs, an average of results taken over a population or sample from a population for approximately the same time interval is best. For a point in time, sample the CMs as closely as possible to that time. For an interval of time, collect samples over the time to cover it well; anything larger than a few hours will require multiple samples from each CM (see above).

Averaging can be done over time using either the statistical average (sample sum over sample size) or an exponentially weighted moving average to weigh current values higher than past values, providing a measure that might better account for random fluctuations yet be a current representation of quality.

9.1.5 Sampling Frequency for Intermittent Problems

Intermittent problems are well known to be difficult to capture. If it is assumed that intermittent issues are essentially random, or that there is no knowledge about when they generally occur, then using observations of variations through periodic sampling much catch them.

If intermittent impairments arrive randomly at E events per day with an expected duration of D seconds per event, then the proportion of the day that an intermittent event is happening is

$$P_{\text{event}} = E \text{ (events/day)} * D \text{ (seconds/ event)} / (24*60*60 \text{ (seconds/day)})$$

Likewise, to capture these events with a probability P_{capture} , then it takes n samples of the data to capture it as

$$P_{\text{capture}} = 1 - (1 - P_{\text{event}})^n$$

Solving for n, results are

$$n = \ln(1 - P_{\text{capture}}) / \ln(1 - P_{\text{event}})$$

where ln is the natural log.

To find the impairment with the given probability P_{capture} over say m days, then it is sampled at a rate of $S = n/m$ per day. Also, finding the minutes between samples is

$$I = 24*60/S.$$

The table below shows a few example results.

Table 10 - Calculate Sample Frequency

E Imp. Freq. events/day	D Imp. Dur. sec	P_{capture} Prob. Catch 0<x<1	m Scan Dur. day	S Scan Freq. sample/day	I Scan interval min./sample
300	0.010	0.50	30	665.410	2.164
100	0.010	0.50	30	1996.252	0.721
300	0.100	0.50	30	66.531	21.644

E Imp. Freq. events/day	D Imp. Dur. sec	P_{capture} Prob. Catch 0<x<1	m Scan Dur. day	S Scan Freq. sample/day	I Scan interval min./sample
300	0.010	0.90	30	2210.443	0.651
300	0.010	0.50	10	1996.229	0.721
100	0.100	0.50	30	199.615	7.214
100	0.010	0.90	30	6631.407	0.217
100	0.010	0.50	10	5988.757	0.240
100	0.100	0.90	30	663.106	2.172
100	0.100	0.90	10	1989.318	0.724

Note that some of these scenarios requires frequent sampling, which is not possible for some measurements such as RxMER per subcarrier, or full spectrum capture.

9.1.6 Stability

If these health indicators are not stable over time, then that lack of stability can also be an indication of poor network or service health or both. There are several methods for adding variability in measurements to a health indicator.

- A measurement range over time is a good indication of variability or lack of stability and can be weighted into the measurement itself.
- Using weighted averages of the measurements, so that poor values are weighted more heavily or for a longer period of time, can also make intermittent poor values weight nearly as much or as much as consistently poor values, as desired.
 - For more heavily weighted values, consider a quadratic loss function, such as the square of the difference from the value to the ideal value.
 - For weighing values over time, consider an exponentially weighted moving average or window average with a weighted function of the values in the average.

9.1.7 Impairments

Consideration of impairments and their severity is worthwhile as well. See Section 9.2 on the network maintenance decision process and the discussion of impairment severity for a severity-based approach worthy of consideration.

Note also that the paper “Profile Management Informed Proactive Network Maintenance” from SCTE Expo 2020 [PMI PNM] provides a comparison of several methods of network and service health for selecting maintenance.

Also, there are recent SCTE papers published which apply machine learning techniques to detect anomalies (impairments) in RxMER per subcarrier or spectrum data: [Zhu, Sundaresan, Rupe], [Ferreira, Harb, et al].

9.1.8 Conclusions

In the downstream, RxMER per subcarrier measurements are a great indicator of network and service health. FEC statistics can be a contributor to health indication too, particularly in the upstream; but FEC errors in the downstream indicate a health problem beyond what profiles are addressing, and RxMER statistics can indicate.

When these health measures are used for identifying proactive work, this section provides high level starting positions for that work. But there are additional features to consider as outlined in Section 9.2, “Network Maintenance Decision Process”, below.

Other measures such as T3 and T4 timeouts may be indicators as well but are not proactive in nature; T3 and T4 timeouts may indicate noticeable service degradation which should be addressed through maintenance. This section focuses on the problem of prioritizing proactive maintenance work, so focus is just on those indicators that inform that problem. FEC statistics and RxMER in combination work very well as a starting point for this. But as found in [PMI PNM], profiles can have a significant impact on the selection of proactive work, though clearly profiles are not the best indicator of network or service health when RxMER and FEC statistics are available.

9.2 Network Maintenance Decision Process

The network maintenance tools often referred to as PNM tools are useful for all types of network maintenance, proactive and otherwise. It is not helpful to differentiate between proactive and reactive maintenance, especially when making use of maintenance tools. Many of the methods outlined in this report work for all types of maintenance use cases.

There are many identifiable elements that contribute to the decision to do maintenance, be that proactive or reactive.

- Problem identification – that it exists, and location – if a customer calls in the problem, then you have an easy place to start. But that can be misleading too. If the problem is broader, the problem may not be near that customer, so additional work may be needed to isolate the problem, and certainly to identify it. If using network information to identify a problem, then you are actually on your way to identifying the issue faster than a customer call.
- Significance of impact – if a customer calls in the problem, you have a good view to at least part of the problem. Others may be impacted, and additional impact may not be felt immediately at or before the customer calls. If using network information to gauge significance, you may be able to quantify the impact in more uniform terms.
- Sufficiency of information to address the problem – when a customer calls, you know you must address the service issue. But if using network data to determine a network issue that may not yet be a service issue, then an important part of the decision process is whether or not enough information is available to make the problem addressable.
- Cost effectiveness to address the problem – sometimes opportunities, like a mandatory reactive repair, will lower the cost of a proactive repair, making it an imperative. When considering the degradation and further damage that is possible, some proactive tools can allow operators to get ahead of costly repairs.
- Resources available to address the problem – when ahead of the problem, operators can plan and schedule maintenance when resources are available and less expensive.
- Success probable – in learning how to associate network data with service and network conditions, data can be collected and tracked, increasing the likelihood of successful resolution. This will inform planning and decision making for maintenance.
- Priority for operations – With all the information outlined above, operators are equipped to prioritize work and optimize operations spend.

The distinction between proactive and reactive maintenance is subtle, and not so cleanly separate.

Operators who are fully occupied with reactive maintenance can still get ahead of the work to be proactive defendants of service. By embracing some of the newer network operations tools which are based on network data that are available through network elements already, they can troubleshoot more efficiently and effectively. Once operations becomes more efficient, and resources are freed up, then those resources can be more focused on the network data that foretells of network problems. Eventually, operations falls down the slippery slope of proactive network maintenance simply by being efficient.

By using the proactive tools, sometimes in reactive ways, operators will slip into doing more proactive work simply because it is more efficient and better for service to do so. There really is no hurdle to get over; simply mix in proactive when and where it makes sense to do, and you will be over any perceived hurdle. By using PNM tools effectively, for reactive problems too, operators find that reactive repairs can be streamlined, freeing resources for proactive work. Mixing proactive with reactive increases the opportunity density, reducing wasted truck rolling time and increasing overall operations efficiency.

Keep in mind that there are network problems that DOCSIS will work around, and thus hide if you aren't using proactive tools to look. These don't necessarily impact service but can contribute to smaller issues causing service impact that otherwise wouldn't and may indicate network damage that must be addressed. Some network damage puts service at risk, and some will get worse and more expensive if left alone. Changes in performance measurements can indicate which problems are new versus which are getting worse versus which are stable.

Note also that the paper [PMI PNM] from SCTE Expo 2020 provides a comparison of several methods of network and service health for selecting maintenance. From that paper, the following conclusions can be used:

“Profile management has a clear influence on the impact to service that results from PNM. Therefore, some consideration of profiles is important when selecting proactive work to do. It is advisable that operators at least look at the impact on profiles when determining whether to conduct PNM, and it is likely a factor in prioritizing proactive work.

Anomaly-driven PNM, finding and eliminating anomalies, appears to be the most popular approach that operators take for their PNM; but modem-driven PNM is an option to consider for some, as the data that most informs PNM comes from the CMs. Likewise, profile-driven PNM is an option to consider as well. But ultimately, each of these factors needs to be considered in some way in deciding which PNM work to take on.

Though this paper does not compare all options, an examination of the options presented here does suggest that impairment-driven identification of maintenance opportunities is a very good approach, that problem severity can be measured by impact on CMs, and that it is necessary and complimentary to PNM decision making to consider the impact on profiles when determining the capacity impact on the network for selecting maintenance. Said differently—find anomalies, investigate their impact on customers’ service, estimate the impact due to profiles and on network impact expected from the work, and use these pieces of information to prioritize and select PNM opportunities. It may at times be best to let profiles drive PNM maintenance decisions, in fact.

Seeing that some network repairs can have negative impact on profiles, it suggests more work on profiles is warranted too. It may be simple enough to set a rule that says, if all CMs improve, then the recalculation of profiles can only result in a net benefit, or a benefit to each CM, or another appropriate rule, before consideration of replacing profiles. In other words, do not harm first. Alternately, it may be useful to develop algorithms that set profiles to be robust against these changes. Likewise, it may be worth investigating different objectives when setting profiles altogether.”

A methodology and framework for building and testing various maintenance decision processes is available from CableLabs and is explained in [Rupe, Zhu]

9.3 FEC, MER and Decoding in the Downstream

Cable operators have for many years tracked pre-FEC and post-FEC bit errors in legacy SC-QAM based digital video and DOCSIS networks. Those legacy SC-QAM signals, based on ITU-T J.83, use concatenated coding with outer-code Reed-Solomon forward error correction. The onset of pre-FEC bit errors for the Reed-Solomon decoder has long been used by many operators to determine that performance degradation is starting to occur, and that the cause of the pre-FEC errors should be identified and fixed before post-FEC bit errors occur. Post-FEC bit errors result in packet loss.

DOCSIS 3.1 introduced a more powerful FEC to the cable industry, known as low density parity check (LDPC). Because of the way LDPC works, it is not possible to rely upon the presence of correctable codeword errors to indicate the onset of potential problems. Indeed, one can have 100% correctable codeword errors and expect generally reliable modem operation, with no real indication of how close to the digital cliff one is. What to do?

One proposed possible solution is to track RxMER per subcarrier (relative to the RxMER threshold for the modulation order in use) along with correctable and uncorrectable codeword errors. The InGeNeOS Working Group commissioned CableLabs to do this testing, and the test results are in Appendix I.

Some plots in the test results of that appendix show FEC decoder input, which is the sum of the correctable codewords and the uncorrectable codewords as a percentage of total codewords. These are the codewords that failed pre-decoder syndrome check. So, one minus the FEC decoder input is the percent of error free codewords after slicing and prior to decoding. It is easiest to tell which data in the figures are from FEC decoder input as they stay flat (and nearly 100%) at below-threshold MER values while the underlying uncorrectable codeword percent increases as MER decreases from above-threshold to below-threshold.

Total codewords is the sum of the number of error-free, correctable, and uncorrectable codewords. FEC decoder input is the sum of the correctable codewords and the uncorrectable codewords. All are reported as a percentage of total codewords.

Note on codeword error terminology: The DOCSIS 3.1 downstream FEC includes LDPC and BCH, while the upstream uses only LDPC. In the downstream, the BCH can be thought of as being analogous to a broom and dustpan, sweeping up residual errors, and providing the ability to get correctable and uncorrectable codeword error

counts. Among the available downstream FEC statistics are: uncorrectable codewords (the number of codewords that failed BCH decoding); correctable codewords (the number of codewords that failed pre-decoding LDPC syndrome check and passed BCH decoding); and the total number of FEC codewords.

In the upstream, because there is only LDPC, one cannot get the same FEC statistics that are available in the downstream. Among the available upstream FEC statistics are: pre-FEC error-free codewords (the number of codewords that passed pre-decoding syndrome check); unreliable codewords (the number of codewords that failed post-decoding syndrome check); corrected codewords (the number of codewords that failed pre-decoding syndrome check, but passed post-decoding syndrome check); and the total number of FEC codewords.

Replicated with permission, from R. Prodan, L. Montreuil, A Kliger, B. Shen, “Demystifying the DOCSIS 3.1 PHY,” SCTE Cable-Tec EXPO, September 22-25, 2014, [Prodan, et al], see the following two unlabeled figures and explanatory text which supports these results.

“DOCSIS 3.1 employs several forward error correction codes for various purposes. This section will discuss both the downstream and upstream forward error correction codes for protecting the payload user data. Several relevant characteristics for these FEC codes are as follows:

Downstream Code:

- DVB-C2¹¹ (16200,14400) inner LDPC + (14400,14232) outer BCH¹²
- Overall Rate 0.8785

Upstream Codes:



- (16200,14400) Quasi-Cyclic LDPC Rate 0.8889
- (5940,5040) Quasi-Cyclic LDPC Rate 0.8485
- (1120,840) Quasi-Cyclic LDPC Rate 0.75

The bit error ratio performance as a function of SNR for the downstream FEC code for all square, non-square, and mixed modulation modes from 64-QAM to 4096-QAM is shown in the figure below. The 10^{-8} bit error ratio SNR threshold for each modulation order is tabulated from this graph. The spectral efficiency is given by the code rate times the number of bits per QAM subcarrier for a given modulation order. The resultant spectral efficiency (bits/s/Hz) vs. SNR is shown in figure 16.¹³ The Shannon capacity limit is also shown in this figure to evaluate the relative proximity of the performance of the downstream FEC code. The figure shows this FEC code to be about 3 dB from the Shannon limit or about one bit/s/Hz below the Shannon limit.”

¹¹ Digital Video Broadcasting – Cable version 2, not spelled out in the original document.

¹² Bose-Chaudhuri-Hocquenghem, not spelled out in the original document.

¹³ The figures that follow were collectively labeled as “figure 16” in the original document, but are unnumbered in this document as part of the quoted material.



A recommended practice is to use the `CmDsOfdmRequiredQamMer` object to determine the MER target values required in each subcarrier, then determine the appropriate bit loading by comparing the `RxMER` values to the target values, per subcarrier. See the following table.

Table 11 - `CmDsOfdmRequiredQamMer` Object from [CM-OSSiv3.1]

Attribute Name	Type	Access	Type Constraints	Units	Default Value
ReqMerQam16	UnsignedInt	R/W		quarterDb	60
ReqMerQam64	UnsignedInt	R/W		quarterDb	84
ReqMerQam128	UnsignedInt	R/W		quarterDb	96
ReqMerQam256	UnsignedInt	R/W		quarterDb	108
ReqMerQam512	UnsignedInt	R/W		quarterDb	122
ReqMerQam1024	UnsignedInt	R/W		quarterDb	136
ReqMerQam2048	UnsignedInt	R/W		quarterDb	148
ReqMerQam4096	UnsignedInt	R/W		quarterDb	164
ReqMerQam8192	UnsignedInt	R/W		quarterDb	184
ReqMerQam16384	UnsignedInt	R/W		quarterDb	208

From a carrier to noise ratio (CNR) performance point of view, the equivalent numbers are stated in the PHY specification as in the following table.

Table 12 - CM Minimum CNR Performance in AWGN Channel from [PHYv3.1]

Constellation	CNR ^{1,2} (dB) Up to 1 GHz	CNR ^{1,2} (dB) 1 GHz to 1.2 GHz	Min P _{6AVG} dBmV
4096	41.0	41.5	-6
2048	37.0	37.5	-9
1024	34.0	34.0	-12
512	30.5	30.5	-12
256	27.0	27.0	-15
128	24.0	24.0	-15
64	21.0	21.0	-15
16	15.0	15.0	-15

Table Notes:
 Note 1 CNR is defined here as total signal power in occupied bandwidth divided by total noise in occupied bandwidth.
 Note 2 Channel CNR is adjusted to the required level by measuring the source inband noise including phase noise component along with transmitter noise and distortion and adding the required delta noise from an external AWGN generator to achieve the desired CNR at the CM F-connector.
 Note 3 Applicable to an OFDM channel with 192 MHz of occupied bandwidth.

From [Sundaresen, Zhu, et al], Section 3, see results that align with these test results where there is overlap, and the work in our new paper confirms and extends the work in this paper. While our work tested a wider range of conditions and CMs, this referenced paper covered the case of profile switchover. If profiles are allowed to switch in a given implementation, the adjustments may continue to keep service stable (avoid uncorrectable codeword errors), though at a cost of data rate. But if profiles are not allowed to be switched, then the uncorrectable codeword errors may impact service.

Outer BCH coding leads to definitive determination of errored codewords, unlike LDPC which only indicates error free assurance. LDPC alone leaves open whether or not codewords not absolutely error-free are errored.

LDPC is used for the inner FEC of both DOCSIS OFDM and OFDMA channels. LDPC differs from traditional Reed-Solomon (RS) Encoding in that it uses much larger code words spread across all of the channel.

Traditionally FEC errors, both corrected and uncorrected, have been used as an indication of network health. With the introduction of DOCSIS 3.1 OFDMA and OFDMA channels, a more aggressive FEC scheme is used, specifically LDPC. With LDPC coding it is normal to have corrected errors on a healthy channel. It can even be argued that if there are no FEC errors that the channel is being underutilized. Seeing uncorrectable errors is still an issue, as it indicates packet loss.

LDPC decoders are iterative: if a codeword is not successfully decoded on first pass, a second pass is performed. The number of decoding passes continues to increase until either the codeword is successfully decode, or the codeword is declared uncorrectable. The number of passes is implementation dependent.

Appendix I Error Correction Margin Solution Analysis

I.1 Introduction

Cable operators have for many years tracked pre-FEC and post-FEC bit errors in legacy SC-QAM based digital video and DOCSIS networks. Those legacy SC-QAM signals, based on ITU-T J.83, use concatenated coding with outer-code Reed-Solomon forward error correction. The onset of pre-FEC bit errors for the Reed-Solomon decoder has long been used by many operators to determine that performance degradation is starting to occur, and that the cause of the pre-FEC errors should be identified and fixed before post-FEC bit errors occur. Post-FEC bit errors result in packet loss.

DOCSIS 3.1 introduced a more powerful FEC to the cable industry, known as low density parity check (LDPC). Because of the way LDPC works, it is not possible to rely upon the presence of correctable codeword errors to indicate the onset of potential problems. Indeed, one can have 100% correctable codeword errors and expect generally reliable modem operation, with no real indication of how close to the digital cliff one is. What to do?

One proposed possible solution is to track RxMER per subcarrier (relative to the RxMER threshold for the modulation order in use) along with correctable and uncorrectable codeword errors.

There is interest in DOCSIS 3.1 OFDM RxMER per subcarrier values at which correctable and uncorrectable codeword errors begin to appear. This is with and without profile changes. There is also interest in the triggering of a profile change as an indicator. Finally, operators want to confirm the effectiveness of a nominal 6 dB of operating headroom above the RxMER threshold for the modulation order in use, and identify risks with operation at lower headroom values.

Several controlled tests were conducted at CableLabs, with the test procedures and results summarized in this document.

I.2 Error Correction Margin Test Plan

I.2.1 Test Set A

Test set up 1 is similar to section 3.3 of [Sundaresen, Zhu, et al], an SCTE PMA paper. An Agilent/Keysight N5182A MXG Vector Signal Generator (VSG) with AWGN option as an AWGN source was used.

A test plant was set up with at least two manufacturers of each chip set (total four CM manufacturers) and at least two CMs of each copy for a total of at least eight CMs, connected to a splitter setup. That splitter was fed by a nearby OFDM source, to avoid the need to use an external amplifier.

First, the RxMER responses for these CMs and test set up to get consistent levels (flat input) were calibrated.

Then configured test conditions were configured so the CMs get equal received power levels. There were four sets of tests: one set each at +10 dBmV, 0 dBmV, -6 dBmV, and -10 dBmV RF input to the modems.

For this first set of tests, all (static) profiles were fixed to 4096-QAM.¹⁴ In most cases, where not indicated, a 192 MHz wide channel was used.

The remaining procedures were as follows. Obtain a baseline reported RxMER per subcarrier from the DOCSIS 3.1 modems (data collected via TFTP from the CmDsOfdmRxMer object) when the CNR is in the high 40s (ideally there should be no correctable codeword errors reported) at the desired RF input level(s). Using a Keysight UXA, measure OFDM RxMER per subcarrier too.¹⁵

¹⁴ It was desired to put together a test set-up that has high enough CNR to be above the point where correctable codeword errors show up. For instance, if testing 4096-QAM, that would require a CNR in the high 40s.

Keep in mind that RxMER per subcarrier values much above the mid 40s (as reported by a cable modem) may be suspect. For such high RxMER values, some significant portion of the measured noise may be contributed by the CM or UXA, and as such may not be a precise indication of the input to the measuring device.

¹⁵ In figures and text showing results, sometimes the Keysight UXA measurements are described more generally as VSA for vector signal analyzer.

Introduce noise level uniformly on all frequencies first. Start injecting wideband AWGN (enough to cover the entire OFDM signal spectrum), increasing its amplitude such that the carrier-to-noise ratio (CNR) at the modems is able to be reduced in 1 dB steps during the tests (smaller increments as the limits were reached). At each step record RxMER per subcarrier, and FEC stats (uncorrected, correctable, and uncorrectable codewords). Uncorrected codewords are codewords that were determined to contain no errors even prior to decoding. Look for the onset of correctable codeword errors. Increase the noise level (decrease the CNR) until uncorrectable codeword errors reach 50%, or if not possible then at least until correctable codeword errors decline for at least three increases are seen¹⁶.

As CNR is reduced in 1 dB steps, correctable codeword errors will eventually show up. Note the value at each 1 dB drop in CNR. Repeat until the onset of uncorrectable codeword errors. Keep dropping CNR until the OFDM signal no longer passes any data at all.

Test set up 2: repeat with an attenuator value to achieve 0 dBmV RF input.

Test set up 3: repeat with an attenuator value to achieve -6 dBmV RF input.

Test set up 4: repeat with an attenuator value to achieve -10 dBmV RF input.

I.2.2 Test Set B

Repeat test set A but with controlled (simulated) LTE ingress instead of AWGN.

I.2.3 Test Set C

Repeat test with AWGN configured to produce a CNR near (about 3 dB above) the RxMER (SNR) threshold (as reported in specifications and repeated later in this section). Then add LTE ingress, and increase the strength of the LTE signal until failure.

I.2.4 Figures for Test Configurations



Figure 135 - Test Setup for Test Sets A and B




Figure 136 - Test Setup for Test Set C

Some plots in the test results show FEC decoder input, which is the sum of the correctable codewords and the uncorrectable codewords as a percentage of total codewords. These are the codewords that failed pre-decoder syndrome check. So, one minus the FEC decoder input is the percent of error free codewords after slicing and prior to decoding. It is easiest to tell which data in the figures are from FEC decoder input as they stay flat (and nearly

¹⁶ We might find that we reach 100% correctable codeword errors fairly quickly, where they remain until uncorrectable codeword errors start to show up. When within about 6 dB of the RxMER threshold for a given modulation order, without ingress or other severe impairment, the correctable codeword errors are near or at 100% all the way down to the RxMER threshold.

100%) at below-threshold MER values while the underlying uncorrectable codeword percent increases as MER decreases from above-threshold to below-threshold.

Total codewords is the sum of the number of error-free, correctable, and uncorrectable codewords. FEC decoder input is the sum of the correctable codewords and the uncorrectable codewords. All are reported as a percentage of total codewords.

I.3 Relating FEC to RxMER in OFDM for avoiding uncorrectable codeword errors

I.3.1 AWGN flat noise only

In the following plots, the MER(dB) values are derived from ratios of signal power and error power at the OFDM demodulator’s slicer for each subcarrier, for both the CM and UXA measurements. The MER for these curves correspond to the average RxMER (averaging the “dB” values of each subcarrier RxMER value in “dB”); because the test condition is intended as a flat channel with flat noise (AWGN) there is some, but minimal, variation in the RxMER per subcarriers. Further, some plots show FEC decoder input, which is the sum of the correctable codewords and the uncorrectable codewords as a percentage of all codewords.

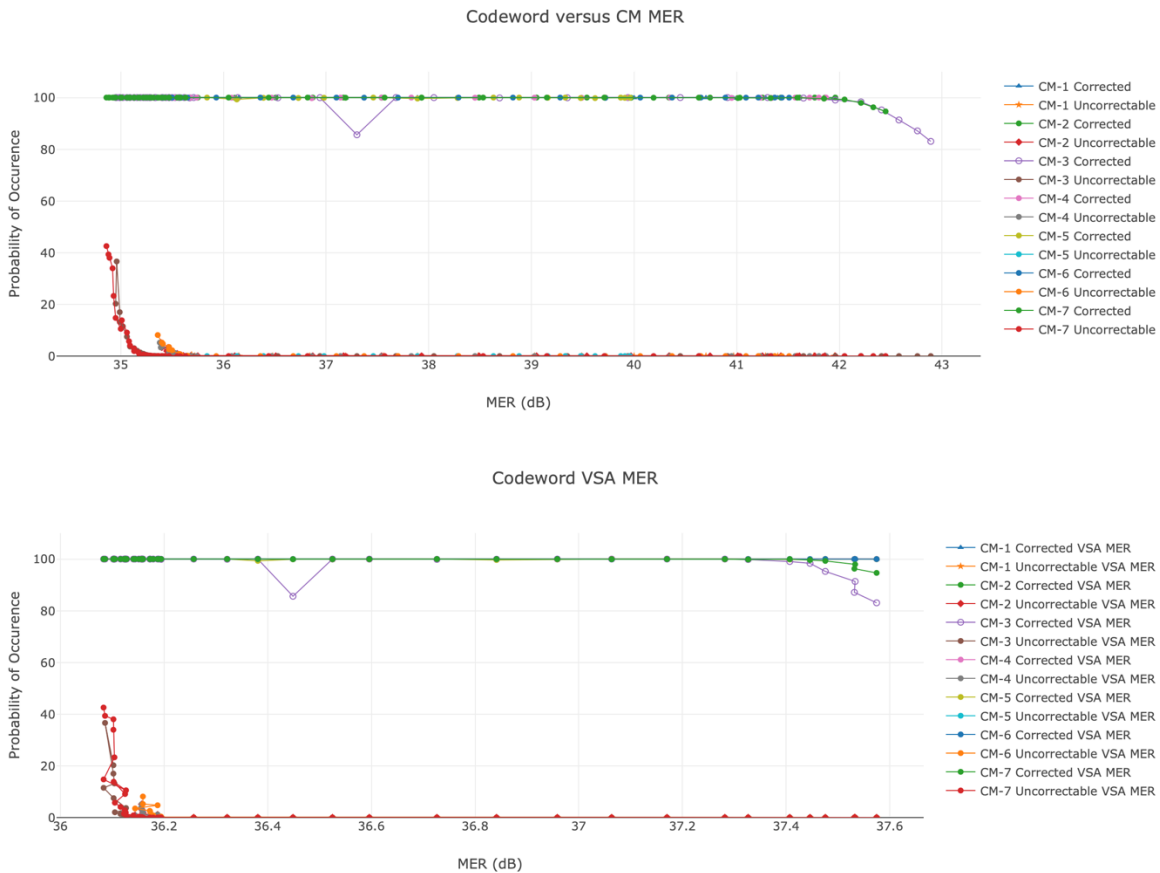


Figure 137 - CM and VSA Input Signal Power -10 dBmV

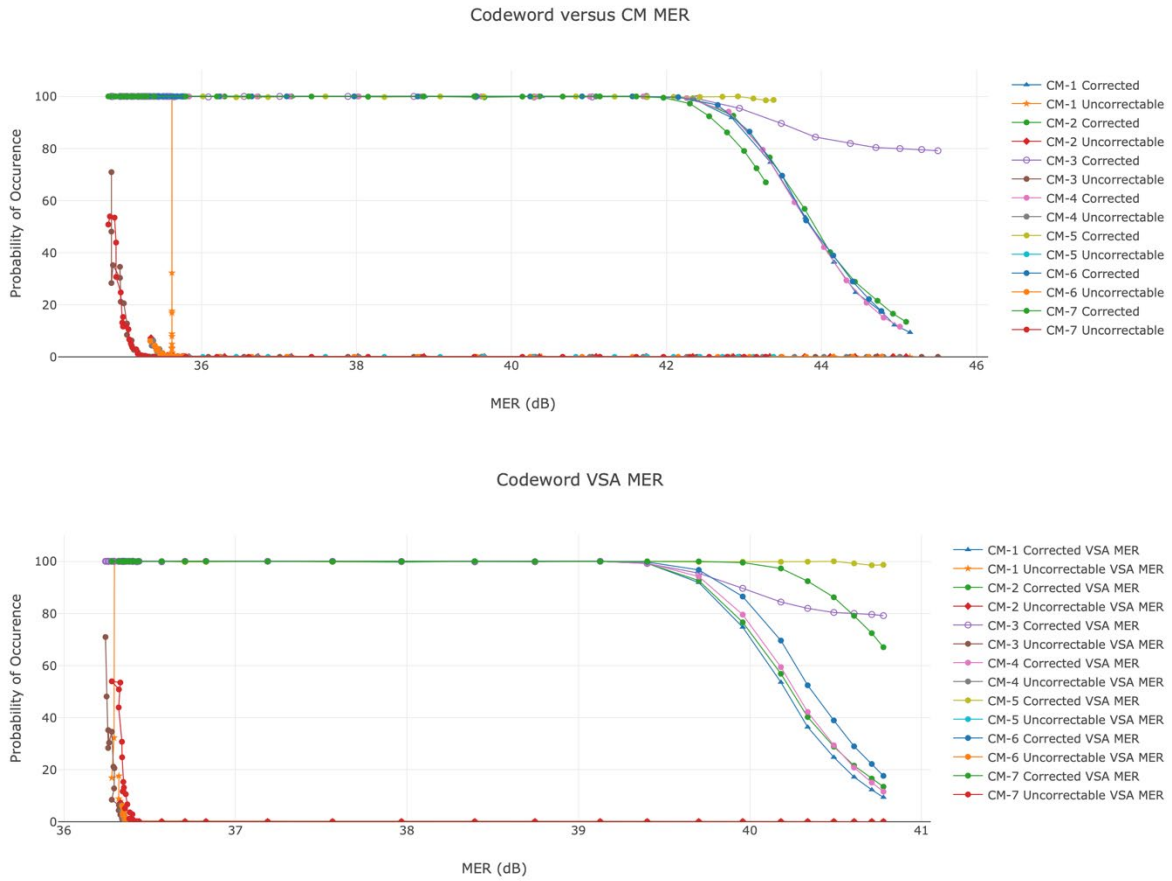


Figure 138 - CM and VSA Input Signal Power -6 dBmV

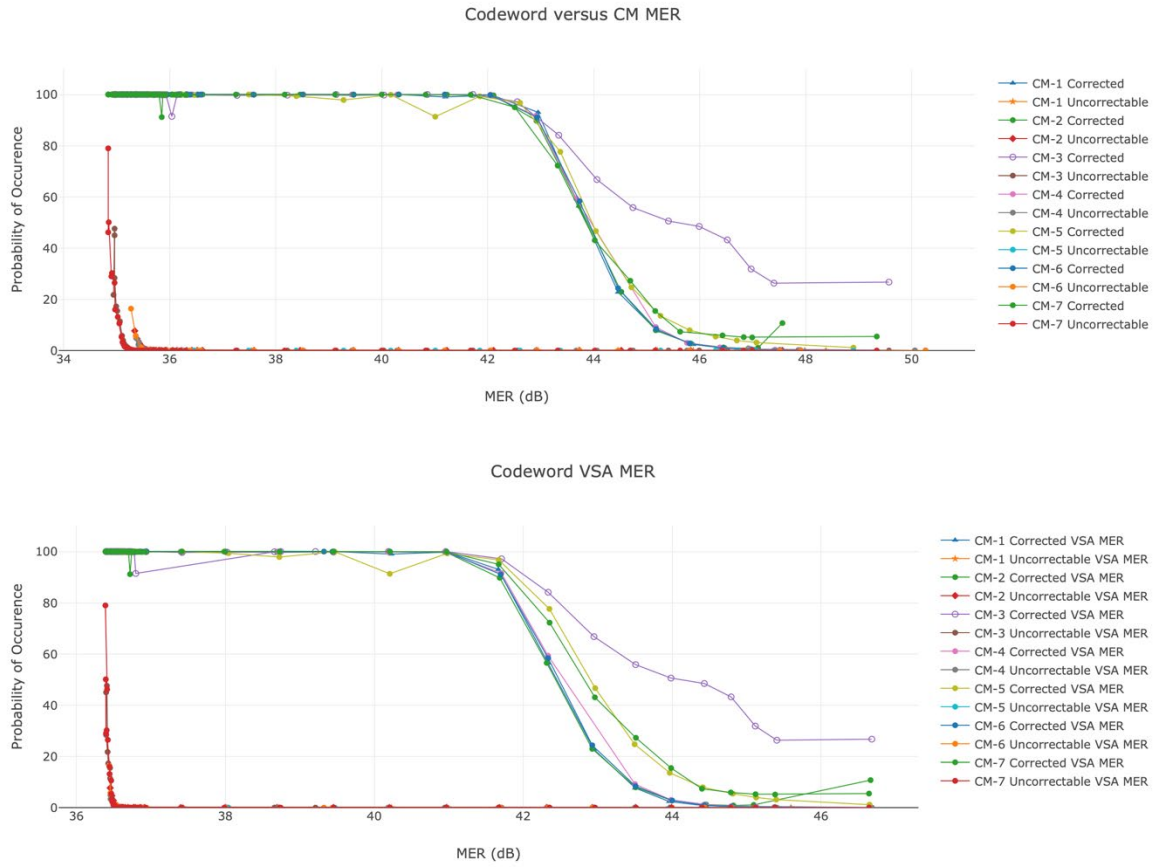


Figure 139 - CM and VSA Input Signal Power 0 dBmV

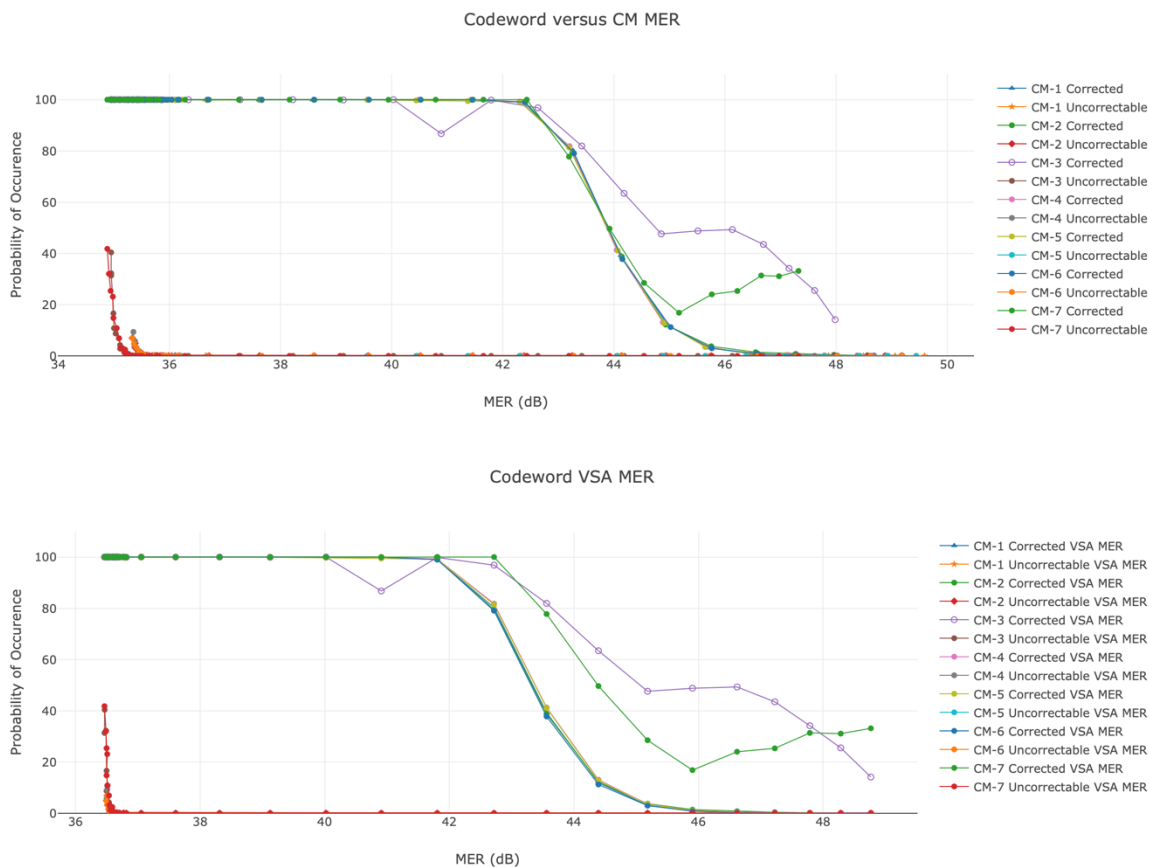


Figure 140 - CM and VSA Input Signal Power +10 dBmV

Refer to Figure 137 through Figure 140. Note that the figures show FEC decoder input, which is the sum total of all correctable and uncorrectable codewords. Also recall the test configuration that the UXA (which is our vector signal analyzer (VSA)) is connected at the same time in the same setup with the CMs, so both measurements are taken at the same time, but resulting in different dB values.

Also note that the UXA measures RxMER using data subcarriers instead of pilots as the CM uses.

- CM and UXA measurements do not agree on the average RxMER values. CMs seem to perform with no uncorrectable codeword errors (100% correct codewords) down to a level just below 36 dB RxMER on their own measure, whereas the UXA reports a higher RxMER value corresponding to the FEC threshold.
- CMs differ somewhat in the RxMER level at which they take on errors (correctable and uncorrectable), though in this set of tested CMs the difference is not very large.
- There is at least a 6 dB range between where a) the CM sends nearly 100% of the data through the decoders and b) the CM has no uncorrectable codeword errors after decoding, resulting in nearly 100% correctable codeword errors.
- As uncorrectable codeword errors appear, increasing the AWGN slightly to reduce the CNR by less than 0.5 dB, the CM loses connection (sync) with the CMTS.
- These results suggest that,
 - As long as the (average) RxMER is above 36 dB (for 4096-QAM) the CM will perform fine under flat AWGN.

- If a given plant has low variability in its noise level, operators could expect service to remain reliable much closer to the 36 dB limit, so could set an operating headroom threshold at a lower RxMER value than if a plant has high variability with noise. In other words, while more reliable operation is expected when RxMER is 6 dB or more above the threshold, one could operate with less headroom if the plant is stable and free of ingress.
- Issues to be studied include
 - What if noise is focused over a subset of frequencies (such as with ingress).
 - Are there differences in these results for different modulation orders, or OFDM channel sizes?
 - Does LTE have a different impact than random noise over the same frequencies.

To address these issues to be studied, bin noise (such as narrow band ingress) to the test will be introduced.

As power level seems not to matter with these tests, only testing at 0 dBmV signal power level in subsequent tests will be done.

1.3.1.1 Results with Correctable and Uncorrectable codeword errors (whale plots)

For convenience and preference, the data from the 0 dBmV tests above are reproduced in Figure 141 with correctable codeword data points (instead of FEC decoder input) and uncorrectable codeword data points, forming what is called “whale plots.” Most of the rest of our results will be showing whale plots, but either can be created from the same data results.

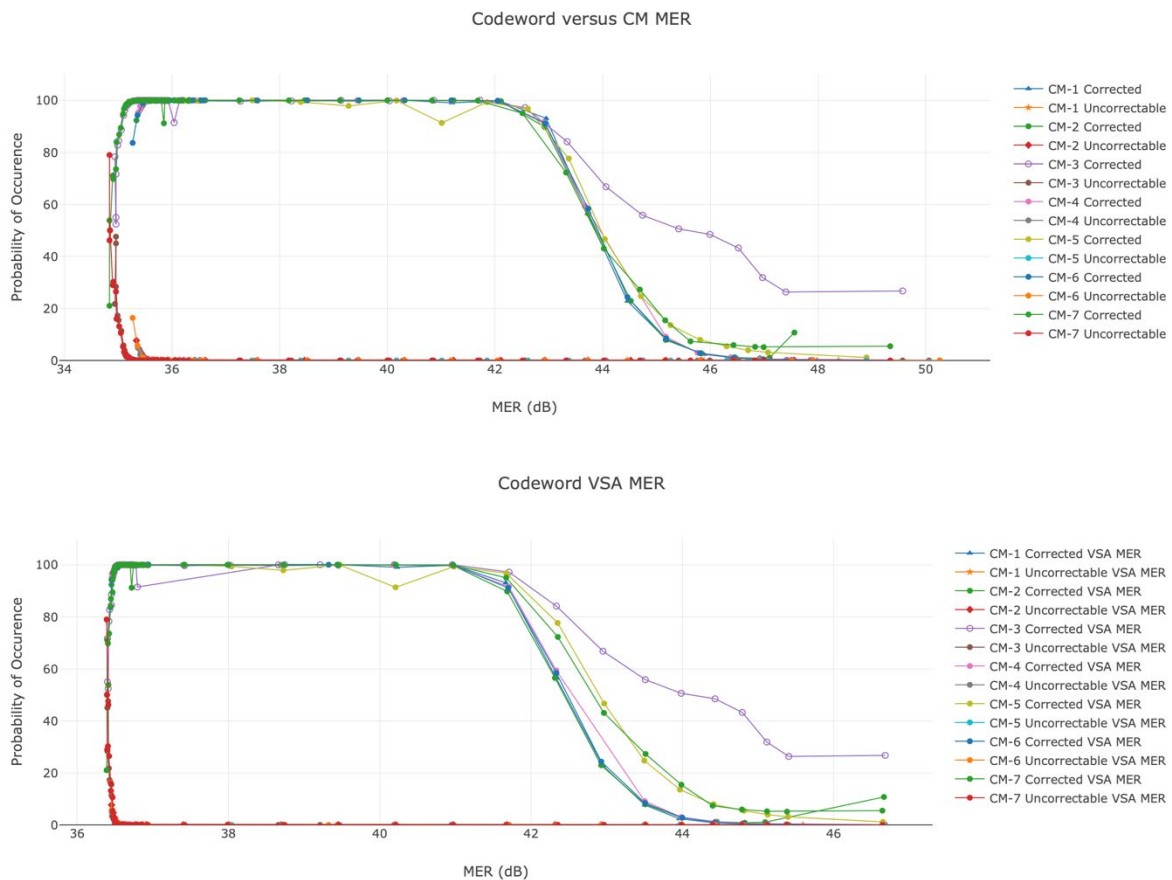


Figure 141 - Whale plots (Correctable and Uncorrectable Codeword Errors) of 0 dBmV Power Level, 4k-QAM, with Increasing Noise to Reduce Average RxMER

I.3.2 LTE Ingress Results with Correctable and Uncorrectable Codeword Errors (Whale Plots)

This subsection reports the results from our testing with simulated LTE ingress, or notch noise. First, testing with an 8 MHz-wide noise inserted into a 96 MHz-wide channel, with 94 MHz modulated spectrum (only band edge excluding subcarriers) demonstrates that CMs have trouble staying synchronized and connected to the CMTS.

First run:

Channel width: 96 MHz

Modulated spectrum: 94 MHz

Profile modulation: 4096-QAM

PLC frequency: 333 MHz

LTE frequency: 375 MHz

LTE width: 8 MHz

Second run:

Channel width: 192 MHz

Modulated spectrum: 190 MHz

Profile modulation: 4096-QAM, and 256-QAM

PLC frequency: 333 MHz

LTE frequency: 375 MHz

LTE width: 8 MHz

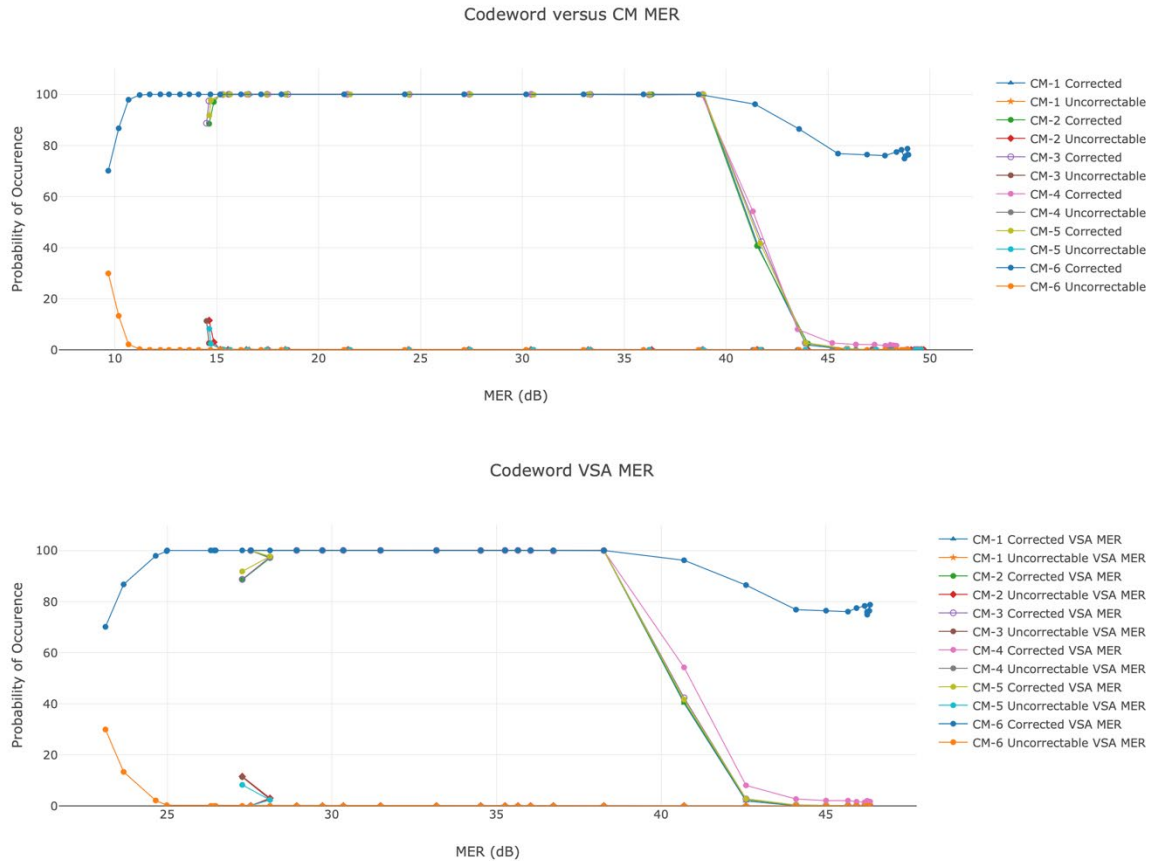


Figure 142 - Whale Plots of 0 dBmV Power Level, 4k-QAM, 96 MHz-wide Channel, 8 MHz-wide LTE Ingress, with Increasing LTE level to Reduce the RxMER in the LTE Band. MER (dB) in these Plots is a Notch MER (dB) Value

The MER (x-axis) values in Figure 142 through Figure 144 are the average RxMER in the spectrum where the LTE signal is present. Figure 142 shows the 96 MHz-wide channel is overtaken by the 8 MHz-wide LTE ingress as the LTE signal power is increased, which decreases the MER value, and results in CMs losing sync at the lower MER values for the notch. Therefore, to avoid losing data from the CMs losing sync, the following tests are done with a wider channel to allow the CMs to remain connected.

Note also that there is a 13 dB difference between the notch MER for approximate threshold in all cases in comparing the CM MER and VSA MER curves. In all cases the approximate threshold MER for the notch is about 13 dB (plus or minus a dB) higher for the “VSA MER” curve set.

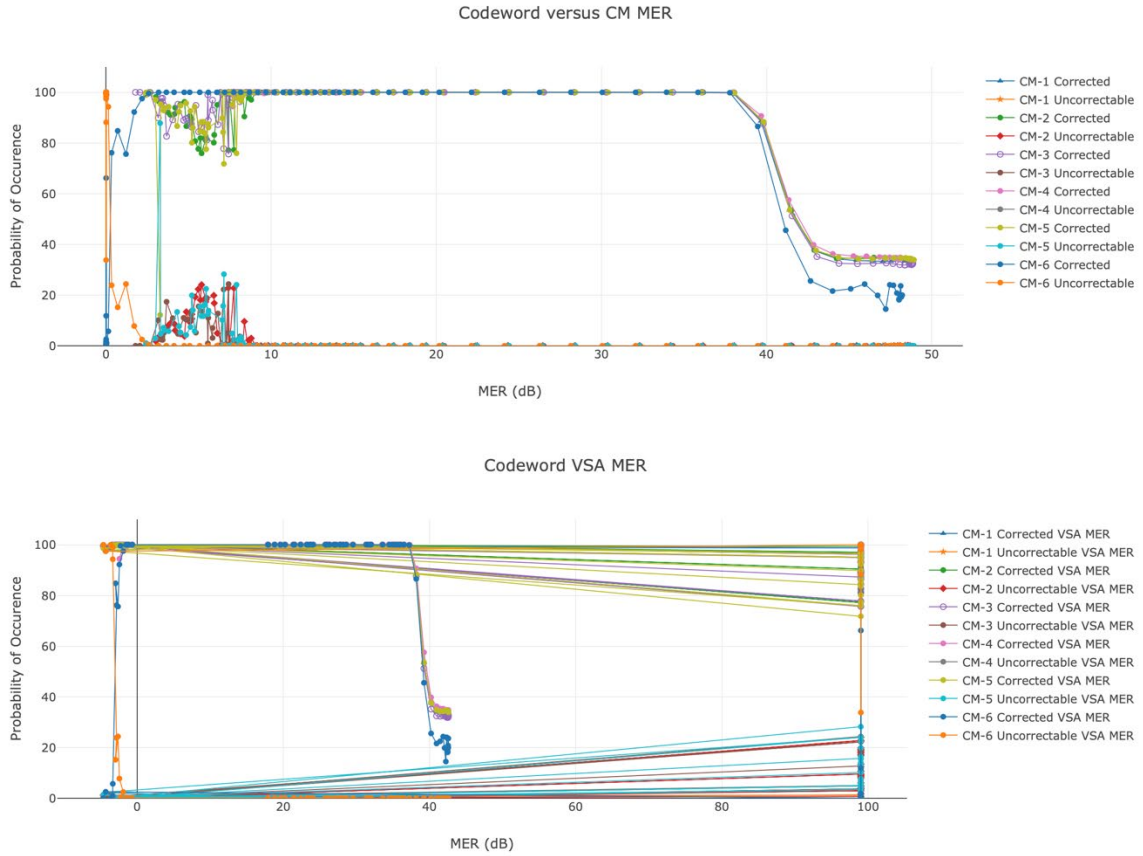


Figure 143 - Whale Plots of 0 dBmV Power Level, 4k-QAM, 192 MHz-wide Channel, 8 MHz-wide LTE Ingress, with Increasing LTE level to Reduce the Average RxMER (x-axis) in the LTE Band, Reporting RxMER over the LTE Band

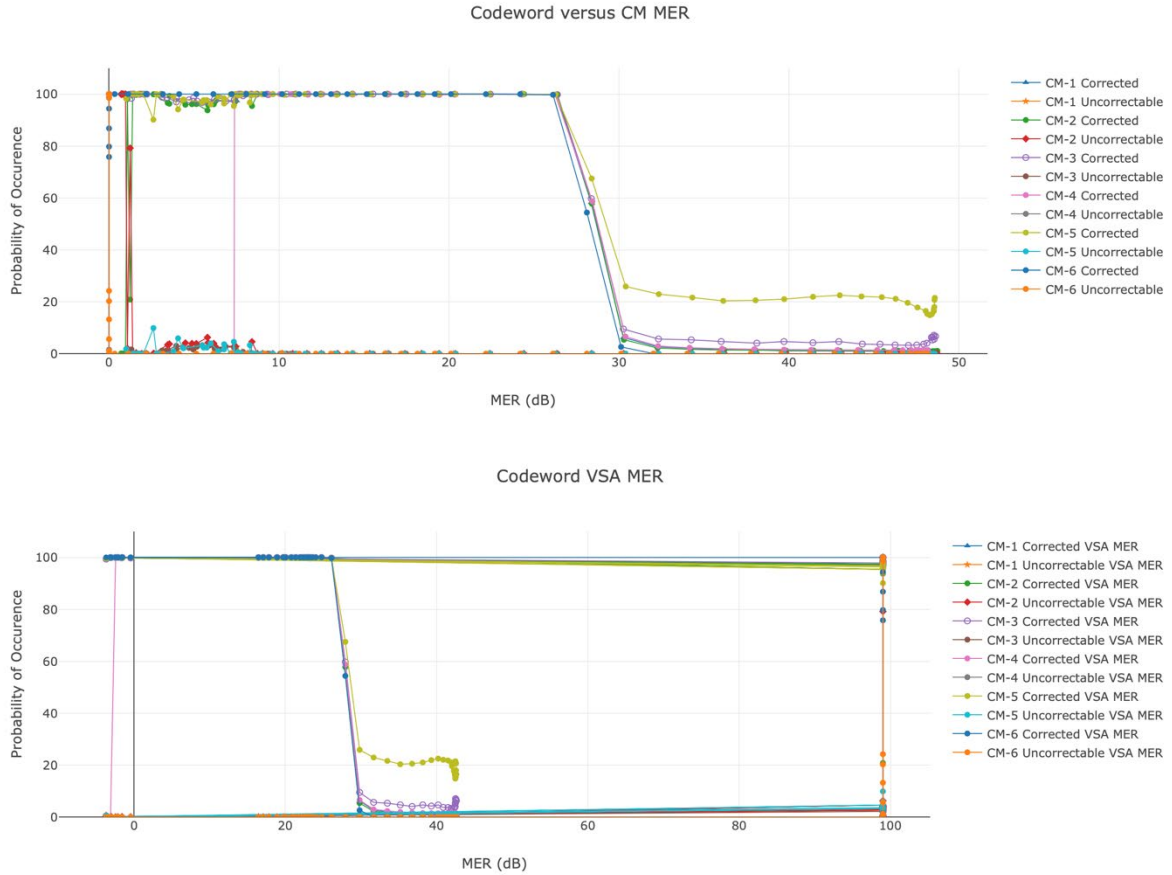


Figure 144 - Whale Plots of 0 dBmV Power Level, 256-QAM, 192 MHz-wide Channel, 8 MHz-wide LTE Ingress, with Increasing LTE Level to Reduce the RxMER in the LTE Band, Reporting Average RxMER over the LTE Band on the x-axis

Figure 142 through Figure 144 indicate the average MER value over the simulated LTE ingress channel, but Figure 143 and Figure 144 have the same 8 MHz-wide notch noise as in Figure 142, but on a 192 MHz-wide channel to allow resiliency mechanisms to keep the CMs connected and working. All of these three figures report the average MER (dB) over the LTE channel. As expected, this wide of an LTE ingress does not disrupt the CM entirely until uncorrectable codewords begin to show, then there is no room for variability in the MER values.

To keep this new test comparable to the previous, subsequent plots of the same data were provided but for the average RxMER value over the full 192 MHz channel instead of the RxMER value in the simulated noise band, to better allow comparisons.




Figure 145 - Whale Plots of 0 dBmV Power Level, 4k-QAM, 192 MHz-wide Channel, 8 MHz-wide LTE Ingress, with Increasing LTE Level to Reduce the RxMER in the LTE Band, this Time Reporting Average RxMER over the Whole Channel




Figure 146 - Whale Plots of 0 dBmV Power Level, 256-QAM, 192 MHz-wide Channel, 8 MHz-wide LTE Ingress, with Increasing LTE Level to Reduce the RxMER in the LTE Band, this Time Reporting Average RxMER over the Whole Channel

If profiles are fixed, new or intermittent ingress or spurious impairments will cause problems sooner than if profiles can adjust to compensate for the presence of impairments, or profiles are set to be robust against known (intermittent) impairments.

Seeing a drop in a profile can indicate that a problem has occurred, and there may be a reason to investigate.

But lacking a drop in the profile, and looking at average RxMER, there is only about 2 dB of room between taking on correctable codeword errors and having a significant number or even full uncorrectable codeword errors (making the CM lose connection and the data not showing in the plot). This observation suggests that any time a CM is experiencing correctable codeword errors it may be vulnerable to LTE or other ingress or channel noise. If an operator is not able to manage a tight plant, it may be important to manage profiles conservatively, or keep a large margin in RxMER.

I.3.3 Testing with raised noise floor, and simulated LTE ingress together

Channel width: 96 MHz

Modulated spectrum: 94 MHz

Profile modulation: 4096-QAM

PLC frequency: 333 MHz

LTE frequency: 375 MHz

LTE width: 5 MHz

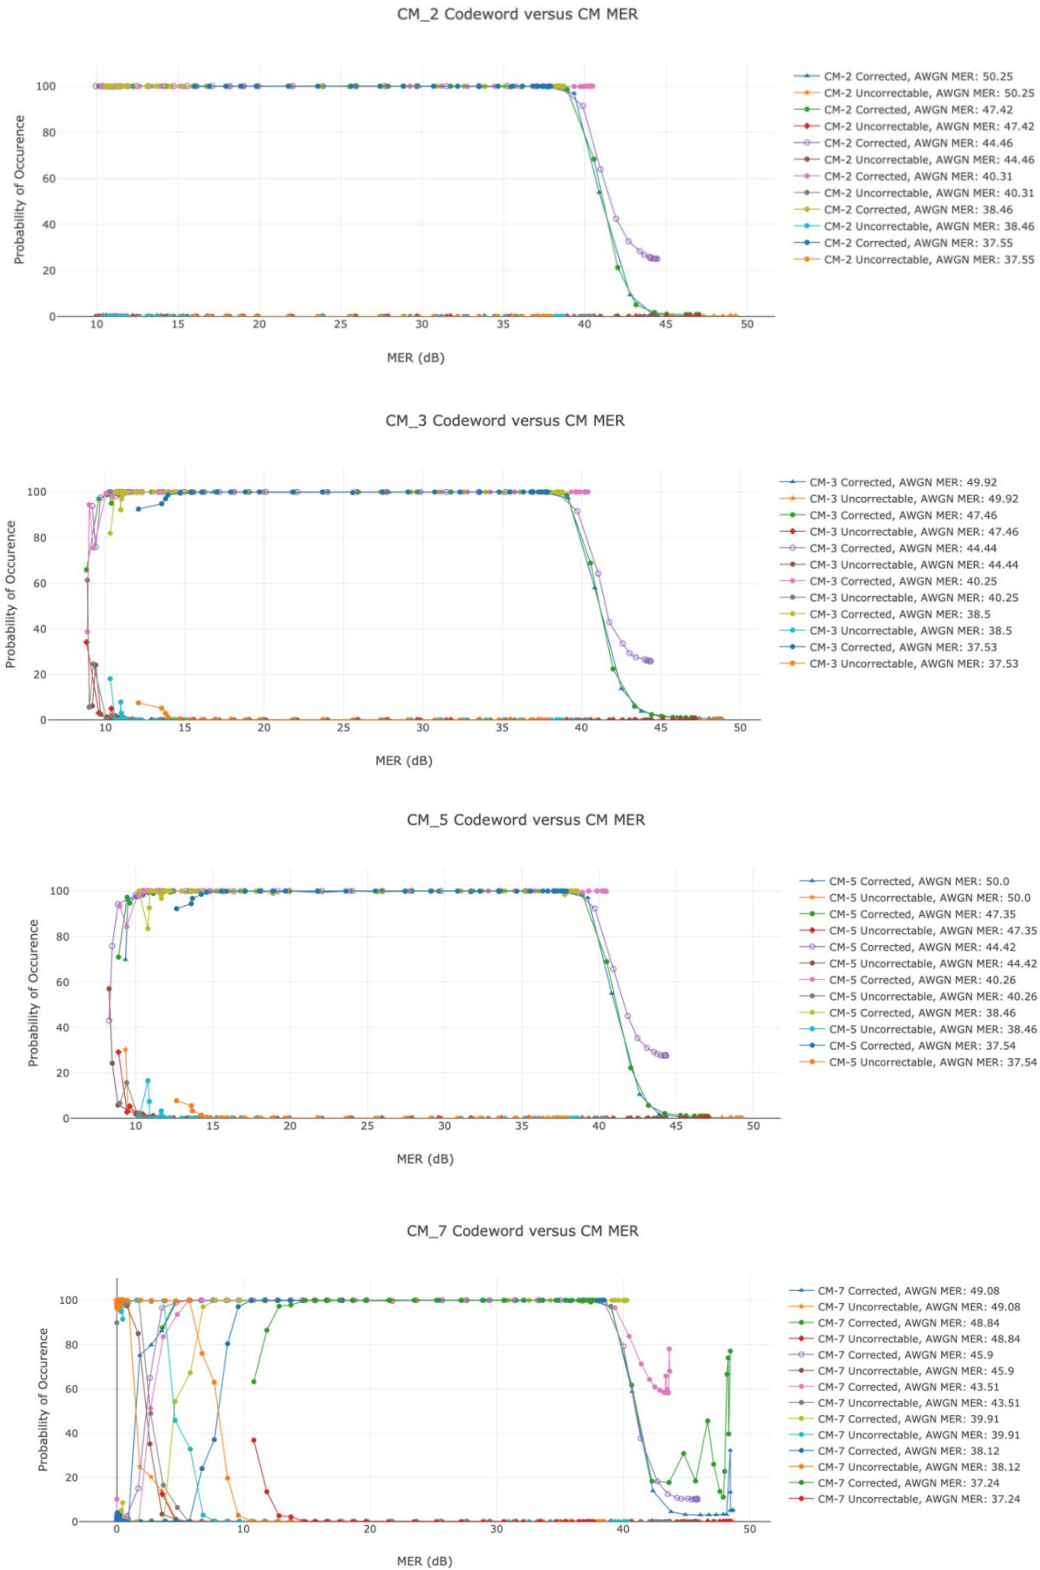


Figure 147 - Individual CMs in each Plot, Showing Various Noise Floors, and Increasing the LTE Level to Lower RxMER Values in the LTE Frequency Range; Correctable Codeword Errors and Uncorrectable Codeword Errors Proportions versus RxMER in the LTE Range

Figure 148 below shows the same results as above (Figure 147), but plotting the average RxMER values instead of the average RxMER in the LTE band.

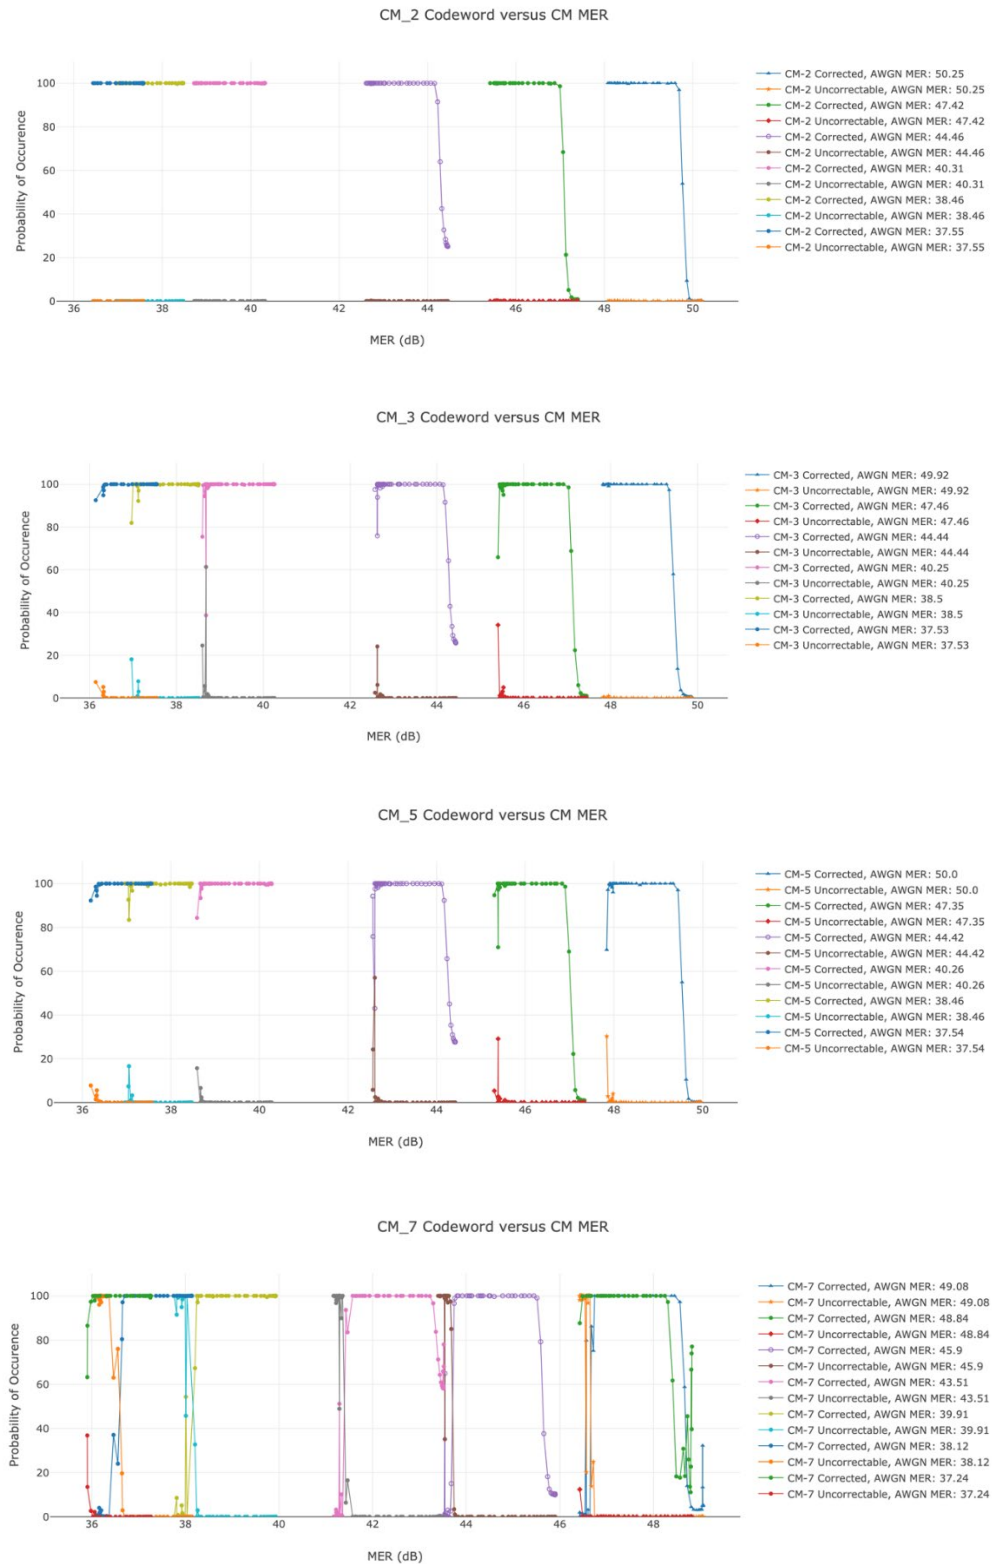


Figure 148 - Individual CMs in each Plot

Figure (above) shows different values of noise floor, and increasing the LTE level to lower RxMER values in the LTE frequency range; this time the results are plotted against average RxMER over the signal, not just the LTE.

It doesn't seem to matter how much noise floor there is; the LTE level dominates.

At 100% correctable codeword errors you are vulnerable (about 2 dB loss overall and uncorrectable codeword errors could begin). Therefore, again, when considering issues like in-channel ingress, to the extent it's possible, you don't want to have CMs operating with correctable codeword errors.

However, it may be helpful to start looking at correctable codewords, then examining the RxMER per subcarrier data for "notch" noise, and measuring the RxMER values in the notch, to determine how close the CM is to taking on uncorrectable codeword errors.

I.4 Conclusions

It is normal and expected behavior to see 100% correctable codeword errors in the downstream when the RxMER per subcarrier is within about 6 dB of the RxMER threshold for the modulation order in use. Good engineering practice suggests maintaining at least 6 dB of operating headroom above the RxMER threshold for reliable operation.

Following this guidance, the minimum recommended operating RxMER for 4096-QAM is about 41~42 dB, which is about 6 dB above the approximately 35~36 dB threshold for that modulation order (see the graphs in the Addendums section). As long as the cable network remains stable and impairment-free – and there are no uncorrectable codeword errors – the correctable codeword errors can be ignored.

Some cable operators choose a smaller amount of headroom, as low as 2 dB or 3 dB, above the RxMER threshold.

While this may be satisfactory in a stable and impairment-free plant, the presence of impairments can increase the risk of degraded performance.

Observations from the lab testing suggest that, while CMs can indeed operate reliably within about 6 dB of the RxMER threshold, the risk of problems increases when in-channel spurious noise and interference are present. Degraded performance may be more likely to occur, such that the issue(s) should be addressed proactively. Even a small bit of ingress could severely impact service. At the very least, if a CM is experiencing any amount of correctable codeword errors, and there is evidence of ingress or similar in-channel interference, even though small, it should be addressed.

In fact, these results suggest that any ingress should be addressed in a cost-effective way.

Further, it is best to have an OFDM signal wide enough that any expected amount of LTE or other interference does not cause CMs to lose connection. For instance, when the bandwidth of LTE ingress exceeds approximately 5% of the bandwidth of the OFDM signal, CM performance is more likely to be impacted.

In all cases, some form of profile management is recommended, whether built-in to the CMTS/CCAP or part of a third-party or standalone tool. With profiles that can tailor to noise issues in some bands, further resiliency can be had, and an imperfect profile can also indicate an issue in the plant that should be addressed before disrupting service too. The combination of profile management and sufficient operating headroom will ensure the most reliable overall OFDM performance.

Appendix II Future Technology

II.1 Technology Evolution

The cable industry has benefitted from seemingly nonstop technology evolution since it began in 1948. Since 2008, that evolution has paralleled developments in PNM. Though this document focuses on PNM for DOCSIS 3.1 systems, PNM will no doubt play a role in future network architectures and technology. The following sections discuss recently introduced and possible future technologies.

II.1.1 Technology Milestones

To help understand where the cable technology might be headed, it is helpful to understand where it has been.

II.1.1.1 *From Vacuum Tubes to Solid State*

The first community antenna television (CATV) systems were built in the late 1940s and served primarily rural communities where over-the-air television reception was difficult or impossible. Vacuum tube amplifiers were used for signal distribution (the amplifiers were installed in large metal boxes on utility poles). Early cable networks carried 1 to 3 television channels; later, up to 12 channels were provided to subscribers, with a maximum operating bandwidth of 220 MHz. The first solid state, i.e., transistorized, amplifiers were "drop-in" replacements for some of the vacuum tube amplifiers. Some solid state amplifier designs of the 1960s more closely resemble their modern counterparts. Today's outside plant actives support operation on frequencies as high as 1.2 GHz and can carry the equivalent of hundreds of video channels, telephony, and high-speed data services.

II.1.1.2 *Satellites*

In 1975 a small regional company called Home Box Office became the first program service to use geostationary satellite technology to deliver television programs to cable systems. Other program services soon followed, and the use of satellites revolutionized the industry, providing justification for cable operators to obtain franchises (licenses) and build networks in larger metropolitan areas. Geostationary satellites are still used today by the cable industry.

II.1.1.3 *Network Architectures*

Until the late 1980s, cable networks used an all-coax tree-and-branch architecture. Tree-and-branch was a cost-effective "broadcast" architecture. On the downside, cascaded active devices degraded signal quality, and cascaded components (actives, passives, power supplies, connectors, etc.) reduced reliability.

The year 1988 marked the beginning of what is now called hybrid fiber-coax (HFC) architectures. Multichannel amplitude modulation-based optical fiber links augmented coaxial cable to distribute services to cable subscribers. HFC architectures segmented cable networks into smaller serving areas, initially with several thousand homes passed per node. The use of fiber reduced cascaded devices, improved quality and reliability, and reduced operating costs. These benefits carry over today as cable operators push fiber deeper into the network, further reducing use of cascaded devices and components and improving overall reliability and availability. What is called a node+0 architecture (no actives after the node) can have as few as 50 homes passed per node.

II.1.1.4 *From One-Way to Two-Way*

Most early cable systems were one-way, which means that signals were transmitted through the cable network in one direction only: from cable company to subscriber. When the cable industry embarked on major network upgrades in the 1990s, two-way operation became widespread. In a two-way cable network, signals are transmitted to and from the subscriber simultaneously but in different frequency ranges. Figure 149 shows an example of a subsplit band plan, in which downstream signals are transmitted to the subscriber from 54 MHz to 1.2 GHz and upstream signals are transmitted from the subscriber in the range from 5 MHz to 42 MHz. Other band plans use different frequency splits, with the upstream spectrum operating to as high as 204 MHz and the downstream spectrum starting at 258 MHz.



Figure 149 - Example of Subsplit Band Plan Used in North America and Other Regions

II.1.1.5 High-Speed Data over Cable

Deployments of high-speed data service over cable networks in the early- to mid-1990s used proprietary technology, which meant there was no interoperability among different manufacturers' products. If a cable operator purchased modems from a given vendor, the operator also had to use that vendor's headend equipment.

In 1996, the cable industry became involved in an effort to develop cable modem standards, primarily to promote interoperability. Similar work had started two years earlier under the auspices of the Institute of Electrical and Electronics Engineers, Inc. (IEEE) 802.14 Cable TV Media Access Control and Physical Protocol Working Group.

In January 1996, four cable operators—Cox Communications, Comcast Cable Communications, Tele-Communications, Inc. (TCI), and Time Warner Cable—formed a coalition known as Multimedia Cable Network System Partners Limited (MCNS). MCNS issued a request for proposal (RFP) for the research and publication of interface specifications for high-speed data services using cable modems. An interoperability specifications consortium comprising Continental Cablevision, Rogers Cablesystems, and Cable Television Laboratories (CableLabs) joined MCNS, and the group announced its preliminary Data-Over-Cable Service Interface Specifications (DOCSIS) in December 1996. The final MCNS specification, issued in March 1997, supported nominal downstream data rates from 27 to 38 Mbps and upstream data rates from 320 kbps to 10 Mbps.

The International Telecommunications Union (ITU) officially approved the DOCSIS 1.0 specification as an international standard in March 1998, and the first DOCSIS 1.0 CMTS and CMs were qualified and certified respectively by CableLabs in early 1999. Given the success of DOCSIS technology, the IEEE's 802.14 efforts were formally disbanded at the end of 1999.

The DOCSIS 1.0 specification was followed by DOCSIS 1.1, 2.0, and 3.0 specifications. The DOCSIS 3.1 suite of specifications was introduced in October 2013. DOCSIS 3.1 technology supports scalability to 10+ Gbps in the downstream and 1+ Gbps in the upstream. It is being used by cable operators to provide gigabit-class data service.

II.1.1.6 Distributed Access Architectures

II.1.1.6.1 Remote PHY

Remote PHY (R-PHY) is a network architecture in which the media access control (MAC) layer is separated from the physical layer (PHY). The MAC layer remains in the CMTS/CCAP core (in the headend or hub) and the PHY electronics are located in a shelf or node. A 10-Gb Ethernet digital fiber link connects the core and a remote PHY device (RPD), which is installed in an R-PHY shelf or R-PHY node. Among the benefits of R-PHY are the ability to overcome limitations of CMTS/CCAP core port scalability as fiber is pushed deeper into the network, as well as improved RF performance because of the switch from analog fiber links to digital fiber links.

Several vendors have R-PHY and related products available today, and interoperability testing has been underway at CableLabs and in the labs of various cable operators. Interoperability testing ensures that RPDs and cores (physical or cloud native) from different vendors work together. Some operators are in (or are planning) lab and field trials; others are already deploying R-PHY technology in fiber-deep (e.g., node+0) and business-as-usual (e.g., node+m) applications and in new builds, rebuilds, and node splits.

II.1.1.6.2 Flexible MAC Architecture

Originally called Remote MAC-PHY, in which both the MAC and PHY electronics are located in a node, the concept has been changed somewhat to Flexible MAC Architecture (FMA). As its name suggests, FMA supports flexibility in the location of the MAC layer. The target architecture includes a remote MACPHY device (RMD), installed in a shelf or node; a remote MAC core (RMC), installed in a shelf or node; a physical MAC core (pCore),

located in a separate chassis; and a virtual MAC core (vCore), located in a server. Benefits are similar to R-PHY architectures and include a claimed reduction in latency between the MAC and PHY.

II.2 What's Next

II.2.1 Full Duplex DOCSIS Technology

Full duplex (FDX) DOCSIS technology is an extension of DOCSIS 3.1 technology and supports symmetrical multi-gigabit service offerings (e.g., 5 to 10 Gbps). The latter is achieved in part by simultaneous transmission of downstream and upstream signals on the same frequencies (see Figure 150). A key part of FDX DOCSIS technology is expansion of the upstream spectrum to as high as 684 MHz (or more) versus 42 MHz or 85 MHz today, yielding the potential for 10 to 50 times more upstream data throughput. FDX DOCSIS technology is optimized for fiber deep (node+0) with R-PHY nodes. The FDX-capable node includes echo cancellation to support simultaneous use of the same frequency ranges by downstream and upstream signals. The reader is directed to the most recent version of the DOCSIS 4.0 Physical Layer Specification [PHYv4.0] for more information.




Figure 150 - Frequency Spectrum Usage in FDX DOCSIS Systems (Bottom) Compared to Traditional Subsplit, Midsplit, or High-Split Band Plans (Top)

II.2.2 FDX Amplifier

In order to use FDX DOCSIS technology in architectures beyond node+0, some vendors have proposed the concept of FDX amplifier technology, in which the amplifier also is equipped with echo cancellation. In principle, an FDX amplifier would support deployment of FDX DOCSIS technology in HFC architectures with amplifier cascades after the node, for example, node+2 (or more).

II.2.3 Extended Spectrum DOCSIS Technology

The [PHYv3.1] specification defines a downstream upper frequency limit of 1,218 MHz (optional to 1,794 MHz). The concept of extended spectrum DOCSIS technology embraces operation at frequencies above 1.2 GHz, to as high as 3 GHz, to support downstream data speeds of up to 25 Gbps, and upstream speeds of 5+ Gbps. Figure 151 shows examples of extended spectrum DOCSIS technology with and without FDX.




Figure 151 - Extended Spectrum DOCSIS Frequency Usage With and Without FDX DOCSIS Technology

II.2.4 10G Initiative

The 10G initiative was announced in early 2019 at the Consumer Electronics Show by NCTA, CableLabs, and Cable Europe. The initiative's speed target for cable networks is symmetrical 10 Gbps and will offer lower latency and better security and reliability. Multiple technologies are part of 10G: DOCSIS 3.1 and 4.0 technology, full duplex, coherent optics, wireless, and more. For additional information on the 10G initiative, see <https://www.cablelabs.com/10g>.

II.2.5 Remote PHY Tap

Moving the remote PHY device to the tap and feeding the tap with digital fiber could potentially support a dedicated RF spectrum to each subscriber. Taking advantage of the relatively short length of the typical subscriber drop and the very high transverse-electromagnetic (TEM) mode cutoff frequency for drop-size coaxial cable, it may be possible for an R-PHY tap to support operating frequency ranges as high as 8 GHz to 12 GHz, with data speeds of 100 Gbps in the downstream and 25+ Gbps in the upstream.

II.2.6 Wireless PNM

Wi-Fi is the most common way customers access cable services, but the overwhelming majority of customer support calls are Wi-Fi related. Whether measured in direct support costs or customer churn, it is an expensive problem to have, making the development of Wi-Fi PNM critical. In June 2019, the Wi-Fi Alliance introduced Wi-Fi CERTIFIED Data Elements™, which provides a standardized data model targeting Wi-Fi performance and gives operators reliable, actionable information. Predictive analytics makes proactive troubleshooting possible and allows remote corrective action to take place before the customer notices an issue. For example, if neighboring Wi-Fi is using overlapping channels, then intelligent analysis of Wi-Fi PNM data can reveal the problem, and instructions for a channel change can be sent to the CPE. Not only does this help the direct customer, but the net impact is improved service for all—especially critical in multiple dwelling units. As cable operators work to improve customer experience and reduce operations costs, Wi-Fi PNM will be a critical pillar in any solution architecture.

II.2.7 Optical PNM

Beyond coherent optics and its unique relationship to cable services, optical networking is becoming a more prominent part of how cable services are provided. Optical technology can be more reliable, but it is often more complex and can still fail. In fact, as optical elements are used nearer to the core of networks, they often carry more traffic, sometimes for more customers, than a single coax cable. This setup creates a larger risk group, which often makes the impact of failures more significant. For example, analog dense wavelength division multiplexing (DWDM) and direct-detection or coherent optics make general monitoring of communication performance on optical links more important. But unlike with DOCSIS systems, operators do not have the same resiliency to problems in some optical networks, so they have fewer chances to be proactive. Still, optical impairments can be monitored, and monitoring often indicates a problem before service is impacted.

Optical networks often have redundancy in the core for protection and restoration, allowing services to be carried with higher reliability and availability than on a single connection. Monitoring these connections affords proactive repair in that though a failure of a single link means higher latency or greater risk, service may be maintained. Thus, repairs on these networks are proactive from a service perspective even though they are reactive from a network perspective.

Even simplex networks, often closer to the edge, have opportunities for proactive maintenance. Optical networking relies on methods like DOCSIS technology to provide bit- and packet-level protection, as well as error correction methods. These measures work like DOCSIS technology to create a gap between a network impairment and a service failure so that operations can receive an indication of a problem before service is impacted. Bit errors can indicate a number of problems, such as a laser or driver circuit in a small form factor pluggable transceiver about to fail or an optical cable stressed by a physical issue. And like DOCSIS systems, optical networks have time domain reflectometry technology that allows impairments to be located and isolated to aid in repair.

For these reasons and use cases, plus many more not described here, the cable industry will rely on PNM solutions for the optical portions of the network. Optical PNM is under development and will continue to progress in the future.

II.2.8 Coherent Optics

The evolution of HFC networks towards fiber-deep strategy and distributed access architectures (DAA) will lead the digital optical portion in cable networks to increase. Digital optics, especially the introduction of coherent optics into

access networks, enables rich intelligence through digital signal processing (DSP) that takes place within optical transceivers. Many of the same PNM processes that the industry leveraged in the RF domain can be applied to the optical domain through advanced trend analysis of characteristic performance parameters. Coherent optical systems provide a better way to support fault forecasting, detection, diagnosis, and localization on the optical terminals and link systems. Additionally, they can also provide signal health information for both traditional analog and intensity-modulated access networks in coexisting optical transmission systems.

II.2.9 Fiber to the Home (FTTH)

Fiber to the home (FTTH) is not new technology, but it has not been deployed widely by the cable industry (at least not in brownfield applications) because of the cost. FTTH has been deployed in some new-build (greenfield) applications over the years, using radio frequency over glass (RFoG) or passive optical network (PON) technology. It is expected that FTTH will become more economical over time, and variations of it could be deployed on a more widespread basis.

II.2.10 Conclusion and Suggestions for Future Work

It is anticipated that PNM can and will be developed for future cable technologies. At its heart, PNM is about service resiliency. A resilient service will continue to deliver with quality despite network impairments and failures. As long as the industry can identify failure modes for network technologies and create resiliency mechanisms for networks, it can identify how to measure and manage these resiliency methods and use them to create opportunities for PNM.

Appendix III Further Reading

The following publications can be useful for learning more about DOCSIS 3.1 technology, proactive network maintenance, and related topics. This list is not exhaustive, nor is it intended to imply an endorsement by CableLabs of third-party resources or reference material.

III.1 CableLabs DOCSIS 3.1 Specifications

- DOCSIS 3.1 Physical Layer Specification
<https://www.cablelabs.com/specifications/CM-SP-PHYv3.1>
- DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification
<https://www.cablelabs.com/specifications/CM-SP-MULPIv3.1>
- CCAP Operations Support System Interface Specification
<https://www.cablelabs.com/specifications/CM-SP-CCAP-OSSIV3.1>
- Cable Modem Operations Support System Interface Specification
<https://www.cablelabs.com/specifications/CM-SP-CM-OSSIV3.1>
- DOCSIS 3.1 Security Specification
<https://www.cablelabs.com/specifications/CM-SP-SECv3.1>

III.2 CableLabs Remote PHY Specifications

- Remote PHY OSS Interface Specification
<https://www.cablelabs.com/specifications/CM-SP-R-OSSI>
<https://www.cablelabs.com/specifications/CM-SP-R-OSSI>
- Remote PHY Specification
<https://www.cablelabs.com/specifications/CM-SP-R-PHY>
- Remote Upstream External PHY Interface Specification
<https://www.cablelabs.com/specifications/CM-SP-R-UEPI>
- Remote Downstream External PHY Interface Specification
<https://www.cablelabs.com/specifications/CM-SP-R-DEPI>
- Remote Out-of-Band Specification
<https://www.cablelabs.com/specifications/CM-SP-R-OOB>
- Remote DOCSIS Timing Interface Specification
<https://www.cablelabs.com/specifications/CM-SP-R-DTI>
- Generic Control Plane Specification
<https://www.cablelabs.com/specifications/CM-SP-GCP>

III.3 SCTE•ISBE Cable-Tec Expo Papers

<https://www.nctatechnicalpapers.com/Paper/2019>; papers from other years available from this page

- Campos, L.A., et al., "Proactive Network Maintenance Evolution to the Optical Domain in Coherent Optics." 2018 SCTE•ISBE Cable-Tec Expo Proceedings
- Topazi, C., Cooper M., "Operational Considerations and Optimization of OFDM Deployments." 2018 SCTE•ISBE Cable-Tec Expo Proceedings
- Wolcott, L., et al., "A PNM System Using Artificial Intelligence, HFC Network Impairment, Atmospheric and Weather Data to Predict HFC Network Degradation and Avert Customer Impact." 2018 SCTE•ISBE Cable-Tec Expo Proceedings

- Downey, J., "DOCSIS® 3.1 Downstream Early Lessons Learned." 2017 SCTE•ISBE Cable-Tec Expo Proceedings
- Druse, S., Miller, J., "DOCSIS® 3.1 Leaves the Lab and Hits the Field with Midco." 2016 SCTE•ISBE Cable-Tec Expo Proceedings
- Salinger, J., "DOCSIS® 3.1—Experiences from Early Deployments." 2016 SCTE•ISBE Cable-Tec Expo Proceedings
- Rupe, J., "A General-Purpose Operations Cost Model to Support Proactive Network Maintenance and More." 2019 SCTE•ISBE Cable-Tec Expo Proceedings
- Hranac, R., Medlock, J., Currivan, B., Fish, R., Kolze, T., Rupe, J., Williams, T., Wolcott, L., "Characterizing Network Problems Using DOCSIS® 3.1 OFDM RxMER per Subcarrier Data." 2019 SCTE•ISBE Cable-Tec Expo Proceedings
- Rupe, J., Zhu, J., "Kickstarting Proactive Network Maintenance with the Proactive Operations Platform and Example Application." 2019 SCTE•ISBE Cable-Tec Expo Proceedings
- Currivan, B., Wolcott, L., "Leveraging DOCSIS 3.1 PNM Measurements for Plant Performance Optimization." 2015 SCTE•ISBE Cable-Tec Expo Proceedings
- Wolcott, L., Heslip, J., Thomas, B., Gonsalves, R., "A Comprehensive Case Study of Proactive Network Maintenance." 2016 SCTE•ISBE Cable-Tec Expo Proceedings
- Campos, A., Cardona, E., Raman, L., "Pre-Equalization Based Pro-Active Network Maintenance Methodology." 2008 SCTE•ISBE Cable-Tec Expo Proceedings

III.4 Articles

- Hranac, R., "What is OFDM?" Communications Technology, November 2012, <https://www.scte.org/TechnicalColumns/12-11-30%20what%20is%20ofdm.pdf>.
- Hranac, R., "An Introduction to PNM." Broadband Library, Spring 2019, <https://broadbandlibrary.com/an-introduction-to-pnm/>.
- Schley, S., "Advancing the Echo-System." Broadband Library, Spring 2019, <https://broadbandlibrary.com/advancing-the-echo-system/>.
- Volpe, B., "What Can DOCSIS PNM Do for Me?" Broadband Library, Spring 2019, <https://broadbandlibrary.com/what-can-docsis-pnm-do-for-me/>.
- Hranac, R., "How Adaptive Equalization Works." Broadband Library, Spring 2019, <https://broadbandlibrary.com/how-adaptive-equalization-works/>.
- Bastian, C., "PNM and Associated Standards." Broadband Library, Spring 2019, <https://broadbandlibrary.com/pnm-and-associated-standards/>.
- Wolcott, L., "PNM: State-of-the-Art and 2019 Roadmap." Broadband Library, Spring 2019, <https://broadbandlibrary.com/pnm-state-of-the-art-and-2019-roadmap/>.
- Rupe, J., "PNM at CableLabs." Broadband Library, Spring 2019, <https://broadbandlibrary.com/pnm-at-cablelabs/>.
- Campos, A., "Frontiers in PNM." Broadband Library, Spring 2019, <https://broadbandlibrary.com/frontiers-in-pnm/>.
- Marchetti, E., "The Business Perspective of Proactive Network Maintenance." Broadband Library, Spring 2019, <https://broadbandlibrary.com/the-business-perspective-of-proactive-network-maintenance/>.
- Finkelstein, J., "Full Band Capture Takes the Stage." Broadband Library, Spring 2019, <https://broadbandlibrary.com/full-band-capture-takes-the-stage/>.

- Dzuban, M., "PNM: In the 10G Era, More Important than Ever." Broadband Library, Spring 2019, <https://broadbandlibrary.com/pnm-in-the-10g-era-more-important-than-ever/>.

III.5 Training Videos and Programs

- Hranac, R., "DOCSIS® 3.1—An Overview." SCTE—Society for Broadband Professionals, Autumn 2016 Lecture Meeting, London (video, approx. 2 hr), <https://www.youtube.com/watch?v=ZwKuPZVh19Y>.
- SCTE•ISBE training videos/programs
 - DOCSIS® 3.1 Essentials
 - DOCSIS® 3.1 Installation
 - DOCSIS® 3.1: Boot Camp for Engineers
 - DOCSIS® 3.1: Boot Camp Advanced Onsite Training
 - Proactive Network Maintenance

III.6 Other Useful References

- White papers by Brady Volpe: <https://volpefirm.com/broadband-consulting-services/docsis/>
- Supporting material including papers, presentations, study results, and examples available on the CableLabs' PNM project sites. For additional information, contact info@cablelabs.com.

Appendix IV Acknowledgements

On behalf of the cable industry and our member companies, CableLabs would like to thank the following individuals that contributed to the development of this document.

Contributor	Company Affiliation
James Medlock	Akleza, Inc.
Roger Fish, Victor Hou, Tom Kolze	Broadcom
Bruce Currivan (contractor) Doug Jones, Jason Rupe, Tom Williams, Jay Zhu	CableLabs
Patricio Latini	Casa
Nick Segura	Charter
John Downey	Cisco
Robin Lavoie, Arni Lundale	Cogeco
Dan Torbet	Commscope
Maurice Garcia, Rich Prodan, Rob Thompson, Larry Wolcott	Comcast
Jeff Finkelstein	Cox
Bernie Cadieux	Devisor
Jack Moran, Ron Hranac	Individual contributors
Satish Mudugere	MaxLinear
Brady Volpe	Nimble This LLC
Kirk Erichsen	OAM Technologies
Nasir Ansari, George Hart	Rogers
Dean Stoneback	SCTE
Mike Darling	Shaw
Patrick Ryan	TDS
Aleksander Soberg, Even Christoffersen	Telia
Mike Gotwals	Viavi
Carlos Bustamante	Zcorum

Jason Rupe, CableLabs
