

JOURNAL OF NETWORK OPERATIONS



SCTE • ISBE™

Society of Cable Telecommunications Engineers
International Society of Broadband Experts

JOURNAL OF NETWORK OPERATIONS

VOLUME 5, NUMBER 1
February 2019

Society of Cable Telecommunications Engineers, Inc.
International Society of Broadband Experts™
140 Philips Road, Exton, PA 19341-1318

© 2019 by the Society of Cable Telecommunications Engineers, Inc. All rights reserved.

As compiled, arranged, modified, enhanced and edited, all license works and other separately owned materials contained in this publication are subject to foregoing copyright notice. No part of this journal shall be reproduced, stored in a retrieval system or transmitted by any means, electronic, mechanical, photocopying, recording or otherwise, without written permission from the Society of Cable Telecommunications Engineers, Inc. No patent liability is assumed with respect to the use of the information contained herein. While every precaution has been taken in the preparation of the publication, SCTE assumes no responsibility for errors or omissions. Neither is any liability assumed for damages resulting from the use of the information contained herein.

Table of Contents

4 From the Editors

Technical Papers

- 6 **An Outlook on Next-Generation Cable Networks**
Dr. Danish Rafique, Senior Manager AI & ML and Principal Engineer, ADVA Optical Networking
Ulrich Kohn, Director Technical Marketing, ADVA Optical Networking
Dr. Stephan Neidlinger, VP Strategic Alliance Management, ADVA Optical Networking
- 19 **Low Power Wide Area Networks**
Mayank Bhatnagar, Accenture Strategy
Byron Scerri, Accenture Strategy
- 30 **MoCA Alternative Channel Assessment: It's More Than Just A Squiggly Line**
Maurice Manuel Garcia, Senior Principal Engineer, Comcast Cable
- 67 **Operational Practices for the Deployment and Maintenance of Wireless Devices on HFC Plant**
Dan L. Whitehouse, Director, Communications Media and Content, Hitachi Consulting Corporation
Cole Reinwand, VP, WiFi Strategy, Comcast Corporation
Matt Zielenbach, Sr. Director, Xfinity WiFi Product, Comcast Corporation
- 92 **Preparing for the Network For 5G And IoT**
Fady Masoud, M. Eng, Principal, Product and Technology Marketing, Infinera
- 102 **The Future of Operations: Building a Data-Driven Strategy**
Sebnem Ozer, Ph.D., Principal Architect, Comcast
Sinan Onder, Executive Director, Comcast
Nagesh Nandiraju, Ph.D., Sr. Director, Comcast

Letter to the Editor

- 111 **The Application of Machine Learning to Alarm Management**
Doug Junkins, Field CTO, Guavus, Inc.

Editorial Correspondence: If there are errors or omissions to the information provided in this journal, corrections may be sent to our editorial department. Address to: SCTE Journals, SCTE•ISBE, 140 Philips Road, Exton, PA 19341-1318 or email journals@scte.org.

Accepted Submissions: If you have ideas or topics for future journal articles, please let us know. Topics must be relevant to our membership and fall under the technologies covered by each respective journal. All submissions will be peer reviewed and published at the discretion of SCTE•ISBE. Electronic submissions are preferred, and should be submitted to SCTE Journals, SCTE•ISBE, 140 Philips Road, Exton, PA 19341-1318 or email journals@scte.org.

Subscriptions: Access to technical journals is a benefit of SCTE•ISBE Membership. Nonmembers can join at www.scte.org/join.

SCTE•ISBE Engineering Committee Chair:

David Fellows, T-Mobile
SCTE Member

SCTE•ISBE Network Operations Subcommittee (NOS) Committee Chair:

Ron Hranac
SCTE Fellow

Senior Editors

Ron Hranac
SCTE Fellow
Daniel Howard
SCTE Senior Member

Publications Staff

Chris Bastian
SVP & Chief Technology Officer,
SCTE•ISBE

Dean Stoneback

Senior Director- Engineering & Standards, SCTE•ISBE

Kim Cooney

Technical Editor, SCTE•ISBE

SCTE • ISBE

From the Editors

Welcome to Volume 5 Issue 1 of the *Journal of Network Operations*, a publication of collected papers by the Society of Cable Telecommunications Engineers (SCTE) and its global arm, the International Society of Broadband Experts (ISBE). While last issue focused on conventional wireline-related topics, in this issue we focus on next generation cable networking technologies such as wireless, MoCA, the Internet of Things (IoT), data-driven management, and artificial intelligence (AI)/machine learning (ML), all of which continue to grow in importance and scale in modern cable networks.

Two papers in this issue provide operational practices for managing the deployment of new technology, especially wireless, in the cable access network at scale: Garcia shows that operators can add more intelligence to managing MoCA coexistence in the face of evolving access networks in “MoCA Alternative Channel Assessment: It’s More Than Just A Squiggly Line.” Garcia provides lessons learned and recommendations for how to make smarter decisions on when to exit a MoCA channel, techniques for point of entry (PoE) filter detection, and recommendations on leveraging MoCA 2.0 RF analysis for upgrading the network to support a DOCSIS 3.1 full-duplex (FDX)-based home assessment. Wi-Fi in the outside plant is also being deployed by many cable operators. In “Operational Practices for the Deployment and Maintenance of Wireless Devices on HFC Plant,” Whitehouse and colleagues give practical observations on deployment planning, installation, maintenance and ongoing operational support for deploying Wi-Fi at scale in a modern cable network. Rounding out the wireless additions to cable networks, in “Low Power Wide Area Networks,” Bhatnagar and Scerri review low power wide area network (LPWAN) technologies, provide examples, identify their advantages and give example strategies for cable operators to use in expanding their services and revenue.

Moving to the AI/ML focus for this issue, the power of AI-based prediction is increasingly showing up in modern cable networks, particularly where it can reduce the cost of prediction, improve the management of diversity, handle incomplete/voluminous data, and speed decision making. Kicking off our highlight of AI and IoT technologies for this issue, Rafique and colleagues will highlight two core technologies to enable digital point-to-point optics in next generation networks: coherent optics and machine learning (ML)-enabled software-defined networking (SDN) operations. In “An Outlook on Next-Generation Cable Networks: Coherent Optics and Machine Learning-Based Operations,” the authors show how AI-enabled, self-driven networking will significantly simplify network operations, leading not only to operational expenditures (OPEX) savings, but also an enhanced customer experience in the highly flexible, adaptive networks of the future. In a letter to the editor titled “The Application of Machine Learning to Alarm Management,” Doug Junkins describes the power of AI-based incident prediction to dismiss the overwhelming majority of network management alarms. Since the majority of these alarms are typically “noise” and non-actionable, AI permits network operations center (NOC) personnel to focus on the 2 to 3 percent of alarms that are related to actual incidents.

Ozer and colleagues show how a unified orchestration platform can process the massive data associated with modern networks, even when the access network includes DOCSIS, Ethernet passive optical networking (EPON), conventional Ethernet, and wireless technologies. In “The Future of Operations: Building a Data-Driven Strategy” they describe how a data-driven and knowledge-defined orchestration platform can be used to improve back office operations when SDN, network functions virtualization (NFV) and cloud-enabled network platforms must manage such diverse

access networks. And in what could be considered the pinnacle of AI/ML use in modern cable networks, Masoud introduces the concept of cognitive networking to enable self-optimized, self-healing, and highly autonomous cable networks. See “Preparing for the Network for 5G and IoT,” to understand how key building blocks, such as software defined capacity (SDC), multi-layer software-defined networking (SDN) control, and transport technology breakthroughs can enable networks that ‘think.’

We are grateful for the individuals who contributed to this issue of the *Journal of Network Operations*, including the authors, reviewers, and the SCTE-ISBE publications and marketing staff. We hope you enjoy this issue of the *Journal*, and that the selected papers provide inspiration for new ideas and innovations in cable network operation. If you have feedback on this issue, have a new idea, or would like to share a success story please let us know at journals@scte.org.

SCTE-ISBE *Journal of Network Operations* Senior Editors,

Ron Hranac

Technical Marketing Engineer – Cable Access Business Unit, Cisco Systems, Inc.

SCTE Fellow

Daniel Howard

Principal, Enunciant LLC.

SCTE Senior Member

An Outlook on Next-Generation Cable Networks

Coherent Optics and Machine Learning based Operations

A Technical Paper prepared for SCTE•ISBE by

Dr. Danish Rafique, Senior Manager AI & ML and Principal Engineer
ADVA Optical Networking
Fraunhoferstr. 9a, 82152
Martinsried/Munich, Germany
DRafique@ADVAoptical.com
+4989890665659

Ulrich Kohn, Director Technical Marketing
ADVA Optical Networking
Fraunhoferstr. 9a, 82152
Martinsried/Munich, Germany
UKohn@ADVAoptical.com
+4989890665867

Dr. Stephan Neidlinger, VP Strategic Alliance Management
ADVA Optical Networking
Fraunhoferstr. 9a, 82152
Martinsried/Munich, Germany
SNeidlinger@ADVAoptical.com
+4989890665826

Table of Contents

Title	Page Number
Table of Contents	7
1. Introduction	8
1.1. Digitization of Cable Networks	8
1.2. Executive Summary	9
2. Point-to-Point Coherent Optics	10
2.1. Technologies	10
2.2. Proof-of-Concept – Short Reach Coherent Systems	10
3. AI-Enabled Cable Networks	12
3.1. Software-defined Networking	12
3.2. Self-driven Networking	12
3.3. Proof of Concept – Predictive Maintenance	13
4. Conclusions	15
5. Abbreviations and Definitions	16
5.1. Abbreviations	16
5.2. Definitions	17
6. References	17

List of Figures

Title	Page Number
Figure 1 – End-To-End Cable Network Architecture Including (1) Coherent And (2) AI-Assisted SDN	8
Figure 2 – Left: High-Speed Coherent Optics. Right: Integrated Coherent Receiver Chip	10
Figure 3 – Left: Point-To-Point Coherent Transmission System. Right: Currently Commercial Component Bandwidth Responses (Transceiver Chain). MZM: Mach-Zehnder Modulator, OFE: Optical Front End	11
Figure 4 – 200 Gb/S And 400 Gb/S Coherent Performance (Single Channel) With Currently Commercial And Next-Generation Specifications. DPE: Digital Pre-Emphasis	11
Figure 5 – Closed-Loop SDN – SEDN: Self-Driven Networking	13
Figure 6 – Evolution of Network Monitoring	13
Figure 7 – Left: T-SDN Controller With PFD Architecture, Right: Definitions For Typical Fault Patterns	14
Figure 8 – (a) Condition Based Fault Detection: Rule-Driven; (b) Data Based Fault Detection (PFD Engine): ML-Driven. Highlighted Symbols (Opacity) Represent Detected Faults.	15

1. Introduction

1.1. Digitization of Cable Networks

Traditional hybrid fiber coax (HFC) networks seamlessly transitioned RF analog transmission technology into the optical domain. Optimized for broadcast services, however, this approach created limitations regarding cable network operators’ aspirations to grow business from residential to enterprise services or tapping into the highly profitable wholesale connectivity market for mobile service providers.

The technology transition towards fiber deep and distributed access architectures in combination with point-to-point fiber and passive optical network (PON) technologies is continuously eradicating bandwidth limitations of legacy HFC networks. A range of new technologies will allow a cable network operator to flexibly grow the network in capacity while improving operational practices. Applying latest innovations in the underlying optical packet transport network as well as software-based control technologies are prerequisites for meeting agility and flexibility requirements of multi-service offerings. A combination of both technologies is the basis for highly economical service production. Recent progress with photonic technologies has made available cost efficient 100G technology, inherently well suited for multi-wavelength applications for the highest utilization of scarce fiber. What’s more, moving from manual processes to software defined networking (SDN) is simplifying control and management, and opening networks for fully automated, artificial intelligence (AI)-assisted intelligent operations.

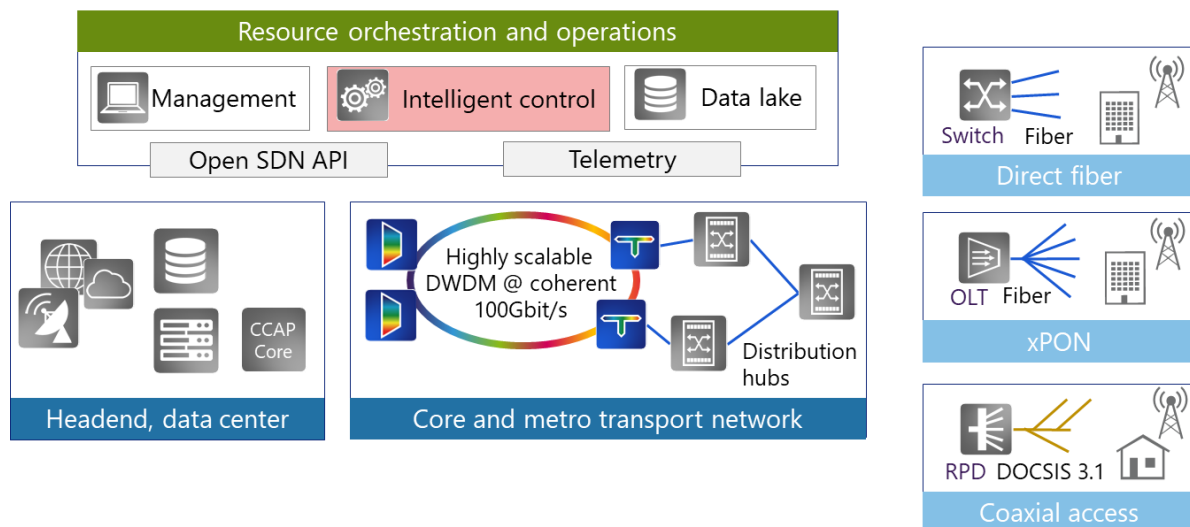


Figure 1 – End-To-End Cable Network Architecture Including (1) Coherent And (2) AI-Assisted SDN

The transition from analog video distribution to the distribution of digital services using technologies such as DOCSIS 3.1, and eventually digital coherent point-to-point optics in cable access networks has a direct impact on several aspects of network design, implementation and operation. CableLabs is helping to evolve the market in a future-proofed direction since the need to use constrained fiber resources efficiently becomes increasingly important as capacity demands increase. Coherent optics technology applies the types of modulation techniques used in RF to transmit multiple bits simultaneously, resulting in orders of magnitude more capacity in the fiber network. CableLabs’ coherent optics specifications have been

developed to help drive down the cost of this technology by increasing market scale and competition through simplification and common, interoperable interfaces [1].

Setting aside optimizations and assurance mechanisms, running applications on such an evolved infrastructure is a challenging task in itself. While an overwhelming undertaking in terms of relevant expertise development, the digitized cable access market dynamics make it mandatory for network providers to run lean and automated operations. In this context, SDN is considered a game-changer by many owing to its flexibility and scalability, in contrast to traditional control and management platforms. SDN distances itself from proprietary, device-specific operation to open, resource-driven control, enabling centralized, programmable, and automated services.

While centralized control and management is one piece of the puzzle, data-driven operations is equally important for cognitive assurance (briefly described as intelligent and automatized network operations, ensuring appropriate service levels and quality of experience, among other aspects) and design mechanisms. These systems can streamline technology innovation by supporting the services teams with a rich set of automated design, activation and maintenance tools leveraging various data analytics solutions. Such frameworks can offload the design teams with self-learning planning systems, field-services with fully automated configuration and activation solutions and operational teams with OSS building on predictive maintenance, among many others.

1.2. Executive Summary

This paper highlights two core technologies to enable next-generation cable networks: coherent optics and machine learning (ML)-enabled SDN operations.

Firstly, the application of coherent optics in cable networks and the potential technologies to re-design and re-engineer the coherent transceiver are discussed, specifically focusing on optoelectronic component integration, driving down costs while delivering reduced footprint, increased capacity and better scalability. Proof-of-concept results demonstrating the operation of silicon photonic IQ-modulator with co-integrated BiCMOS segmented linear driver [2], termed as electronic photonic integrated circuit (ePIC) are shown for up to 64 Gbaud systems for short-reach cable access applications. In particular, the goal is to enable scalable and high capacity coherent optics with currently commercial high-volume and low-cost components.

The results show that 64 Gbaud coherent transceivers may be operated with tolerable performance penalties in the context of cable networks, owing to the available high OSNR. While performance may be improved using next-generation components, such choices are more suitable to data center interconnect (DCI) and metro application scenarios. In particular, we show that current generation components may incur a worst-case OSNR penalty, with regard to theoretical bounds (assuming SD-FEC), of 3.3 dB for 64 Gbaud DP-4-QAM (200 Gb/s) applications, whereas such penalties increase to ~5.1 dB for 64 Gbaud 16-QAM (400 Gb/s) scenarios.

Secondly, on the network control and management side, one of the key challenges in cable networking is how an operator can efficiently manage the network and utilize all the data made available by the network devices. The key components here relate to meaningfully gathering network data via advanced telemetry platforms, using ML for generating insights, and eventually using the produced knowledge for optimized network operations via SDN. This contribution will provide an ML-driven predictive maintenance use case, focusing on the architectural aspects and proof-of-concept results [3-4], establishing a robust business case for ML-driven networking, enabling low entry cost and immediate savings.

2. Point-to-Point Coherent Optics

2.1. Technologies

In order to enable higher bandwidth in next generation cable networks new initiatives have emerged in the recent past, targeting coherent optics initiatives that utilize higher order modulation formats such as quadrature amplitude modulation (QAM) [5]. These next-generation architectures necessitate lower cost per Gbit/s, smaller footprint (size) and lower power consumption, while at the same time an increase in net data rate per wavelength is expected. Dense integration of components such as modulators, receivers, drivers and transimpedance amplifiers (TIAs) can help to address the aforementioned requirements, while enabling higher RF performance due to decreased trace lengths. Silicon photonics can potentially leverage the mature fabrication and packaging processes established in the electronic IC industry and facilitate low cost transceiver assemblies suitable for next generation coherent transceivers. However, unlike traditional lithium niobate (LiNbO₃) based modulators that are widely used today, silicon photonic modulators exhibit strong nonlinear behavior in their phase response as a consequence of carrier dispersion effects. In addition, this mechanism of modulation entails voltage dependent losses due to free carrier absorption.

To address these challenges, the Silicon Photonics Enabling Exascale Data Networks (SPEED) project was started in November 2015 [6]. Sponsored by Germany's Federal Ministry of Education and Research (BMBF) and coordinated by ADVA Optical Networking, the project creates a platform for development, manufacturing and packaging of application-specific electro-photonic integrated circuits on silicon. ePICs combine electronic and optical functions on a single semiconductor chip, delivering better performance, smaller footprint and lower cost than competing solutions. The project currently develops up to 400 Gb/s transceivers (Figure 2: left), and test chips containing key blocks are already available for test and verification (Figure 2: right).

In context of access networks, the above initiative ensures not only repurposing long-haul and metro network system solutions for such new applications, but also considers scalable designs exploiting state-of-the-art lower cost initiatives for high-speed coherent transmission systems.




Figure 2 – Left: High-Speed Coherent Optics. Right: Integrated Coherent Receiver Chip

2.2. Proof-of-Concept – Short Reach Coherent Systems

Figure 3 shows the system setup, where pseudo-random binary sequences (2^{18}) were mapped to 16-QAM or 4-QAM constellations, and 2.5% training overhead was added. The digital signal processing (DSP) prior to the digital to analog converter (DAC) included root-raised-cosine pulse shaping (roll-off factor of 0.2), and – optionally – digital pre-emphasis (DPE) [7]. The signals were then fed, together with a 100 kHz laser (1550 nm), to the ePIC driver-modulator structure, resulting in DP-QAM signals. At the receiver, noise loading was varied to enable optical signal-to-noise ratio (OSNR) sweeps, followed by coherent detection,

employing 90° hybrid, photodiodes and TIAs at the frontend. The signals were then digitized using the analog to digital converter (ADC). Finally, DSP was performed using conventional stages, and bit error ratio (BER) was calculated. Figure 3 also shows the S_{21} responses for the individual components and the complete transmit and receive portions of the entire transceiver chain, modeling current commercial bandwidth specifications, and experimentally measured component responses. The total transmitter and receiver -3 dB bandwidth is 11.5 GHz and 15.5 GHz, respectively. Note that while the -3 dB bandwidth is the typically reported specification, S_{21} roll off is also a critical design parameter, captured by the -6 dB bandwidth in presented curves.

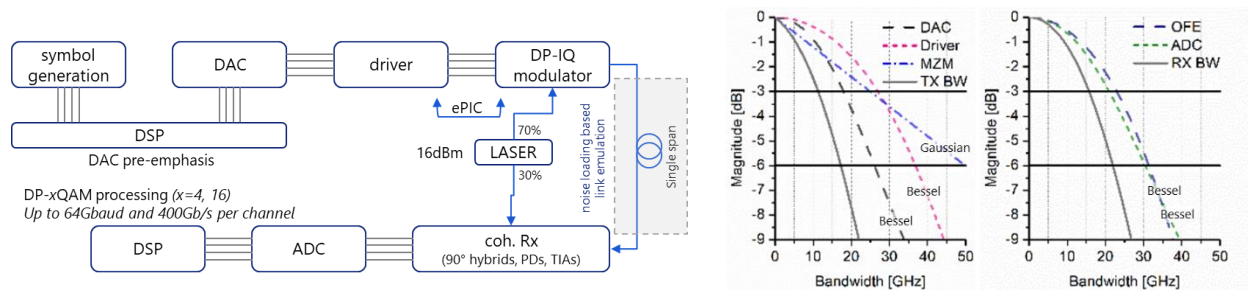


Figure 3 – Left: Point-To-Point Coherent Transmission System. Right: Currently Commercial Component Bandwidth Responses (Transceiver Chain). MZM: Mach-Zehnder Modulator, OFE: Optical Front End

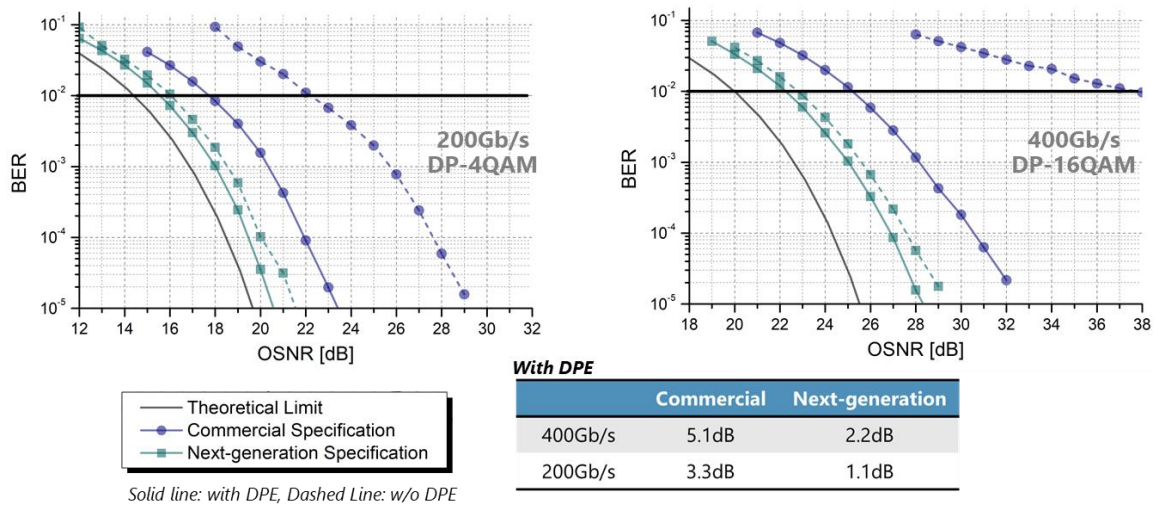


Figure 4 – 200 Gb/S And 400 Gb/S Coherent Performance (Single Channel) With Currently Commercial And Next-Generation Specifications. DPE: Digital Pre-Emphasis

Figure 4 depicts waterfall curves for current commercial components (shown in Figure 3), and high-end next-generation components as well (see [2, 8] for specifications), potentially targeting cable and DCI applications, respectively. As expected, OSNR penalties, with regard to theory, reduce with decreasing modulation order, best performance is enabled by next-generation components, and that DPE enables substantial improvements, especially for higher order modulation at lower bandwidth specifications. It is worth mentioning that we have only pre-emphasized the DAC, as the extent of DPE (considering other

components) significantly trades off with transmit output power; and effectively transceivers may not be fully pre-emphasized due to amplifier-free transmission constraints.

The specific OSNR penalties, defined as the different between measured and theoretical performance, at a BER of 10^{-2} (a typical pre-FEC SD-FEC threshold assuming some system margin) are summarized in the tables. It can be seen that while high-end next-generation components enable penalties as low as 2.2 dB for 400 Gb/s and 1.1 dB for 200 Gb/s, the component specifications may be potentially relaxed for cable networks, at the cost of ~ 3 dB and ~ 2 dB worst-case penalty for 400 Gb/s and 200 Gb/s, respectively.

3. AI-Enabled Cable Networks

3.1. Software-defined Networking

Open communities such as Open Networking Foundation (ONF) and Open Network Operating System (ONOS) are investing major effort in separating the physical network hardware from network control by means of open SDN control interfaces. Those standardized interfaces are streamlining integration of multivendor networks and enabling service providers to orchestrate any resource in their data center as well as in their network in a consistent way, opening unprecedented opportunities for new service offerings as well as more efficient ways to operate the network. The value of SDN control has been explored in a multi-technology, multi-vendor demonstration showcasing the ability to provision services with simple definition of service intents [9].

Clearly, SDN architectures are not limited to traditional networks, and will help to simplify operations in cable networks as well. Fiber deep structures including remote PHY devices introduce many new network elements which need to be managed and controlled. Manual and distributed control approaches will not scale in such network environments. Consequently, centralized management and orchestration will be required in order to operate next generation cable access networks efficiently and OPEX optimized.

3.2. Self-driven Networking

The discussions around SDN have almost exclusively focused on separation of data and control planes, with little to no attention on an overall operational feedback loop that includes monitoring, intelligence and management functionalities. Figure 5 captures this theme in a self-driven networking (SeDN) architecture, where network resources—physical or virtual—are continuously monitored using a telemetry engine, exposing real-time network states to the analytics stage, which in turn feeds into the control and management planes. This holistic platform not only caters for centralized and programmable control, but also makes ML-driven decisions to trigger actions, essentially connecting data-driven automation with policy-based orchestration and management.

It is worth highlighting the importance of closed loop operation (CLO) for SeDN, as it fundamentally changes the way networks are operated today, empowering truly dynamic and autonomous operation. Consider the example of the path computation element, which identifies paths through the network based on a given set of constraints. SeDN, through its closed loop interactions and a consolidated view of both multi-layer infrastructure and services would aim to discover, and enforce, intent-driven application-aware routes based on historical, current and predicted network state information.

SeDN-enabled networks would represent a wide variety of rapidly evolving application stacks, consuming both physical and virtual resources. Traditional network management tools are unable to efficiently tap into this goldmine since they lack the capabilities to probe network states in real-time as well as other issues

related to scalability, vendor lock-in, etc. For instance, SNMP based monitoring largely relies on data access at fixed intervals, managed using trap-based alarms. While SNMP has served the industry long and well, network monitoring needs to rise to the new visibility requirements.

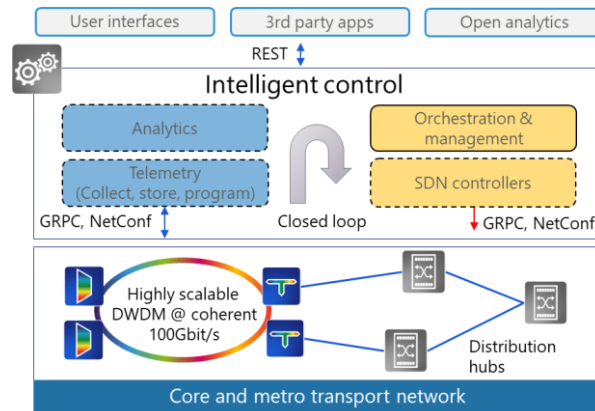


Figure 5 – Closed-Loop SDN – SEDN: Self-Driven Networking

On the other hand, model-driven streaming telemetry is not about a fancy way of data monitoring. Rather, it's defined by operational needs and requirements that are set by telecom network operators (e.g., OpenConfig, TAPI). It enables vendor-agnostic network state monitoring on a continuous basis using time series streams, and abstracts data modelling from data transport, leveraging unique open source initiatives, like YANG (model), GRPC (transport), etc. Furthermore, one of the key differentiators in moving from legacy monitoring to model-driven telemetry is the use of subscription based data access, as opposed to request- or trap-based desired data selection. Figure 6 depicts and enlists a few key differences between legacy and streaming telemetry solutions.

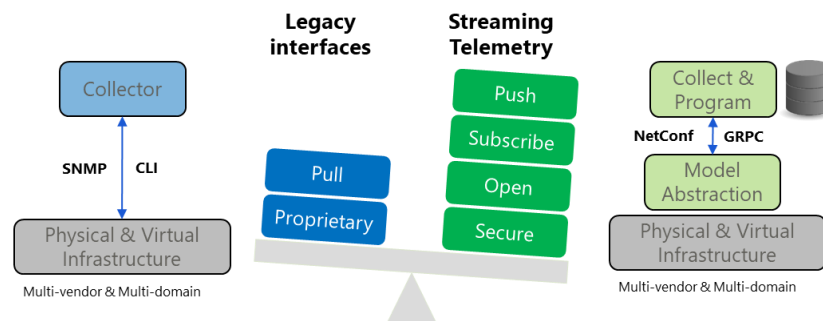


Figure 6 – Evolution of Network Monitoring

3.3. Proof of Concept – Predictive Maintenance

The capability to achieve a stateful view of the network based on streaming data monitoring is a must for SeDN. However, it's the analytics engine which enforces network behavior, and enables use cases like traffic engineering, predictive maintenance, capacity planning, etc. The goal of this engine is to incorporate intelligence via various ML and big data toolsets, learn hidden relationships, discover traffic patterns, find anomalous events, and recommend actions. In this subsection, we consider the predictive maintenance use case, and elaborate on relevant details.

Network disruptions are typically caused by underlying faults of network equipment, optical transmission lines, etc., where the behavior of individual entities contributes to ensemble network behavior. The principal challenge in such a complex interactive environment is not only to detect singular faults, but also to isolate and pinpoint the source of a failure. Furthermore, network downtimes can be prohibitively expensive, and any potential fault needs to be proactively discovered and addressed. Such high availability requirements have traditionally resulted in pessimistic service level agreements, necessitating over-designed transport networks with expensive protection schemes at multiple layers, and underutilization of available network resources. On the other hand, network operators have access to a plethora of data sources (SNMP, syslog, etc.), however, assurance solutions have remained limited to localized processing based on fixed threshold-triggered alarms, with little or no global context.

While service configuration and resource utilization frameworks have been proposed including intelligent network slicing and planning [10], integrated within flexible SDN frameworks, next-generation architectures need to incorporate a centralized cognitive operational assurance framework including streaming telemetry, fault detection, and proactive root cause analysis (RCA) for high network availability [11].

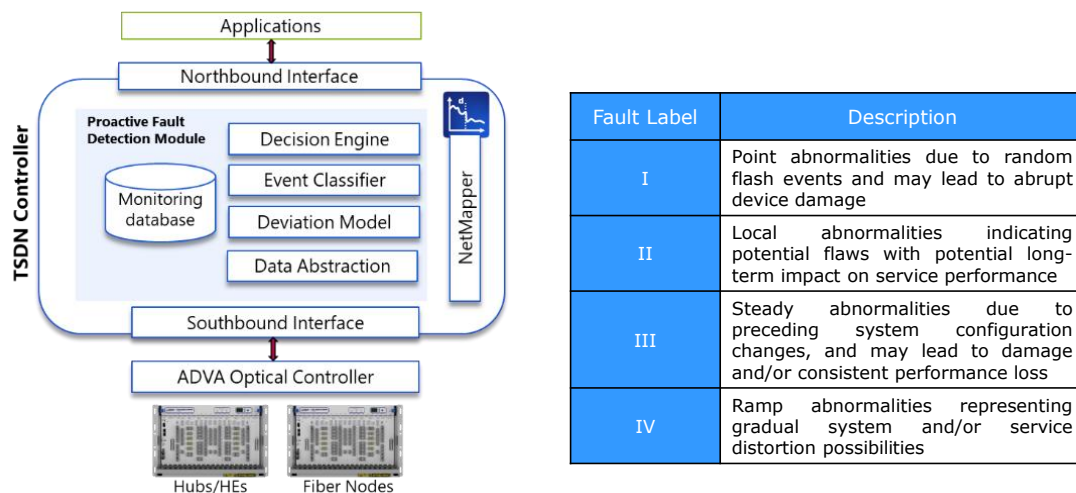


Figure 7 – Left: T-SDN Controller With PFD Architecture, Right: Definitions For Typical Fault Patterns

We developed a proof-of-concept prototype of a transport SDN (T-SDN)-integrated cognitive fault detection architecture, incorporating data analytics based on neural networks. In particular, we disclose various types of real-life network fault use cases and identify them using the proactive fault detection (PFD) framework. Figure 7 depicts the software architecture for our approach, where the PFD framework is located within the T-SDN controller. The monitoring data is collected every 15 minutes through the southbound interface (SBI) – NETCONF, which is abstracted and stored in a database. The engine performs fault detection and classification, generates fault layer information (by mapping the metadata), fault locations, and maps the machine learning outcome to an internal decision engine followed by respective application (not specified here), exposed via RESTCONF to the northbound interface (NBI). Also shown in the figure are the types of faults considered in our prototype.

The physical layer received optical power levels were retrieved, and the proposed architecture was executed on a host server. The monitored data was normalized to same scale (-1 dBm) for better visualization, and

this action did not have any impact on the performance. We evaluated two scenarios: one employing the PFD engine, where the true potential failure scenarios were cognitively predicted, and other where PFD was replaced by a fixed pre-determined fault determination threshold. Note that since we normalized our data, a single threshold was used for comparison, whereas in practice it is necessary to define and maintain several thresholds for various network configurations. Typical worst-case tolerable optical power level for coherent transceivers is -16 dBm,

Figure 8 shows various fault patterns, as described previously, for different time traces. Figure 8a depicts the performance of threshold-triggered fault detection, where label I is fully detected, labels III and IV are only partially detected, whereas label II remains undetected. In comparison, Figure 8b illustrates results from the proposed cognitive architecture, where all labels are largely detected, except for three (one) potential faults in label II (IV). It is worth mentioning that while the detection rate of faults is an important metric, the fundamental operational requirement is the ability to react to such potential failures in time, which is intrinsically enabled by the PFD framework.

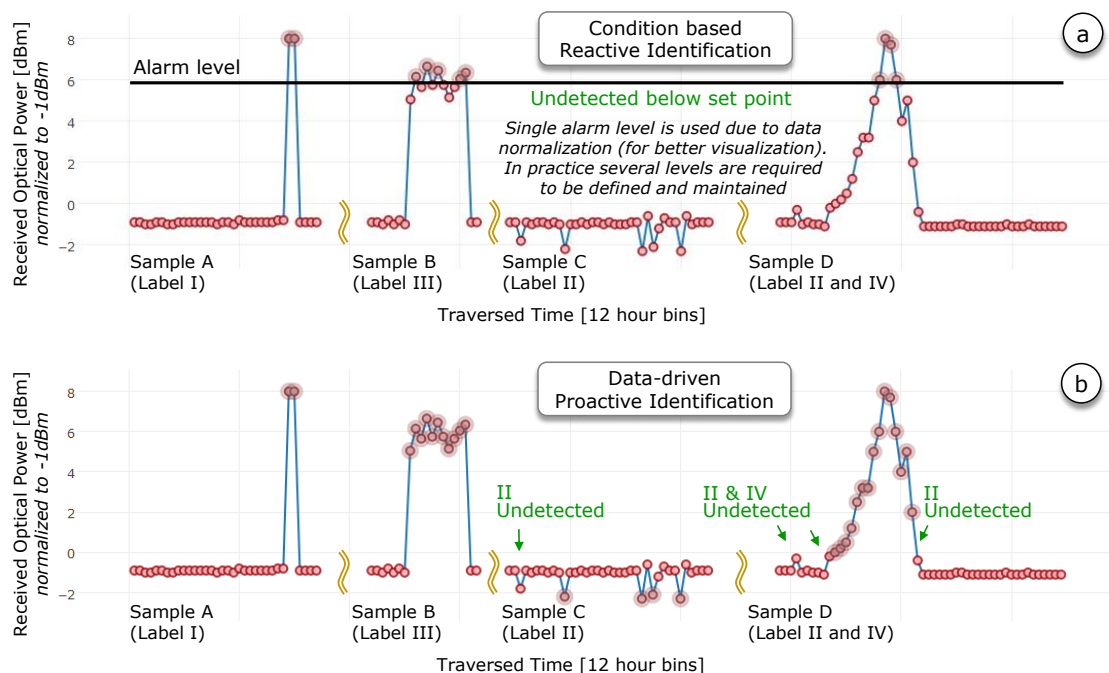


Figure 8 – (a) Condition Based Fault Detection: Rule-Driven; (b) Data Based Fault Detection (PFD Engine): ML-Driven. Highlighted Symbols (Opacity) Represent Detected Faults.

4. Conclusions

With the digitization of cable access networks, analog fiber becomes replaced by high bitrate Ethernet packet transport. This network transformation results not only in a need for cost efficient optical high-speed transport but also in a need for a new means to operate those networks in an efficient and easy way. Two emerging technologies addressing these digital fiber access network needs are:

- Ultra-high bit rate transceivers will drastically improve equipment density, power consumption and cost-efficiency of next generation cable networks. By integrating the silicon photonics, electronics

and active indium phosphide (InP) photonics on a common platform, the initiatives are poised to give cable operators what they need to meet tomorrow's demands.

- The diversity of the modern networking stack, and the dynamic nature of both the application and infrastructure layers require that conventional SDN solutions are augmented with advanced monitoring and ML-driven analytics' functionalities for end-to-end network management. The challenges remain in terms of lack of standardizations or reference designs on how these technologies may be coupled. Furthermore, issues related to data integrity, multi-vendor coexistence, and common data models need to be addressed, together with security aspects.

All in all, it can be appreciated that AI-enabled self-driven networking will significantly simplify network operations, not only leading to OPEX savings, but an enhanced customer experience. This complex and heterogeneous architecture will almost certainly be multi-domain, compelling policy and analytics to become mission critical in context of holistic network management [11].

5. Abbreviations and Definitions

5.1. Abbreviations

ADC	analog to digital converter
AI	artificial intelligence
AP	access point
API	application programming interface
BER	bit error ratio
BiCMOS	bipolar complementary metal oxide semiconductor
BMBF	Germany's Federal Ministry of Education and Research
CLO	closed loop operation
DAC	digital to analog converter
DCI	data center interconnect
DPE	digital pre-emphasis
DSP	digital signal processing
ePIC	electronic photonic integrated circuit
FEC	forward error correction
Gb/s	gigabits per second
GRPC	Google remote procedure call
HFC	hybrid fiber coax
HD	high definition
Hz	hertz
InP	indium phosphide
ISBE	International Society of Broadband Experts
MZM	Mach-Zehnder modulator
ML	machine learning
NETCONF	Network Configuration Protocol
NBI	north bound interface
OFE	optical front end
ONF	Open Networking Foundation
ONOS	Open Network Operating System
OPEX	operational expenditure

OSNR	optical signal-to-noise ratio
PON	passive optical network
PFD	proactive fault detection
RCA	root cause analysis
RESTCONF	Mapping a YANG specification to a RESTful interface
SBI	south bound interface
SD-FEC	soft decision forward error correction
SDN	software defined networking
SeDN	self-driven networking
SCTE	Society of Cable Telecommunications Engineers
SNMP	Simple Network Management Protocol
SPEED	Silicon Photonics Enabling Exascale Data Networks [project]
Syslog	system log
TAPI	transport API
TIA	transimpedance amplifier
T-SDN	transport SDN
YANG	yet another next generation

5.2. Definitions

4-QAM or QPSK	The terms are interchangeable, and refer to 2-bit modulation encoding (4-state quadrature amplitude modulation and quadrature phase shift keying)
closed loop operation	End-to-end data-driven network monitoring, control, and management
digital pre-emphasis	Refers to transmit-side pre-compensation of electro-optic component bandwidth limitations

6. References

- [1] **P2P Coherent Optics Physical Layer 1.0 Specification**, <https://specification-search.cablelabs.com/P2PCO-SP-PHYv1.0>, 2018
- [2] **Enabling 64Gbaud Coherent Optical Transceivers**, Danish Rafique et. al.; Th2A.19, ECOC 2017
- [3] **TSDN-enabled network assurance: a cognitive fault detection architecture**, Danish Rafique et al.; p. 1-3, ECOC 2017
- [4] **Analytics-Driven Fault Discovery and Diagnosis for Cognitive Root Cause Analysis**, Danish Rafique et al.; W4F.6, OFC 2018
- [5] **Optical Internetworking Forum white paper**, <http://www.oiforum.com/wp-content/uploads/OIF-Tech-Options-400G-01.0.pdf>, July 2015
- [6] **Silicon Photonics Enabling Exascale Data Networks – SPEED**, Jörg-Peter Elbers, https://www.photonikforschung.de/media/kommunikation/pdf/speed_silizium_photonik-bf_c1.pdf, 2016
- [7] **Digital preemphasis in optical communication systems: On the DAC requirements for terabit transmission applications**, Danish Rafique, et al., p. 3247-3256, IEEE JLT 2014

[8] **Modeling and Design Aspects of a Monolithically Integrated Optoelectronic Chip enabling 64Gbaud Operation**, Danish Rafique et. al.; Th2A.19, OFC 2018

[9] **Intent-Based In-flight Service Encryption in Multi-Layer Transport Networks**, Mohit Chamania et al.; Tu3L.10, OFC 2016

[10] **Field trial of Machine-Learning-assisted and SDN-based Optical Network Planning with Network-Scale Monitoring Database**, Shuangyi Yan et al.; ThPDPB4, ECOC 2017

[11] **Machine Learning for Network Automation: Overview, Architecture, and Applications [Invited Tutorial]**, Danish Rafique, et al., p. D126-D143, IEEE JOCN 2018

Low Power Wide Area Networks

Best Suited for Cable Operators

A Technical Paper prepared for SCTE•ISBE by

Mayank Bhatnagar, Accenture Strategy
Byron Scerri, Accenture Strategy

Accenture Offices
415 Mission St, Floor 35
San Francisco, California 94105
mayank.a.bhatnagar@accenture.com
byron.scerri@accenture.com
(415) 537-5000

Table of Contents

Title	Page Number
Table of Contents	20
1. Introduction	21
2. Low Power Wide Area Networks Landscape	22
2.1. LoRaWAN – Unmanaged Unlicensed Network	23
2.1.1. Business Overview	23
2.1.2. Technology Overview	23
2.2. Sigfox – Managed Unlicensed Network	24
2.2.1. Business Overview	24
2.2.2. Technology Overview	25
2.3. NB-IoT – Managed Licensed Network	26
2.3.1. Business Overview	26
2.3.2. Technology Overview	26
3. Conclusion	27
4. Abbreviations	28
5. References	28

List of Figures

Title	Page Number
Figure 1 – Global Number of Active IoT Connections (Installed Base) in Billions	21
Figure 2 – Typical Range “Connectivity”	22
Figure 3 – LoRAWAN Architecture	24
Figure 4 – Sigfox Network Topology	25
Figure 5 – NB-IoT Spectrum Deployment Options	27

1. Introduction

With the increased demand and popularity of Internet of Things (IoT) devices and more capable in-home routers, the cable industry has opportunities to enhance and deliver connected IoT services to their customers using scalable standards. This allows them to expand their services and revenue into this rapidly expanding and competitive field, along with wireless carriers who have an edge over low power wireless technologies.

The number of IoT devices will grow significantly over the coming years, up from approximately 7 billion in 2018* to 10 billion by 2020, and 22 billion by 2025. ¹ According to IDC, the worldwide IoT market will grow from \$631.4 billion in 2017 to \$1.19 trillion in 2022, representing a CAGR of 13.6%. ² In the United States, connectivity demand in the smart home market alone will nearly double by 2022, rising to 31% of total households, from 16% in 2017. ³




Figure 1 – Global Number of Active IoT Connections (Installed Base) in Billions

Source: [LPWAN Market Report 2018-2023 for IoT Analytics](#)

This significant increase in connected devices will require a comprehensive network planning strategy to successfully address a diverse set of use cases and end device requirements. This paper discusses three prominent low power wide area networks (LPWAN) technology standards that are well suited to support a wide array of IoT applications and are currently being deployed globally.

In this paper, we review prevailing LPWAN technologies in the marketplace with real-world examples and identify their strategic advantages. This will offer cable providers insights as to how next generation in-home Wi-Fi routers, and other hyper-connected devices pave the new way to gain substantial consumer traction and increase long term ARPU.

* Number does not include smart phones, tablets, laptops or fixed line phones

2. Low Power Wide Area Networks Landscape

We are seeing an unprecedented explosion in connected device form factors driven by key market and technology trends. Also, IoT brings some very specific connectivity requirements.

Connected applications with low power sensors are sustainable with very low per-device capital cost, and very low operational cost. Many IoT applications must operate on a single battery for long periods of time, and only need to transmit small data packets on an occasional basis. Such applications include smart metering, asset tracking, advanced agriculture, smart parking, and many others. However, the scaling and geolocation challenges needed to support these applications leads to complexity in managing and associating edge devices.

The current landscape of IoT deployments leverages a heterogeneous mix of wireless technologies to meet each specific use case. Since network infrastructure deployment is capital intensive, use of long-range wireless network can significantly minimize this cost.

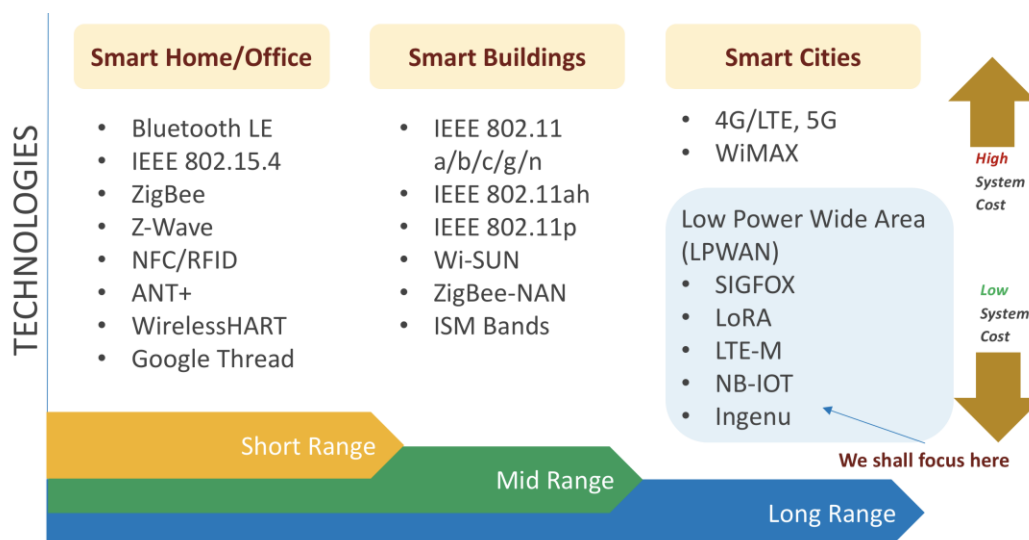


Figure 2 – Typical Range “Connectivity”

Source: Accenture Strategy Analysis, 2018

LPWAN technology provides low power connectivity along with sufficient bandwidth for IoT applications. The LPWAN market is evolving with a mix of open and proprietary technologies, and several operators and vendors are providing niche IoT applications at scale. Providers are taking different approaches to deployment technologies and associated business structures.

The most widely deployed network types are:

- 1) LoRaWAN
- 2) Sigfox
- 3) NB-IoT

These LPWANs are discussed in detail in the subsequent sections.

2.1. LoRaWAN – Unmanaged Unlicensed Network

Long range (LoRa) WAN is an LPWAN standard managed by the non-profit LoRa Alliance™. LoRa is a physical radio interface using radios from semiconductor vendor SemTech. LoRaWAN provides the overall communications link and network architecture for a local, regional or national network, including multi-tenant networks.

2.1.1. Business Overview

With rapid global adoption by cable and telecom operators, the LoRaWAN specification has been gaining significant traction as a standard LPWAN protocol. LoRaWAN supports many low bandwidth IoT use cases, such as smart waste management, smart parking, fleet tracking and management, air pollution monitoring, water metering, and shipment tracking.

Currently, LoRa networks are deployed and managed across over 80 public network operators in approximately 60 countries, with many other locations actively developing networks. As of June 2018, there were over 50 million connected LoRa-based devices, growing at >100% CAGR. ⁴

LoRa growth targets for financial year 2019 include deployments in 70 countries, involving 200,000 gateways (macro and picocell) with the capacity to support over 1 billion end nodes. In the same timeframe, over 80 million LoRa end nodes (a 60 percent increase from approximately 50 million end nodes deployed at the end of last fiscal year) are expected to be deployed. ⁵

In the United States, several Tier 1 operators and service providers have deployed LoRa networks across many U.S. metropolitan areas for their IoT platform offering. ⁶ Deployments span across multiple use cases including asset tracking, utility and water metering, facilities management, location tracking, power management and agriculture.

International adoption of LoRa has been strong with near full network coverage across the European Union either completed or underway, and there has been significant traction in Asia Pacific countries such as India, Japan and China.

2.1.2. Technology Overview

The LoRaWAN specification defines the device-to-infrastructure of LoRa physical layer parameters and the LoRaWAN protocol and provides seamless interoperability between devices. The LoRa Alliance™ drives standardization and global harmonization of the LoRaWAN protocol.

Figure 3 shows the LoRaWAN architecture used for implementing the standard with applications.

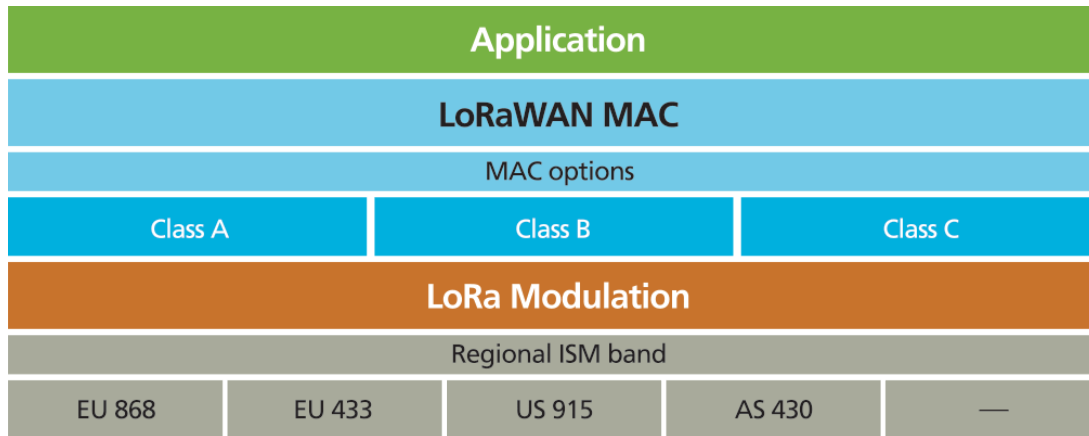


Figure 3 – LoRAWAN Architecture

Source: [2018 Semtech LoRa® Documentation](#)

LoRaWAN uses unlicensed frequencies, also referred to as ISM (industrial, scientific & medical) radio bands. ISM refers to a group of radio bands or parts of the radio spectrum that are internationally reserved for the use of radio frequency energy that is intended for scientific, medical and industrial requirements rather than for communications.

LoRa uses ultra-narrow band for upstream (half or full duplex) communication for low power, long range (up to 30 km) with data rates from 3 kbps to about 200 kbps. While the spectrum used is unlicensed and unmanaged, LoRa does have a security layer. LoRa has mechanisms for handling retransmits due to high noise environments, though this comes at the cost of battery life. It is recommended that for critical applications, gateways be physically located in places that minimize interference.

2.2. Sigfox – Managed Unlicensed Network

Sigfox is a cellular style, long range, low power, low data rate form of wireless communication that has been developed to provide wireless connectivity for devices like remote sensors, actuators and other M2M and IoT devices. Sigfox operates a model where partner companies (Sigfox Network Operators – SNOs) provide the network infrastructure and commercial offerings in each country, and end users access the network via their SNO.

2.2.1. Business Overview

The Sigfox network technology has gained significant momentum in the global market, and a number of aspects should be taken into consideration when deciding whether to integrate Sigfox into a cable operator’s IoT network strategy. Among the leading LPWAN technologies, Sigfox offers the longest range (upwards of 40 km) and broadest connectivity (a municipal area may be covered by a single base station), while maintaining the best power efficiency for some use cases.⁷

Sigfox has a proven market for its technology and services. In addition to its low power and long range capabilities, Sigfox also offers customers geolocation services ⁸ and integration with Microsoft Azure IoT Hub for real time analytics. ⁹

In the United States, Sigfox U.S.A. is a regionally owned and operated division of Sigfox. The network is deployed in 24 major metros, covering approximately 30 percent of the national population ¹⁰, and is also available in over 40 of the 50 busiest U.S. airports. ¹¹ Major cable operators are looking to Sigfox technology as a LPWAN connectivity option for IoT applications. ¹² A few applications in progress include weather tracking for farms and parking monitoring.

Global adoption for the Sigfox technology has reached over 50 countries through local partnerships that oversee network buildout and management. ¹³ Sigfox coverage reaches over 1 billion people worldwide ¹⁴, and 6 million objects are estimated to be connected to date. ¹⁵ Applications that are served range from temperature and humidity sensors in Barcelona, to the real-time tracking of sea-freight containers. ¹⁶

2.2.2. Technology Overview

Sigfox’s wireless specifications have been developed to minimize power consumption levels and enable IoT communications over extended distances. The network design follows a cellular style approach, where remote nodes are able to communicate with Internet-connected base stations that enable control and data collection.



Figure 4 – Sigfox Network Topology

Source: [Sigfox 2017 “Make things come alive in a secure way”](#)

Sigfox uses ultra narrow band (UNB) radio technology and operates in the unlicensed ISM bands. Radio messages handled by the Sigfox network are small, on the order of 12 bytes payload in the uplink and 8 bytes in the downlink. The exact frequencies can vary according to national regulations – in Europe the 868 MHz band is used; the U.S. uses 915 MHz; and Asia uses 433 MHz.

For UNB, Sigfox is using 192 kHz of the publicly available band to exchange messages over the air, and the modulation is UNB. Each message is 100 Hz wide and transferred with a data rate of 100 or 600 bits per second depending on the region. The signal is susceptible to high levels of interference, due to the FCC's time-on-air limits (400 ms), which weakens the link and limits the area the technology can cover in the U.S.

Sigfox provides a standard way of collecting data from sensors and devices with a single, standard-based set of APIs, which supports data security based on end-point encryption.

2.3. NB-IoT – Managed Licensed Network

NB-IoT or narrowband IoT is a low power wide area network radio technology standard developed by 3GPP to enable a wide range of cellular devices and services. The technology specification was established in 3GPP Release 13, in June 2016.

2.3.1. Business Overview

Narrowband-IoT has witnessed substantial adoption and deployment through major telecom operators worldwide¹⁷ and the global market for NB-IoT is expected to grow at ~50 % CAGR between 2016 and 2022.¹⁸ As most of the world does not have substantial wireless network coverage, NB-IoT is expected to be the dominant licensed network technology globally since it can leverage existing network infrastructure.

In Europe, international NB-IoT network roaming trials were conducted in 2018, and efforts are moving forward to support coverage across the continent.¹⁹ Chinese operators are deploying NB-IoT for smart city, industrial and smart home applications.²⁰ Tier 1 telecom operators in the United States have launched NB-IoT services nationwide.

2.3.2. Technology Overview

While standardized by the 3GPP and specially designed for IoT devices, NB-IoT's focus is to leverage existing GSM and LTE spectrum to provide indoor coverage at a lower cost with minimal impact to user equipment battery life.

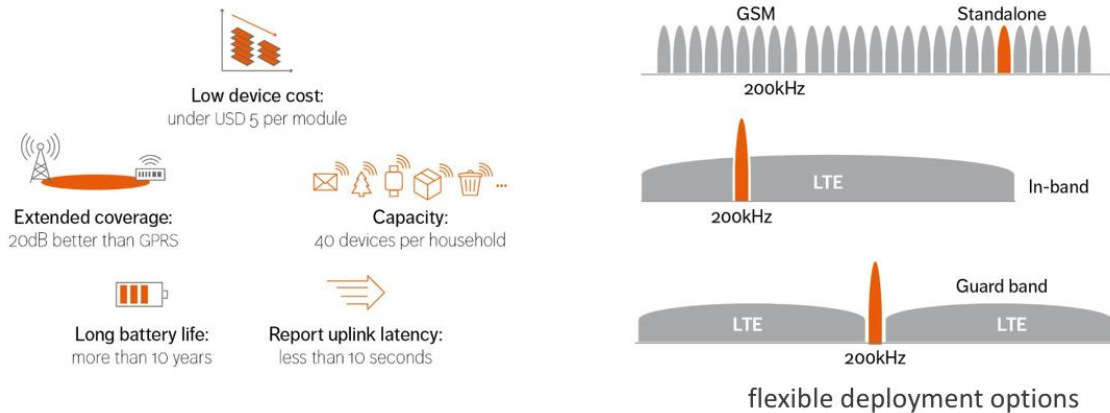


Figure 5 – NB-IoT Spectrum Deployment Options

Source: [Ericsson, 2016 “NB-IoT: A sustainable technology for connecting billions of devices”](#)

Unlike other existing LPWAN technologies, NB-IoT is designed to operate in licensed bands. To achieve spectrum efficiency, NB-IoT has been designed with a number of deployment options for GSM, WCDMA, or LTE spectrum, which are illustrated in Figure 5.

NB-IoT is directly comparable to other LPWANs with a 200 kbps narrowband architecture. Simplified modems can exist in guard bands and re-farmed cellular spectrum. Full duplex communication is used for low power and long-range data rates. While NB-IoT has internal challenges within LTE standards which can produce interference within bands, it benefits from existing carrier LTE infrastructure that helps to boost signals and avoid interference issues. The 3GPP standard also specifies other NB-IoT signal structure features like supported modulation types, frequency error tolerances, and EVM values that allow for synchronization and control channels. NB-IoT supports GSMA IoT security guidelines that are mostly adopted by wireless carriers worldwide, including network, data, and subscriber device end-points.

3. Conclusion

Cable providers must consider a flexible and scalable approach towards selecting the right IoT radio interface technology to complement their network design. Choosing an optimal IoT strategy will depend on a number of factors, including but not limited to the following:

1. Long range with sustained connectivity
2. Low power consumption for long term operation
3. Security between devices and networks
4. Managed ecosystem from physical layer to service layer
5. Managed TCO with sustained maintenance plan

Beyond technical challenges, there are compliance and interoperability issues that need to be addressed based on the target deployment strategy. As an example, if a network operator controls its target geographical region and will provide IoT connectivity within that region – the operator can choose a technology like Sigfox which is unlicensed but tightly managed within the network – and does not need interoperability with other standards.

The previous example will not work for regions where a mobile network operator has licensed wireless spectrum infrastructure and can interoperate with existing mobile standards.

As with any new technology, the market is evolving with a mixture of open and proprietary technologies and a complex mix of operators and vendors. Operators are taking a different approach both on the technology and business structure, thus a direct comparison of technology capabilities does not tell the whole story. The market is maturing with time and suggests evolution, possible conciliation and standardization over coming years.

4. Abbreviations

3GPP	3rd Generation Partnership Project
API	application programming interface
ARPU	average revenue per user
bps	bits per second
CAGR	compound annual growth rate
GSM	global system for mobile
Hz	hertz
EVM	error vector magnitude
IoT	Internet of Things
ISM	industrial, scientific & medical
LTE	long-term evolution
LPWAN	low power wide area networks
M2M	machine to machine
NB-IoT	narrowband IoT
SNO	Sigfox Network Operator
TCO	total cost of ownership
UNB	ultra narrow band
WCDMA	wideband code division multiple access

5. References

- ¹ [State of the IoT 2018: Number of IoT devices now at 7B – Market accelerating](#). IoT Analytics. August 8, 2018.
- ² [Worldwide Internet of Things Forecast, 2018–2022](#). IDC. September 25, 2018.
- ³ [US consumer internet of things market: growth, but still not explosive growth](#). SNL Kagan. November 23, 2018.
- ⁴ [LoRa Technology & LoRaWAN: Mass Adoption – Key Takeaways from All Members Meeting](#). June 6, 2018.
- ⁵ [Transcript of Semtech Corporation, Q3 2019 Earnings Call, November 28, 2018](#). SNL Kagan.
- ⁶ [Momentum is building for LoRaWAN](#). Embedded Computing Design. August 17, 2018.
- ⁷ [A comparative study of LPWAN technologies for large-scale IoT deployment](#). September 2017.
- ⁸ [The State of LPWANs in 2018](#). SNL Kagan.
- ⁹ [SigFox Collaborates with Microsoft to Connect Internet of Things Devices to Microsoft Azure IoT Hub](#). Business Wire. March 16, 2016.
- ¹⁰ [Sigfox U.S.A. President Christian Olivier to Present at IoT Summit Chicago](#). Factiva. November 28, 2018.
- ¹¹ [Sigfox U.S.A. Enters 2018 Ready to Build on Gains in Network Coverage and Partner Ecosystem](#). Press Release. March 20, 2018.
- ¹² Accenture Strategy analysis of IoT market.

-
- ¹³ [SigFox Coverage](#). January 6, 2018.
- ¹⁴ [Sigfox celebrates covering One Billion People across the globe](#). Press Release. October 25, 2018.
- ¹⁵ [Sigfox U.S.A. Enters 2018 Ready to Build on Gains in Network Coverage and Partner Ecosystem](#)
Press Release. March 20, 2018.
- ¹⁶ [Fira de Barcelona adds cutting-edge IoT technology to its Gran Via venue](#). Press Release. December 13, 2018.
[Michelin tracks inter-continental sea-freight flows with an IoT-based solution developed with Sigfox France and Argon Consulting](#). Press Release. May 24, 2018.
- ¹⁷ [Mobile IoT Deployment Map](#). GSMA. January 2019.
- ¹⁸ [Narrowband - IoT Market Research Report - Global Forecast to 2022](#). Market Research Future. January, 2019.
- ¹⁹ [GSMA announces completion of first European NB-IoT roaming trial](#). GSMA. June 4, 2018.
- ²⁰ [China is the land of NB-IoT, state support, and IoT platforms](#). Stacey on IoT. July 2, 2018.

MoCA Alternative Channel Assessment: It's More Than Just A Squiggly Line

Going Beyond the MoCA Mesh PHY Rate as the Exit Criteria

A Technical Paper prepared for SCTE•ISBE by

Maurice Manuel Garcia, Senior Principal Engineer, Comcast Cable
SCTE•ISBE Member
maurice_garcia@cable.comcast.com

Table of Contents

Title	Page Number
Table of Contents	31
1. Introduction	34
2. Caveat	34
3. ACA, ScMod and PHY Rate Quick Refresher	34
3.1. Alternative Channel Assessment	34
3.1.1. ACA Requirements	34
3.1.2. Active EVM Probe	35
3.1.3. Passive Quiet Probe	35
3.2. Per Subcarrier Modulation	35
3.3. Mesh PHY Rate	35
4. Lessons Learned	36
4.1. Point-of-Entry Filter - Location, Location, Location, it Does Matter	36
4.1.1. Customer Use Case - PoE too Far Away	36
4.1.2. The Home Configuration	37
4.2. MoCA Home Isolation Detection	38
4.2.1. PoE Filter Detection Using ACA EVM Probe	39
4.2.2. Visually Recognizing Reflections and Calculating its Signal Delay	39
4.2.3. Fast Fourier Transform Method	40
4.2.4. Calculating the Ripple Frequency from the FFT Frequency Response	42
4.3. Determining the Signal Delay and PoE Distance by Example	43
4.3.1. Pilot Carrier Signal Detection Using ACA Quiet Probe	44
4.4. Bad Splitter/Port Detection (Field Use Case)	46
4.5. 4.5. Node Max Transmit Power [2]	47
4.5.1. 4.5.1. MoCA Transmission Power Specification Requirements	47
4.5.2. Splitter Input Port to Output Port Insertion Loss	48
5. MoCA Home Readiness - Key Performance Indicators	49
5.1. MoCA Channel Suck-Out Detection	50
5.1.1. Skewness	50
5.2. MoCA Home Assessment KPI Recommendations	53
5.3. Introduction to the use of Machine Learning for POE Detection	54
5.3.1. High-Level Review of Conditional Probability	54
5.3.2. Total number of connected devices (DOCSIS and MoCA)	55
5.3.3. ScMod/VLPER/NPER: Average bit load of active subcarriers	55
6. Use Case for a Home Pre-Check for Full Duplex DOCSIS Using MoCA 2.0	55
7. Conclusions	56
8. Acknowledgments	57
9. Abbreviations	57
10. Bibliography and References	58
11. Appendix	59
11.1. Fourier Transform	59
11.2. Discrete Fourier Transform	59
11.3. Fast Fourier Transform Software Package	59
11.4. Normalizing the ACA for FFT Analysis	59
11.4.1. ACA Sample Size vs. FFT Points	59
11.4.2. Sliding Window Moving Average	60
11.4.3. Converting ACA points from dBm to Linear Scale	62
11.5. Calculating FFT Results	62
11.5.1. Complex Number Magnitude Power	63

11.5.2.	Graphing Magnitude Response	63
11.5.3.	Understanding the FFT Results	63
11.6.	Normalizing ACA Profile Data	63
11.7.	OSS SNMP Call Flow	63
11.7.1.	ACA Data Acquisition	63
11.7.2.	MeshScMod Data Acquisition	66

List of Figures

Title	Page Number
Figure 1 - Active EVM Probe	35
Figure 2 - MoCA Chan D1 - D9 EVM Probe (With/Without Poe) [1]	36
Figure 3 - ACA and ScMod Signal Impact of the PoE	38
Figure 4 - ACA EVM Response Showing 12 Null Valley Points	39
Figure 5 - Echo Detection Topology	41
Figure 6 - PoE + 75' + 3-way MFS 512 EVM points	41
Figure 7 - 2046 FFT Points - Frequency Response	42
Figure 8 - 184 FFT Points - Frequency Response of Figure 7	42
Figure 9 - DOCSIS 3.0 + MoCA CM Internal LP + HP Filters	44
Figure 10 - ACA Quiet Probe – Pilot Carrier Disable	45
Figure 11 - ACA Probe - Pilot Carrier Enable	45
Figure 12 - MoCA Mesh - Pilot Carrier Enable	46
Figure 13 - Customer Use Case Bad Port Topology	46
Figure 14 - Customer Use Case (Bad Port)	47
Figure 15 - Customer Use Case (Good Port)	47
Figure 16 - Summary of MoCA Transmit Power [3]	48
Figure 17 - 8-Way NMF Splitter Breakout + PoE	48
Figure 18 - Two-node – 2-Way Splitter Scenario	49
Figure 19 - General Example of Skewness [4]	50
Figure 20 - Histogram or Distribution of ACA Data Points	51
Figure 21 - Raw ACA Data MoCA Channel	51
Figure 22 - Raw ACA Data/Mean Left Block	52
Figure 23 - Raw ACA Data/Mean Right Block	52
Figure 24 - 3-Point Moving Average	60
Figure 25 - 10-Point Moving Average (Over-Average)	60

List of Tables

Title	Page Number
Table 1 - MoCA Mesh PHY Rate in Mbps	36
Table 2 - NMF Splitter Port to Insertion Loss + PoE	48
Table 3 - Two-node – 2-way Splitter Scenario Insertion Loss Matrix	49
Table 4 - Frequency Response – Suck-Out Detection – KPI	53
Table 5 - Frequency Response - MoCA Friendly Splitter - KPI	53
Table 6 - PoE Detection – KPI	54
Table 7 - MoCA MeshScMod – KPI	54
Table 8 - Sliding Moving Average Pseudo Code (Java)	61
Table 9 - ACA EVM quiet SNMP calls	63
Table 10 - MoCAIfAcaEntry SNMP call	64
Table 11 - Subcarrier Modulation SNMP calls	66

List of Equations

Title	Page Number
Equation 1: Calculate Delay Signal in Time	39
Equation 2: 20 <i>ns</i> Index Frequency Threshold	42
Equation 3: Echo Round Trip Delay from FFT Response	43
Equation 4: Series 6 Cable Signal Delay Correction Coefficient	43
Equation 5: Suck-out Indicator using Skewness	51
Equation 6: Conditional Probability Based on Dependent Events	54
Equation 7: Discrete Fourier Transform (DFT)	59
Equation 8: FFT Point Zero-Padding Technique	59
Equation 9: FFT Scale Calculation	60
Equation 10: Sliding Window Moving Average	61
Equation 11: dBm to Linear	62
Equation 12: Linear to dBm	62
Equation 13: Complex Number Magnitude Power Calculation	63
Equation 14: Linear to dB Calculation	63

1. Introduction

The advancement of the home wireless networks, including multi-access point adaptive Wi-Fi solutions, is improving customer Internet experience as the delivery of video, data, voice, and security service over Wi-Fi is becoming more prevalent. A number of these emerging services continues to be carried over MoCA-to-MoCA only video clients. With millions of MoCA 2.0 capable devices in customers' homes, cable operators must continue to leverage MoCA technology to deliver services while still updating the in-home networking solutions and DOCSIS to provide higher Internet speeds.

Migrating to MoCA 2.0 not only offers the added throughput in the MoCA home network, but it also provides additional network robustness. To reduce the cost of truck rolls, MoCA 2.0 delivers the ability to perform RF analysis via alternative channel assessment (ACA) and subcarrier modulation bit loading (ScMod) reporting. Additional reporting metrics can be leveraged in determining the MoCA network health and be used in point-of-entry (PoE) filter detection. With MoCA 2.0 network RF analysis, operators can leverage these measurement cues to extend existing PNM monitoring and score a customer's home for advanced Internet services such as DOCSIS 3.1 and full duplex (FDX) DOCSIS.

This paper provides the lessons learned with MoCA 1.1 and MoCA 2.0 coexistence challenges, details the techniques for PoE detection, and provides recommendations on leveraging MoCA 2.0 RF analysis for upgrading to DOCSIS 3.1 FDX home assessment.

2. Caveat

All ACA signal plots are not normalized for actual power unless otherwise specified. The normalization of the ACA probe data requires additional steps outlined in Section 11.6.

3. ACA, ScMod and PHY Rate Quick Refresher

This section briefly describes the three principal metrics used in this paper: ACA, ScMod, and mesh physical layer rate. Details on how to perform and collect each operation can be found in Section 11.7.

3.1. Alternative Channel Assessment

There are two types of ACA operations: alternate channel quiet line assessment and alternate channel EVM probe assessment. The purpose of the ACA is to perform a measured signal test, similar to a vector network analyzer. Although a vector network analyzer can measure both amplitude and phase, the ACA only measures amplitude power at the center frequency of the orthogonal frequency division multiplexing (OFDM) subcarriers.

In executing an ACA, one or more nodes may perform the alternate channel quiet line assessment. During the ACA, the network coordinator (NC) quiets the MoCA network and perform the ACA operation.

3.1.1. ACA Requirements

The NC must be a MoCA 2.0 device.

- The NC is not required to participate in the ACA probe, only to coordinate.

- At least two nodes are required to complete the alternate channel EVM probe assessment operation.
- At least two nodes must support MoCA 2.0

3.1.2. Active EVM Probe

In the following example, Node:0 will broadcast a high-power error vector magnitude (EVM) pilot probe into the coaxial network. During this time, Node:1 is actively measuring each subcarrier's power.




Figure 1 - Active EVM Probe

3.1.3. Passive Quiet Probe

Unlike the active EVM probe, the quiet probe measures the energy of each subcarrier index. It does not require a locked pilot carrier. The process measures any noise sources found at the OFDM individual subcarrier indexes.

3.2. Per Subcarrier Modulation

The per-subcarrier modulation describes the OFDM indexing for 512 subcarriers' bit load allocation, where 32 subcarriers are used for the guard band that is unavailable for data transport. The total number of usable subcarriers is 480.

There are three types of bit-load profiles: very low packet error rate (VLPER), nominal packet error rate (NPER) and standard ScMod. The VLPER and NPER are for the MoCA 2.0 100 MHz OFDM per-subcarrier bit load allocation. The ScMod is for the MoCA 1.1 50 MHz OFDM per-subcarrier bit load allocation.

3.3. Mesh PHY Rate

Home network MoCA diagnostics primarily leverage the full mesh rate (FMR) table, also known as the mesh PHY rate, which represents throughput capacity between all MoCA mesh devices.

Table 1 - MoCA Mesh PHY Rate in Mbps

Node	0	1	2
0	-	201	208
1	149	-	270
2	213	205	-

4. Lessons Learned

The following sections summarize the lessons learned regarding MoCA 1.1 and MoCA 2.0 in mix mode home networks. At the time of this document, a number of service providers are completing the upgrade of their MoCA 2.0 capable devices from MoCA 1.1 to MoCA 2.0 firmware.

4.1. Point-of-Entry Filter - Location, Location, Location, it Does Matter

The PoE filter is more than just a low pass filter (LPF) preventing the MoCA signal from entering the HFC network; it is also a way to reclaim the MoCA energy back into the home coaxial network and improve the signal-to-noise ratio (SNR). Due to the rejection band of the LPF filter having less than 2 dB return loss, essentially the PoE filter acts as a reflector at the MoCA channel frequencies as depicted in Figure 2.




Figure 2 - MoCA Chan D1 - D9 EVM Probe (With/Without Poe) [1]

The reflection of the MoCA signal also has a negative implication; if the reflection is too far away from the source of the MoCA signal and has enough reflected channel power, it can cause ripples on the MoCA channel and, in some cases, cancel out many subcarriers in a predictable way. Fortunately, MoCA is designed to dynamically bit load around these issues as depicted in Figure 3, but in some cases, the MoCA client RF front end can have some difficulties, depending on the interoperation between different MoCA chipsets.

4.1.1. Customer Use Case - PoE too Far Away

This section illustrates a scenario where the effects of the PoE filter are different depending on the cable configuration. Cable operator markets can have different standard operating procedures (SOP) when it

comes to where a PoE is to be installed, and in some cases, one can have situations where it is not installed by mistake. Also, the location and distance of the PoE filter relative to the MoCA client can inadvertently cause service issues, including service interruption.

The following scenario has commonly been reported in the field; however, there was not enough information to find the root cause. The anecdotal effect was video tiling which is characteristic of packet loss. The cause of this problem was elusive because the PHY rates were within the operational specification in a mix mode environment of 220 Mbps.

MoCA mix mode is defined as a MoCA network containing both MoCA 1.1 and MoCA 2.0 devices operating with different physical layer throughput rates (PHY rates). MoCA 1.1 has an average PHY rate of 220 Mbps to 250 Mbps and MoCA 2.0 has a PHY rate > 600 Mbps. Lab testing revealed that the issue was related to a specific device, but this discovery was not sufficient to determine the root cause of the problem.

Ultimately, it required a customer home visit, where it was determined that the location of the PoE filter relative the MoCA client was causing the video impairment. The MoCA client had an issue with its channel estimation related to how aggressively it managed deep nulls in the MoCA channel, causing signal canceling of some subcarriers and reducing discrete bit loading in MoCA mix mode.

Unfortunately, this issue could not be remedied in software and required an in-home cabling change. One or more of the following changes can help to address this type of issue:

- Reduce the distance between the PoE and the farthest MoCA client.
- Add attenuation at the PoE to increase the return loss, without compromising video or DOCSIS common path signal power dynamic range requirements.
- Change to a cable home-run configuration where the PoE is at the input of the root splitter.
- Upgrade all devices to a MoCA 2.0 network.

4.1.2. The Home Configuration

The EVM probe from Node1 | Node3 demonstrates a ~5 dB_{Pk-Pk} ripple and a second delay signal of roughly 60 ns. This has a minimum impact on the overall throughput as depicted in the bits-per-symbol Node1_{Tx} | Node3_{Rx}

The EVM probe from Node2 | Node3 demonstrates a more impressive ~15 dB_{Pk-Pk} ripple and a second delay signal of 180 ns. This has an impact on the overall throughput as depicted in the bits-per-symbol Node2_{Tx} | Node3_{Rx}.

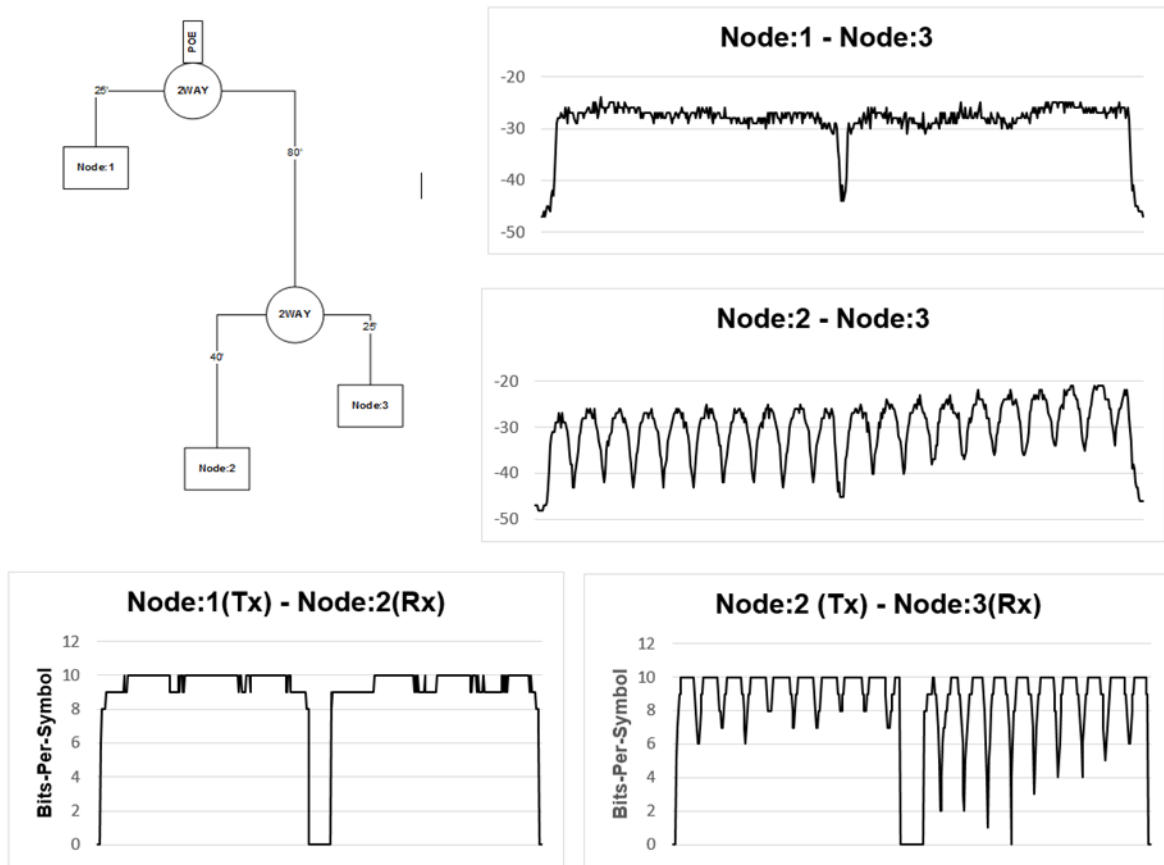


Figure 3 - ACA and ScMod Signal Impact of the PoE

4.2. MoCA Home Isolation Detection

MoCA home isolation is critical for preventing the bridging of MoCA networks between single-family homes and multiple dwelling units (MDUs).

Home isolation can be achieved by the following:

- Enabling MoCA security with unique keys per home, similar to Wi-Fi.
- Installing a MoCA PoE filter, preferably inside the home.
- Installing a unity gain amplifier (UGA), which by its design, prevents MoCA signals from returning into the HFC network. To improve the output-to-output port MoCA signal path quality, MoCA friendly UGAs should be used.
- MoCA friendly splitters that are designed with an embedded PoE filter.

4.2.1. PoE Filter Detection Using ACA EVM Probe

Determining if a PoE filter is installed is one of the methods to isolate the MoCA signal inside the home. An operator may want to check for home isolation at the following events:

- Scheduled service visit.
- Perform an assessment during a remote troubleshooting session.
- Proactive monitoring for PoE filter detection to verify that the customer continues to maintain in-home MoCA isolation.

This section describes two methods of detecting an installed PoE filter at the ground block and a PoE filter that is not connected at the root splitter. These methods can also be applied to DOCSIS 3.1 PNM features:

- Spectrum analyzer
- DS RxMER
- Channel estimation

If one has visibility into the sampling rate and the number of ripples within a given bandwidth, one can estimate the reflection point relative to the measured position.

4.2.2. Visually Recognizing Reflections and Calculating its Signal Delay

By way of a visual inspection of the ACA EVM probe using a line plot to represent the data, in Figure 4 there are 12 nulls or valleys. By using the following equations, one can approximate the signal delay and distance of the PoE filter from the input of the root splitter.




Figure 4 - ACA EVM Response Showing 12 Null | Valley Points

Equation 1: Calculate Delay Signal in Time

$$\left[\frac{\text{OFDM Occupied Bandwidth}}{\text{Number of (Nulls | Peak | Valley)}} \right]^{-1} = \text{Delay Signal in time}$$

$$\left[\frac{100 \text{ MHz}}{12 \text{ (Null)}} \right]^{-1} = 120 \text{ ns} = 120 \text{ feet}$$

$$\text{RG6 Cable Speed of Light} \sim \frac{1 \text{ ns}}{\text{feet}}$$

$$\frac{120 \text{ ns}}{2} = \text{PoE Distance} \sim \mathbf{60 \text{ feet}}$$
 assuming 120 ns of round trip delay

4.2.3. Fast Fourier Transform Method

Determining the ripple in the ACA EVM probe using a fast Fourier transform (FFT) is an easy process. Detailed steps can be found in Section 11.1. The following example is an FFT of an ACA EVM probe where the PoE filter is 75 feet from the input of the splitter. An FFT is used to determine the ripple frequency across the one MoCA 100 MHz channel. The minimum delay that can be analyzed is 20 ns.¹

Figure 5 and Figure 6 depict the EVM response of a PoE 75 feet from the input of a 3-way MoCA friendly splitter (MFS). No moving average was applied, nor additional interpolation to increase the sample size.

¹ To analyze ripples less than 20 ns, it requires multiple ACAs of adjacent MoCA channels to observe an elongated ripple response.

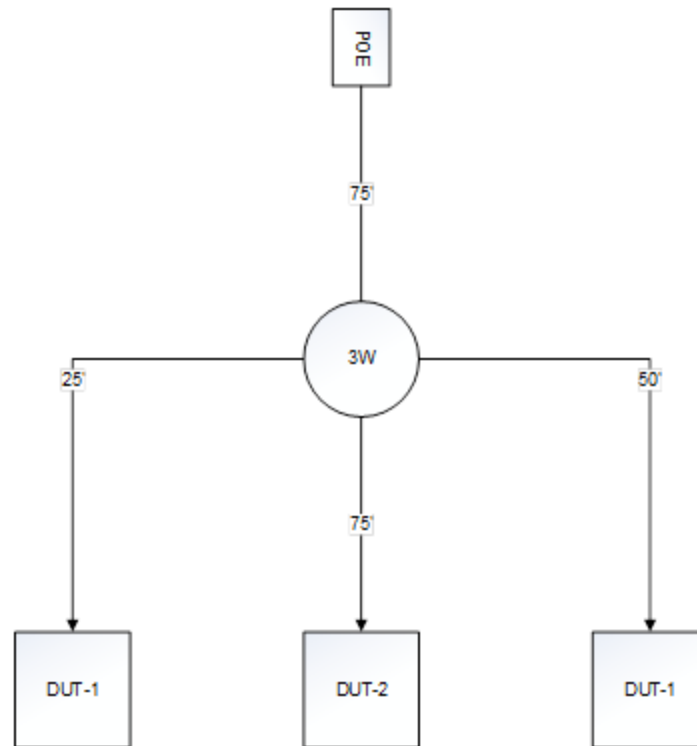


Figure 5 - Echo Detection Topology




Figure 6 - PoE + 75' + 3-way MFS 512 EVM points

Figure 7 depicts the EVM FFT frequency response. The actual number of FFT points is 4096, with a sample size of 512. Figure 7 is the first 2048 points or $\frac{4096}{2}$, this is due to the mirroring effect of the FFT.

Using Equation 9: *FFT Scale Calculation*, the $FFT_{Scale} = .125$, Figure 7 shows two major peaks in the lower frequencies.




Figure 7 - 2046 FFT Points - Frequency Response




Figure 8 - 184 FFT Points - Frequency Response of Figure 7

Figure 8 depicts the first 184 FFT points or (1 Hz to 184 Hz) found in Figure 7. The shaded region represents the guard band low magnitude points of the EVM response in Figure 6. The shaded region or threshold region is to be excluded when searching for the ripple frequency in the EVM response.

If only one ACA EVM probe is used any delay of 20 ns or less would not render a ripple on the OFDM carriers, and cannot be detected. In Figure 8, the threshold region is from 1 Hz to 16 Hz.

The calculation to determine the 20 ns threshold region is:

Equation 2: 20 ns Index Frequency Threshold

$$Index_{FrequencyThreshold(FT)} = \frac{20 \cdot 10^{-9}}{FFT_{Scale}([T_s]^{-1})}$$

$$T_s = SampleRate = NumberOfMoCAChannels (100 \cdot 10^6)$$

4.2.4. Calculating the Ripple Frequency from the FFT Frequency Response

Figure 6 depicts a ripple on the MoCA OFDM carriers. A micro-reflection or standing wave can create the ripple due to the PoE filter reflecting the MoCA signal to the receiver of the MoCA device. The

reflection reduces the attenuation at specific frequencies based on the linear cable distance from the PoE filter to the MoCA device. Determining the ripple frequency, one can calculate the echo or signal delay to determine the distance from the PoE filter or reflection point to the MoCA device.

In this step, one must exclude the threshold region and find the highest magnitude peak (HMP) as depicted in Figure 8, then use the following equation to obtain the ripple frequency.

Equation 3: Echo Round Trip Delay from FFT Response

$$SignalDelay_{RoundTrip} = [FFT_{[HMP]} \cdot FFT_{Scale}] \cdot [T_s]^{-1}$$

The next step is to calculate the actual propagation delay assuming Series 6 cable (commonly called RG-6 cable) is used in a customer's home. The nominal velocity of propagation (NVP) differs depending on the type of coaxial cable. The percentage coefficient needs to be multiplied to normalize the signal delay results.

Equation 4: Series 6 Cable Signal Delay Correction Coefficient

$$SpeedOfLight = 1.016703362E^{-9}$$

$$RG6_{NVP} = 84\% = .84$$

$$SignalDelay_{RG6(NVPCorrection)} = [SpeedOfLight \cdot RG6_{NVP}] = .854030824$$

$$SignalDelay_{RoundTrip(RT)} = SignalDelay_{RG6(NVPCorrection)} \cdot SignalDelay$$

4.3. Determining the Signal Delay and PoE Distance by Example

$$FFT_{Scale} = \frac{512}{4096} = .125$$

$$Index_{FrequencyThreshold(FT)} = \frac{20 \cdot 10^{-9}}{.125([100 \cdot 10^6]^{-1})} = 16$$

$$FFT_{[HMP]} = 149$$

Check: $Index_{FrequencyThreshold(FT)} < FFT_{[HMP]}$

$$\text{Signal Delay}_{RT} = [149 \cdot .125] \cdot ([100 \cdot 10^6]^{-1}) = 186.25 \cdot 10^{-9}$$

$$\text{SignalDelay}_{(RT)\text{Corrected}} = .854030824 \cdot 186.25 \cdot 10^{-9} = 159.06 \cdot 10^{-9}$$

$$\text{POE}_{\text{ReflectionPoint}} = (159.06 \cdot 10^{-9}) \cdot .5 = 79.5 \cdot 10^{-9} \approx 79.5 \text{ feet}$$

4.3.1. Pilot Carrier Signal Detection Using ACA Quiet Probe

This section describes another approach to remotely determine if a PoE filter or 1 GHz UGA is installed using the ACA quiet probe. The goal of the process is to determine if the MoCA network is isolated to the customer's residence.

In developing the technique to use a pilot carrier to detect whether a home is isolated, two principal design constraints limited the pilot carrier to be placed in a specific region of the spectrum to avoid interfering with the MoCA OFDM channel.

The following scenarios required the pilot carrier frequency to be placed at 1.125 GHz, but a pilot carrier frequency of 1.175 GHz can also be safely used.

- Depending on the remote PHY (R-PHY) CMTS, it may only have a max transmit frequency of 1.2 GHz.
- Due to HFC upgrades in the field, HFC taps may only have a 1.2 GHz passband.
- The most common HFC tap/splitter passband is 1 GHz, and the pilot should have a minimum loss in the rejection band.
- Inserting the pilot carrier between the MoCA OFDM channels or near the guard bands can cause MoCA network formation instability if a PoE filter is not installed.




Figure 9 - DOCSIS 3.0 + MoCA CM Internal LP + HP Filters

In a coaxial network where there is an absence of a neighboring customer's MoCA signal as depicted in Figure 10, the ACA quiet probe should show a flat response. Figure 10 is the expected response if no pilot carrier is detected inside the MoCA channel.




Figure 10 - ACA Quiet Probe – Pilot Carrier Disable

Figure 11 depicts the presence of a pilot carrier at 1.125 GHz. A high amplitude algorithm that is later discussed in Section 5.1.1 can quickly detect the pilot and indicate that the customer's home is not isolated.²

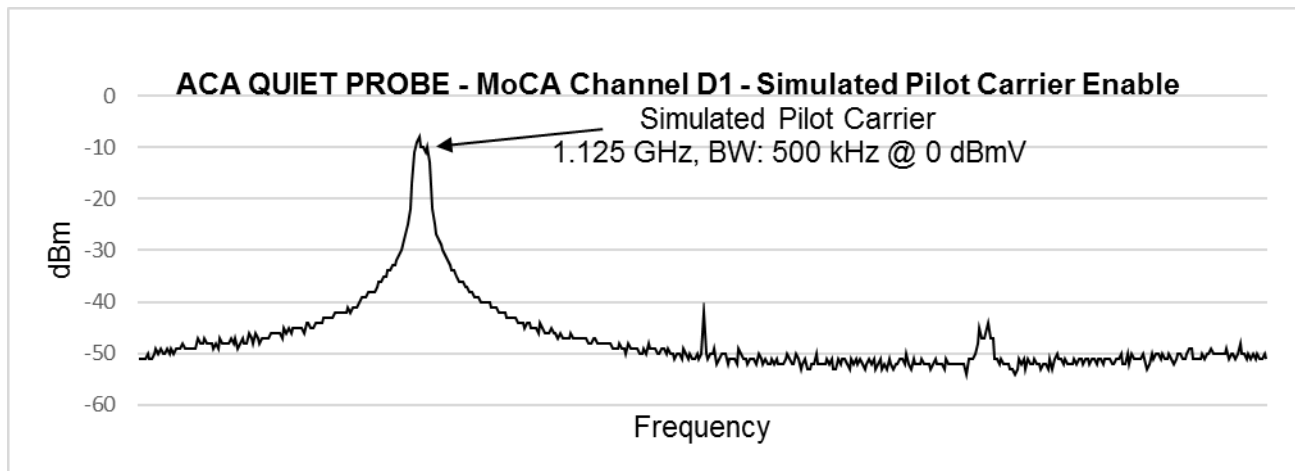


Figure 11 - ACA Probe - Pilot Carrier Enable

Figure 12 indicates no impact to the lower 50 MHz OFDM MoCA channel. This is because the MoCA transmit power in normal operation has an upper limit of $\sim +49$ dBmV to $\sim +57$ dBmV or 0 dBm to +8 dBm³. The strong MoCA signal overpowers the pilot carrier. The estimated LDPC code gain SNR requirement would be ~ 32 dB SNR at 1024-QAM and ~ 25 dB SNR at 256-QAM.

² The ACA quiet probe analysis assumes that a 1.2 GHz UGA is not used. Further ML analysis is provided in Section 5.3.

³ Figure 16 summary of MoCA Transmit Power [3]

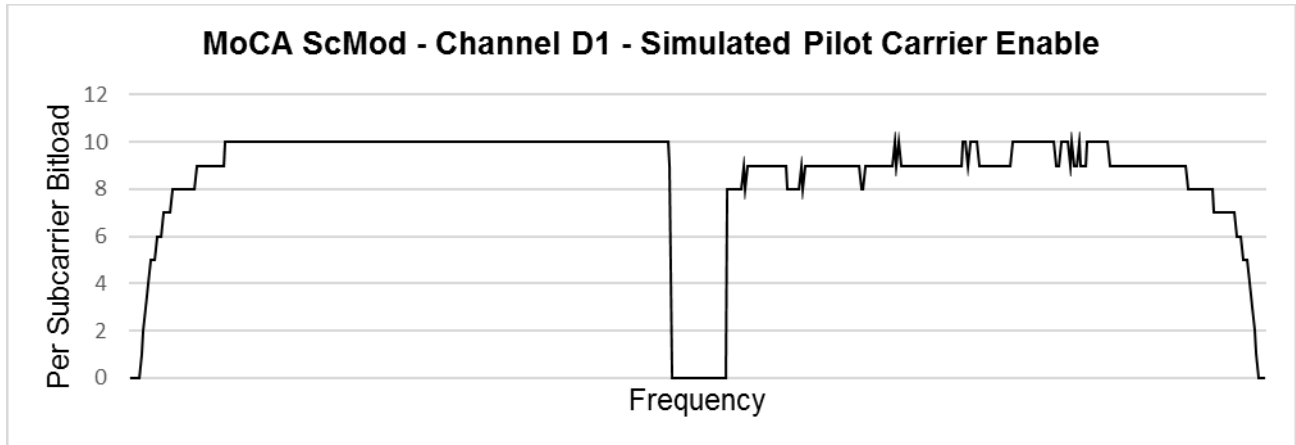


Figure 12 - MoCA Mesh - Pilot Carrier Enable

4.4. Bad Splitter/Port Detection (Field Use Case)

The following is a field use case where the customer reported video stability issues while the device was reporting a slightly lower than average but passable PHY rate of 210 Mbps. Unfortunately, at the time of this writing, the ScMod reporting was not available, and the option to report the ACA was only available from Node:0.

The customer’s splitter was an authorized non-MoCA friendly (NMFS) splitter, initially installed by the technician.

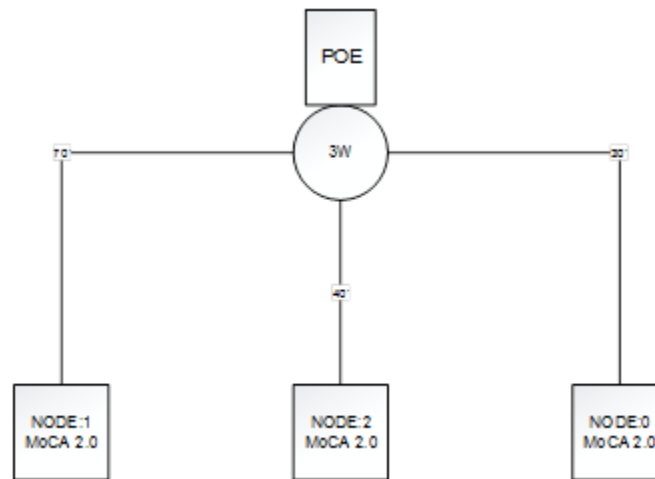


Figure 13 - Customer Use Case Bad Port Topology

Figure 13 depicts the home network topology. Note that although there is a PoE at the root of the splitter, no ripple was found in both Figure 14 and Figure 15. The absence of the ripple indicates no delayed signal present in this topology.




Figure 14 - Customer Use Case (Bad Port)

Figure 14 depicts a physical impairment at the splitter port associated to Node:1. There does not appear to be an impairment between Node:2 and Node:0. Replacing the splitter resolved the issue.




Figure 15 - Customer Use Case (Good Port)

4.5. 4.5. Node Max Transmit Power [2]

This section details the MoCA specification related to the maximum transmit power, insertion loss of the in-home standard splitter networks, and the effects of the PoE filter on the MoCA signal receive power from the perspective the MoCA device. Understanding the physics of the splitter network one may determine the presence of a PoE filter connected to the input port (root) of the splitter network.

A PoE filter connected to the input port of a splitter that is configured as a home-run coaxial network will not cause a ripple on the MoCA OFDM carrier when exercising an ACA EVM. A PoE filter will improve the Rx power and negate the effects of the splitter insertion loss because the PoE filter is returning the signal power into the network. Knowing the number of connected cable devices, one can speculate the number of splitter ports and estimate the expected losses while checking against the measured MoCA Rx power from the ACA EVM probe. This technique is not a certainty, but is another tool in determining the likelihood of an installed PoE filter that is discussed in Section 5.3.

4.5.1. 4.5.1. MoCA Transmission Power Specification Requirements

A transmitting node must have a maximum total output power between -3 dBm and +5 dBm at every supported MoCA channel frequency within the frequency band of 100 MHz around the center frequency of the transmitted signal when transmitting in MoCA 2.0 PHY. [2]

Version	Number of Channels Transmitted	Channel Bandwidth (MHz)	Maximum Output Power per Channel	Maximum Output Power Total for all Transmitted Channels
1.0/1.1	1	50	-1 dBm to +7 dBm	-1 dBm to +7 dBm
2.0	1	100	-3 dBm to +5 dBm	0 dBm to +8 dBm
	2	100	-4.5 dBm to +3.5 dBm	+0.3 dBm to +8.3 dBm
2.5	3	100	-5.3 dBm to +2.7 dBm	+0.7 dBm to +8.7 dBm
	4	100	-6 dBm to +2 dBm	+1 dBm to +9 dBm
	5	100	-1 dBm to +7 dBm	-1 dBm to +7 dBm

Figure 16 - Summary of MoCA Transmit Power [3]

4.5.2. Splitter Input Port to Output Port Insertion Loss

The typical insertion loss of a 2-way splitter @ 1150 MHz is ~7 dB.

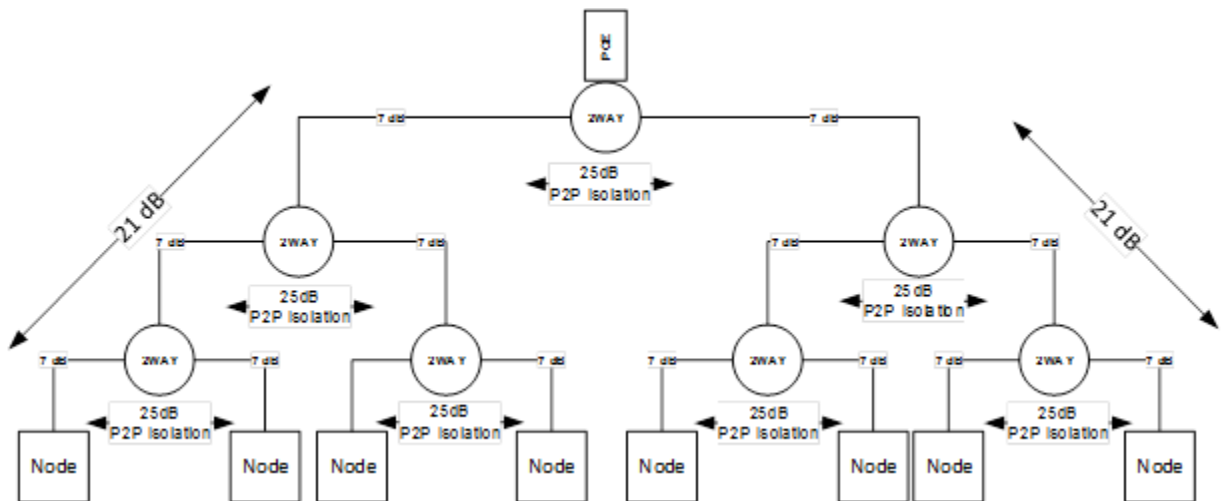


Figure 17 - 8-Way NMF Splitter Breakout + PoE

Table 2 - NMF Splitter Port to Insertion Loss + PoE

Number of Ports	Insertion Loss @ 1150 MHz + PoE
2	~14 dB
4	~28 dB
8	~42 dB

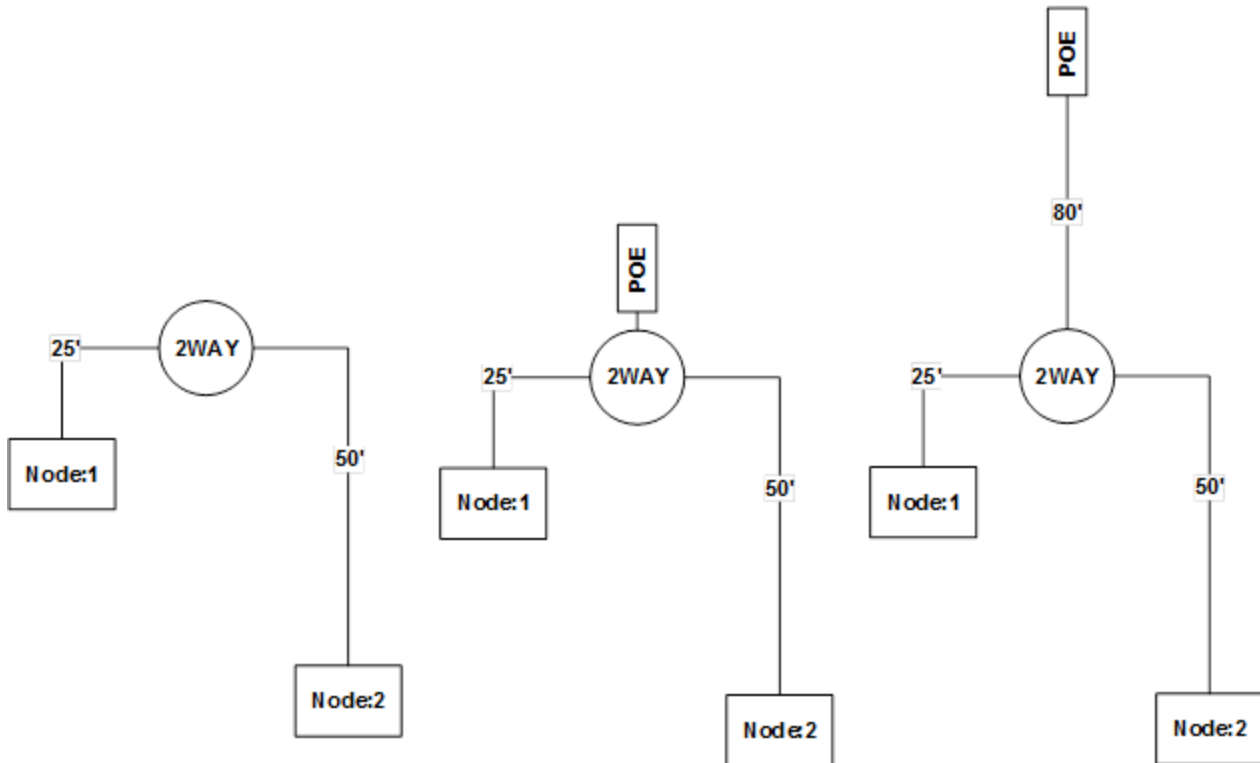


Figure 18 - Two-node – 2-Way Splitter Scenario

Using the ACA EVM probe data and known splitter characteristic, the data in Table 3 shows the difference between the absence of a PoE and the presence of a PoE connected to the input of a 2-way splitter. Understanding the relationship between the number of nodes or splitter ports and the received power, one can conclude the likelihood that a PoE filter is connected to a splitter or not.

Table 3 - Two-node – 2-way Splitter Scenario Insertion Loss Matrix

<i>Series 6 Cable Loss = $\frac{6.5 \text{ dB}}{100 \text{ Feet}}$ @ 1150 MHz</i>			
	2-WAY	PoE+2-WAY	PoE+80+2-WAY
Insertion Loss	25 dB	14 dB	24.4 dB
Measure Rx Power via EVM Probe	-21 dBm	-11 dBm	-18 dBm
Estimated Insertion Loss	-24 dB ... -18 dB	-15 dB ... -7 dB	-25.4 dB ... -17.4 dB

5. MoCA Home Readiness - Key Performance Indicators

The section provides some key performance indicators (KPI) that can be used as an exit criterion scoring for technicians to complete their work order. It can also be used as a tool for proactive maintenance of the MoCA health in a customer’s home.

5.1. MoCA Channel Suck-Out Detection ⁴

This section describes a technique to detect suck-outs within the MoCA channel. In general, a suck-out is a discontinuity of a uniform frequency response.

5.1.1. Skewness

An indication of a suck-out can be determined by calculating the skewness using the magnitude points of the ACA data. Skewness is a measure of symmetry, or more precisely, the lack of symmetry. A normal distribution is symmetric if the data points look the same to the left and right of the center point. Skewness is calculated by using the *Pearson's Coefficient of Skewness* algorithm, and it can be found in most general purpose statistical software programs.

A skewness of less than -1 is an indication of the presence of a significant suck-out, where most of the data points are pushed to the right. An example of data that is skewed to the right is provided in Figure 20. A skewness of greater than +1 is an indication of a significant noise source, where most of the data points are pushed to the left.




Figure 19 - General Example of Skewness [4]

When evaluating the ACA EVM data, it is required to remove the guard band subcarriers to determine the standard deviation (SD), mean and skewness. A normalization method is required to normalize the data so as not to misrepresent the actual meaningful information to determine suck-outs.

As shown in Figure 21, by leaving the guard band magnitude values in the distribution model, it skews the population of data to the right as indicated with a skewness of -3.63. The model shows data points with little population frequency, starting with points lower than -29 dBm.

⁴ KPI




Figure 20 - Histogram or Distribution of ACA Data Points

$$x \leq -1 \text{ or } 1 \geq x$$

Equation 5: Suck-out Indicator using Skewness

In Figure 22 and Figure 23 the skewness value indicates no suck-out.

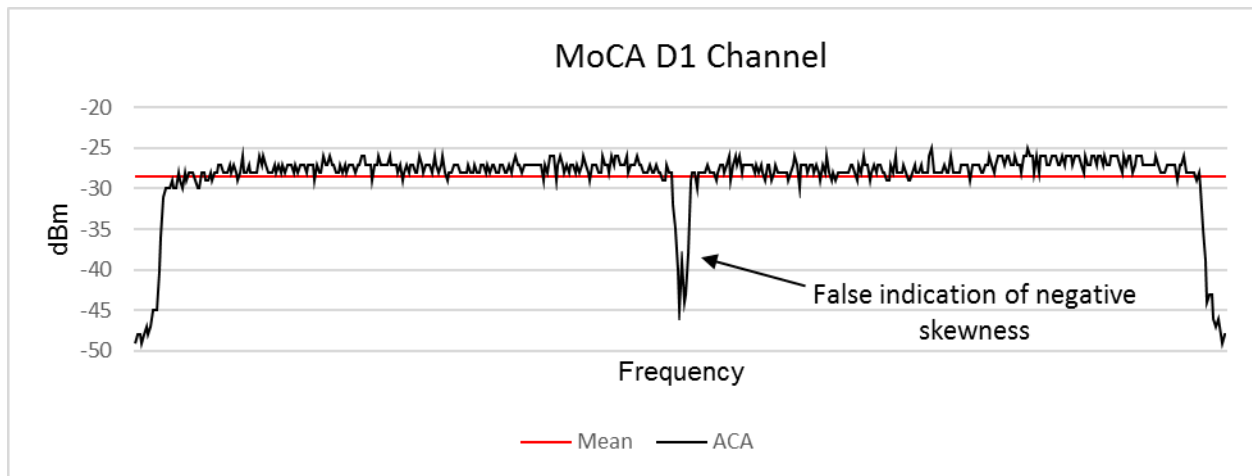


Figure 21 - Raw ACA Data MoCA Channel

Mean: -28.48 SD: 4.30 Skewness: -3.63




Figure 22 - Raw ACA Data/Mean Left Block

Mean: -27.45 SD: 0.85 Skewness: -0.50

As with Figure 22 and Figure 23, when calculating the SD, and observing the plots, the points are no higher than ~ 3 dB. The SD for both plots two times the SD is less than ± 3 dB and greater than 95% of all the points. Both figures indicate a relatively flat response for each OFDM block but do not indicate a zero slope tilt. Although this is the case, the figures indicate a desirable result.




Figure 23 - Raw ACA Data/Mean Right Block

Mean: -27.25 SD: 0.93 Skewness: -0.11

5.2. MoCA Home Assessment KPI Recommendations

This section is a summary of the metrics discussed in this paper. The following should be used as a guideline to determine the MoCA health. The variances are mostly determined by the known splitters and coaxial cable characteristics and should be a part overall conditional decision.

Table 4 - Frequency Response – Suck-Out Detection – KPI

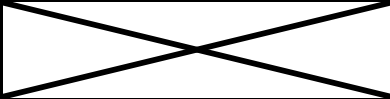
	Frequency Response (Good) + MA		Frequency Response + MA Suck out Detection	
	OFDM Left Block	OFDM Right Block	OFDM Left Block	OFDM Right Block
Standard Deviation	$x < 0.8 \text{ dBm}$		$x > 1.5 \text{ dBm}$	
Skewness	$x \geq -1 \text{ and } x \leq 1$		$x \leq -1$	
Rx Channel Power	$x > \{x\} \text{ dBm}$			

Table 5 - Frequency Response - MoCA Friendly Splitter - KPI

	Frequency Response (Good) + MA	Frequency Response (Poor) + MA	Frequency Response (Good) + MA
	Standard Splitter		MoCA Friendly Splitter
Tilt	NA	Negative Tilt	Positive Tilt
Standard Deviation	$x < 0.8 \text{ dBm}$		$x < 0.8 \text{ dBm}$
Slope Delta	$\pm 3 \text{ dB}$	$x < -9 \text{ dB or } -3 \text{ dB} > x$	$3 \text{ dB} > x < 9 \text{ dB}$

Table 6 - PoE Detection – KPI

	PoE Detection (Strong Reflection) + MA		PoE Detection (Weak Reflection) + MA	
	OFDM Left Block	OFDM Right Block	OFDM Left Block	OFDM Right Block
Standard Deviation	$x > 2.0 \text{ dBm}$		$1.5 \text{ dBm} > x < 2.0 \text{ dBm}$	

Table 7 - MoCA MeshScMod – KPI

	MoCA MeshScMod (Bit Load) Excellent		MoCA MeshScMod (Bit Load) Minimum (Good)	
	OFDM Left Block	OFDM Right Block	OFDM Left Block	OFDM Right Block
Mean	9 bpsym		7-8 bpsym	
Standard Deviation	$x < 1$		$x > 1$	
Skewness			$x \geq -1 \text{ or } 1 \leq x$	
Calculated PHY Rate	$x > 848.75 \text{ Mbps}$		$x > 656.25 \text{ Mbps}$	
Reported PHY Rate	$x > 650 \text{ Mbps}$		$x > 600 \text{ Mbps}$	

5.3. Introduction to the use of Machine Learning for POE Detection

Machine learning (ML) has continued to impress technologists in analyzing big-data with the use of neural networks. This section is a high-level review of the ML algorithm conditional probability by quantifying dependent events of MoCA/DOCSIS RF metric cues.

5.3.1. High-Level Review of Conditional Probability

Conditional probability is the likelihood of an event or outcome occurring based on the occurrence of a previous event or outcome. Conditional probability is calculated by multiplying the probability, a value from .01 to .99, of the preceding event by the updated probability of the succeeding, or conditional, event. [5]

When the probability of an event A is given, and we have to find the probability of the other event B based on that event A or $P(A \cap B)$, then the probability obtained is called conditional probability of B given A or $P(A | B)$.

$$P(A|B) = \frac{P(A \cap B)}{P(B)} = \frac{\text{Probability of the occurrence of A and B}}{\text{Probability of B}}$$

Equation 6: Conditional Probability Based on Dependent Events

Understanding the physics of the splitter and makeup of the coaxial network, one can implement an ML algorithm to detect a PoE is present at the root of the splitter or a UGA.

The following are metric cues and the likelihood of occurrences that can be weighted to affect the total probability of occurrence of a PoE and or UGA.

5.3.2. Total number of connected devices (DOCSIS and MoCA)

Typically, a UGA is present when the number of linear video or DOCSIS enabled devices is greater than four. Knowing the number of connected devices, one can calculate the minimum insertion loss and apply the maximum transmit power during an EVM probe. Review Section 4.5.2 for further details.

Legacy deployments of DOCSIS 3.0 CMs are most likely to be associated with a UGA having a bandpass (BP) amplification of up to 1 GHz. This can affect the result of a pilot carrier passing thru the UGA and giving a false positive if a UGA has a BP of 1.2 GHz.

Note: In 2018, a new line of MoCA friendly splitters (MFS-MO) that are designed with the intent to improve support of MoCA-only video clients can have an internal design that may not follow the traditional internal splitter network as depicted in Section 4.5.2. This can affect the expected total insertion loss and may lead to false positives. An independent event weight could be added to affect the probability based on if a technician visited the customer. The developer of the algorithm can include the likelihood that a technician replaced an existing splitter with the MFS-MO.

5.3.3. ScMod/VLPER/NPER: Average bit load of active subcarriers

- MoCA 1.1 uses the ScMod bit load reporting and has a maximum bit load of 8 bits or 256-QAM per subcarrier bit load reporting.
- MoCA 2.0 uses the VLPER/NPER bit load reporting and has a maximum of 10 bits or 1024-QAM per subcarrier.

By calculating the bitload's mean and SD, one can conclude the overall signal quality relative to the Rx power from the **mocaIfAcaTotalRxPower** SNMP object obtained from the ACA EVM probe.

A high receive power along with a high bit load mean and low SD (absent of the guard band subcarrier indexes) can indicate low insertion loss along the signal path. Depending on the number of devices in the MoCA network, one can estimate the total insertion loss of the coaxial system and develop a probability of the effects of a PoE would have if it was installed or not.

6. Use Case for a Home Pre-Check for Full Duplex DOCSIS Using MoCA 2.0

One of the most attractive features of full duplex (FDX) DOCSIS is its ability to scale to support upstream service offerings greater than 1 Gbps. Many customers may not want the additional upstream (US) speeds and may opt only for the increase of downstream (DS) speeds. In this case, the customers with a D3.1 CM would only upgrade the DS tier service, but customers with a D3.0 CM would be required to upgrade to the latest FDX CM.

To reduce the number of professional installs, operators may offer the customer a self-install kit (SIK) option. Before that option can be given, a home health check is needed. One of the most important checks

is the ability to determine if a UGA is installed. For FDX to operate, the customer's home must not have a UGA in place.

Electrically, a UGA may not be easily detected because electrically it may not offer any RF metric cues using only DOCSIS 3.0 PNM options. The current implementation at the time of this paper, FDX works on the premise of an entirely passive coaxial system, from the R-PHY node down to the FDX CM. Using the home isolation strategies detailed in Section 4.2, one could determine with high confidence if a UGA is present if:

- A ripple is present with an ACA EVM probe; the likelihood that a UGA is present is low to close to zero depending on the $\text{ripple}_{\text{pk-pk}}$, because current UGA designs do not have an internal MoCA PoE filter.
- No pilot carrier is detected during an ACA quiet probe, a UGA may be present, the PoE is at the root of the splitter, or there is a PoE filter embedded in the splitter.

7. Conclusions

As operators increase their investment in wireless protocols for delivering in-home video IP services, MoCA continues to be the primary workhorse for in-home video distribution and will be for years to come. Operators should take advantage and leverage the advances in the MoCA technology not just in video service delivery, but its ability to improve their RF signal analysis of the home coaxial network.

MoCA Coexistence Challenges Lessons Learned

Prior to the introduction of the ACA, operators only had one reliable metric cue called the PHY rate. Although it provided a pass/failure exit criteria for the technician, it falls short as a tool in determining the root cause when troubleshooting MoCA issues.

Section 4 demonstrated the advantages of the ACA in determining and resolving elusive customer impacting issues using the following implementations:

- The use of the ACA probe as a network analyzer, and signal processing techniques to determine linear distortion such as micro-reflections and splitter frequency response.
- The usage of the per-subcarrier bit load reporting can be used to determine the effects of the linear distortions.
- The location of the PoE filter can cause undesirable implications to the MoCA signal.
- Key performance indicators that provide a better understanding of the overall health of the in-home coaxial network.

Point of Entry Detection for MoCA Home Isolation

Preserving the customer's MoCA home isolation is essential in the maintaining of account CPE assets and the inadvertent ingress of MoCA signals into the HFC network. In Section 4.2 and Section 4.3.1 this paper explored the various methods and calculations to detect a PoE filter using both the EVM and quiet ACA probes by means of the following methodologies:

- Signal processing analysis to detect MoCA OFDM echoes using the ACA EVM probe.
- Using a pilot carrier and performing a statistical description analysis to determine the presence of a PoE filter using the ACA quiet probe.

MoCA Home Assessment Key Performance Indicators

The ACA by itself is not enough to determine the overall health of both the MoCA channel and the in-home coaxial network, because each customer's installation practice differs since the introduction of MoCA years ago. The availability of MoCA's diagnostic technology has caught up to the legacy installation of the in-home coaxial network.

Section 5 detailed the use of the statistical description attributes of skewness, mean and standard deviation, along with the ACA EVM receive power as a valid indicator of the MoCA physical layer health in the home coaxial network.

Finally, Section 5.3 provided a high-level introduction to the use of machine learning, specifically the algorithm of conditional probability. The section provides the conditions to develop a probability to determine the presence of a PoE filter or UGA.

In conclusion, as the title of this paper indicates, operators should expand their use of the ACA and the per-subcarrier bit load reporting. The use of the PHY rate will become less relevant as the gold standard to determine if a home is MoCA-ready. The ACA continues to open the door to provide richer data to improve proactive network analysis and the technician confidence of a high performing home coaxial network.

8. Acknowledgments

With great appreciation, I want to acknowledge Jon Cave, Rick Vetter, Paul E. Schauer and Jaci Bates from Comcast Cable along with John Rak, a freelance editorial expert, for their thoughtful review and guidance in helping to prepare this material for publication.

9. Abbreviations

ACA	alternative channel assessment
bps	bits-per-second
bpsym	bits-per-symbol
BPSK	bi-phase shift keying
CNR	carrier-to-noise ratio
CM	cable modem
dB	decibel
DS	downstream
DStat	descriptive statistics
DUT	device under test
EVM	error vector magnitude
FEC	forward error correction
HFC	hybrid fiber/coax
HP	high pass
HTTP	hyper text transfer protocol
Hz	hertz
kbps	kilobits per second
KPI	key performance indicator
LCI	LDPC code iteration

LDPC	low-density parity check
LP	low pass
OFDM	orthogonal frequency division multiplexing
OSS	operation software support
PDF	probability density function
PNM	proactive network maintenance
PoE	point-of-entry [filter]
MER	modulation error ratio
MoCA	Multimedia over Coax Alliance
MR	micro-reflection
NA	network analyzer
NVP	nominal velocity of propagation
RF	radio frequency
RL	return loss
RT	round trip
SA	spectrum analyzer
SCTE	Society of Cable Telecommunications Engineers
SD	standard deviation
SNR	signal-to-noise ratio
SNMP	simple network management protocol
SWR	standing wave ratio
TDR	time domain reflectometer
VSA	vector signal analyzer
VSWR	voltage standing wave ratio
WMA	window moving average
WS	window size

10. Bibliography and References

- [1] R. Thompson and M. Garcia, "MoCA PNM Paradigm Shift," *SCTE Winter Journal*, 2017.
- [2] M. o. C. Alliance, "MoCA MAC/PHY SPECIFICATION v2.0," 2010.
- [3] MoCA Multimedia over Coax Alliance, "Final_Best-Practices-for-Installation-of-MoCA_170516rev01," March 2017. [Online]. Available: http://www.mocalliance.org/technology/Final_Best-Practices-for-Installation-of-MoCA_170516rev01.pdf. [Accessed 2 2018].
- [4] Wikipedia, "Picture Of Skewness," 2018.

- [5] M. Manohar, "Machine Learning 101: What is a conditional probability," [Online]. Available: <https://medium.com/@mithunmanohar/machine-learning-101-what-the-is-a-conditional-probability-f0f9a9ec6cda>. [Accessed Nov 2018].
- [6] MOCA Alliance, "Remote Management of MoCA Interfaces using SNMP MIB," MOCA, 2018.

11. Appendix

11.1. Fourier Transform

This section describes the process to perform a Fourier transform (FT). The Fourier transform decomposes a function of time (time domain) into the discrete frequencies (frequency domain). Although the ACA is a frequency domain representation, this process interprets the signal as a time domain signal.

11.2. Discrete Fourier Transform

Discrete Fourier transform (DFT) has a higher computational complexity than the FFT. There is plenty of open source implementation, in various languages. See [Implementing the DFT](#)

$$X(k) = \sum_{t=0}^{n-1} x(t)e^{-2\pi tk/n}$$

Equation 7: Discrete Fourier Transform (DFT)

11.3. Fast Fourier Transform Software Package

The [Apache Commons Fast Fourier Package](#) was used for its FFT analysis. The suite is the base package, and additional software development is needed to normalize the ACA data for processing.

11.4. Normalizing the ACA for FFT Analysis

The following section describes the steps to normalize the ACA data for FFT analysis.

11.4.1. ACA Sample Size vs. FFT Points

The ACA sample size consists of 512 log-magnitude points. The FFT points are the number of sample points that are submitted to the FFT. The FFT requires the total number of points to be a power of 2.

For a better resolution, a technique called zero-padding is used to extend the number of FFT points.

$$FFT[i]_{Points(1024)} = \{x_{[0]} = x_{ACA_{dBm}}, x_{[1]} = x_{ACA_{dBm}}, \dots, x_{[511]} = x_{ACA_{dBm}}, x_{[512]} = 0_{ZeroPad}, \dots, x_{[1023]} = 0_{ZeroPad}\}$$

Equation 8: FFT Point Zero-Padding Technique

The ratio between sample points and FFT points SHOULD be [1:8 or .125], but no more than [1:2 or .5]

$$FFT_{Scale} = \frac{ACA_{sampleSize}}{FFT_{Points}} = 1 : 8 = .125$$

Equation 9: FFT Scale Calculation

For signal delay detection, add 512 zeros to the end of the ACA list of values, for a total of 1024.

11.4.2. Sliding Window Moving Average

A conventional signal processing technique to remove high frequencies or spikes in a given time-domain signal is the use of a sliding window moving average.




Figure 24 - 3-Point Moving Average

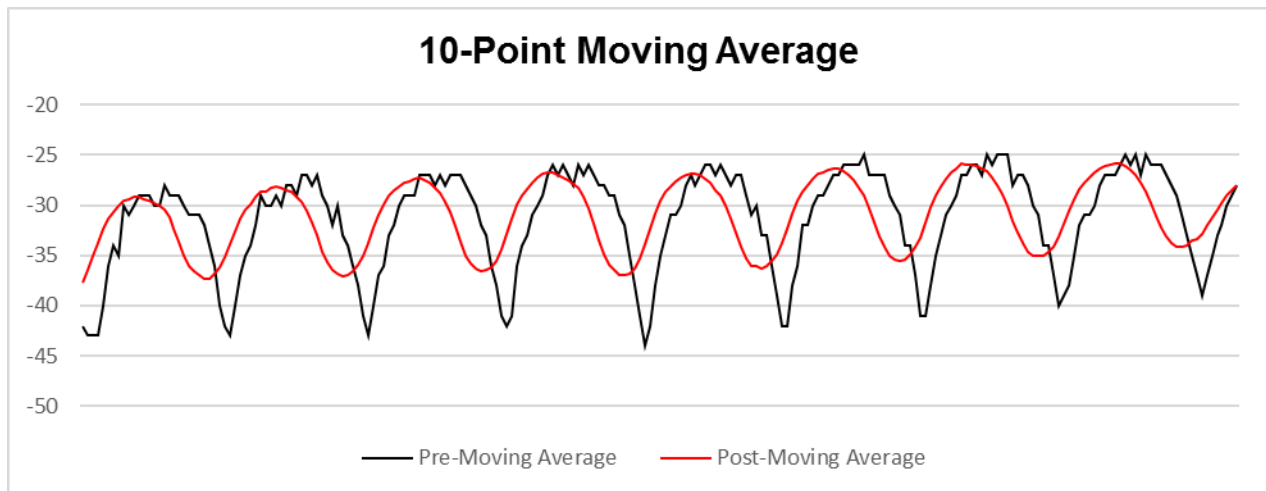


Figure 25 - 10-Point Moving Average (Over-Average)

x[0]	x[1]	x[2]	x[3]	x[4]	x[5]	x[6]	x[7]	x[8]	x[9]
------	------	------	------	------	------	------	------	------	------

$$x[k] = \sum_{k=0}^n \frac{x[k] + x[k + (n - 1)] + \dots}{n}, \dots$$

k = array index, x = ACA array, n = window size

Equation 10: Sliding Window Moving Average

$$x[0] = \frac{x[0]+x[1]+x[2]}{3}, x[1] = \frac{x[1]+x[2]+x[3]}{3}, \dots$$

Table 8 - Sliding Moving Average Pseudo Code (Java)

```
// Note: The following is not optimized for efficiency
public class MovingAverage {
    private List<Double> ldEntry = new ArrayList<Double>();

    public MovingAverage(List<Double> ldData) {
        this.ldEntry.addAll(ldData);
    }

    public List<Double> getAverage(int iWindow) {
        List<Double> ldAverage = new ArrayList<Double>();

        double dSum = 0;
        double dAverage = 0;

        int iWindowDivisor = 0;

        /* Traverse thru Data Set */
        for (int iIndex = 0; iIndex < ldEntry.size(); iIndex++) {

            /* Reset Average/Sum */
            dSum = 0;
            dAverage = 0;
            iWindowDivisor = 0;

            for (int iWindowIndex = iIndex, iWindowIndexCount = 0;
```

```

        iWindowIndexCount < iWindow;
        iWindowIndex++, iWindowIndexCount++) {

            if (iWindowIndex >= ldEntry.size()) {
                break;
            }

            /* Update Sum */
            dSum += ldEntry.get(iWindowIndex);

            /*
            This is require so to adjust the denominator value when the
            Index reaches the end of the list size
            */
            ++iWindowDivisor;
        }

        /* Moving Average apply Sliding Window Number */
        dAverage = dSum / (double) iWindowDivisor;

        ldAverage.add(dAverage);
    }
    return ldAverage;
}
}

```

11.4.3. Converting ACA points from dBm to Linear Scale

Before submitting the 1024 sample points, each value requires being converted to a linear scale.

$$x_{Linear} = 10^{\frac{x_{dBm}}{10}}$$

Equation 11: dBm to Linear

$$x_{dBm} = 10 \cdot \text{Log}(x_{Linear})$$

Equation 12: Linear to dBm

11.5. Calculating FFT Results

The following section describes the steps used to calculate the FFT results to obtain the ACA ripple signal found in the ACA data.

11.5.1. Complex Number Magnitude Power

After executing the FFT, the algorithm will produce a list of complex numbers. To calculate the magnitude power of the given Index_Freq in hertz (Hz), perform the following calculation for each complex number index.

$$\text{Magnitude Power}_{Linear} = \sqrt{\text{Real}^2 + \text{Imaginary}^2}$$

Equation 13: Complex Number Magnitude Power Calculation

11.5.2. Graphing Magnitude Response

When graphing, it is recommended to convert the linear magnitude power to dBm or dB to have a clearer FFT frequency response. To convert to dBm, use Equation 12.

$$dB = \text{Log}(\text{MagnitudePower}_{Linear})$$

Equation 14: Linear to dB Calculation

11.5.3. Understanding the FFT Results

The FFT is converting a time-domain signal to a frequency-domain response. The ACA consist of two OFDM channels, each consisting of 256 points for a total of 512 points. From a time-domain signal point of view, the ACA looks like a square wave with a 90% duty-cycle.

To appreciate the process of the Fourier-series of a square wave, please review: [Fourier-Series Square-Wave Video](#).

After the FFT, the following algorithm depicted in Figure 8 is to “pick-off” the ripple, by determining the index location before it is scaled. The scaling process determines the actual frequency of the ripple.

11.6. Normalizing ACA Profile Data

The procedure to normalize the ACA profile data can be found the SCTE Journal paper “The MoCA PNM Paradigm Shift” [1]

11.7. OSS SNMP Call Flow

11.7.1. ACA Data Acquisition

The following describes the step-by-step CLI procedure for operating and obtaining the ACA.

Table 9 - ACA EVM | quiet SNMP calls

Source NodeID for EVM Probe 0 - 15	snmpset -v 2c -c <RW_String> -m all -M < MIB-DIR > udp: <IPv4> udp6:<IPv6> mocaIfAcaNodeID.1 i <NodeID>
---------------------------------------	--

0 = Quiet 1 = EVM	<pre>snmpset -v 2c -c <RW_String> -m all -M <MIB-DIR> udp: <IPv4> udp6:<IPv6> mocaIfAcaType.1 i <ACA_Probe_Type></pre>
1150MHz = MoCA Chan <46> 1150MHz/25MHz = 46	<pre>snmpset -v 2c -c <RW_String> -m all -M <MIB-DIR> udp: <IPv4> udp6:<IPv6> mocaIfAcaChannel.1 u <MoCA_Channel></pre>
<p>BitMask Convert to Binary</p> <p>$2^{\text{NodeID}} = \text{BitMask}$</p> <p>0000000000000001 = NodeID: 0 Reporting HEX: 00 00 00 01</p> <p>0000000000000010 = NodeID: 1 Reporting HEX: 00 00 00 02</p> <p>0000000000000100 = NodeID: 2 Reporting HEX: 00 00 00 04</p> <p>.</p> <p>.</p> <p>1000000000000000 = NodeID: 15 Reporting HEX: 80 00 00 00</p> <p>Multi-Report-Node</p> <p>000000000001001 = NodeID 0 + 3 00 00 00 09 = HEX (NodeID:0 + 3)</p>	<pre>snmpset -v 2c -c <RW_String> -m all -M <MIB-DIR> udp: <IPv4> udp6:<IPv6> mocaIfAcaReportNodeMask.1 x "<00 00 00 01>"</pre>
	<pre>snmpset -v 2c -c <RW_String> -m all -M <MIB-DIR> udp: <IPv4> udp6:<IPv6> mocaIfAcaInitiate.1 i <1 = Start_ACA_Process></pre>
Poll until success(0)	<pre>snmpwalk -v 2c -c <RW_String> -m all -M <MIB-DIR> udp: <IPv4> udp6:<IPv6> mocaIfAcaStatus.1</pre>
	<pre>snmpwalk -v 2c -c <RW_String> -m all -M <MIB-DIR> udp: <IPv4> udp6:<IPv6> MocaIfAcaEntry</pre>

Table 10 - MoCAIfAcaEntry SNMP call

```
snmpwalk -v 2c -c <RW_CommString> -m all -M <MIB-DIR> udp: <IPv4> | udp6:<IPv6>
MocaIfAcaEntry
```


	MOCA20-MIB::mocalfAcaNodeID.1 = Gauge32: 1
	MOCA20-MIB::mocalfAcaChannel.1 = Gauge32: 46
	MOCA20-MIB::mocalfAcaReportNodeMask.1 = BITS: 00 00 00 01 31
	MOCA20-MIB::mocalfAcaInitiate.1 = INTEGER: 0
<p>MocaAcaStatus</p> <p>DESCRIPTION "Represents the status of the last Aca probe"</p> <p>SYNTAX INTEGER { success (0), fail-BADCHANNEL (1), fail-NOEVMPROBE (2), fail (3), in-PROGRESS (4) }</p>	MOCA20-MIB::mocalfAcaStatus.1 = INTEGER: success(0)
Result in dBm	MOCA20-MIB::mocalfAcaTotalRxPower.1 = INTEGER: -8
<p>Two's Complement Representation.</p> <p>Each return value is positive, but the actual value is negative.</p> <p>Implementer requires to multiply by -1</p> <p>0x2F = 47 (-1) = -47dBm</p>	MOCA20-MIB::mocalfAcaPowerProfile.1 = Hex-STRING: 2F 2F 2E 2F 2E 2D 2E 2D 2B 2B 2B 26 21 1C 1B 1B 1C 1B 1B 1C 1A 1A 1C 1A 1B 1B 1A 1B 1B 1D 1C 1A 1A 1B 1B 1A 1B 1A 1A 19 1A 1B 1B 19 1B 19 1A 1B 1A 18 1A 1A 1A 1B 1A 1A 1A 19 19 19 19 1B 1A 1B 1A 1A 1A 1A 1A 1B 1A 1A 1A 1B 1A 1A 1A 1A 1A 1A 1A 1B 1B 1A 1A 1C 1B 1C 1B 1B 1B 1B 1B 1B 1B 1C 1D 1C 1C 1B 1C 1B 1B 1C 1B 1C 1C 1B 1A 1B 1B 1C 1D 1C 1C 1B 1C 1C 1B 1C 1C 1B 1C 1C 1B 1C 1C 1B 1A 1C 1B 1B 1C 1C 1B 1C 1B 1C 1D 1C 1D 1C 1B 1C 1C 1B 1D 1C 1B 1B 1D 1C 1B 1C 1C 1B 1C 1D 1B 1A 1C 1D 1C 1C 1C 1C 1C 1D 1D 1C 1D 1B 1D 1D 1C 1D 1D 1C 1D 1C 1C 1D 1C 1C 1D 1D 1D 1C 1D 1D 1C 1E 1C 1C 1C 1D 1D 1D 1D 1C 1C 1D 1C 1D 1C 1D 1E 1C 1D 1C 1C 1C 1E 1C 1C 1C 1D 1D 1E 1C 1D 1D 1C 1D 1D 1C 1D 1D 1C 1D 1D 1E 1D 1C 1C 1D 1C 1C 1C 1C 1E 1C 1D 1C 1C 1C 1D 1C 1B 1D 1D 1D 1C 1B 1D 1C 1E 1D 1D 1D 1D 1E 1D 1D 1E 1E 1E 1C 1D 1D 22 24 28 2C 26 2C 2A 27 1F 1D 1D 1E 1E 1E 1D 1D 1C 1D 1C 1D 1E 1D 1C 1C 1D 1C 1B 1E 1B 1B 1C 1B 1D 1C 1C 1C 1C 1D 1C 1D 1C 1E 1D 1C 1C 1D 1C 1C 1C 1C 1D 1C 1C 1E 1C 1C 1B 1B 1F 1C 1C 1C 1C 1B 1C 1C 1C 1C 1D 1C 1D 1A 1D 1E 1D 1E 1C 1D 1D 1D 1C 1C 1B 1C 1D 1E 1A 1C 1A 1B 1B 1D 1C 1B 1D 1B 1B 1D 1C 1D 1D 1C 1D 1A 1C 1D 1C 1B 1D 1D 1D 1C 1B 1D 1C 1C 1C 1D 1C 1A 1A 1C 1C 1C 1C 1D 1C 1D 1C 1A 1C 1C 1C 1D 1C 1A 1B 1B 1D 1B 1A 1B 1B 1C 1C 1B 1C 1A 1B 1C 1C 1B 1A 1B 1B 1A 1A 1B 1B 1A 1B 1B 1C 1A 1A 1A 1A 1A 1B 1A 1B 1A 1A 1A 1B 1A 1A 1A 1A 1A 19 1A 1B 1A 1A 1A 1A 1C 19 1A 1A 1A 19 1A 19 19 1A 19 1A 1A 1A 19 19 19 1A 1B 19 1A 19 1A 1A 19 1A 19 18 1A 19 1A 1A 19 1A 19 19 1A 1A 1A 1A 19 19 1A 1A 1B 19 19 18 1A 1A 1A 1A 1B 1A 1C 20 25 2A 29 2A 2C 2C 2C 2E 2E 2E 2E

Operational Practices for the Deployment and Maintenance of Wireless Devices on HFC Plant

Practical Observations on Deployment Planning, Installation, Maintenance and Ongoing Operational Support

An Operational Practice prepared for SCTE•ISBE by

Dan L. Whitehouse, Director, Communications Media and Content
Hitachi Consulting Corporation, SCTE•ISBE Member
3524 Pineywood Trace
Birmingham, AL 35242
dwhitehouse@hitachiconsulting.com
(205) 541-7742

Cole Reinwand, VP, Wi-Fi Strategy
Comcast Corporation
1701 John F Kennedy Blvd
Philadelphia, PA 19103
Cole_Reinwand@cable.comcast.com
(215) 286-4107

Matt Zielenbach, Sr. Director, Xfinity Wi-Fi Product, Wi-Fi Business Operations
Comcast Corporation
1701 John F Kennedy Blvd
Philadelphia, PA 19103
Matthew_Zielenbach@cable.comcast.com
(215) 286-8832

Table of Contents

Title	Page Number
Table of Contents	68
1. Introduction	70
2. Considerations for Deployment Planning	71
2.1. Workflow Management	71
2.2. Device Management	73
2.2.1. Benefits of Staging Devices	76
2.2.2. Provisioning Considerations	76
2.3. Site identification information	77
2.4. Tools and telemetry	78
2.5. Preparing for deployment at scale	79
2.5.1. Identify functional roles, responsibilities and required staffing levels	79
2.5.2. Assess the ability to leverage existing tools	80
2.5.3. Establish clear field troubleshooting steps and escalation paths	81
2.5.4. Document staging processes and expected boot times	81
2.5.5. Determine device power consumption for plant design	81
2.5.6. Device tracking and return accountability	82
2.6. Planning for bi-directional communication	82
2.7. Measurements, metrics and uptime	83
2.8. Operational Support	83
3. Deployment management	84
3.1. Identify service champions for each local market	84
3.2. Identify local market challenges	86
3.2.1. Local market legal considerations	86
3.2.2. Coordination with other programs in the market	86
3.3. Segment markets for parallel deployment	87
3.4. Training Considerations	88
4. Summary of Operational Efficiencies Gained Through Common Practices	88
5. Conclusion	91

List of Figures

Title	Page Number
Figure 1 - Operational Efficiency Through Common Planning and Deployment Practices	70
Figure 2 - HFC strand mount options for wireless devices	73
Figure 3 - Associating a Radio Device With the Wrong HFC Cable Modem	75
Figure 4 - Managing Deployment with a Service Champion	85
Figure 5 - RF Device Notification Drop Tag	86
Figure 6 - Market Segmentation for Parallel Deployment	87

List of Tables

Title	Page Number
Table 1 – Workflow States	73
Table 2 – Functional Roles and Responsibilities	80
Table 3 – Device Offline Aging Report	84
Table 4 – GPS Delta Report	84

1. Introduction

The purpose of this document is to identify operational efficiencies that can be gained through a common set of deployment and operations support practices for wireless services deployed on hybrid fiber coax (HFC) plant. There are many technologies in varying stages of maturity that utilize wireless communications such as Wi-Fi, small cell, LoRa and Citizens Broadband Radio System (CBRS) that allow a cable operator to reach and provide new or enhance existing services to residential, business and wholesale customers. Devices supporting these wireless services act as an aggregation point on the cable network, and there are commonalities across the different service lines that can be consolidated into a consistent and common approach for deployment and operational support.

Although many devices supporting these technologies can utilize non-traditional mounting points such as exterior building mounts, off-network billboards or municipal assets such as light poles, the HFC plant itself continues to serve as an essential resource for strategic mounting points outside of the customer premise to achieve full coverage. The HFC plant also provides power and data backhaul from the edge network via an onboard or separate HFC-powered cable modem. But anytime the outside plant is touched, there is the potential for negatively impacting customer services.

Thus planning, deployment, and operations activities in support of installing wireless devices on the outside plant must be carefully thought out and coordinated to ensure customer satisfaction is achieved as the coverage of new wireless services is expanded. As depicted in Figure 1, the impact on ongoing operational support can be substantially reduced through common planning and deployment practices across multiple wireless services.



Figure 1 - Operational Efficiency Through Common Planning and Deployment Practices

While each of these technologies has the potential to drive new services and sources of revenue, they also share many similar attributes that can be consolidated into a common deployment planning process and operational support model. Topics addressed by this paper are as follows:

- Considerations for deployment planning
- Deployment management
- Summary of operational efficiencies gained through common practices

2. Considerations for Deployment Planning

The primary emphasis of this paper is on up-front preparation for deployment of multiple wireless devices and services on HFC plant. Advance preparation and planning that consolidates similar processes and requirements across multiple services or lines of business can streamline the installation and deployment process. It can also reduce downtime and help improve the corresponding time to market, as well as greatly simplifying the level of operational support required in the local market to maintain the devices on plant over time.

Regardless of the wireless technology being deployed to reach the end customer, there are several elements each of these services has in common that should be considered in the initial planning stages and field trials. These considerations should be followed through to incorporate lessons learned for production deployment at scale, as well as to provide consistency in product maintenance and operational support for the devices supporting each service. The main considerations are listed below:

- Workflow management
- Device management
- Site identification
- Tools and telemetry
- Preparing for deployment at scale
- Planning for bi-directional communication
- Measurements, metrics and uptime
- Operational support

Each of these is described in the following subsections.

2.1. Workflow Management

A key initial decision point is to determine how to manage the workflow and overall end-to-end deployment process. This can be done either through traditional billing systems or through a custom solution and can be either developed in-house or use off-the-shelf software which is then customized to meet the operator's specific needs.

Wireless devices mounted on HFC plant often do not fit easily into traditional support models. They are neither customer premise equipment, nor are they equipment essential to the function of the HFC network itself like plant actives such as fiber optic nodes, power supplies or amplifiers. Each radio device is essentially a collection point that funnels off-net traffic for a specific service over the HFC plant. Nonetheless, the workflow for installation of wireless devices can be systematized in one of two ways:

1. Managing workflow through a traditional billing system: This approach leverages existing infrastructure to manage workflow, mounting locations, device inventory and provisioning. Specific equipment types and accounts should be created for each location or service address, and device provisioning processes should follow similar methods used for residential customer premise equipment. However, traditional billers are geared towards installation in a customer premise. Processes are not easily adapted towards support of equipment directly mounted on HFC strand. The actual installation location may not align to a specific service address.

2. Managing workflow through a custom in-house application or off-the-shelf workflow management software: Using a custom solution has the benefit of tailoring the system to meet an operator's specific needs for the deployment of multiple services utilizing the HFC plant for mounting and backhaul. This approach also lends itself to the development of repeatable processes to support the specific needs of each new service as it is deployed on plant.

In either scenario, the primary purpose of the workflow manager is to maintain an accurate correlation of all devices currently deployed at any given location. It should also maintain the current status of each location as it progresses through the various workflow states. The workflow management system should maintain site information, and either internally contain or have transparent access to current and accurate device inventory information.

Whether internally developed, or using customized off-the-shelf software, the workflow management system should be able to perform the following functions at a minimum:

- Distinguish between the different high-level services supported by the workflow manager (e.g., Wi-Fi, LoRa, CBRS). While many of the deployment processes are similar, the workflow manager needs to be able to track and report out on workflow separately by the type of business service being deployed.
- Each service should have the capability to be segmented further by the operator's administrative or geographical management areas (e.g., market, region, division). This allows workflow for any given service to be managed at the local market level.
- Ingest and store site-specific information with potential locations to mount an asset.
- Store device inventory internally or have integrated access to existing inventory management systems.
- The workflow manager should be able to accurately correlate which device is currently installed at any given site, as well as the swap history of previous devices at the same location.
- The workflow manager should require the minimum amount of information necessary to track and manage the deployment process through the states listed in Table 1:


Table 1 – Workflow States

Workflow Status	Workflow Description
Site Review	Is the location suitable for placement of the device and supported service?
Site Hold	Site cannot be installed, either temporarily or permanently
Temporary Hold	Site is viable, but cannot be installed at the current time
Permanent Hold	Site is not viable
Site Under Design	Site is in the process of being designed on plant
Site Under Construction	Pre-installation construction activities (e.g., power supply upgrades, plant extensions)
Site Ready for Install	The device is ready to be placed onto the plant with a power passing tap or directional coupler
Site has Device on Plant	A device has been installed at the site and documented in the workflow manager
Site Flagged Investigation	A device installed on plant that requires some level of additional follow up
Site Decommissioned	Removal of devices from a site and flagging the site as no longer suitable


2.2. Device Management

Management of intelligent devices on HFC strand involves more than simple provisioning and configuration of the radio and cable modem on a device or set of paired devices. The following considerations can apply to devices shown in Figure 2, including those with an integrated board with a radio and cable modem, devices with an onboard cable modem and radio in the same chassis, and to radio devices connected to a separate HFC cable modem via an external power over Ethernet (PoE) cable.

Integrated Radio and HFC Cable Modem



Separate Radio and HFC Cable Modem with PoE mounted end-to-end



Separate Radio and HFC Cable Modem with PoE mounted as a side-car




Figure 2 - HFC strand mount options for wireless devices

Regardless of whether the radio media access control(s) (MAC) and cable modem MAC reside in the same or different chassis, there are several operational implications to consider in storing and correlating information about these devices that should be accessible through the workflow manager:

- Wireless devices contain additional information that should be stored and maintained for ongoing operational support. Minimum device information the workflow manager should maintain or have transparent access to via an inventory management system:
 - Vendor
 - Model
 - Serial number
 - Cable modem MAC address (where applicable)
 - Radio MAC addresses (where applicable)
 - Basic service set identifier (BSSID)s or virtual MAC addresses (where applicable)
 - Date of purchase
 - Warranty expiration date
 - Original purchase order number or RMA (return merchandise/material authorization) number

- Correlation of the radio MAC address(s) of the device with the MAC addresses of the cable modem that is providing backhaul is critical. This relationship must be tightly maintained, accurately documented and kept current because of the following:
 - Accuracy of this data can drift over time even with integrated devices in the same chassis due to RMA repairs and refurbished replacements (cable modem and radio MACs can return in different devices with new serial numbers). This can be offset by utilizing the same process to ingest and overwrite device data for new purchases, RMA repairs and replacements. Regardless of whether it's a new purchase or a repaired device that's been returned, an electronic file should be provided by the vendor with all relevant device information in the same format.

 - It becomes even more difficult to maintain this relationship when the radio and cable modem are two physically separate devices, with the radio\ connected via a powered Ethernet cable to the HFC strand mount cable modem. To effectively maintain the logical relationship between two physically separate devices, the following practices can assist in keeping this relationship current and correct:
 - Operationally, both devices should be viewed and treated as a single piece of equipment in the local market.
 - Pre-staging and assembly of both pieces of equipment should be performed at a centralized or third-party location. Electronic documentation should be created at time of assembly to correlate radio MACs with the cable modem MAC.
 - It is inadvisable to assemble components in the field at time of install. Radio and HFC modem device pairs should be staged and documented at a centralized or third-party location, then installed as a single cohesive unit.
 - It is also inadvisable to swap or troubleshoot individual components in the field. Instead, treat both devices as a single unit, swapping both devices out and returning them to the warehouse as a single unit connected by Ethernet.

- It is vital that electronic documentation be updated when changing either device of a paired radio and HFC cable modem at a centralized or third-party location.

A common error that is preventable by the above procedures is associating the radio MAC address(s) with the wrong cable modem MAC address, as depicted in Figure 3.




Figure 3 - Associating a Radio Device With the Wrong HFC Cable Modem

- Operational which result from not maintaining this logical relationship in inventory or the workflow manager are as follows:
 - Backoffice systems gathering and processing radio data can be matched to an incorrect location
 - Loss of where a radio device is located on plant
 - Technicians are dispatched to the wrong location, returning with no problem found
 - Elongated service downtime
 - No ability to correlate local plant issues or maintenance with connectivity problems between the radio and back-office service
 - Inability to correlate customer reported issues to the correct device
 - Inability to correlate service metrics and utilization to correct locations
- The device should be able to provide visual confirmation that it is functioning properly. There are pros and cons to having visual LEDs or other indicators on the device. While lights on the device have the potential to attract vandalism or customer complaints, they are also the simplest and most direct way for a technician to confirm if a device is functioning at time of install or when dispatched to correct or swap a failed device.
- Either the HFC cable modem, the radio device, or both should ideally have an internal GPS to provide accurate location information about where it resides on plant. While this is an additional method to locate equipment on plant, not all locations are able to receive a clear GPS signal due to line of sight or other interference that can impair the accuracy of the GPS in the device. The state or accuracy of GPS coordinates provided by the device can vary depending on the number of satellites it has within line of sight, sometimes separated by hundreds of feet. Maintaining precise records of the location where a device is designed and installed on the plant, in addition to storing GPS readings from the device itself, can greatly reduce downtime and truck rolls to an incorrect location.

- Obtain the maximum permitted exposure (MPE) documentation from the hardware vendor for each type of radio device being deployed. This will almost certainly be required either by a municipality or a pole owner within the operator's footprint.

2.2.1. Benefits of Staging Devices

Unlike customer premise equipment (CPE), these devices are intended for mounting on HFC plant, in most cases on aerial strand or a pedestal. Staging of these devices prior to installation can result in substantial savings on installation time as well as post-installation support and decreased downtime. Staging involves the following activities:

- Pre-stage devices in the warehouse or installer's facility by connecting the device to a powered drop to pull down the latest firmware and code. This can reduce install time in the field by 10 to 15 minutes per site by avoiding code upgrades in the field at time of install. This requires close coordination between software release schedules and deployment planning. Staging will also identify any devices that are non-functional, so they can be immediately returned to the vendor for repair without ever going to the field.
- If not already provided by the vendor, label the bottom of the device with the last four to six characters of the serial number in non-metallic mailbox type letters that can clearly be seen from the ground. This allows technicians called out to check or swap a location to easily identify the correct device from the ground without having to use a bucket truck to verify the serial number or MAC address printed on the device label in a small font.
- Cover any exterior labels that have the serial number, cable modem MAC or radio MAC addresses with clear weather- and UV-resistant tape. This can keep information legible over a longer period despite weather conditions and extend the practical use life of the equipment.
 - If a device with a weathered label is non-functional and has been swapped, there is no way to identify the device or verify if it is under warranty other than returning to the vendor.
 - This also impacts functional devices returned to the warehouse that have weathered labels. For practical purposes, this can make it impossible to clearly identify the device and associate it to another site for redeployment.

2.2.2. Provisioning Considerations

In provisioning the wireless devices, note that there are two provisioning processes going on in tandem, one for the HFC cable modem, and the other for the wireless radio and service. Even in instances with an integrated board, having separate MAC addresses for the radio(s) and cable modem helps to clearly distinguish between HFC connectivity and transport problems vs. issues with the radio device or service.

Wireless devices receive their portion of a boot file after the cable modem has fully come online. While the configuration of any radio is device specific, there are some common steps that rely on commonalities between these devices that can help provide a more consistent customer experience and reduce one-off issues over the life of the device. These steps include:

1. Explicitly state the values for default parameters critical for remote access to the device and functions critical to the service as part of the device boot process. Default values can change between versions of code over time, and taking the time up front to explicitly state values critical to the service in the boot process can provide a few advantages:
 - Consistency in device performance and customer experience in the production environment over time.
 - Devices returned through RMA may originate from another cable operator and may not utilize the same parameters or be factory reset before shipment from the vendor. Explicitly stating critical default values in the boot process reduces this liability.
2. When the radio and cable modem are physically separate devices connected by an Ethernet cable, ensure any timeout parameters on the radio device allow sufficient time for the cable modem to fully boot and establish back office connectivity. That way, the radio device can complete its own boot process after the cable modem is fully online and active. This can help avoid inconsistent behavior at initial power up and reduce the potential post-install truck rolls to reset the devices in the field, such as after routine plant maintenance or power loss due to weather.
3. If using a traditional biller, it's possible to provision the HFC cable modem and radio device on an account and service address created for that location. Correlation of these devices at the correct site is dependent on the installation tech for accuracy.
4. If using a workflow management system, there are a few options to support these devices as network equipment rather than CPE and still avoid a “walled garden” situation:
 - Manual provisioning – While hardcoding the device MAC into the provisioning system can be used to quickly provision devices to test basic functionality, it is not a sustainable long-term process for production deployment.
 - Automated provisioning (e.g., using DHCP discover packets with Option 43 to validate the device) provides a scalable way to provision an HFC-powered, strand-mounted cable modem. This provides the most flexibility for deployment, since each device type set up in this manner boots with a valid configuration and goes live on the network without any manual configuration or intervention.

2.3. Site identification information

The process for identifying and prioritizing individual locations for deployment will vary depending on the technology and type of services deployed. For example, plant-mounted Wi-Fi access points have a practical limit of approximately 300 to 350 feet, whereas a LoRa device can cover several kilometers depending on the environment (urban vs. suburban or rural). Regardless of the type of service provided to the end customer, there are common elements for any location that should be maintained.

At a minimum, site information should contain the type of service it will be supporting (e.g., Wi-Fi, small cell, CBRS), site proximity to plant, basic plant information (e.g., aerial, pedestal or vault, pole number, node name, available clearance to the next provider on aerial strand) and location information such as the latitude and longitude for device placement on plant and the closest approximate street address. Information stored for each site should also identify the corresponding cable company administrative area the site resides in (i.e., market, region or division).

In addition to data internal to the operator, information for these sites should also tie back to public data, such as county, city, state or ZIP Code. If a site is targeting a specific area or type of business (e.g., a park or restaurant), that information should be captured as well, using standard definitions such as a Standard

Industrial Classification 4 (SIC4) or North American Industry Classification System (NAICS) code. Inclusion of public information in the site detail allows for correlation of device data for reporting and other purposes, (e.g., performance or usage information by city, geographic distribution and density, type of business, easier correlation with competitive data).

2.4. Tools and telemetry

Tools supporting wireless devices mounted on plant should include not only the status on communication between the radio device and the back-office systems supporting the service, but should also identify where the radio and the device providing data backhaul are physically located on plant and detect plant conditions that can impact the service.

Devices are not always installed in the exact location where they are designed on the plant. This may not have an immediate impact on coverage for services that utilize longer range radio devices, but accurate placement on plant is essential to services such as Wi-Fi, that can have extremely short distance limitations.

There are longer term implications of inaccurate device placement that can affect the uptime and manageability of any service that are not immediately apparent with initial deployments, including:

- Inability or elongated time required to locate a device on plant for swap or repair
- Repeat truck rolls to the wrong location
- Swapping the wrong device
- Inability to meet third party service level agreements or commitments (e.g., uptime and availability)
- Negative impact to customer experience

Device placement issues can be addressed in several ways:

- The simplest solution is to require the technician to call in the site and device information to the operations center at time of install, while he or she is physically present on site. This allows the operations support staff to associate the device to the site in the workflow manager, as well as to verify that the device is functioning correctly before releasing the tech for the next install.
- The technician should provide both the serial number and MAC address to avoid errors in assigning the device to the wrong site. Requiring two pieces of information to identify a device will significantly reduce human error in instances where the MAC address on two different devices can vary by a single character, such as a “b” or a “d”.
- **NOTE:** Allowing a technician to call in devices and sites in groups or batches post-install, to associate a device to a site, can quickly cause errors to accumulate in the workflow manager. Associating a device with a site at time of install, while the tech is physically at the location, greatly reduces the opportunity for error.
- In addition to having the technician call in an install or a swap in real time while onsite, devices with an internal GPS can in many instances assist in identifying and resolving location errors and reduce the time required to locate a device on plant for post-install replacement or repair.

Tools supporting the service should periodically poll and store the GPS coordinates and GPS state (accuracy level) of the device on plant. This provides the following benefits:

- Ability to compare the distance between where devices are physically located on plant vs. the location identified in the workflow management system.
- Take appropriate action, either by updating the workflow management system to correlate the device to the correct location or by rolling a truck to reposition the device at the correct location it was designed for.
- Ability to locate an offline device on plant from the last valid GPS reading stored.
- While having a GPS in the device can be a valuable troubleshooting tool, it is not a solution for all sites and situations. Not all locations can receive a GPS signal (e.g., urban canyons, blocked line of sight, overhangs), or have line of sight to enough satellites to establish a solid “lock” and provide coordinates that are accurate within a few feet.

The tools used to operationally support the service should be able to provide real-time or near real-time status of both the radio and the associated backhaul cable modem, whether the device is an integrated unit, or the radio and HFC cable modem are separate devices. Just because the cable modem is up and running correctly doesn’t guarantee that the radio is also functioning properly, or that the radio has connectivity to the end customer device or back office systems supporting the service.

2.5. Preparing for deployment at scale

The requirements to support an initial field trial and the requirements to support a larger scale production deployment can be very different. Often the scope of an initial field trial is limited to testing the radio device and the back-office applications that make up the service. Besides validating basic service functionality, an initial field trial is also an opportunity to gain insight into issues that can affect the deployment process and impact longer term operational support. Examples include:

2.5.1. Identify functional roles, responsibilities and required staffing levels

Identify the functional roles and associated responsibilities listed in Table 2, and determine the staffing levels required to stage, deploy, maintain and operate the service, and include where the responsibility for each role resides (i.e., corporate, local market or a third-party). Upfront identification of the roles and responsibilities can often reveal when the local market may require additional headcount to support the deployment and ongoing operations.

Table 2 – Functional Roles and Responsibilities

Functional Role	Description
Deployment management	Program management and coordination of device deployment and operational support activities
Support documentation	Operational and training documentation required to install, maintain and support the service
Training	Creation of training material and conduct training for each role
Site identification	Location where a radio device will be placed on plant
Procurement	Purchase of production and spare devices, ancillary plant equipment required for installation
Staging/Kitting	Label, assemble and boot equipment to current code levels
Plant design	Update plant maps with design for power requirements, plant upgrades
Construction	Any required build activity, e.g., power supply upgrades
Installation	Install device on cable plant
Validation at time of install	Verify that the radio device and backhaul are functioning correctly
Ongoing service monitoring	Monitoring at the service / application level
On site repair or swap	Dispatch of a technician to a device location to correct plant issues, reset or replace a device
RMA processing	Return to vendor for repair or replacement
Site decommission	Return equipment to warehouse and decommission the site in the workflow manager
Operational reporting	Device status, uptime, offline and aging, mean time to repair
Service reporting	Metrics on the performance of the service, typically outside the scope of operations support

2.5.2. Assess the ability to leverage existing tools

Tools that provide service to a customer are not necessarily the same as those required for operational support to maintain the service once it's been deployed in the field. Determine whether existing support tools can meet the needs to deploy, manage and support the service. Document any gaps that need to be addressed either by application development or by process definition. Tool considerations during an initial trial include:

- Element managers are typically focused on a specific vendor's device, not an end-to-end service. The radio device is only one component of the solution that comprises a service to the end customer. Element managers are basically single function products, and while a potential tool for a Tier 2 / Tier 3 escalation, front line operational support tools should be able to support multiple devices and services.
- Determine which processes should be developed to leverage existing tools for operational support of the service.
- Determine whether existing tools should be modified, or if new toolsets need to be developed to support the service:
 - a. Can the tool validate that both the radio device and the associated backhaul device are functioning normally in near real-time?

- b. Do existing tools support outage identification, both at a service and the device level? (i.e., the ability to determine not only if the end device is up and functional, but also if it is connected to the back-office applications that comprise the service)
- c. Can existing tools identify where the equipment is located on plant?
- d. Can the tools correlate a customer reported issue to a specific device on plant (e.g., can the toolset correlate a BSSID as seen by a customer to the physical Wi-Fi access point to which they are connected)?

2.5.3. Establish clear field troubleshooting steps and escalation paths

Identify and document troubleshooting steps for operations centers and technicians installing or swapping a device on plant, and develop a clearly defined process for local market escalations to Tier2/Tier3 support. An initial trial is an excellent proving ground to identify issues that can impact the service and document processes to preempt or correct potential problems at the local market level.

2.5.4. Document staging processes and expected boot times

Document any staging requirements, including pictures of the device and steps for any physical assembly that may be required, including labeling and weatherproofing serial numbers or MAC address information for longer legibility over time. Identify the device information and electronic format required for ingestion into the workflow or inventory management system (e.g., CSV, XML).

Document the time required for the device to fully boot and for the service to become active. In addition, validate the average time from the initial boot sequence to the time the device shows online and active in operational monitoring tools. This will be needed for installer and operations center training to help prevent unnecessary swaps or RMA of equipment that may simply take a longer time to boot than might have been expected otherwise.

2.5.5. Determine device power consumption for plant design

Measure the device's actual power consumption and draw on plant both at rest and under full load. Radio devices over HFC plant consume power traditionally budgeted for running and operating the HFC cable plant. Understanding real world power consumption by the device and any deviation from the manufacturer's device specifications can help avoid overestimating power requirements and performing unnecessary power supply upgrades.

Services requiring a two-device solution (a separate radio device and HFC-powered cable modem) can place a substantial load on the network, depending on the required power draw of the radio device supplied over Ethernet from the HFC cable modem.

IEEE 802.3af Power over Ethernet (PoE) can draw up to 15.4 watts for the radio device, not counting the power requirements for the HFC modem itself. The IEEE 802.3at (PoE+) standard supports up to 25.5 watts per Ethernet port, again not including the power requirements for the HFC cable modem. The IEEE PoE standards allow the device supplying power (e.g., an HFC cable modem) and the device receiving power (e.g., a radio device) to negotiate the actual amount of power required to boot and run the powered device.

Conservatively estimating for the full power draw of multiple devices, simultaneously, from the same power supply, can quickly exhaust the available power budget. However, depending on the requirements of the radio device being powered, savings can be achieved by understanding the average and peak power draw of the radio device in real world conditions and under normal usage -- rather than conservatively designing to a sustained maximum PoE or PoE+ output for every device.

2.5.6. Device tracking and return accountability

Strand mount radio devices are not CPE, nor are they required for proper operation of the cable plant itself, however they can have a potential for downtime and maintenance issues. Identify inventory management requirements and internal SLAs not only for swap or repair of defective equipment on plant, but also for the technician to return defective devices to the warehouse for repair or RMA.

Enforcement of accountability for devices that are not in the warehouse or active on the plant can assist in the timely return of defective devices to the warehouse while still under warranty. This reduces stranded equipment on technician's trucks and improves tracking and distribution of spare devices.

2.6. Planning for bi-directional communication

Many times, communication plans for new product deployments are heavily focused on top-down communication – or what is colloquially known as “shout down communication” – which typically conveys information on the service, as well as deployment, care and support training to the local market, with little opportunity for upward feedback or participation in deployment planning. The importance of bi-directional communication is often overlooked in planning for deployment of a new device or service.

Direct and immediate feedback between the local market and the corporate service owner is crucial, both for initial deployment and ongoing operational support for the life of the service. Appointing a service champion in the local market to manage the deployment and serve as the local market advocate to the corporate service owner can significantly improve time to market, speed of deployment and communication of issues identified in the field.

Regardless of the effort and diligence put into initial trials and deployment plans, there are always unforeseen issues that will impact the service over time. The local market's ability to identify issues and directly communicate them upward to the corporate service owner should be included in the communications plan. Benefits of bi-directional communication include:

- Rapid identification and communication of issues – Where the local market has responsibility for installation and maintenance, the local operations center can observe and report on issues that are not readily visible at a centralized service management level. Software or hardware errors with a device or tool can manifest as the density of devices deployed in a local market increases. What may have been considered as a one-time anomaly during an initial trial may affect a demonstrable percentage of devices as field deployments increase.
- Recommendations for device, tool or process improvements – The local market may have insights that can provide a positive impact on customer experience, as well as recommendations that enhance the service or increase tool functionality. For example, a recommendation to update the ticketing system with a clickable map link with the latitude and longitude to route a technician to the exact location of a device on plant can result in quicker outage resolution, reduced

downtime, and, in areas where devices are densely deployed, improved accuracy when rolling trucks to the appropriate location.

2.7. Measurements, metrics and uptime

An issue that often gets overlooked or deliberately deferred until a significant number of devices have been deployed is the defining of the local market's level of accountability and response time for deploying and maintaining devices on the plant that support the service. As previously stated, radio devices and HFC cable modems are neither CPE nor critical to the function of the cable plant itself and can thus fall outside of the normal metrics by which a market is measured. This can lead to delayed responses to outages, and failure to correct non-functional devices on plant. Accumulated over time, this can impair service availability and disaffect the customer experience.

Failing to include device availability (uptime) as a local market metric by which they are measured can introduce multiple problems, such as:

- Tickets opened by the local operations center, but not worked in the field, can lead to multiple tickets for the same issue. Without action in the field, the operations center can become desensitized to issues over time. This results in a failure to create tickets or jobs for resolution.
- Without measurement and accountability, there is no incentive for technicians to swap or repair non-functional devices. Technicians and maintenance staff will focus on what they are measured on, in order of priority.
- Devices can be deliberately disconnected from plant weeks or years after install as a “temporary” fix to resolve powering issues related to non-related plant extension or maintenance. Once disconnected, there is no incentive to bring the device back online.

2.8. Operational Support

Garnering effective operational support in the local market is key to maintaining service reliability and a consistent customer experience. Typically, the primary role of local market operation centers is to validate the device at the time of installation, update the workflow manager, and schedule any swap or replacement activity. In addition to basic offline reporting, additional information can be provided by Tier 2 / Tier 3 support to assist the local operations center in prioritizing and resolving issues, which prevents offline devices on the plant from accumulating over time. Examples of operational reporting that can be used to support multiple device types and different services are provide next.

Aging reports (see Table 3) can identify both the number of devices offline in a given market, as well as the average days a device is offline, and the longest time a device in the market has been offline. The actual impact to the service can be deceptive if the number of offline devices isn't also associated with the duration they have been offline.

Table 3 – Device Offline Aging Report

Market	Total Devices	Devices Offline	Percentage offline	Average days Offline	Max Days Offline
Florence	5,892	32	0.54%	3	7
Rome	4,795	40	0.83%	6	10
London	6,400	100	1.56%	1	1
Paris	500	22	4.40%	20	200

In the example in Table 3, London has a relatively large number of devices offline, but they have only been down for a day, which can point to a weather event or other issue that caused an outage that affected a number of devices. Paris, while having the fewest offline devices, has the highest percentage of offline devices, takes a longer amount of time to correct an offline condition, and has even had a device offline for 200 days.

Table 4 – GPS Delta Report

Site Name	Delta in feet	GPS Lock?	Site Lat/Long	Device Lat/Long	Device Serial	Device CM MAC	Device Radio MAC
A1	3	Yes	42. 1111, -92. 944	42. 1111, -92. 943	abc123	ef:ba:01:3c:58:56	ab:06:45:8f:db:50
A2	1,000	Yes	42.8055, -92. 666	42. 8060, -92. 672	abc456	ef:ba:01:87:a4:29	ab:06:45:90:5d:50
A3	5	Yes	42. 8333, -92. 555	42. 8334, -92. 556	abc789	d2:89:3d:d0:45:29	ab:06:45:9a:8f:74
B4	20,000	No	42. 4166, -92. 555	42. 5267, -92. 555	def123	ef:ba:01:1a:48:a6	ab:06:45:7d:ec:f4
B6	2	Yes	42. 5277, -92. 888	42. 5278, -92. 888	def456	ef:ba:01:44:05:fc	ab:06:45:7c:67:68

Maintaining accurate device placement records – Comparing the latitude/longitude of where the site was designed in the workflow manager, to the latitude/longitude as reported by the device, can provide information to keep the device correlated with the proper location on plant. In the example in Table 4, Site B4 appears to be 20,000 feet away from where the device was installed. However, the device does not have a GPS lock, so the coordinates provided by the device may be invalid (insufficient line of sight, or a previously cached GPS value from a prior location). Site A2 may indicate a device that was installed at the wrong location.

Similar techniques can be used to identify active devices on plant that are not associated with a location in the workflow management system, or devices associated with a site in the workflow manager that are in an improper workflow state (e.g., sites under design should not have a device associated to them).

3. Deployment management

While the creation of new services and products are typically developed at the corporate level, physical deployment usually depends on the local market to install and manage devices on plant. Given that radio devices are neither CPE nor network equipment necessary for the plant to function, the importance and the business impact of an outage to these devices may not be apparent to local market leadership.

3.1. Identify service champions for each local market

Appointing a service champion with the responsibilities shown in Figure 4 at the local level is one of the most effective ways to simplify the deployment and ongoing operational support of wireless devices on

cable plant. Ideally, a service champion reports to the corporate service owner, but works physically within the local market to manage all aspects of the deployment and facilitate direct communication between the local market and the corporate service owner.

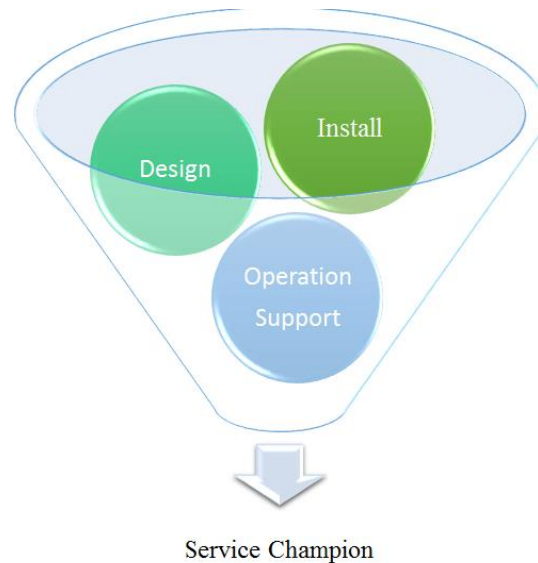


Figure 4 - Managing Deployment with a Service Champion

The establishment of a local service champion provides several benefits:

- Accountability at the corporate level for executing a standardized deployment methodology over the operator's footprint, across multiple markets. This leads to a consistent customer experience across the footprint.
- Understanding of and compliance with local market design, construction, installation and change management standards and processes, while continuing to manage the deployment within corporate guidelines.
- Conducting and coordinating training activities for all functional roles necessary for the initial deployment and ongoing operational support of the service.
- Coordinating and managing the daily design, construction, installation and post-install operational support activities in the local market, to meet corporate objectives and the deployment plan.
- Rapid escalation of issues identified in the field to corporate for resolution.
- Direct lines of communication between the local market leadership and the corporate service owner, informing and updating the local market leadership about the service and serving as an advocate for the local market to raise concerns and support local market needs (e.g., prioritizing deployment in specific areas to offset local competition, and addressing issues and concerns from legal, governmental affairs, public relations or marketing representatives, or supporting special requests).

3.2. Identify local market challenges

There are several areas to consider in preparing the local market for deployment of any service with a wireless component:

3.2.1. Local market legal considerations

These can vary depending on the type of service and radio device, how the device is deployed or the point of demarcation between the operator and the customer, for services such as small cell. Pole owners have concerns about RF-emitting equipment on plant, and the risk of exposure to their workers. Tagging the device or drop cable to the device with a warning label such as the one shown in Figure 5 alerts utility workers that a radio device is in proximity.



Figure 5 - RF Device Notification Drop Tag

Some questions to consider:

- Do local pole attachment agreements specifically exclude installation of RF emitting devices?
- Does the local utility or municipality require permits or recurring fees associated with installation of RF emitting devices on plant?
- Is a non-ionizing electromagnetic radiation (NIER) report required by the local pole owner or utility?
- Does the pole owner require that the radio device be clearly identified with a tag or label as in the example in Figure 5?

3.2.2. Coordination with other programs in the market

It's very important to know what other programs may be ongoing that impact the cable plant. Many operators are executing on initiatives to upgrade plant and push fiber deeper into the network, and a sizeable number of municipalities are requiring plant to be relocated to back-alley or buried underground as part of urban renewal or beautification projects.

Without knowing what other initiatives are ongoing within the local market, it's possible to install a device at a site one month, only to have another program move or bury the plant the following month -- rendering the site no longer capable of supporting the service. Cross-program communication and coordination can avoid this type of confusion by managing internal competition for design, construction and installation vendors across different programs.

3.3. Segment markets for parallel deployment

Like states are divided into counties, a local market can be segmented into groups of contiguous geographical areas, or sub-market groups according to local market needs, as shown in Figure 6. This provides flexibility in allowing the local market to participate in the prioritization of deployment, while facilitating a more rapid deployment, in an organized, parallel manner.

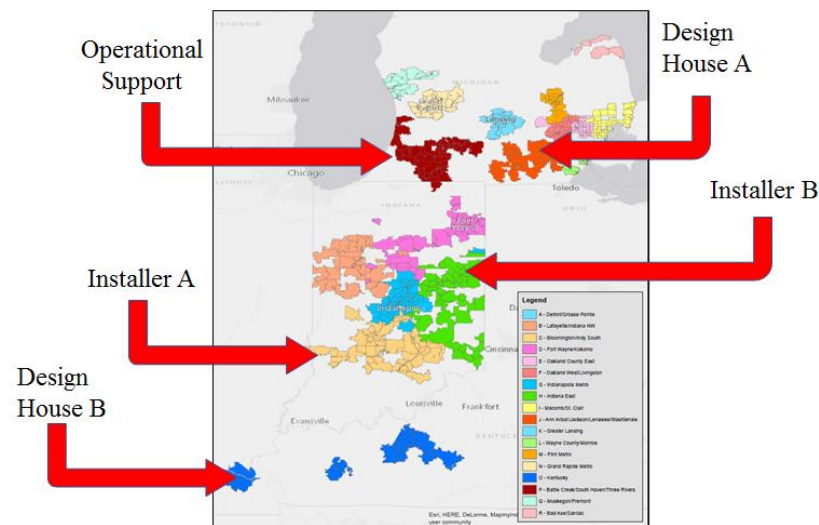


Figure 6 - Market Segmentation for Parallel Deployment

In the example above, the local market was segmented into sub-market groups using contiguous ZIP Codes. Each group was prioritized in a sequential order by local market leadership with support from the service champion. Within each sub-market group, a ZIP Code prioritization was established to structure and manage the deployment process within each sub-group and across the local market. The example also shows the sequential order of work, from design to install to ongoing operational support in completed areas.

Structuring the deployment in this manner allows the local market to predict where and when deployment activities will occur, simplifying coordination with other programs and plant activities.

Using public data sets such as ZIP Codes or counties, as opposed to data sets internal to the cable operator (e.g. headend or node boundaries) to segment the market provides several advantages:

- Cable plant is under constant change. Using shifting data or data sets that are not current increases the complexity and difficulty of managing a deployment against a changing geography.
- Using public data to segment the local market allows correlation with competitive data and supports metrics for future reporting needs (e.g., 100 access points installed in Anytown, USA). By contrast, reporting by node would change with every node split, while the geography of a headend changes with every plant extension and expansion into newly developed areas.

3.4. Training Considerations

Before start of deployment, conduct training for staffing in all functional roles. For deployment of wireless devices on plant, a “train the trainer approach” will not result in clean execution. Direct training sessions, in person, or at a minimum via videoconference, are necessary if the local market deployment is expected to meet corporate guidelines for the service.

The installation of wireless devices on plant introduces a requirement that may not be typical of other deployments on plant: Device design and placement at the exact location previously specified in the site identification process. This should be a heavily repeated theme across all training sessions.

- Locations for device placement should be identified during the site selection process. The importance of exact location accuracy when designing, installing or swapping a device at a site cannot be over stated. Designing or installing a device as little as 100 feet from the site specification can prevent the radio device from providing service to the intended target location.
- Design and placement should be by the latitude and longitude from the site specifying where the device should be placed on plant, not by the closest approximate service address.

Refresher training should be conducted on a quarterly basis. Staffing does change over time, and training is an ongoing effort that should be sustained for the life of the service, not just for the initial deployment. Functional roles for deployment training include:

- Design house
- Construction
- Installation
- Plant maintenance
- Operations center

Plant maintenance can often be overlooked or only trained after devices have begun deployment on plant in volume. Maintenance technicians and operations support need to understand the device, installation and swap processes before the first device is placed on plant in their area. Training logistics need to account for at least two sessions for design, construction and installers.

If the operations center is a 24-hour operation, all three shifts should be specifically trained, not just the day shift. Plant maintenance technicians are dispersed over the footprint, and the training schedule should account for this. Direct and periodic training at the tech level is essential to proper maintenance of the service over time.

4. Summary of Operational Efficiencies Gained Through Common Practices

Deployment of radio devices on cable plant introduces new requirements and processes at the local market level, but with advance preparation, these processes are scalable and repeatable across different service lines and different radio device types and configurations.

Establishing and following common practices and a common methodology for deployment of radio devices on cable plant eliminates duplication of tools, reduces separate processes across different

service lines, and allows for a more rapid deployment and a consistent model for ongoing operations for similar services at the local market level. And perhaps most importantly, establishing these practices can reduce downtime and thus improve customer satisfaction, both for existing customers as well as the new wireless services customers. A summary of common practices to be used is as follows:

- **Use a common work flow management system across multiple service lines** – Regardless of the wireless service deployed, there are common steps all will need to follow for deployment on plant. Deployment and operational support processes are driven to a great degree by the workflow management system and can be standardized to streamline deployment efforts and operational support across local markets.
- **Device Management** – Tightly manage the correlation of radio and cable modem MAC addresses, whether in an integrated device or separate devices connected via PoE. This avoids disconnects in troubleshooting transport and service issues. Specifically:
 - Device information should include any virtual addresses to associate a customer-reported issue to a specific device on plant.
 - Devices should have visual indicators to confirm if the device is functioning properly.
 - Devices should have an internal GPS to help identify where the device is located on plant.
 - Two-device solutions (separate HFC cable modem and radio device) should be treated as a single unit in the field. Any assembly or repair should be conducted at a central or third-party location, not at time of install or swap.
 - Stage devices at a central location prior to deployment:
 - Label the bottom of the device with the last four to six characters of the serial number in non-metallic mailbox letters that can easily be seen from the ground.
 - Boot the device to upgrade to the current version of code or firmware. Any non-functional devices can be returned for repair prior to going to the field for install.
 - Cover exterior labels with clear, UV resistant weatherproof tape. It will extend the legibility of serial numbers and MAC addresses over a longer period of time. Weathered, illegible labels can prevent a functioning device from being re-deployed, or properly accounted for in inventory.
 - Configure the device for automated provisioning as a piece of network equipment, not as CPE.
- **Site information** – Information for a specific site should include public data such as county, ZIP Code, city and state as well as information internal to the operator. This allows for easier operational reporting and correlation with competitive or other public data to predict trends and feedback into site selection processes.
- **Tools** – A common set of tools should be utilized to manage the deployment and confirm that devices are active in near real-time, and that the service is functioning properly over the device at time of install or post-install swap. The same tools and associated processes should be able to support multiple wireless services deployed on cable plant.

- Tools should be able to compare where a device is designed on plant, and any delta between the intended location and where the GPS in the device is reporting.
- Tools should be able to regularly poll and record the last known GPS reading of a device and whether the GPS was in a “locked” state (which impacts whether the coordinates provided by the device are accurate)
- **Device identification at time of install** – Whether a swap or a new installation, the technician should call the operations center while on site. The operations center staff will associate the device with the site and validate that the device is functioning properly before releasing the tech from the site.
- **Deployment Planning** – Use initial field trials to identify processes for production deployment, gather information for training documentation, assess existing tools for applicability and to identify gaps. Determine if application development is needed or if gaps can be addressed through additional process definition.
 - Identify functional roles and responsibilities, and where those roles reside, either corporate or the local market (e.g., site selection, design, install and monitoring).
 - Document the time span between when power is supplied to the device to when the device shows online in operational monitoring toolsets. This is required for operations center, installer and maintenance training.
 - Determine the actual power consumption of the device – Conservatively designing to full PoE power output may exhaust the power budget. Determine the actual draw of the device(s) under rest and under load, in real world conditions, then confirm actual power recommendations for design.
 - Plan to manage inventory not only in warehouse or on plant, but also equipment techs returned from swaps, and where spares are distributed. Enforce accountability for equipment and establish internal SLAs for prompt return to the warehouse for repair.
 - Communication plans should be bi-directional and direct. The local market should be able to escalate issues directly to the corporate service owner and provide feedback to improve the service.
 - Establish uptime metrics and accountability in the local market. Failure to address this in initial planning stages will result in delays in repair or replacement of offline devices.
 - Provide operation centers with tools and reporting necessary to manage the service, including offline reports, offline aging, and site vs. device GPS delta reporting.
- **Deployment Management** – Identify a service champion to lead and manage deployment efforts in the local market and serve as an advocate to corporate for local market needs and requests, to:

- Identify and escalate local market challenges for corporate awareness and accommodation. Be aware of any pole attachment restrictions, permitting or recurring fees from the pole owner or local municipalities.
 - Be sensitive to RF exposure concerns by utility workers and understand the MPE limits for each type of wireless device type being deployed on cable plant.
 - Label devices or cable drops with RF warning tags and a phone number for the pole owner or utility to contact.
- Understand the impact of other programs ongoing in the local market and cross-coordinate efforts where possible.
- Segment the local market into smaller geographical groups to help prioritize areas and manage the deployment in parallel across the market.
- Develop training and educational materials such that they address all functional roles, from local market leadership to design, construction, installation, maintenance and operations. Refresher training should be periodically conducted in the local market as staff and the service change over time.

5. Conclusion

Many operators are actively deploying or trialing multiple wireless services and technologies to provide new services to new customers, as well as enhancing the reach and value of existing services to their traditional customer base. However, radio devices and HFC cable modems do not represent customer premise equipment, nor are they critical to the function of the cable plant itself. For those reasons, they can often exist outside of the normal metrics by which a market is measured.

A great deal of attention and focus is typically placed on each individual service, and the specific radio devices required to support each unique wireless service. As important is advanced planning and development of a consistent deployment strategy, as well as a set of common operational practices. Such a coordinated tactical strategy can eliminate multiple variations of similar tools and support processes for wireless services in the field, resulting in a shorter time to market, greater service uptime and operational processes that can scale effectively across multiple service lines and markets.

Preparing for the Network For 5G And IoT

A Technical Paper prepared for SCTE•ISBE by

Fady Masoud, M. Eng.

Principal, Product and Technology Marketing
Infinera

555 Legget Drive, Suite 222, Tower B, Ottawa, ON, Canada K2K 2X3

fmasoud@infinera.com

Table of Contents

Title	Page Number
Table of Contents	93
1. Introduction	94
2. What Is Cognitive Networking?	94
3. Paving the Way for Cognitive Networking	95
3.1. Evolve the Network Architecture	95
3.2. Unlock the Network's Full Capacity Through Super-Channels	96
3.3. Leverage Software Defined Capacity	97
3.4. Automate Network Operations Through Software-Defined Networks	98
4. Cognitive Networking – Not in a So Distant Future	98
4.1. Maximized Capacity/Reach Ratio	98
4.2. Analytics and Machine Learning-triggered Network Capacity	99
4.3. Proactive Network and Traffic Protection	99
5. Conclusions	100
6. Abbreviations	100
7. Bibliography and References	101

List of Figures

Title	Page Number
Figure 1 - Elements of a Cognitive Network	95
Figure 2 - Evolution from OSI to Layer T and Layer C model	96
Figure 3 - High Capacity Through Super-Channels	97
Figure 4 - Software Defined Capacity	97
Figure 5 - Analytics and Machine Learning-Triggered Network Capacity	99
Figure 6 - Proactive Network and Traffic Protection	100

1. Introduction

The Internet of Things (IoT) is shaping our day-to-day lives where it is estimated that 30 billion devices¹ are expected to connect to the Internet in 2020. Moreover, cloud-based applications are also changing the landscape of today's enterprises, from the products they manufacture and the services they offer to the way their employees interact with each other, or with customers and partners. As a matter of fact, enterprise applications are doubling every 2.5 years and the global cloud traffic is expected to increase almost fourfold between 2015 and 2020². Moreover, the world is on the verge of massive deployment of 5G mobility, which is poised to radically transform our mobile connectivity and start a new era of high-performance mobile applications and machine-to-machine real-time communication. The deployment of 5G networks will dictate an unprecedented level of performance from the underlying optical transport networks, such as ultra-low latency, network sliceability and scalability. This fast-paced migration to the cloud, forthcoming 5G deployments, and the proliferation of Internet-connected devices driven by IoT are fueling the migration toward a new breed of intelligent and highly autonomous optical networks called cognitive networks. They use advanced analytics, machine learning, and artificial intelligence techniques to help build self-optimized, self-healing and highly autonomous transport networks, setting new benchmarks in scalability, agility and automation. This paper describes the journey to cognitive networking by explaining its key building blocks, such as software defined capacity (SDC), multi-layer software-defined networking (SDN) control, and transport technology breakthroughs that make the creation and the deployment of cognitive networks a fast-approaching reality.

2. What Is Cognitive Networking?

Cognitive networking is the ultimate goal for the intelligent transport layer that underpins all cloud-based digital communications. By definition, a cognitive network is multi-layer, self-aware, self-organizing, and self-optimizing, and can take predictive and/or prescriptive action based on what it has gleaned from its collected data and experience. Realistically, no network can completely plan or run itself; however, cognitive networking will dramatically reduce the number of manual tasks required across a multi-layer network. This goal can be achieved by leveraging advanced software, streaming telemetry, big data with machine learning, and analytics to autonomously conduct network operations to meet the demand for connectivity, maximize network resources, and increase reliability. There are multiple important elements in a cognitive network (see Figure 1), such as:

- Advanced analytics designed to parse streams of machine data to monitor network health and raise awareness of any anomalies
- Machine learning software tools that leverage advanced analytics to understand and identify trends in operations
- Autonomous hardware and software capable of executing various tasks and conducting required maintenance
- Predictive intelligence tools capable of identifying potential problems before they happen
- Prescriptive software tools designed to proactively recommend new solutions for maximizing capacity, enhancing reliability, and optimizing assets.

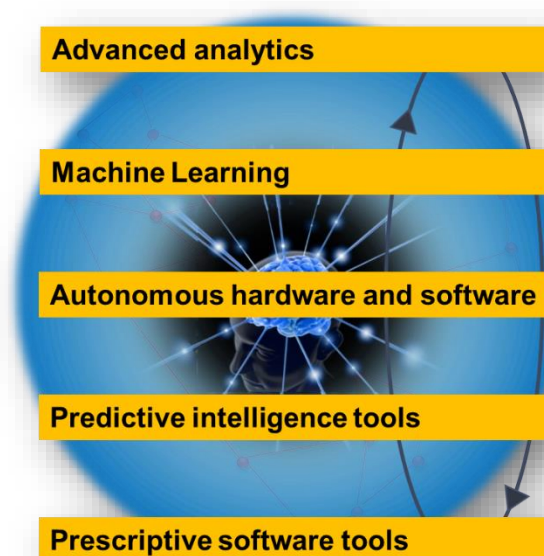


Figure 1 - Elements of a Cognitive Network

Cognitive networking is the result of seamless and highly dynamic interaction between software and hardware assets across network layers, and brings optical networking to a new level of scalability, flexibility, and automation. The following section describes how to build the foundation for cognitive networking by using the latest technology breakthroughs.

3. Paving the Way for Cognitive Networking

The journey to cognitive networking starts by building the foundation of a highly scalable, flexible, and programmable network architecture. The building blocks for this foundation are described in the following:

3.1. Evolve the Network Architecture

A well-defined architecture dictates how networks are planned, operated, and evolved. Today, when content is king and must be accessible anywhere, anytime, and on any device with the highest level of quality, it is clear that the 1980s-era seven-layer Open Systems Interconnection (OSI) model has reached a tipping point. It needs to support the transformation in networks (e.g., network functions virtualization [NFV], SDN, etc.) and the new service delivery model based on cloud applications, service virtualization, etc. The OSI model's heritage of function- and layer-specific network definitions led to closed and proprietary protocols, rigid networking capabilities, and high operational costs. These limitations have sparked an urgent need to evolve toward a simpler, more efficient and agile architectural model to accelerate the adoption of cloud-based networking. Thus, the first step in paving the way for cognitive networking consists of evolving the network architecture to a simpler cloud-powered model: a cloud services layer, Layer C, and an intelligent transport layer, Layer T, as depicted in Figure 2. This new model consolidates and simplifies cloud service delivery and networking into two layers, wherein all the OSI networking layers (Layer 3 and below) are represented by Layer T, while all the application layers

(Layer 4 and above) are grouped under Layer C. Layer T sets the guidelines and principles for the transport of data streams, whether between end users and data centers or among multiple data centers with bursty and often unpredictable traffic patterns. It also defines the features and capabilities that increase network agility and performance and sets new benchmarks for service delivery and cost-effectiveness, key ingredients to the successful deployment of any cloud application. Layer C contains all the applications, functions, and services that run in the cloud, including consumer and business applications, SDN-based service creation and orchestration tools, software frameworks and applications for big data and machine learning, virtualized network functions (VNFs), and many others, to enable the large-scale task automation and programmability that streamline operations, eliminate human error, and reduce operating costs.

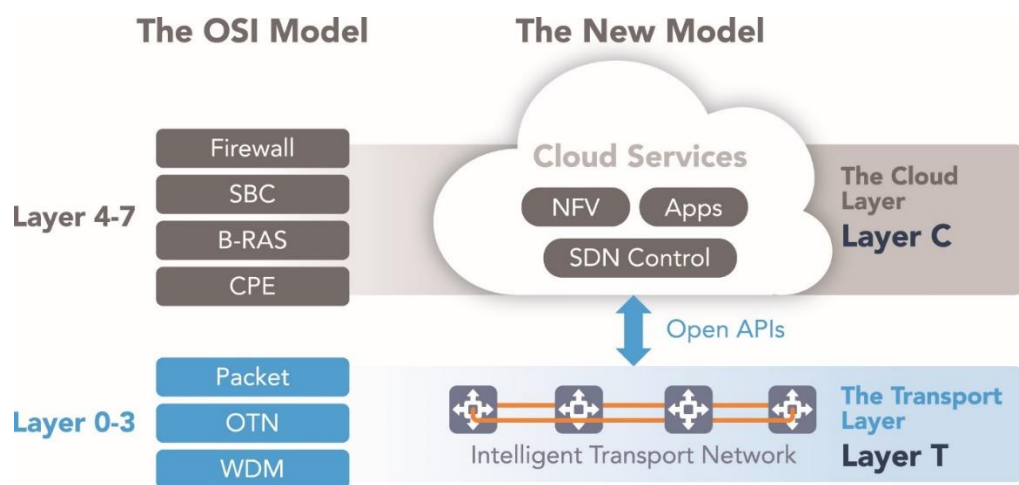


Figure 2 - Evolution from OSI to Layer T and Layer C model

3.2. Unlock the Network's Full Capacity Through Super-Channels

DWDM technology disrupted the telecommunication industry by enabling multiple optical carriers to travel in parallel on a fiber, thus increasing capacity and maximizing fiber utilization. However, the current growth in Internet traffic and enterprise migration to the cloud is demanding a new level of scalability. An innovation, called super-channel, evolved to take DWDM networks to a new era of high capacity and optical performance, all without increasing operational complexity. A super-channel includes several optical carriers combined to create a composite line-side signal of the desired capacity that is provisioned in one operational cycle, as depicted in Figure 3. Super-channels overcome three fundamental challenges:

- optimizing DWDM capacity and reach
- scaling bandwidth without scaling operational procedures
- supporting next-generation high-speed services such as 100 gigabit Ethernet (GbE), 400 GbE, etc.

The use of super-channels increases spectrum efficiency and thus network capacity by reducing spectrum waste due to guard bands. It also enables seamless capacity growth without the need for network re-engineering or major disruption to current operating processes.




Figure 3 - High Capacity Through Super-Channels

3.3. Leverage Software Defined Capacity

A key steppingstone toward cognitive networking is to break away from the current methods of optical capacity planning, engineering, and hardware-based deployment that require numerous truck rolls, extensive manual labor, and human interaction at multiple points in the network. The road to cognitive networking starts with allowing intelligent software tools to dynamically add, modify, move, and retire optical capacity based on the real-time requirements of upper-layer applications (Layer C), as depicted in Figure 4.

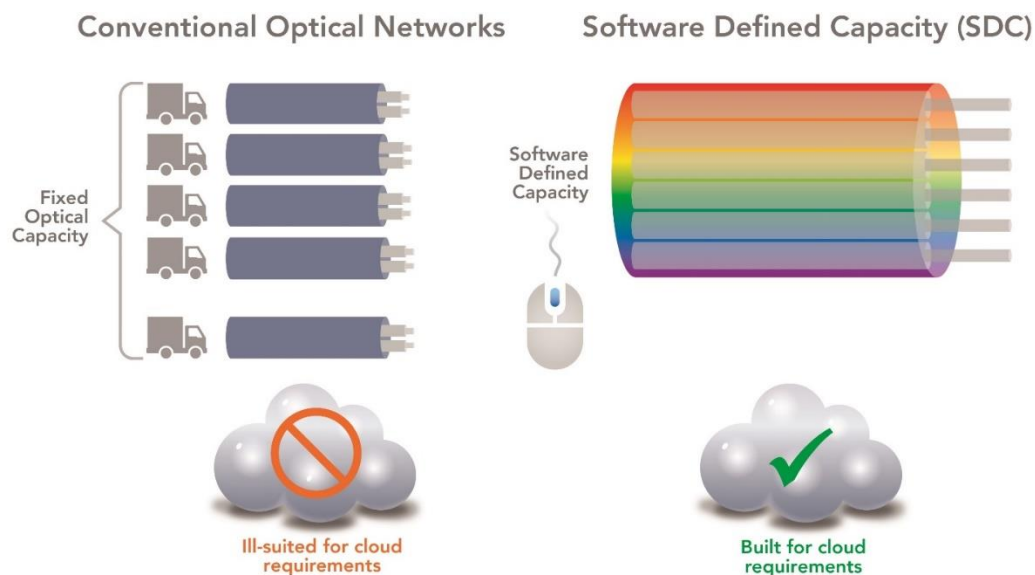


Figure 4 - Software Defined Capacity

SDC provides instant software activation of additional capacity, creating a pool of bandwidth that can be dynamically allocated based on traffic demand. SDC extends the principles of SDN, which has primarily focused on Ethernet and packet layers, to the optical transport layer. With intelligent software tools, a

network can become an integral part of the rest of the information technology infrastructure, enhancing service turn-up and management.

SDC is a true game-changer from both business and operational perspectives. It enables a perfect match between the timing of capital expenditure (capex) and service revenue, thus accelerating time-to-revenue from months to minutes. SDC also reduces operational expenditure (opex) by streamlining operations and eliminating truck rolls. Moreover, SDC is a key enabler of automation throughout the network and across all operational levels, which is a vital element in building the foundation for cognitive networking.

3.4. Automate Network Operations Through Software-Defined Networks

Building automation and intelligence across all network layers and various operating tasks is central to cognitive networks. Multi-layer SDN controllers and frameworks span both Layer C and Layer T to act as the brain of the network by providing cloud-based intelligence to plan, monitor, and conduct various network operations without human intervention. SDN frameworks also provide various ready-to-deploy network applications, such as real-time capacity planning, bandwidth on demand, network virtualization, and many others, to unlock full network potential and take advantage of the dynamic transport layer (Layer T). For example, very sophisticated algorithms and data models can be used to build a microservices-based path computation engine (PCE). The PCE replaces manual offline route and capacity planning processes with highly automated, real-time service planning and activation over optimal routes across both layers, to overcome multiple and often-challenging fiber impairments. As cognitive networking relies heavily on streaming telemetry, big data with machine learning, and analytics to learn and adjust, it is crucial to ensure a seamless flow of information between the various parts of the network elements and the upper layer software tools and SDN controllers. This flow is enabled by open application programming interfaces (APIs) such as RESTCONF, NETCONF/YANG, and other northbound interfaces that connect to upper-layer orchestration systems or same-layer intelligent tools and scripts to coordinate and optimize resources across the network as well as to orchestrate VNFs. This bidirectional flow of data serves as the bloodstream of the network, delivering predictive and prescriptive real-time recommendations and taking actions to enable maximum performance.

4. Cognitive Networking – Not in a So Distant Future

As discussed earlier, many key building blocks toward cognitive networking have already implemented across various parts of the network and its operating processes, thus making a steady and well-paced progress toward the end-goal of a completely autonomous network. Nonetheless, and as technology evolves, new functions and capabilities will further accelerate this journey. A few examples are mentioned below:

4.1. Maximized Capacity/Reach Ratio

Despite the fact that today's WDM line transponders offer the ability to change the modulation scheme (such as QPSK, 8-QAM, 16-QAM, etc.) through software, it is still process that network planners and operators must manually perform to maximize the capacity/reach ratio. In other words, and select the modulation schemes with the higher capacity that can satisfy link budgets (reach). Cognitive networks will be able to parse various real-time network data points such as required bandwidth, fiber impairments, link spans and many others to automatically adjust modulation schemes and – data rates to maximize capacity over each span. Finer granularity of WDM line rate can also be achieved through the implementation of hybrid modulations of super-channels and subcarriers. With the ability to program

different modulation and dynamically adjust constellation shaping gains on each subcarrier, a variety of spectral efficiencies can be derived [3][4]. This provides higher fiber capacity and a better flexibility for diverse network applications.

4.2. Analytics and Machine Learning-triggered Network Capacity

While SDC has already forever changed the way services are planned, provisioned and activated, cognitive networks will allow real-time network analytics, microservice-based engines and machine-learning algorithms to dynamically increase or decrease network capacity on specific routes based on past trends, spontaneous changes in traffic demand or an anticipated spike in capacity. This data-driven real-time traffic engineering will be fully autonomous, programmable and proactive to identify potential sources of failures before they happen and take the necessary steps to maintain the network at the highest levels of reliability and efficiency (Figure 5).

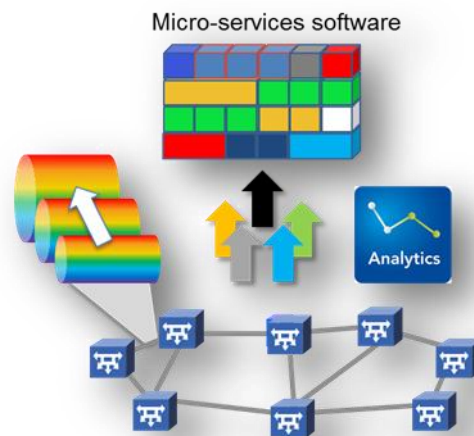


Figure 5 - Analytics and Machine Learning-Triggered Network Capacity

4.3. Proactive Network and Traffic Protection

The use of control plane has significantly increased service and network availability by automatically restoring traffic after a failure over alternative paths while maintaining SLAs. Moreover, Layer 1 or Layer 2 traffic encryption embedded into optical transmission platforms has proven to be an effective way to protect in-flight data from intruders and hacking tools. Cognitive networks will evolve traffic protection by combining control plane capabilities, real-time traffic and network topology engineering as well as traffic encryption – all functioning as parts of an intelligent and highly proactive network protection mechanism. Advanced, accurate and fast intrusion detection capabilities could trigger proactive measures to virtually isolate suspicious network areas and automatically encrypt the traffic over specific links (Figure 6). The same proactive capabilities can also be applied to minimize the impact of network maintenance operations by automatically and proactively isolating affected network areas and setting up backup plans during the maintenance activities.

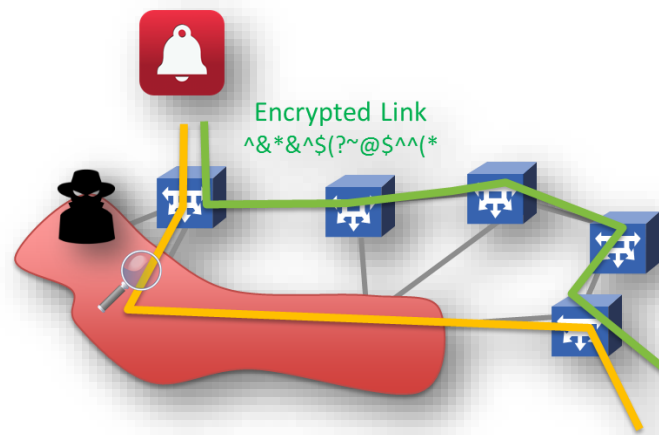


Figure 6 - Proactive Network and Traffic Protection

5. Conclusions

It is not science fiction that today's hyperconnected world is fueling the fast-paced evolution toward cognitive networks that are self-optimized, self-healing, and highly autonomous. These cognitive networks will set new benchmarks in scalability, agility, and automation dictated by the cloud. This journey starts today by laying down the foundation for cognitive networking by evolving the network architecture toward Layer T and Layer C, unleashing full network capacity through the use of super-channels, accelerating service operation through game-changing SDC, and automating tasks throughout the network using SDN controllers. The integration of these next-generation hardware and software tools is making the deployment of cognitive networking a fast-approaching reality.

6. Abbreviations

AI	artificial intelligence
API	application programming interface
CapEx	capital expenditures
DWDM	dense wavelength division multiplexing
GbE	gigabit Ethernet
IoT	Internet of Things
NETCONF	network configuration
NFV	network functions virtualization
OpEx	operational expenditures
OSI	Open Systems Interconnection
OTN	optical transport network
PCE	path computation engine
REST	representational state transfer
SDC	software defined capacity
SDN	software-defined networking
VNF	virtualized network functions
YANG	yet another next generation

7. Bibliography and References

[1] <http://www.mckinsey.com/industries/high-tech/our-insights/the-internet-of-things-sizing-up-the-opportunity>

[2] <http://siliconangle.com/blog/2016/11/11/global-cloud-traffic-to-increase-by-3-7-fold-by-2020-cisco-says/>

[3] *Design Considerations for a Digital Subcarrier Coherent Optical Modem*; David Krause, Ahmed Awadalla, Abdullah S. Karar, Han Sun, Kuang-Tsan Wu, OCF 2017

[4] *Subcarriers having different modulation formats*; US patent pending; A. Awadalla, et al.

The Future of Operations: Building a Data-Driven Strategy

A Technical Paper prepared for SCTE•ISBE by

Sebnem Ozer, Ph.D., Principal Architect, Comcast, SCTE•ISBE Member
1800 Arch Street
Philadelphia, PA 19103
sebnem_ozar@comcast.com
215-2868890

Sinan Onder, Executive Director, Comcast, SCTE•ISBE Member
1800 Arch Street
Philadelphia, PA 19103
sinan_ozar@comcast.com
267-2600964

Nagesh Nandiraju, Ph.D., Sr. Director, Comcast, SCTE•ISBE Member
1800 Arch Street
Philadelphia, PA 19103
nagesh_nandiraju@comcast.com
215-2861941

Table of Contents

Title	Page Number
Table of Contents	103
1. Introduction	104
2. Service Provider's Network and Service Transformation	104
2.1. Data-driven and Knowledge-Defined Architecture	106
2.2. Example Access Network	108
3. Conclusions	109
4. Abbreviations	109
5. Bibliography and References	110

List of Figures

Title	Page Number
Figure 1 - Data-driven and knowledge-defined service provider architecture	106
Figure 2 - Service lifecycle management	108
Figure 3 - Example SDN/NFV/cloud Enabled Access Network	109

1. Introduction

Software-defined networking (SDN), network functions virtualization (NFV) and cloud applications (CA) enable cable operators to deploy automated and programmable network architectures with big data telemetry and analytics. Today, common network orchestration, controller platforms and networking fabrics can integrate different distributed access technologies. This architecture opens the door for service automation and agility, self-optimization for better customer experience, cost optimization, green networking, real-time network and product planning and novel business intelligence. To reach these objectives, cable operators must become more collaborative by adopting organizational transformations in networking, services, operational support systems (OSS) and business support systems (BSS).

This transformation relies on an efficient, knowledge-defined, data-driven and collaborative service and network architecture. In this architecture, big data analytics platforms get data from both network components and end services for a unified modeling of network components and application components. Together, the network and application components represent the quality of service (QoS) and quality of experience (QoE), respectively. The analysis provides high-level relationships between programmable resources, services and products. The hierarchical control plane has the scope and context to manage resources in a continuous way, based on an operator's business policies and intent definitions.

In this paper, we discuss a common platform that can support different access technologies (e.g., DOCSIS/EPON/Ethernet/wireless) by providing:

- An overview of today's operator network and service operations to define the areas that need the organizational transformation
- Relationship models among products, services and resources based on data-driven architectures and hierarchical abstraction structure
- Key challenges and solutions for an orchestration platform with a closed-loop approach to manage the end-to-end system, based on inputs from processing of big telemetry data and intent/policy control
- Guidelines on how to apply these solutions for SDN, NFV and cloud enabled network platforms

2. Service Provider's Network and Service Transformation

In traditional systems, network resources are designed, provisioned and managed in a rigid way. These are defined by initial estimations and projections of the services, service group sizes and traffic scaling requirements. Customer equipment is shipped to the warehouse for staging and initial configuration and then shipped to the customer premises for installation. Then, the service is activated and tested, often done by a dispatched technician. Once the service is activated, billing is initiated. Services follow a predefined workflow and path over the network resources defined by a manual design process. The service assurance is enabled by poll-based monitoring of data and sometimes by manual troubleshooting. If a service needs to be upgraded, moved or deleted, some of these steps are repeated. If network resource capacity needs to be increased due to new service requirements, another manual design and network resource installation and configuration is required. Therefore, traditional systems do not support service agility, because they require complex operations for resource and service lifecycle management as well as to resolve customer experience issues.

The ultimate objective is to automate service workflows, such that a service provider can deliver a new service, or a customer can request a new or upgraded service through a portal. In those scenarios,

resources (both physical and virtual function components) are dynamically mapped, deployed and provisioned based on service characteristics and policies, services are activated and billing is initiated, all through an automated process. With the new model, end-to-end service delivery is validated at both the design-time and operations-time continuously.

To achieve these goals, service providers have started to adopt SDN, NFV, CA and digital transformation. However, a survey conducted by the TM Forum in October 2018 [1] shows that communications service providers (SPs) are struggling with business cases for this transformation and also with adopting new ways of working within their OSS/BSS.

According to the survey, which involved 160 people working for 66 service providers, “legacy OSS/BSS systems remain the biggest challenge to network transformation, according to 60% of respondents, while 56% ranked security vulnerabilities as a major issue and 52% were concerned with how long it’s taking standards to mature.” In the early stages of transformation, the first two challenges can be addressed by increasing collaboration between DevOps, NetOps and SecOps personnel. While DevOps introduces innovative and disruptive technologies for a speedy delivery, traditional NetOps functions are focused on availability and stability.

Depending on the SP’s objective, a balance may be found by applying incremental changes – for example, defining failover systems and double operations (i.e., both traditional and new monitoring systems) at the initial stages. A more aggressive transformation path may be adopted for new technologies if the impact on the existing network and services can be avoided. The third challenge depends on how the industry adopts the transformation, considering that traditional network standardization and certification initiatives were designed for purpose-built hardware and embedded software requirements with long update cycles. SDN and NFV, by contrast, require continuous development, integration and testing with focus on software interfaces and data models.

The collaboration and business models among service providers, suppliers, integrators and open networking groups are also at an early stage. Today, we see fragmentation among both the cable and telco industries related to SDN, NFV and digital transformation. Different open source groups and standard bodies work on the same concepts by creating different platforms. A good survey on standards, open source projects, and open reference architectures along with some harmonization initiatives may be found in the Linux Foundation’s white paper, titled “Harmonization 2.0: How Open Source and Standards bodies are Driving Collaboration Across IT” [2]. Similarly, CableLabs’ Access Network Virtualization Group was formed to identify standards projects and architectures that are critical to the cable industry.

In a previous *SCTE-ISBE Journal of Network Operations* paper, titled “EPON Architecture Considerations for Intent-based Networking” [3], we discussed how an SDN/NFV/cloud-enabled access network architecture provides automation and programmability – two key features for intent-based networks that offer service agility, simplified and effective operations and improved customer experience. Through a remote EPON example, we addressed the areas operators need to transform. Those include onboarding of physical and virtual components; telemetry and analytics; and dynamic and adaptive control. We concluded that the new architecture requires a well-planned roadmap of disruptive technologies and new ways of design, development, test and operations.

In this paper, we focus on how the current service provisioning and management architecture can be transformed by applying a programmable and automated network infrastructure as defined in the earlier paper [3]. Here, we discuss a data-driven and knowledge-defined architecture roadmap to enable service agility, decrease operations complexity and increase customer satisfaction.

2.1. Data-driven and Knowledge-Defined Architecture

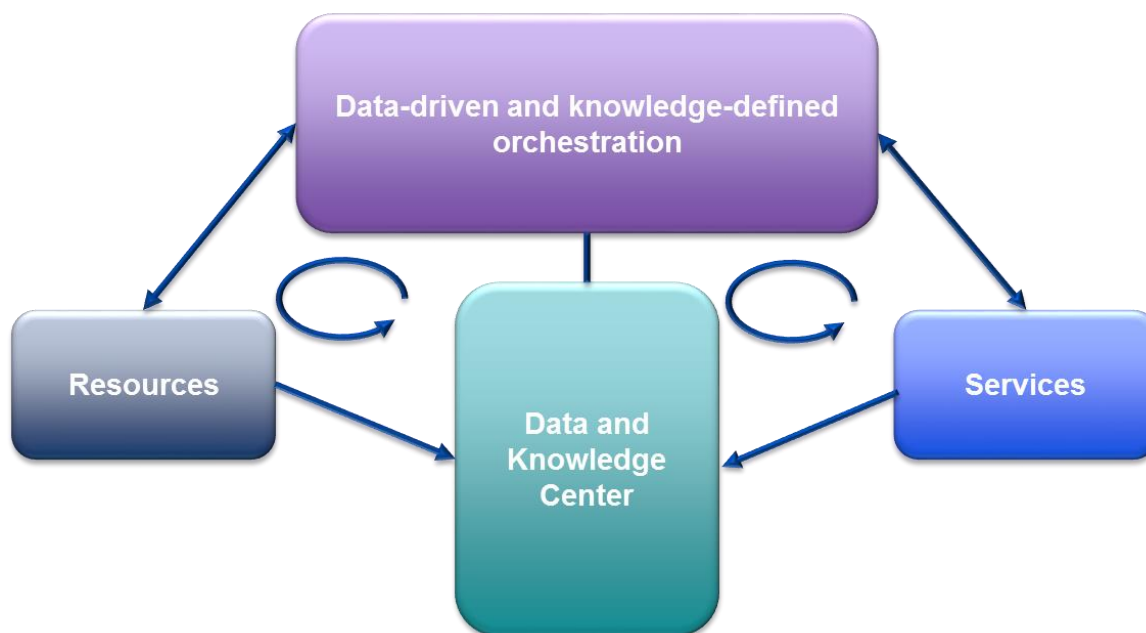


Figure 1 - Data-driven and knowledge-defined service provider architecture

Figure 1 displays a simplified view of a data-driven and knowledge-defined architecture. Resource blocks consist of physical and virtualized network and service functions and components, compute and storage nodes. Depending on the service and operations objectives, including real-time and scale requirements of the functions, resources and network components may be virtualized or implemented in purpose-built hardware. They can be located in the customer premises, the field, edge network, core network or in the cloud.

One critical design choice is the ability to reuse the resources for functionalities common to all access technologies (e.g., HFC, PON, wireless). Resources should be defined through models that are agnostic of access technology and that can be abstracted from low layer attributes. Physical and virtual resource topologies are modeled with their network status. Inventory provides a real-time, end-to-end view of those resources, along with mapped services and products. The Open Network Automation Platform (ONAP, [5]) and the Open Source MANO (OSM, [6]) are two industry initiatives that are developing such automation platform for telecommunications, cable, and cloud service providers.

Service blocks consists of a service provider's application and services catalog, used by subscribers and other customers. A design framework specifies a service with multiple attributes. These attributes define service requirements such as functionalities, policy rules to achieve target service behavior, and monitoring and alarm events for closed loop control through analytics. Resources and services are modeled with data models (e.g., YANG/ yet another next generation) and defined before the architecture

deployment, as a design process, and are updated in an automated way during the deployment and runtime of resources and services.

The orchestrator and controller domain uses design, runtime data and knowledge to map services to resources, automatically instantiate the service and provision the resources based on the service policies as well as system status (e.g., load balancing). The design and run time data is also applied to manage services and resources in an elastic manner, based on the analytics results and demand variations. Analytics monitor the services and resources, and apply techniques such as artificial intelligence and machine learning [4]. Because training is an important part of these techniques, collaboration from the operations team, especially on the corner cases, is crucial. Another important point is to create sufficient abstraction between resources and services so that a service-resource relationship is established to allow programmability and automated configuration of virtual network functions (VNFs) and physical network functions (PNFs). QoS based on the abstracted status of resources can then be mapped to QoE for the service that uses these resources, making self-optimization and self-healing feasible. Until the transformation becomes mature, the role of current NetOps and SecOps functions is significant to keep the system stable while training the analytics process for knowledge-based automation.

The mapping is related to service chaining as a sequence of interconnected network service functions, through which a data stream must pass to support an application. These functions may reside in a physical or virtual network component (i.e., PNF and VNF). In traditional systems, this is done manually in the design phase. With SDN/NFV-enabled networks, a service and resource orchestrator can define dynamic service chaining. Virtualization, augmented with digital transformation (in this sense, the digitization of core infrastructure elements and actions against them), will enable the rapid creation and deployment of microservices that facilitate service integration. In other words, service integration is improved through more flexible and simple blocks, enabling fault isolation. This also helps eliminate delays attributable to monolithic software and slow or latent release cycles to achieve rapid feature implementation.

In addition, cable operators can develop their own operational, analytical systems/services and new, value-added products more easily and iteratively. Containers are widely preferred in microservice architectures because they provide faster initiation, efficient execution and better isolation. Orchestration software such as Kubernetes or Docker Swarm can be used to automatically deploy, scale and manage containers inside a cluster. A service mesh like Istio and Linkerd can be used to describe the network of microservices that a service goes through and the interactions between them. They support load balancing, applying traffic policies, service discovery and monitoring, failure recovery, routing and secure microservice-to-microservice communication.

The transformation of the SP's architecture will lead to a service lifecycle management, as shown in Figure 2, with continuous design/development and integration through analytics-based feedback. This approach will provide the ability to generalize, harden, elastically scale and operationalize the system. As a result, service providers gain the ability to introduce new resources or services without affecting the existing platform base and operations. Modular and flexible design with reusable components for different access technologies can unify the operations and lifecycle management.

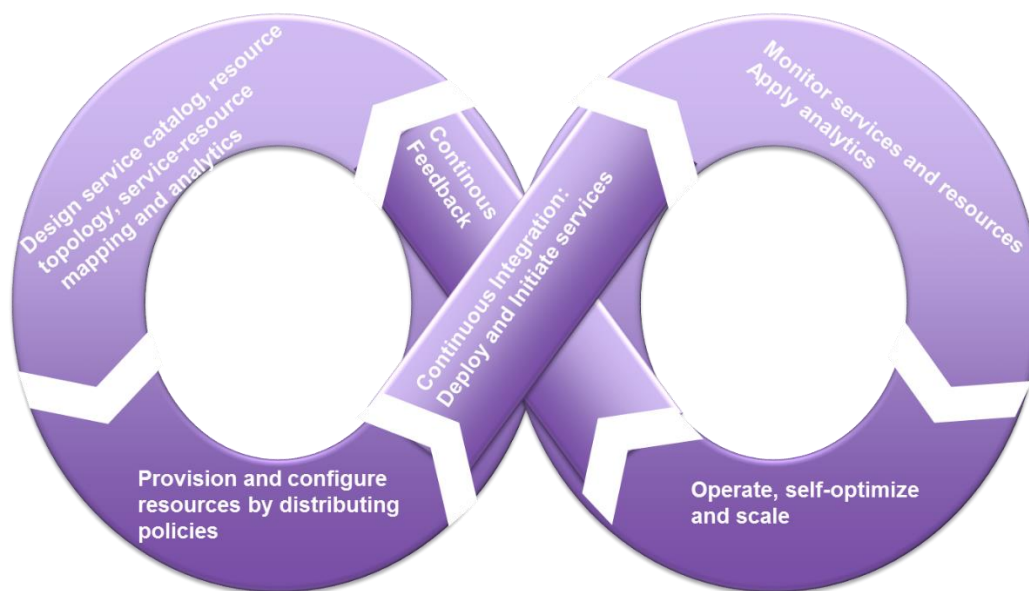


Figure 2 - Service lifecycle management

2.2. Example Access Network

Figure 3 displays an example access network that consists of physical (e.g., R-PHY, R-MACPHY, R-OLT, CM, ONU devices) and virtual (e.g., virtual router, subscriber and traffic management VNFs) network resources. The access functionalities are distributed between the distributed access nodes and virtual functions, running on common servers as containers. The control domain is separated from access nodes and implemented as VNFs. The access nodes are connected through a switch fabric that consists of whitebox switches controlled by an SDN controller. VNFs are controlled and orchestrated by an NFV orchestrator. Each resource and service is monitored based on a push-based telemetry system that collects, stores and analyzes data and that is hosted in the cloud. The service and network resources lifecycle orchestration is also hosted in the cloud and interacts with the data and knowledge center.

In this architecture, service and network orchestrator gets service information through the service catalog, resource topology and state information, along with corresponding policies and rules. The orchestrator then maps services to resources, and distributes the mapping and policies to access, NFV and underlay controllers that provision and manage the physical and virtual resources. Continuous monitoring and analytics provide feedback to the orchestrator for self-optimization, self-healing and scaling, based on services and resources performance indicators and demand changes.

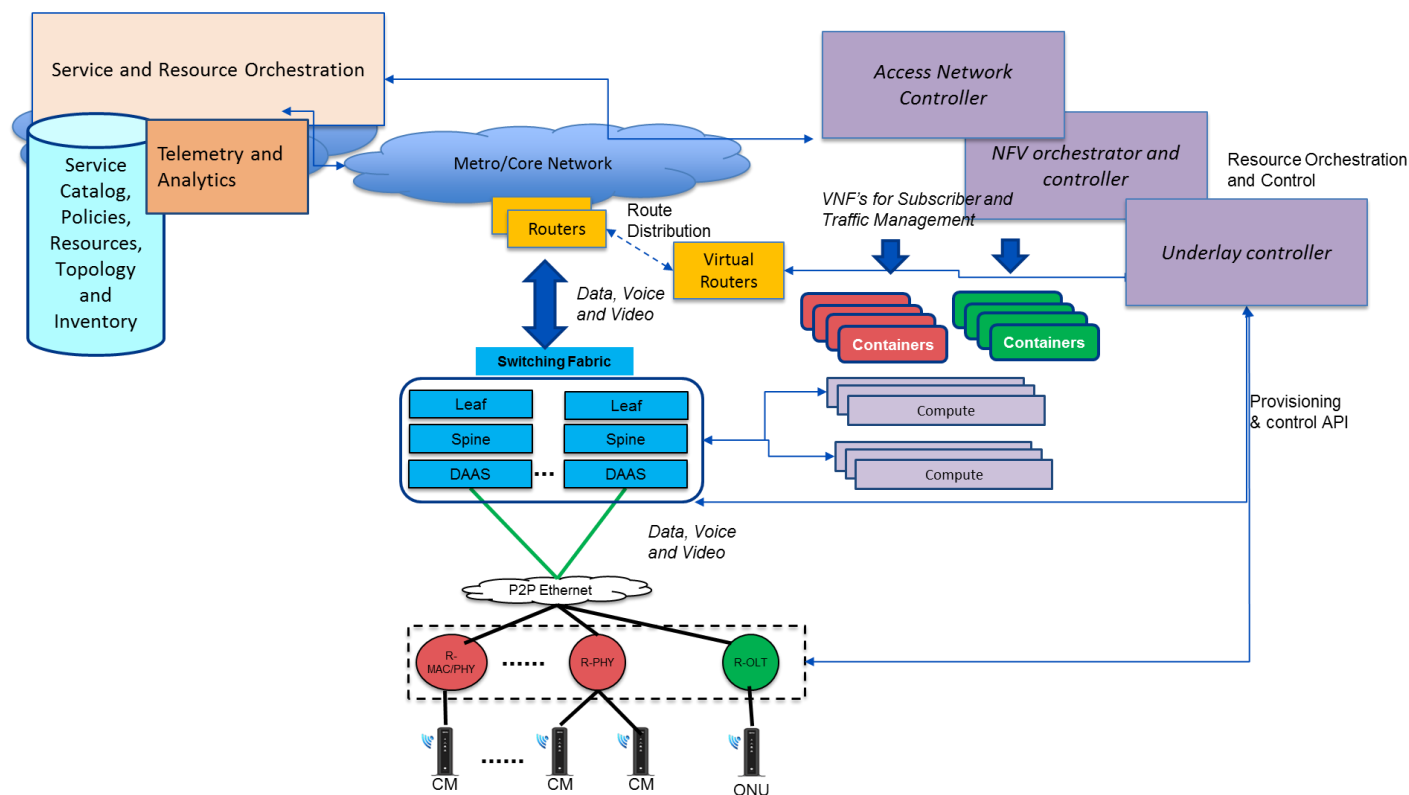


Figure 3 - Example SDN/NFV/cloud Enabled Access Network

3. Conclusions

Although many challenges exist in the network and service architecture transformation for service providers, the long-term benefits of a data-driven and knowledge-defined system make solid business cases. Among those cases are 1) increased speed of service deployment and modification by automating deployment and configuration; 2) decreased operations complexity; 3) easier consumption of services and new business revenues; and 4) enhanced service assurance and faster innovation.

4. Abbreviations

BSS	business support systems
CA	cloud applications
CM	cable modem
DAAS	distributed access aggregation switch
DOCSIS	Data-Over-Cable Service Interface Specifications
EPON	Ethernet passive optical network
HFC	hybrid fiber coaxial
NFV	network function virtualization
ONU	optical network unit
OSS	operational support systems

PNF	physical network function
QoE	quality of experience
QoS	quality of service
R-MACPHY	remote media access control and physical layer
R-OLT	remote optical line terminal
R-PHY	remote physical layer
SCTE	Society of Cable Telecommunications Engineers
SDN	software defined networking
VNF	virtual network function

5. Bibliography and References

[1] <http://techblog.comsoc.org/2018/10/19/tm-forum-survey-communications-service-providers-struggle-with-business-case-for-nfv-dx/>

[2] *Harmonization 2.0: How Open Source and Standards Bodies Are Driving Collaboration Across IT*, Linux Foundation Networking and Orchestration White Paper, March 2018
https://www.linuxfoundation.org/wp-content/uploads/2018/03/lfn_harmonization2_0_os_standards.pdf

[3] *EPON Architecture Considerations for Intent-Based Networking*, Sebnem Ozer and Nagesh Nandiraju, SCTE Journal of Network Operations, Vol 3, No 1, October 2017

[4] *Machine Learning for Network Automation: Overview, Architecture, and Applications*, Danish Rafique and Luis Velasco; Journal of Optical Communications and Networking, Vol. 10, Issue 10, 2018

[5] https://www.onap.org/wp-content/uploads/sites/20/2018/11/ONAP_CaseSolution_Architecture_112918FNL.pdf

[6] <https://osm.etsi.org/images/OSM-Whitepaper-TechContent-ReleaseFOUR-FINAL.pdf>

The Application of Machine Learning to Alarm Management

Noise Identification and Incident Prediction

Letter to the Editor prepared for SCTE•ISBE by

Doug Junkins

Field CTO, Guavus, Inc.
San Jose, CA
Douglas.Junkins@guavus.com

1. Introduction

The strategic application of machine learning (ML) to existing network event management system (EMS) data flows can lead to a dismissal of the overwhelming majority of alarms that are simply noise and allow network operations center (NOC) personnel to focus on the 2-3 percent of alarms that are related to actual incidents.

2. Background and Case Study

2.1. Alarms

Networking monitoring was once a simple matter. Twenty-five years ago, a cable operator needed to be sure that video receivers and processors at the headend were up and running, and that any active outside plant devices (e.g., fiber optic nodes) were not only operational but also had adequate battery backup. Then we saw the advent of high-speed data, telephony, headend interconnects, business services, industry consolidation, enhanced residential services and cloud computing, all adding complexity, devices and scale.

Today, any EMS or network monitoring system (NMS) that tracks faults and aims to assure service delivery must understand dozens of communication technologies and protocols – from Ethernet to DOCSIS to MPLS. Additionally, it needs to be able to probe a thousand or more devices, each of which may generate alarms when failing to meet operational thresholds. The upshot is not surprising: NOCs are inundated with an immense number of alarms. In our work in the industry, we have found that a typical NOC receives one alarm per second, 3,600 per hour, and 86,400 per day.

2.2. Machine Learning

The rise of the ML branch of artificial intelligence (AI) is another current trend. According to a *Heavy Reading* report, amongst its drivers are advances in neural network theory, improved computing capacity and readily available public cloud infrastructure. The report also notes that massive data sets “are fertile ground for ML” and that success stories have fueled further interest [1]. Whether through Google’s AlphaGo or the launch of image recognition, lip-reading or other such applications, public awareness of this category is now widespread.

Network engineers are also aware of ML and its ability to process data without predetermined rules. The question of how to leverage ML in the cable industry is several years old. CableLabs took an early look at the problem, but over the past two years cable operators and technology partners have also contributed to the discussion. This includes looking into how ML could apply to field operations, network virtualization, device modification or upgrades, and set-top box error messages. [2, 3, 4, 5, 6]

Network alarm data is one of those massive data sets that is well-suited to ML’s capabilities, but the application of ML faces some common resistance. First, while awareness has grown, expertise in the algorithms and techniques of ML is relatively scarce. Data scientists are in short supply. Second, constraints and culture tend to argue against change. Budgets at NOCs are not expanding, and NOC directors – and ops leaders in general – typically resist deploying new technology until they know precisely how it is going to impact existing systems and networks.

2.3. Noise and Predictions: Case Study

Existing EMSs have to deal with the large number of alarms, whether 3600 per hour or more. The standard approach is to reduce the volume by filtering or de-duplication. Filtering, however, can be difficult to manage: Every time a new piece of equipment is added to the network, the rules must be reviewed and updated. In any case, even if you can filter out 80 percent of the alarm traffic, it still leaves more alarms to attend to than a reasonably sized NOC staff can handle.

The capabilities of ML technology are more powerful than rules-based filters, and its potential is open-ended; yet it makes sense to start with a few simple but powerful tasks, such as reducing noise and predicting real incidents. All you need is an ML engine to tap into the existing EMS-bus, which can be achieved via an XML-gateway. Once the ML analytics are completed, the annotated alarms are fed back into the bus, with minimal changes to existing systems and procedures. (See Figure 1.) The upshot is an EMS that is smarter about which alarms truly matter.




Figure 1 – Machine Learning Engine and Alarm Bus

A test that Guavus conducted with a large North American cable operator in early 2018 proved the value of this strategy. Starting with historical alarm data labeled with incidents that had been opened in the operator’s trouble ticket system, we used this data to train an ML predictive model to identify new alarms in a real-time stream that were likely to lead to tickets.

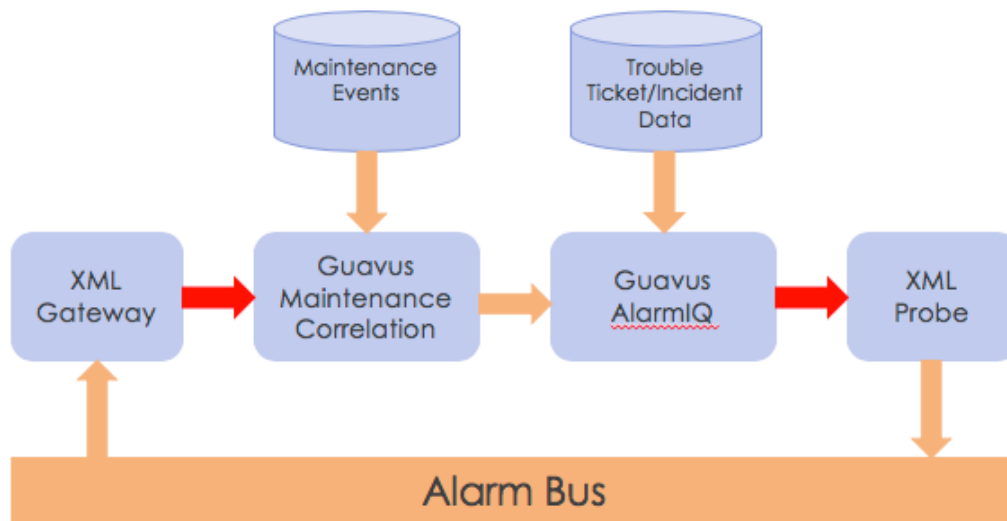
Then handling 447,989 live alarms over the course of a week, we classified alarms in three ways: 1) alarms related to already opened maintenance tickets, 2) alarms likely to lead to incidents, and 3) alarms unlikely to lead to incidents or customer-generated tickets. In this test, we classified 97.6 percent of them as noise, i.e., unlikely to lead to a ticket. That meant that only about 2 percent of alarms were related to unplanned incidents, thus requiring immediate attention. The accuracy of that prediction, compared with actual results, proved high: 99.5 percent of alarms were correctly identified as not being associated with tickets. Note that achieving these insights required no new UI, training, workflow or automation changes. Nor will additional training or rules be necessary going forward, because once an ML model is trained, subsequent actions taken by operators further reinforces the model and keeps it up to date.

Being able to dismiss nearly 98 percent of all alarms leads to operational savings. Assuming that the existing EMS filtered as much as 80 percent of alarms, and that investigating each alarm takes 1 minute, costing \$0.50, we calculated the total annual savings of dismissing the remaining noise at \$2.3 million. The value of achieving faster mean-time-to-repair (MTTR) is even higher. Chasing noise allows real incidents to escalate, generating calls, tickets and truck rolls. By focusing more quickly on real issues instead, it is estimated that the operator could realize further annual savings of \$15 million.

3. Conclusions

Alarm classification based on incident prediction is a solid, initial use case for applying ML to network alarm data. The potential scope for ML in overall cable operations covers considerable territory, from the access plant to virtualization to set-top impairments; and given the rich set of ML algorithms, tools and techniques, there are likewise many ways to generate further value from the application of smarter analytics to network alarm data sets. Incident prediction is good place to start.

Increased complexity has led to ever more powerful “alarm storms.” Existing EMSs can mitigate their impact on NOCs, to some extent. The promise of ML is to go further and change the alarm climate at a fundamental level. More specifically this means to identify the 2-3 percent of alarms that require attention; to eliminate the cost of pursuing noisy alarms; to save potentially millions of dollars more through more rapid network repairs; and to eliminate the need to continuously update EMS rules manually. All of this can be achieved with a light touch on existing systems and procedures.



4. Bibliography and References

1. *AI in Telecom Operations: Opportunities and Obstacles*, James Crawshaw, Heavy Reading, September 2018
2. *Applications of Machine Learning in Cable Access Networks*, K., Sundaresan, N. Metts., G. White, A. Cabellos-Aparicio. NCTA, 2016
3. *Simplifying Field Operations Using Machine Learning*, S. Dorairaj, C Bastian, B. Burg, N. Pinkernell. SCTE/ISBE, 2017.
4. *Leveraging Machine Intelligence and Operational Analytics to Assure Virtualized Networks and Services*, A. Sundelin. SCTE/ISBE, 2017

5. *Embracing Service Delivery Changes with Machine Learning*, A. Sundelin. SCTE/ISBE, 2018

6. *Predicting Service Impairments from Set-top Box Errors in Near Real-Time and What to Do About It*, J. Watson, R. Brooks. SCTE/ISBE, 2018

