

SCTE ISBE

# JOURNAL OF NETWORK OPERATIONS

Volume 1, Number 2  
September 2016



[www.scte.org](http://www.scte.org) | [www.isbe.org](http://www.isbe.org)



**Society of Cable  
Telecommunications  
Engineers**



**International  
Society of Broadband  
Experts**

# JOURNAL OF NETWORK OPERATIONS

**VOLUME 1, NUMBER 2  
September 2016**

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

© 2016 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

5 From the Editors

### Access Network

- 6 **Strategies for More Efficient HFC-FTTx Transitions in the Last Mile**  
Mark Alrutz, Sr. Director, CommScope
- 16 **Improving OSP Network Reliability Using Predictive Alarming**  
Rob Anderson, Sr. Director of Product Management, Alpha Technologies  
Jeff Burgett, Senior Partner, Olivaw Technologies
- 36 **Beyond RFoG**  
Yuxin (Eugene) Dai, Ph.D, Cox Communications
- 44 **Next-Generation Node Splitting Architecture with Unidirectional DWDM Digital Return Path**  
Gaylord Hart, Director MSO Market Segment, Infinera
- 58 **Optimizing Cable Access Network Evolution for Cost, Performance, and Competition**  
Jean-Philippe Joseph, Lead Consultant, Bell Labs Consulting  
Ton Van Caenegem, Bell Labs Consulting  
Luc Ongena, Bell Labs Consulting  
Rex Coldren, Bell Labs Consulting
- 78 **Reducing the cost of Fiber-To-The-Home Brownfield Deployment through the Implementation of Core Extraction on an HFC Network**  
Franklin M. Lartey, Ph.D., Senior Manager Technology, Cox Communications  
Rod McGinn, Executive Director, Cox Communications  
Nicholas Diponzio, Vice President, Cox Communications
- 91 **Interference-Aware Spectrum Resource Scheduling for FDX DOCSIS**  
Tong Liu, Technical Leader & CTO Cable Access, Cisco  
John T. Chapman, Fellow & CTO Cable Access, Cisco  
Hang Jin, Distinguished Engineer & CTO Cable Access, Cisco
- 109 **Cable's Success is in its DNA: Designing Next Generation Fiber Deep Networks with Distributed Node Architecture**  
Venk Mutalik, Engineering Fellow, ARRIS  
Zoran Maricevic, Engineering Fellow, ARRIS

### Advanced Networking

- 128 **Framework for the Evaluation of SDN WAN Controllers**  
Javier Ger, Expert, Sr. Network & Services Architect, Cablevision Argentina SA  
Eduardo Panciera, Expert, Head of Technology and Architecture, Cablevision Argentina SA

#### **SCTE/ ISBE Engineering Committee Chair:**

**Bill Warga**  
VP- Technology, Liberty Global

#### **SCTE/ISBE Network Operations Subcommittee (NOS) Committee Chair:**

**Ron Hranac**  
Technical Leader- Cable Access Business Unit, Cisco Systems, Inc.  
SCTE Fellow

#### **Senior Editors**

**Ron Hranac**  
Technical Leader- Cable Access Business Unit, Cisco Systems, Inc.  
SCTE Fellow

#### **Daniel Howard**

Director- Energy Solutions, Hitachi Consulting Corp.  
SCTE 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





## Wireless Networking

- 153 **Deploying Holistic Wi-Fi: Using Self Optimizing-Network Techniques to Provide a 360-degree View of Wi-Fi QoE**  
Patricio S. Latini, SVP Engineering, Intraway
- 179 **In-Home Wi-Fi Performance Monitoring and Management: Ensuring QoS Across Managed and Unmanaged Devices**  
Metin Taskin, CTO and Co-Founder, AirTies Wireless Networks

## Letters to the Editor

- 189 **Orchestrating the Software-Enabled Operator: An Overview of Key Concepts in Network Function Virtualization Orchestration as Applied to the Cable Network**  
Alon Bernstein, Distinguished Engineer, Cisco Systems
- 195 **PNM Developments: Improving Customer Satisfaction with Better Proactive Network Maintenance**  
Tom Williams, Principle Architect, CableLabs

### **SCTE/ ISBE Engineering Committee Chair:**

**Bill Warga**

VP- Technology, Liberty Global

### **SCTE/ISBE Network Operations Subcommittee (NOS) Committee Chair:**

**Ron Hranac**

Technical Leader- Cable  
Access Business Unit, Cisco  
Systems, Inc.  
SCTE Fellow

### **Senior Editors**

**Ron Hranac**

Technical Leader- Cable  
Access Business Unit, Cisco  
Systems, Inc.  
SCTE Fellow

**Daniel Howard**

Director- Energy Solutions,  
Hitachi Consulting Corp.  
SCTE 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



**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](mailto:journals@scte.org).

**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](mailto:journals@scte.org).

**Subscriptions:** Access to technical journals is a benefit of SCTE/ISBE Membership. Nonmembers can join at [www.scte.org/join](http://www.scte.org/join).

**Note:** DOCSIS is a trademark of CableLabs

## From the Editors

Welcome to the second issue 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). Several of the cable industry's best and brightest have contributed to this issue, covering topics that include the access network, advanced networking, and even wireless networking.

If there is a constant in the cable industry, that constant is change. From the early days of all-coax tree-and-branch networks to the more reliable hybrid fiber-coax (HFC) networks and finally to today's networks that combine HFC, fiber-to-the-premises and wireless technologies the evolution of the access network shows no sign of slowing down. Fiber continues to be pushed deeper toward all customers and making it all the way to increasing numbers of subscribers. Seven papers discuss this evolution and ways to ensure its reliability, providing strategies for the migration to what has been called fiber-deep and fiber-to-the-x, where "x" is the curb, building, home, premises, and so forth. One paper touches on a new concept: full duplex DOCSIS operation, which has the potential to as much as double the capacity of a cable network by allowing simultaneous and symmetric downstream and upstream transmissions in the same spectrum.

One of our papers describes the main aspects comprising the framework used by Cablevisión Argentina to evaluate software defined network wide area network controllers, a must-read if you have been thinking about SDNs.

While one might at first glance think the cable industry is a wired industry, wireless technology is playing a big role in delivering services to our subscribers. Indeed, Wi-Fi is in many instances an extension of our networks. From the perspective of many end users, Wi-Fi in the home *is* cable service. Two papers highlight the importance of quality of experience and quality of service in wireless networks, both critical to subscriber satisfaction.

We would like to thank 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 the readers enjoy this issue and that the selected papers stimulate new ideas and innovations in cable network operation. In closing, if there is any editorial information or topics that you would like us to consider for the third issue of SCTE/ISBE Journal of Network Operations, please refer to the "editorial correspondence" and "submissions" sections at the bottom of the table of contents for further instructions.

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



Ron Hranac  
Technical Leader – Cable Access Business Unit, Cisco Systems, Inc.  
SCTE Fellow



Daniel Howard  
Director – Hitachi Energy and Environmental Efficiency, Hitachi Consulting Corp.  
SCTE Member

## **Strategies for More Efficient HFC-FTTx Transitions in The Last Mile**

A Technical Paper prepared for SCTE/ISBE by

Mark Alrutz, Senior Director, Service Provider Solutions  
CommScope, Inc., SCTE/ISBE Member  
6519 CommScope Road  
Catawba, NC 28609  
malrutz@commscope.com  
828.241.6492

## Table of Contents

Table of Contents	7
1. Abstract	8
2. Introduction	8
3. Fundamental Technologies	8
4. Cost Drivers	10
5. Solutions	11
5.1. MDU	11
5.2. SFU	12
6. Conclusions	14
7. Abbreviations and Definitions	15
7.1. Abbreviations	15
7.2. Definitions	15

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - PON diagram	9
Figure 2 - Fusion Splicing (as described above)	10
Figure 3 - MDU graphic with riser	12
Figure 4 - Splice closure	13
Figure 5 - Full Plug-and-Play	14

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - MDU Installation Comparison	12
Table 2 - SFU Installation Efficiency per Home Passed	14

## 1. Abstract

Many in the multiple system operator (MSO) community are planning to build fiber to the x (FTTx) networks where required to satisfy customer demand for higher capacity and as an extension or overlay to hybrid fiber-coax (HFC) architectures. Naturally they want to invest wisely and obtain the fastest return on investment. Planning for the last mile of fiber deployment is especially important where the network architecture and equipment designs can greatly impact the speed of deployment, construction costs, and flexibility for future services. Key challenges for MSOs are determining when and how much of the network to build out. For example, do you invest more in upfront fiber equipment and construction to serve every customer on day one, or do you build the network to serve a smaller percentage of customers in line with the anticipated take rate? And if you build the network to reach fewer customers, what do you do when more customers want service? Finally, once a network design is selected, it's important for the MSO to explore the most efficient methods for connecting each subscriber, whether it be in a single family or multi-tenant unit or other building type. This paper addresses strategies and techniques for building more efficient FTTx networks in the last mile, using existing HFC networks as a solid foundation. The reader will learn how to right-size their FTTx network designs and extensions to maximize return on investment (ROI), and methodologies for expanding the network over time, minimizing day one costs.

Additional learning points include:

- Common challenges in last mile HFC - FTTx transition and construction
- Network architecture considerations for initial fiber build-outs
- Planning for future network upgrades and services for higher take rates
- Key technologies that enable faster FTTx construction and scalability

## 2. Introduction

Many papers have been written addressing various aspects of FTTx and fiber to the home (FTTH) architectures and technologies. The sheer volume of these papers begs the question, "Why?" Why are we considering building another network? This is a particularly compelling question to the cable operator, having a robust, capable and evolving HFC network already in place.

The answer is as varied as the networks, communities and subscriber needs we deal with every day. In some cases, subscriber demand for symmetrical data rates is the driver. In others, the presence of competitive forces is the primary cause. Yet others desire to lower operating expense (OPEX) and upgrade an aging infrastructure can influence the build decision. Whatever the cause, the MSO community is driving ahead with plans to build FTTx networks. This paper discusses various transition architectures, and how making some decisions earlier in the deployment of the physical plant can dramatically impact build efficiency as well as total cost of deployment (TCD), and total time to deploy (TTD).

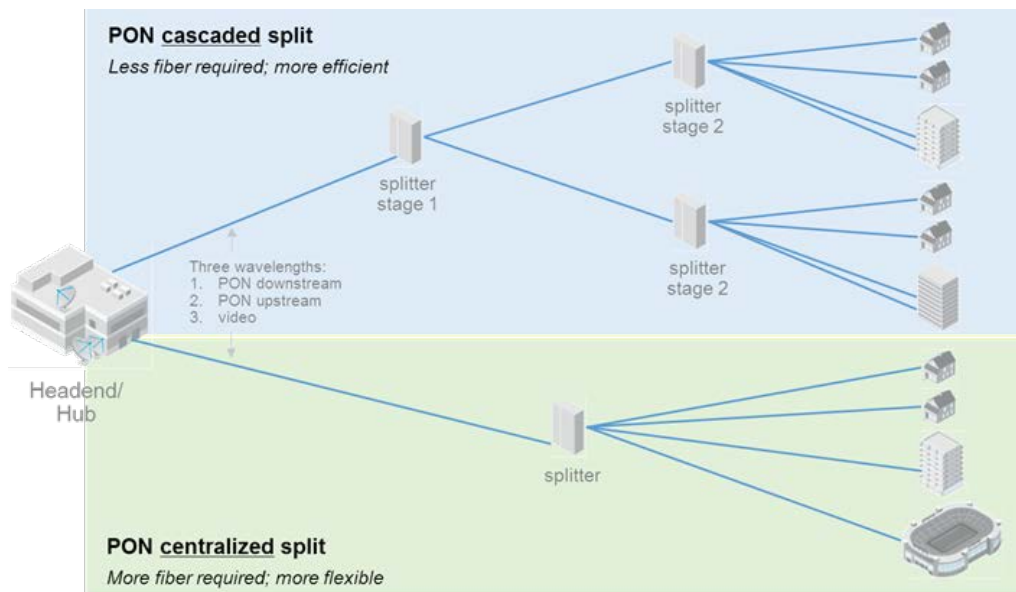
## 3. Fundamental Technologies

There are a few fundamental technologies and networking methodologies to understand as we discuss transitioning from HFC to FTTx networks. The key reason to understand these varying technologies is that they may impact how a fiber network is laid out, and can certainly impact TCD/TTD. As we walk through various cost drivers and solutions for FTTx, these technologies will form a backdrop worthy of consideration.



The Data-Over-Cable Service Interface Specifications (DOCSIS<sup>®1</sup>) protocol is the defining technology for data transmission in the HFC network. As DOCSIS technology evolves, it continues to provide competitive performance with many other network types, including fiber. Gigabit downstream services are available today using DOCSIS 3.1 specifications; however, gigabit speeds require significant HFC network upgrades, which may include reducing node sizes and amplification. Upstream speed enhancement can also require a change of spectrum split. Full Duplex DOCSIS may well answer the need for symmetrical services over passive coaxial plant, as the roadmap continues to evolve.

Passive optical networks in various flavors (EPON/GPON/etc.) offer ever increasing symmetrical and asymmetrical speeds over passive fiber links. A PON service area typically comprises 32 or 64 subscribers, but additional variations exist. Splits can be made centrally, or distributed as needed, which will be discussed later. PON solutions are typically associated with all-Internet protocol (IP) operation, although PON can be used to deliver analog video services and DOCSIS-based services as well.



**Figure 1 - PON diagram**

Radio frequency over glass (RFoG) is a technology that allows traditional HFC video and DOCSIS services to be transported over PON. One key advantage to RFoG is that it allows existing HFC transmission gear, business support and operational support systems (BSS and OSS) to be utilized over a fiber-to-the-subscriber infrastructure, where IP-based PON solutions may require additional changes in the way these HFC-centric services are offered. RFoG can serve as a permanent solution where fiber access is available and DOCSIS services are deemed adequate, or as a transition technology where fiber is placed and IP-PON will eventually be made available.

<sup>1</sup> DOCSIS is a trademark of CableLabs

## 4. Cost Drivers

A quick review of the various cost drivers in network build yields some very familiar categories. This section focusses on the major categories, since our ability to influence these has the largest impact on total cost and total time of deployment (TCD/TTD). The four categories/drivers discussed here are splicing, construction, truck rolls, and inventory. Before we begin exploring these areas, we need to consider two unique construction situations: overlay and greenfield.

Overlay or brownfield construction involves building a new network in place of or to augment an existing network. The challenges associated with this situation include space availability, potential penetration rate for new services, access by different crews and skill sets, and potential for disruption of existing services. Greenfield builds involve constructing a new plant where services are currently not offered. Disruption, access issues and space may be significantly less in this scenario, which is a reason why many providers are electing to build full fiber networks for all-greenfield areas today.

Fusion splicing is an ever present cost in the implementation of fiber networks. Tools and technology have vastly improved over the years, but the process remains a cost driver, in both labor and materials. Splices must be properly formed, and then properly protected. In order to facilitate splicing, excess cable is often stored in the vicinity of a splice closure to allow easier relocation of the closure by craft. Modern splice closures offer many features and benefits including multiple ports for cable access and cable attachment; as well as routing features to secure, protect and store exposed cable elements. While splicing and closure sealing are well understood technologies, and are quite reliable when properly done, care must be taken to properly outfit, train and calibrate crews and equipment. A faulty splicer or a closure left open after access can significantly drive quality of service (QoS) and truck roll costs. On the positive side, splicing does offer great flexibility and a low loss connection, at the previously described cost of splicing.



**Figure 2 - Fusion Splicing (as described above)**

Construction itself is a cost driver that must be considered. Each foot of cable installed comes with not only a material charge, but also a labor charge. Each splice made, each cabinet set, and every connection made can impact the overall labor charges. Some of these costs are unavoidable, as in a location where cables must be placed underground or where an interface box must be placed, but many of these costs can be minimized by choosing solutions that minimize time to install and/or material usage.

Truck rolls are another well-known cost driver. Each provider has metrics to track and programs to minimize truck rolls. For the purposes of this paper, consider truck rolls as a function of either network component reliability or as a function of the difference between complete vs incremental build scenarios. It is also worth noting that reduced maintenance expense in all-fiber networks, a well-documented subject, can impact both truck rolls and potential headcount.

## 5. Solutions


Each of the cost drivers just described can now be considered and applied to common build scenarios to help determine the most efficient and cost effective means to transition from HFC to FTTx. These examples are generic, and are not meant to fully describe every possible build situation. The logic used can be applied broadly to your individual needs and network challenges.

### 5.1. MDU

Multiple dwelling units (MDUs) come in many shapes and sizes, but are generally a good target for fiber connectivity due primarily to density. Fiber can typically be extended to the building or apartment complex without difficulty, and many providers are doing this today, as well as extending fiber to local businesses. The challenge with an MDU or apartment complex is then transitioning from fiber to the building/business (FTTb) to fiber to the unit (FTTu), or pulling fiber to each unit.

In a greenfield scenario, hybrid coax and fiber cable can be readily pulled to apartments to allow maximum flexibility in the transition from HFC to FTTu. This solution allows traditional HFC, RFoG and/or PON to be activated in any order, either separately or in combination. The fiber in this example can be spliced or field terminated at the apartment and fiber distribution hub (FDH) or collector box location.

In brownfield MDU scenarios, vertical runs are made by pulling cable through risers to intermediate or main distribution frame (IDF or MDF) and terminal locations throughout the building. These distribution frames and terminals can be spliced, or pre-terminated using multi-fiber connectors. Using pre-terminated solutions reduces labor significantly, as shown in Figure 3 and Table 1. Also, deploying connector solutions in the wiring closet allows the use of pre-terminated drops in the horizontal, further reducing time and labor.



**Figure 3 - MDU graphic with riser**

Finally, in MDU vertical environments a combination of pre-terminated patching, full plug-and-play and distributed splitting can be used to eliminate multiple cable pulls in crowded risers, as well as to minimize the number of unique parts required to keep on hand, also called stock keeping units (SKUs). This lowest cost solution is also presented in Table 1.

**Table 1 - MDU Installation Comparison**

	<b>Spliced</b>	<b>Connector</b>	<b>Full Plug-and-Play</b>
Splices	196 splices 13 locations	3 splices 1 location	3 splices 1 location
Feet of installed cable	~200 feet (includes slack)	~ 650 feet	~ 180 feet

## 5.2. SFU

Like MDU areas, single family unit (SFU) construction can come in many different flavors. The solutions reviewed here are applied to a generic SFU model and are not meant to capture all scenarios, but to guide your logic in analyzing your own build situation.

We discussed earlier the impact and issues around greenfield vs. brownfield areas. Another build variation to consider is success based building vs. full build out. In a success based build, infrastructure is not built until a subscriber or group of subscribers request service. The benefit of this approach is that costs are delayed until revenue can be applied to them. On the negative side, the first subscriber must bear the entire infrastructure cost rather than sharing that cost among the wider network. While this might be a semantic issue, these figures often drive monetary decisions. Full build out approaches allow the network costs to be spread more thinly across homes passed and connected, but on the negative side pull build costs up ahead of known revenue and take rates.



Like MDU builds, splicing is a viable solution to use in deploying an FTTH network. As described earlier in this paper, splicing offers flexibility and a low loss connection, but it remains labor intensive and closures can be a failure point if accessed often due to risk of being left unsealed.



**Figure 4 - Splice closure**

Another couple of options to consider are pre-engineered and pre-terminated solutions. Pre-engineered solutions rely on careful plant analysis to build custom cable and connectivity solutions. These solutions are highly dependent on the accuracy of mapping and plant walk-outs, but can reduce installation labor. Another means of reducing labor is the use of pre-terminated solutions, where terminals with hardened connector ports are factory assembled with cable tails that can be spliced into a collector or FDH location.

The lowest cost and most flexible solution from a labor perspective involves a combination of plug-and-play pre-terminated cable and distributed splitting. With this scenario, SKUs can be minimized (potentially to one), splicing can be minimized or even fully eliminated, and installation time can be reduced by half or more, as shown in Table 2. Flexibility in this solution can be achieved by providing spare fiber(s) at each terminal location for future support of business services, growth or emerging needs like 5G. 5G may be a very unique case for service providers, since conceptually the systems will require many distributed wireless terminals, each requiring both backhaul and power. Installing additional fiber will help manage this proliferation and aid with backhaul requirements.

There are several cost and management considerations related to pre-terminated solutions and distributed splitting. First, fully pre-terminated cable solutions can lead to situations where excess cable must be managed. Creative cable management solutions are available to help minimize this issue. Second, the use of distributed splitters may add complexity to documentation and troubleshooting, and cost both in terms of number of splitters and port utilization. Continued improvements in documentation and troubleshooting systems overcome most concerns in this area. Also, material and port utilization costs can be readily modelled and shown to be lower than the benefits of deploying the system. Finally, in addition to the cost benefits shown, expensive and time consuming engineering and site surveys are potentially not required when using fully plug-and-play solutions.



**Figure 5 - Full Plug-and-Play**

**Table 2 - SFU Installation Efficiency per Home Passed**

<b>Resource Time per home passed (hours)</b>	<b>Spliced</b>	<b>Pre-terminated</b>	<b>Full Plug-and-Play</b>
Fiber Technician	0.8	0.6	0
Cable Technician	0.9	0.9	0.5
Total	1.7	1.5	0.5

## 6. Conclusions

As varied as individual network scenarios are, solutions to make the transition from HFC to FTTx are equally varied. In MDU and SFU environments, splicing solutions offer flexibility but can add cost, time and risk. Pre-terminated and pre-engineered solutions reduce cost and risk to some degree, but may limit flexibility. Enhanced plug-and-play pre-terminated solutions that include distributed splitting can have the largest impact on cost and efficiency, minimize SKUs and allow transition work to occur at scale by providing the fastest installation using the broadest labor force.

When considering a transition to FTTx, the following questions will help guide the decision process regarding types of solutions to deploy.

1. How much stability is in the build area? Is it an established or growing area? (flexibility)
2. How quickly does the network area need to be built? (efficiency)
3. How large an area is to be built? Are there adequate and qualified resources? (scale)

After considering the required flexibility, efficiency and scale the best solution will become apparent.

## 7. Abbreviations and Definitions

### 7.1. Abbreviations

BSS	business support systems
DOCSIS	Data-Over-Cable Service Interface Specifications
FDH	fiber distribution hub
FTTb	fiber to the building/business
FTTH	fiber to the home
FTTu	fiber to the unit
FTTx	fiber to the x (where “x” is curb, home, etc.)
HFC	hybrid fiber-coax
IDF	intermediate distribution frame
IP	Internet protocol
MDF	main distribution frame
MDU	multiple dwelling unit
MSO	multiple system operator
OPEX	operating expense/expenditure
OSS	operational support systems
PON/EPON/GPON	passive optical network/Ethernet PON/gigabit PON
QoS	quality of service
RFoG	radio frequency over glass
ROI	return on investment
SFU	single family unit
SKU	stock keeping unit
TCD	total cost of deployment
TTD	total time to deploy
SCTE	Society of Cable Telecommunications Engineers

### 7.2. Definitions

5G	5 <sup>th</sup> generation mobile technology, being defined
----	---

# Improving OSP Network Reliability Using Predictive Alarming

## A Study of Advanced Battery Analysis Methods

A Technical Paper prepared for SCTE/ISBE by

Rob Anderson, Sr. Director of Product Management, Alpha Technologies, SCTE/ISBE Member  
3767 Alpha Way  
Bellingham, WA, 98226  
randerson@alpha.com  
+ 1 (360) 647-2360

Jeff Burgett, Senior Partner, Olivaw Technologies, SCTE/ISBE Member  
Alpharetta Technology Commission - Innovation Center  
2972 Webb Bridge Rd.  
Alpharetta, GA 30009  
jburgett@olivawtech.com  
+1 (678) 231-2420



## Table of Contents

Title	Page Number
Table of Contents	17
1. Introduction	18
2. Intra-String Monitoring	18
2.1. Effects of Charge on Battery Capacity	18
2.2. Battery Chemistry Results in Charge Mismatch	19
2.3. Battery Charge Case Study	21
3. Predictive Alarming	23
3.1. A Case for Proactive Maintenance	23
3.2. A New Analysis Approach – Predictive Alarms	24
3.2.1. Data Collection	26
3.2.2. Data Pre-Processing	26
3.2.3. Data Analysis	28
3.2.4. Extracting and Implementing Predictive Alarms	29
3.3. Example: Predictive Analysis Using Live Hybrid Fiber Coax Data	30
4. Conclusions	33
5. Abbreviations and Definitions	34
5.1. Abbreviations	34
5.2. Definitions	34
6. Bibliography and References	34

## List of Figures

Title	Page Number
Figure 1 - Battery Charge Effect	19
Figure 2 - Battery Series Resistance	20
Figure 3 - Battery $\Delta$ Resistance	20
Figure 4 - Battery $\Delta$ Charge	22
Figure 5 - Battery Corrosion	23
Figure 6 - Predictive Process	26
Figure 7 - Coup De Fouet	28
Figure 8 - Sample Prediction, 8 Minute Discharge	31
Figure 9 - Sample Prediction, 20 Minute Discharge	32
Figure 10 - Sample Prediction, 45 Minute Discharge	33

## List of Tables

Title	Page Number
Table 1 - Distribution of Battery Voltage Variations	22

## 1. Introduction

Power is fundamental to all aspects of the broadband network. Because utility power disruptions are unpredictable, operators have historically utilized uninterruptible power supply (UPS) systems for reliable power. Status monitoring systems provide real-time notification of critical events, enabling timely, manual intervention to avoid disrupting customer service. Although usually effective, outages are still possible. But now, existing data can be analyzed in new ways and data from multiple sources can be combined to provide advanced notification of situations that are not yet, but can eventually impact services. We present two approaches for utilizing existing data which is analyzed in new ways, to provide advanced predictive alarming to increase network reliability by eliminating future, service-impacting events. These two approaches are summarized as follows:

1. **Intra-String Monitoring.** Current outside plant (OSP) battery testing methods focus on battery string voltage change over time. Although important, this is not enough. Variations in intra-string battery voltages during charge and discharge cycles can be telling. One study found that 37% of an operator's outside plant batteries were at less than 80% of expected capacity. This capacity reduction could have been avoided through advanced alarming on unanticipated data combinations.
2. **Predictive Alarming.** Vast amounts of OSP data is regularly collected but seldom used for proactive maintenance. For example, one U.S. operator processes over 20 million OSP power related events daily. Event processing is real-time but analysis is done at a future date, if at all. Automated machine learning techniques have been used across multiple industries with applications that predict outcomes based on data patterns and trends where large volumes of data are to be reviewed. Applied to OSP battery data, these techniques should be able to report, on an ongoing basis, a prioritized list of sites that need to be maintained based on a set of priorities determined by the operator.


This paper discusses these two approaches to OSP battery analysis with the objective of providing operators with new methods and insights for increasing OSP reliability.

## 2. Intra-String Monitoring

Multiple factors impact battery runtime and health. For clarity, we define runtime as the instantaneous runtime available from a battery or bank of batteries irrespective of environment or history. Battery health is defined as the present maximum capacity of a battery or bank of batteries. Five primary factors affect battery runtime, battery state of health or both. These are: state of charge, ambient temperature, temperature history, battery age and charge history. This section discusses charge history and its effect on both battery runtime and battery health or capacity.

### 2.1. Effects of Charge on Battery Capacity

Overcharging and/or undercharging batteries has a significant effect on battery life. Using charging specifications from battery manufacturers, Figure 1 illustrates how overcharging or undercharging batteries will cause premature battery failure.




**Figure 1 - Battery Charge Effect**

In Figure 1 the X-axis identifies specific voltages under and over the optimum charging voltage. The Y-axis shows the acceleration factor on battery degradation. For example, if a battery were undercharged by 0.6 volt (V), its effective age would accelerate by a factor of 2x resulting in a useful life that is now only 50% of that battery's specified life expectancy. Likewise, if a battery were overcharged by 0.5 V, its useful life would be reduced to 30% of its specified life expectancy.

It should be noted that any charge-related degradation effect occurs only while the overcharge or undercharge condition is applied. For example, undercharging a battery by 0.6 V for a period of 4 months (perhaps the duration until the next preventative maintenance cycle) would result in a loss of 2 months to the overall life of the battery (using a 2x acceleration factor for 0.6 V undercharge.) Once the charge problem is corrected, no further damage will occur. Nonetheless, capacity has been permanently diminished.

## 2.2. Battery Chemistry Results in Charge Mismatch


Typical OSP power supply installations include one or multiple battery strings. Each battery string consists of 12 V batteries connected in series. Three series batteries are combined to achieve a 36 V string or four batteries are configured for a 48 V string. The connection between the power supply and the battery string(s) is through a wire harness connected across the entire string (i.e., one connection is at the negative terminal of the first battery (ground), the second connection at the positive terminal of the third battery for 36 V strings or the fourth battery for 48 V strings). In this configuration, power from the batteries and charge to the batteries is routed through the wire harness for the entire string. A three battery, 36 V battery string is shown in Figure 2.



**Figure 2 - Battery Series Resistance**

Figure 2 also lists the optimal float voltage for this 36 V string as 13.62 Vdc per battery. This value may vary among different battery manufacturers and technologies. The power supply battery charger is configured to supply 40.86 Vdc across the entire battery string. From the charger’s perspective, each battery appears in the circuit as a fixed resistance. Per Ohm’s law, the three batteries, acting as resistors in series, create a voltage divider and the 40.86 Vdc is distributed equally across each of the batteries at 13.62 Vdc per battery.

Batteries operate by electro-chemical reactions. Time, temperature and charge history can affect this chemistry, altering a battery’s internal resistance. If each battery in Figure 2 has internal resistance values that change at different rates, the 40.86 Vdc charge voltage will not be applied equally across each battery. The result of this unequal internal battery resistance is illustrated in Figure 3.



**Figure 3 - Battery Δ Resistance**

As the internal resistance of each battery changes at different rates over time, the actual circuit where  $R1 \neq R2 \neq R3$  causes the charge voltage to be distributed unequally across the three batteries. The result is



that some batteries will be undercharged while others are overcharged. Figure 3 illustrates an example where one battery is undercharged at 12.79 Vdc while the other two batteries are overcharged at 14.20 Vdc and 13.78 Vdc. The total charge voltage of 40.86 Vdc is correct but the battery chemistry has caused an internal string charge variation that will shorten the life of all three batteries over time if not corrected.

Now however, OSP monitoring software can be configured to identify disparate voltages within individual battery strings for maintenance before extended battery damage occurs. Operators should consult battery manufacturers for specific voltage threshold parameters to trigger alerts and initiate corrective action. Multiple options exist to mitigate the effects of time on internal battery resistance. Modern OSP power supplies are equipped with charge balancing technology that will re-direct charge within a battery string to offset the effect of changes to internal battery resistance. This charge management technology is available from various manufacturers in a variety of configurations. Operators should be aware of the effects of charge imbalance on battery life and determine the best course of action for their situation [1].

### **2.3. Battery Charge Case Study**

In 2005, a field study was undertaken with a cable operator in Midwestern United States to examine the effects of battery age on internal battery resistance in a live OSP network. The study and findings are summarized here:

A total of 2,538 individual VRLA type batteries in 846 battery strings of 36 Vdc were identified for the study. The entire sample set of batteries were installed between 2-3 years prior to the study. The cable operator provided access to status monitoring data enabling the average float voltage for each battery to be identified. An acceptable float voltage range was defined as the manufacturer's optimal specified float voltage  $\pm 100$  millivolt (mV). Battery float voltages equal to or greater than  $\pm 200$  mV from optimal were considered undesirable.

The study identified 928 batteries, representing 37% of the sample set, with an average float voltage outside of the acceptable range. Figure 4 shows the float voltage distribution across the complete sample set.

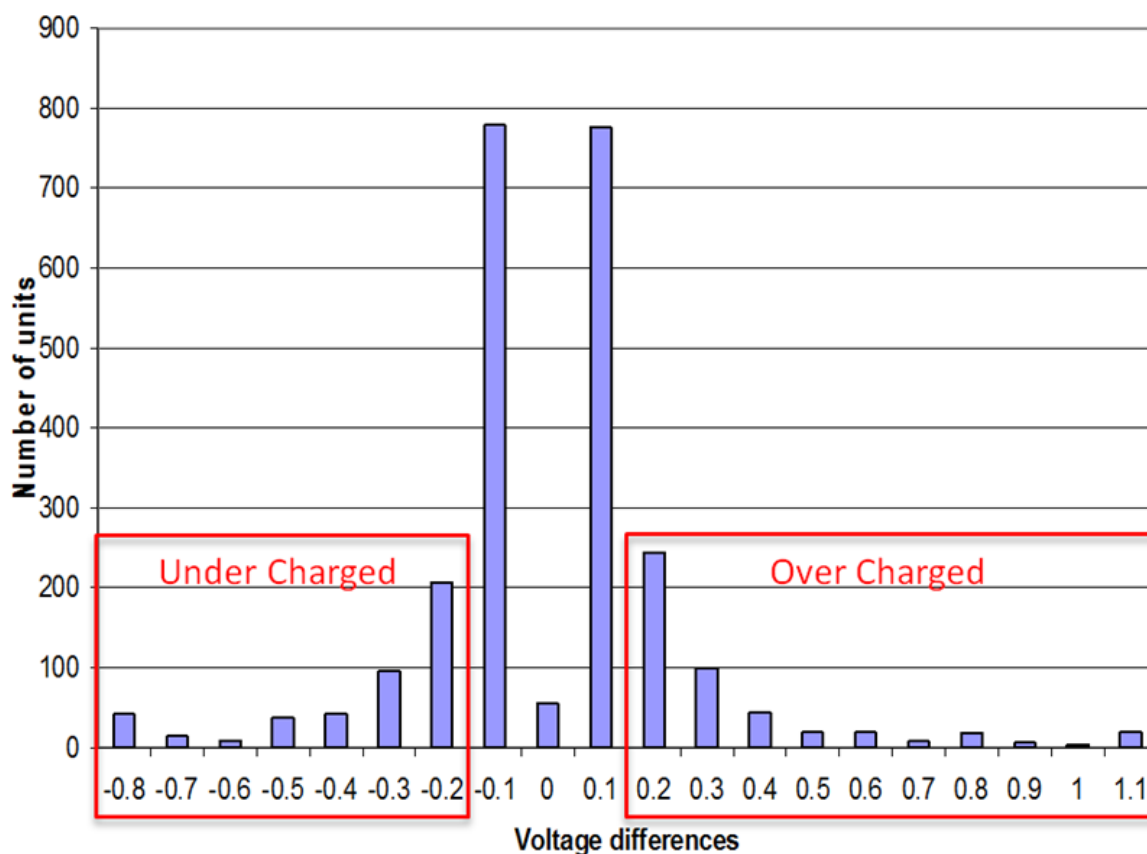


Figure 4 - Battery Δ Charge

The breakdown of number of batteries in specific voltage ranges is shown below in Table 1:

Table 1 - Distribution of Battery Voltage Variations

Batteries	Δ(V)
1610	± 0.1
451	± 0.2
195	± 0.3
282	> ± 0.3

The data from this study represents one instance in time. As mischarge is applied to batteries over time, the acceleration effect shown in Figure 1 continues to degrade battery life resulting in premature end of life from irreversible capacity reduction. The operator in this study elected to install battery charge management devices which were shown later to have corrected the mischarge, eliminating further

damage. Operators can initiate similar reviews of their battery charge condition through their individual status monitoring systems.

### 3. Predictive Alarming

For discussion within this paper we define the terms proactive maintenance and preventative maintenance as follows: Proactive maintenance is intelligence-directed maintenance that is complete prior to any service deprivation or loss. Preventative maintenance is schedule-driven maintenance intended to be completed prior to any service affecting conditions. Preventative maintenance must be utilized for activities that require on-site inspection and that cannot be measured practically through remote data collection. This may include inspection and repair of an enclosure's physical damage, cleaning battery terminals, snow and ice removal or removing pest infestations.

As a practical example, the following case study illustrates an obvious need for regular preventative maintenance. In addition, the study illustrates an opportunity for proactive maintenance to be utilized and identifies associated operational cost savings.

#### 3.1. A Case for Proactive Maintenance

In 2016, a U.S. cable operator conducted an investigation to determine the root cause of their growing number of OSP power supply alarms. Standby fail alarms were of particular interest due to the critical nature of this alarm. Across three (3) U.S. cities, this operator identified that 22% of their power systems had failed standby tests due in part, to battery cable corrosion. FIGURE 5 shows representative battery corrosion at one of these installations.



**Figure 5 - Battery Corrosion**

The center battery shows excessive corrosion around both the positive and negative battery posts. This type and level of corrosion can occur with some styles of OSP batteries that have not been properly maintained. At the conclusion of this study, the operator identified deficiencies in their local OSP preventive maintenance (PM) practices that led to these potentially service-impacting results.

Corrosion damages battery power cables and battery voltage monitoring sense wires. As battery power cables degrade, electrical resistance increases and the ability of the batteries to provide sufficient current diminishes. As this operator experienced, eventually the corrosion will increase cable resistance sufficiently to cause a standby test to fail. Prior to a standby test failure, backup capacity had diminished and expected runtimes during actual power outages were reduced. Had an actual utility outage occurred, standby power would be compromised and customer loads dropped, potentially before the standby test indicated any problem.

Could this liability have been avoided through more diligent preventative maintenance practices? The answer, of course, is yes. “How frequently should each power supply be visited?” is a frequently debated question. Responses vary between operators and even between systems within the same operator. These answers range from six-months to two-years when technicians are queried. Conducting PM visits and finding nothing to correct is wasting valuable service resources. Waiting too long between PM visits could result in service-impacting situations going unchecked. We often hear reports of PM visits to multiple installations within a geographic area with some locations checking out OK while other sites require extensive maintenance. Clearly, there is no one right answer to the question of PM frequency. Can anything be done to reduce unneeded PM visits while focusing limited resource on locations needing physical intervention?

The answer to this question is two-fold:

- First, there are some issues that require on-site inspection to identify and correct. Examples include pest infestations and water intrusion due to physical enclosure damage. Because this category of problems exists, scheduled preventative maintenance visits will always be required.
- Second, it may be possible to identify a category of future service-impacting problems through analysis of data available from status monitoring systems. This would enable us to prioritize site visits and focus on locations at high risk of causing future service disruptions. Low risk sites could be visited less frequently and only to inspect and correct issues that are undetectable any other way.

Using our power supply case study as an example, could status monitoring data have been analyzed to forewarn of these corrosion problems and eventual service impacting results? We believe the answer is yes as described in the following sections.

### **3.2. A New Analysis Approach – Predictive Alarms**

The prior section mentioned analyzing power supply data to predict future events. For the specific problem of battery corrosion in our case study, there are several powering characteristics that can be indicators of future problems, these indicators include:

- An increase in the rate of change of the discharge voltage over time with a fixed output load
- Significant change in the charge and discharge voltage trends across batteries in a string over time



- Decrease in battery conductance as a trend or an absolute conductance value relative to the battery specifications
- Failed standby test due to batteries failing to provide anticipated runtime

Failed standby events are clear problem indicators and typically receive priority attention by operators. These events need no further discussion. Likewise, for operators utilizing battery conductance, excessively low readings indicate batteries have gone bad. Operators typically treat these situations with high priority and no additional discussion is needed here either.

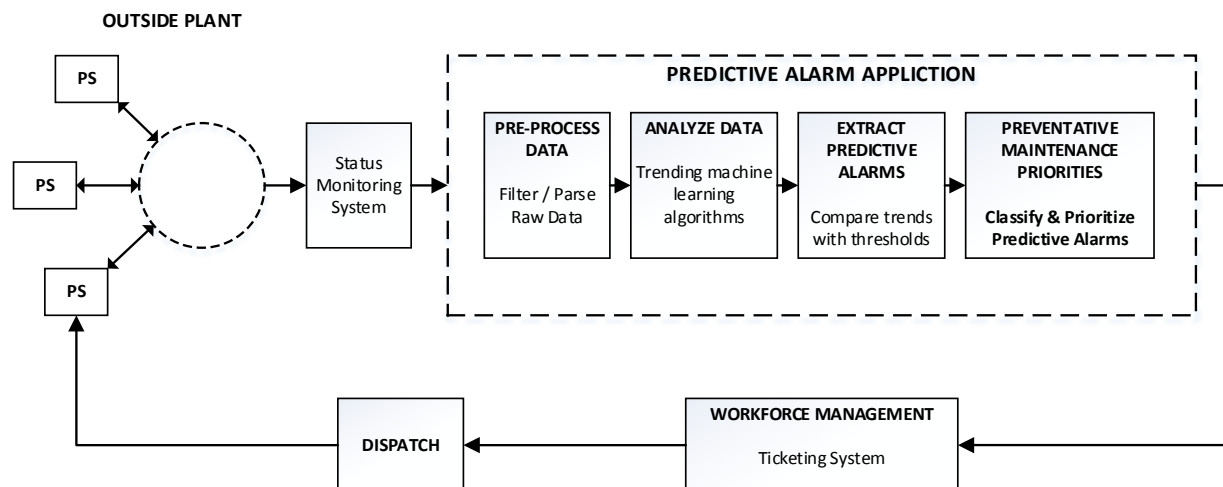
For changes over time in charge voltage, discharge voltage and conductance, a single snapshot in time of this data may not produce actionable information on its own. However, when identifying correlations in the trends, patterns can emerge that provide valuable information for heading off network troubles before they can cause significant problems.

Current status monitoring systems allow the trending of some battery data over time. A healthy battery string will support a typical load for several hours often without an appreciable decrease in battery voltage. As the batteries in a string age, the duration over which batteries can sustain a steady voltage under a relatively constant output current decreases. While this indicator does not provide accurate information as to the remaining runtime provided in the battery string, an accelerated reduction in this duration does provide an indicator that the battery string is experiencing a negative trend and should be observed more carefully.

A significant change in the charge and discharge voltage trends for one or more batteries in a string relative to the other batteries is also an indicator that the battery or batteries are experiencing a negative trend. Like battery string voltage, this alone does not provide a quantitative result regarding the battery state of health. However, it does indicate the battery or batteries should be more closely monitored.

To produce meaningful results, we utilize machine learning techniques and data analytics that have been utilized for predictive decision making for many years in a variety of industries and applications. The objective of the analysis process is to look into the current and historical data derived from monitored network elements to determine future trends that could be identified as a predictive alarm. For our example, the predictive analysis will focus on battery voltages. Historical numerical data is analyzed over time. This requires a time-series analysis approach typically using functions for prediction instead of rules or decision trees.

An overview of this approach is shown in Figure 6 and described in the following section.



**Figure 6 - Predictive Process**

Our data analysis approach outlined in Figure 6 is described in Section 3.2.1:

### 3.2.1. Data Collection

Data from OSP elements, including power supplies (PS), is collected by a status monitoring system. Data is collected using a polling scheme with a typical polling cycle of 15 minutes to 30 minutes, depending on the network architecture, number of polled devices, and other factors. The polling rate usually increases during utility outages to facilitate higher data granularity of critical items such as battery voltage. Alarms and other important notifications are transmitted from the network elements asynchronously in the form of traps or alerts.

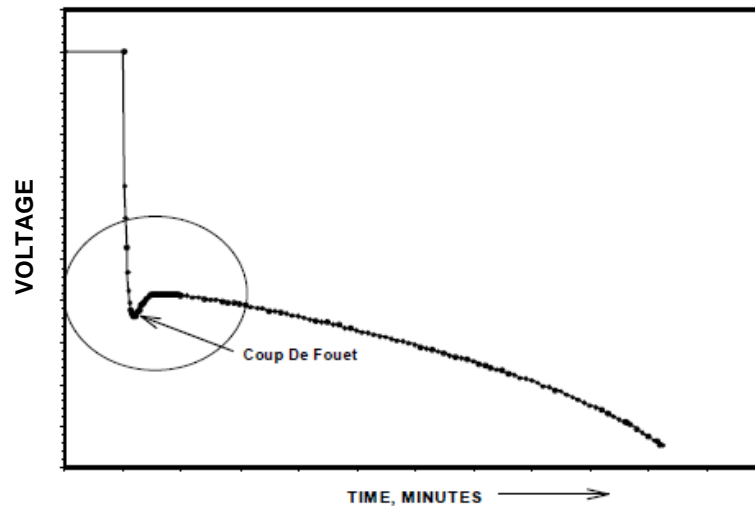
### 3.2.2. Data Pre-Processing

The status monitoring system will provide OSP information in the form of reports or logs. The results will need to be regularly (e.g., weekly) retrieved and converted into a file that can be filtered, parsed, and output to a predictive data file for the analysis stage. In machine learning or data mining terminology this would be considered the data cleansing step. To more accurately identify potential failure indicators for batteries, some preprocessing of the data will require field conditions to be considered. Some data will be included in the analysis phase and some will be excluded. The following data selection criteria were used for this study:

- Data from failed standby tests was eliminated from this study. These events would have triggered alarms in the status monitoring systems and thus would have resulted in immediate action.
- Data from standby events that occur prior to batteries being fully recharged from a previous standby event may not provide meaningful information - depending on the type of analysis conducted. As a result, this data was filtered from the resulting “cleansed” predictive data file.
- Data from power supplies with consistent loads was selected. Load current often remains fairly consistent for power supplies in fixed environments so most samples were included. More sophisticated analysis can be conducted that takes into consideration a change in load, but that is a smaller percentage of existing power supplies and out of the scope of this paper.

- In order to identify trends, three time stamps were selected for data analysis. The first time stamp was the 8<sup>th</sup> minute in each standby event, the second time stamp was the 20<sup>th</sup> minute in each standby event and the third was the 45<sup>th</sup> minute in each standby event. The reasons for these selections were:
  - Most scheduled standby tests have a 10-minute duration. A negative trend detected within this time frame is a strong indicator of an imminent battery problem. An 8-minute timestamp provides a consistent datum within each successful standby test.
  - Because of the *coup de fouet* (“crack of a whip”, or “whiplash” in French, after the shape in the curve) at the beginning of a battery discharge cycle, the battery needs to “settle” prior to extracting meaningful data for trend analysis. The 8-minute time stamp is outside the *coup de fouet* range and meets the criteria. Although there is research being conducted around the use of *coup de fouet* data for battery health (see for example [2]), this research is outside the scope of this study. Refer to Figure 7 - *Coup De Fouet*
  - A battery string that fails within 20 minutes typically exhibits a significant voltage drop early in the discharge cycle. There is not much prediction required for batteries with such diminished capacity. Any failures will trigger immediate alarms and associated action. The 20-minute mark was selected as point outside of an early battery failure window that could provide meaningful information about future behavior of a battery string.
  - The 45-minute mark was selected because there are a significant number of outages that last 45 minutes or longer. These outages allow a large number of batteries to be analyzed. Also, changes in the battery behavior at the 45-minute mark are typically subtle and gradual, providing additional data for the machine learning algorithms to use for trending.

**A note on *coup de fouet*:** *Coup de fouet* is the initial voltage dip and recovery experienced when discharging a fully charged lead acid battery. A voltage vs. time representation is shown in Figure 7. Variables for *coup de fouet* include the load, temperature, the magnitude of the voltage dip, and its duration. [3]



**Figure 7 - Coup De Fouet**

### **3.2.3. Data Analysis**

In the analysis process, cleansed data passes through a series of functions to attempt to establish future trends that could be used for predictive alarms. This requires a time-series analysis approach typically using functions for prediction instead of rules or decision trees. There are many machine learning algorithms that can be used for predictive models. The goal of the predictive model is to make the best prediction possible with the smallest error in the model. Accuracy is increased if unnecessary attributes are eliminated prior to the analysis. For example, we eliminated power supplies with variations in output current during the data cleansing phase, thus eliminating output power as an attribute in the model. The algorithms selected to test for a good predictive model were linear regression, Gaussian processes, multilayer perceptron, and sequential minimal optimization for regression (SMOreg). These four algorithms are discussed briefly in this section.

#### **3.2.3.1. Linear Regression**

Linear regression is a well know and simple technique that has been around since the 1800s. In its most basic form it uses a least-squares approach to prediction where it finds the best line through the given data points. The linear regression algorithms used in typical data mining tools are significantly more sophisticated than basic least-squares functions with built in capabilities to reduce overfitting and the ability to stabilize degenerate cases.

For our testing, we utilized the standard multiple linear regression form of linear regression since this seemed to best fit the multi-attribute instances and time-series characteristics of the predictive data.

The disadvantage of the linear regression model relative to predicting battery behavior is that it is most effective when working with data that produces a more linear prediction. Batteries do not typically degrade linearly [4].

#### **3.2.3.2. Gaussian Processes**

Gaussian processes are a probabilistic non-linear regression approach that utilizes a Bayesian Gaussian process technique and has similarities to kernel ridge regression used in statistical analysis. Like linear

regression, there are a variety of adjustments that can be made to this algorithm to control the closeness of fit for the projected results. Unlike linear regression, the prediction model does not assume that the trend of future events is more linear relative to past events. Gaussian processes can model any type of continuous behavior, whether it is linear, polynomial, or even non-polynomial. This fits well with what is expected by the behavior of battery degradation.

One limitation of Gaussian processes is the computational complexity. For very large data sets this approach may not be feasible; however, the data points used for our sample analysis were in the thousands, not millions so the analysis was feasible on a small scale. This approach often produces very good results [5].

### **3.2.3.3. *Multilayer Perceptron***

The multilayer perceptron (MLP) is a feedforward artificial neural network and is the most widely used type of neural network. The term perceptron describes a mathematical model of a biologic neuron. A multilayer perceptron adds hidden layers of nodes – analogous to neurons - to the input and output layers. Given enough hidden layers, the MLP is able to approximate any continuous function to any given accuracy.

The advantage of a MLP approach is that it can be a very good predictor of future behavior of the network elements, especially for battery degradation since it generally has a consistent trend, i.e. battery health will only decline over time. The disadvantage of the MLP approach is that it can get stuck in a local minima within the trend. Since this is a potential issue, close attention must be given to the output of the forecast during training and production [6].

### **3.2.3.4. *Sequential Minimal Optimization for Regression***

SMOreg is a sequential minimal optimization algorithm for training support vector machines (SVC). The SVC approach uses learning via quadratic programming and was derived from statistical learning theory primarily for binary classification problems. While a very effective approach for classification, its limitation is that it is not very scalable due to performance and memory constraints. The primary advantages of the SVC over other regression techniques like MLP are the reduced chance of over-fitting and issues with local minima. Like the MLP approach, SMOreg can provide a very good predictor of future behavior of the network elements, especially for battery degradation since it generally has a consistent trend, i.e., battery health will only decline over time [7].

## **3.2.4. *Extracting and Implementing Predictive Alarms***

After completing the battery analysis with each of the four analysis approaches, linear regression, Gaussian processes, MLP, and SMOreg, the results from each are compared with our historic field data to determine which approach models the behavior most accurately. This step is part of the algorithm evaluation process. Once the most accurate model is determined, that model is implemented for use in actual predictive alarm generation. In our example, the results of the analysis produces projected battery voltages threshold failures. These results are programmatically compared against thresholds for each power supply to determine if a failure is predicted to occur. For our sample analysis a battery voltage failure threshold of 10.5 volts was used.

Once the potential battery failures are identified, the predictive alarms are generated. This can take many forms, but should be automated programmatically by either sending the alarms directly to the workforce management system, the ticketing system, or the network management systems via their application



programming interfaces (APIs) or file transfer where they can be flagged as actionable events and associated maintenance can be scheduled.

### 3.3. Example: Predictive Analysis Using Live Hybrid Fiber Coax Data

Historical data from 200 power supplies from the southeastern United States was collected. Often the maintenance history of an individual power supply is not maintained to the extent that the date of battery replacement is recorded along with the make and model of the batteries. To demonstrate the robustness of the modeling approach, a power supply was selected for example that had a battery string replacement during the timeframe the batteries were analyzed. As part of the data preprocess (i.e., cleansing) step, standby events of shorter than 8 minutes were excluded as were standby events that occurred within 24 hours of a previous standby event to allow for charging.

**Summary of Results:** The power supply was analyzed using each of the four machine learning algorithms: linear regression, Gaussian processes, multilayer perceptron, and sequential minimal optimization for regression (SMOreg). In order to determine the relative accuracy of the models, a power supply was selected that had exhibited a failure at 53 minutes during its last standby event. All four models predicted that the batteries would not sustain a 45-minute standby event within the next five months. Also all models predicted that a 20-minute standby event would be supported within its forecast window of eight months.

The results from the SMOreg algorithm are shown on the next three pages. The final results for each forecast are provided at 8, 20 and 45 minutes into the standby events. The forecasted results in the tables are in **blue**, with the forecasted failure results in **red**.

There are a number of techniques that can be applied to this approach to increase accuracy such as increasing the amount of test data and training steps, adjusting weights within the algorithms, and providing more sophisticated data cleansing on the front end.

The advantage of using machine learning algorithms over more traditional software approaches (e.g., rules-based or decision trees) is the robustness and broad accuracy of the results. It would be difficult to write software that could effectively accommodate the different conditions faced in an operational network. For example, the algorithms provided in this study all handled the new battery installation at the power supply in the example without any adjustment to the algorithm. It would be very complex to write software that could cover all the cases experienced in the network.

With predictive alarms, operators can more effectively schedule maintenance visits to correct future potential failed batteries rather than just respond to unexpected outages.

Analysis using the SMOREg algorithm:

=== Future predictions from end of training data at 8 minutes into standby events ===

Time	BatteryVoltage1	BatteryVoltage2	BatteryVoltage3
2013-12-05	12.1	12.2	12.3
2014-03-26	12.2	12.2	12.3
2014-07-15	11.8	12.0	12.0
2014-11-03	12.1	12.2	12.3
2015-02-22	12.2	12.3	12.4
2015-06-13	12.2	12.3	12.4
2015-10-02	12.3	12.3	12.4
2016-01-21	12.2	12.3	12.3
2016-05-12	12.1	12.2	12.3
2016-08-31*	12.1	12.2	12.3
2016-12-20*	12.0	12.1	12.3




Figure 8 - Sample Prediction, 8 Minute Discharge

SMOreg

=== Future predictions from end of training data at 20 minutes into standby events ===

Time	BatteryVoltage1	BatteryVoltage2	BatteryVoltage3
2013-12-05	12.0	12.1	12.2
2014-03-26	12.1	12.2	12.3
2014-07-15	11.7	11.8	11.9
2014-11-03	12.1	12.2	12.3
2015-02-22	12.2	12.3	12.4
2015-06-13	12.2	12.3	12.4
2015-10-02	12.3	12.3	12.4
2016-01-21	12.1	12.2	12.3
2016-05-12	12.0	12.1	12.2
2016-08-31*	12.0	12.1	12.2
2016-12-20*	12.0	12.1	12.2




Figure 9 - Sample Prediction, 20 Minute Discharge

SMOreg

=== Future predictions from end of training data at 45 minutes into standby events ===

Time	BatteryVoltage1	BatteryVoltage2	BatteryVoltage3
2013-12-05	11.6	12.0	12.1
2014-04-11	10.7	10.6	10.8
2014-08-16	10.2	11.7	12.0
2014-12-20	11.9	12.0	12.2
2015-04-27	11.8	12.1	12.2
2015-09-01	12.0	12.1	12.2
2016-01-05	11.6	11.9	12.0
2016-05-11	10.6	11.4	11.6
2016-09-15*	10.1	11.3	11.6
2017-01-20*	10.2	11.3	11.6

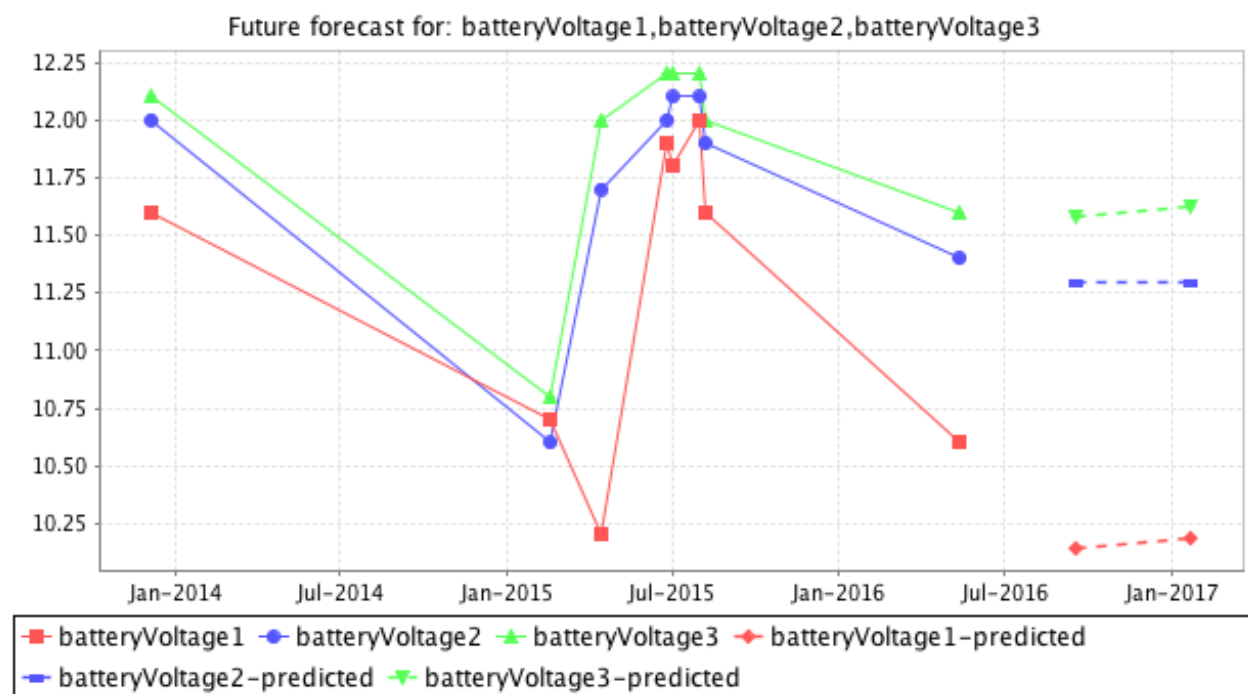


Figure 10 - Sample Prediction, 45 Minute Discharge

#### 4. Conclusions

Two approaches were presented for increasing OSP network reliability through advanced battery monitoring and analysis. The first method reviews how battery chemistry can result in disparate charge being applied to individual batteries within a battery string. The effects of this charge delta include reduced battery capacity and runtime. Remedies for detection, correction and prevention were discussed. The second method that was presented analyzed historic battery data using machine learning algorithms to determine future trends that could be used to create predictive alarms. Predictive alarms provide advanced indicators of imminent battery failures. With predictive alarms, directed PM visits can be

scheduled and appropriate action taken. Although additional algorithmic validation is required, this approach appears promising for use by operators to better direct limited PM resources.

Maintaining the integrity of the OSP network and specifically batteries, is one of the single largest operational expenses that cable operators face today. Any approach that increases network reliability, lowers operational costs, or both, in connection with OSP batteries, has the potential of significantly improving operational efficiency.

## 5. Abbreviations and Definitions

### 5.1. Abbreviations

API	application programming interface
HFC	hybrid fiber-coax
MLP	multilayer perceptron
mV	millivolt
OSP	outside plant
PM	preventive maintenance
PS	power supply
SCTE	Society of Cable Telecommunications Engineers
SVC	support vector machine
UPS	uninterruptible power supply
V	volt
Vdc	volt(s) direct current
VRLA	valve regulated lead acid

### 5.2. Definitions

<i>Coup de fouet</i>	The initial voltage dip and recovery experienced when discharging a fully charged lead acid battery
----------------------	---

## 6. Bibliography and References

- [1] Anderson-Burgett, R. a. (2014). Improving Network Reliability Through Effective Power Management. Cable-Tec Expo. Denver: SCTE - Society of Cable Telecommuncations Engineers.
- [2] Bose, C.S.C. and Beaird, T., “Battery state of health estimation through *coup de fouet*: field experience,” BattCon 2003.
- [3] IEEE Standards Association. (2012). IEEE Guide for Selection and Use of Battery Monitoring Equipment in Stationary Applications. IEEE Power & Energy Society, 16.
- [4] Witten, E. F. (2011). *Data Mining: Practical Machine Learning Tools and Techniques*, 3rd Ed. Morgan Kaufmann.
- [5] Rasmussen, C. E. (n.d.). *Gaussian Processes in Machine Learning*. Max Planck Institute for Biological Cybernetics.



- [6] Honkela, A. (2001). Nonlinear Switching State-Space Models.  
<https://www.hiit.fi/u/ahonkela/dippa/dippa.html>.
- [7] Platt, J. (1998). *Fast Training of Support Vector Machines Using Sequential Minimal Optimization*. MIT Press.

## Beyond RFoG

### PON DOCSIS Backhaul, for Smoother Migration to PON

A Technical Paper prepared for SCTE/ISBE by

Yuxin (Eugene) Dai Ph.D, Cox Communications, SCTE/ISBE Member  
6305 Peachtree Dunwoody  
Atlanta, GA 30328  
Eugene.Dai@cox.com

## Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents	37
1. Introduction	38
2. RFoG in a Nutshell	38
3. Lessons Learned from RFoG Deployments	39
4. PON DOCSIS Backhaul Architecture	41
5. Conclusions	42
6. Abbreviations	42
7. Bibliography and References	43

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - RFoG Architecture and System	38
Figure 2 - RFoG Deployment Scenarios	39
Figure 3 - Receiver Noise Floor Cause by OBI	40
Figure 4 - PON DOCSIS Backhaul	41

## 1. Introduction

The RFoG (Radio Frequency over Glass) standard was developed at SCTE under the IPS (Interface Specification) 910 around 2008, and later adopted by ANSI (American National Standard Institute) as ANSI/SCTE174 [1] in 2010. RFoG defines transmission of DOCSIS<sup>®</sup> i (Data-Over-Cable Service Interface Specifications) downstream and upstream RF signals via optical fiber directly to the customer premises. RFoG delivers DOCSIS services, e.g. HSI (High Speed Internet), Voice, OTT (over the top video), etc., over fiber ODN (optical distribution network) to the FTTH (Fiber to the Home) customers. The RFoG signals can also coexist with some of the other PON (Passive Optical Network) signals. It enables fast deployment of FTTH using existing DOCSIS headend architecture. Therefore, RFoG is considered a stepping-stone for PON-based FTTH. Now, it seems natural to assume that anything beyond RFoG would be PON. However, considering the successes and the problems observed in RFoG deployments in recent years along with the progress in PON technologies and standards, there may be another technology beyond RFoG that leads to an even smoother migration to PON. In this short paper, we first introduce RFoG in a nutshell. Then we'll briefly discuss the problems with RFoG based on deployment experiences, and finally we'll introduce a new technology and architecture that can be the next step of RFoG, the PON DOCSIS Backhaul (PDB) architecture. PDB is designed to produce lower cost, better performance, and smoother migration to PON.

## 2. RFoG in a Nutshell

RFoG can be understood as DOCSIS over fiber ODN. Figure 1 shows the RFoG protocol architecture. In HFC, the physical layer signals of DOCSIS is transmitted over HFC (hybrid fiber coax) trunk fibers in optical formats and converted to electrical signals at an optical node and then transmitted over coaxial cable to the customer premises. At a customer home, a CM (cable modem) receives the DS (downstream) DOCSIS RF signal from the coaxial cable plant and transmits the US (upstream) DOCSIS RF signal to the coaxial cable plant towards the headend. An RFoG system replaces the HFC with a fiber ODN as physical media, as shown Figure 1 (a and b). The function R-ONU (RFoG Optical Network Unit) serves to transmit DOCSIS US signals towards the headend in optical format. Since legacy set top boxes use proprietary upstream communication schemes that vary among vendors, many of the details of the RFoG specification concern how US signals are carried over the fiber.

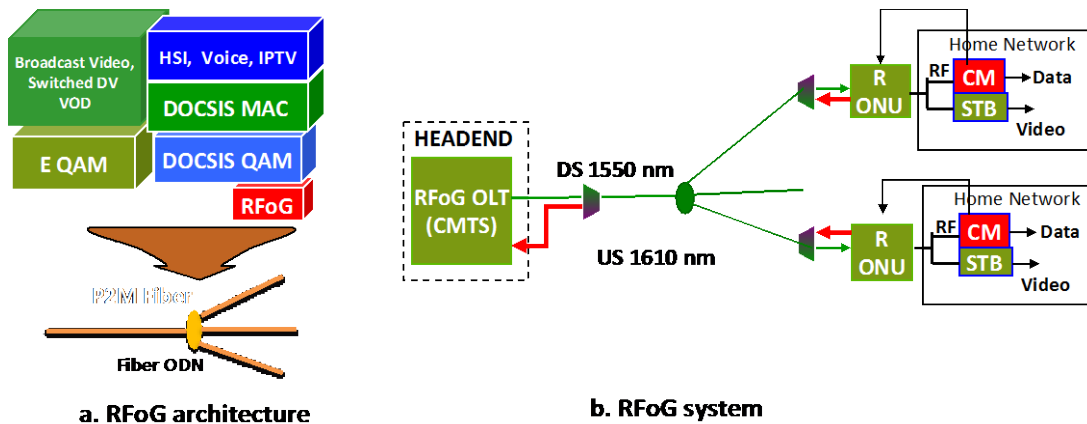


Figure 1 - RFoG Architecture and System

The RFoG system enables FTTH deployment without changing the existing expensive DOCSIS headend architecture and equipment and is thus a lower cost method of delivering FTTH services to neighborhoods that require it. Considering that DOCSIS was recently updated from DOCSIS 3.0 to 3.1 and full duplex DOCSIS is now under consideration, we predict that DOCSIS headends will be with us for a fairly long time. Therefore, the RFoG type of solution may be needed longer than might have been previously expected.

### 3. Lessons Learned from RFoG Deployments

There have been many successful RFoG deployments in the past few years. Figure 2 shows an example deployment scenario. An RFoG system can be deployed either as a standalone system, or as an overlay with other PON technologies, such as EPON (Ethernet Passive Optical Networks) or GPON (Gigabit Passive Optical Network) When deployed standalone, an RFoG system delivers full DOCSIS services, including HSI, cable voice, OTT and RF (Radio Frequency) video, over an ODN, as shown in the upper portion in Figure 2. For simplicity, only video services are shown in Figure 2. This is the scenario for some RFoG deployments where is main motivation is fast FTTH rollout with an existing DOCSIS headend. The data rate of this type of deployment is limited primarily by the capacity of DOCSIS protocols. Another scenario for RFoG deployments is the overlay architecture. In this architecture, an RFoG system is wavelength-overlaid with PON, such as GPON as shown in Figure 2. In this type of deployment, the RFoG system mainly provides legacy DOCSIS video services, including QAM (quadrature amplitude modulation) video, OTT, with DOCSIS STB (set top box). The GPON (or EPON) provides Internet data, SIP voice, on-demand IPTV, etc. In spite of its high cost, there is still a fair amount of RFoG & PON overlay deployments due to the slow progress in migrating to all IP Headend and the large numbers of DOCSIS customer premises equipment (CPE) in the field.

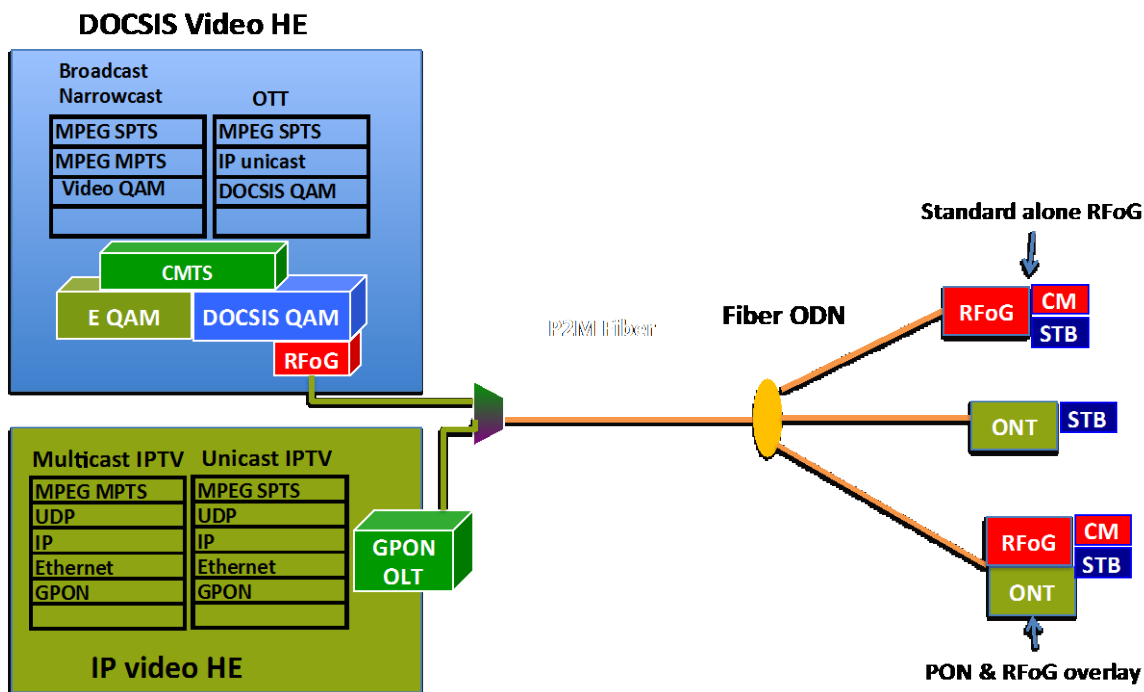


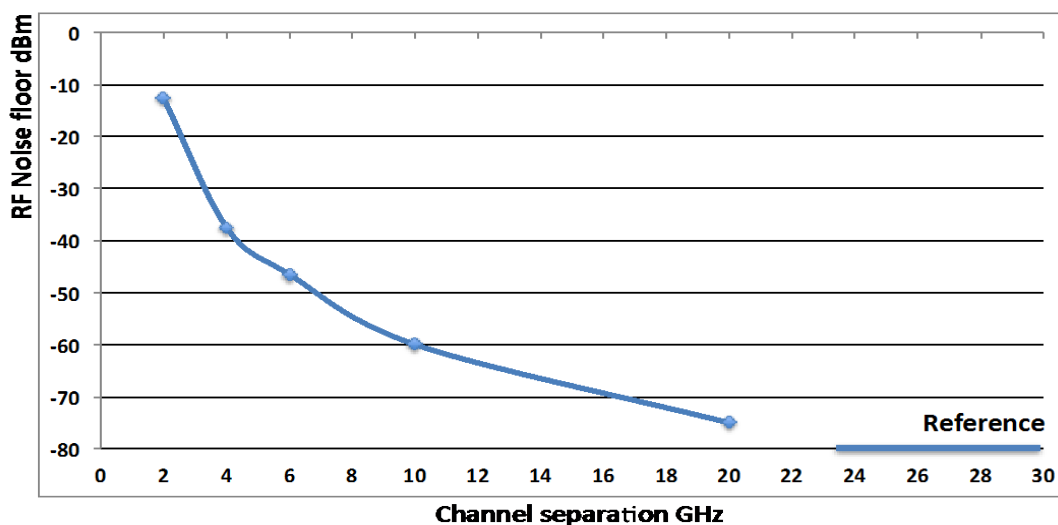
Figure 2 - RFoG Deployment Scenarios



As we discussed above, a stand-alone RFoG system, even though it is a FTTH solution, cannot provide more data rate beyond the underlying DOCSIS system capacity. And the overlay solution is very high in CPE cost; the cost of R-ONU is as high as a PON ONU if not higher, due to the low production volume. However, it may be possible to robustly transmit higher orders of QAM over RFoG than over conventional HFC, especially compared to long amplifier cascades. Again, the limit is set by the underlying DOCSIS technology, it is just that the RFoG channel may have fewer impairments than a legacy HFC channel.

Another problem with the RFoG solution is the performance in high traffic conditions. It is known that an RFoG system could have OBI (Optical Beats Interference) that is caused by optical signals colliding from the R-ONU transmissions when traffic loads are heavy. OBI is an impairment where optical signals nearby in frequency from two or more sources are received by a photo detector at the same time. The root cause of OBI in an RFoG system is the two-dimensional multiple channel DOCSIS scheduling in time and frequency and the one-dimensional R-ONU transmissions. The fundamental nature of the DOCSIS scheduler, i.e. 2-dimensional scheduling, is designed for control of multiple electrical channel transmissions. Therefore, it makes OBI hard to avoid in an optical RFoG system when traffic is heavy.

Figure 3 shows a simulation result on the receiver noise floor caused by OBI. We simulated two simultaneously transmitting R-ONUs with different frequency separation. We can see from Figure 2, at about 24 GHz separation, the OBI noise floor approaches the reference noise floor. At 10 GHz separation, the OBI penalty is about 20 dBm. Therefore, the OBI effect will be significant for separation less than 10 GHz. Since the wavelengths of all the R-ONUs are centered at 1610 +/-10 nm with a Gaussian-like distribution, the chance for a high noise floor OBI event is not insignificant. And OBI will be worse when there are more than 2 simultaneous transmissions.



**Figure 3 - Receiver Noise Floor Cause by OBI**

The impacts of OBI have been observed in the field. When the high noise floor event happens, modulation errors at the DOCSIS PHY layer are observed. If uncorrectable, DOCSIS MAC layer errors will occur. At this point, higher layer errors such as TCP/IP could cause TCP errors. The TCP protocol could wrongly interpret the TCP error as the network becomes congested, and would subsequently

reduce TCP window size. This is analogous to a snowball effect and can seriously reduce service layer throughput.

Various preparatory OBI mitigation solutions have been proposed, and vendors have implemented a few. But none of them solve the root problem without introducing other problems. An extreme example is a method that involves OEO (Optical Electrical Optical) at the optical power when the signals from R-ONUs combine.

#### 4. PON DOCSIS Backhaul Architecture

RFoG, as we have discussed, is an overlay solution with PON. That is to say it is an out-of-band solution. Now however, consider the following in-band solution: Figure 4 shows the PON DOCSIS Backhaul (PDB) [2] architecture. In a PDB system, in one implementation example, the upstream RF signal from a CM would be digitally sampled with an A/D (Analog to Digital) unit and packed in a proper frame, for example an Ethernet frame, for transmission. The Ethernet frames that contain digitalized CM US signals are then transported upstream as ONU data payloads to the PON OLT where a D/A (Digital to Analog) unit restores the digital signal back to the original analog waveform and sends it to a CMTS receiver for further processing.

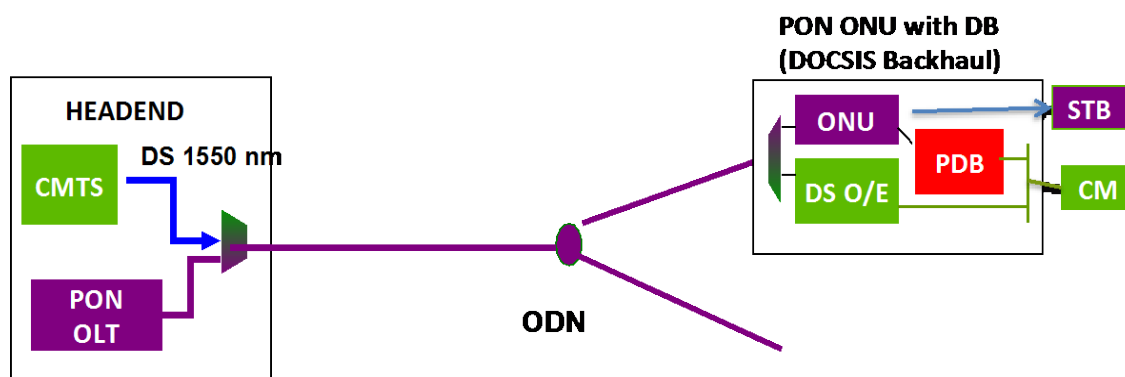


Figure 4 - PON DOCSIS Backhaul

The PDB provides the same functions as RFoG. However, from performance, cost, and convergence point of view there are major differences.

First, the PDB architecture completely eliminates OBI. Since the CM DOCSIS RF signal is backhauled to the headend through the PON data channel and TDM PON has built-in mechanisms to prevent optical signals from colliding, therefore OBI doesn't exist in PDB.

Secondly, PDB has the potential to have much lower cost than the RFoG and PON overlay solution. The RFoG system becomes a module in the PDB architecture.

Thirdly, the PDB architecture aligns better with PON migration than RFoG, since RFoG is an overlay solution and PDB is an integrated solution. The PDB provides a smooth migration from DOCSIS infrastructure to PON.

As the industry moves to symmetric 10 Gb/s PON, for example IEEE 10G EPON and FSAN/ITU-T XGS-PON, there is enough upstream bandwidth for an ONU to carry the digitalized DOCSIS technology upstream signal. The PDB method could be implemented to backhaul DOCSIS with RF carriers (40 MHz - 80 MHz bandwidth) as previously discussed. This method is similar to the CPRI (Common Public Radio Interface) mobile front-haul. Since the data rate needed for one service group (one 6 MHz channel) is about 18 Mb/s, it is even feasible for 2.5 Gb/s PON. The PDB method could also be implemented to backhaul DOCSIS without RF carriers. This method is similar to mobile backhaul in that it is more efficient in using PON upstream bandwidth. Although this method it is more complicated, it is nonetheless technically feasible.

The PDB method is largely PON protocol agnostic, it can be implemented in EPON family of standards (10G EPON and the coming 25G/100G EPON) and GPON family of standards (GPON, XG-PON, XGS-PON and NG-PON2). The PDB architecture can be standardized at various standard originations such as SCTE, and so on.

## 5. Conclusions

The PON DOCSIS Backhaul architecture is equivalent with RFoG in function, but has many advantages over traditional overlay solutions. It completely eliminates OBI, provides a smoother migration from DOCSIS to PON, and has the potential to have a much lower system cost than RFoG and PON overlay. There are two methods to implement PDB, either backhauling DOCSIS signal from a cable modem with, or without RF carriers. The PDB architecture and method can be standardized for interoperability.

## 6. Abbreviations

A/D	analog to digital
ANSI	American National Standard Institute
CPE	customer premises equipment
CPRI	common public radio interface
CM	cable modem
D/A	digital to analog
DOCSIS	Data-Over-Cable Service Interface Specifications
DS	downstream
EPON	Ethernet passive optical networks
FTTH	fiber to the home
GPON	Gigabit passive optical network
HFC	hybrid fiber-coax
HSI	high speed Internet
IPS	[SCTE] Interface Practices Subcommittee
OBI	optical beat interference
ODN	optical distribution network
OTT	over-the-top
PDB	PON DOCSIS backhaul
PON	passive optical network
QAM	quadrature amplitude modulation
RF	radio frequency

RFoG	radio frequency over glass
R-ONU	RFoG optical network unit
STB	set top box
US	upstream

## 7. Bibliography and References

[1] “Radio Frequency over Glass Fiber-to-the-Home Specification” ANSI/SCTE 174 2010

[2] Yuxin Dai, “Unified Residential and Business Services Access Network Architecture for MSOs”, SCTE Cable-Tec EXPO proceedings 2010

---

<sup>i</sup> DOCSIS is a trademark of CableLabs

# Next-Generation Node Splitting Architecture with Unidirectional DWDM Digital Return Path

A Technical Paper Prepared for SCTE/ISBE by

Gaylord Hart  
Director MSO Market Segment, Infinera  
140 Caspian Ct.  
Sunnyvale, CA 94089  
ghart@infinera.com  
303.840.9381



## Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents	45
1. Introduction	46
2. HFC Nodes and Node Splits	46
3. The Primary Hub and Digital Return Architecture	47
4. Challenges with Passive Digital Return Design	49
5. Transponder-based Digital Return Architecture	50
6. Commercial Services Overlay and DAA Architecture	54
7. Migration to 100G Metro Core	55
8. Conclusions	56
9. Abbreviations	56
10. Bibliography and References	57

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - HFC Access Architecture With Secondary Hubs and HFC nodes	47
Figure 2 - HFC Architecture: Primary Hub CMTS with Digital Return	48
Figure 3 - Passive WDM Digital Return Architecture	49
Figure 4 - Node Splitting: Passive Digital Return Design	50
Figure 5 - Node Splitting with Transponder Based Digital Return	51
Figure 6 - Transponder Digital Return Architecture Reduces Fiber Construction	52
Figure 7 - Transponder Digital Return Architecture Improves Improves Network Manageability	53
Figure 8 - Transponder Digital Return Architecture Improves Optical Performance	53
Figure 9 - Network Migration to Commercial Services and Remote PHY	54
Figure 10 - Migration to 100GE Transport with Integrated L2 Aggregation and Switching	55

## 1. Introduction

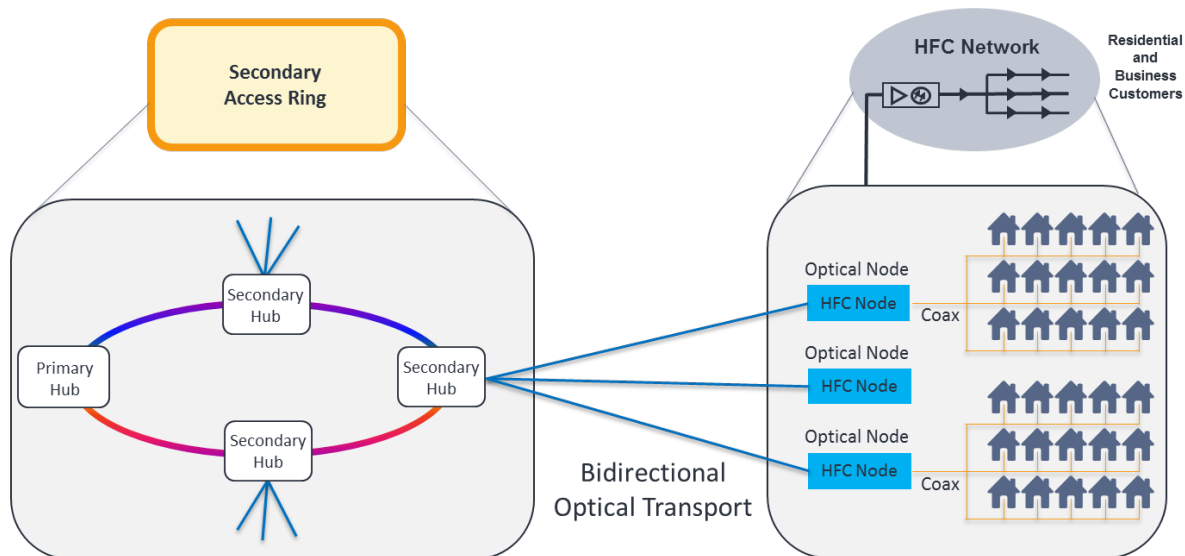
As available HFC (hybrid fiber-coax) bandwidth is consumed ever more quickly by residential users, node splitting (where a single OTN, or optical transition node, is split into two OTN nodes with the net impact of doubling the amount of narrowcast bandwidth available per home) has been a common technique used by cable operators to relieve the bandwidth bottleneck for narrowcast services as nodes begin to reach congestion. Legacy access plant architectures, however, can limit the range of options or feasibility of node splitting, requiring either expensive fiber builds or plant upgrades.

This paper examines HFC architectures where the CMTS (cable modem termination system) or edge QAM (quadrature amplitude modulation) modulator is located in a primary hub (not in a secondary hub closer to the subscriber) and where a digital return path is used from the OTN to the primary hub. In this architecture, the secondary hub serves as an optical aggregation and pass-through hub between the primary hub and the OTN, and the digital return path is often limited by a lack of fiber resources between the primary and secondary hubs and between the secondary hubs and the OTNs. Additional limitations are frequently imposed by a lack of support at the OTN for DWDM (dense wavelength division multiplexing) or a sufficiently large number of DWDM lambdas to provide efficient usage of existing fibers. Optical loss budgets may suffer from additional limitations as the digital return path lambdas are multiplexed for more efficient fiber transport.

This paper explores a novel architectural solution which mitigates the legacy architecture constraints by using a unidirectional, transponder based DWDM digital return solution between the primary and secondary hubs. This architecture supports a greater number of lambdas per fiber between these hubs (even with a limited set of digital return lambdas from the HFC node itself) and overcomes loss budget constraints often encountered when a node must be split. This architecture has the additional benefit of supporting wavelength based commercial services with the addition of a transponder based downstream digital transmission path in parallel to the upstream path, and this architecture can support the migration to Remote PHY and MAC/PHY architectures by providing low cost bi-directional 10G links to HFC nodes when needed and without the need for additional chassis or plant upgrades.

## 2. HFC Nodes and Node Splits

The HFC node is the point at which optically transmitted signals are converted to electrical signals and then fanned out to subscribers over coaxial cable. Over the years, the number of subscribers served per HFC node has dropped. Where MSOs once fed 2,500 homes from a single node, a more common range now in the U.S. is 400-500, with fiber-deep architectures tending toward 100. A typical HFC access network is shown in Figure 1.



**Figure 1 - HFC Access Architecture With Secondary Hubs and HFC nodes**

HFC node splitting is a plant upgrade whereby a single HFC node is split into two new nodes, in effect doubling the amount of narrowcast bandwidth available per subscriber. A node need only be split when the traffic on it is sufficiently high. Node splits can be challenging based upon the type of node hardware already in place and the number of fibers serving that node. If the proper hardware and fibers are in place, a node split can be inexpensive and quick.

Node splitting may be accomplished in two ways:

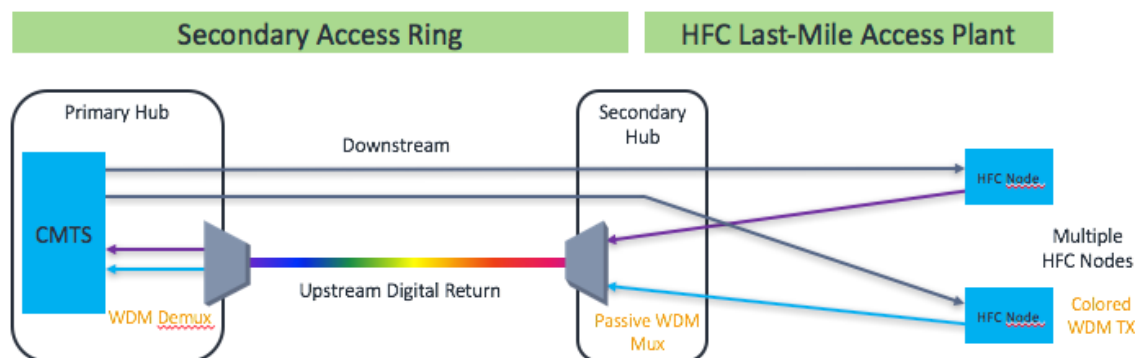
- In the first case, a new node is added adjacent to an existing node, and the subscribers on the coaxial side of the original node are then split between the original and new node. This is typically the approach used when an existing legacy node does not support segmentation.
- In the second case, a segmentable node (either two-way or four-way) is used to achieve node splitting. A segmentable node contains the internal electrical connections for splitting the coaxial subscribers on the coaxial side of the plant and internal slots for inserting additional optical transmitters and receivers each time the node is split.

The basic procedure itself entails connecting the split node to new fiber or adding additional coarse or dense wavelength division multiplexing (CWDM or DWDM) wavelengths to an existing fiber, activating an additional transmitter and receiver internally, and then splitting the existing subscribers at the node between the coaxial cable ports.

### 3. The Primary Hub and Digital Return Architecture

On the subscriber side of the HFC node is coax. On the other is fiber, which extends back to a CMTS or converged cable access platform (CCAP). Some MSOs position this equipment in a secondary hub, while others position it deeper into the core network in a primary hub. The architecture discussed here is optimized for the latter architecture, where the HFC node's optical signal transits the secondary hub to

reach the CMTS at the primary hub, and more specifically, where a digital upstream return path is utilized for the high-speed data return. This is shown in Figure 2, below. For node splitting applications, this architecture can pose challenges between the HFC node and the secondary hub and between the secondary hub and primary hub. It should be noted that there are other variations possible for this architecture, each with tradeoffs, and this paper will explore one of these approaches.

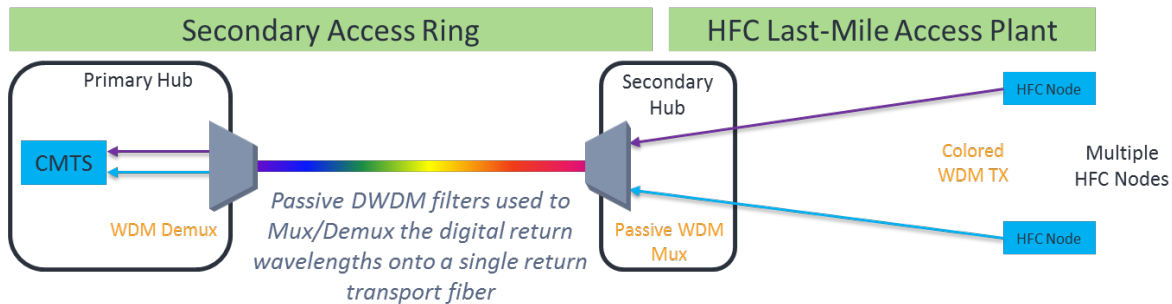


**Figure 2 - HFC Architecture: Primary Hub CMTS with Digital Return**

The CMTS uses RF analog DOCSIS<sup>®1</sup> (Data-Over-Cable Service Interface Specification) transport downstream and was originally designed for analog upstream return, as well. Early on, however, the industry saw value in replacing the HFC node's analog return lasers with digital technology. In this case, the analog-return signal is digitized at the HFC node and optically transmitted to the CMTS using digital technology, and is then converted back to an analog signal at the primary hub for handoff to the CMTS.

The benefits of digital return over analog return include enhanced noise immunity and longer transmission distances between the CMTS and HFC node. If WDM is used for the digital return path, additional efficiencies can be realized in fiber utilization between the HFC node and the secondary hub and between the secondary and primary hubs. Many MSOs have deployed a passive WDM multiplexing system to transmit multiple colored wavelengths across a single return transport fiber to the primary hub. (See Figure 3) For simplification, the subsequent sections of this paper will focus only on the digital return architecture, and not the downstream.

<sup>1</sup> DOCSIS is a trademark of CableLabs



**Figure 3 - Passive WDM Digital Return Architecture**

At the node itself, the supported WDM capabilities may vary significantly depending on node type and manufacturer. Early digital-return nodes did not use WDM, but only supported grey optics, which only allow one return signal per fiber. Current models come in different CWDM and DWDM varieties, supporting anywhere from 4 to 40 return wavelengths per fiber.

#### 4. Challenges with Passive Digital Return Design

In some legacy HFC nodes, digital return path transmitters use grey optics, which allow only one wavelength per fiber. This requires a substantial amount of fiber between the secondary and primary hubs, and between the node and secondary hub, one for each node. Other HFC node types support WDM, though they may be limited to somewhere between 4 to 40 wavelengths. A WDM digital return path greatly increases fiber efficiency, especially between the secondary and primary hubs. If WDM return is not used between the primary and secondary hubs, this architecture does not scale effectively.

The traditional digital return design also entails complex channel and node planning. Channel budgeting for each link between the node and primary hub is required because the node itself typically supports a limited number of WDM channels, and these must not overlap since they share a common fiber from the primary to the secondary hub.

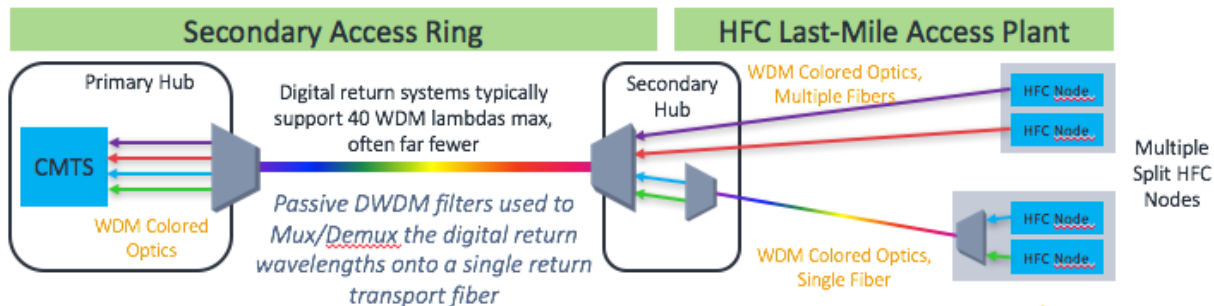
Another issue is the distance between the node and CMTS. Node splitting can adversely impact the optical power budget required to reach a primary hub, and this may result in unacceptable bit errors or even a link that is no longer supportable by the currently deployed transport solution. Finally, in the traditional architectural approach, performance monitoring (PM) can take place only in the HFC node itself and the CMTS. This makes it difficult to diagnose problems when they occur and to perform fault isolation.

Two scenarios illustrate these limitations. Where a node has spare fiber but no WDM muxing capabilities, the operator may split the node and utilize a second fiber back to the secondary hub. This case uses up scarce fiber, or may even require an expensive build to put a second fiber in place between the HFC node and the secondary hub. This is shown in the upper right section of Figure 4, below, where WDM is used to deliver the return signal to the secondary node, but over separate fibers since a local multiplexing capability is not supported at the node. These return signals are then in turn multiplexed at the secondary hub for transport to the primary hub over a single fiber.

When a node has a WDM mux, it gains efficiency by transmitting WDM-colored optics over a single fiber from the node to the secondary hub (see Figure 4, lower right section). But the multiplexing impacts



optical signal strength, as the mux and demux introduce additional attenuation. Insufficient optical power level makes a signal vulnerable to bit-errors and may result in an unacceptable optical loss budget. In these two scenarios, nothing changes from the secondary to primary hub, except for the need to transmit more wavelengths on the single fiber, as the return traffic bandwidth doubles with each node split.



**Figure 4 - Node Splitting: Passive Digital Return Design**

At some point, node splitting may lead to wavelength blocking between the secondary and primary hubs. For example, because of the limited number of return wavelengths supported at the node, two green or two red wavelengths may arrive at the secondary hub from different nodes, but these are now unable to fit on the same fiber to the primary hub because they occupy identical optical frequencies. To avoid that contention, an operator would need to build or otherwise provide additional fiber between the primary and secondary hubs – an expensive proposition.

To sum up, as currently designed, the node-splitting systems built with passive digital return design do not scale well. The drawbacks include the following:

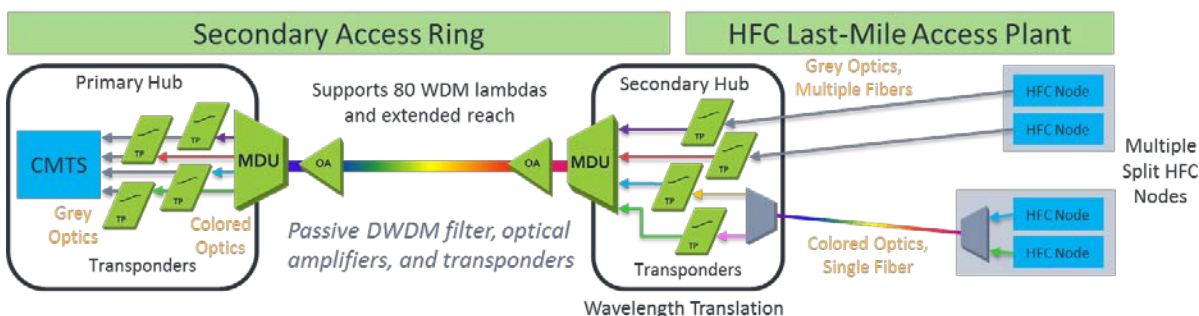
- Expensive fiber build-outs may be required between the node and secondary hub and between the secondary and primary hubs.
- Wavelength-blocking may decrease fiber utilization and require additional fibers.
- WDM multiplexing may introduce additional signal loss, with implications for increased bit error rates and reduced distance between the CMTS and node.
- Poor visibility of the optical layer and service performance makes fault isolation more difficult.

## 5. Transponder-based Digital Return Architecture

Cable operators face this dilemma: one of their most effective bandwidth enhancement tools has significant challenges when the CMTS is located in the primary hub. To overcome the limits and complexity of node splitting in this case, the digital return architecture can be adapted to address these challenges by using a unidirectional, transponder-based DWDM architecture.

In this case, unidirectional upstream transponders are installed in a chassis at the primary and secondary hubs. At the secondary hub, these transponders receive and regenerate grey or WDM optical signals from the installed base of newly split HFC nodes and then transport these return path signals to the primary hub

over a single fiber. At the primary hub, matching transponders receive these signals and translate them for delivery to the CMTS. (See Figure 5) Using transparent multi-rate DWDM transponders supporting a wide range of transport rates (e.g., from 650 Mb/s to 14 Gb/s), a wide range of digital return rates and protocols can be supported using a single transponder and architecture type.



**Figure 5 - Node Splitting with Transponder Based Digital Return**

In between the secondary and primary hubs are optical amplifiers (OAs) and multiplexer-demultiplexer units (MDUs). At the secondary hub, one MDU takes the regenerated signals from the transponder and muxes them onto long-reach DWDM optics. Boosted by an optical amplifier, up to 80 wavelengths travel upstream across a single fiber to a reciprocal MDU and chassis with transponders at the primary hub. This approach not only increases optical reach, but greatly increases spectral utilization of the fiber.

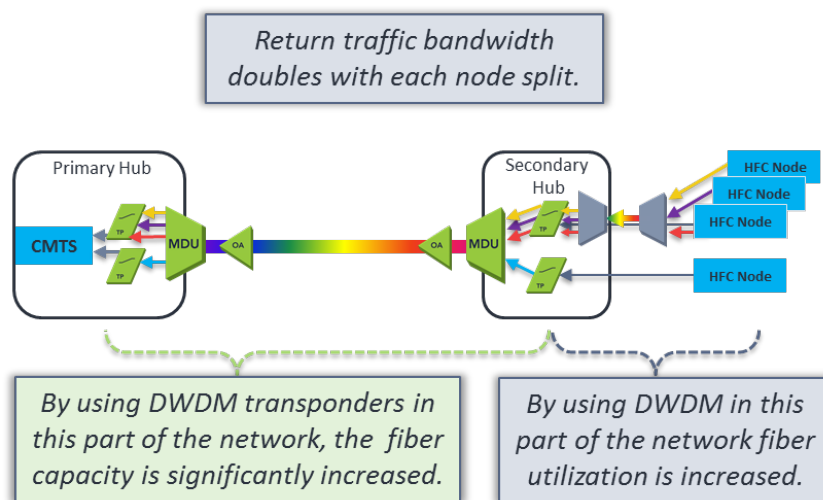
The transponder’s wavelength translation capability (the ability to accept one wavelength at the input and put this signal on a different wavelength at the output) solves the problem of HFC nodes that only support grey optics in the digital return. The transponders at the secondary hub are capable of accepting grey optics from any HFC nodes that support only a single 1550 or 1310 nm wavelength per fiber. The transponder then translates each of these waves to a unique DWDM wavelength for multiplexing and transport to the primary hub on a single fiber. Supporting up to 80 return wavelengths, this conserves a significant amount of fiber between the primary and secondary hubs.

The transponder’s wavelength translation capability also solves the pressing problem of blocked wavelengths at the secondary hub. While many HFC nodes support multiple WDM wavelengths in the digital return, most of these nodes only support a small number of wavelengths (typically 40 or much less), far below the overall capacity of the fiber between the secondary and primary hubs, and wavelength blocking is quickly reached.

For example, if the HFC node only supports 8 wavelengths in the digital return, wavelength blocking on the fiber between the secondary and primary hubs will occur when all 8 of these wavelengths are used. Using the transponder solution, 80 wavelengths may be transported on a single fiber, an increase of tenfold capacity.

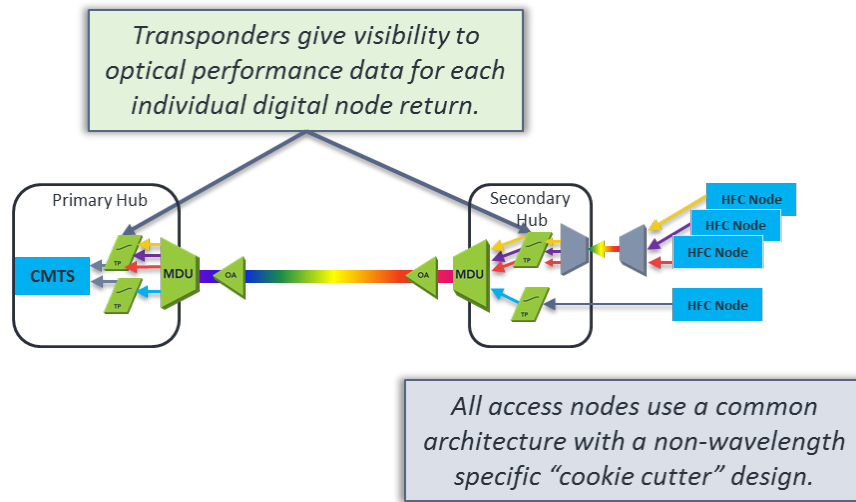
The technology required to implement this architecture is cost-effective, simple, and available today. This architecture leaves existing nodes alone and creates no impediments to other network upgrades. In fact, it is very well suited with minor upgrades to support commercial service delivery and DAA (distributed access architecture) solutions. It achieves three primary benefits:

**Reduced fiber construction.** The transponder based digital return architecture reduces or eliminates the need to install fiber in two key areas. Transponders in the secondary hub support both grey and highly efficient colored optics from the nodes, removing the need to deploy more fiber in the last mile. They also increase the capacity of secondary-to-primary hub optical transport, potentially boosting the number of available wavelengths from 8 to 80. The result is improved network economics and an obviated need for more fiber on the long hub-to-CMTS run. (See Figure 6)



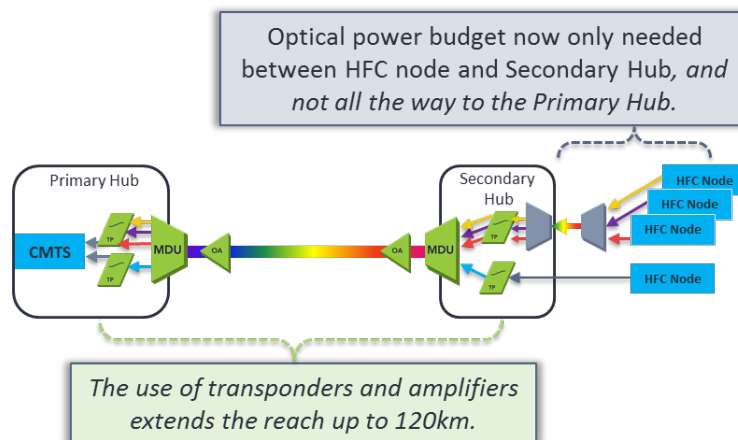
**Figure 6 - Transponder Digital Return Architecture Reduces Fiber Construction**

**Improved network manageability.** In place of complex wavelength maps unique to each network segment, the transponder based digital return solution provides a common architecture that simplifies planning. Whereas network monitoring was previously available only at the HFC node and CMTS, this solution delivers monitoring data for each digital node return at the secondary and primary hub transponders, enabling easier troubleshooting and fault isolation in areas where performance may be substandard or where failures or fiber cuts have occurred. (See Figure 7)



**Figure 7 - Transponder Digital Return Architecture Improves Network Manageability**

Improved optical performance. Without transponders, the initial optical reach budget extends from the HFC node to the CMTS in the primary hub. The transponder based digital return architecture splits the reach budget between the node-to-secondary and secondary-to-primary hub spans, enabling greater reach in each. Using 3R (regenerate, reshape, and retime) transponders, even better optical performance can be achieved. Using amplification between the secondary and primary hubs, the solution increases the available budget even further, providing a reach between secondary and primary hubs of up to 120 km. (See Figure 8) And finally, this architecture can readily support route-diverse ring architectures between the secondary and primary hubs, thus assuring a significantly higher degree of reliability against fiber cuts.




**Figure 8 - Transponder Digital Return Architecture Improves Optical Performance**

## 6. Commercial Services Overlay and DAA Architecture

The same components that enable the high-capacity transponder based digital return architecture can be extended to deliver bi-directional commercial services and support migration to Remote PHY and MAC/PHY architectures. This will not only support new revenue-generating commercial services on existing infrastructure today, but will flexibly allow future migration to distributed DOCSIS architectures (Remote PHY or MAC/PHY) without forklift upgrades.

This is accomplished by simply installing the necessary optical amplifiers and MDUs at the primary and secondary hubs and activating a new fiber between the endpoints to support downstream digital transport between the secondary and primary hubs, as shown in Figure 9. Migrating to this architecture creates a bidirectional digital path between the hubs. There is no need for a forklift upgrade because the added components plug into the existing chassis deployed to support the digital node-splitting application. This architecture supports a wide range of services, including residential DAA and commercial services.



**Figure 9 - Network Migration to Commercial Services and Remote PHY**

In what began a decade ago as a headend consolidation application to digitally interconnect hubs and headends, 10GE (10 gigabit Ethernet) deployment will soon move into the access plant as well. Driven by skyrocketing consumer bandwidth demand, transport costs that have dropped significantly over time, and a disaggregated CMTS, 10GE bidirectional digital transport will be driven closer to the customer endpoint over the next two or three years as DAA is deployed.

The enabling DOCSIS DAA technology, Remote PHY or Remote MAC/PHY, is often considered the next phase in HFC access architecture evolution. DAA moves the QAM modulator and PHY layer (and for some architectures the MAC layer, too) from the CMTS to the HFC node, replacing the current analog optical link with a fully digital optical link, thus driving digital delivery closer to the customer (see Figure 9, above). The Remote PHY nodes will typically be fed with 10GE circuits.

Remote PHY or Remote MAC/PHY is likely to be deployed within two-to-three years, not as an overlay, but once all services are digitally delivered over a common DOCSIS 3.1 architecture. According to a forecast by industry analyst IHS [1], global HFC optical node shipments are expected to more than double in the 5 years from 2014 to 2019. By 2019, IHS expects 35 percent of new physical nodes to be remote CCAP devices, 27 percent to be R-PHY units, and 23 percent to be traditional digital return nodes.

As more DAA nodes are deployed, they displace the installed base of traditional HFC nodes. As this occurs, the unidirectional digital return transponders at the secondary and primary hubs are simply



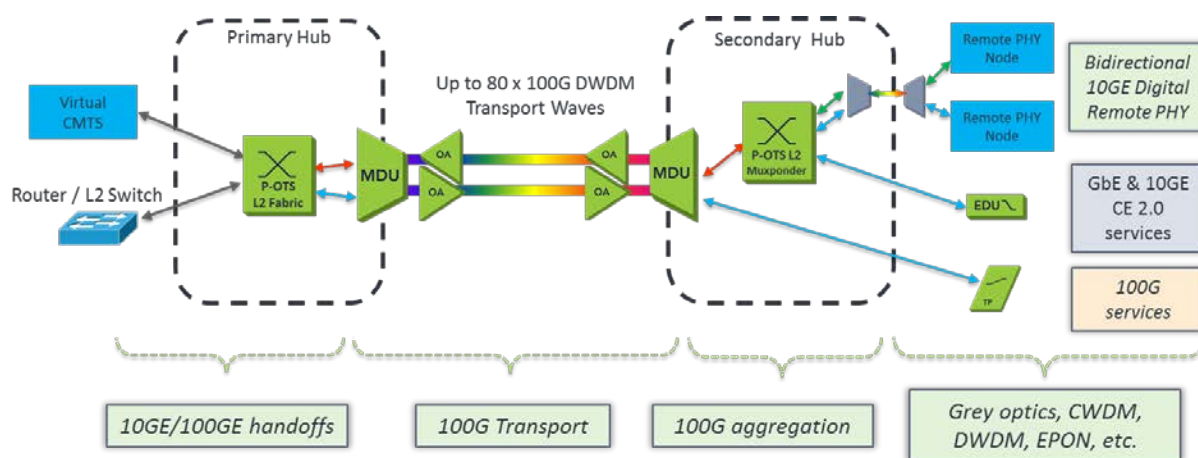
replaced with new bidirectional Ethernet muxponders. Over time the HFC network is fully converted to a DAA network, with each DAA node fed from a 10GE circuit off of a 10GE muxponder. This bidirectional capability, which pushes the digital domain much deeper into the network, will extend the life of the existing network while improving bandwidth availability and performance.

In addition to DAA, this architecture is particularly well-suited for fiber delivered Ethernet commercial services, especially Carrier Ethernet 2.0 (CE 2.0), as well as cell tower backhaul and fronthaul services. These are shown being delivered in Figure 9, above, with an Ethernet Demarcation Unit (EDU) at the enterprise premises or cell tower.

As an adjunct to the commercial services and DAA architecture shown in Figure 9, above, Layer 2 (L2) switching in the primary and secondary hubs (Layer 2 switching fabric at the primary hub and Ethernet multiplexing transponders at the secondary hub) allow Packet Optical Transport Services (P-OTS) to be efficiently aggregated, statistically multiplexed, switched, and provisioned to different types of customers, both residential and business, throughout the metro network and on a common P-OTS platform as indicated in Figure 9, above.

## 7. Migration to 100G Metro Core

As the HFC nodes are being replaced with Remote PHY nodes, the logical parallel move in the optical transport layer is to replace the 10GE waves with 100GE waves to handle the increased bandwidth capacity. This will initially and primarily take place in the core between the primary and secondary hubs, but 100GE services may also be extended to enterprise customers that require high bandwidth. As with the initially deployed architecture, there will still be 80 wavelengths available in the core network, but now each wave will carry ten times as much capacity as shown in Figure 10. Enabled by the L2 switching fabric at the primary hub P-OTS platform and Ethernet muxponders at the secondary hub P-OTS platform, cable operators will be able to efficiently aggregate and statistically multiplex a variety of Ethernet services at the metro edge, transport these via 100G waves to the primary hubs at the metro core, and then switch Ethernet traffic wherever needed in the network or hand it off to a local router or virtual CMTS.



**Figure 10 - Migration to 100GE Transport with Integrated L2 Aggregation and Switching**

## 8. Conclusions

MSOs have several ways to cope with growing competition and intense bandwidth demands on their network. These include plant upgrades, migration to IPTV and DOCSIS 3.1 specification, and deployment of DAA. Most MSOs also are continuing to split HFC nodes, a well-established, economical and effective way to increase narrowcast bandwidth available to subscribers. Depending on the architecture, this popular technique may have drawbacks, especially on the return path.

In architectures where the CMTS is located in the primary hub, existing digital return path systems do not scale well when splitting nodes. This may result in expensive fiber build-outs, increased optical insertion loss which may in turn reduce link performance, wavelength-blocking scenarios, and poor visibility of the optical layer and service performance.

The transponder based digital return architecture described in this paper addresses these shortcomings. Using an efficient and flexible transponder-based DWDM upstream design, this solution reduces or eliminates the need to install fiber in the last mile and between primary and secondary hubs. It improves network manageability, simplifying wavelength planning and delivering additional performance monitoring data for easier trouble-shooting and fault isolation. The solution also improves optical performance in node-to-secondary and secondary-to-primary spans, with attendant boosts in distance, performance, and reliability.

This solution can also easily be migrated, with the addition of a small number of components to the existing chassis and activation of a second downstream DWDM fiber, to fully bidirectional digital transport, thus enabling the ability to offer commercial services and to seamlessly migrate to a DAA access solution in the future—and all without a forklift upgrade. Additional enhancements to the architecture support 100GE waves in the metro core for increasing spectral efficiency on existing fiber and Ethernet switching and statistical multiplexing for routing and aggregating services efficiently throughout the network.

## 9. Abbreviations

3R	regenerate, reshape, and retune
10GE	10 gigabit Ethernet
100GE	100 gigabit Ethernet
CCAP	converged cable access platform
CMTS	cable modem termination system
CWDM	coarse wavelength division multiplexing
DAA	distributed access architecture
DOCSIS	Data-Over-Cable Service Interface Specifications
DWDM	dense wavelength division multiplexing
EDU	Ethernet demarcation unit
HFC	hybrid fiber-coax
ISBE	International Society of Broadband Experts
L2	layer 2
MAC	media access control
MDU	mux / demux unit

OA	optical amplifier
OTN	optical transition node
PHY	physical layer
PM	performance monitoring
QAM	quadrature amplitude modulation
SCTE	Society of Cable Telecommunications Engineers

## 10. Bibliography and References

[1] *IHS Infonetics HFC Optical Nodes Market Share and Forecast Report*, July 29, 2015. Jeff Heynen.

# Optimizing Cable Access Network Evolution for Cost, Performance, and Competition

## A CASE STUDY

A Technical Paper prepared for SCTE/ISBE by

Jean-Philippe Joseph, Lead Consultant, Bell Labs Consulting, SCTE/ISBE Member  
600-700 Mountain Avenue  
Murray Hill, NJ 07974-0636  
jean-philippe.joseph@bell-labs-consulting.com  
+1 908 582-2903

Ton Van Caenegem, Bell Labs Consulting

Luc Ongena, Bell Labs Consulting

Rex Coldren, Bell Labs Consulting

## Table of Contents

<b>Title</b>	<b>Page Number</b>
1. Abstract	60
2. Introduction	60
3. Cable Access Network Evolution Drivers	61
4. Cable Access Network Evolution Options	62
4.1. Fiber Node Splits	63
4.2. Extension of RF Spectrum and Migration to DOCSIS 3.1	64
4.3. Distributed Access Architecture (DAA)	65
4.4. Next Generation PHY Technologies	66
4.5. Fiber to the Premises (FTTP)	66
5. Cable Access Network Evolution Case Study	67
5.1. Study Objectives and Overall Methodology	67
5.2. Scenarios Considered	67
5.3. Key Inputs and Assumptions	68
5.4. Capacity Analysis Results	70
5.5. TCO Analysis	73
6. Discussion	73
6.1. Benefits	74
6.2. Challenges	74
6.3. Recommendations	74
7. Conclusion	75
8. Abbreviations	76
9. Bibliography & References	77

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Billboard/Advertised HSD Speeds for European Cable Operators	62
Figure 2 - Logical Segmentation and Physical FN Split	64
Figure 3 - Maximizing DOCSIS 3.0 Capacity	64
Figure 4 - Expanded Spectrum with High Split and DOCSIS 3.1 with Legacy Video	65
Figure 5 - Full Spectrum DOCSIS 3.1 and IP Video	65
Figure 6 - DOCSIS 3.0 Capacity, Throughput and Usage	71
Figure 7 - DOCSIS 3.1 Capacity, Throughput and Usage	72

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - HSD Service Roadmap	69
Table 2 - Case Study TCO Results Summary	73



## 1. Abstract

Cable operators globally are evolving their access networks to cost-effectively meet market demands for higher speed services. Insights derived from numerous cable access network evolution total cost of ownership (TCO) analyses indicate that a combination of architectures and technologies is often required to meet cost and performance targets and to address market expectations and growing competitive threats. The architectures, technologies and evolution alternatives are many. Centralized hybrid fiber-coax (HFC) architectures with integrated CCAP (I-CCAP) or legacy CMTS and EQAM are being contrasted by new fiber-deep distributed access architectures (DAA) with R-PHY or R-MAC/PHY. Technology alternatives such as DOCSIS<sup>1</sup> 3.0 and DOCSIS 3.1 are continuing to grow in importance and to evolve. Full-duplex DOCSIS (FDX) technology is on the horizon. Passive optical networks (PONs) are in play as HFC overlay or greenfield new-build solutions. Point-to-point optical Ethernet links serve many high end business customers today. Corresponding outside plant (OSP) and headend upgrades are occurring with increasing frequency and at considerable cost. This paper discusses the drivers for access network evolution, options operators have to evolve their networks, and a use case analysis which considers architectural and technological alternatives to identify optimal evolution scenarios that minimize one cable operator's access network TCO.

## 2. Introduction

The need to increase capacity is an absolute driver for ongoing architectural and technological evolution of cable access networks. Historical growth rates project near-term capacity exhaust on large portions of existing HFC plant. Competitive pressures are driving operators toward high speed data (HSD) product roadmaps with aggressively increasing downstream and upstream speeds. Growth in cloud-based services and the need to serve small to medium size business customers requires cable access networks to support more symmetrical downstream/upstream speeds. Continued growth of video streaming traffic, increased demand for high definition formats, and the advent of video-oriented applications like augmented reality/virtual reality (AR/VR) remind us that a continuation or exacerbation of the historical growth trends and competitive pressures on cable operators will continue into the foreseeable future.

Cable operators have an array of alternatives to consider when dealing with these realities. Many involve extending the life of the HFC plant in an evolutionary fashion, while some are more revolutionary in nature. Logical segmentation and physical splitting of fiber nodes (FNs) reduces the HSD service group (SG) size by reducing the number of households passed (HHP) by a logical or physical FN. This must be done in some cases to avoid capacity exhaust. Increasing HSD SG capacities by increasing the number of allocated DOCSIS channels can enable operators to both expand the available capacity and increase speeds per subscriber. Upgrading the usable RF spectrum to 1 GHz or 1.2 GHz may help in this regard, as can less drastic measures such as reclaiming analog video channels. Moving the RF downstream/upstream frequency split will improve upstream capacity and throughput. Implementing the more efficient DOCSIS 3.1 technology can incrementally improve throughput in a given amount of RF spectrum. DAA and fiber deeper will further improve throughput and reduce HSD SG sizes. FDX promises symmetrical downstream/upstream speeds in fiber deep architectures. And of course fiber deeper gets operators closer to fiber all the way, with PON and optical Ethernet appearing at the customer premise.

We will first discuss in more detail the drivers behind the need for access network evolution. Then we examine the most important mechanisms cable operators have or will have available to contend with this

---

<sup>1</sup> DOCSIS is a trademark of CableLabs


evolution. Finally, we present results of one case study where Bell Labs Consulting performed a capacity exhaust assessment and TCO analysis to identify optimal evolution scenarios for a cable operator.

### 3. Cable Access Network Evolution Drivers

Numerous papers have made predictions regarding growth of residential broadband traffic in various regions of the world. Video consumption is currently the primary driver, resulting from increased demand for high definition (HD) and ultra-high definition (UHD) formats, high penetration of tablets, and more time spent consuming video. Compound annual growth rates (CAGR) as high as 50 percent have been witnessed in busy hour sustained downstream traffic consumption for the North America region. Recent demonstrations of AR/VR applications, combined with 360-degree video experiences, have generated enthusiasm. Wide angle VR streams can consume 20 to 45 Mbps, but higher resolutions and bitrates would be needed for a better user experience in the future. Maturation of these new technologies and their increased acceptance, coupled with increased video consumption and the potential scale of traffic generated by the Internet of things (IoT), will drive high capacity and throughput needs and low latency requirements into 2020 and beyond.

The consequence of increased bandwidth consumption for spectrum constrained, HFC-based access networks is a progressive reduction in available SG capacity. Depending upon the amount of spectrum allocated to each, the limiting factor can be either downstream or upstream capacity. This trend eventually triggers the cable operator to execute FN splits to reduce the number of subscribers sharing upstream and downstream capacity in a SG. The FN splitting process may be repeated numerous times as capacity consumption continues to increase.

Similar to consumption growth rates, advertised HSD speeds have increased over the last two decades, coincidentally with a CAGR of 50% [1], and mainly to keep up with increased demand. Subscribers expect things such as shorter download/upload times that are associated with faster HSD service, which then improves customer satisfaction. From the operator perspective, the main driver for positioning faster HSD services in the residential market is due to competition. Higher speeds continue to be an effective way to attract new broadband customers and retain existing customers, especially in premium market segments.



**Figure 1 - Billboard/Advertised HSD Speeds for European Cable Operators**

Figure 1 shows a sample of 2015 European cable operator advertised HSD tiers. 1Gbps offers are now becoming common among the larger operators in Europe, and in North America as well. Upstream HFC capacity remains scarce due to the minimal RF spectrum allocated to it. The introduction of DOCSIS 3.1 technology includes requirements for support of expanded upstream spectrum, as it has become apparent to operators that upstream capacity available from the limited available spectrum will soon be breached. Throughput improvements of DOCSIS 3.1 also enable incremental downstream and upstream capacity gains, given the same amount of spectrum as DOCSIS 3.0 technology. But even with the requirement for support of expanded upstream spectrum in DOCSIS 3.1, symmetrical services, which Bell Labs Consulting expects to grow in coming years, will be difficult and costly to provide using DOCSIS 3.1 with the current frequency division duplex (FDD) technologies employed. Today, premium symmetrical tiers are offered almost exclusively toward enterprises – often via optical access technologies. However, recent innovations in full duplex DOCSIS, including a working demonstration by Nokia with its XG-Cable [2] are projected to enable symmetrical Gigabit services over coax deployment by the 2020 time frame.

#### 4. Cable Access Network Evolution Options

Cable operators have a variety of options to enable higher HSD speed tiers and to support their growing traffic capacity requirements. These options may be applied independently or in combination. The best alternative or set of alternatives depends not only on the operator or region of that operator’s system, but upon the situation in a given fiber node serving area (FNSA) within a particular market. That is to say, a given cable operator will potentially leverage a number of different combinations of the tools described below.

The best options for a FNSA depend upon a combination of technological and economic factors and also on an operator’s strategic vision. Operators concerned with touching customers will avoid scenarios where CPE replacements are required, or where technician visits to the home or business are necessary.

They may take a conservative approach which constrains available DOCSIS spectrum to allow continued support of legacy video, as an example. On the other end of the scale, some cable operators will go full-in to one technology for simpler operations, but at the expense of significant customer interaction, such as a full-spectrum DOCSIS and IP video strategy where every video customer requires a new set-top box. What follows are options that can be leveraged by all cable operators, but of which only a subset will normally apply to any individual operator.

#### 4.1. Fiber Node Splits

Logically segmenting and physically splitting FNs are common practices to increase the average available capacity per subscriber in a given SG. Splitting a FN provides additional capacity, reducing the number of subscribers served by the FN by the scale of the split. For example, a two way split doubles capacity by halving the number of subscribers per SG on the FN. However, what splitting nodes does not do is enable the operator to offer faster HSD speeds. The maximum speed supported on a SG remains the same as it was prior to the split. The only way to increase the maximum speed supported in a SG is to increase the throughput by raising the amount of spectrum allocated to the SG or improving the technology used to deliver the capacity (e.g., go from DOCSIS 3.0 to 3.1).

FN splits can be achieved in two ways. *Segmentations (logical splits)* reconfigure an existing physical FN to appear as a number of logical nodes. For example, a logical split can enable individual coax branches of a physical FN to be fed by their own dedicated SGs. For a physical FN with four coax segments, the maximum logical split is typically four, resulting in the creation of four logical nodes at the original FN location. Other combinations are possible, depending upon the configuration of the node and the size of the FNSA.

If the segmentation approach is not possible (e.g., maximum segmentation has already occurred at the FN) and node splitting is the preferred strategy, then *physical splits* are required. A physical split implements new FNs at locations deeper in the FNSA than the original FN, and potentially removes or reverses the direction of existing amplifiers. Physical splits require significantly higher investments since network redesign and sometimes permit acquisition are required, fiber must be extended to the new FN locations, and labor-intensive construction activities must take place.

Segmentations and physical node splits both incur corresponding capacity increases on the network side of the FN as well. Fiber capacity between the headend/hub and the FN needs to be increased according to the split ratio, both in downstream and upstream directions. In the analog feeder fiber used today for HFC systems, either dark fiber is leveraged or wavelength division multiplexing (WDM) is deployed, where each wavelength carries traffic associated with a specific SG. Additionally, for each new SG created, capacity must be added to headend/hub gear such as CMTS and EQAM, or CCAP, as well as aggregation routers north of them.

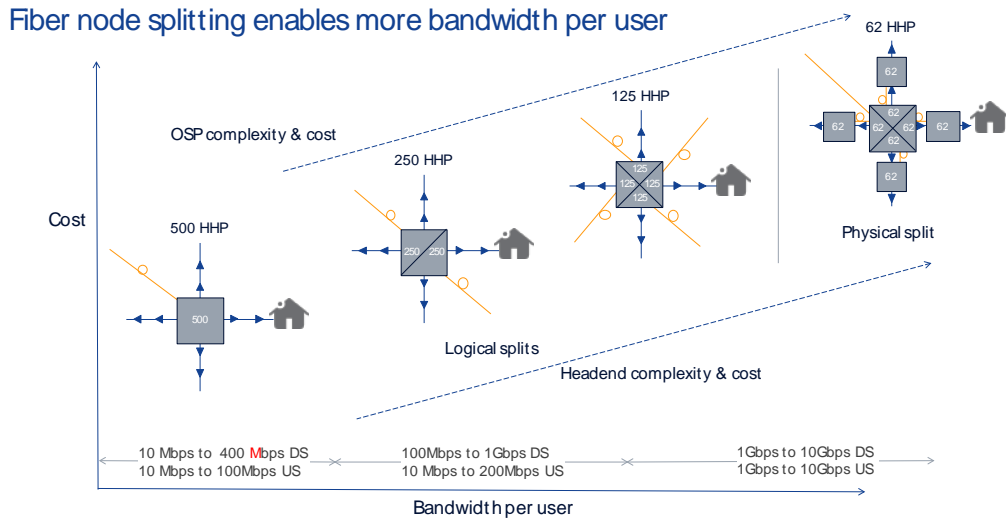


Figure 2 - Logical Segmentation and Physical FN Split

#### 4.2. Extension of RF Spectrum and Migration to DOCSIS 3.1

In practice, different vintages of HFC plants have different spectral upper bounds, limited by frequencies that FNs, amplifiers and taps in the network support. The most common upper bound today has grown from the 750 MHz range and below to around 860 MHz, with some plants already upgraded to be capable of 1002 MHz. In HFC networks the majority of FDD spectrum is dedicated to downstream transmission. Upstream spectral ranges are typically limited to the low-split mode of operation, which is up to 42 MHz in North America and 65 MHz elsewhere. DOCSIS 3.0 specifications require support for upgrade of the HFC plant to 85 MHz in the upstream for North America. This is called “mid-split”. The 42 MHz low-split enables a maximum capacity of approximately 108 Mbps (4x27 Mbps) with DOCSIS 3.0 single carrier QAM (SC-QAM) in the upstream. Mid-splits to 85 MHz deliver up to 216 Mbps (8x27 Mbps) to DOCSIS 3.0 modems. The maximum available throughput in DOCSIS 3.0 downstream HSD SGs is approximately 1.2 Gbps, this due to the practical limit of bonding 32 SC-QAM channels in a given SG. These upstream and downstream capacities are shared between all voice, HSD and IP video services of a given SG.

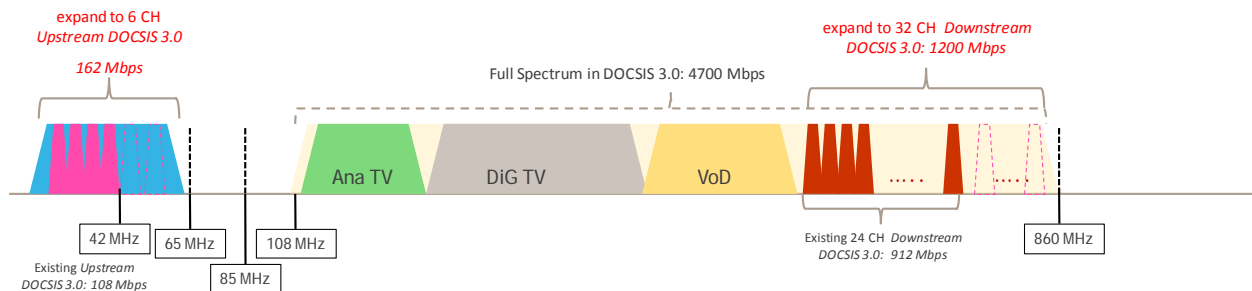



Figure 3 - Maximizing DOCSIS 3.0 Capacity

Capacity and throughput per user can be generally increased by extending the amount of RF spectrum available for a given HSD SG. In the downstream this can be realized by extending the upper frequency in the HFC outside plant to 1002 MHz/1 GHz or 1212 MHz/1.2 GHz, and/or by reclaiming video




channels within the existing downstream spectrum range and reallocating them for DOCSIS. The operator may also move the downstream/upstream frequency split to enable additional upstream capacity and throughput.



**Figure 4 - Expanded Spectrum with High Split and DOCSIS 3.1 with Legacy Video**

Some cable operators may have strategies where the full available RF spectrum is eventually allocated to DOCSIS, effectively retiring the use of RF and QAM-based video platforms and transmission gear.

Figure 5 illustrates a 1.2 GHz plant with high-split, five 192 MHz orthogonal frequency division multiplexing (OFDM) channels for downstream traffic, and two 96 MHz orthogonal frequency division multiple access (OFDMA) channels for upstream traffic. The downstream channels can be bonded together to increase capacity and throughput of a given SG, and the upstream channels can be as well.



**Figure 5 - Full Spectrum DOCSIS 3.1 and IP Video**

Upgrading from DOCSIS 3.0 to DOCSIS 3.1 enables operators to leverage additional available spectrum to increase SG capacity and throughput maximums, primarily due to the ability to perform what is essentially full spectrum bonding within the allocated upstream and downstream frequency ranges. DOCSIS 3.1 can also support the high-split upstream spectrum up to 204 MHz. At an average limit of 1024-QAM, the high split could provide as much as 1.4Gbps of throughput to a SG with OFDMA.

In addition, DOCSIS 3.1 downstream provides potential improvements in capacity and throughput due to support for modulations up to 4096-QAM, OFDM, and better error correction with low density parity check (LDPC) technology. Realistically, the maximum modulation achievable in HFC plant with the most typical of North American amplifier cascades (i.e. N+5, or “node plus 5”; a 5 amplifier cascade average per coax leg) is expected to be 1024-QAM, as determined by the carrier-to-noise ratio (CNR) levels that can be realized in such HFC plants. Reaching 4096-QAM requires a CNR enabled by deeper fiber which can support cascades of N+1 or N+0. At a theoretical limit of 4096-QAM, DOCSIS 3.1 could provide as much as 9.7 Gbps of throughput, whereas an average HFC plant would deliver 7.2 Gbps.

### 4.3. Distributed Access Architecture (DAA)

The natural evolution of HFC networks is driving fiber deeper as a consequence of the reduction in SG sizes and corresponding physical FN splits. DOCSIS 3.1 is designed to enable major gains in modulation

order that require shorter coax runs with fewer or no amplifiers in the cascade. Legacy FN technology is adequate to leverage these realities. However, the increase in number of SGs drives a need for additional space, power and cooling in the head-end and hub locations where the CMTS+EQAM and/or I-CCAP equipment resides. There is also the issue of analog optics between the hub and FN being a capacity and operational expenditure (OPEX) concern as capacity requirements on feeder fiber increase along with SG growth in a FNSA.

DAA has arisen as the next architectural advance in cable access networks. It considers all of the reasons mentioned previously as well as the conclusion that digital feeder fiber provides a ubiquitous platform which can support all fiber deeper requirements of the cable operator access network, including fiber to the premise cases. A common requirement in all DAA variants is digital fiber between the headend/hub and FN location. Where DAA variants differ is in the distribution of the functional elements of I-CCAP. One variant is R-PHY, which remotes the common PHY used by all HFC-based voice, data and digital video services into the next generation FN. R-PHY leaves the MAC at the existing head-end/hub CCAP location. The other main DAA alternative remotely locates both MAC and PHY in the FN. This alternative enables the majority of remaining CCAP functionality, which is largely implemented in software, to be placed anywhere in the cable operator network.

DAA in its different flavors is currently maturing in the specification stage at CableLabs®. Product should be appearing in operator networks by 2017.

#### **4.4. Next Generation PHY Technologies**

A spectral top-end growth to 1.2 GHz, high split operational mode, and DOCSIS 3.1 deployed via DAA still suffers from asymmetry that greatly favors downstream over upstream capacity. Reusing HFC coax assets for symmetrical Gbps services requires a new coax transmission technology that can break through the FDD limitations of today's HFC plant. In the interest of solving this problem, CableLabs has initiated the DOCSIS full duplex (FDX) workgroup to define technologies and interfaces required for high bandwidth symmetrical services. In addition, initial working demonstrations of one candidate technology for FDX have been shown recently by Nokia with XG-Cable technology. Even with these significant events, it should be realized that standards-compliant FDX technology is not expected to become commercially available for the next several years.

#### **4.5. Fiber to the Premises (FTTP)**

FTTP would be an ideal architecture for a cable operator access network, but economic and speed-of-deployment realities prevent that from happening in wholesale fashion in the near term. Cable operators are instead taking an opportunistic approach to FTTP deployments. Many business customers are served with optical Ethernet links or PON, which are also prevalent in mobile backhaul applications. Residential applications involving multiple dwelling units (MDUs) are easy targets for FTTP as well. PON deployments are becoming common in greenfield new builds and as overlays in cable operator networks where fiber competition is appearing and threatening to erode cable operator high end customer bases.

PON-based FTTP architectures have two main advantages over HFC access architectures. First, because of minimal to non-existent equipment powering requirements in the field and no impact of aging on fiber transmission quality, PON technologies exhibit very low OPEX compared to HFC networks. Second, capacity upgrades can be cost effectively and easily implemented. For example, if 2.5 Gbps downstream capacity of GPON is becoming insufficient, one can overlay GPON with XGS-PON or TWDM PON without touching the OSP. This only requires the operator to activate a XGS/TWDM PON OLT port

(WDM multiplexed with the GPON wavelength inside the hub) and to migrate some number of FTTP subscribers to the new PON variant by providing an ONT upgrade or replacement.

## 5. Cable Access Network Evolution Case Study

Following are the details of a recent case study performed for a Tier 1 North American cable operator to assess various access evolution scenarios and determine the relative TCO of each. This operator wanted a cost-effective access evolution strategy and solution roadmap to beat competitors with faster broadband service offerings in its residential markets. The methodology employed and results recognized are equally applicable to access evolution studies done for nearly any cable operator currently deploying centralized DOCSIS 3.0 HFC networks and facing increasing traffic growth, as well as competition from other broadband providers.

### 5.1. Study Objectives and Overall Methodology

The study developed an overall technology roadmap for the cable operator to cost-effectively evolve its access infrastructure, taking into account its HSD service speed tier roadmap and diverse system geography and demographics. A number of market-specific technology roadmaps were addressed in the study to characterize the most prevalent access network profiles in the operator's system, since for diverse and geographically dispersed markets it is generally not possible that a single solution fits all.

Bell Labs Consulting and the operator first came to a common understanding of the operator network architectural and technological current mode of operation. Then a list of the most probable evolution scenarios was developed. Bell Labs Consulting laid out a set of common assumptions used for studies of this nature, gleaned from past experience and associated industry knowledge, and collected input data specific to the operator. A capacity analysis was then done for each scenario to discover trigger points driving significant capital expenditure, for the purpose of developing optimal evolution strategies for each scenario. The capacity analysis' optimal evolution strategies staged the TCO analysis, which provided a feed-back loop into the capacity analysis for further evolution strategy refinements. In the end, the evolution strategies and TCOs for each scenario were compared side-by-side to determine the most efficient technology roadmap from a TCO perspective.

### 5.2. Scenarios Considered

The cable operator's baseline architecture was centralized, hub-based CMTS and EQAM with legacy, channelized video and DOCSIS 3.0-based HSD and voice services. Analog video services were still prevalent. The operator wanted to address five residential market evolution scenarios considering this common starting point, all of which maintain legacy video and were therefore prohibited from evolving toward a full-spectrum DOCSIS future.

- Scenario 1 - A centralized CMTS-based DOCSIS 3.0 architecture with 32 bonded SC-QAM channels allocated for downstream transmission per SG, with one SG per FN. Low-split (5-42 MHz) bonded channels are used for upstream transmission in the SG. This scenario keeps all subscribers on DOCSIS 3.0 for the 10-year study period. It assumes FN segmentation and physical splits to accommodate traffic growth up to the capacity limit of the downstream and upstream channels.
- Scenario 2 - A variation of Scenario 1 that introduces an I-CCAP platform in place of the CMTS and EQAM. It continues with DOCSIS 3.0, and delays the onset of FN segmentation and physical

splits by increasing the number of downstream channels at the FN to 64 after reclaiming analog video channels for that purpose.

- Scenario 3 - An evolution from CMTS-based DOCSIS 3.0 to I-CCAP based DOCSIS 3.1. When DOCSIS 3.0 capacity is nearing exhaust, I-CCAP and DOCSIS 3.1 are introduced to provide capacity relief. This is achieved by migration of enough subscribers off of DOCSIS 3.0 onto DOCSIS 3.1 so that DOCSIS 3.0 is not exhausted over the remainder of the study period. An important element of this scenario is the use of reclaimed analog video channels to provide spectrum for adding DOCSIS 3.1 downstream. As DOCSIS 3.0 subscriber counts diminish, unused SC-QAM channels are added to the DOCSIS 3.1 OFDM frequency stable.
- Scenario 4 - A two-step evolution from DOCSIS 3.0 to DAA with DOCSIS 3.1 to GPON, realized by means of an overlay PON network. Each successive architecture and technology provides capacity relief to the predecessor to the point of avoiding FN segmentations and physical splits for the study period, and subscribers are migrated in the same sequence.
- Scenario 5 - Represents a direct migration from DOCSIS 3.0 to GPON. When DOCSIS 3.0 capacity is nearing exhaust, GPON overlay is introduced to provide capacity relief such that FN segmentations and physical splits are avoided for the study period.

To deal with varying access network profiles, the study covers deployments in 12 different morphologies. The morphologies are characterized by two population density types: urban and suburban. They can be 100% single dwelling unit (SDU), 100% MDU, or a mix. Terrain types of rocky and top soil are contrasted as well. The terrain morphology impacts the TCO as it determines the OSP construction costs for scenarios that require significant new fiber extension/deployment (FN physical splits, DAA and GPON). Population densities and the mix of SDU and MDU also co-determine the roll-out of GPON fiber infrastructure costs.

### 5.3. Key Inputs and Assumptions

A key input to the case study was the cable operator's HSD service roadmap. In Table 1, the HSD service roadmap is shown to consist of five service tiers, labeled S1 to S5, offered to the residential and small and medium business (SMB) markets. The downstream speed and upstream speed advertised for each service tier are shown per year as "DSxUS." As an example, in 2015 the S1 tier has advertised downstream and upstream speeds of 5 Mbps and 1 Mbps, respectively. Each tier is shown with projected growth over time to counter expected competitive pressure on the operator.

Note that the take rates across the different tiers (not shown) are also subject to change over time, and depend very much on the target market.



**Table 1 - HSD Service Roadmap**

Downstream/Upstream	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024
S1	5x1	5x1	5x1	5x1	10x1	10x1	10x2	15x2	20x5	30x5
S2	15x2	15x2	15x2	25x5	25x5	25x5	40x8	40x8	50x15	100x15
S3	50x5	50x5	100x10	100x10	100x10	100x10	100x10	200x20	200x20	400x40
S4	200x20	300x30	300x30	300x30	500x50	500x50	600x60	600x60	800x80	800x80
S5	N/A	N/A	1Gx100	1Gx100	1Gx100	1Gx100	1.5Gx150	1.5Gx150	1.5Gx150	1.5Gx150

The list below illustrates the primary traffic and engineering inputs and assumptions which have the most significant impact on study results.

- Busy hour sustained usage per subscriber is 650 kbps downstream and 55 kbps upstream in study Year 1 (2015)
- Busy hour sustained usage per subscriber growth is at a CAGR of 55% downstream and 25% upstream
- Sustained usage and growth include Internet video streaming (e.g. Netflix, YouTube) as well as TV-everywhere video service offered by the cable operator
- Linear/broadcast video is delivered in legacy fashion on dedicated channels in the frequency spectrum
- DOCSIS traffic is spectrum constrained, with varying amounts of spectrum over time as projected by the operator and as described for each scenario
- Target (maximum) capacity utilization per SG is 70% for downstream and 60% for upstream traffic (e.g., FN segmentation/split is triggered by downstream usage reaching 70% of the utilization target for the respective SG)
- Engineering rules require that a SG have available throughput of double the maximum tier speed for downstream and 2.5 time the maximum tier speed for upstream before offering a maximum tier in a given market (NOTE: This engineering rule is not used to trigger node splits)
- Initial average FN size is 450 HHP with a HSD service take rate of 60%
- Upstream frequency use consists of two cases:
  - Low-split at 42 MHz - DOCSIS3.0 and DOCSIS3.1 users share the available spectrum using DOCSIS 3.0 technology with SC-QAM
  - Mid-split at 85 MHz – DOCSIS 3.0 uses SC-QAM in the 5-42 MHz range, and DOCSIS 3.1 uses OFDMA in the 5-85 MHz range; DOCSIS 3.1 can time division multiplex OFDMA with D3.0 SC-QAM in the 5-42 MHz range, as necessary to optimize capacity
- Downstream lower frequency depends upon low-split/mid-split condition (54 MHz for low-split and 108 MHz for mid-split); upper frequency is a maximum of 860 MHz or 1 GHz depending on the market
- DOCSIS 3.0 capacities and throughputs per SG:
  - 32 SC-QAM channels for a maximum capacity of 1.2 Gbps downstream
  - Four SC-QAM for a maximum capacity of 108 Mbps upstream
  - Channel bonding of all downstream/upstream channels means throughput = capacity
- DOCSIS3.1 capacities and throughputs per SG:
  - 32 SC-QAM channel equivalent (192 MHz) for a maximum capacity of ~1.6 Gbps downstream, assuming an average modulation approaching 1024-QAM



- Operator gradually adds OFDM channels to increase available capacity as the DOCSIS 3.1 cable modem penetrations increase and DOCSIS 3.0 modem counts decrease
- Combined DOCSIS 3.0 and DOCSIS 3.1 downstream capacities approach 2.8 Gbps

#### 5.4. Capacity Analysis Results

HFC capacity exhaust triggers cable operator action to mitigate the exhaust situation in some manner. The operator for whom this study was done projected an enormous financial impact to segment and split FNs. Delaying or eliminating FN splits was considered a key objective so that the operator could compare the anticipated impact of significant FN splitting with the most logical alternatives. Therefore, the objective of the capacity analysis was to optimize scenarios to avoid exhaust.

The approach used in the analysis was as follows. First, we determined the capacity available from each technology/architecture. For instance, for HFC/DOCSIS, the deployed capacity is based on the characteristics (e.g., number of channels, spectral efficiency, and target utilization) of the HSD SG. Second, we estimated the overall traffic/usage generated by the subscribers served on the particular technology/architecture. We then applied the cable operator's engineering rules to identify when the available capacity will be exhausted. This determined the timing for the introduction of new DOCSIS channels, fiber node split execution, and/or for the implementation of the new technology (DOCSIS 3.1, DAA, and GPON) within the different evolution scenarios.

For illustration purposes, we show the capacity exhaust assessment performed for Scenario 4 (CMTS-based DOCSIS 3.0 to DAA with DOCSIS 3.1 to GPON), as it addresses a number of important points on DOCSIS 3.0 and DOCSIS 3.1 co-existence, as well as how an overlay GPON helps the cable operator avoid incurring costly HFC OSP upgrades associated with frequent node splits. Provisioned capacity of HFC/DOCSIS in this scenario was optimized using the following guidelines:

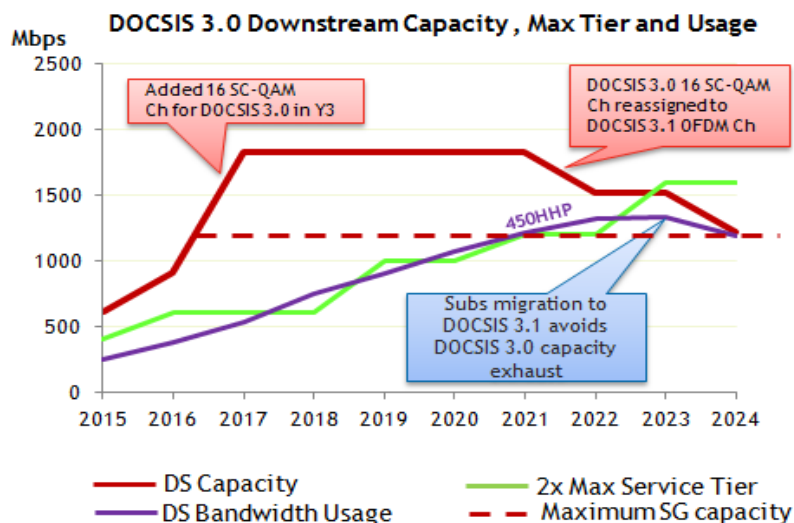
- One 32 channel DOCSIS 3.0 downstream SG per FN (16 channels in Year 1, 24 channels in Year 2, and 32 channels in Year 3) provides downstream capacity needs for the first three years
- One 16 channel DOCSIS 3.0 downstream SG was added to the FN in Year 3 for capacity relief
- One 96 MHz channel DOCSIS 3.1 downstream SG was added in Year 3 for capacity relief
- GPON overlay was added in Year 5 to completely avoid FN splits and support symmetrical downstream/upstream speed tiers, mainly offered to the enterprise market, and not shown in above service roadmap.

Low upstream usage and growth rates meant that the shared DOCSIS 3.0 /3.1 upstream capacity available with the low-split was sufficient to avoid FN splits given the projected optimal subscriber migration between technologies. However, total available bandwidth of 108Mbps did not satisfy the tier-based capacity requirement of 2.5x the maximum service tiers, which per the operator roadmap exceeded 40 Mbps in Year 4. Thus, in this scenario the operator would need to rely upon the new technology (i.e., GPON) with higher throughput to meet the projected upstream service tier requirement.

Figure 6 depicts the scenario's DOCSIS 3.0 downstream available capacity, projected usage, and required capacity from the maximum tier offering perspective. The red curve provides the available capacity based on the above assumptions. The total capacity in Year 2 is estimated at 1.2 Gbps (=32 SC-QAM channels). However, starting in Year 3, when the operator begins deploying DOCSIS 3.1 downstream, 32 additional SC-QAM equivalent channels are allocated for downstream capacity. These channels are reclaimed from

spectrum made available by discontinuing analog video transmission. These additional 32 channels are initially shared between DOCSIS 3.0 and DOCSIS 3.1 as follows:

- 16 SC-QAM channels are assigned for DOCSIS 3.0, which results in a total of 48 channels available for DOCSIS 3.0 to be split between multiple DOCSIS 3.0 SGs on the FN
- 16 SC-QAM channel equivalents assigned for DOCSIS 3.1 OFDM, which allows the cable operator to get a critical mass of DOCSIS 3.1 modems into the field before needing to turn on the full 192 MHz that is ultimately planned for DOCSIS 3.1 OFDM
- Starting in Year 8, the 16 SC-QAM channels initially assigned for DOCSIS 3.0 are gradually re-assigned to expand the OFDM spectrum
- At the end of the study period, the cable operator has the original 32 SC-QAM channels for DOCSIS 3.0, and another 192 MHz OFDM channel for DOCSIS 3.1

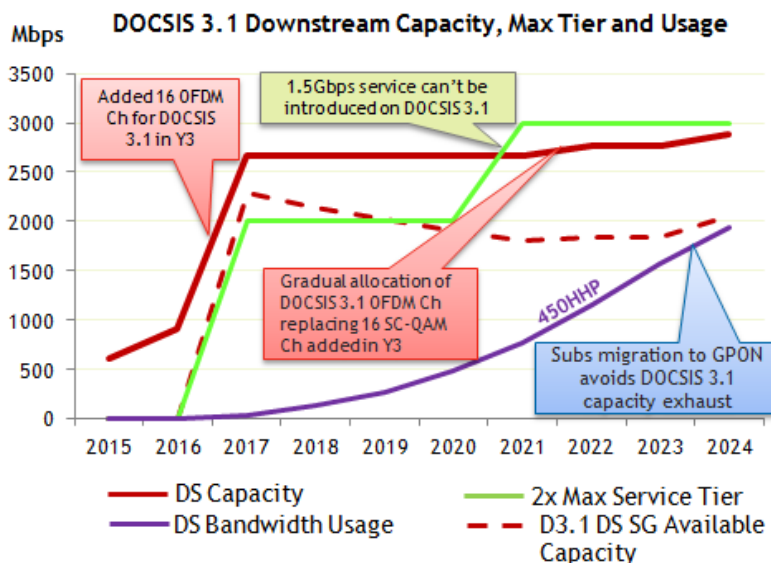


**Figure 6 - DOCSIS 3.0 Capacity, Throughput and Usage**

In Figure 6, the purple line shows the projected usage for the DOCSIS 3.0-based subscribers of a 450 HHP SG with a 60% HSD take rate. Due to migration away from DOCSIS 3.0, the usage curve stays beneath the available capacity line (the red one) for the duration of the study period. That means that there is no FN split triggered by the DOCSIS 3.0 capacity needs of the connected subscribers.

The dashed red line indicates actual maximum *throughput* per DOCSIS 3.0 SG on the FN. The 1.2 Gbps is because DOCSIS 3.0 technology caps out at 32 bonded channels downstream. The green line shows what throughput is necessary to meet projected operator speed tier roadmap speeds. By the time the green line crosses the DOCSIS 3.0 per-SG downstream capacity line, the operator would be advised to have its subscribers requiring maximum tier speeds to be migrated onto either DOCSIS 3.1 or GPON.

Figure 7 depicts the overall downstream capacity available to DOCSIS 3.0 and DOCSIS 3.1 subscribers in the SG with the solid red line. The portion of the overall downstream capacity that is available for DOCSIS 3.1 subscribers after DOCSIS 3.0 subscriber capacity requirements have been met is shown by the dashed red line.



**Figure 7 - DOCSIS 3.1 Capacity, Throughput and Usage**

The purple line shows the projected usage for the DOCSIS 3.1-based subscribers of a 450 HHP SG with a 60% HSD take rate. The usage curve stays beneath the available DOCSIS 3.1 capacity line (the dashed red one) for the duration of the study period. This means that there is no FN split triggered by the DOCSIS 3.1 capacity needs of the connected subscribers.

The solid red line in the diagram indicates actual maximum *throughput* per DOCSIS 3.1 SG on the FN; 2.8 Gbps is because DOCSIS 3.1 can bond both the OFDM and SC-QAM channels available at the FN. The green line shows what throughput is necessary to meet projected operator speed tier roadmap speeds. By the time the green line crosses the DOCSIS 3.1 per-SG downstream capacity line, the operator would be advised to have its subscribers needing maximum tier speeds migrated onto GPON.

Note that downstream capacity of GPON is 2.5 Gbps, shared among (typically) up to 32 users. PON technology has available XGS-PON and TWDM PON wavelength overlay capabilities, which can provide the next step in capacity improvement in a cost-effective and simple upgrade. PON capacity can be improved by +10 Gbps and +40 Gbps, respectively. PON upgrade was not specifically considered in this study.

General observations from the overall capacity analysis were as follows:

- The DOCSIS 3.0-only scenarios (1 and 2) could not avoid capacity exhaust and could not meet the cable operator's speed tier roadmap.
- The DOCSIS 3.0 to DOCSIS 3.1 scenario (3) could avoid capacity exhaust, but with the allocated RF spectrum could not meet the operator's projected speed tier roadmap.
- The DOCSIS 3.0 to DOCSIS 3.1 to GPON scenario (4) could avoid capacity exhaust and meet the operator's projected speed tier roadmap. However, it is the most complex scenario since it requires multi-stage migration.
- The DOCSIS 3.0 to GPON scenario (5) could avoid capacity exhaust and meet the operator's projected speed tier roadmap.

- Mid-split is not needed to avoid upstream usage triggered splits, but it is required to support maximum speed tier upstream speeds greater than 40 Mbps on DOCSIS (per 2x max speed rule).

Optimizing each scenario to minimize or eliminate FN splits prepared the study for corresponding TCO analysis.

### 5.5. TCO Analysis

**Error! Reference source not found.** presents comparative TCO per household connected (HHC) results for all five scenarios, at the end of a 10-year study period. The TCO numbers are normalized to correspond to Scenario 1 TCO, which rates the relative cost per HHC as “1”. The “% Subscribers in Yr10” columns show the percentage of subscribers that are served by a given technology in Year 10. Scenarios 1, 2 and 3 only support the cable operator’s services roadmap for up to 77% of the subscribers, which is evidenced in the total subscriber count in these columns. Hence, these scenarios will not help the cable operator achieve its business objectives given its capacity and engineering rules. The low/high normalized TCO/HHC ranges when GPON is involved (Scenario 4 and Scenario 5) are explained by the fact that different morphologies and population densities were taken into account.

We further observe that although no GPON overlay is built in Scenarios 1-3, these exhibit the highest cost per HHC. This is mainly caused by the multiple node splits associated with these scenarios over the course of the 10-year study. Scenarios 3-5 are comparable from a TCO viewpoint. However, the double subscriber migration of Scenario 4 adds risk as the cable operator will in many cases need to access the customer premise multiple times, which could potentially lead to increased churn with these subscribers.

**Table 2 - Case Study TCO Results Summary**

Evolution Scenario	TCO/HHC (Normalized)	% Subscribers in Yr10			Remarks
		DOCSIS3.0	DOCSIS3.1	GPON	
<b>Scenario 1</b> DOCSIS3.0 Only (32CH DS)	1	77%	0%	0%	<b>High Cost</b> – Roadmap not supported for 23% subs
<b>Scenario 2</b> DOCSIS3.0 Only (64CH DS)	0.83	77%	0%	0%	<b>Medium Cost</b> – Roadmap not supported for 23% subs
<b>Scenario 3</b> DOCSIS3.0 - DOCSIS 3.1	0.65	31%	46%	0%	<b>High Cost</b> – Roadmap not supported for 23% subs
<b>Scenario 4</b> DOCSIS3.0 - DOCSIS 3.1 - GPON	0.62-0.68	11%	21%	68%	<b>100% subs – Multiple-step migrations</b>
<b>Scenario 5</b> DOCSIS3.0 - GPON	0.58-0.65	22%	0%	78%	<b>100% subs – single-step migration</b>

## 6. Discussion

The case study establishes that in spectrum constrained HFC/DOCSIS environments such as those which maintain legacy QAM-based video delivery, GPON overlay and subscriber migration directly to GPON from DOCSIS 3.0 can be a relatively attractive access evolution strategy from a TCO perspective (Scenario 5). Such a scenario offers a number of benefits, but also creates challenges for the cable operator.



## 6.1. Benefits

Direct evolution from HFC/DOCSIS 3.0 to an overlay GPON helps the cable operator to:

- Avoid FN splits and related costs caused by downstream HFC/DOCSIS capacity exhaust
- Eliminate the need for mid/high-frequency splits for supporting future upstream speed tier offerings
- Adopt a technology with a cost-effective, simple upgrade path that can enable significant increases in capacity and throughput, symmetrical services, and minimal OPEX that extend well beyond the end of the 10-year study period
- Establish and/or maintain a competitive advantage with FTTP, helping to protect market share in areas where it faces competitive threats from FTTP-based service providers
- Avoid multiple subscriber migrations (e.g. from DOCSIS 3.0 to DOCSIS 3.1 to GPON)
- Moderate the pace of subscriber migration (78% of current HFC subscribers should be migrated over an eight year period starting from Year 3)

## 6.2. Challenges

Direct evolution from HFC/DOCSIS 3.0 to an overlay GPON creates certain challenges for the operator:

- Requires upfront capital investment to support GPON deployment in Year 3; this includes investment to build the fiber-based OSP and a decision on whether video service will be delivered via RF overlay or based on IPTV
- Requires new CCAP platforms in order to achieve DOCSIS 3.0 64 SC-QAM channel expansion to create a second HSD SG per FN (assumption being that DOCSIS 3.0 CMTS equipment is limited to 32 downstream channels per SG)

## 6.3. Recommendations

There is no one-size-fits-all solution alternative for cable access evolution. There will be markets where competition is lower, and there will be markets where customers are not willing or able to pay for premium offerings. Moves toward FTTP and/or DOCSIS 3.1 can be targeted first to markets where this is not the case. Our recommendations to the cable operator were as follows:

1. Leverage DOCSIS 3.0 in areas with low take rates for higher service tiers and low traffic growth
2. Leverage DOCSIS 3.1 to delay/reduce FN splits where fiber deployment costs are high or forecasted take rates for higher tier services are low
3. Low-split upstream with DOCSIS 3.0 and DOCSIS 3.1 is generally sufficient from a capacity standpoint, but will become a limiting factor on throughput without a mid-split or technology solution such as GPON
4. Unless a spectrum allocation plan can be developed to provide major increases for DOCSIS capacity and throughput growth, consider the deployment of GPON overlay, migrating the higher speed tiers to GPON, without need for transitioning first to DOCSIS 3.1 (Scenario 5)

Scenario 5 was most attractive from a TCO, risk and competitive perspective and it supports the full HSD speed tier roadmap for all of the operator's subscribers. Above all, it minimizes subscriber migrations (i.e. supports 1-step migration) and enables the cable operator to direct investments towards fiber deployment as opposed to making frequent OSP upgrades for FN splits and mid/splits.



## 7. Conclusion

The need to increase capacity is an absolute driver for ongoing architectural and technological evolution of cable access networks. Historical growth rates project near-term capacity exhaust on large portions of existing HFC plant. Competitive pressures are driving operators toward HSD product roadmaps with aggressively increasing downstream and upstream speeds. Numerous factors suggest these drivers will continue and perhaps accelerate.

Cable operators have an array of alternatives to consider when dealing with these realities. Many involve extending the life of the HFC plant in an evolutionary fashion, while some are more revolutionary in nature. In the low end of the scale are FN splitting approaches to reduce the HSD SG size by reducing the number of HHP by a FN. Increasing capacity at a FN can be done by adding more spectrum for DOCSIS. Increasing throughput on available spectrum can be done by implementing DOCSIS 3.1 and moving to DAA. Emerging technology called full-duplex DOCSIS promises symmetrical downstream/upstream speeds in fiber deep architectures. FTTP technologies such as GPON provide the most disruptive solution for upgrading capacity in areas currently serviced by HFC.

Bell Labs Consulting performs access network evolution analyses for service providers of all kinds. We presented here the methodology and results of one study performed for a Tier 1 North American cable operator. The study revolved around an HFC capacity exhaust assessment and associated TCO analysis which identified optimal scenarios with the most cost-effective evolution paths for the operator. One very striking result was that building out a GPON overlay next to a centralized DOCSIS 3.0 HFC network was an appealing solution from cost and simplicity of migration perspectives, as well as in superior support of projected HSD service tier offerings. However, experience with these types of studies reveals that there is definitely no one-size-fits-all access evolution path, as there are numerous considerations that differ between operators and even within the system and access networks of a given operator.

## 8. Abbreviations

AR/VR	augmented reality/virtual reality
bps	bits per second
CAGR	compound annual growth rate
CCAP	converged cable access platform
CMTS	cable modem termination system
CNR	carrier-to-noise ratio
DAA	distributed access architecture
DL	downlink
DOCSIS	Data-Over-Cable Service Interface Specifications
DS	downstream
EQAM	edge QAM modulator
FDD	frequency division duplex
FDX	full-duplex DOCSIS
FN	fiber node
FNSA	FN serving area
FTTP	fiber to the premises
GPON	gigabit PON
HD	high definition
HFC	hybrid fiber-coax
HHC	households connected
HHP	households passed
HSD	high speed data
Hz	hertz
I-CCAP	integrated CCAP
IoT	Internet of things
IP	Internet protocol
IPTV	IP television
ISBE	International Society of Broadband Experts
LDPC	low density parity check
MAC	medium (or media) access control
MDU	multiple dwelling unit
MSO	multiple system operator
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
OLT	optical line termination
ONT	optical network terminal
OPEX	operational expenditure
OSP	outside plant
PHY	physical layer
PON	passive optical network
QAM	quadrature amplitude modulation
RF	radio frequency
R-MAC/PHY	remote MAC/PHY
R-PHY	remote PHY

SCTE	Society of Cable Telecommunications Engineers
SC-QAM	single carrier QAM
SDU	single dwelling unit
SG	service group
SMB	small and medium business
TCO	total cost of ownership
TWDM	time and wavelength division multiplexing
UHD	ultra-high definition
UL	uplink
US	upstream
WDM	wavelength division multiplexing
XGS-PON	10-Ggigabit-capable symmetric PON (ITU-T G.9807.1)

## 9. Bibliography & References

- [1] <https://www.nngroup.com/articles/law-of-bandwidth/>
- [2] “XG Cable for HFC networks”, Nokia White Paper : <http://resources.alcatel-lucent.com/asset/195765>

The authors would like to thank several colleagues at Bell Labs Consulting for their contributions to the content and the editing of this paper. In particular, we are grateful to: Marty Glapa, Jeroen Wellen, and Enrique Hernandez-Valencia.

# Reducing the cost of Fiber-To-The-Home (FTTH) Brownfield Deployment through the Implementation of Core Extraction on an HFC Network

A Technical Paper prepared for SCTE/ISBE by

Franklin M. Lartey, Ph.D., Senior Manager of Technology, Cox Communications,  
SCTE/ISBE Member  
6305 Peachtree Dunwoody Rd  
Atlanta, GA 30328  
Franklin@Lartey.net  
305-900-6601

Rod McGinn, Executive Director, Cox Communications, SCTE/ISBE Member  
6305 Peachtree Dunwoody Rd  
Atlanta, GA 30328  
Rod.McGinn@cox.com  
404-269-3433

Nicholas DiPonzio, Vice President, Cox Communications, SCTE/ISBE Member  
6305 Peachtree Dunwoody Rd  
Atlanta, GA 30328  
Nicholas.DiPonzio@cox.com  
404-269-3451

## Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents	79
1. Abstract	80
2. Introduction	80
3. FTTH Network Deployment	81
4. Core Extraction as a Solution for FTTH Brownfield Deployments	83
4.1. Coaxial Cable structure	83
4.2. Core Extraction or Core Ejection	84
5. Case Study: Deployment of FTTH Using Core Extraction in the U.S. Midwest	85
5.1. Scenario	85
5.2. Core Ejection Implementation	85
6. Analysis and Evaluation of the Implementation of Core Ejection	86
6.1. Execution Timeframe	86
6.2. Cost per foot	86
6.3. Execution Process	87
6.4. Powering	87
6.5. Permitting	87
6.6. Weather Conditions	87
6.7. Landscape	87
6.8. Financial Summary	88
Conclusion	88
Abbreviations	88
Bibliography & References	89

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Representation of an HFC Network	81
Figure 2 - Representation of an FTTH Network	82
Figure 3 - General Constitution of a Coaxial Cable	83
Figure 4 - Core Ejection Process	84
Figure 5 - Blowing Fiber in the Aluminum Shield Serving as New Conduit	85



## 1. Abstract

Deploying fiber to the home (FTTH) to hybrid fiber-coax (HFC) serviceable households (brownfield) requires considerable investments from cable companies and suggests the abandonment of existing coaxial plants in the long run. The costs involved in such operations include those related to trenching, boring, conduit installation, fiber installation, splicing, as well as materials and labor among others. Recently, a new technology became available to cable providers who can now use their existing coaxial cables as conduits for fiber deployment. This technology is known as core extraction or core ejection.

Core extraction consists of ejecting the center conductor and dielectric of a coaxial cable and using the remaining hollow tube (aluminum tube with plastic jacket) as conduit to deploy fiber. This process strives to reduce the cost of fiber deployment as compared to a traditional build that uses trenching and boring to install fiber in conduits. The current study presents the principle of core extraction and analyzes the cost of fiber deployment using this new technology as compared to the cost of traditional FTTH deployment using trenching and boring. The findings of the study suggest that the utilization of core extraction allows a less costly deployment of fiber to homes previously passed with coaxial cable than a traditional fiber build. The study also identifies some limitations and suggests opportunities for further research.

## 2. Introduction

The deployment of FTTH is now a ubiquitous activity in the telecommunications and cable broadband industry. One main driver of FTTH deployment is the increasing demand for bandwidth to support contemporary applications. A second driver is the competitive broadband environment (Lam, 2010; Lartey, Hargiss, & Howard, 2015; Van Loggerenberg, Grobler, & Terblanche, 2012). In addition, various studies have shown that customers served with fiber are generally more satisfied than those served with other delivery media such as twisted pairs and coaxial cables (Ross, 2011). Indeed, many broadband providers such as AT&T, Centurylink, Cox Communications, Google, and Verizon offer residential Internet connections over fiber optics cable at the very high speed of one thousand megabits per second (Mbps), also known as one gigabit per second (Gbps). The deployment of this gigabit Internet service over fiber requires excessive funding as well as the installation of parallel infrastructure in addition to existing broadband plants made of twisted copper wires or coaxial cables.

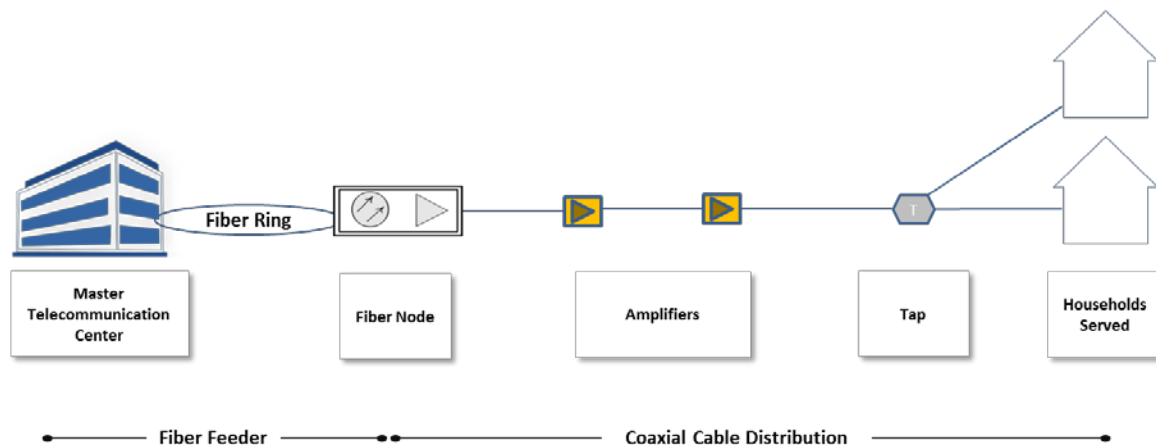
The existing cable plants of broadband companies constitute their most valuable assets in terms of cost of acquisition. In most cases, these plants were deployed over many decades, making it difficult and expensive to overlay fiber and abandon the cable plant. Sure enough, once fiber is deployed and all services are moved onto the fiber network the previous cable network becomes obsolete and is generally abandoned. This constitutes a loss of assets and money for the company. To prevent such asset loss while reducing the cost of fiber deployment, this paper discusses a method for repurposing existing coaxial plants of broadband companies through core ejection to deploy new fiber optics infrastructures.

To achieve the above stated goal, this paper will first make a brief review of FTTH network deployment, followed by a discussion of deployment challenges. It will later present a solution to these challenges in the form of core extraction. For that, the paper will do a brief presentation of the coaxial network followed by a presentation of the core extraction process. This will be followed by a case study of the implementation of core extraction and a discussion of the lessons learned from the case study, before ending with concluding remarks including opportunities for further studies.

### 3. FTTH Network Deployment

Fiber to the home consists of delivering services such as Internet, video, and voice over optical fiber going from an operator’s technical facilities to the end users’ homes (Lartey, 2015). The technical facility is also known as hub, hut, master telecommunication center (MTC), or central office (CO). There are two main scenarios for an FTTH construction, greenfield and brownfield. In a greenfield scenario, the broadband operator deploys fiber in an area where it does not already provide services. In other words, the operator initially has no existing customers in a greenfield deployment. A greenfield area can be a new franchise, a new housing development, or even the extension of an existing neighborhood.

In a brownfield scenario the operator has existing customers in the covered area. For that reason, the brownfield deployment is also known as “overbuild”. Because this study is focused on the FTTH brownfield deployment in the multiple system operator (MSO) industry, it suggests the existence of an HFC network currently serving Internet, video, and voice services to existing customers. The study also assumes that the reader has a basic understanding of the HFC network as presented on Figure 1.



**Figure 1 - Representation of an HFC Network**

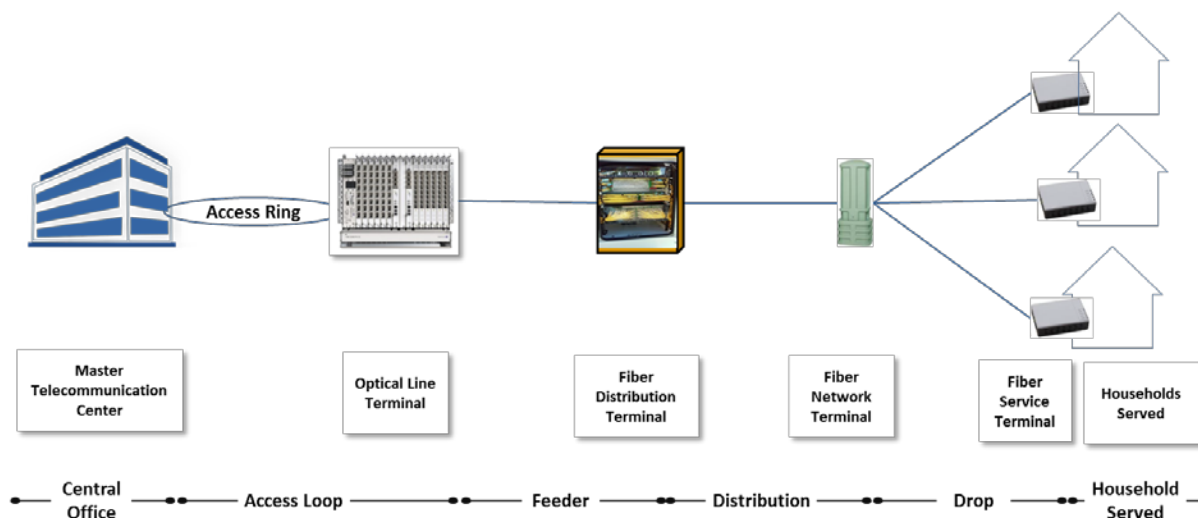
In an HFC network the coaxial cable is used to provide Internet, video, and voice services to customers. Information goes from the central office to the fiber node using fiber optic cables. In the fiber node, the optical signals are converted into electric signals and transmitted to residences using coaxial cables.

In a simplified representation, a fiber network serving homes can be organized in three sub-categories. They are presented by Brouse (2004) as headend, outside plant (OSP), and drop. Emmendorfer (2016) sees them as facility, outside plant, and subscriber. In their paper presenting the design and implementation of FTTH networks, Kadhim and Hussain (2013) identify four sub-categories constituting a FTTH network namely the MTC, feeder, distribution, and terminal subsystems.

Even though there are different architectures for FTTH deployment, this study focuses on the gigabit passive optical network (GPON) architecture. It subdivides the FTTH network into six sections as shown on Figure 2. These sections include (1) the MTC; (2) the access loop that goes from the MTC to the optical line termination (OLT); (3) the feeder which goes from the OLT to the fiber distribution terminal (FDT); (4) the distribution that goes from the FDT to the fiber network terminal (FNT); (5) the drop that goes from the FNT to the fiber service terminal (FST) where the optical network unit (ONU) could reside

to convert the optical signal on the fiber to an electrical signal for the home as explained by Wilson et al. (1999); and (6) the households served (HHS) representing homes that contain equipment providing end services to the customers through copper cable such as cat5, cat6, or even coaxial cable connected to the ONU or to a radio frequency over glass (RFoG) micro-node.

The access loop can reside within the MTC when the OLT is located in the same facility. The OLT can also be installed in the field in order to reach customers further away from the MTC as well as to reduce the number of fibers in and out of the MTC. In this case, the OLT could be installed on a pole, in a pedestal, or in a sophisticated air-conditioned cabinet. The logical representation of the FTTH network is shown on Figure 2.



**Figure 2 - Representation of an FTTH Network**

When Verizon Communications launched its FTTH network branded FiOS, it projected a capital cost of \$717 per subscriber (McGarty, 2006). After a quick analysis of FTTH costs, McGarty (2006) estimated the real cost of such deployment to be around \$3,200 per subscriber. Assuming a 30% penetration rate, the author breaks down this cost into \$2,000 for the fiber, \$700 for the drop and optical network unit (ONU), \$200 for headend and optical line termination (OLT) electronics, and \$300 for set top converters assuming two per household (HH). This was in alignment with other analysts who estimated the cost of the FiOS network to be around \$4,000 per subscriber (Farzad, 2011). Since then, companies deploying FTTH have been looking for ways to reduce their cost per subscriber in general, and most particularly the cost per household passed (HHP) with fiber.

In a subsequent 2006 report, Verizon presented a target cost to pass a home of \$850 and the cost to connect the home of \$880, with the goal of dropping them by 2010 to \$700 and \$650 respectively (Nyquist Capital, 2006). The company acknowledged that its costs in 2004 were at \$1,400 to pass a home and \$1,200 to connect a home passed, taking the cost per subscriber over \$2,600. In a report published by Sanford C. Bernstein, Kirjner, et al. (2015) estimate the cost of home served by Google Fiber to be between \$1,500 and \$2,500 depending on penetration assumptions. Similar to the FiOS deployment by Verizon Communications, FTTH deployments by cable broadband providers such as Comcast, Cox Communications, or Charter Communications which recently acquired Time Warner Cable in 2016, require significant capital investments (Kulkarni & El-Sayed, 2010).


In an article presenting the fiber access network from the perspective of a cable operator, Brouse (2004) organizes the FTTH deployment costs into three main categories, namely outside plant (OSP) deployment, headend equipment deployment, and drop. The author estimates the cost of OSP material and labor to be \$26,084.01 per mile of fiber, and that of the headend equipment and labor at \$16,118.77 per mile of OSP. The drop is estimated at \$748.00 and covers drop material and labor, including the optical network unit (ONU) and power. Brouse’s analysis suggests that the cost per mile of fiber deployment excluding drops is estimated at \$42,202.78. Reducing the cost of FTTH deployment requires the reduction of the cost per mile of fiber built. Le (2014) confirms that while the cost of electronic devices decreased gradually, that of the fiber deployment in the outside plant has stayed constant over the years. Hence, if a technology can help reduce the cost of building a mile of fiber, it would help reduce the overall cost of deploying FTTH. One such technology exists and is known as core extraction or core ejection.

#### 4. Core Extraction as a Solution for FTTH Brownfield Deployments

In order to understand the concept of core ejection, it is necessary and important to first understand the constitution of a coaxial cable plant.

##### 4.1. Coaxial Cable structure

There are many different types of coaxial cable. This study organizes them into two main categories: hardline and flexible cables. Hardline coax cables such as CommScope’s P3 500 are used for the distribution and trunk (Vikimatic, nd), while flexible coaxial cables such as Series 6, Series 11, Series 59, etc. are generally used as drop cables. Because core ejection of the drop cable is still in the research and development (R&D) process, this study focuses on the hardline coax that has a constitution as shown on Figure 3.




**Figure 3 - General Constitution of a Coaxial Cable**

In general, a hardline coaxial cable has a single copper clad aluminum wire in its core, a dielectric, a shield, and an outer jacket. The coaxial cable carries signal for Internet, video, and voice services from the HFC node to customers’ homes. As shown in the representation of the HFC network on Figure 1, fiber goes from the MTC to the fiber node. Converting the HFC plant into all fiber network such as

GPON requires conversion of the coaxial cable distribution into optical fiber. Coaxial cable distribution is generally mounted on poles or buried underground. It is easier to install aerial fiber plant, but cost prohibitive to dig out underground cables. To achieve the goal of deploying fiber underground, companies dig parallel trenches and install conduits where there are no available conduits, then push, pull, or blow fiber in the conduits while deploying supporting equipment and infrastructure such as optical line terminals (OLT), splice cases, pedestals, cabinets, and cross-connects among others. An easier solution for deploying fiber when there exists a coaxial plant consists of extracting the core of the distribution coaxial cable and replacing the ejected core with fiber.

#### 4.2. Core Extraction or Core Ejection


Also known as core ejection, core extraction of the coaxial cable is a technique consisting of removing the center conductor and dielectric of the coaxial cable to make it hollow, allowing it to serve as conduit for fiber deployment. Indeed, inside the aluminum shield of coaxial cables used in the distribution plant such as P3 500 series cables, there is a polyethylene insulation foam dielectric as shown on Figure 3. Core extraction consists of ejecting the polyethylene foam and center conductor by injecting a high velocity fluid between the aluminum shield and the polyethylene foam. This process compresses the foam to over 90 percent of its initial size, allowing it to be pushed out through the other end of the cable. After the ejection of the core, the cable becomes a hollow tube as shown on Figure 4.



**Figure 4 - Core Ejection Process**

With the aluminum shield of the coaxial cable serving as conduit after the core ejection, micro-fiber can now be blown or pushed into the previous cable plant. Depending on the application to deploy, the optical cable sheath can contain a count of up to 144 fiber strands or even more, along with copper wires that can carry power to devices on the network when deploying technologies such as RFoG. In other words, not only fiber but a hybrid or composite cable consisting of fiber and copper wires can be pushed and blown through the newly formed conduit as shown in Figure 5.





**Figure 5 - Blowing Fiber in the Aluminum Shield Serving as New Conduit**

The process of core extraction and fiber deployment presents various challenges depending on the state of the current network and the type of services provided. For instance, cable operators constantly upgrade their plants or acquire other operators thus creating duplicate plants, some of which are abandoned. These duplicate or abandoned plants are very good candidates for coax core ejection and fiber deployment.

When there is no duplicate plant, the cost of deploying fiber can be prohibitive due to the lack of conduit in the ground and the need for trenching and boring after the fact (ATF). In such cases, there is always the possibility of performing core extraction on live plants. Just like any other activity, live plant core ejection has specific challenges requiring proper planning and execution. The key to success is that a temporary cable solution utilizing fiber or coax is deployed along side of the live plant. The customer drops are moved to the temporary network which allows the coax ejection, fiber blowing and activation in the remaining coax shield or plant. Customers are then moved to the activated fiber network and the temporary cable is removed and used again on the next target area. A real example of core extraction is presented in the case study that follows.

## 5. Case Study: Deployment of FTTH Using Core Extraction in the U.S. Midwest

### 5.1. Scenario

In 2015, GPON FTTH was deployed by a broadband cable operator in a Midwestern city of the United States. The deployment was considered a brownfield application as the broadband operator chose an HFC node with existing customers. This node was selected because it had a good proportion of abandoned coaxial plant. This allowed the team to perform core extraction and customer activation with minimum disruption of existing services, thus offering a better opportunity to learn this new technology.

The selected HFC node had 510 HHP with a penetration rate around 77%. The node had 9.14 miles of coaxial plant, all of which was underground. In addition, there were over 8 miles of parallel or abandoned cable plant. This node constituted a good test case for coax core ejection and a deployment plan was created and executed for its implementation.

### 5.2. Core Ejection Implementation

The deployment of FTTH using core ejection was performed with the assistance of Deep Fiber Solutions, Inc. (DFS), a company offering a new patented method for removing the dielectric and center conductor

of a hardline coax cable. The execution of the core ejection consisted of four main phases. The first phase comprised 131 HHP, all accessible through unused parallel coaxial cables. As such, the redundant plant could be fully ejected and replaced with fiber optic cables. This allowed the company to provide gigabit services to customers over fiber while keeping their existing video and voice services on the coaxial cable.

The second phase of the implementation of core extraction involved a section of plant with 241 HHP. Only 5% of the plant had redundant coaxial cable and the rest was non-redundant, live plant. The non-redundant plant was divided into multiple single spans and core extraction was done on each of these spans one at a time. A temporary cable was installed in parallel with each single span and used to maintain customer connectivity while the core extraction was taking place. This approach gave the team adequate time to gain experience and feel comfortable performing coax core ejection.

Phase 3 of the implementation of core extraction covered 75 HHP with 10% of live plant requiring core ejection. This time, multiple spans of coaxial cable existed as opposed to single short spans in phase 2. As such, the temporary coaxial cables covered longer spans of coaxial segments requiring the use of temporary amplifiers.

Finally, with 63 HHP, phase 4 was entirely made of live plant with active customers. It also had very long spans requiring the use of temporary fiber cables as well as the use of temporary RFoG micro-nodes at every existing pedestal. The micro-nodes convert the optical signal on the fiber to RF signal transmitted to the homes using existing drop cables. Once a temporary cable was installed and customers connected to it, the crew worked from the FDT toward the end of line, lighting homes on the permanent newly blown fiber as they went. When the entire HFC segment could not be fully migrated to the temporary cable, a mini-node was used to provide HFC service continuity to the homes downstream. The completion of this phase also completed the FTTH deployment using core extraction on the selected node boundary.

## 6. Analysis and Evaluation of the Implementation of Core Ejection

### 6.1. Execution Timeframe

Once this new technology was mastered, the implementation of core extraction to deploy FTTH became a straight-forward process. Contrary to standard trenching and boring, core extraction requires far less deployment time, thus allowing savings in time and in labor costs. For instance, during the case presented here, a crew of four members ejected 4,000 feet of coaxial core in one day. During that same timeframe, the crew blew-in and spliced all required fiber. Such performance is difficult to achieve in traditional fiber construction where the crew needs to trench, install conduits, install fiber, then hand over to another crew for fiber splicing. Because core ejection occurs for the most part on live plant, it requires a quick turn-around to minimize customer downtime, hence the need for the installers to be knowledgeable splicers as well. The crew generally targets a 2,500 foot area of plant and performs all the key functions including temporary cable installation, core ejection, fiber blowing, splicing, and micro node activation in a single day. The time saved by using core extraction was not estimated in terms of cost, but the process still provided savings associated with labor cost per foot of deployed fiber.

### 6.2. Cost per foot

Overall, the cost of a foot of coax core ejection and fiber install came below 50% that of traditional construction, and even up to 85% cheaper if the crew met rock in the ground or needed to trench across

asphalt. In addition, traditional construction requires the purchase and installation of duct while core extraction uses the existing aluminum shield as conduit, eliminating the need to purchase and install conduits, thus saving costs.

### 6.3. Execution Process

While core extraction saves time and conduit costs, it has the propensity of offsetting the savings with the need for additional material and cost. Until Internet protocol (IP) video is deployed, this technology requires RFoG equipment in order to continue serving customers on live plant segments. The RFoG equipment allows legacy customers to continue having their quadrature amplitude modulation (QAM)-based services on coaxial cable from the pedestal as they did prior to core ejection. In this case, fiber runs from the MTC up to the pedestal near the customers' homes, but their services are still delivered using the same coaxial drop cable as before. When any of these customers subscribes to the gigabit Internet, a fiber drop is connected from their home to the pedestal. The RFoG equipment is also needed when the cable operator does not have an IP solution for video and phone services. In this case, the RFoG micro-nodes convert the signal from the fiber optic cable to the coaxial cable providing the appropriate signal to existing set top boxes and embedded multimedia terminal adapters (eMTA).

### 6.4. Powering

Another requirement of core extraction on live plant is the need for electrical power. The RFoG equipment is made of active devices needing power at the FDT to power the RFoG splitters that also serve as optical beat interference (OBI) mitigation devices to reduce impediments on the legacy HFC network (Aurora Networks, 2013; Al-Banna, et al., 2014). Power is also needed at the FST or pedestal for the RFoG micro-nodes. In the deployment presented here, power was provided using the existing power supply that previously fed the HFC node. In order to transport power from the FDT to the FST, composite fiber with two 16 gauge copper conductors were used. The cost of the composite fiber is higher than that of traditional fiber thus reducing some of the savings previously identified.

### 6.5. Permitting

One of the pain points in traditional constructions is the requirement to obtain various types of permits such as barricade permits, trenching permits, road crossing permits, rail crossing permits, etc. In deploying core extraction, most of these permits are not needed and in some cases, the process can be classified under plant maintenance and performed without the time and cost prohibitive process of obtaining permits. Furthermore, additional savings are gained by avoiding asphalt trenching as well as railroad and highway crossings through the use of coaxial cables already crossing those paths.

### 6.6. Weather Conditions

While traditional construction comes to a halt when the ground is hardened due to frost, core extraction can still be performed. The deployment presented here continued being executed when temperatures went below freezing levels in December 2015 and January 2016. During that same timeframe, other ongoing traditional fiber construction projects were temporarily delayed.

### 6.7. Landscape

Overall, the experience with core ejection demonstrated that customers' landscapes stayed undamaged. Compared to traditional construction of the distribution fiber in the ride of way, there were very few

customer complaints and an increase in homeowners' satisfaction. The smoothness of operations could potentially result in the gain of additional customers in the neighborhoods served and this can constitute a subject for further research.

### 6.8. Financial Summary

A financial analysis of the deployment of core ejection on the selected node demonstrated that while the material for core ejection was 35% higher than a traditional build, the associated labor cost was 64% lower. This resulted in an overall favorability of over 20% for core ejection compared to traditional FTTH deployment through PON overlay. In other words, the deployment of core ejection allowed the broadband operator to save over 20% on the overall deployment costs of FTTH. This difference is expected to be much more favorable in major metropolitan markets where conventional boring prices are higher compared to the Midwest.

## Conclusion

Over the years, cable broadband service providers have deployed hybrid fiber-coax networks that evolved to meet customers' needs for faster and more reliable Internet, voice, and video services. With fierce competition from telecommunication companies as well as new bandwidth requirements from contemporary applications, a cost effective method to deploy fiber to the home to serve residential cable broadband customers is now a reality.

This paper presented a method for deploying FTTH using core ejection in a brownfield scenario in order to reduce cost, improve network performance, and increase customer satisfaction. The paper first presented a brief review of FTTH network deployment, followed by a discussion on deployment challenges. It later presented a solution to these challenges in terms of ejecting the core and dielectric of coaxial plants and replacing them with fiber. Finally, the paper presented a case study of the implementation of core ejection and discussed lessons learned through an evaluation of the implementation.

The findings of the study confirmed that the deployment of core ejection generated cost savings of over 20% when compared to a traditional FTTH PON overlay. The savings were generated by reduced labor, time, and permitting costs. One pending question consists of finding the point at which traditional construction could become cost effective compared to core ejection, considering factors such as available conduit, aerial versus underground fiber deployment, trenching mileage, and labor costs among others. This can constitute a topic for further research in the cable broadband industry.

## Abbreviations

ATF	after the fact
CO	central office
DFS	Deep Fiber Solutions, Inc.
FDT	fiber distribution terminal
FNT	fiber network terminal
FST	fiber service terminal
FTTH	fiber to the home
Gbps	gigabits per second

GPON	gigabit passive optical network
HFC	hybrid fiber-coax
HH	households
HHP	households passed
HHS	households served
IP	Internet protocol
Mbps	megabits per second
MSO	multiple system operator
MTC	master telecommunication center
OBI	optical beat interference
OLT	optical line termination
ONU	optical network unit
OSP	outside plant
PON	passive optical network
QAM	quadrature amplitude modulation
R&D	research and development
RF	radio frequency
RFOG	radio frequency over glass

## Bibliography & References

- Al-Banna, A., Spaete, L., Vieira, A., Zorluozer, S., Schemmann, M., Cloonan, T., . . . Mutalik, V. (2014). *Handling optical beat interference in RFOG systems*. Suwanee, GA: ARRIS Enterprises.
- Aurora Networks. (2013). *Defining the evolution: Industry's first standards-compliant OBI-free RFOG solution*. Alpharetta, GA: Aurora Networks.
- Brouse, J. A. (2004). *Fiber access network: A cable operator's perspective*. Geneva: ITU-T.
- Emmendorfer, M. (2016). An economic analysis of brownfield migration CTTH vs. FTTH. *INTX: The Internet & Television Expo, 2016 Spring Technical Forum Proceedings*. Boston: NCTA, CableLabs® and the Society of Cable Telecommunications Engineers (SCTE).
- Farzad, R. (2011, March 17). *Is Verizon's bet on FiOS paying off?* Retrieved July 8, 2016, from Bloomberg Businessweek: <http://www.bloomberg.com/news/articles/2011-03-17/is-verizons-bet-on-fios-paying-off>
- Kadhim, D. J., & Hussain, N. A.-R. (2013). Design and implementation of a practical FTTH network. *International Journal of Computer Application*, 72(12), 50-56.
- Kirjner, C., de Sa, P., Rosso, A., Betcher, B., Chun, I., & Zheng, J. (2015). *U.S. Internet & U.S. Telecom: What is behind Google Fiber's continued expansion? Does it matter for incumbents?* New York: Sanford C. Bernstein & Co.
- Kulkarni, S., & El-Sayed, M. (2010). FTTH-based broadband access technologies: Key parameters for cost optimized network planning. *Bell Labs Technical Journal*, 14(4), 297-310. doi:10.1002/bltj.20417



- Lam, C. F. (2010). *FTTH look ahead - Technologies & architectures*. Mountain View, CA: Google Inc. Retrieved from <http://static.googleusercontent.com/media/research.google.com/en//pubs/archive/36936.pdf>
- Lartey, F. M. (2015). *Increasing promoters in the residential broadband service industry: Relationship between customer satisfaction and loyalty using ordinal logistic regression*. ProQuest LLC. doi:UMI 3682580
- Lartey, F. M., Hargiss, K., & Howard, C. (2015). Antecedents of customer satisfaction affecting broadband loyalty: An implementation of SERVQUAL and NPS. *International Journal of Strategic Information Technology and Applications*, 6(1), 27-42. doi:10.4018/IJSITA.2015010103
- Le, H. N. (2014). FTTH network optimization. *Journal of Telecommunications and Information Technology*, 88-99.
- McGarty, T. P. (2006). *Fiber to the home: Capital costs and the viability of Verizon's FiOS*. Florham Park, NJ: The Telmarc Group.
- Nyquist Capital. (2006, September 27). *Nyquist Capital: More signal, less noise*. Retrieved from Verizon opens the FiOS kimono: <http://www.nyquistcapital.com/2006/09/27/verizon-opens-the-fios-kimono/>
- Ross, S. S. (2011). FTTH: Life after the FiOS build. *Broadband Communities*, 34-40.
- Van Loggerenberg, S. P., Grobler, M. J., & Terblanche, S. E. (2012). Optimization of PON planning for FTTH deployment based on coverage. *Southern Africa Telecommunication Networks and Applications Conference (SATNAC)*. Fancourt in George, Western Cape, South Africa: Telkom. Retrieved from [http://www.satnac.org.za/proceedings/2012/papers/1.Access\\_Network\\_Technologies/38.pdf](http://www.satnac.org.za/proceedings/2012/papers/1.Access_Network_Technologies/38.pdf)
- Vikimatic. (nd.). *Commscope P3 500 series cable*. Retrieved from *Vikimatic Outside Plant Networks*: <http://www.vikimatic.com/coaxial-cable---connectivity/commscope-p3-500-series-cable/>
- Wilson, G. C., Wood, T. H., Stiles, A. J., Feldman, R. D., Delavaux, J.-M. P., Daugherty, T. H., & Magill, P. D. (1999). FiberVista: An FTTH or FTTC system delivering broadband data and CATV services. *Bell Labs Technical Journal*, 300-322.

# Interference-Aware Spectrum Resource Scheduling for FDX DOCSIS

A Technical Paper prepared for SCTE/ISBE by

Tong Liu, Technical Leader & CTO Cable Access, Cisco, SCTE/ISBE Member  
300 Beaver Brook Road, Boxborough, MA 01719  
tonliu@cisco.com

John T. Chapman, Fellow & CTO Cable Access, Cisco, SCTE/ISBE Member  
375 East Tasman Drive, San Jose, CA 95134  
jchapman@cisco.com

Hang Jin, Distinguished Engineer & CTO Cable Access, Cisco, SCTE/ISBE Member  
2250 East President George Bush Highway.  
Richardson, TX 75082  
hanjin@cisco.com

## Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents	92
1. Abstract	94
2. Introduction	94
3. Passive HFC Plant Model	95
4. Interference Group	97
5. IG Discovery	98
6. Duplex-Simplex Scheme	98
7. Transmission Group	99
8. Resource Sharing Model	99
1. FDX Throughput Multiplier	100
2. Factors Affecting FDX Throughput	101
3. Dynamic Load Balancing	102
3.1. Static FDX Load-Sharing	102
3.2. IG Pool and Dynamic IG-to-TG Mapping	103
3.3. Dyanmic Load-Sharing Performance	104
4. Conclusion	107
5. Abbreviations and Definitions	107
5.1. Abbreviations	107
5.2. Definitions	108
6. Bibliography and References	108

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - FDX spectrum	94
Figure 2 - A Node +0 Passive HFC Plant	96
Figure 3 - US to DS Signal Isolation Within a Segment	96
Figure 4 - CM IG Formation	98
Figure 5 - Duplex-Simplex Scheme	99
Figure 6 - FDX Resource Sharing Example	100
Figure 7 - Segment-Based Static TG Mapping	103
Figure 8 - IG Pool and Load-Based IG-to-TG Mapping	104
Figure 9 - Imbalanced Subscriber Traffic Demand	104
Figure 10 - Traffic Load-Based Dynamic IG-to-TG Mapping	105
Figure 11 - FDX Multipliers for Static vs. Dynamic Load Sharing	107

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - CM-to-CM Interference and IG Formation	97
Table 2 - Average Bit-loading in Bits per Subcarrier	101
Table 3 - Per TG Throughput with FDX Resource Sharing	101
Table 4 - FDX Throughput Multiplier	101
Table 5 - Static vs. Dynamic Load Sharing Traffic Demand and Throughput	105
Table 6 - Average Bit-loading in Bits per Subcarrier	106

## 1. Abstract

FDX (full duplex) DOCSIS technology can up to double the usage of the spectrum in cable network by allowing simultaneous downstream (DS) and upstream (US) transmissions over the same spectrum. One of the key challenges for FDX DOCSIS is the DS co-channel interference caused by neighboring CMs' US transmission which is known to be difficult to solve at pure physical layer. A MAC layer scheme has been proposed in previous work [1] to mitigate such interference with intelligent FDX scheduling. This paper further explores this solution and develops an FDX resource-sharing model. A dynamic FDX load-balancing scheme is presented that allows the resource scheduling to adapt to various user traffic patterns. Simulation result shows significant FDX throughput gain over a static FDX resource sharing approach.

## 2. Introduction

The bandwidth capacity of the cable network is a composite of many factors, including the total RF spectrum in each DS and US direction, the information efficiency with which signals are modulated onto RF carriers, the cable distribution network architecture, and efficient sharing among users with varying traffic demand. Many technologies have been adopted by the cable industry to increase bandwidth capacity. These include the recent DOCSIS 3.1 OFDM technologies for improved spectral efficiency and N+0 HFC network architecture for noise reduction and increased bandwidth per home passed.

FDX DOCSIS is the latest effort to substantially increase the US throughput using the same spectrum occupied by the DS. It is based on the fact that energy can travel simultaneously in opposite directions on a passive coax network. So instead of splitting the available spectrum into dedicated US and DS bands as in today's DOCSIS<sup>® 1</sup>, FDX DOCSIS allows the same spectrum to be used for both US and DS transmission at the same time, as shown in Figure 1. Therefore, it has the potential to double the spectral efficiency and promises the ability to deploy true symmetric services.

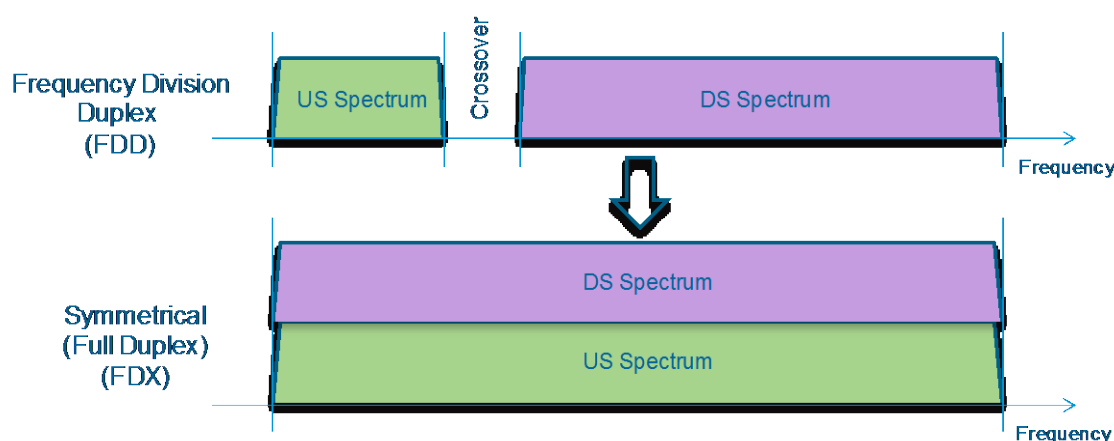


Figure 1 - FDX spectrum

The challenge for FDX DOCSIS is to mitigate the interference when the DS and US signal need to be separated, as the signal from one direction become noise/interference in the other. The following are the two key FDX technologies proposed in [1].

<sup>1</sup> DOCSIS is a trademark of CableLabs



- Echo Cancellation (EC) at RPD

To suppress the co-channel and adjacent channel interference at RPD (remote PHY device), EC is needed at RPD to suppress the DS energy from the combined DS and US energy from the coax entering the node upstream receiver. EC may be implemented in both the RF domain and the digital domain. EC in the RF domain cancels out the echo before it hits the receiver ADC. EC in the digital domain further reduces the echo with state-of-the-art DSP algorithm.

- Interference Mitigation at CM

At the CM side, the US signal interference to the received DS signal may be sourced from multiple CMs. EC cannot effectively remove the US interference from the desired DS signal. Since not all CMs' US signals may interfere with the received DS signal, the CM-to-CM interference can be effectively mitigated through intelligent RF resource scheduling. Specifically, CMs that interfere with each other are placed into an interference group (IG) and operate in the non-FDX mode as in DOCSIS today. CMs from different IGs, however, can transmit and receive over the same spectrum and at the same time. With proper IG assignment, the CM-to-CM co-channel interference is negligible relative to the desired DS signal.

This paper explores the solution for mitigating the CM-to-CM co-channel interference using intelligent resource-scheduling technique to maximize FDX system throughput. Sections 2-4 describe IGs, which are created based on CM-to-CM interference levels. Sections 5-7 explain the FDX resource-sharing model based on the duplex-simple FDX operational principle. Sections 8-9 define a FDX throughput gain quantifier and analyses the throughput affecting factors. Section 10 presents a FDX resource optimization scheme using a dynamic load-sharing scheme, and evaluates its efficiency with simulations. Section 11 concludes the paper by summing up the value of interference aware scheduling, specifically its ability to adapt to interference and traffic variations to double the usage of the spectrum in the cable network.

### 3. Passive HFC Plant Model

FDX DOCSIS assumes a passive HFC plant, which is also referred to as N+0 (node plus zero amps), as there is no amplifier after the node. An optical node typically has four inputs/outputs, each of which feeds an independent segment of the plant. Each plant segment is typically composed of a series of five to seven taps. Each tap is connected to one to eight (typically four) homes via drop cables. Drop cables are typically 50 feet to 150 feet in length. The drop cable is followed by a two-to-one splitter that does a home run to the cable modem.

Figure 2 shows an ideal model of a cable segment with five taps along the distribution line.

DS output power: 39 dBmV/6 MHz  
 US Input power: 8 dBmV/6.4 MHz

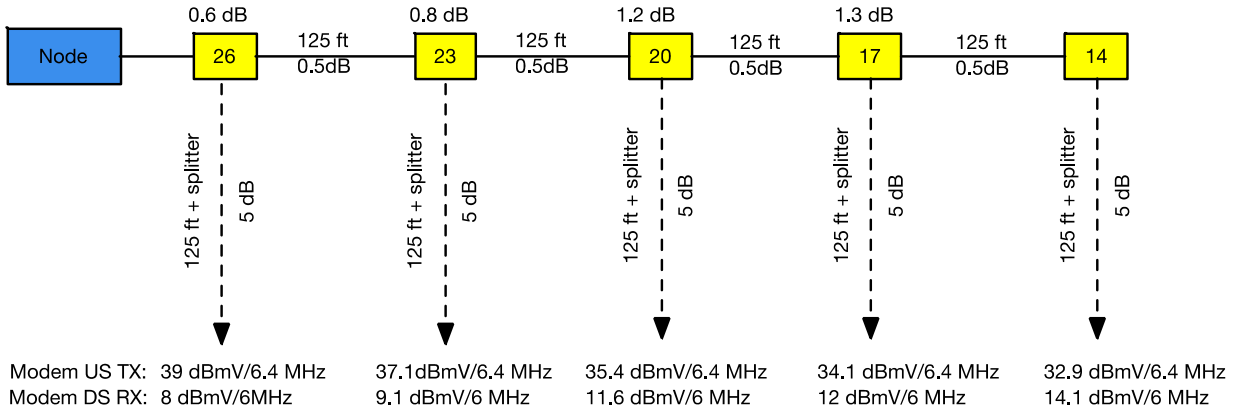


Figure 2 - A Node +0 Passive HFC Plant

Within a segment, the US signal from one home may leak through the distribution line and cause co-channel interference to other homes, as shown in Figure 3. If there is sufficient isolation between the neighboring homes, such interference might be negligible relative to the modulation order used by the DS. The signal-to-interference ratio depends on the received DS signal power level, the transmitted US signal power level, and the path loss from the interferer CM to the receiver CM. The path loss is a function of the frequency dependent attenuation and reflections at various network elements along the distribution line including taps, splitters, trunk and drop cables.




Figure 3 - US to DS Signal Isolation Within a Segment

#### 4. Interference Group

The concept of IG and the later concept of transmission group (TG) as explained in Section 7 were originally proposed by the authors of this paper in [1].

A CM's IG is defined as the set of CMs whose DS reception may be interfered by the CM's US transmission. A typical IG is the group of CMs in adjacent homes that are connected to the same tap; it may also contain CMs under adjacent taps in case of poor tap-to-tap isolations. Quantitatively, given a set of CMs,  $cm_1, cm_2, \dots, cm_N$  of the distribution network,  $cm_i$ 's IG group,  $IG_i$ , is a function of a signal-to-interference ratio threshold,  $SIR$ , such that,

$$\text{for } cm_j \in IG_i, SIR_{ij} < SIR; \tag{1}$$

or,

$$\text{for } cm_j \notin IG_i, SIR_{ij} \geq SIR; \tag{2}$$

Where,  $SIR_{ij}$  is the DS signal-to-interference ratio measured by  $cm_j$  when  $cm_i$  transmits in the US direction over the same spectrum. The SIR threshold used for forming the IG also affects the DS modulation order on which the group can operate. The SIR value can be used to determine if a CM can receive the DS signal. Hence, we can use the SIR value to sort CMs into interference groups.

As an example, the following tables show the average SIR measured in the 200-400 MHz range and the corresponding modulation order, using the plant shown in Figure 2. The shaded cells have the SIR or the modulation order lower than the target threshold, 35 dB for SIR or 10 bits/subcarrier for the modulation order.

From the data, we can observe the following:

1. Low SIRs for CMs under the same tap; for example, the SIR is 8 dB for CMs under Tap1, as the RF path loss under the same tap is less when compared to the inter-tap case.
2. Low SIRs for CMs under the taps close to the end of distribution line; for example, CMs under Tap 3 through Tap5 all have SIR below 35 dB, due to the poor coupling loss of the lower-value taps.
3. Symmetrical CM-to-CM interference indicating reciprocal path loss of the passive plant.

**Table 1 - CM-to-CM Interference and IG Formation**

RECEIVE	SIR (dB)	TRANSMIT				
		Tap1	Tap2	Tap3	Tap4	Tap5
RECEIVE	Tap1	6	39.9	39.9	39.9	39.9
RECEIVE	Tap2	39.9	9.8	37.5	7.5	37.5
RECEIVE	Tap3	39.9	37.5	13.2	34.5	34.5
RECEIVE	Tap4	39.9	37.5	34.5	15.8	31
RECEIVE	Tap5	39.9	37.5	34.5	31	18.2

RECEIVE	Mod Order	TRANSMIT				
		Tap1	Tap2	Tap3	Tap4	Tap5
RECEIVE	Tap1	-	11	11	11	11
RECEIVE	Tap2	11	-	10	10	10
RECEIVE	Tap3	11	10	-	9	9
RECEIVE	Tap4	11	10	9	-	8
RECEIVE	Tap5	11	10	9	8	-

In this example, 3 IGs are identified using a 35 dB SIR threshold. This is also as shown in Figure 4.

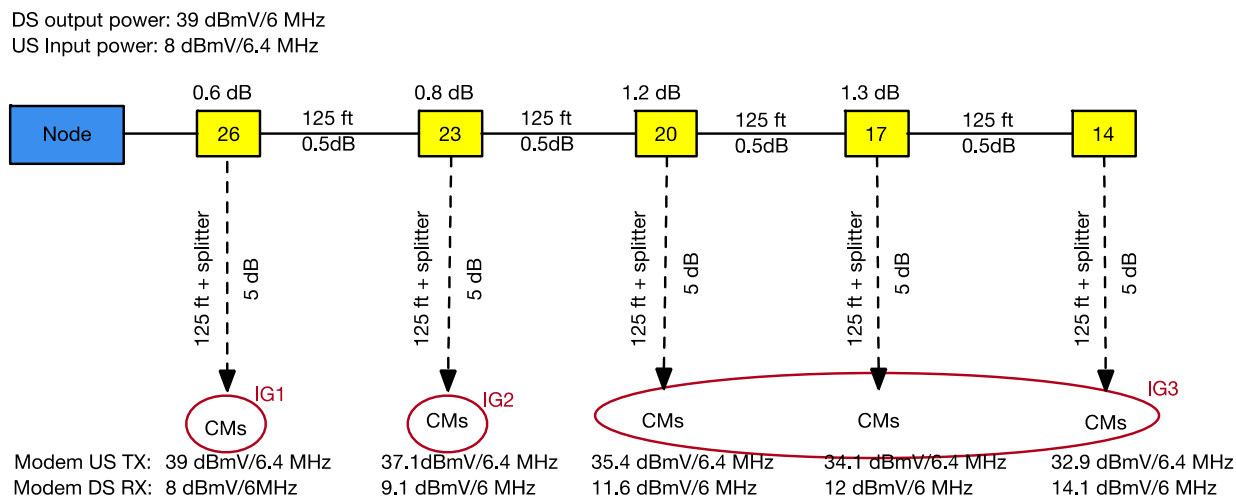


Figure 4 - CM IG Formation

## 5. IG Discovery

In general, the CMTS is not aware of the physical topology of the HFC network, in terms of which tap a CM is connected to and how long the cable is. The only way to determine the interference level is through testing. In [1], a measurement-based IG discovery procedure has been proposed to dynamically measure the CM-to-CM interference. Specifically, to conduct the interference test, the CMTS would instruct one CM to send a test signal in the upstream while asking other CMs to measure on the DS receive path.

The measured interference level and the received DS signal power measured at the CM, when combined, allow the CMTS to derive the signal-to-interference ratio in order to identify IGs. If the signals are sent at different frequencies, the measured data further allow the CMTS to quantify the interference variations throughout the FDX frequency band.

IG discovery can be achieved by using either existing DOCSIS3.1 profile / PNM functionalities or new techniques to be specified by CableLabs.


## 6. Duplex-Simplex Scheme

A CM's IG separates the neighboring CMs into two categories:

1. CMs within the IG that need to operate in simplex mode (uni-directional at any particular point of time and frequency) with the specific CM to avoid co-channel interference;
2. CMs outside the IG that have sufficient RF isolation therefore can operate in full duplex mode with the specific CM.

In this nomenclature, simplex refers to a unidirectional path at a given instance of time and frequency. Thus FDD is referred to as simplex by this definition even though FDD is really two separate unidirectional paths that exist at the same time. Duplex would refer to a bidirectional path at a given instance of time and frequency.

Although the CMs in an IG operate in simplex mode, the overall plant as shown in Figure 4 operates in duplex mode over the collection of IGs. This is illustrated in a simple way in Figure 5.



**Figure 5 - Duplex-Simplex Scheme**

Within an IG, the CMs cannot transmit and receive at the same time, on the same frequency. They may transmit and receive on different frequencies (FDD) or at different times (TDD). Meanwhile, while one IG is only receiving, another IG can only transmit. This is what is implied with the simplex view per IG. When scheduled properly the entire downstream spectrum, at each point in time, is used for CM reception and the entire upstream spectrum can be used for CM transmission, to achieve full duplex, also shown as the CMTS view in Figure 5.

## 7. Transmission Group

While IG represents CM interference relationship at the PHY layer, TGs are used for FDX resource allocation at the MAC layer. A minimum of 2 TGs are required for an FDD spectrum based design to be used for both DS and US transmissions. Specifically, each TG is assigned with a set of non-overlapping DS and US spectrum resource blocks; between the TGs, the DS and US resource blocks are paired up so the same spectrum can be used for DS and US transmissions.

Since the resource allocated to a TG may be shared by multiple IGs, the TG-based resource sharing allows the FDX scheduler to scale with a large number of IGs while keeping a simple resource allocation scheme at the root level.

## 8. Resource Sharing Model

Figure 6 illustrates the FDX resource-sharing model based on the “duplex-simplex” FDX operation principle. Vertically, the spectrum resource is shared in non-FDX mode, just as in DOCSIS today. Within each TG, the non-overlapping DS and US resource blocks are shared among a number of IGs, with each IG representing a group CMs. Horizontally, the spectrum resource is shared between the TGs for both DS and US transmissions to achieve FDX. The IG-to-TG mapping in Figure 6 uses the IG discovery result shown in Figure 4. Specifically, IG1 and IG2 are assigned to TG1, and IG3 is assigned to TG2. Since an IG is included in a TG in its entirety, there are no co-channel interference concerns between the TGs.



The spectrum resource can be allocated in the unit of a DS / US channel as in FDD (FDX-FDD), or in the unit of a DS / US frame as in TDD (FDX-TDD), or can be a combination of both. Today's DOCSIS technology is based upon FDD. Supporting TDD would require significant hardware and protocol changes at both CMTS and CMs. For the rest of the discussion, we will focus on FDX-FDD (read as "FDX with FDD"), as this is likely to be the first FDX resource-sharing scheme to deploy.

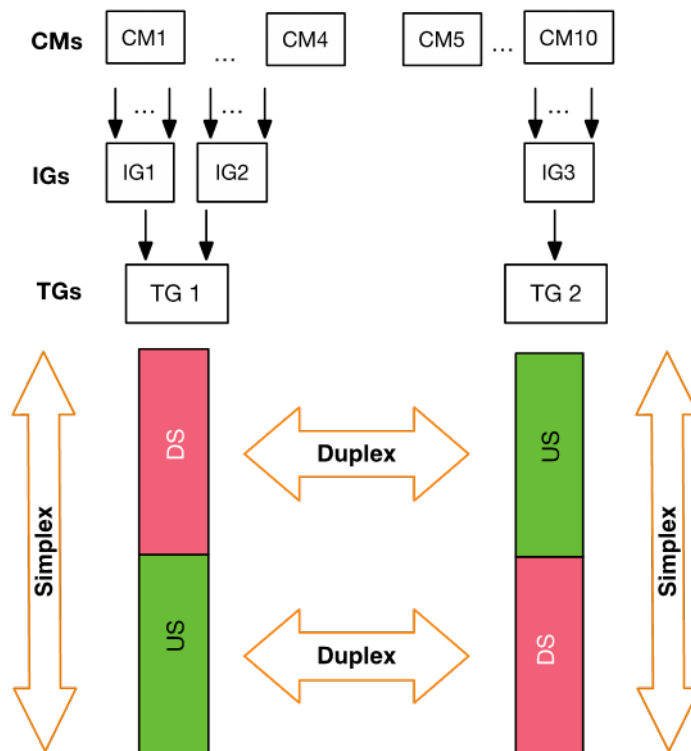


Figure 6 - FDX Resource Sharing Example

## 1. FDX Throughput Multiplier

The ultimate goal for FDX DOCSIS is to improve DS and US data throughput using DOCSIS3.1 spectrum. To evaluate the throughput gain we introduce a quantifier, defined as the FDX throughput multiplier, as below:

$$FDX \text{ Throughput Multiplier} = \frac{DS \text{ and } US \text{ throughput with FDX}}{DS \text{ and } US \text{ throughput without FDX}}$$

Using Figure 6 as an example, we assume the average DS subcarrier bit-loading is set as below, over a 400 MHz FDX frequency band. We also assume there are two CMs attached to each tap along the distribution line.

**Table 2 - Average Bit-loading in Bits per Subcarrier**

Transmission Direction	FDX			Non-FDX
	TG1		TG2	
	IG1	IG2	IG3	
DS	11	10	10	11.5
US	10	10	10	10

With FDX, each of the two TGs will be assigned with one 192 MHz OFDM channel for the DS, and two 96 MHz OFDMA channels for the US. Without FDX, the entire 400 MHz band will operate in simplex mode with non-overlapping DS and US channels as in existing DOCSIS. A slightly lower DS bit-loading is used for the FDX case to consider the potential loss due to the US-to-DS interference.

Table 3 shows the bandwidth capacity in terms of maximum DS and US throughput and per CM average DS and US rate of each TG involved for FDX resource sharing.

**Table 3 - Per TG Throughput with FDX Resource Sharing**

TG	Spectrum Width in MHz (DS xUS)	Average Bit-Loading (DS xUS)	FEC Efficiency	Maximum Throughput in Mbps (DS xUS)	Number of CMs	Per CM Average in Mbps (DSxUS)
TG1	192 x192	10.5 x10	0.8	1613 x 1536	4	403 x 384
TG2	192 x192	10 x10	0.8	1536 x 1536	6	256 x 256

Table 4 shows the FDX throughput gain with the two TGs combined. With FDX resource sharing, the entire spectrum is used twice, one for each DS and US directions. With the assumed DS and US bit loading values, the FDX throughput multiplier is 1.88 in this example.

**Table 4 - FDX Throughput Multiplier**

Resource Sharing Mode	Maximum Throughput in Mbps (DS xUS)	Per CM Average in Mbps (DSxUS)	FDX Throughput Multiplier
Non FDX	1766 x 1536	177 x 154	1
FDX	3149 x 3072	TG1: 403 x 384 TG2: 256 x 256	1.88

## 2. Factors Affecting FDX Throughput

The FDX throughput can be expressed as the sum of all the TGs sharing the spectrum resource, shown as below:

$$FDX\ throughput = \sum_{i=1}^N D_i * S_{D_i} * X_{D_i} + U_i * S_{U_i} * X_{U_i}$$

Where  $N$  is the total number of TGs sharing the spectrum. For each TG,  $TG_i$ ,  $D_i$  and  $U_i$  are the DS and US resource block size;  $S_{Di}$  and  $S_{Ui}$  are the DS and US spectrum efficiency in terms bits/s/Hz;  $X_{Di}$  and  $X_{Ui}$  are the DS and US bandwidth utilization percentage within the TG.

This equation reveals the following factors affecting FDX throughput:

- Resource block size
- Resource block utilization
- Spectrum bit-loading

At given spectrum resource limitations and spectrum bit-loading, improving a TG's spectrum resource utilization is the key to maximize the overall FDX throughput.

### 3. Dynamic Load Balancing

The following two approaches could be used to improve a TG's spectrum resource utilization:

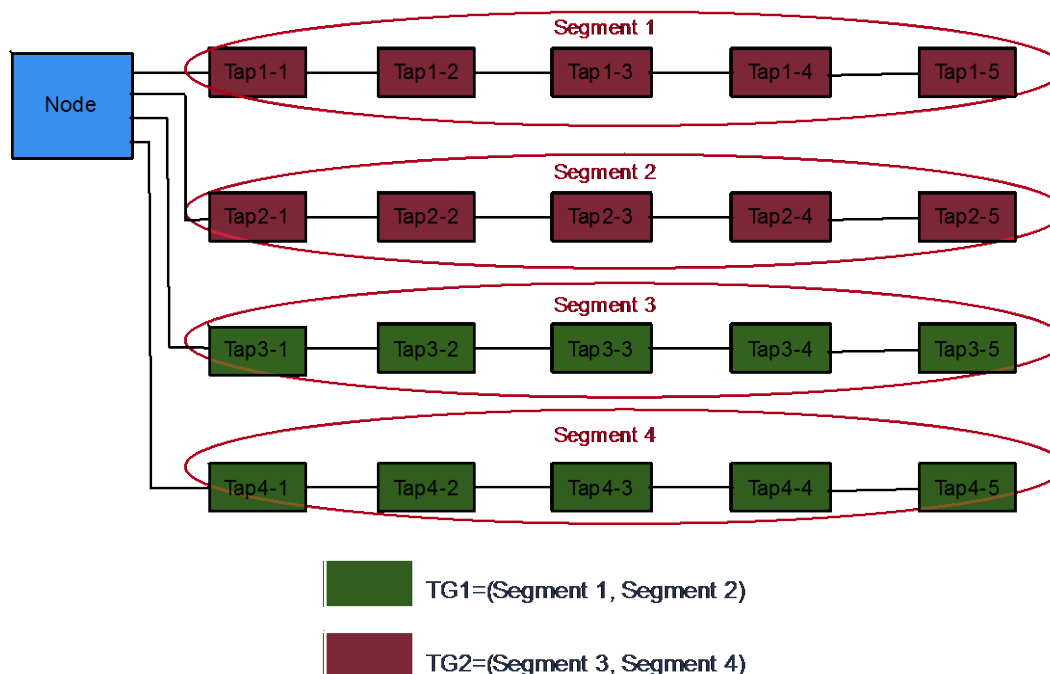
- Adjusting a TG's resource allocation to match traffic load
- Balancing the traffic load to match the TGs' resource allocation

For FDX-FDD, the bandwidth partition likely needs to be done at the initial provisioning phase as changing the channel width impacts the OFDM physical layer operation and causes service interruption. Adding or removing an OFDM/OFDMA channel from a TG, on the other hand, doesn't necessarily provide the granularity for the traffic need. Due to these difficulties, it is more economical and efficient to match the TGs' resource allocation via load balancing.

#### 3.1. Static FDX Load-Sharing

One approach to construct the TGs is to use the physical port or plant topology information. Figure 7 shows an example that has two TGs constructed over 4 segments, with each TG containing two adjacent segments. In this scheme, a segment is equivalent to an IG in the sense that the segments are isolated from each other in RF domain. Duplex operation can therefore be enabled between the TGs.

This allocation scheme meets the duplex-simplex FDX requirement but at the cost of high rigidity in load-sharing. Since TGs are fixed to the network physical topology, traffic load cannot be balanced between the TGs. For example, when TG1 is highly congested, TG2 can be well under-utilized due to lack of traffic demand. The FDX spectrum resource therefore cannot be fully utilized.




**Figure 7 - Segment-Based Static TG Mapping**

### 3.2. IG Pool and Dynamic IG-to-TG Mapping

In order to gain more dynamic behavior, it is desirable to have the TGs dynamically sharing the traffic load. This requires dynamic IG-to-TG mapping, as an IG has to be included in a TG in its entirety to avoid co-channel interference between the TGs.

For this we introduce the concept of an “IG Pool,” as shown in Figure 8, that includes all IGs in the node’s service group. The IG pool represents the common traffic load that needs to be shared between the TGs. Instead of a fixed IG-to-TG mapping, IGs are assigned to TGs dynamically during operation to maximize the combined traffic throughput. For example, consider the case that one TG has its US resource block fully congested while the other TG has its DS resource block fully congested. Balanced load could be achieved by trading the IGs between the TGs such that both TGs’ DS and US resource budget can be fully utilized with matching traffic.


For FDX-FDD, moving an IG from one TG to another implies changing the CM receive and transmit channel sets. This can be done gracefully with DOCSIS DBC (dynamic bonding change) mechanism without forcing the CM offline.



**Figure 8 - IG Pool and Load-Based IG-to-TG Mapping**

### 3.3. Dyanmic Load-Sharing Performance

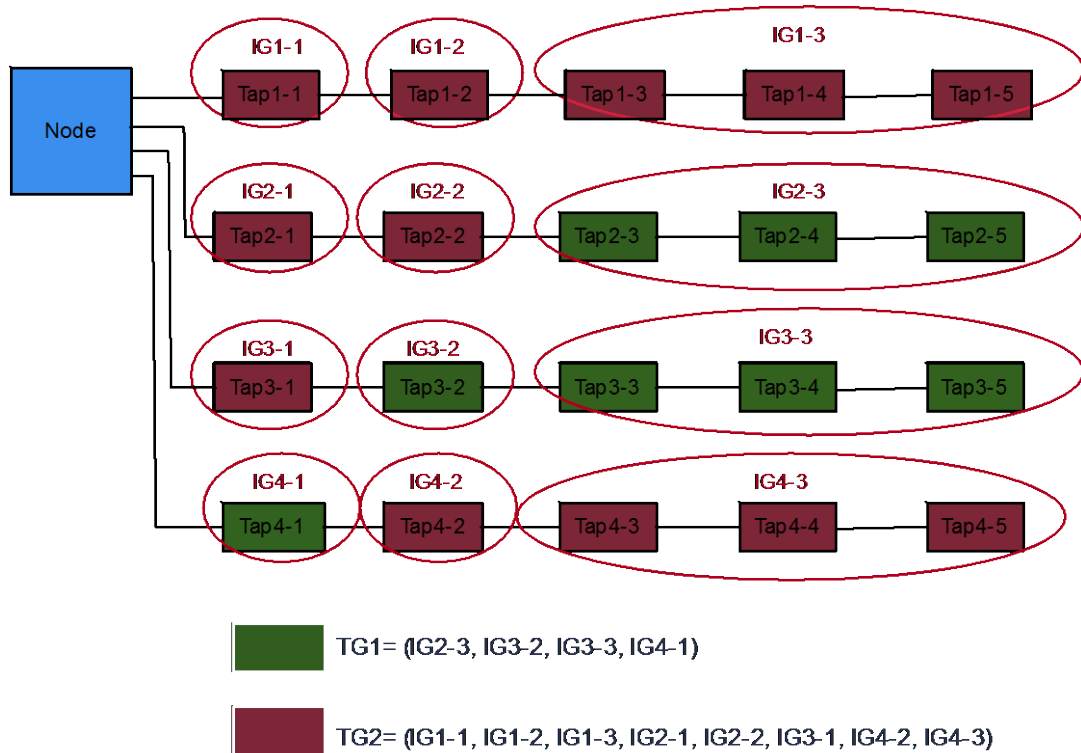
To study the dynamic load balancing performance, we simulated a N+0 HFC network with four segments using the same tap topology as shown in Figure 7. By assuming two CMs under each tap and numbering the CMs sequentially along each distribution line, e.g., CM0 and CM1 attached to Tap1-1, and CM38 and CM39 attached to Tap4-5, we created an arbitrary imbalanced traffic pattern as shown in Figure 9. As shown, most of the traffic load is contributed by the latter twenty CMs with various DS-to-US traffic ratio, while the first twenty CMs have a much lighter traffic load.



**Figure 9 - Imbalanced Subscriber Traffic Demand**

Based on the load distribution, the dynamic load-balancing scheme produced an IG-to-TG mapping as shown below.





**Figure 10 - Traffic Load-Based Dynamic IG-to-TG Mapping**

By applying the static load-sharing scheme to the same traffic pattern, we are able to compare the FDX throughput gain between the two approaches. The result is shown in Table 5. For this traffic pattern, since the total traffic demand is less than the maximum FDX spectrum capacity, we use the demand-to-throughput percentage as a quantifier to measure the FDX performance. As shown, at the given traffic load, TG2 is fully congested for the static scheme while TG1 is well underutilized. Because of this, even the total traffic demand is less than the FDX spectrum capacity; the static scheme can only satisfy 69% of the DS traffic demand and 74% of the US traffic demand. The dynamic load-sharing scheme can fully stratify the traffic demand by distributing the load between TG1 and TG2.

**Table 5 - Static vs. Dynamic Load Sharing Traffic Demand and Throughput**

Load Sharing Scheme	Total FDX Spectrum Width in MHz (DS x US)	Total FDX Bandwidth Capacity in Mbps (DS x US)	Total Traffic Demand in Mbps (DS x US)	Per TG Throughput in Mbps (DS x US)		Demand to Throughput Percentage (DS x US)
				TG1	TG2	
Static	192 x 192	3512 x 3072	2919 x 2351	283 x 214	1736 x 1536	69% x 74%
Dynamic	192 x 192	3149 x 3072	2919 x 2351	1600 x 980	1319 x 1317	100% x 100%

The above results assumed a 400 MHz FDX band with following average subcarrier bit-loading settings:

**Table 6 - Average Bit-loading in Bits per Subcarrier**

Transmission Direction	Static Load Sharing	Dynamic Load Sharing (n = 1,2,3,4)		
		IG n-1	IG n-2	IG n-3
DS	11.5	11	10	10
US	10	10	10	10

In the following study, we further evaluate the two resource-sharing schemes under different traffic patterns. By sequentially activating the CMs along the distribution line for each segment, we are able to create various degrees of unbalanced traffic distributions. Among the activated CMs, traffic load is evenly generated per CM based on the maximum DS and US spectrum capacity and the number of active CMs as below:

$$\text{Per CM DS rate} = \text{Max DS Bandwidth Capacity} / \text{number of active CMs}$$

$$\text{Per CM US rate} = \text{Max US Bandwidth Capacity} / \text{number of active CMs}$$

For the simulation, the maximum DS and US bandwidth capacities assumed for both the static and dynamic load-sharing are listed in Table 5.

Figure 11 shows the resultant FDX throughput multipliers for both the static and dynamic load-sharing schemes. Before the traffic is added to the 3<sup>rd</sup> segment, there is no FDX gain for the static approach, as there is only TG1 active, and the resource allocated to TG2 is not utilized at all, simply because there is no traffic on the latter two segments.

For the dynamic load-sharing approach, the measurement based IG discovery allows multiple IGs to be formed within a single segment. For this example, IG1-1 contains the first two active CMs, IG1-2 contains the third and fourth CMs. Therefore, as soon as the traffic starts on the third CM, both IG1-1 and IG1-2 become active. The dynamic load-sharing scheme then assigns IG1-1 to TG1 and IG1-2 to TG2, and achieves the FDX throughput gain. As more CMs are activated, the dynamic load-sharing scheme continues to adjust the IG-to-TG mapping and to maximize the overall FDX throughput.

The static throughput line crosses the dynamic throughput line when traffic load among the distribution lines become fully symmetrical. The slightly better throughput for this case is due to the higher average bit-loading assumed. Overall, this simulation shows that the dynamic load-sharing scheme can effectively distribute the traffic load to the TGs and achieve high FDX throughput gain under various traffic conditions.




Figure 11 - FDX Multipliers for Static vs. Dynamic Load Sharing

## 4. Conclusion

FDX DOCSIS can up to double the usage of the spectrum in cable network by allowing simultaneous downstream and upstream transmissions over the same spectrum. This paper explained how an intelligent resource-scheduling scheme could be used to maximize FDX throughput gain while simultaneously mitigating the US to DS interference sourced from multiple CMs across the cable access network. First, we formulated the CM-to-CM interference with IGs. Second, we developed a hybrid FDX resource-sharing model by sharing non-overlapping spectrum resource within a TG, and by matching the DS and US resource allocations between the TGs. Finally, we presented an FDX throughput optimization scheme using dynamic TG load-sharing, and demonstrated the its effectiveness with simulations.

## 5. Abbreviations and Definitions

### 5.1. Abbreviations

ADC	analog to digital convertor
CM	cable modem
DBC	dynamic bonding change
DS	downstream
DOCSIS	Data-Over-Cable Service Interface Specifications
EC	echo canceller
FDD	frequency division duplex
FDX	full duplex
IG	interference group
RPD	remote PHY device
TDD	time division duplex
TG	transmission group

US	upstream
----	----------

## 5.2. Definitions

Interference group (IG)	A CM's interference group includes the set of CMs whose DS reception may be interfered by the CM's US transmission. An IG is a PHY concept that is based upon path attenuation measurements.
Transmission group (TG)	A number of interference groups that share a common set of non-overlapping DS and US resource blocks. A TG is a MAC concept that is used for scheduling of traffic.

## 6. Bibliography and References

[1] John T.Chapman, Hang Jin (2016). *Full Duplex DOCSIS*, INTX 2016, May 18, 2016

[2] DOCSIS 3.1 MAC and Upper Layer Protocol Interface Specification, CM-SPMULPIv3.1-I08-151210, December 10, 2015

# **Cable's Success is in its DNA:**

## **Designing Next Generation Fiber Deep Networks with Distributed Node Architecture**

A Technical Paper prepared for SCTE/ISBE by

Venk Mutalik, Engineering Fellow, ARRIS  
15 Sterling Drive  
Wallingford, CT 06492  
Venk.Mutalik@arris.com  
+1 203 494 6556

Zoran Maricevic, Engineering Fellow, ARRIS  
15 Sterling Drive  
Wallingford, CT 06492  
Zoran.Maricevic@arris.com  
+1 203 494 0877



## Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents	110
1. Introduction	111
2. The Fiber Deep and All Fiber Migration	111
3. Peak and Average Traffic Rates and Segmentation Options	113
4. Distributed Node Architecture – A Spectrum of Options	114
4.1. The HFC Reference Architecture	115
4.2. The N+0 Architecture	116
4.2.1. Fiber Deep N+0 Networks	116
4.2.2. N+0: Efficient and Effective	118
4.2.3. Remote PHY in the N+0 Architecture	119
4.2.4. N+0 and Preparing for FTTH	120
4.3. The FTLA Architecture	121
4.3.1. Aging RF Plant Discussion	121
4.3.2. Remote PHY in the FTLA Architecture	123
4.3.3. Preparing FTLA for the FTTH Migration	124
5. Conclusions	124
6. Acknowledgements	125
7. Abbreviations and Definitions	125
7.1. Abbreviations	125
7.2. Definitions	126
8. Bibliography and References	127

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Optical Beat Interference in Traditional RFoG	112
Figure 2 - Typical N+3 Reference Architecture	115
Figure 3 - Typical Fiber Deep N+0 Architecture	117
Figure 4 - Fiber Deep N+0 to R-PHY Migration	119
Figure 5 - R-PHY and PON in N+0 Architecture	120
Figure 6 - Amplifier Failures per Year	121
Figure 7 - Basic FTLA Architecture	122
Figure 8 - Fiber Deep N+0 to R-PHY Migration	123
Figure 9 - N+0 to R-PHY, FTTH Selective Subscriber Migration	124

## List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Basic Summary of the MSO Case Study	113

## 1. Introduction

Relentless traffic growth in the upstream (US) and downstream (DS) is a prime driver for fiber deep (FD) networks. A move to an all fiber network is a journey spanning several decades [1]. Anticipating this, multiple system operators (MSOs) have researched progressively complex schemes to increase capacity and drive fiber deeper. Thus node segmentation, active drop-in (ADI), fiber deep node plus zero (N+0), fiber to the last active (FTLA) N+0, extended spectrum DOCSIS<sup>®1</sup>, and full duplex DOCSIS (FDX) are all techniques that can help while fiber is being driven deeper into the network, ultimately leading to fiber to the home (FTTH).

Each MSO will make changes to their own hybrid fiber-coax (HFC) plant to optimize for the attributes that they deem to be the most important – premised on where their networks start from and where they'd like to take them. Different MSOs will likely prioritize the many attributes in different ways. In situations where the network and deployed technology maybe relatively old, MSOs may choose to optimize their network evolution by moving as rapidly as possible to end-state technologies of the future. These MSOs maybe be more likely to move rapidly towards a passive optical network (PON) or Point-to-Point Ethernet solutions. Other MSOs may be more constrained by their capital and choose to take a more paced approach by optimizing their network evolution to reduce headend power and rack-space requirements by moving towards fiber deep architectures with distributed access architecture (DAA) sub-systems that remove functionality from the headend. These MSOs will likely deploy remote PHY (R-PHY) or remote MACPHY (RMACPHY) sub-systems within their nodes. Other MSOs will want to preserve much of their current architectures while capitalizing on improved technologies.

While all of these approaches are valid and justifiable approaches moving forward into the future, we will look at each of the approaches indicated above and introduce a new approach to fiber deep networks called the Distributed Node Architecture. Utilizing well known multi-wavelength optics, digital daisy chaining, and recent technology breakthroughs in optical beat interference (OBI)-free optics, we show how the node functionality can be distributed across the service groups, with limited impact on the existing infrastructure, and help match capacity to customer demand.

## 2. The Fiber Deep and All Fiber Migration

Companies offering 1,000 megabits per second (Mbps) over fiber have emerged as new competitors to established MSOs [2]. Essentially operating in a greenfield environment and with relationships forged with the local governments, one company has been on an aggressive fiber build program in the United States. For the MSOs, however, the issue is not so much the greenfield builds, but more of how to upgrade the vast majority of brownfield HFC plants to all fiber. We have discussed before that a migration to an all fiber network for most MSOs is a journey and not an event. The industry has typically proceeded around the rate of 37,500 plant miles per year for the past several decades. Even a significant acceleration of this build rate would entail a transition time of several decades for the nation to convert to all fiber networks.

Many equate an all fiber network necessarily with one of the traditional passive optical networks (PON). However, a traditional PON network such as 10G Ethernet passive optical network (EPON) or XG-PON would not only require different equipment and transmission formats, but also require a system such as DOCSIS provisioning of EPON (DPOE) to conserve the existing infrastructure at the home and at the headend. Furthermore, it will also require an effective Internet protocol television (IPTV) system to be in

---

<sup>1</sup> DOCSIS is a trademark of CableLabs



Challenges of fiber exhaustion and long links are well known and are likely to get more visible while driving fiber deeper. It may seem a bit counter intuitive that driving fiber deeper should exacerbate fiber exhaustion, after all how could adding fiber to the network militate against itself? It can because the fiber exhaustion effects are felt primarily in trunk fibers, these fibers laid down or leased many years ago have served limited numbers of fiber nodes. Now with the explosion of fiber nodes across the landscape, trunk fibers are at a premium. MSOs today have the capability to cumulatively deploy up to 44 wavelengths (WLs) on one strand of fiber. WL plans optimize stimulated Raman scattering (SRS), cross phase modulation (XPM), and four wave[length] mixing (FWM) and other non-linear effects. Furthermore, the use of C-Band directly modulated lasers (DMLs) for moderately long links with electronic dispersion compensation and optical passives that can accommodate these lasers significantly enhances cost effectiveness of fiber deep optics deployments.

### 3. Peak and Average Traffic Rates and Segmentation Options

Presented in Table 1 is a summary of a case study [6] that looked at the composition of transmission demand among major MSOs over several years preceding 2014. The table lists the peak transmission rate (Tp) of the four service tiers and the average transmission rates (Ta) including their compound annual growth rate (CAGR). It is interesting to see that the top billboard rate is subscribed to by less than 1% of the subscribers. While the billboard rate grows at an insatiable 50% year over year, performance and basic tiers form a big chunk of subscribers, have modest peak bandwidth requirements, and have more moderate peak bandwidth growth rates (15% to 32%). The measured average rate for the subscriber service group was around 0.5 Mbps/subscriber during peak hour and grew at 45%.

**Table 1 - Basic Summary of the MSO Case Study**

2014 MSO Case Study Downstream Service Tiers			
Service Tiers	% Subs	Tp CAGR	Tp Mbps 2014
Top Billboard Tier Tp	<1%	50%	300
Performance Tier Tp	14%	32%	75
Basic Tier Tp	65%	26%	25
Economy Tier Tp	20%	15%	5
		Ta CAGR	Ta Mbps 2014
Average rate Ta		45%	0.5

It is important to note in this context that average transmission rates per subscriber can be enhanced either via increasing total available capacity or by decreasing the number of subscribers per group. However, only increasing the total available capacity can enhance peak transmission rates.

Therefore, in the past, when service groups and node sizes were large, billboard rates were modest, and only a sliver of spectrum was allocated to high-speed data (HSD), it made sense to segment nodes and de-cluster the headend so that each service group had its transmitter and its own cable modem termination system (CMTS) port. This allowed each service group to receive its own sliver of HSD spectrum and boosted the average data rate available to each subscriber. As service groups were segmented, the headend space requirements increased and this led to denser optics and CMTS resulting in the converged cable access platform (CCAP) standards in vogue today.



MSOs today are in a race to offer increasingly higher billboard rates, among other things to compete with fiber deployments. However, higher peak transmission rates are possible only by having higher total capacity per the service group. This has been the basis for analog harvesting, spectrum enhancement, driving fiber deeper, and the migration to D3.1.

A simplified human factors formula that addresses quality of experience (QoE) for evolving traffic needs [5] is shown below:

- $T = \alpha \cdot T_p + \beta \cdot T_a \cdot N_s$ 
  - $T$ : Total Transmission Capacity in Mbps or gigabits per second (Gbps) needed for a given service group
  - $T_p$ : Billboard rate or Peak transmission advertised rate in the service group
  - $T_a$ : Measured Transmission average rate per subscriber used at the busy hour  
(Note that  $T_a$  is a measured quantity during busy hour and is different from the Average Data per Sub, which is  $T/N_s$ )
  - $N_s$ : Number of subscribers in the service group
  - $\alpha$  and  $\beta$  are variables that depend upon  $N_s$  and slowly diverge from 1 as  $N_s$  increases

An expression of the type above, restates the general dictum discussed earlier, that when peak rates dominate further reduction of node sizes by reducing  $N_s$  is unlikely to lead to significant reduction in the total capacity requirement per service group. Therefore, should the  $T_p$  rates continue growing at 50% CAGR for the next decade and beyond,  $T$  will be dominated by  $T_p$ . If this happens,  $T$  is relatively unaffected by the subscriber count, therefore there might be benefits to aggregating service groups and sharing the  $T$  amongst a larger subscriber base.

However, should the  $T_p$  rates stabilize and grow much more modestly for the bulk of the subscriber base, then the  $T_a$  growth rate could outstrip the  $T_p$  growth rates. This is especially true when IPTV comes online and the viewing moves from streaming Moving Picture Experts Group (MPEG) to HSD, sustained  $T_a$  rates could become a very critical parameter for the service groups. In this case, node segmentation would indeed be very important.

Realizing this, one strategy that the MSOs could follow is to size the node group small enough that should the average rates dominate, each node could serve as its own service group. But, for the duration when peak rates dominate, there is adequate opportunity to aggregate several node groups into one service group. Simultaneously, the MSOs would continue to increase total capacity per service group by appropriate use of spectrum enhancement, D3.1, and other advanced techniques. This yields an optimal solution for the long term and enables seamless migrations towards matching capacity needs of each service group as the demand increase either for peak or for average data rates.

#### 4. Distributed Node Architecture – A Spectrum of Options

Distributed node architecture is an appropriate positioning of downstream amplification and of upstream aggregation in the network to provide higher capacity per service group, support increasing nodes in the field, and facilitate driving fiber deeper in the network.

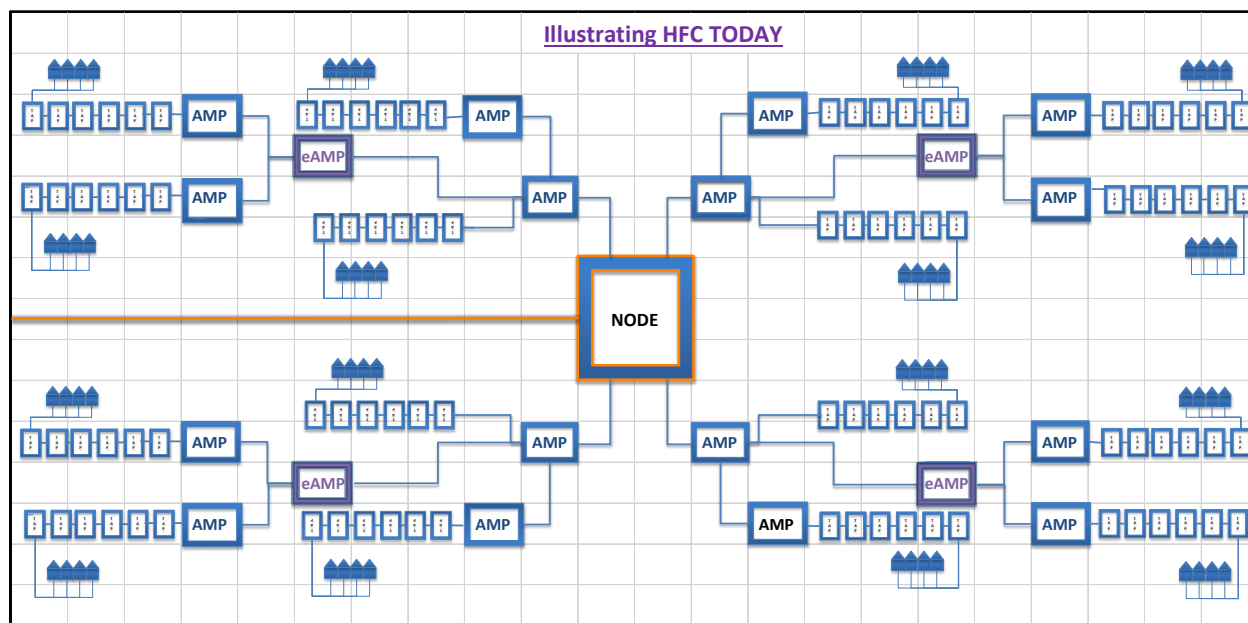
Since peak rates and average rates can increase unevenly over the next decades, we have noted previously that nodes should be adequately sized to provide needed average rates per subscriber, while the service group sizes should be adequately big to provide to support peak rates economically.



Over the years, MSOs have made decisions on the use of either analog reverse path or baseband digital reverse path technologies in their network. Therefore, the process of node segmentation and service group aggregation must support either technology.

Keeping this mind, we describe several options that accommodate business as usual staples such as node segmentation and newer fiber deep architectures such as N+0 and FTLA that also poise the cable companies to move into FTTH with relative ease.

#### 4.1. The HFC Reference Architecture



**Figure 2 - Typical N+3 Reference Architecture**

Presented in Figure 2 is an idealized N+3 node service area. In this figure, we have one fiber node followed by 16 RF amplifiers designated as AMP. Some actives labeled eAMP are express RF amplifiers and serve no homes directly. For simplification, in this example we have shown no homes served by the optical fiber node, although the nodes in the field might serve homes. For the remaining part of the paper, this node location is referred to as the parent node location. One can therefore see that there are 21 actives, serving 384 households passed (HHP). Each AMP serves six taps and each tap serves four homes. Although nodes sizes vary from MSO to MSO and across the same cable system in response to density and the state of where they find themselves in the upgrade cycle, we use the aforementioned to illustrate various options as we migrate towards a fiber deeper and ultimately an all fiber network.

In a typical node service area, such as that in Figure 2, there are about 30,000 plant feet, not counting the drops. Of these, approximately 33% (i.e., 10,000 plant feet) cover the node to the last RF amplifier. The remaining 67% (i.e., 20,000 plant feet) cover the distance from the last amplifier to the last tap with approximately 200 feet separating each of the six taps on a line. Finally, the drop distance is approximately 150 feet per home, bringing the total drop length in this example to approximately 60K

feet. In the U.S. alone, there are 250,000 to 300,000 of such nodes and 4 million to 6 million RF amplifiers accommodating upwards of 100 million HHP.

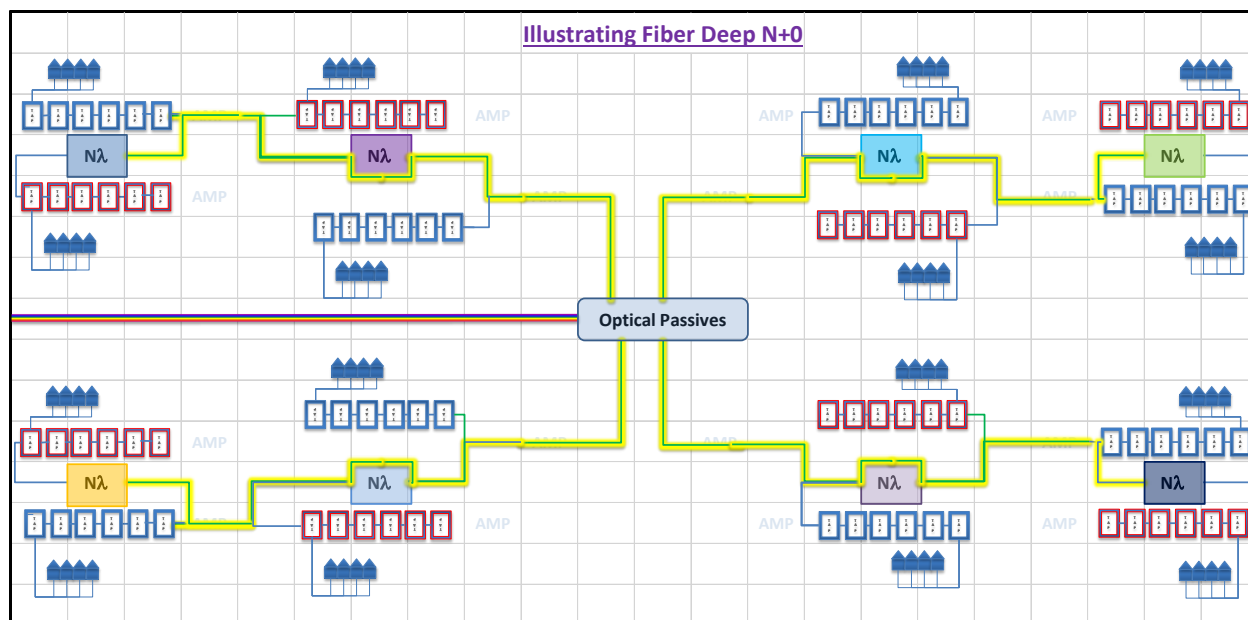
About a decade ago, node service areas were considerably larger and served many more HHP. Pervasive presence of the analog channel line-up limited available RF spectrum that could be dedicated to HSD and other interactive services. Over time however, analog channels were harvested and RF amplifiers in the field replaced by 860 MHz or 1 GHz gear via active drop in (ADI). This helped liberate more of the RF spectrum for streaming MPEG QAM interactive services such as switched digital video (SDV) and videon demand (VoD) and to DOCSIS-based HSD transmission. Requirements of average rates dominated, and with HSD spectrum still at a premium, it made perfect sense to segment the nodes.

Today, most MSOs already operate all-QAM systems or are close to making that transition. Node segmentation is particularly effective in reducing RF amplifier cascade sizes, improving signal to noise ratios, and managing RF ingress from the upstream. Earlier, one transmitter and one CMTS port served many fiber nodes, and that was sufficient to fill up capacity on the CMTS port. Over time, the transmitters and CMTS ports were de-clustered, such that typically one transmitter, one reverse receiver, and one CMTS port are dedicated to a single fiber node.

## 4.2. The N+0 Architecture

### 4.2.1. *Fiber Deep N+0 Networks*

In view of the previous discussion, the industry has created a blue print to migrate fiber deeper, with fewer RF actives and a quicker path to FTTH when required. This architecture is called fiber deep N+0 (FD N+0), in recognition of the fact that there are no RF amplifiers subsequent to the fiber node. Innovations in gallium nitride (GaN) technology offer higher RF output levels right out of the node and in turn allow for longer cable lengths and higher splits to serve more HHP than before. This is illustrated in Figure 3, where it is seen that the 21 actives of the previous design have now been replaced by just eight actives in the form of fiber deep nodes. In typical designs, there is a 70% reduction in the number of actives as compared to traditional HFC plant. These new nodes are generally designed to accommodate around 50-150 HHP depending upon plant density.



**Figure 3 - Typical Fiber Deep N+0 Architecture**

In the plant, each of the new nodes is placed strategically to maximize the number of HHP to the limit indicated and minimize the numbers of new nodes. The point to note here is that this architecture gives an opportunity for the designers to re-balance the designs of an HFC system done decades ago.

Often times (although not always) the new location is on the available strand. Often, when the new nodes are laid out, the orientation of taps must be turned to accommodate RF passing from the opposite side to minimize coaxial build. When this happens, the tap values change in addition to their orientation, a process called 'flipping/turning the tap'. Note however that if the objective is also to enhance spectrum to 1.2 GHz, then this effort to replace taps is required anyway, as most current taps only go up to 1 GHz.

Typically, each new node has its own transmitter, receiver, and CMTS port at the headend, with a 1.2 GHz/85 MHz split. Since there is an order of magnitude increase in the number of nodes, the trunk fibers are heavily utilized and therefore need multi-wavelength (MWL) transport. It is in this context that developments in MWL transport and the lower cost of DMLs enable a cost effective deployment of this architecture.

It is important to notice that the parent node location is now just a passives splice enclosure. The parent node location is thus turned into a virtual hub of sorts and it accommodates the optical multiplexers and de-multiplexers. A reduction in the number of actives in the field improves reliability and reduces maintenance and operating costs.

A move towards N+0 drives fiber deeper in the network and closer to the home, additionally the elimination of all actives provides MSOs a capability towards symmetric Gbps service by FDX and higher capacity by extended spectrum DOCSIS. FDX would use the ability of DOCSIS scheduling algorithms to enable full duplex bi-directional transmission over the cable plant in the mid RF region. N+0 thus lets cable operators emulate the success of telcos who also have driven fiber deeper and innovated on services like G.fast by leaving the last few hundred feet of twisted pair in the ground

undisturbed. Cable can also deliver higher capacity by driving fiber closer to the home, but stopping short of replacing all the millions of feet of drop cable across to the homes and thus extend cable infrastructure.

#### **4.2.2. N+0: Efficient and Effective**

The N+0 idea is a rich source for innovation in evolving future architectures. We discussed earlier that a typical HFC node service area has a hardline cabling of around 30,000 feet of which around 10,000 feet links the node to the amplifiers. In the N+0 environment, a strategic placement of nodes not only partitions the HHPs appropriately, but will also considerably reduce the hardline cabling length, due to the limited number of actives that need to be powered. Although each N+0 node may consume more power than an amplifier, the total power in an N+0 design with 1 GHz operation could be 50% of that consumed by the traditional HFC plant. This reduces the number of power supplies needed, saves substantial power cost over the decades in which the networks operated, and also reduces the carbon footprint.

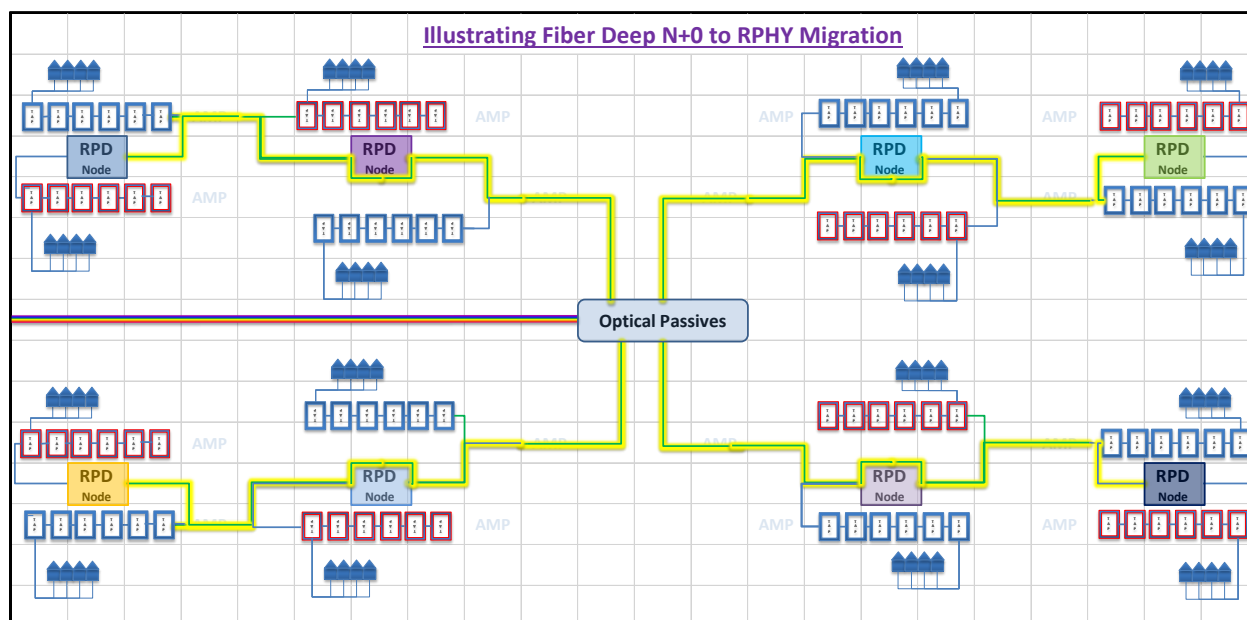
Furthermore, as we discussed in previous sections, when peak rates dominate it is more important to allow for more spectrum and capacity than it is to target ever smaller service groups. Therefore, rather than have one transmitter, receiver, and CMTS port per node, one could cluster up to four nodes (or more) per transmitter if the cumulative HHP are still modest. Note here that the clustering can be done at the headend or in the parent node location in the optical domain or at the headend in the RF domain.

If the clustering is done optically, optical amplification may be placed at the headend or at the parent node location, just before it is distributed in the network of nodes. Doing the former would have the erbium doped fiber amplifier (EDFA) at the headend and increase the launch powers, which can increase fiber non-linear effects but maintain the parent node location as a completely passive splice enclosure. Placing the EDFA at the parent node location would enable lower optical launch levels, reduce optical non-linearities, and possibly enable longer trunk links.

Clustering is especially suited for the digital return technology in the upstream path. When digital return technology is exclusively used, daisy chaining of the return path would enable MSOs to conserve fiber and limit the number of optical wavelengths needed on the trunk fiber, while continuing to provide excellent signal-to-noise ratio (SNR) and one way network monitoring.

A judicious application of upstream and downstream aggregation with optical passives and EDFAs at the parent node location can further reduce cost by enabling the use of ‘grey’ optics rather than DWDM optics at each node location, further improving the cost efficiency of the design. In addition, when the optical clustering of the DS and the daisy chaining of the US are taken together, it serves to throttle down the severe expansion necessary at the headend that could otherwise be needed.

### 4.2.3. Remote PHY in the N+0 Architecture



**Figure 4 - Fiber Deep N+0 to R-PHY Migration**

Remote PHY (R-PHY) has come to mean a set of technologies and devices that enable the CMTS to be logically divided into a central location that does higher layer activity and transmits over conventional (say 10 GbE) Ethernet type binary links, while the node accepts this information and converts that to OFDM/QAM transmissions more suitable for coaxial networks. A node that incorporates a remote PHY device (RPD) is illustrated in Figure 4. R-PHY therefore envisions a reduction in the headend space since the centralized location will not need to take on the power and space requirements for generating the physical layer (PHY). Furthermore, it is anticipated that with the use of cost effective binary optics, distances exceeding 80 km can be robustly covered by multiple wavelengths in the binary realm, while the distribution from the node to homes is covered via the coaxial cable as before.

Presented previously is one way in which the node may be distributed in the N+0 case. Here, each of the fiber deep nodes is replaced by an RPD. The entire optical passive infrastructure remains the same and the optical transmission is now replaced by 10 GbE links. At the RPD location, the node converts 10 GbE to its constituent OFDM/QAM and lets it pass into the coax as before. With a move to R-PHY, the headend to the node location can be vast and thus eliminate intervening hubs.

One thing to keep in mind is that the power to generate the PHY, of OFDM/QAM is transferred from the headend to the node. Therefore, it is now imperative for the node to be of an appropriate size and efficiency to be able to absorb this additional requirement. The benefits of R-PHY could be immense as it could help reduce the number of hubs needed by spanning longer links and provide more capacity by its robust SNR of digital transmission and regeneration.

Up to 44 WLs can be used, connecting clusters of nodes if 100 GHz C-Band DWDM optics are used. The number of links could be doubled if 50 GHz optics were used. This is a significant benefit of these architectures, that for a bit of initial effort at optimizing the service group areas, the potential migration to



R-PHY could have many options of spanning long distances and serving a large number of nodes with excellent SNR.

#### 4.2.4. N+0 and Preparing for FTTH

We discussed before that approximately 1% of the subs apply for the high billboard rates with a CAGR of 50%. If these growth rates hold, as year 2023 approaches these peak rates would approach and perhaps exceed 10 Gbps. If N+0 becomes prevalent by that time, it now becomes possible to selectively migrate subscribers that demand billboard rates approaching 10 Gbps to perhaps an FTTH PON system. When that happens, the next highest tier, the performance tier, is likely to be in the 1 Gbps range and therefore it now becomes possible to size the network appropriately and provide a longer runway for the current infrastructure for an additional 4 to 6 years.

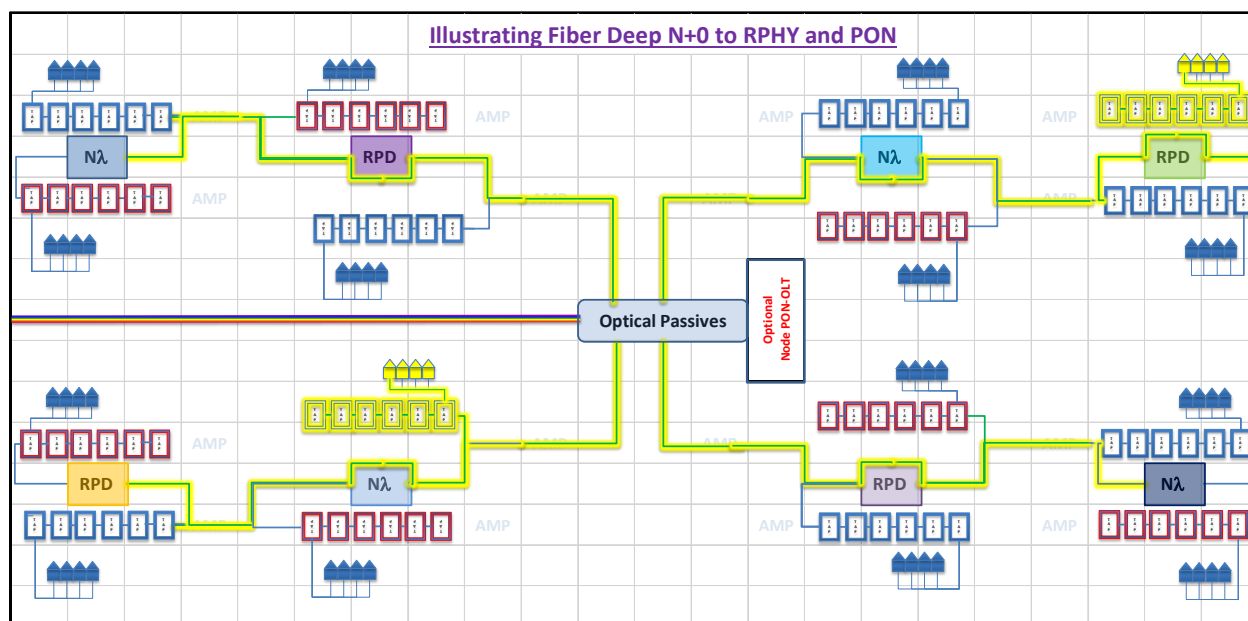


Figure 5 - R-PHY and PON in N+0 Architecture

Figure 5 illustrates the process of selective subscriber migration. For illustration purposes, we have shown a service group comprising a mixture of RPDs and fiber deep nodes. When customers requesting the high billboard rates are identified, it is possible to add a node PON optical line termination (OLT) next to the optical passives splice enclosure case and construct fiber to the required homes. When this is done, the specific home is served by the node PON OLT and will need to have PON ONUs. Note that this fiber construction can be done either when RPD or regular fiber nodes are deployed. For the vast majority of homes, regular cable modems and set top boxes continue to function.

As an aside, looking at the FTTH migration question in this context also allows us an additional insight in to how to deploy fiber. In a conventional PON system, it is advantageous to deploy fiber via distributed splitting. By this we mean that a fiber is not just split at one central location, but split at continuous intervals at discrete but different locations. It makes little difference in a PON environment if the split be centralized or distributed, as long as the ultimate link budget is conserved. However, if the aim is to ease future transition, then it is better to have centralized splitting from the parent node location as much as

possible. On the face value and in the limit, it appears as if it requires one fiber per HHP, but when considered that fiber deployment follows the RF cabling and there are limited numbers of HHP per port of the node, it appears that access to a 24-48 fiber bundle per port would reliably support most fiber deep deployments. The advantage of the centralized split is that it now enables one to go point-to-point over the life of the network. This can have major benefits throughout the life of the networks.

### 4.3. The FTLA Architecture

#### 4.3.1. Aging RF Plant Discussion

The HFC plant has been in continuous use for the better part of two decades, riding the big upgrade waves of 1996-2000. Due to the age of the plant and based on other factors, the 6 million to 8 million amplifiers in the plant can experience varying but increasing rate of failures year over year. A mere replacement of all the amplifiers would consume a fair amount of resources and end up decreasing failure rates, but not really contribute to driving fiber deeper to be prepared for what is next by way of competition. Perhaps realizing that there is less of ‘bang for the buck’, the industry has reluctance to a wholesale deployment of ADI.


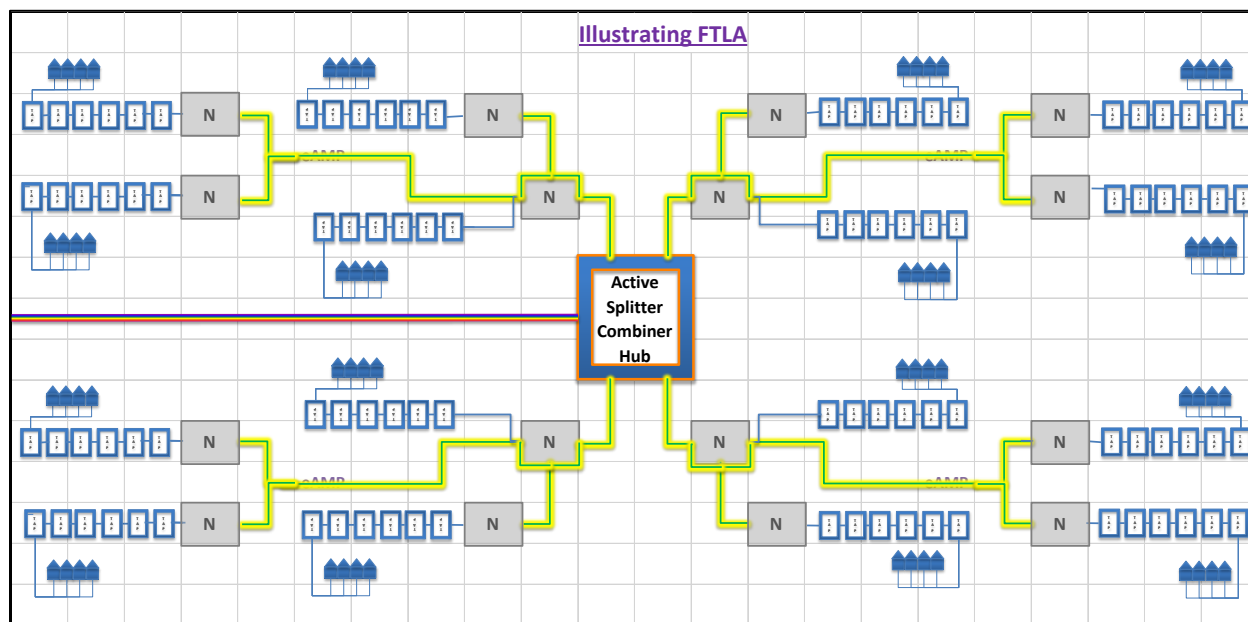


Figure 6 - Amplifier Failures per Year

Presented above is a typical estimate for amplifier reliability year over year. What this shows is that the first eight years of the amplifier has a cumulative failure rate of less than 2%. By 10 years, the cumulative failure rate is around 4%. However, by the 15<sup>th</sup> year, the failure rate is around 5% per year and by year 20, that rate is around 13% per year. Clearly this would degrade system reliability and as discussed earlier, faced with this situation, there is an option of replacing the current crop of amplifiers with newer amplifiers, or use this as an opportunity to future proof the system and invest in fiber deep architectures and replace the current set of amplifiers with fiber nodes.

We have previously seen that for a typical node of 30,000 hardline cable plant feet, around 10,000 hardline cable plant feet connect the node to the last active. The remaining 20,000 hardline plant feet is the cable plant that connects the last active to the last tap position along the curb. For the same system, we

have seen that there is around 60,000 feet of drop cable. Therefore, replacing 33% of hardline cable plant with fiber is a much more achievable target.



**Figure 7 - Basic FTLA Architecture**

For those MSOs that use analog reverse path, FTLA can be an attractive option. It enables MSOs to drive fiber deeper gradually by replacing the RF amplifiers with nodes at the same location as current amplifiers. The express amplifiers that do not serve any of the homes can be eliminated as illustrated above. Since the nodes are at the same spot as current amplifiers, there is no need to flip or replace any of the taps if staying at 1 GHz. By going all the way to the last active, fiber will also need to be driven deeper into the network and the fiber construction cost must be considered. Finally, if 1.2 GHz capability across the network is desired, the taps will have to be replaced as indicated before in earlier sections.

However, note that the more numerous new nodes could substantially increase the number of wavelengths, transmitters, receivers, and CMTS ports in the headend currently in use, with higher average capacity per household but no consequent improvement in peak capacity as we discussed in earlier sections. A way to significantly increase the ability to cluster these new nodes while maintaining adequate SNR in the DS and US would make the move to FTLA easier and cost effective.

An active splitter combiner (ASC) technology used in RFoG networks that uses the 1550 nanometer (nm) band as the DS and the 1610 nm band as the US, when applied at the virtual hub in the parent node location, offers optical amplification capability for the downstream and the capability to aggregate optical upstream paths at the same location. This will significantly reduce the required number of wavelengths on the trunk fiber and conserve valuable headend space. Furthermore, the performance of the link is significantly enhanced in the upstream as the OBI free active combiner would aggregate the reverse paths right at the parent node location and retransmit to the headend. Since the optical amplification is provided at the end of the link in the parent node location, this has the net effect of being able to launch modest amounts of light in the trunk fiber thus reducing optical non-linearities and increasing link distances.

In the case of FTLA, the parent node location comprises the ASC, which is an active device, but consumes much less power than a node would in that location. Also note that the express amplifiers that do not directly connect to the home can be bypassed in the FTLA scheme. Thus the total number of actives is 17 in the FTLA system to the original 21 in the HFC diagram, around 80% of the HFC actives.

Since new nodes are all at the same location as the previous RF amplifiers, they can mimic previous RF levels, and therefore can have similar power requirements. Since the node is replaced by an ASC which draws considerably less power, and also due to the marginal reduction in actives count, the power draw of the FTLA network is around 80% of the HFC plant.

All in all, this architecture marginally reduces the actives count and power requirements, enables the use of current active locations, accommodates very small node sizes (16-48 HHP/FTLA nodes), and with the ASC provides significant optical aggregation capability thus limiting wavelengths on the trunk fiber and transmitters, receivers, and CMTS ports in the headend. When time is right for segmentation, the single ASC that fed the entire system across the four ports of the parent node location can be replaced by up to four individual ASCs each operating at different US/DS WL combinations and connected back to the headend, while the individual FTLA node returns themselves function on the standard 1610 nm WL. When time comes to move towards FTTH, there are a few options and will be discussed in the next sections.

#### 4.3.2. Remote PHY in the FTLA Architecture

Keeping in mind the description of RPDs mentioned earlier, migrating to Remote PHY is accomplished by using the parent node location with an RPD. The RPD output is fed to a DS transmitter in the parent node location, which feeds the ASC. The ASC aggregates all the reverse path analog signals and combines them in the RF domain and presents to the US of the RPD.

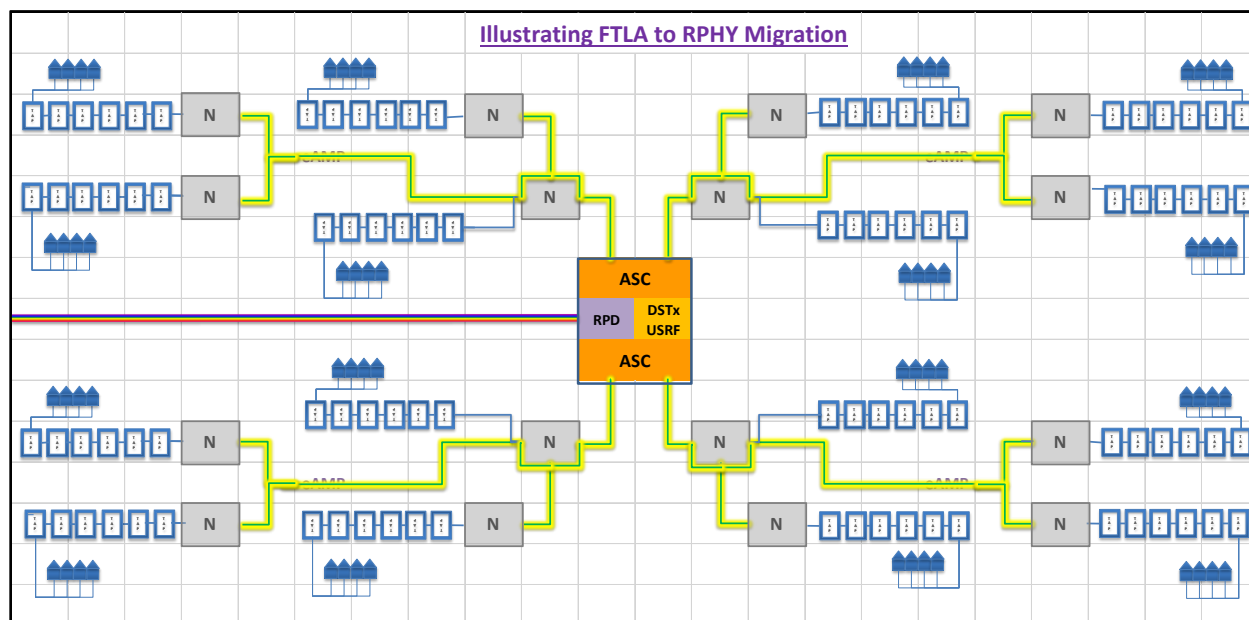


Figure 8 - Fiber Deep N+0 to R-PHY Migration

In the illustration presented in Figure 8, it is seen that the parent node location is now a proper sized node that enables long spans from the FTLA service group to the headend, while driving fiber deeper in the network. Note that the reverse path of all nodes continues to be analog. Up to 44 WLs can be used, connecting clusters of nodes if 100 GHz C-Band DWDM optics are used for the RPD. The number of links could be doubled if 50 GHz optics were used. The potential migration to R-PHY could have many options of spanning long distances and serving a large number of nodes.

### 4.3.3. Preparing FTLA for the FTTH Migration

To build out FTTH in an FTLA environment for the subset of customers that need high capacity, a node PON OLT is dropped at the parent node location. Fiber is constructed from the node location to the home and a PON ONU dropped there. The remaining HHP are fed using the RPD or just the ASC as before.

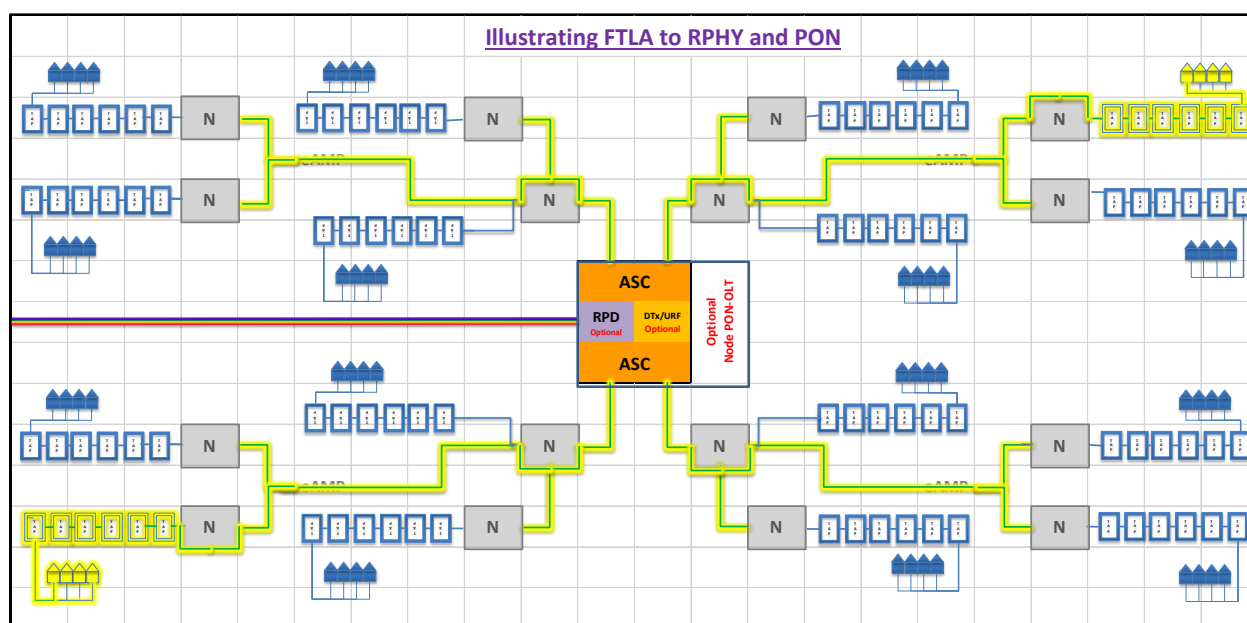


Figure 9 - N+0 to R-PHY, FTTH Selective Subscriber Migration

Figure 9 illustrates this point with an RPD and ASC combination. However, it can be done the same way if just the ASC existed at the parent node location. In cases where the MSO gets ahead of the FTTH build process before all of the customers necessarily need high billboard rates, the MSO can place RFoG ONUs at each home and reuse the infrastructure at the HE and home. This design is then very similar to RFoG, except that the fiber build process may be more gradual than an immediate all fiber requirement for a new RFoG design.

## 5. Conclusions

As indicated at the beginning of this paper, going fiber deep is a journey, not an event. For cable operators, in order to reach any of the above fiber deep architectures, new fiber construction to starts from the parent node location and progresses towards the home, with multiple strategic and logical stops. The use of DWDM DML multi-wavelength capability helps relieve fiber exhaustion cost effectively while the ASC devices help de-link fiber construction from a need to provide for PON. This approach is important



since it provides a way to conserve infrastructure so that fiber construction all the way to FTTH can proceed without a wait for all the IPTV issues to be resolved.

The prime virtue of fiber deep N+0 is the total elimination of all RF amplifiers and a reconfiguring of the network in a way to deploy the minimum number of new nodes. We have seen in earlier illustrations that the total number of new actives is reduced by almost 80% relative to the HFC plant and provides for power savings in the vicinity of 50%. N+0 is especially well suited to digital return systems as the efficiency and cost effectiveness of N+0 can be considerably enhanced by clustering the downstream and daisy chaining the upstream.

Fiber to the last active is well suited for analog return systems and affords 20% reduction in actives and power. Since the current position and spacing of amplifiers is not re-examined there are more numerous actives in the plant. When the FTLA is use in conjunction with the ASC it enables the MSO to cost effectively replace ageing HFC plant while driving fiber gradually into the HFC plant.

Given the aforementioned analysis, there are many factors and circumstances that will favor one architecture over another, and more so for one operator than another. Construction costs dominate in the quest for driving fiber deeper. While there are technologies now available to fight trunk fiber exhaustion, it is an ideal opportunity to build the fiber deep plant with the ability to support point-to-point links to the home from the node onwards. Such ability would unleash an almost unlimited bandwidth capacity per home.

## 6. Acknowledgements

It is a pleasure to acknowledge David Job and Harj Ghuman of Cox Communications for the many discussions and insights in fiber deep architectures and deployments. Our thanks to John Ulm and Ayham Al-Banna for their insights into traffic modeling and full duplex DOCSIS architecture. We thank Wim Mostert and Stewart Eastman for discussions on N+0 and FTLA designs.

## 7. Abbreviations and Definitions

### 7.1. Abbreviations

ADI	active drop in
ASC	active splitter combiner
bps	bits per second
CAGR	compound annual growth rate
CCAP	converged cable access platform
CMTS	cable modem termination system
DML	directly modulated lasers
DOCSIS	Data-Over-Cable Service Interface Specifications
DS	downstream
DWDM	dense wavelength division multiplexing
EDFA	erbium doped fiber amplifier
EPON	Ethernet passive optical network
FD	fiber deep
FD N+0	fiber deep N+0
FDX	full duplex DOCSIS

FEC	forward error correction
FTLA	fiber to the last active
FTTC	fiber to the curb
FTTH	fiber to the Home
FTTP	fiber to the premises
FWM	four wave[length] mixing
Gbps	gigabits per second
GHz	gigahertz
GPON	gigabit passive Optical Network
HFC	hybrid fiber-coax
HHP	households passed
HPON	hybrid passive optical network
HSD	high-speed data
IPTV	Internet protocol television
Mbps	megabits per second
MHz	megahertz
MPEG	Moving Picture Experts Group
MSO	multiple system operator
MWL	multi wavelength optics
N+0	node plus zero
nm	nanometer
OBI	optical beat interference
OFDM	orthogonal frequency division multiplexing
OLT	optical line termination
ONU	optical network unit
PHY	physical layer
PON	passive optical network
QAM	quadrature amplitude modulation
QoE	quality of experience
RF	radio frequency
RFoG	radio frequency over glass
RPD	remote PHY device
R-PHY	remote PHY
SDV	switched digital video
SNR	signal-to-noise ratio
SRS	stimulated Raman scattering
US	upstream
VoD	video on demand
WL	wavelength
XPM	cross phase modulation

## 7.2. Definitions

Downstream	Information flowing from the hub to the user
Upstream	Information flowing from the user to the hub

## 8. Bibliography and References

- [1] *The Yin and the Yang of a Move to All Fiber: Transforming HFC to an all Fiber Network While Leveraging the Deployed HFC Assets*; Venk Mutalik, Marcel Schemmann, Zoran Maricevic, John Ulm; INTX 2015 Spring Technical Forum
- [2] <https://fiber.google.com/newcities/#viewcities>
- [3] *When Wavelengths Collide, Chaos Ensues*; Venk Mutalik, Fernando Villarruel, Kuang-Yi Wu, Marcel Schemmann; 2011 SCTE Cable-Tec Expo
- [4] *How Network Impairments affect Customer Perceptions: The role of Leading and Lagging Indicators as We Evolve HFC to FTTP.*; Sebnem Zorlu Ozer, Ph.D., Venk Mutalik, A. Vieira, Ph.D., J. Chrostowski; From OBI and SNR to Ookla and Netflix: SCTE Cable-Tec Expo 2015
- [5] *Simulating the Impact of QoE on Per-Service Group HSD Bandwidth Capacity Requirements.* Tom Cloonan, Jim Allen, Mike Emmendorfer' 2014 SCTE Cable-Tec Expo
- [6] *Is Nielsen Ready to Retire? Latest Developments in Bandwidth Capacity Planning*; John Ulm, T. Cloonan, M. Emmendorfer, J. Finkelstein, JP Fioroni; 2014 SCTE Cable-Tec Expo
- [7] <http://www.broadbandtvnews.com/2016/04/28/g-fast-technology-brings-1-gbs-over-copper/>
- [8] *Giving HFC a Green Thumb, a case study on Access Network and Headend Energy & Space Considerations for Today and Future Architectures*; John Ulm and Zoran Maricevic; SCTE 2016
- [9] *Adding the Right Amount of Fiber to Your HFC Diet: A Case Study on HFC to FTTx Migration Strategies*; John Ulm, Zoran Maricevic; 2016 SCTE Cable-Tec Expo
- [10] *Distributed Node Architecture and More*; Venk Mutalik; FTTH Connect 2016

# Framework for The Evaluation of SDN WAN Controllers

A Technical Paper prepared for SCTE/ISBE by

Javier Ger, Expert, Sr. Network & Services Architect, Cablevision Argentina SA,  
SCTE/ISBE Member  
Hornos 690  
CABA, C1272ACL - Argentina  
jger@cablevision.com.ar  
+54-1155304531

Eduardo Panciera, Expert, Head of Technology and Architecture, Cablevision Argentina SA,  
SCTE/ISBE Member  
Hornos 690  
CABA, C1272ACL - Argentina  
epanciera@cablevision.com.ar  
+54-1155306562

## Table of Contents

Title	Page Number
Table of Contents	129
1. Abstract	130
2. Introduction	130
3. Aspects to be evaluated	133
3.1. Interfaces	133
3.1.1. D-CPI/Southbound	134
3.1.2. A-CPI/Northbound	135
3.1.3. East/West	137
3.2. Reliability and Redundancy	139
3.2.1. Transactional Integrity	139
3.2.2. Redundancy	139
3.2.3. Integrated/Single Multi-Layer Controller	140
3.3. Performance and Scalability	140
3.4. Applications	141
3.4.1. Topology Discovery and Maintenance	143
3.4.2. Path Provisioning	143
3.4.3. Multi-Layer Optimization	143
3.5. Security	144
3.6. Controller Deployment, Maintenance and Support	146
4. Conclusions	147
5. Abbreviations	148
6. Bibliography and References	150

## List of Figures

Title	Page Number
Figure 1 - SDN Architecture by ONF (TR-502 - SDN Architecture 1.0)	131
Figure 2 - Multi-Service Converged Networks Hierarchies Considered In This Paper	132
Figure 3 - East/West Interfaces Integrating Controllers from Different Domains And Working with Redundancy (clusters)	137
Figure 4 - PCE Hierarchies	138



## 1. Abstract

Currently, software defined networks (SDNs) are booming in the industry. This new paradigm along with the advantages it offers and the architecture recommended by standards development organizations (SDOs), is applicable to multiple scenarios and use cases.

The previously-mentioned architecture is made up of the application plane, the controller plane and the data plane. In the controller plane resides the controller itself, which constitutes a key element in this architecture.

Whenever a new paradigm arises, it is important to define the aspects to bear in mind to evaluate its components as there is not enough experience in the industry and some of these aspects may not be necessarily evident at first. Although these aspects can evolve as the innovation and knowledge of the new proposal progress, their definition should enable greater chances of achieving the full advantages and benefits that give meaning to the proposed changes.

The aspects just mentioned constitute the framework of the new paradigm. Among these aspects, we can mention features, functionalities and related blocks, expected behaviors, interconnections and interfaces, impact or integration in the current models, and others.

In this paper, we specifically describe the main aspects comprising the framework used by Cablevisión Argentina to evaluate SDN wide area network (WAN) controllers (aggregation/distribution/core hierarchies of multi-service converged networks). This framework can be applied both to solutions developed by vendors and integrators as well as to open-source or individual telecommunications service providers' versions.

This framework can be used by service providers regardless of the access technology used, even if the framework includes some specific details for those companies that have deployed Data-Over-Cable Service Interface Specifications (DOCSIS<sup>®1</sup>) technology.

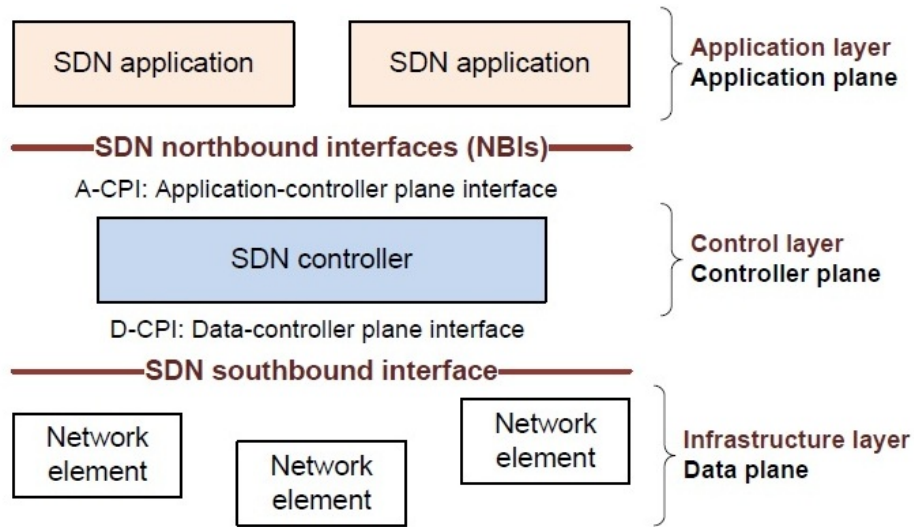
## 2. Introduction

As discussed, SDN proposes a redefinition of current network architecture (see Figure 1<sup>2</sup>).

---

<sup>1</sup> DOCSIS is a trademark of CableLabs

<sup>2</sup> Figure 1 introduces the SDN architecture proposed by ONF (TR-502 - SDN Architecture 1.0). Both IRTF (RFC 7426 - SDN: Layers and Architecture Terminology) and ITU-T (Y.3300 - Framework of SDN) have also proposed very similar architectures with three main layers or planes. Recursion described by ONF, differentiation between control plane and management plane from the IRTF and even the new proposals of the IRTF (still draft, draft-irtf-sdnrg-layered-sdn-00) that divide the control plane into service and transport stratum are supported by this basic architecture.

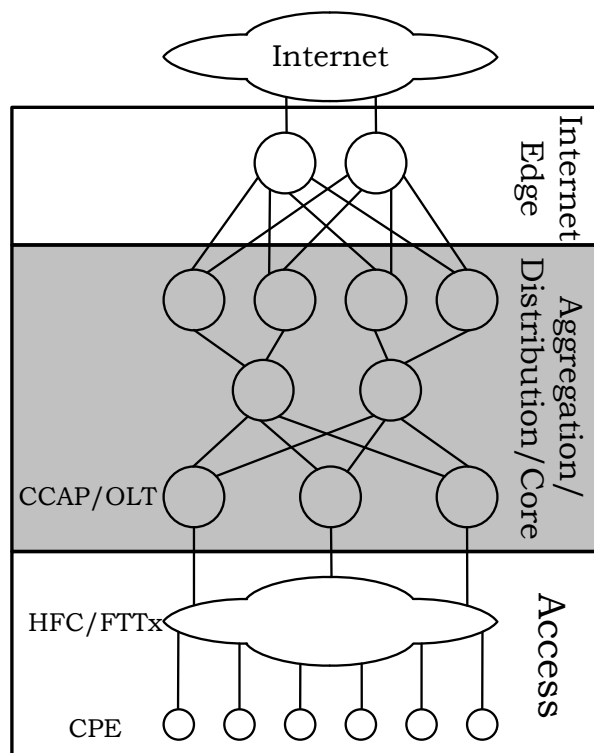


**Figure 1 - SDN Architecture by ONF (TR-502 - SDN Architecture 1.0)**

This new architecture is applicable to several sectors of current networks, such as datacenters, access networks with specific characteristics (e.g., radio access – Wi-Fi, 3G, Wimax, LTE, etc.) and the aggregation/distribution/core hierarchies of multi-service converged networks. In each one of these scenarios, the use cases and issues to solve to enable the ultimate SDN goals<sup>3</sup> are diverse. For these reasons, the industry makes explicit differentiations regarding platforms/products for them.

Specifically, when the case of hierarchies of aggregation/distribution/core is considered (see Figure 2), where devices and elements working in transport of multi-service converged networks and mainly in the physical, link and network layers of the ISO/OSI model are located, controllers residing in controller plane/control layer of SDN architectures are commonly referred to as WAN controllers within the industry. In some cases, there are applications that are directly included in these controllers and make it possible to address typical use cases of this sector of the network.

<sup>3</sup> In this paper it is considered that these ultimate goals are more flexibility, programmability, automation, openness and interoperability of networks and their components that will, in turn, generate higher revenues from new services and less CAPEX/OPEX. To achieve them the architecture of Figure 1 is proposed.



**Figure 2 - Multi-Service Converged Networks Hierarchies Considered In This Paper**

Based on what is being presented, this paper has the following scope:

- Hierarchies of aggregation/distribution/core of multi-service converged networks (See Figure 2) and, within them, devices and elements in the physical (L1), link (L2) and network layer (L3) of ISO/OSI model (also including what is sometimes known in the industry as photonic layer and/or layer 0).
  - As the access networks aggregation devices can represent the beginning of transport in multi-service converged networks, they are included in the mentioned hierarchies. Moreover, considering the diversity of functions performed by these devices, they could be the SFC<sup>4</sup> initiation points as discussed in the next paragraph. On the other hand, although in section 3.4 some use cases are proposed considering networks with DOCSIS access technology (CMTS/CCAP)<sup>5</sup>, they can be replaced or extended to other types of access aggregation devices (DSLAM, OLT<sup>6</sup>, etc.).
  - Network elements from upper layers of the ISO/OSI model are not currently included but may be in the future. This is so SFC cases can be considered where transportation is determined by traffic classification based on the information that resides in ISO/OSI model upper layers.
- Controller plane/control layer within the SDN architecture where the controller resides and has centralized network intelligence (see Figure 1). Focus is made specifically on WAN controllers.

4 ONF - TS-027 - L4-L7 Service Function Chaining Solution Architecture.

5 ODL controller (<https://www.opendaylight.org>) supports COPS as D-CPI since Helium release.

6 ONOS controller (<http://onosproject.org/>) has defined a use case called CORD which runs a virtualized OLT on general purpose hardware. This scenario could be controlled by a WAN controller.

- Following the aforementioned, in this paper a number of aspects<sup>7</sup> that constitute a framework for evaluating<sup>8</sup> SDN WAN Controllers are suggested. This proposal is based, in turn, on the current Cablevision Argentina framework (version 1.0), considering its most significant aspects.

### 3. Aspects to be evaluated

The suggested criterion to analyze and evaluate WAN controllers is considering them a black box defined by their external behavior and abstracting from implementation details. This is consistent with the definitions made by ONF<sup>9</sup> related to SDN architecture and allows certain consistency in some comparisons, although this approach may have limitations<sup>10</sup>. The following six aspects constitute the proposed framework.

- 1- Interfaces.
- 2- Reliability and Redundancy.
- 3- Performance and Scalability.
- 4- Applications.
- 5- Security.
- 6- Controller Installation, Maintenance and Support.

#### 3.1. Interfaces

This aspect defines conditions as interoperability with applications (A-CPI/Northbound), network elements (D -CPI/Southbound) and even other controllers (East/West). Thus, interfaces that a WAN controller has are an essential aspect to consider, since they determine its ability to run with platforms, solutions, architectures and services from different origins (e.g., open-source, other vendors, white boxes, integrators).

In short, this interoperability option is what determines the possibility to reach open solutions (i.e., those that integrate these different origins and are also neutral to any of them) which is one of the means defined by SDN to achieve their goals.

---

<sup>7</sup> In the context of this paper, aspects are understood as the features and functionalities set and related blocks, expected behaviors, interconnections and interfaces, impact or integration into existing models, etc., to be included in the analysis and evaluation of a WAN controller. These aspects enable the advantages and benefits of this new architecture in the mentioned sectors of multi-service converged networks.

<sup>8</sup> It can be of use in individual controllers evaluations, controllers eligibility for use cases, consistent comparisons among different controllers, etc.

<sup>9</sup> ONF - TR-502

<sup>10</sup> These limitations could be given mainly with platforms arising from open-source community or, even, developed by individual cable operators. In these cases, there could be greater granularity in the analysis and evaluation of aspects of the WAN controller, with also a greater choice of defining implementation details that can impact on its features, functionalities, performance and/or advantages.

### 3.1.1. D-CPI/Southbound

The controller should support a set of standard interfaces to network elements as they enable the architecture integration with multiple vendors (even white boxes if applicable). In general, there is a relationship between the amount of supported interfaces and the extension of the ecosystem that can be integrated.

The weighting of specific interfaces that should support a WAN controller is conditioned by the target infrastructure, particular use cases and solutions to cover, along with their evolution. It is important to consider covering an ecosystem as broad and open as possible.

Currently, some interfaces can be of particular interest.

NETCONF<sup>11</sup> – It can work with a wide range of architectures, from those aimed at a SDN pure model to hybrid models with different levels and combinations of centralized and distributed control in the network elements. This is usually accompanied by a data model known as YANG<sup>12</sup>.

OpenFlow/OF-CONFIG – Original interface of the SDN model known as pure, standardized by the ONF. Despite not having an extensive reception in use cases related to WAN controllers and network hierarchies considered in this paper<sup>13</sup>, their developments and medium-term acceptance should be monitored<sup>14</sup>.

OF-PI<sup>15</sup> /P4<sup>16</sup> /POF-FIS<sup>17</sup> – The most revolutionary proposals related to D-CPI interfaces<sup>18</sup>, which allow additional levels of programmability – provided network elements contain programmable hardware. Actually, new protocols, actions, instructions or network behaviors can be defined and applied<sup>19 20</sup>, independently from device vendors' innovation cycles. They are not standardized by SDOs yet, but P4<sup>21</sup> and POF<sup>22</sup> have specifications (and code already developed).

SNMP – It essentially enables information collection (also configuration, although it is a much less common scenario and with certain limitations<sup>23</sup>) from network elements. This information (interfaces

11 Considering current leading vendors (February 2016) of optical and MPLS/IP network elements for transport networks having a 70% aggregate market share, all of them provide in their WAN controllers support to NETCONF, and only about 60% of them support OpenFlow. ODL and ONOS support both.

12 IETF RFC 6020 - YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF). This data model has a large number of versions. Also see OpenConfig (<http://www.openconfig.net/>) and IETF NETMOD WG.

13 Mainly in the case of OF-CONFIG.

14 At present (February 2016) virtually no optical platform (DWDM) in the market supports OpenFlow, despite the extensions that are being done to the protocol since version 1.4 for configuring parameters such as transponders wavelength or power. Consideration is then needed as to whether the definition of the protocol is suitable for this type of technology. Additionally, ONF-Open Datapath WG will be working in 2016 on OpenFlow data models (TTP 1.1) to improve this protocol interoperability.

15 ONF – TR-505 – OF-PI: A Protocol Independent Layer and ONF Protocol Independent Forwarding and <https://www.opennetworking.org/protocol-independent-forwarding>.

16 P4 Language Consortium (<http://p4.org/>).

17 POF (<http://www.poforwarding.org/>).

18 Depending on the architecture some components can be located north of the controller, but they are not a typical business application but intrinsic part of the controller.

19 ONF – TR-505 – OF-PI: A Protocol Independent Layer.

20 <http://p4.org/p4/clarifying-the-differences-between-p4-and-openflow/>

21 P4 Language Consortium - The P4 Language Specification Version 1.1.0 ([http://p4.org/wp-content/uploads/2016/03/p4\\_v1.1.pdf](http://p4.org/wp-content/uploads/2016/03/p4_v1.1.pdf)).

22 POF (<http://www.poforwarding.org/documents/>).

23 IETF - RFC 3535 - Overview of the 2002 IAB Network Management Workshop.



utilization, QoS queues, interfaces/equipment failures, etc.), supplemented with those that can be obtained through XFLOW is very valuable for certain use cases such as multi-layer capacity optimization, TE, etc.

**PCEP** – It enables the learning of LSP state information from network elements on a MPLS/GMPLS network and even their update from a centralized PCE. It also includes extensions for SR<sup>24</sup>. The use case can determine the architecture of WAN controllers to use, considering the recursion details proposed by the ONF architecture<sup>25</sup> and the work being developed in the IRTF<sup>26</sup> (some in draft form).

**BGP-LS** – It collects the information stored in L3 topological databases of the link-state IGP and integrates it with PCE/PCEP (and SR<sup>27</sup>). There might be other solutions as effective as this one depending on the particular network architecture<sup>28</sup>.

**COPS** – It is a specific interface for the case of DOCSIS technology (CMTS) and/or convergent (CCAP<sup>29</sup>) devices in HFC networks and can cover specific use cases associated with telecommunications service providers coming from the cable industry (originally MSOs). Therefore, it may be valuable to support it<sup>30</sup>, though it is an interface that has not had great momentum and its productive use cases are relatively few.

There are other alternative interfaces that may have some effect on this evolution path, such as TL1 (for the optical layer devices, mainly legacy), SSH/CLI, I2RS, LLDP, LISP, etc.

Additionally, proprietary interfaces provide very specific value as they allow the integration of very specific devices.

On the other hand, given the aforementioned recursion proposed by the ONF for SDN architectures it may be necessary, depending on the particular use cases and controllers architecture, to support A-CPI interfaces (see Section 3.1.2) as D-CPI interfaces.

### **3.1.2. A-CPI/Northbound**

The WAN controller should support interfaces to applications (API). These interfaces are very important because ultimately they enable application portability between controllers.

Presently, their definition tends to be discretionary since there are no standards covering them. This can lead to great fragmentation, according to particular interests, without interoperability between different solutions and implementations and requiring that many of these interfaces be developed to meet different use cases. Some SDOs<sup>31</sup> and open-source communities propose developing concrete implementations of

24 SPRING - <https://tools.ietf.org/wg/spring/> and draft-ietf-pce-segment-routing-07.

25 ONF - TR-502

26 IETF PCE WG – draft-dhodylee-pce-stateful-hpce-00 and draft-chen-pce-h-sdns-00. Also IETF - RFC 6805 - The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS.

27 IETF draft-ietf-idr-bgpls-segment-routing-epe-05.

28 Centralized control and definitions of restrictions and path attributes (latency, economic cost, bandwidth, etc.) working with some combination of southbound interfaces.

29 Including future support to PON networks.

30 As previously mentioned, ODL supports this D-CPI since Helium release.

31 ONF - Northbound Interfaces Charter – NBI WG.

the interfaces in question prior or parallel to this standardization process<sup>32</sup>. In any case, both groups are working on this topic<sup>33</sup>.

There is some consensus in the industry regarding WAN controllers support for REST<sup>34</sup> interfaces due to the flexibility they provide (given their architecture, operations, syntax, etc.) for integration with applications. They work with data models based on YANG<sup>35</sup> (XML or JSON) and there is a draft in progress in the IETF regarding REST working with NETCONF and YANG data models<sup>36</sup> (again, XML or JSON). There are also platforms that support XMPP as A-CPI.

The industry is also considering having "Intent" type interfaces that enable applications to indicate their needs to the network but not the details of what to do to meet them, while being independent of protocols, vendors and media. The rationale behind this is that in an environment of dynamic and programmable networks, it should be possible to generate infrastructure changes from business changes. Moreover, policies that generate these changes are not defined by technical people, so it should be also possible to formulate them in natural language to improve their use and deployment options<sup>37</sup>. Although it is beyond the scope of this paper, it is relevant to note that in some cases this can have implications even in some functional blocks within the controller, such as the Virtualizer mentioned by the ONF architecture<sup>38</sup>.

Taking into account the importance of these interfaces to support a proper ecosystem of applications and WAN controllers, their simplicity and, if complete standardization is not possible, at least the existence of a framework in the short/medium term are essential. This framework should minimize incompatibilities between applications and WAN controllers, facilitating their interoperability, portability and integration. As a side benefit, developers will focus on improving and differentiating their applications and on controllers' performance and features/functionality diversity.

In any case, the evolution of these interfaces should be monitored, considering there is still abundant work to be done and the industry attention on them.

Finally, as mentioned in Section 3.1.1 and considering the recursion proposed by the ONF for SDN architectures, it may be necessary, depending on the particular use cases and controllers architecture, to support D-CPI interfaces (see Section 3.1.1) as A-CPI<sup>39</sup> interfaces.

---

32 This is a clear example, but not the only one, of the discussion taking place in the industry about the balance between standardization and open-source or innovation in general.

33 ONF NBI WG, IETF NETCONF WG, IETF NETMOD WG and IETF ANIMA WG are examples. Also the improved interfaces to applications of open-source solutions, such as ODL and ONOS, through its successive versions.

34 Considering current leading vendors (February 2016) of optical and MPLS/IP network elements for transport networks having a 70% aggregate market share, almost all of them provide in their WAN controllers support to REST as A-CPI. Both ODL and ONOS support this interface.

35 Again, this data model has a large number of versions. Also, see OpenConfig (<http://www.openconfig.net/>) and IETF NETMOD WG.

36 IETF - draft-ietf-netconf-restconf-13.

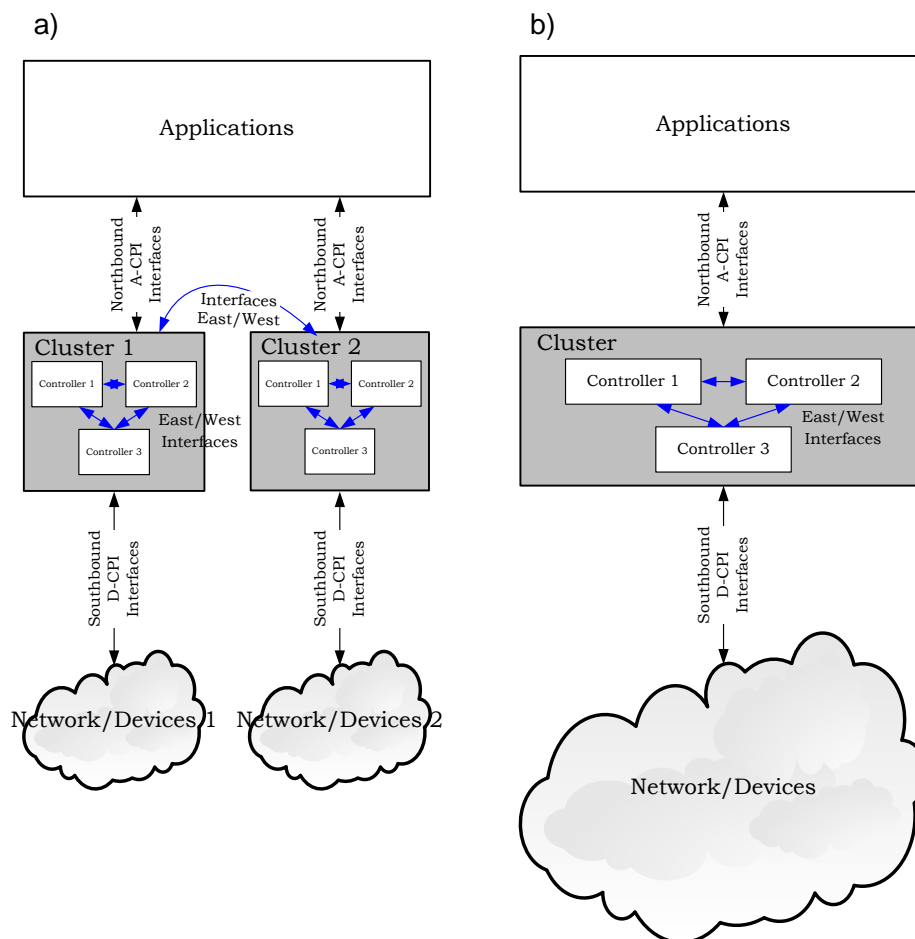
37 ONF NBI WG ONF will be working on this in 2016. Furthermore, there is a project called BOULDER based on open-source. IETF also has some drafts on this topic, e.g. draft-bi-supra-problem-statement-00 and draft-du-anima-an-intent-03.

38 ONF - TR-502

39 ONF - TR-502 and ONF - Northbound Interfaces Charter – NBI WG - API Abstractions and Latitude of NorthBound Interfaces.

### 3.1.3. East/West

The SDN WAN controller should support interfaces to other controllers to enable integration of different control domains or redundancy (see Figure 3). Similar to what has been described in the preceding section, these interfaces and the aforementioned scenarios have no standardized solutions yet. The main difference is that in this case the industry dedication currently evidenced to resolve the A-CPI interfaces is not observed. Both examples show that SDN is an evolving concept.



**Figure 3 - East/West Interfaces Integrating Controllers from Different Domains And Working with Redundancy (clusters)**

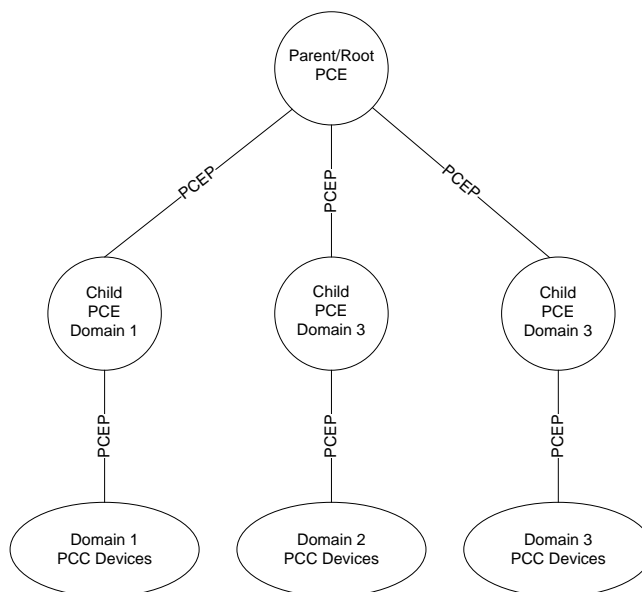
In the case of different controllers and/or clusters that control different administrative domains (see Figure 3 - a) an interface is needed to exchange information and define the end-to-end path<sup>40</sup> for both, the services and the transport – overlay and underlay. There are some approaches in this regard based on PCE hierarchies (see Figure 4) and PCEP<sup>41</sup> or MDSO<sup>42</sup> architectures. They are consistent with the concept of recursion proposed by the ONF, as discussed in Section 3.1.1, so the previously mentioned case is

<sup>40</sup> Business issues have to also be addressed -commercial terms, pricing, consumption, SLA, etc.

<sup>41</sup> IETF PCE WG – draft-dhodylee-pce-stateful-hpce-00 and draft-chen-pce-h-sdns-00. Also IETF - RFC 6805.

<sup>42</sup> ACM SIGCOMM 2015 - Multi-Domain Service Orchestration over Networks and Clouds: A Unified Approach.

resolved with the existing A-CPI and D-CPI interfaces or extensions to them. One issue to be revised in such architectures is possible inefficiencies noticeable when very strict response time from the centralized control plane is needed (e.g., reconvergence or others). Additionally, it would be desirable to extend their coverage, mainly in the PCE hierarchies cases, to other transport technologies besides MPLS/GMPLS.



**Figure 4 - PCE Hierarchies**

There are also proposed solutions using MP-BGP to exchange L3 routes between different domains<sup>43</sup>. It is very important that the industry moves forward in its standardization and identifies alternatives for the same information exchange for lower layers (L1 and photonics)<sup>44 45</sup>.

Another scenario, perhaps much more demanding in terms of complexity, is the redundancy of WAN controllers within the same domain (see Figure 3 – a and b). The industry has not yet achieved an acceptable level of maturity in this aspect, even with controllers from the same origin (source code, implementation, configuration, etc.)<sup>46 47</sup>. It would be desirable to address future standardization to allow the composition of redundancy schemes. These schemes should be possible among the mentioned same origin controllers and even platforms from heterogeneous or non-identical sources. This will be possible based on the resolution of certain behaviors and capabilities among them (e.g., replication, synchronization and consistency of persistent information and states, algorithms for selection of work

43 ODL Controller ([https://wiki.opendaylight.org/view/ODL-SDNi\\_App:Main](https://wiki.opendaylight.org/view/ODL-SDNi_App:Main)).

44 Ideally, through existing protocols and integrating information from different layers into a single protocol, considering the alternative to define a single multilayer controller.

45 Considering current leading vendors (February 2016) of optical and MPLS/IP network elements for transport networks having a 70% aggregate market share, very few of them provide in their WAN controllers support to BGP as East/West interfaces. ODL supports BGP.

46 Considering current leading vendors (February 2016) of optical and MPLS/IP network elements for transport networks having a 70% aggregate market share, almost all of them support WAN controllers clustering among the identical controller instances, however maturity of the solutions varies. The same applies to ODL and ONOS.

47 There are different approaches, in some cases relying on control plane of databases and/or storage solutions to achieve high availability, validity and integrity of persistent information. Then, network elements control availability must also be ensured.

mode and roles -active/active, active/passive, etc.) to create a cluster<sup>48 49</sup>. That is, through concepts of distributed systems<sup>50</sup>.

As mentioned, it is desirable that the industry moves forward to standardize or define these interfaces. They are of remarkable importance for the scenarios specified throughout this section (and possibly other scenarios that may arise in the future), allowing the integration of heterogeneous architectures and providing the necessary behavior certainty, given the sensitivity of the new architecture, mainly in its purest version, to the WAN controller availability.

### 3.2. Reliability and Redundancy

This aspect includes considerations such as transactional integrity and redundancy of WAN controllers (the latter is also conditioned by the East/West interfaces mentioned in Section 3.1.3). They are critical in any architecture of service providers to obtain consistent behavior and high availability, especially in the particular case of SDN, since its architecture becomes vulnerable, mainly in its pure conception, to the behavior or unavailability of the centralized controller.

It may also be considered within this aspect the option of a single multi-layer WAN controller from a logical point of view. This option facilitates transactions integrity, simplifies the exchange and synchronization of information and states, which in turn facilitates redundancy, and enables reliable integration through consistent information between the different layers of the ISO/OSI model (L1 to L3 and, even, what the industry knows as layer 0 or photonics). This last advantage should streamline the implementation of some applications and use cases mentioned in Section 3.4, such as multi-layer optimization.

#### 3.2.1. Transactional Integrity

This functionality is intended to allow the consistency of information and network conditions to ensure reliability. This requires the application of the "*all- or- nothing*"<sup>51</sup> principle, i.e., a certain requirement and/or event must be run on all network elements and with its complete sequence or, alternatively, return to the initial state to avoid unwanted or unpredictable behavior.

Transactional integrity is particularly important in recursive scenarios, including multiple domains, where each controller should provide this functionality within the domain under its responsibility.

#### 3.2.2. Redundancy

As previously mentioned, this concept is related to the issues discussed in Section 3.1.3 regarding interfaces between controllers (see Figure 3 - b). These interfaces are the means used to achieve redundancy and high availability of controller plane/control layer, with its associated controllers.

---

48 In this case a cluster is considered a single logical entity controlling a single administrative domain and having homogeneous behavior, information, states and answers, regardless of the node that takes the request. Thus, the cluster ensures its permanent availability and control over network elements. It is also a function of a cluster to guarantee the continued availability of controller states and persistent information, which are the basis of subsequent responses to the network elements.

49 IRTF - RFC 7426 – CAP Theorem (Consistency, Availability, Partition Tolerance).

50 ONF – TR-502 – Distributed controller considerations.

51 ONF – TR-502 – ACID (Atomicity, Consistency, Isolation, Durability).



Considering that a WAN controller is a critical element and its unavailability can lead, in the worst case scenario, to complete unavailability of the network (for example, in a pure SDN architecture without distributed control plane or with it being limited), it is essential to have a redundancy scheme. This scheme will ensure high availability in order to minimize the aforementioned events.

As part of this redundancy, it is proposed that each individual controller be part of a cluster, consisting of a group of controllers, or nodes, of identical or different origins. This cluster will be a logical entity that behaves as a single controller and, as indicated previously in Section 3.1.3, allows permanent availability of network elements control and both controller information and states.

### **3.2.3. Integrated/Single Multi-Layer Controller**

The ONF in its SDN architecture proposal<sup>52</sup> explicitly mentions it does not specify integration or separation of controllers based on ISO/OSI model layer. Other SDOs do not mention this topic at all. However, it would be convenient to propose as a working model an integrated/single multi-layer controller. This concept means a single logical entity (again taken as a black box) that is responsible for controlling the layers considered in this paper<sup>53</sup> within hierarchies of aggregation/distribution/core of multi-service converged networks.

The benefits of this proposal include the following.

- Transactions Integrity – Coordination of fewer elements, or at least all of them under only one control domain, in order to monitor transactions and potential application of roll-back processes (back to the initial state).
- Redundancy – Similar to the previous case, intervention and coordination of fewer elements, exchanging and synchronizing information and states, is required, or at least all of them under only one control domain, to guarantee permanent availability of network control.
- Applications and use cases – Improved integration with consistent and reliable information and states between different layers of the ISO/OSI model (L1 to L3, including layer 0 or photonics). This enables the implementation of some applications and use cases mentioned in Section 3.4. An example is the multi-layer network design optimization that applies complex models based on the analysis of statistical multiplexing, grooming, links and elements costs, etc., but with the peculiarity of a holistic approach through the different layers<sup>54</sup>.

### **3.3. Performance and Scalability**

This aspect analyzes response times, requests processing rates, capacity/sizing of the controller and how expansion takes place once thresholds are reached, which establishes the modularity and scalability of the platform. Considering the sector of multi-service converged networks that WAN controllers are covering, it is essential their performance and ability to scale, especially with the increase that traffic and services are presenting today and their projection in the coming years, creating an enormous pressure over the mentioned networks. This scenario becomes even more critical when combined with the dynamism SDN

<sup>52</sup> ONF – TR-502

<sup>53</sup> Considering current leading vendors (February 2016) of optical and MPLS/IP network elements for transport networks having a 70% aggregate market share, not all their WAN controllers allow multilayer integration. Normally each vendor supports the layer from which its core business (L3 or L0/L1) comes. Vendors who have presence in several of these previous layers support integration, but with some limitations.

<sup>54</sup> An example of this approach can be found in The International Journal of Computer and Telecommunications Networking 2008 - Multi-layer MPLS Network Design: the Impact of Statistical Multiplexing.

proposes, related to requirements of real-time responses to make services available while optimizing resources.

There are some IETF drafts describing methodologies for benchmarking performance of SDN controllers<sup>55</sup>, both standalone and cluster modes. They include relevant values<sup>56</sup> as topology discovery and change detection times, messages processing time and rate, path provisioning time and rate, network discovery size, control sessions and forwarding table capacity and controller failover and re-provisioning time. The same measurements should be performed in exceptional situations, such as significant increases in error messages coming to the controller at the same time as valid ones and under DoS attacks to the controller.

Other relevant values can be capacity for sessions to applications, flows, routes<sup>57</sup> and traffic or, even, execution of some tests on concurrent failure and reconvergences scenarios. There could be additional limits, thresholds or restrictions in the controllers operation. It is essential in each case to identify and include them in the analysis.

Additionally, the complexity and reliability of growth options once reaching the limits identified will define the WAN controller modularity to facilitate its scalability. This has to be evaluated along with the likelihood of achieving the mentioned limits by the network in question.

Some efficiency factor should also be included as part of the analysis. This factor can be calculated as a ratio of some of the previously mentioned items over some solution requirements, such as physical configuration (CPU, memory, disc, etc.) and dimensions (RU), power consumption, cooling needs, etc.

These items should be weighted based on the size of the controlled network, the specific use cases and solutions to be covered and their estimated evolution<sup>58</sup>.

Finally, SDOs are working on the definition of performance test methodologies for controllers<sup>59</sup>.

### 3.4. Applications

As part of this aspect the use cases currently covered by the controller under analysis should be considered, either through own developments (identical source as the controller) or from third parties. These applications can reside in the controller itself as part of the controller plane/control layer or, even, in the application plane/layer<sup>60</sup>. In the latter case, the advantage would be simplicity for their portability between controllers versus a theoretical better performance offered in the first case<sup>61</sup>.

While it is desirable that standard A-CPI/Northbound interfaces be defined, favoring application portability between different controllers, today this scenario is not yet materialized (as mentioned in section 3.1.2). For this reason, it is important to examine applications and use cases already developed for

55 IETF - draft-bhuvan-bmwg-sdn-controller-benchmark-term-01 and draft-bhuvan-bmwg-sdn-controller-benchmark-meth-01.

56 Although some of these functionalities and/or applications can reside in the controller or application plane, they are always executed in the network via the controller.

57 ITU-T Y3.300.

58 The defined security architecture should also be considered in this weighting (see section 3.5).

59 Besides the IETF drafts (draft-bhuvan-bmwg-sdn-controller-benchmark-term-01 and draft-bhuvan-bmwg-sdn-controller), the ONF announced that it will publish a methodology for evaluating SDN controllers performance in Q3 2016 and it will be coordinated with the work being done by the IETF.

60 ONF – TR-502, IRTF - RFC 7426, ITU-T Y.3300.

61 The chosen option is discretionary, considering the advantages and disadvantages of both approaches.

the controller in question. In general, the more applications and use cases available for a particular controller, the greater its potential.

Existing third party applications for the controller are a significant component. They are indicating a tendency to be integrated into an ecosystem with elements of various origins, still in the indicated absence of clear definitions regarding the A-CPI interfaces. These applications might also indicate certain acceptance level in the industry or some maturity level of a particular controller. This last item is a very indirect indicator and must be evaluated in conjunction with the other aspects discussed throughout this paper.

Considering what has just been mentioned, it would be appropriate to consider giving applications developed by third parties a higher weighting than those of identical origin as the controller, for the same use case.

The following WAN controllers use cases, each one self-explanatory, exemplify some that could be remarkable for telecommunications services providers in general and cable companies (HFC/DOCSIS technology) in particular. In any case, each operator should define its own weighting based on business needs and their evolution.

- Topology Discovery.
- Path Provisioning.
- Multi-Layer Optimization.
- TE (SR, RSVP) – Can be included or be part of the Multi-Layer Optimization case.
- Network Planning (growth trends, simulations of new products/failures/maintenance – “*what if*”-scenarios) - Can be included or be part of the Multi-Layer Optimization case.
- Bandwidth on Demand.
- SFC<sup>62</sup> - Can even consider optimizing the location of the middleboxes<sup>63</sup>.
- DoS attacks mitigation<sup>64</sup>.
- Software Defined Internet Exchange Point (SDX)<sup>65</sup>. Multidomain routing based on multiple fields packet processing rules for traffic forwarding.
- Integration to IaaS platforms<sup>66</sup> as a foundation for NFV MANO<sup>67</sup> scenarios.
- Given the characteristics of companies with an HFC/DOCSIS access networks, some particular applications/use cases related to this type of scenarios are mentioned.
  - PS/PCRF with CMTS/CCAP acting as PEP via COPS<sup>68</sup>.
  - CCAP distributed architectures with L1/L2 functionalities in the fiber nodes and L2/L3 connection to hub/headend.
  - Similar to the previous scenario but modified through L2 DOCSIS technology virtualized (NFV) and again centralized in the hub/headend.
  - FTTH/PON services with the OLT being part of the CCAP device.

62 ONF - TS-027 - L4-L7 Service Function Chaining Solution Architecture.

63 IEEE 2013 - STEERING: A Software-Defined Networking for Inline Service Chaining.

64 ODL and Defense4All use case

([https://nexus.opendaylight.org/content/sites/site/org.opendaylight.docs/master/userguide/manuals/userguide/bk-user-guide/content/\\_defense4all\\_overview.html](https://nexus.opendaylight.org/content/sites/site/org.opendaylight.docs/master/userguide/manuals/userguide/bk-user-guide/content/_defense4all_overview.html)).

65 Princeton University proposed use case (<http://sdx.cs.princeton.edu/>). Code available for ONF open source community ([https://www.opennetworking.org/?p=2069&option=com\\_wordpress&Itemid=442](https://www.opennetworking.org/?p=2069&option=com_wordpress&Itemid=442)).

66 An example can be found at OPNFV Brahma Putra - <https://www.opnfv.org/brahmaputra>.

67 ETSI Network Functions Virtualisation (NFV) Management and Orchestration - ETSI GS NFV-MAN 001 V1.1.1 (2014-12).

68 CableLabs - PacketCable Multimedia Specification.

- Others

As a continuation, some of the most common use cases of WAN controllers are described, since they are the basis for subsequent use cases.

### **3.4.1. Topology Discovery and Maintenance**

This functionality is considered critical since topological information is key to develop any following use cases of network programmability, therefore almost all of them rely on its existence.

Given that this programmability should be end-to-end, between any pair of sources and destinations, through all layers of the ISO/OSI model within the network in question, this topology feature should be available in the various layers mentioned (photonic to L3).

It includes the initial discovery of the network and its elements, interconnections in the ISO/OSI layer of relevance and updates based on events and changes that occur in them. To do this, it relies on D-CPI interfaces described in section 3.1.1.

This information should be persistent.

### **3.4.2. Path Provisioning**

This functionality follows the same terms as the previous one, since it allows the programming of end-to-end routes and paths in the network through multiple ISO/OSI layers to achieve the desired behavior. Again, it is essential for any later use case. Unavoidably its operation is based on a previous topology discovery and maintenance. To do this, it consumes D-CPI interfaces described in Section 3.1.1.

It should be an intrinsic controller functionality since, in its most basic vision, it involves translating the applications requirements into rules and instructions on devices and network elements.

### **3.4.3. Multi-Layer Optimization**

Multi-layer optimization is based on the two previous cases. In the context of this paper, this functionality includes (but is not limited to) the following concepts:

- Improve network utilization moving traffic to paths with available resources (defrag, re-grooming, bin-packing, etc.).
- Analyze multi-layered integrated redundancies minimizing resources in the connection between IP/MPLS and optical devices, etc.
- Identify the best design or multi-layer topology based not only on technical attributes (availability, latency), but also business (economic costs), etc.

It can also include capacity planning (projection of expected growth based on new products, customers, portfolio, etc.) and fault/maintenance events simulation ("*what if*", restoration of optimal paths, etc.) scenarios recently mentioned in the list of remarkable applications and use cases.

Some of these options are currently available as part of capacity planning tools. However, in these tools the work is done offline with a post-implementation of its results and very few do it with a multi-layer vision. Additionally, they have to face the difficulty and complexity in their set-up and maintenance.

Thus, the benefit of a WAN controller is to perform optimizations on the network at the same time they are identified. This gives greater granularity to schedule or apply changes based on time ranges, a single centralized multi-layer view and, finally, provides a certain dynamism to offer capacity and services that would otherwise not be possible.

In the case of the first option, it could be decomposed, according to each particular case, founded on the specific features that allow such optimization, from relatively trivial options such as modifying IGP costs to more complex as TE (RSVP and/or SR), CDC/GMPLS on the photonic layer, etc.

In the second option, this single, centralized and multi-layer control enables to optimize additional attributes that today in many cases are not possible or, in case they are, its application is limited. For example, optimization of multi-layered integrated redundancies from the photonic layer to L3 (e.g., MLR-O/MLR-P and even more ambitious options) and/or ports considering the deployment of optical transponders on routers. Presently, these examples require distributed protocols, such as GMPLS-UNI, to coordinate tasks of diverse elements and show interoperability limitations between different vendors' versions.

In the third option, the multi-layer optimization is a fundamental use case in network hierarchies which are part of the scope of this paper and should be covered by WAN controllers. This option allows investment optimizations (links, devices or network elements, ports, etc.) based on a holistic approach through the different ISO/OSI layers<sup>69</sup> with reference to the target network availability and respective SLA. Hence the importance of its inclusion in the analysis.

Finally, it is remarkable that the multi-layer optimization could include CMTS/CCAP devices, considering that they perform, or can do, aggregation and transport functions based on specified technologies (IP, MPLS, etc.).

### 3.5. Security

Security is an essential concept in any architecture and should not be underestimated, as it can jeopardize the architecture's behavior<sup>70</sup> and even its availability.

This is aggravated in the case of the SDN architecture<sup>71</sup> (see Figure 1) and, particularly, with WAN controllers scenarios (given criticality of the network sector where they are working -hierarchies of aggregation/distribution/core). As mentioned, they are very sensitive to the behavior of the centralized controller, at least in its purest version without distributed control planes or with very limited ones. Additionally, the network enables programmability through new interfaces to applications or other controllers, which raises further threats even in hybrid versions of the architecture mentioned. This is so because the operation of the network is exposed to these new connections and the anomalous behavior of an application. Furthermore, the communication between the controller and devices also generate additional threats. These interfaces can also be available for communication between different domains, both technical and administrative (examples are some controllers architectures shown in section 3.1.3 and recursive ones<sup>72</sup>).

<sup>69</sup> See already cited example in The International Journal of Computer and Telecommunications Networking 2008 - Multi-layer MPLS Network Design: the Impact of Statistical Multiplexing.

<sup>70</sup> Including reliability.

<sup>71</sup> ONF – TR-502, ONF – TR-511 – Principles and Practices for Securing Software-Defined Networks, IRTF - RFC 7426, IETF – RFC 7149 - Software-Defined Networking: A Perspective from within a Service Provider Environment, ITU-T Y.3300.

<sup>72</sup> ONF – TR-502 and ONF – TR-511.



In the center of these added threats is the WAN controller. For this reason, its protection is critical and can be addressed using a multi-level approach.

A first level within this multi-level approach can be performed with different functional blocks applying traditional perimeter security best practices (including controller isolation and/or sectorization, IP address assignments, jumpboxes, DoS attacks mitigation, acceptable behavior baselines, rate-limit for certain events, firewall, IPS, etc.). These protections should not impose limitations or restrictions on the performance and scalability of the controller (see Section 3.3) in particular, and architecture in general.

Likewise, the controller should send its own security data/information as well as network status to security incidents analysis platforms<sup>73</sup>. These platforms contain databases of threats, vulnerabilities, malicious sites, etc., and allow their correlation in real-time, while also auditing accesses, behaviors and events. This could even identify risky events that are not necessarily obvious at first. In any case, once the threats are detected, the centralized controller can take action for immediate remediation.

This last point leads to a paradox and is the reason why, in some instances, it is said that SDN architectures improve network security. On the one hand, there is a controller centralized view of the entire network that performs actions immediately when there are deviations from the expected behavior (there are security use cases included in the list of examples given in Section 3.4<sup>74</sup>). On the other hand, as mentioned at the beginning of this section and throughout the paper, this architecture is very vulnerable to the behavior of this centralized controller, its interfaces and applications and creates new threats and risks not previously existing.

Another level of protection could be given consideration on how SDN architectures should be safeguarded, focusing not only on the network infrastructure itself, but also on applications given the aforementioned scenarios<sup>75 76</sup>.

Specifically regarding the security evaluation that should be made on the WAN controller, it may be appropriate that this controller gives options of additional intrinsic or built-in protections. They allow an extra level of security that should be evaluated as part of the presented framework.

Some of these built-in features should include controller protection from users and interfaces (applications, devices, elements, other controllers or domains) accesses<sup>77</sup>. Identities, credentials, authentication and authorization management should be performed with high granularity and flexibility that allows the definition of different profiles and associated permissions. In turn, connections should be encrypted<sup>78</sup>. Also, accounting is essential for the purposes of analysis of behaviors previously mentioned. In the case of interfaces, protocols described in Section 3.1 should support SSH, TLS or at least MD5 where there are no better options to encrypt connections<sup>79 80</sup> and information validation (e.g., BGP-RPKI), especially in cases of communication between different domains.

---

73 FCAPS platforms are also an option, and not mutually exclusive.

74 ODL and Defense4All use case.

75 IETF – RFC 7149 – Security Considerations and ONF – TR-511.

76 <https://wiki.onosproject.org/display/ONOS/Security-Mode+ONOS>

77 Several options, e.g. AAA (IETF – RFC 2903, 2904, 2905, 2906, 2989, 3127, 3539).

78 Ideally, PKI architectures.

79 It is required that the other architecture elements support them.

80 Security definitions are described in their respective documents and specifications. Additionally, it is essential to consider the design of secure protocols for these interfaces in the work SDOs are currently doing, as mentioned in Section 3.1. This is to avoid the same mistakes of the past in current and future protocols design (ONF – TR-511).

Other features that can be considered as an intrinsic part of the controller are OS and platform hardening (e.g., configuration best practices, disabling or removing unnecessary ports, services or files, software update for patching identified vulnerabilities only for the existing services, etc.), policies configuration (e.g., event logging<sup>81</sup>, export of flows information, sending data and information to security incidents analysis platforms, identifying patterns of access, etc.), DoS attacks mitigation or rate-limits for certain types of connections, and others. Some more advanced protections such as firewalls, IPS, validating connections against blacklists or threats databases should be analyzed carefully because they can generate the same effects they intend to prevent<sup>82</sup>.

In all cases these controller intrinsic security functionalities and features must permit their application in a scalable manner. Yet, the weighting of these elements will depend on each individual case, based on the SDN architecture as well as the specific security architecture.

### 3.6. Controller Deployment, Maintenance and Support

This is another essential aspect and perhaps one of the most strategically important, since the chosen controller option may involve a structural change in the current business model for several industry players. Even the choice of the controller can be completely conditioned by this. An example would be telecommunications service providers developing their own controllers, from the beginning or from open-source software. There may be cases where these operators develop or deploy this software and adapt, modify, configure and integrate it in their networks, being directly responsible for post-maintenance and support. It might also occur that subsequently they make the controller available to the community<sup>83</sup>.

The evaluation of this aspect is very particular to each service provider and its business model in the short and medium term.

Currently, according to the source of the software and related tasks for deployment, configuration, commissioning and integration with the existing network, support/maintenance and evolution, there are basically three controllers choices.

- **Open-source – Pure-Play.** The service provider takes care of all the previously mentioned tasks, having all the advantages and disadvantages of working with open -source software. The same applies in case of software developed internally. Among other considerations, it will need new skill sets for their engineers to take these responsibilities and new processes to ensure service availability.
- **Open-source – Turnkey.** The controller is based on open-source but has adaptations and optimizations made by vendors or integrators. They are in charge of the tasks mentioned initially with different options for service contracts. It may involve changes in the service providers' engineers skill sets and processes in a more open perspective, as some integrators offer. There are some advantages related with products based on open-source – e.g., innovation cycles, product continuity, etc., but there is also a risk of ending up with closed solutions and vendor dependency (and lock-in)<sup>84</sup>.

81 Several options, e.g., Syslog (IETF – RFC 5424).

82 Also overlapping of functions among different architectural blocks must be analyzed and the functionalities mentioned should be deployed on the most suitable sector/block.

83 Same as with RDK/CPE, for example.

84 Considering current leading vendors (February 2016) of optical and MPLS/IP network elements for transport networks having a 70% aggregate market share, slightly more than half of them offer WAN controllers based on open-source versions (primarily ODL).

- **Proprietary Options.** The controller is based on proprietary code and depends entirely on the vendors' decisions and innovation cycles. In the best-case scenario, it has interfaces to interoperate with third party applications. Therefore, some call them open, but they are not neutral to vendors (as an example, integrations must be done when adapting interfaces) and, in many cases, do not support multiple vendors' devices and network elements. It is a closed and proprietary environment having a high risk of vendor lock-in.

The SDN architecture, in the most disruptive scenario, which would be the first choice described and including network elements with thinner distributed control layers and/or white boxes, could involve a structural change in the current business model of several industry players, not only the service providers themselves.

Finally, as mentioned, the choice of the controller could be completely conditioned by this aspect and the examples considered, such as telecommunications services providers developing their own controllers from the beginning or from open-source software, taking over the entire product life cycle and possibly making it available to the community.

## 4. Conclusions

Currently SDN is having a boom in the industry. This new paradigm along with the advantages it offers and the architecture recommended by the SDOs, is applicable to multiple scenarios and use cases.

Whenever a new paradigm arises, it is important to define the aspects to bear in mind when evaluating its components to have greater chances of achieving the full advantages and benefits that give meaning to the proposed changes. These aspects constitute the framework of the new paradigm.

In this paper, a framework has been proposed to evaluate SDN WAN controllers working on hierarchies of aggregation/distribution/core of multi-service converged networks. It can be of use in individual controllers evaluations, controllers eligibility for use cases, consistent comparisons among different controllers, etc. This framework describes six main aspects that have to be considered.

- 1- Interfaces
- 2- Reliability and Redundancy
- 3- Performance and Scalability
- 4- Applications
- 5- Security
- 6- Controller Deployment, Maintenance and Support

The aforementioned are the most significant aspects of the current Cablevision Argentina framework (version 1.0) and are considered fundamental, since they affect technical and business perspectives and determine the use cases and applications that can be finally supported to make SDN promises a reality.

As a last note, this framework is not intended to be an exhaustive or final definition as further aspects may be included derived from each specific scenario under analysis while bearing in mind this framework is evolving.

## 5. Abbreviations

AAA	authentication, authorization and accounting
ACID	atomicity, consistency, isolation, durability
A-CPI	application-controller plane interface
ANIMA	autonomic networking integrated model and approach
ANIMA WG	Autonomic Networking Integrated Model and Approach Working Group – IETF
BGP	border gateway protocol
BGP-LS	border gateway protocol–link state
BGP-RPKI	border gateway protocol–resource public key infrastructure
BoD	bandwidth on demand
CAP	consistency, availability, partition tolerance
CAPEX	capital expenditures
CCAP	converged cable access platform
CDC	colorless, directionless, contentionless
CLI	command line interface
CMTS	cable modem termination system
COPS	common open policy service
CORD	central office reimaged as a datacenter
DSLAM	digital subscriber line access multiplexer
D-CPI	data-controller plane interface
DDoS	distributed denial of service
DOCSIS	Data-Over-Cable Service Interface Specifications
DoS	denial of service
ETSI	European Telecommunications Standards Institute
FCAPS	fault, configuration, accounting, performance, security
FTTH	fiber to the home
GMPLS	generalized multi-protocol label switching
GMPLS-UNI	general multi-protocol label switching-user network interface
HFC	hybrid fiber-coax
IaaS	infrastructure as a service
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IGP	interior gateway protocol
IPS	intrusion prevention system
IRTF	Internet Research Task Force
ISBE	International Society of Broadband Experts
ISO/OSI	International Organization for Standardization / Open Systems Interconnection
ITU-T	International Telecommunication Union - Telecommunication Standardization Sector
I2RS	interface to the routing system
JSON	Javascript object notation
LISP	locator/id separation protocol

LLDP	link layer discovery protocol
LSP	label switched path
LTE	long term evolution
L0	layer 0 – photonic layer
L1	layer 1 – physical layer ISO/OSI model
L2	layer 2 – link layer ISO/OSI model
L3	layer 3 – network layer ISO/OSI model
MDSO	multi-domain service orchestration
MD5	message-digest algorithm 5
MLR-P	multi layer restoration - port
MLR-O	multi layer restoration - optical
MP-BGP	multi-protocol border gateway protocol
MPLS	multi-protocol label switching
MSO	multiple system operators
NBI	north bound interface
NBI WG	North Bound Interface Working Group - ONF
NFV	network function virtualization
NGFW	next generation firewall
NETCONF	network configuration protocol
NETCONF WG	Network Configuration Working Group – IETF
NETMOD WG	Network Configuration Data Modeling Language Working Group – IETF
ODL	OpenDaylight
OF-CONFIG	openflow management and configuration protocol
OF-PI	protocol independent layer
OLT	optical line termination
ONOS	open network operating system
ONF	Open Networking Foundation
OPEX	operational expenditures
OPNFV	open platform for network functions virtualization
PCC	path computation client
PCE	path computation element
PCE WG	Path Computation Element Working Group – IETF
PCEP	path computation element protocol
PCMM	PacketCable Multimedia
PCRF	policy and charging rules function
PEP	policy enforcement point
PKI	public key infrastructure
PS	policy server
POF	protocol oblivious forwarding
POF-FIS	protocol oblivious forwarding-flow instruction set
PON	passive optical networks
P4	programming protocol-independent packet processors
QoS	quality of service
RDK/CPE	reference design kit/customer premises equipment
REST	representational state transfer
RFC	request for comments



RSVP	resource reservation protocol
RU	rack unit
SCTE	Society of Cable Telecommunications Engineers
SDN	software defined networks
SDO	standard development organization
SFC	service function chaining
SNMP	simple network management protocol
SPRING	source packet routing in networking
SR	segment routing
SSH	secure shell
STEERING	SDN inline services and forwarding
TE	traffic engineering
TL1	transaction language 1
TLS	transport layer security
TTP	table type patterns
WAN	wide area network
WG	working group
WIMAX	worldwide interoperability for microwave access
XML	extensible markup language
XMPP	extensible messaging and presence protocol
YANG	yet another next generation
3G	third generation mobile

## 6. Bibliography and References

ACM SIGCOMM Computer Communication Review SIGCOMM'15 2015, Volume 45 - *Multi-Domain Service Orchestration Over Networks and Clouds: A Unified Approach*. Balázs Sonkoly, János Czentye, Robert Szabo, Dávid Jocha, János Elek, Sahel Sahhaf, Wouter Tavernier, Fulvio Rizzo - Association for Computing Machinery - Special Interest Group on Data Communications

*draft-bhuvan-bmwg-sdn-controller-benchmark-meth-01 - Benchmarking Methodology for SDN Controller Performance*. Internet Engineering Task Force

*draft-bhuvan-bmwg-sdn-controller-benchmark-term-01 - Terminology for Benchmarking SDN Controller Performance*. Internet Engineering Task Force

*draft-dhodylee-pce-stateful-hpce-00 - Hierarchical Stateful Path Computation Element (PCE)*. Internet Engineering Task Force

*draft-chen-pce-h-sdns-00 - PCE Hierarchical SDNs*. Internet Engineering Task Force

*draft-ietf-idr-bgpls-segment-routing-epe-05 - Segment Routing BGP Egress Peer Engineering BGP-LS Extensions*. Internet Engineering Task Force

*draft-ietf-netconf-restconf-13 - RESTCONF Protocol*. Internet Engineering Task Force

*draft-ietf-pce-segment-routing-07 - PCEP Extensions for Segment Routing*. Internet Engineering Task Force

*draft-irtf-sdnrg-layered-sdn-00 - Cooperating Layered Architecture for SDN*. Internet Research Task Force

GS NFV-MAN 001 V1.1.1 (2014-12) - *ETSI Network Functions Virtualisation (NFV) ; Management and Orchestration* - European Telecommunications Standards Institute

<http://onosproject.org/> - ONOS Controller

<https://wiki.onosproject.org/display/ONOS/Security-Mode+ONOS> – ONOS Controller

[https://wiki.opendaylight.org/view/ODL-SDNi\\_App:Main](https://wiki.opendaylight.org/view/ODL-SDNi_App:Main) - OpenDaylight Controller

<https://www.opendaylight.org/> - OpenDaylight Controller

<https://www.opennetworking.org/> - Open Networking Foundation

<https://www.opennetworking.org/protocol-independent-forwarding> - Open Networking Foundation

<https://www.opnfv.org/> - Open Platform for Network Functions Virtualization

<https://www.opnfv.org/brahmaputra> - Open Platform for Network Functions Virtualization

<http://www.poforwarding.org/> - POF

<http://www.poforwarding.org/documents/> - POF

<http://p4.org/> - P4 Language Consortium

<http://p4.org/p4/clarifying-the-differences-between-p4-and-openflow/> - P4 Language Consortium

<http://sdx.cs.princeton.edu/> - Princeton University

*Northbound Interfaces Charter* - Open Networking Foundation

*PacketCable Multimedia Specification*. CableLabs

*The P4 Language Specification Version 1.1.0* - P4 Language Consortium -

RFC 3535 - *Overview of the 2002 IAB Network Management Workshop*. Internet Engineering Task Force

RFC 6020 - *YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)*. Internet Engineering Task Force

RFC 6805 - *The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS*. Internet Engineering Task Force

RFC 7149 - *Software-Defined Networking: A Perspective from within a Service Provider Environment*. Internet Engineering Task Force

RFC 7426 - *SDN: Layers and Architecture Terminology*. Internet Research Task Force

SPRING - <https://tools.ietf.org/wg/spring/> - Internet Engineering Task Force

*StEERING: A Software-Defined Networking for Inline Service Chaining*. 2013. Institute of Electrical and Electronics Engineers

The International Journal of Computer and Telecommunications Networking 2008, Volume 52 – *Multi-layer MPLS Network Design: the Impact of Statistical Multiplexing*. Pietro Belotti, Antonio Capone, Giuliana Carello, Federico Malucelli - Computer Networks

TR-502 - *SDN Architecture 1.0* – Open Networking Foundation

TR-505 – *OF-PI: A Protocol Independent Layer and ONF Protocol Independent Forwarding* – Open Networking Foundation

TR-511 – *Principles and Practices for Securing Software-Defined Networks* – Open Networking Foundation

TS-027 - *L4-L7 Service Function Chaining Solution Architecture* - Open Networking Foundation

Y.3300 - *Framework of SDN* - International Telecommunication Union - Telecommunication Standardization Sector

## Deploying Holistic Wi-Fi

### Using Self-Optimizing-Network Techniques to Provide A 360-degree View of Wi-Fi QoE

A Technical Paper prepared for SCTE/ISBE by

Patricio S. Latini, SVP Engineering, Intraway, SCTE/ISBE Member  
Campillo 2541  
Buenos Aires (1427)  
Argentina  
patricio.latini@intraway.com  
+1 (305) 504-9250

## Table of Contents

Title	Page Number
Table of Contents	154
1. Introduction	156
2. Current Situation	156
2.1. Wi-Fi is Booming	156
2.1.1. 2.4 GHz Wi-Fi Spectrum	157
2.1.2. 5 GHz Wi-Fi Spectrum	159
2.1.3. Other Interference Sources	162
3. The Wi-Fi 360 QoE	162
3.1. Self Optimizing Network	163
3.1.1. Gateways data access	164
3.1.2. SON Operations examples	164
3.2. Installation and Service Wizard	166
3.2.1. Technician with Free-Apps running on laptops or smartphones	166
3.2.2. Technician with Hardware solutions	167
3.2.3. A Third Way – A Specialized App in the Technician Device	167
3.3. Self-Management	173
3.4. Call Center Toolbox	175
4. Conclusions	177
5. Abbreviations	177
6. Bibliography and References	177

## List of Figures

Title	Page Number
Figure 1 - 2.4 GHz – 20 MHz Non-Overlapping Channels	157
Figure 2 - 2.4 GHz – 20 MHz Non-Overlapping Channels	157
Figure 3 - Throughput vs SNR for 802.11n Channels with 3 Spatial Streams	158
Figure 4 - Throughput vs SNR for 802.11n Channels with 3 Spatial Streams	159
Figure 5 - 5 GHz – Available Channels	160
Figure 6 - Throughput vs SNR for 802.11ac Channels with 3 Spatial Streams	161
Figure 7 - Wi-Fi 360 Architecture	163
Figure 8 - Non-Optimized Neighbor AP Setup	164
Figure 9 - Optimized Neighbor AP Setup	165
Figure 10 - Non-Optimized Neighbor AP Setup	165
Figure 11 - Optimized Neighbor AP Setup	166
Figure 12 - Technician Free-App Capture	167
Figure 13 - Hardware Solutions	167
Figure 14 - Installation and Service Wizard Flow	168



Figure 15 - Installation and Service Wizard Cheat Sheet	169
Figure 16 - Wi-Fi Installation Mobile App	170
Figure 17 - Wi-Fi Installation Mobile App	171
Figure 18 - Wi-Fi Certification Report	172
Figure 19 - Birth Certificates History	173
Figure 20 - Self Management Flow	174
Figure 21 - Wi-Fi Installation Mobile App	174
Figure 22 - Statistics of Calls Reduction after deploying an MSO App	175
Figure 23 - CSR Management Flow	176
Figure 24 - CSR Troubleshoot Portal	176

### List of Tables

<b>Title</b>	<b>Page Number</b>
Table 1 - Attenuations for Common Materials	159
Table 2 - 802.11ac Available Non-Overlapping Channels	161
Table 3 - Non-Wi-Fi Typical Interference Sources	162

## 1. Introduction

In the last few years most of the cable industry has been deploying Wi-Fi gateways on a large scale. But most of the monitoring and troubleshooting has been on traditional Quality of Service measurements, not on the customers' actual quality of experience (QoE).

This paper analyzes some of the lessons learned from Wi-Fi residential deployments done by cable operator and proposes a set of new tools to have both real time monitoring and historical dashboards of key parameters of Wi-Fi services. These new tools will allow the cable companies to have a much more accurate and real-time understanding of the end-user QoE and dramatically improving their ability to react quickly to problems in the wireless loop portion of the network.

A full section is dedicated to analyzing and proposing specific solutions for the home network environment. Particular attention is paid to methods of applying self optimizing network (SON) techniques to automatically monitor, troubleshoot and predict Wi-Fi network behavior. This is especially important today, given the prevalence of Wi-Fi connected Internet protocol television (IPTV) and over-the-top (OTT) set-top boxes and the radio frequency (RF) interference that can potentially degrade service to these devices.

A new set of OSS tools will be described that give real time control of the service key performance indicators (KPI). This control will not only support the improvement of the end-user QoE but, by better understanding end user problems, particularly in the home network environment, also help reduce operational costs.

The resulting conclusions will allow the cable operator to understand how to leverage existing tools and experiences learned to enhance their Wi-Fi service deployments.

## 2. Current Situation

Today every cable operator is deploying Wi-Fi enabled gateways with DOCSIS<sup>®1</sup> 3.0 access bandwidth well above 100 megabits per second (Mbps). With DOCSIS 3.1 deployments being readied for launch in the near future, bandwidth speeds will increase to 1 gigabits per second (Gbps) level. With these developments, having automated control of the in-home wireless environment is now a critical factor for success.

According to some estimates, there are on average 13 connected devices per household in the U.S. [1]. Most of these devices utilize Wi-Fi. The growing number of Wi-Fi devices combined with a finite amount and limitation of unlicensed spectrum and the need to coexist with other technologies sharing these same bands turns the problem of assuring the Quality of Experience of the wireless loop into a serious one.

### 2.1. Wi-Fi is Booming

As mentioned above, the number of connected devices is rapidly growing year after year. At the same time customers are using Wi-Fi for bandwidth intensive activities such as streaming video, gaming and more. The signal levels required for this increase in bandwidth demand can require the installation of

---

<sup>1</sup> DOCSIS is a trademark of CableLabs

multiple access points or repeaters. Twisted-pair Ethernet and Multimedia over Coax Alliance (MoCA) technology are being used as to backhaul interconnections for all the access points in the home.

Wireless telephone providers have started to implement Wi-Fi Calling services, off-loading all the voice traffic from the mobile network to the local Wi-Fi networks. This is causing a significant increase in the minimum quality requirements of the standard Wi-Fi deployment.

In densely populated environments competition for airtime in unlicensed Wi-Fi spectrum is growing rapidly.

### 2.1.1. 2.4 GHz Wi-Fi Spectrum

The 2.4 gigahertz (GHz) band provides much better spatial coverage than the 5 GHz band, but has only three non-overlapping 20 megahertz (MHz) channels and just one when using a 40 MHz channel width (Figure 1 and Figure 2). This increases the chance of co-channel interference with neighboring access points, and with the non-Wi-Fi technologies than use the same band. This is one of the main drivers for why intelligent and flexible channel management will be a requirement for high density Wi-Fi environments. Additionally, 802.11n has implemented a technology called multiple input multiple output (MIMO) which allows an increase in the performance of a channel by n times by the addition of extra antennas in the access point and the end device. Currently the most common access points in the market support a 3 transmitters and 3 receivers (3x3) configuration however the 802.11ac standard defines the possibility to support up 8x8 configurations in the future. On the end device side newer laptops and tablets support 3x3 capabilities but smaller devices have only 1x1 support.




Figure 1 - 2.4 GHz – 20 MHz Non-Overlapping Channels




Figure 2 - 2.4 GHz – 20 MHz Non-Overlapping Channels

Typically, a 20 MHz channel is used in the 2.4 GHz band, which under optimal conditions (e.g., very close to the AP), yields a maximum raw capacity of 450 Mbps for a 3x3:3 access point (Figure 3). However, that result becomes affected by several factors such distance, obstacles, remote device sensitivity, transmission power and MIMO antenna capabilities. All of these factors can result in a worst-

case performance that is reflected in a poor received signal strength indication (RSSI) that ends up using a very strong modulation and coding scheme (MCS), these transmission formats are denominated and are defined in the different 802.11 standards. In general, the strongest MCS is the index 0 which for a 1x1 device like a smartphone, and throughput as low as 7 Mbps of raw capacity or about 3.5 Mbps of IP bandwidth. This means that single device with a poor RSSI trying to watch a streaming video can consume *all* of the available airtime for a local 2.4 GHz spectrum. (Figure 4).





Figure 3 - Throughput vs SNR for 802.11n Channels with 3 Spatial Streams



**Figure 4 - Throughput vs SNR for 802.11n Channels with 3 Spatial Streams**

**2.1.2. 5 GHz Wi-Fi Spectrum**

Using the following formula, it can be shown that 5.8 GHz has 7.66 decibels (dB) greater free space attenuation than 2.4 GHz assuming the same distance:

$$\text{free space attenuation of 5.8 GHz versus 2.4 GHz} = 20 * \log(5.8/2.4) = 7.66 \text{ dB}$$

Furthermore, the solid material attenuation of the 5 GHz band is also higher than 2.4 GHz as seen in Table 1.

**Table 1 - Attenuations for Common Materials**

Building Material	2.4 GHz Attenuation	5 GHz Attenuation
Solid Wood Door 1.75"	6 dB	10 dB
Brick Wall 3.5"	6 dB	10 dB
Concrete Wall 18"	18 dB	30 dB
Cubical Wall 2.25"	18 dB	30 dB
Interior Glass Window 1"	3 dB	6 dB
Double Glass Window 1"	13 dB	20 dB
Marble 2"	6 dB	10 dB



These two considerations cause 802.11ac devices to operate in the 5 GHz band at short distances from the access point and then switch to 2.4 GHz as the distance increases.

Another important factor is that the 5 GHz band more available spectrum. This allows for the operation of 20, 40, 80 and even 160 MHz channels as seen in Figure 5. The number of available channels, and whether dynamic frequency selection (DFS) is required is shown in Table 2.

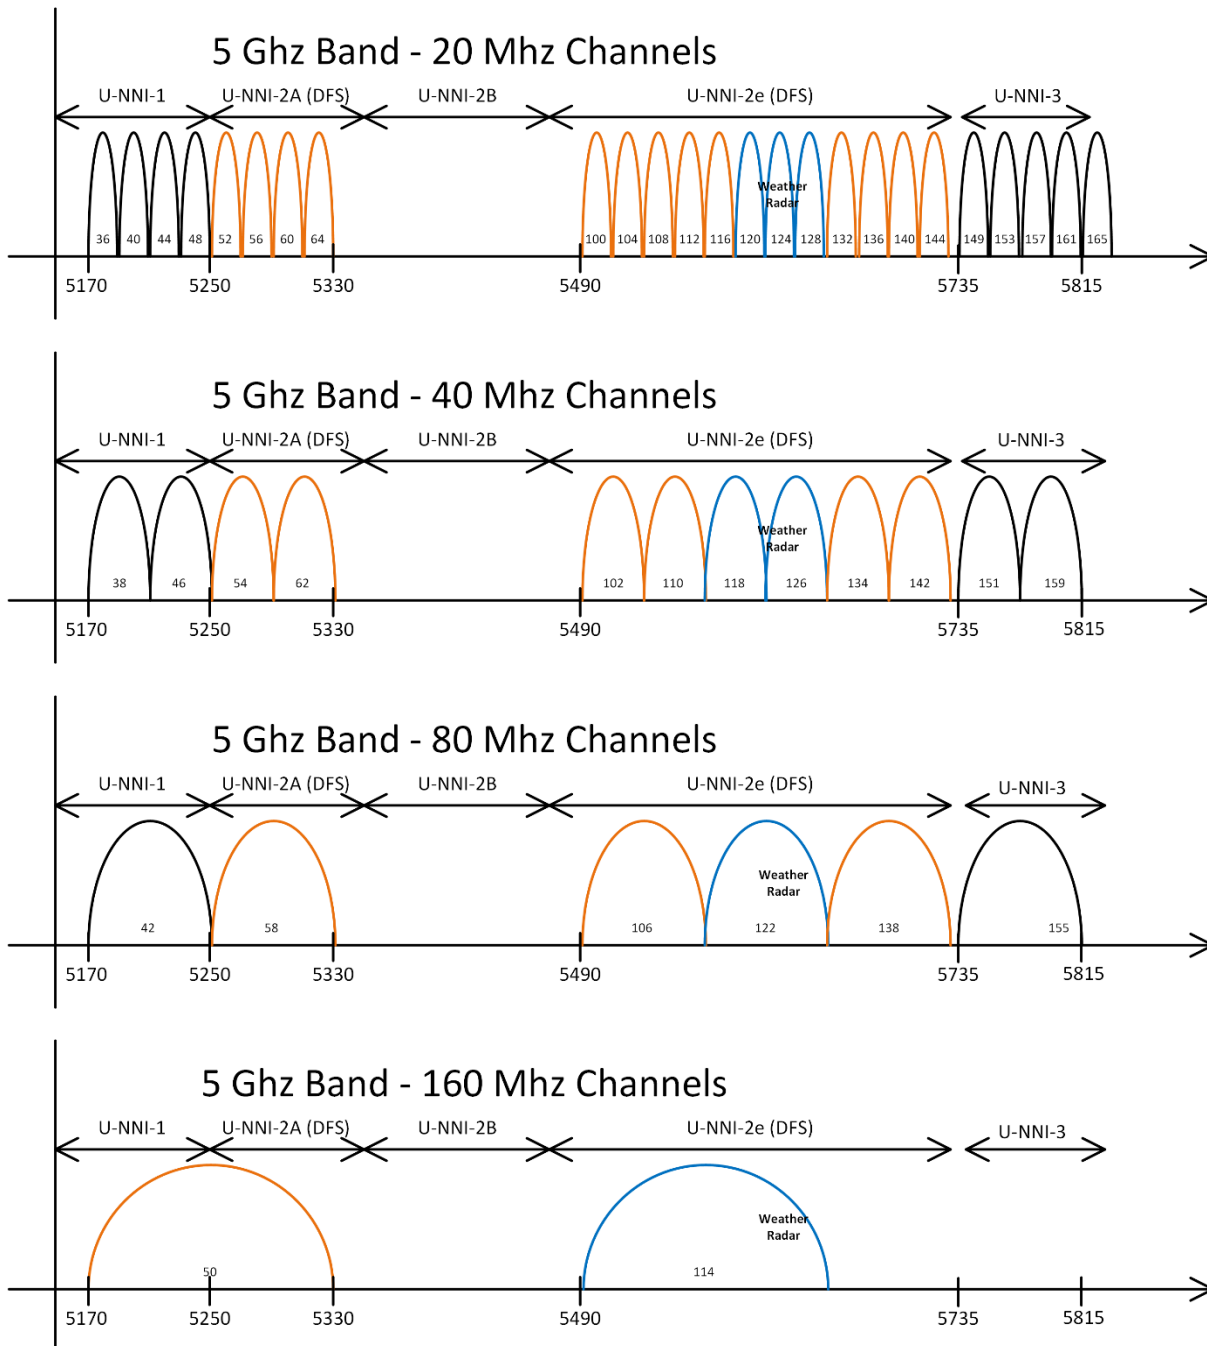



Figure 5 - 5 GHz – Available Channels

**Table 2 - 802.11ac Available Non-Overlapping Channels**

	<b>DFS</b>	<b>Non DFS</b>
<b>20 MHz</b>	22	9
<b>40 MHz</b>	10	4
<b>80 MHz</b>	5	2
<b>160 MHz</b>	1	0

From Figure 6, it is evident how much more capacity is available for 802.11ac-capable devices. A 3x3:3 Access point with an 80 MHz Channel width, for example, can provide up to 1300 Mbps of raw bandwidth! However, it is important to mention that even in the 5 GHz band, the scarcity of non-overlapping channels is still present. There are only two channels available for zones with DFS detect and only five where radar-like signals will definitely not be present. So again, intelligent and flexible channel management will be a requirement for dense Wi-Fi environments.



**Figure 6 - Throughput vs SNR for 802.11ac Channels with 3 Spatial Streams**

### 2.1.3. Other Interference Sources

Interference from non-Wi-Fi sources is an important consideration. A list of potential interference sources is listed in Table 3. These interferers are much harder to mitigate as they don't have the "listen before talk" media contention format that Wi-Fi has. Additionally, radar sources can trigger the DFS detection algorithm, forcing the access point to immediately stop transmission on those channels

**Table 3 - Non-Wi-Fi Typical Interference Sources**

2.4 GHz	5 GHz
Bluetooth	Surveillance cameras
Smart Home (Zigbee/Zwave)	Radar (DFS)
Cordless phones	Microwave backhaul links
Microwave ovens	Weather stations
2-way radios	

## 3. The Wi-Fi 360 QoE

When considering the 'bar-room brawl' landscape described previously, it quickly becomes obvious that an automated management-system that could remotely adjust the radio configuration of each access point would significantly improve the health of the whole network. This type of approach is called a self-optimizing network or "SON".

But it is possible to go further.

Once the required components for deploying a SON architecture are in place, a complementary architecture can be overlaid on it to provide a more holistic management of the network. This approach is called "Wi-Fi 360 QoE Assurance" and involves the following integrated components.

- 1 – Self-Optimizing Network
- 2 – Installation and Service Wizard
- 3 – Self-Management App
- 4 – Call Center Toolbox

The software architecture of a Wi-Fi 360-enabled backend (Figure 7) includes the components described in the CableLabs RRM/SON Proposal, such as the radio resource management (RRM) Engine and a remote-management gateway (an SNMP gateway and/or TR069 ACS<sup>1</sup>). This architecture adds a cloud web-Services front-end which will be described later when the new functions are presented in point 3.2.2.<sup>2</sup>

<sup>2</sup>Given the still lacking standardization of the managing protocol of Wi-Fi enabled gateways in the cable industry

## Wi-Fi 360 Architecture

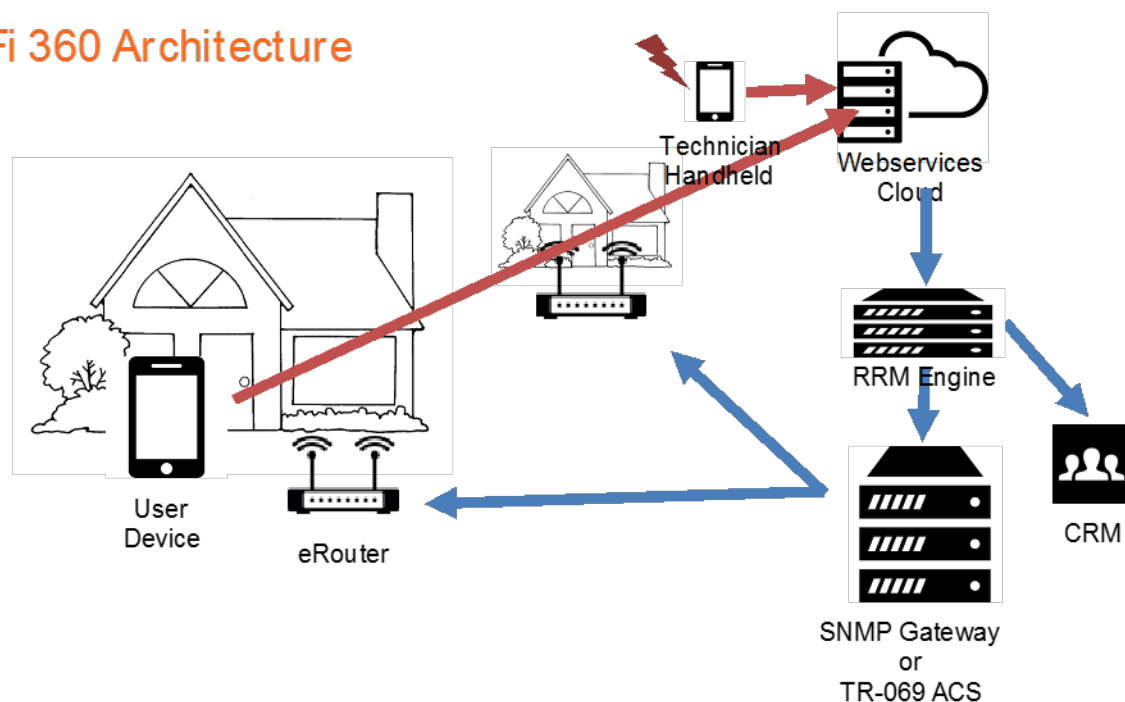


Figure 7 - Wi-Fi 360 Architecture

### 3.1. Self Optimizing Network

A SON architecture ensures the best *network* performance by addressing changes in the environment in an automated manner. It does this by managing the parameters of each controlled Wi-Fi gateway. Each gateways' parameters can be adjusted to provide a better overall network operation. The adjustable parameters are:

- Transmit power
- MCS rates
- MIMO and MU-MIMO configurations
- Beam forming configurations
- Channel bandwidth
- Maximum throughput per device
- Carrier sense thresholds
- Multi band configuration and steering of devices
- Channel assignments among others.

At the same time some of the useful KPIs that can be retrieved from the access points include:

- Received signal strength from devices
- Noise and interference levels
- Packet error rates and packet loss rates
- Throughput of uplink and downlink per device
- Device location

- Load threshold indicators
- Channel utilization
- Band utilization
- Rogue AP detection
- Neighbor AP detection
- Delays, latencies and jitter for traffic uplink and downlink

### 3.1.1. Gateways data access

Cable operators have deployed Wi-Fi gateways that can provide the wireless information and control objects in three different manners.

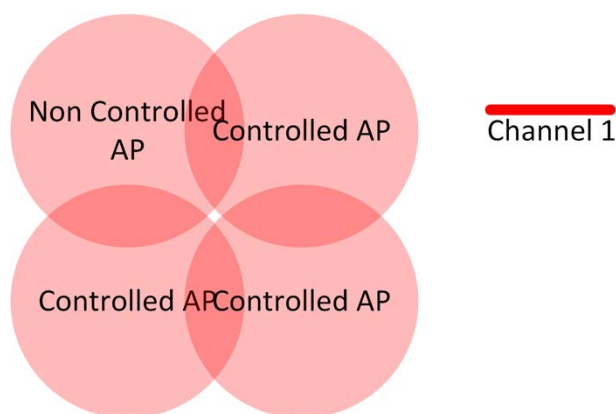
- SNMP Proprietary MIBs
- SNMP CableLabs-Standard Defined MIBs
- TR-069 access using the TR-181 Data-model

As Gateways of the three different types typically may coexist in a network, the RRM engine should support all the method described above.

### 3.1.2. SON Operations examples

#### 3.1.2.1. Channel Selection

Except for login and SSID names, it is common for newly deployed Wi-Fi access points to be left in the default configuration. In this situation, the channel selection algorithm is performed during the installation, and will not be performed until the next reboot. This often results in the access point staying in the same channel until that reboot, which might be months or years in coming. The problem is that while the algorithm may detect interference from neighboring access points, it can also choose the same channel if they are not transmitting at the moment that the access point boots, such as in Figure 8.



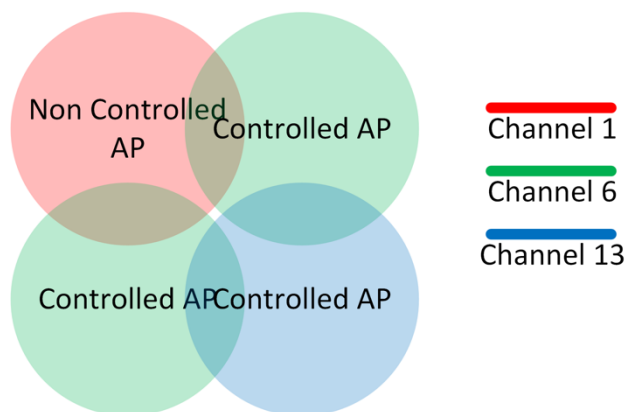
**Figure 8 - Non-Optimized Neighbor AP Setup**

Most Wi-Fi gateways are first turned up during one of the lower-usage points of the day (i.e., ‘normal working hours’). This means that the channel selection algorithm does *not* see the neighboring access point interference that can be prevalent during the actual high-usage times (e.g., evening hours).



The result is that it can be very common for *all* neighboring Wi-Fi devices to be competing for airtime on the same channel (channel 1, for example).

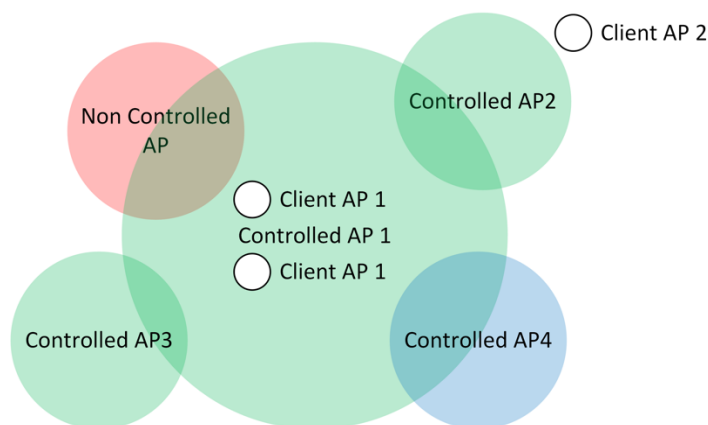
One very effective SON operation would be to periodically scan (i.e. one hour periods) all the APs available on the network and cross-reference the Neighbor AP list. The RRM engine would detect which access points are controllable. Then, using the non-controllable access points as channel constraints, create the most efficient channel distribution to avoid interferences (See Figure 9).



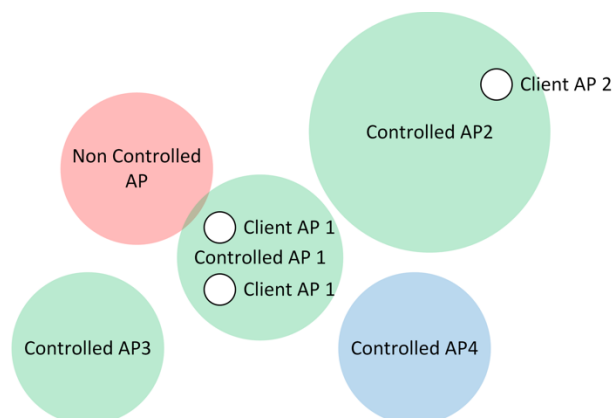
**Figure 9 - Optimized Neighbor AP Setup**

### 3.1.2.2. Power Adjustment

Another common case is where access points are configured for full-power transmission by default. But full power is often not being needed, especially as the client devices might be located close to the transmitter (See Figure 10). In those cases, the RRM can reduce the transmit power of access point 1 in order to avoid interference with access point 2. At the same time, it can increase the power of access point 2 to better serve a borderline client on that access point as shown in Figure 11.



**Figure 10 - Non-Optimized Neighbor AP Setup**



**Figure 11 - Optimized Neighbor AP Setup**

### 3.2. Installation and Service Wizard

Cable operators have done an excellent job with their wireline installations. Their standard practices and tests ensure that the cable-RF or optical signals at the customer premise meet well defined standards. However, now a paradigm change is required. Most end-users no longer use the Ethernet ports. Instead, they use Wi-Fi as the main way of accessing the broadband service.

There are several proposed and/or implemented plans to incorporate Wi-Fi measurement tools in the installation or service process. They can be separated in two groups:

#### **3.2.1. Technician with Free-Apps running on laptops or smartphones**

Advantages

- Cost – Free or a couple of dollars per technician
- Training – Basic training required
- Quick Deploy – Just install the App in the Technician Field Device

Disadvantages

- Overly Simplistic – Usually just a simple signal level test
- Data – No centralized data or birth certificate records
- Support – No support and lack of new features




Figure 12 - Technician Free-App Capture

### 3.2.2. Technician with Hardware solutions

#### Advantages

- Extensive Testing – Very detailed Wireless RF Environment Measurements
- Support – Hardware Vendors provide support for that tools

#### Disadvantages

- Expensive – Need to add another testing device in the range of thousands of dollars per technician
- Complex – Requires a lot of specialized training
- Data – Most of them still lack centralized data or birth certificate records
- Under-utilized – The full capabilities are required just 5 % of the time, and only in the most complex cases.



Figure 13 - Hardware Solutions

### 3.2.3. A Third Way – A Specialized App in the Technician Device

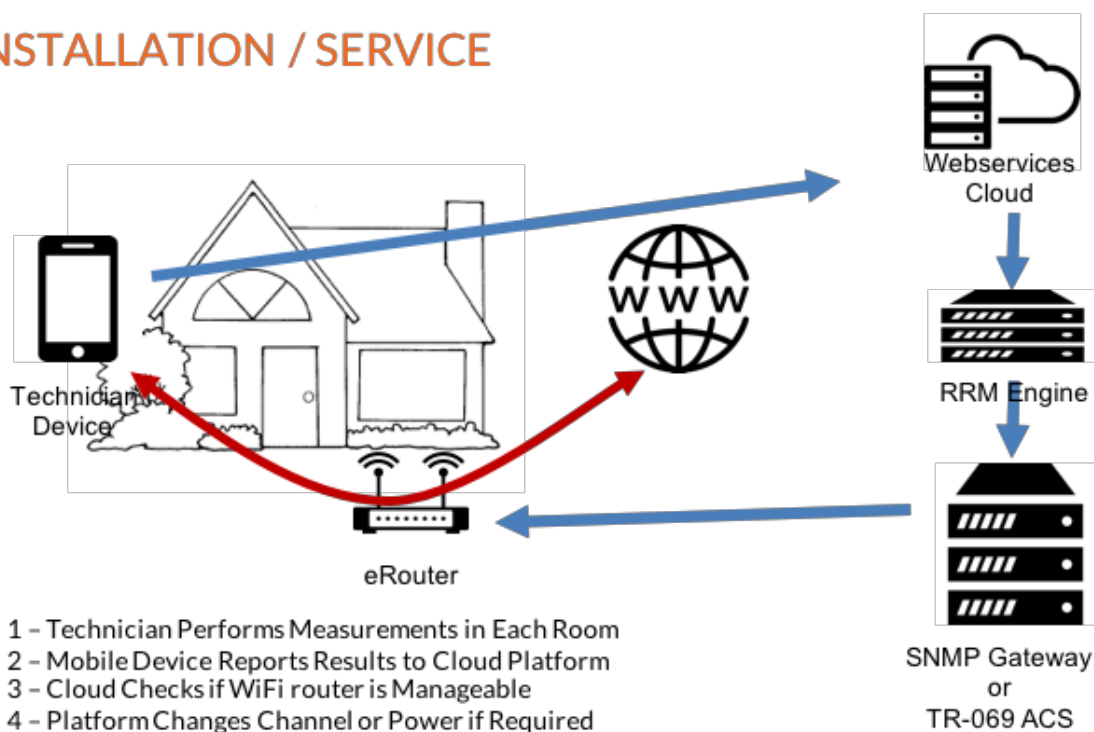
While both of previous approaches have some advantages and disadvantages, they both still miss a key aspect: the validation of the actual customer experience.

A dedicated application with programmable flows can give the advantages of both approaches while also greatly expanding its capabilities.

- Cost – No need for a specialized device, runs on existing technician devices (smartphones/tablets/PCs).
- Training – Basic Training Required, but the application flow guides the process
- Extensive Testing – The application can not only test Wi-Fi RF parameters but also emulate actual user data download and uploads
- Data – Can report all the data automatically to a centralized server

One key aspect is that this architecture can easily be integrated with an RRM engine. This allows it to reuse the backend components to perform real-time network changes such as Radio Channel, SSID Name or Encryption Type (See Figure 14).

## INSTALLATION / SERVICE



**Figure 14 - Installation and Service Wizard Flow**

With these capabilities available, the following seven-step process for good Wi-Fi installations cannot only be done, but done quickly and reliably:

- 1 – Assess
- 2 – Install
- 3 – Test and Record
- 4 – Adjust and Retest
- 5 – Review
- 6 – Certify
- 7 – Record, Store and Report

### 3.2.3.1. Assess

The technician looks over the customer's residence, identifying the type of house, layout, and how the Wi-Fi will be used. In general, technicians have to be trained on the assessment process. It is also recommended to provide a “cheat sheet” similar to the one show in Figure 15.

Basic Wi-Fi Installation: A Tech's "Cheat Sheet"	Common Wi-Fi Issues																								
<ol style="list-style-type: none"> <li>1) <b>Assess the building</b> <ol style="list-style-type: none"> <li>a) Where is the cable feed?</li> <li>b) What kind of home construction? Age? Layout?</li> <li>c) How close are other houses?</li> <li>d) External RF sources?</li> </ol> </li> <li>2) <b>Assess the Customer's Usage Pattern</b> <ol style="list-style-type: none"> <li>a) What kind of applications will be used?</li> <li>b) What rooms will have what applications?</li> <li>c) How many devices and what kinds?</li> <li>d) What does the customer anticipate in the next 3-5 years?</li> </ol> </li> <li>3) <b>Determine the Wi-Fi hub location using the "POACH" method:</b> <ol style="list-style-type: none"> <li>a) Power – how close and accessible?</li> <li>b) Obstructions – what items or materials could block or weaken the Wi-Fi signal?</li> <li>c) Access – how easy to pull cable/fiber there?</li> <li>d) Centralized – is it centered in the area?</li> <li>e) Height – Five to seven feet off the floor?</li> </ol> </li> <li>4) <b>Test!</b> <ol style="list-style-type: none"> <li>a) Test at the Cable Modem with Ethernet</li> <li>b) Check Wi-Fi SSID 'sniffer map'. Select channel with the least interference.</li> <li>c) Perform bandwidth tests in each room. Record the results.</li> </ol> </li> <li>5) <b>Confer and Adjust</b> <ol style="list-style-type: none"> <li>a) If needed, adjust channel settings and/or location to improve performance.</li> <li>b) Show the results to the customer.</li> <li>c) If needed, <b>upsell</b> Wi-Fi extension options, including: Moca, Wireless Repeaters, and Ethernet runs.</li> </ol> </li> <li>6) <b>Acceptance</b> <ol style="list-style-type: none"> <li>a) Get customer acceptance of system</li> <li>b) Issue 'Birth Certificate' for Wi-Fi set up</li> <li>c) Explain what can change the Wi-Fi results</li> <li>d) Provide the customer with a support number for questions.</li> </ol> </li> </ol>	<div style="background-color: #d4edda; padding: 5px; border: 1px solid #c3e6cb; margin-bottom: 5px;"> <p style="text-align: center; margin: 0;"><b>Traffic Congestion/Contention</b></p> <ul style="list-style-type: none"> <li>OTT Video Streaming on &lt;25Mbps</li> <li>Multiple simultaneous OTT Video streams on &lt;50Mbps</li> <li>Online Gaming Computers/ Consoles on &lt;25Mbps</li> <li>Remote Office/Video Conferencing &lt;50Mbps</li> <li>Combinations of the above (&lt;100Mbps)</li> </ul> </div> <div style="background-color: #fff3cd; padding: 5px; border: 1px solid #ffee58; margin-bottom: 5px;"> <p style="text-align: center; margin: 0;"><b>Software/Hardware Problems</b></p> <ul style="list-style-type: none"> <li>Old Wi-Fi equipment may not support 802.11n (pre-2010) or 802.11ac (pre-2014).</li> <li>Old software or firmware in Wi-Fi device.</li> <li>Positioning of device next to or behind material with high RF impact.</li> <li>Positioning of device next to or behind other RF device.</li> <li>Android devices auto-switch between 2.4GHz &amp; 5GHz networks with same SSID.</li> </ul> </div> <div style="display: flex; justify-content: space-between;"> <div style="background-color: #d1ecf1; padding: 5px; border: 1px solid #bee5eb; width: 45%;"> <p style="text-align: center; margin: 0;"><b>RF Interferers (2.4GHz &amp; 5GHz)</b></p> <ul style="list-style-type: none"> <li>Other Wi-Fi hubs</li> <li>2.4GHz and 5GHz Cordless Telephones</li> <li>Bluetooth devices</li> <li>Microwave Ovens (loose door)</li> <li>Satellite Dish (poorly wired)</li> <li>ZigBee devices (IoT)</li> </ul> </div> <div style="background-color: #d1ecf1; padding: 5px; border: 1px solid #bee5eb; width: 45%;"> <p style="text-align: center; margin: 0;"><b>Physical Interferers</b></p> <ul style="list-style-type: none"> <li>Large Kitchen Appliances</li> <li>Aquariums</li> <li>Brick and Plaster/Lathe Walls</li> <li>Tile/Stone surfaces</li> <li>Metal Screens</li> <li>Large Screen TVs</li> <li>Steel or Security Doors</li> </ul> </div> </div> <table border="1" style="width: 100%; border-collapse: collapse; margin-top: 5px;"> <thead> <tr style="background-color: #d1ecf1;"> <th>Material</th> <th>RF Impact</th> <th>Material</th> <th>RF Impact</th> </tr> </thead> <tbody> <tr> <td>Wood</td> <td>Low</td> <td>Marble</td> <td>Medium</td> </tr> <tr> <td>Most Synthetics</td> <td>Low</td> <td>Plaster</td> <td>High</td> </tr> <tr> <td>Glass</td> <td>Low</td> <td>Concrete</td> <td>High</td> </tr> <tr> <td>Water</td> <td>Medium</td> <td>Bulletproof Glass</td> <td>High</td> </tr> <tr> <td>Bricks</td> <td>Medium</td> <td>Metal</td> <td>Very High</td> </tr> </tbody> </table>	Material	RF Impact	Material	RF Impact	Wood	Low	Marble	Medium	Most Synthetics	Low	Plaster	High	Glass	Low	Concrete	High	Water	Medium	Bulletproof Glass	High	Bricks	Medium	Metal	Very High
Material	RF Impact	Material	RF Impact																						
Wood	Low	Marble	Medium																						
Most Synthetics	Low	Plaster	High																						
Glass	Low	Concrete	High																						
Water	Medium	Bulletproof Glass	High																						
Bricks	Medium	Metal	Very High																						

Figure 15 - Installation and Service Wizard Cheat Sheet

### 3.2.3.2. Install

The tech decides on the best location and install the Wi-Fi gateway.

### 3.2.3.3. Test and Record

The technician then uses the Wi-Fi Application on his Android/IOS phone or tablet and runs the six typical automatic standard tests in each room.

- Wi-Fi SSID Sniffer
- RSSI level
- Upstream
- Downstream Bandwidth Tests
- Latency
- Packet Loss



The complete testing in each room is automated and should last no more than 40 seconds. The test results are automatically uploaded to central server, providing real-time monitoring and permanent records of the installations.

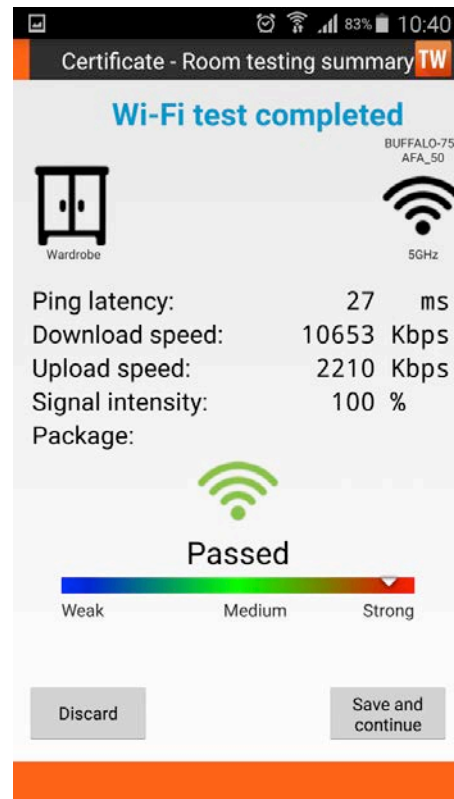


Figure 16 - Wi-Fi Installation Mobile App

#### 3.2.3.4. *Adjust and Retest*

The technician checks the results and makes any needed changes or adjustments, such as switching to a less congested Wi-Fi channel or moving the gateway to a better location, then re-tests to make sure that changes were effective.




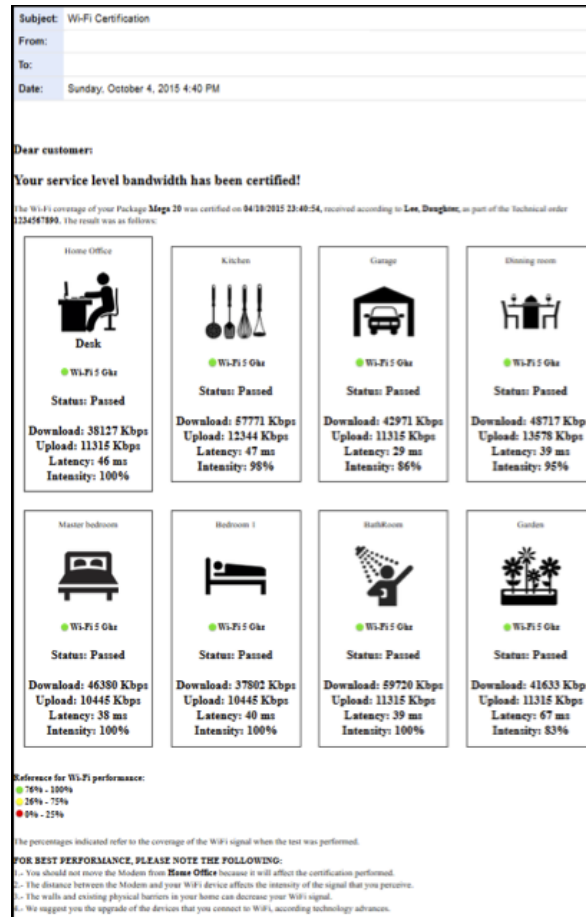
Figure 17 - Wi-Fi Installation Mobile App

**3.2.3.5. Review**

With all the rooms now passing the tests, the tech gets the customer’s approval and *certifies* the installation. If need be, this is also when the tech can recommend Wi-Fi extenders, a higher end service, or other adjustments.

**3.2.3.6. Certify**

The customer immediately receives an email with a Wi-Fi ‘Birth Certificate’, documenting the test results in each room (See Figure 18). This also furthers the customers Wi-Fi education A very high percentage of support calls are resolved by simply explaining how Wi-Fi works.



**Figure 18 - Wi-Fi Certification Report**

### 3.2.3.7. Record, Store and Report

A real-time dashboard can show the status of all certification efforts (See Figure 19). While the test results of each Wi-Fi certification effort is recorded for future reference.

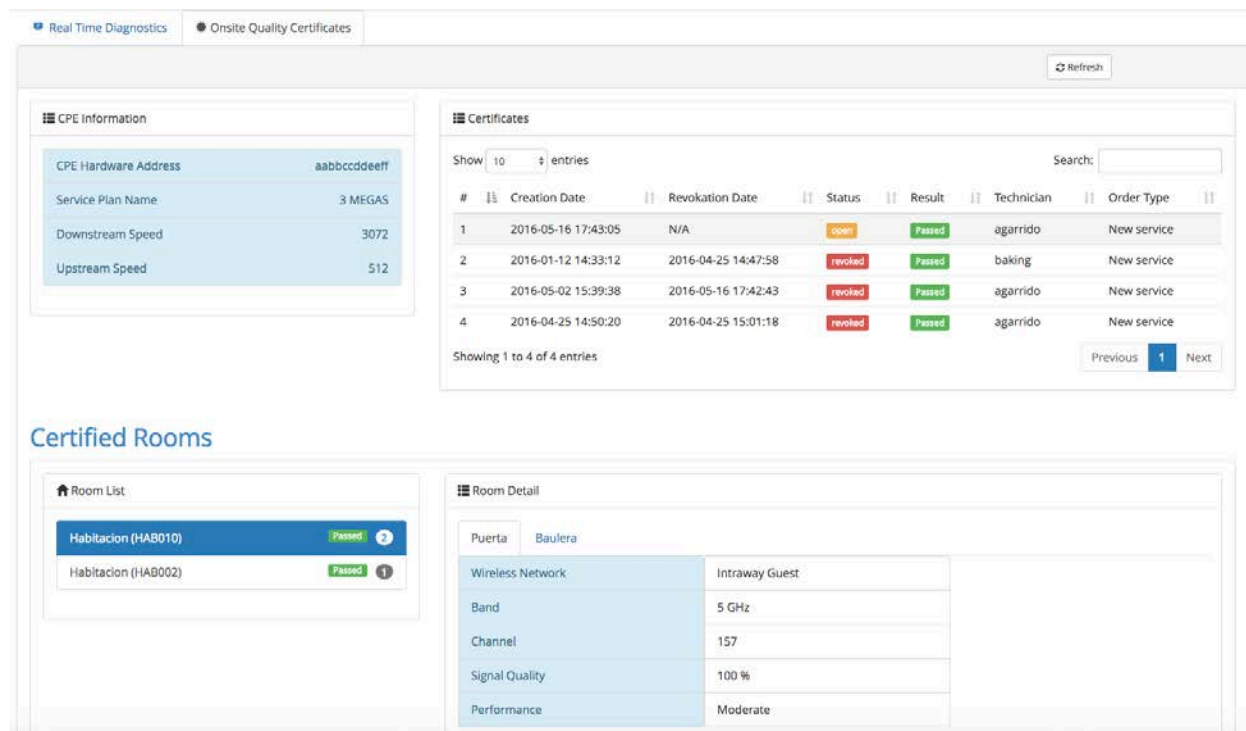


Figure 19 - Birth Certificates History

### 3.3. Self-Management

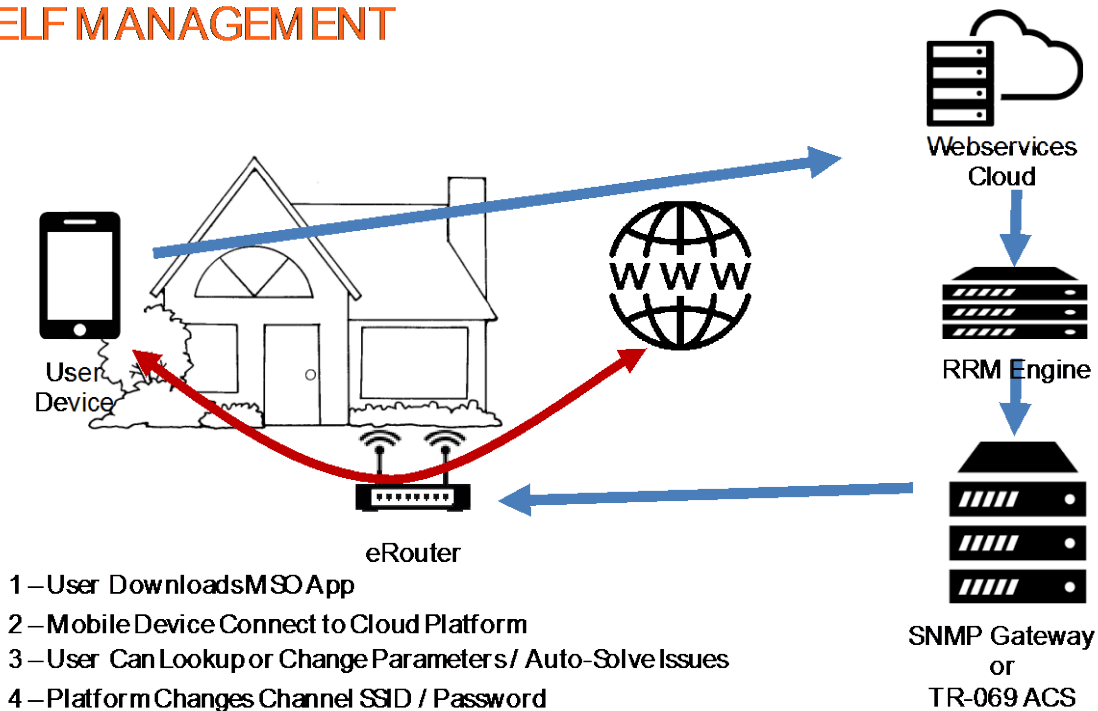
It has been widely reported that issues related to Wi-Fi are the single biggest reason an end-user call a cable operator call center. In papers published by MSOs [4] and our own empirical data we find that the two top reasons for a Wi-Fi related call are 1) educating customers on how Wi-Fi works, and 2) performing simple Wi-Fi configuration tasks, such as resetting a password or changing an SSID name.

These simple requests cost the operator in terms of call-center time and resources to help address issues that the customer could otherwise remedy themselves. There obviously needs to be a better way of handling this situation.

While section 3.2 addresses customer education, a very effective solution for reducing the number of simple configuration adjustment calls could be for the cable operator to provide the customer with a self-management application that could be downloaded to a mobile phone or tablet.

An application of that kind could easily integrate into the RRM architecture as shown in Figure 20.

## SELF MANAGEMENT



**Figure 20 - Self Management Flow**


Figure 21 is the screen capture of a replica of an Android application developed by a service provider.



**Figure 21 - Wi-Fi Installation Mobile App**



In its initial field trial deployment with a South American cable operator, call center results showed that after the deployment and implementation of this application in mid-March, there was a significant drop in three tracked call reasons (See Figure 22).



**Figure 22 - Statistics of Calls Reduction after deploying an MSO App**

From the data in Figure 22 it is easy to see that implementing such an application can result in significant cost savings at the call center level. What can be inferred is a similar reduction in customer frustration by being able to quickly and easily resolve a simple issue *without* a support call.

### 3.4. Call Center Toolbox

As discussed in section 3.3, one reason for customers contacting the call center is to perform simple operations. But not all simple operations can be done by an end-user. For those operations it is equally beneficial to have an easy-to-use intuitive portal for the customer service representatives. Such a portal can help to significantly reduce the number of truck roll and at the same time reduce the call durations. A sample portal that can easily be integrated into the RRM architecture is shown on the Figure 23 and Figure 24.

### Call Center Toolbox

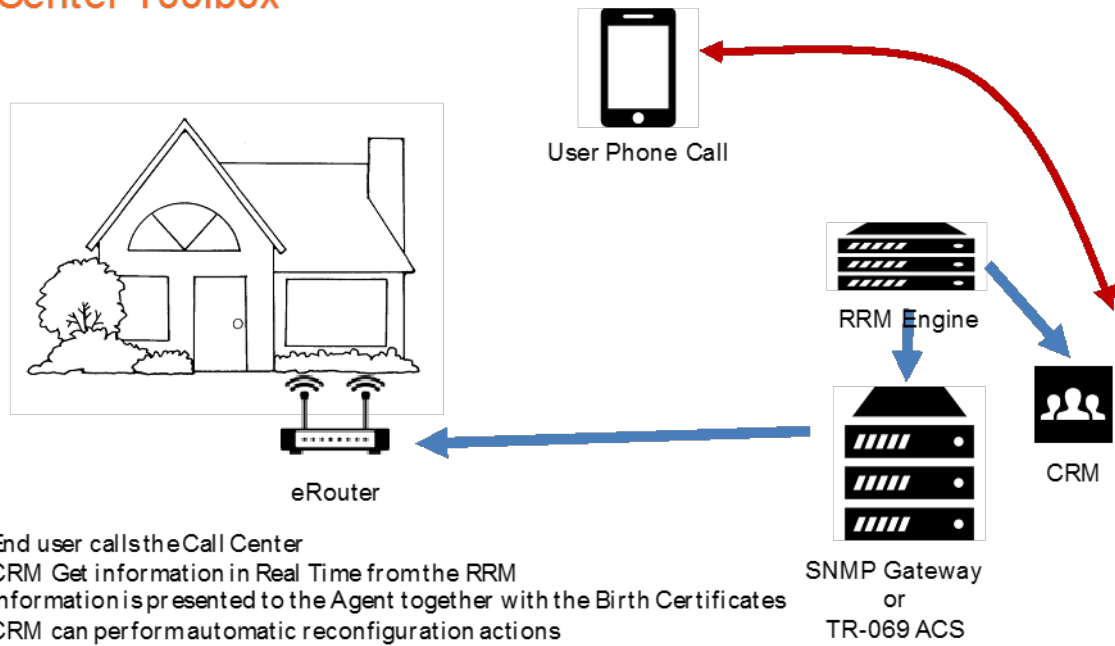


Figure 23 - CSR Management Flow

The screenshot shows the CSR Troubleshoot Portal interface. At the top, there are tabs for 'Real Time Diagnostics' and 'Onsite Quality Certificates'. Below this are four status panels: 'Clients Quality' (PASS), 'Air Quality' (PASS), 'Link Quality' (FAIL), and 'Device Compliance' (FAIL). Each panel has a 'View Details' link and an 'Optimize' button. The 'Device Information' section shows details for a ZHONE device, including vendor, model, type, hardware address, and IP addresses. The 'WiFi Clients' section displays 'Clients Not Found'. The 'WiFi Networks in Range' section shows a table of nearby networks.

#	Network Address	Network Name	Channel	Signal Quality	Noise
1	88:f0:31:ef:44:60	Intraway-Corp	7		-88 dBm
2	a8:b1:d4:22:70:21	Intraway Guest	5		-92 dBm

Figure 24 - CSR Troubleshoot Portal

By providing the customer service representative with a simple, *unified*, interface not only can simple call center functions be performed more easily. A unified, vendor-independent interface would simplify training, reduce mistakes, and speed up problem resolution.

## 4. Conclusions

This paper analyzed the current status of Wi-Fi technology as well as the problems related to the exponential growth of devices in the home and the increased sources of interference.

A detailed explanation of a proposed end-to-end diagnostic and troubleshooting architecture was presented with the objective to reduce not only operational cost but also improve the end-user QoE

## 5. Abbreviations

ACS	auto configuration server
AP	access point
CM	cable modem
CPE	customer premises equipment
dB	decibel
DFS	dynamic frequency selection
DOCSIS	Data-Over-Cable Service Interface Specifications
Gbps	gigabits per second
GHz	gigahertz
IPTV	Internet protocol television
KPI	key performance indicator
Mbps	megabits per second
MCS	modulation coding scheme
MHz	megahertz
MIB	management information base
MIMO	multiple input multiple output
MoCA	Multimedia over Coax Alliance
MSO	multiple system operator
OTT	over-the-top
OSS	operations support system
QoE	quality of experience
RF	radio frequency
RRM	radio resource management
RSSI	received signal strength indication
SNMP	simple network management protocol
SON	self optimizing network
SSID	service set identifier

## 6. Bibliography and References

- [1] IHS Connected Device Market Monitor
- [2] 3Com Wireless Antennas Product Guide

[3] CableLabs Wi-Fi Radio Resource Management (RRM)/Self Organizing Networks (SON) Technical Report

[4] Maintaining Customer Experience in the Face of Increasing Wi-Fi Congestion – SCTE Cabletec Expo 2015 – Christopher Green

# **In-Home Wi-Fi Performance Monitoring and Management**

## **Ensuring QoS Across Managed and Unmanaged Devices**

A Technical Paper prepared for SCTE/ISBE by

Metin Taskin, CTO and Co-Founder, AirTies Wireless Networks, SCTE/ISBE Member  
Istanbul, Turkey  
metin@airties.com



## Table of Contents

<b>Title</b>	<b>Page Number</b>
Table of Contents _____	180
1. Abstract _____	181
2. In-Home Wi-Fi Challenges _____	181
2.1. Consumer Environment _____	181
2.2. Technology Issues _____	182
2.3. Operator Dilemma _____	182
3. Solutions to Ensuring QoS _____	183
3.1. Wi-Fi Mesh Networking _____	183
3.2. Monitoring System _____	184
4. Conclusions _____	187
5. Abbreviations _____	187
6. Bibliography and References _____	188

## List of Figures

<b>Title</b>	<b>Page Number</b>
Figure 1 - Wi-Fi Performance Metrics _____	184
Figure 2 - Wi-Fi Mesh Home Topology _____	185
Figure 3 - Three AP, Nine-Device Wi-Fi Mesh Network _____	186
Figure 4 - Three-AP, 10-Device Wi-Fi Mesh Networks _____	186
Figure 5 - Device/Client-level Detailed Information _____	187

## 1. Abstract

Cable operators have limited visibility into the home Wi-Fi network beyond the router/gateway. This limitation is an issue, because as Wi-Fi devices proliferate and video streaming gains more traction, operators face mounting pressure to ensure that Wi-Fi can deliver the full speeds and range of broadband and video services within the home. To do so, they need the ability to monitor and manage the in-home Wi-Fi network itself.

This paper discusses how operators can identify, monitor and address issues with consumers' home Wi-Fi using a cloud-based data management system. After surveying a number of Wi-Fi challenges and introducing the elements of Wi-Fi Mesh networking, the paper shows how a Mesh-based monitoring platform can arm operators with a full dashboard of detailed analytics on historical events and data, as well as real-time feedback on active Wi-Fi connections, traffic and throughput. Several examples illustrate how this kind of platform provides key performance data on managed and unmanaged devices and delivers actionable, privacy-ensured insights to engineers or field technicians tasked with improving, maintaining or ensuring quality of service (QoS) within subscribers' homes.

## 2. In-Home Wi-Fi Challenges

Wireless local area networking (WLAN) has become almost ubiquitous. The preeminent WLAN technology, specified by the Institute of Electrical and Electronics Engineers (IEEE) 802.11 standard, is now associated with a global Wi-Fi market that reached \$15 billion in 2015 and is estimated to more than double by 2020.<sup>1</sup> It has transformed consumer electronics (CE) and public behavior. Yet it is becoming a victim of its own success, especially on the home front.

Wi-Fi faces challenges from several angles. Almost two-thirds (66 percent) of the 19,000 worldwide consumers polled in the ARRIS 2015 Consumer Entertainment Index said they have issues with Wi-Fi.<sup>2</sup> Each subsequent release of 802.11 has boosted Wi-Fi performance, but the variety of implementations and equipment has led to inconsistency in QoS. Meanwhile, Wi-Fi has become a leading source of trouble tickets for multiple system operators (MSOs), who in turn are limited by their inability to see beyond the home gateway.

### 2.1. Consumer Environment

Consumers love Wi-Fi, almost too much. Wi-Fi-enabled video display devices alone now average globally six per home.<sup>3</sup> Households with over-the-top (OTT) video subscriptions have more. The total number of in-home Wi-Fi devices of all kinds is even higher, and growing, especially as Internet of things (IoT) applications gain steam. But video is a clear driver, and as it grows, home networks can become congested.

Apart from video traffic, Wi-Fi faces other challenges. The home itself can pose problems. Some construction materials impede the flow of Wi-Fi signals, and interference from appliances is another issue. Total square footage may exceed coverage, especially in the higher frequencies. These impediments – combined with the demands of video – have led many consumers to take action. Some buy new routers.

<sup>1</sup> “Global Wi-Fi Market by Business, Product, Service, Vertical, Regional- Global Forecast to 2020,” Market and Markets, July 2015.

<sup>2</sup> “The ARRIS 2015 Consumer Entertainment Index – Global Results,” p. 7.

<sup>3</sup> Ibid., p. 9.

Many invest in Wi-Fi extenders or repeaters. About three out of every 10 consumers globally are now using these devices, solutions that often create problems of their own.<sup>4</sup>

## 2.2. Technology Issues

Few consumers know – or need to know – the ins and outs of Wi-Fi or its various implementations. But most of the issues that arise are technical. Here are a few:

**Standards.** At its base, Wi-Fi is a series of variations on a standard. Just to name three: 802.11ac operates in the 5 gigahertz (GHz) band; 802.11n, in both 5 GHz and 2.4 GHz; and 802.11g, in 2.4 GHz alone. Both “ac” and “n” devices use multiple-input multiple-output, orthogonal frequency division multiplexing (MIMO-OFDM) but “g” devices use only OFDM. These and other differences can impact performance.

**Speeds and range.** The 5 GHz band along with the greater channel options of “n” and “ac” yield much greater throughput than earlier standards. Advertised speeds, however, may overstate the case. The protocol overhead and interference in operating channel result in much lower actual throughput than advertised speed. Meanwhile, higher frequencies, while usually faster, are less able to penetrate objects or traverse distance. Conversely, moving a device away from the home’s main router also has consequences, in some cases drastically impacting the performance of every device.<sup>5</sup>

**Repeaters.** These devices cause trouble about half the time. They may suffer from a poor link with the access point (AP) or router, meaning that a client links up with a strong repeater, but ends up with very low speed. They are not very smart. In-home tests conducted by CableLabs have shown that client selection of Wi-Fi nodes tend to be “solely based on the received signal strength, which may result in sub-optimal or even very poor network performance.”<sup>6</sup>

**Sticky clients.** Mobile or tablet clients are sticky when they retain their connection to a particular AP, even if the signal strength has dropped significantly or a better device is closer. Like “bad apples,” these slowest-performing devices can spoil the whole bunch. By hogging the available airtime, they may degrade the overall performance for all other devices in the home.

These issues tend to overlap. The shorter range of the more advanced – and faster – “ac” standard, for example, may lead a consumer to purchase a repeater; yet that device could be inadequately connected to the AP or attract clients that won’t let go. At any point along the way, of course, this customer may end up calling his or her broadband service provider.

## 2.3. Operator Dilemma

Cable operators who deploy broadband wireless gateway/routers have several Wi-Fi problems. One is their subscriber base. If two-thirds of all consumers have issues with Wi-Fi, as noted previously, then you can be sure that many of them have contacted their broadband service provider. For some operators, Wi-Fi may account for more than half of all call-center traffic. For others, it is still steady and persistent.

<sup>4</sup> Ibid., p. 16. See also Vikas Sarawat, “Are Wi-Fi Repeaters and Extenders Beneficial?” CableLabs, Feb. 23, 2016. Hereafter in this paper, we refer simply to repeaters, instead of both devices.

<sup>5</sup> For examples, see Bulent Celebi, “Modernizing In-Home W-Fi for Video: Multiple Access Points, Intelligence, and Mesh,” NCTA, 2016 Spring Technical Forum.

<sup>6</sup> Sarawat, CableLabs.

Resolving these calls can be arduous. According to one MSO, it involves a mix of customer education, wireless gateway management, provisioning, on-site technical work and device support.<sup>7</sup>

A related problem is that service providers must resolve these calls with limited visibility. Despite massive investments to get ever-faster broadband to the home, MSOs have only rudimentary tools to see within. Simple network management protocol (SNMP) and TR-069 offer basic support. Yet beyond rebooting the gateway, changing its channel assignments, and restoring or updating its settings, there is little else that operators can do, especially to the extensive network of unmanaged devices beyond their range of view. The agent responsible for QoS thus is not empowered to maintain it.

### 3. Solutions to Ensuring QoS

A good answer to these frustrations, challenges and limitations is Wi-Fi Mesh. Adding intelligent APs modernizes the home network in two ways. First, it replaces the static gateway or gateway-plus-repeater model with smart and resilient routing. Second, it enables the collection of operational data that provides for visibility and greater ability to assure QoS.

#### 3.1. Wi-Fi Mesh Networking

Wi-Fi Mesh combines devices, PHY (physical layer) technologies and software. By definition, Mesh enables alternate routes. In-home, that translates to a Wi-Fi router/gateway, along with two or more APs. You can also port Wi-Fi Mesh software onto other devices, such as Internet protocol (IP)-enabled set-top boxes or digital video recorders (DVRs). Sky Europe has done so, transforming its Sky Q family of products into coordinated Wi-Fi hotspots.

The addition of APs extends the reach of Wi-Fi signals, potentially eliminating dead zones. But Wi-Fi Mesh accomplishes much more. Unlike simple repeater devices, multiple Mesh APs can overcome obstacles and manage mobile clients to connect at their maximum capability, while creating resilience and extra capacity. It does so several ways. One is by treating multiple PHYs, 802.11 Wi-Fi and Ethernet, power line communications (PLC) and Multimedia over Coax Alliance (MoCA), as individual point-to-point Mesh connections.

The primary driver of Wi-Fi Mesh performance, however, is routing software. Similar to how IP packets are routed across the Internet, traffic on a Wi-Fi Mesh network is routed by algorithms that determine the best path. The Mesh operates logically as a single service set identifier (SSID) and continuously calculates routes according to the number of active devices, the traffic profile and any QoS required by the subscriber's agreement with the service provider. The decision for best or least-cost path takes into account factors such as source and destination location, number of hops required, and the best end-to-end speed.

This kind of smart routing also enables client steering, which helps devices incapable of roaming to a better AP by themselves. A device may choose the stronger 2.4 GHz signal, for instance, when the intelligent choice is 5 GHz. Another element is channel selection. Instead of a gateway/router using its limited perception of interference, a Wi-Fi Mesh network taps into clear channel assessment (CCA) data and collects information from all APs before making that selection.

---

<sup>7</sup> Christopher Green, "Maintaining Customer Experience in the Face of Increasing Wi-Fi Congestion," SCTE Cable-Tec Expo, 2015, pp. 8, 9.

### 3.2 Monitoring System

Wi-Fi Mesh achieves two goals at once. It delivers much greater performance to the consumer in the home, and it exposes large quantities of data, making the network at last visible to the operator. Once collected and then processed by a cloud-based analysis engine, the data can be displayed in insightful applications that give operators the tools to ensure QoS across both managed and unmanaged devices.

It is possible to create a remote monitoring platform using intelligent APs to collect and share key Wi-Fi performance data from all devices. No download of client-side software on subscribers' personal devices is necessary. Some data is ready for the taking. Every wireless chip already generates streams of data. By default, the software requires subscribers to grant permission to enable monitoring, even though the data collected concerns network connections, not browser-level information about site visits.

A well-designed data structure allows for a lightweight but continuous flow of statistics. One approach is to use TR-069, a common protocol used between gateways and other managed devices. Designed for on-demand connectivity, however, it is relatively heavy, takes more time to retrieve data and does not cover multi-nodes, APs or mesh networked devices. A better approach, instead, is generating massive amounts of current data, which then is processed in large batch files and displayed across several user interfaces (UIs), such as "AirTies Remote View." (See Figure 1) This dashboard offers engineers and operations staff a range of information, including:

- Real-time and historical data on active Wi-Fi connections, traffic and throughput
- Connection rates of wireless APs and third-party devices
- Distribution of 802.11g, 802.11n or 802.11ac clients/devices
- Speed capability and distribution of client brands
- Band connection durations of all clients
- 5 GHz vs. 2.4 GHz connection rates and client steering occurrences
- Airtime consumption

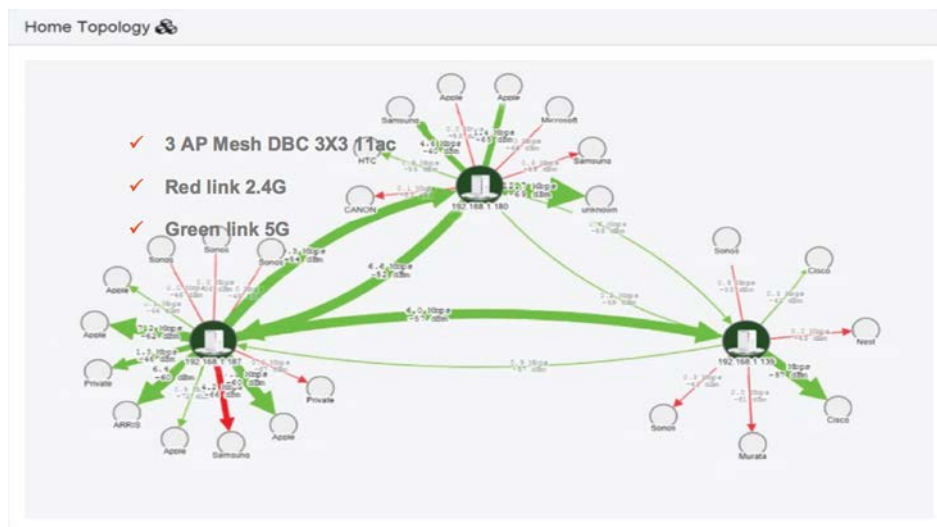


Figure 1 - Wi-Fi Performance Metrics

The "Remote View" platform also displays home topologies, in visual maps of the in-home Wi-Fi Mesh network that show real-time quality and speed conditions across APs. What surfaces in these images sheds light on what lies beyond the gateway. one immediate takeaway, at least from the example seen in Figure 2, is that nearly half of all clients are operating at 2.4 GHz. (see red arrows.) The total number of



clients – 25 – itself is an insight. Twenty-five may seem high. It is, indeed, far removed from the original, two-device personal computer (PC)-printer model of home networking, but the addition of laptops, smartphones, tablets, speakers, thermostats and other Smart Home devices can easily exceed two-dozen.



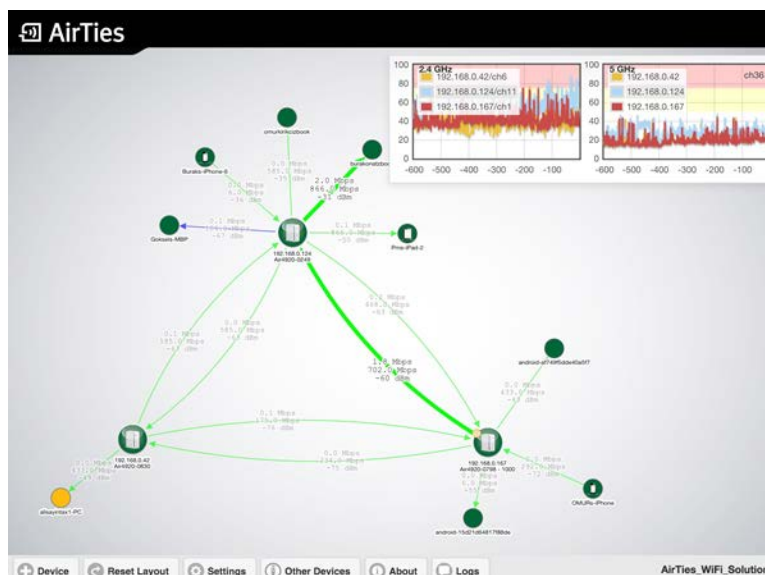
**Figure 2 - Wi-Fi Mesh Home Topology**

What this home topology also reveals is the strength of Wi-Fi Mesh. First, imagine by contrast a single AP serving 25 busy clients. Or what would happen if a repeater were stuck with all of the slow clients. Instead of a potentially feeble link between those two devices, the 5 GHz links backhauling traffic between these three APs all show strong signals, in the mid -50 decibel milliwatt (dBm) range.

Further insight into the Wi-Fi Mesh network is available via an application such as AirTies Network Visualizer (ANV), downloadable to an iPad as a troubleshooting tool for technicians. This app presents another version of the Mesh network maps with real-time quality and speed conditions across the APs within a home, and “AirTime Graphs” displaying how much time (as a percentage) an AP used to send or receive data and any additional interference from other neighboring APs over a shared channel.

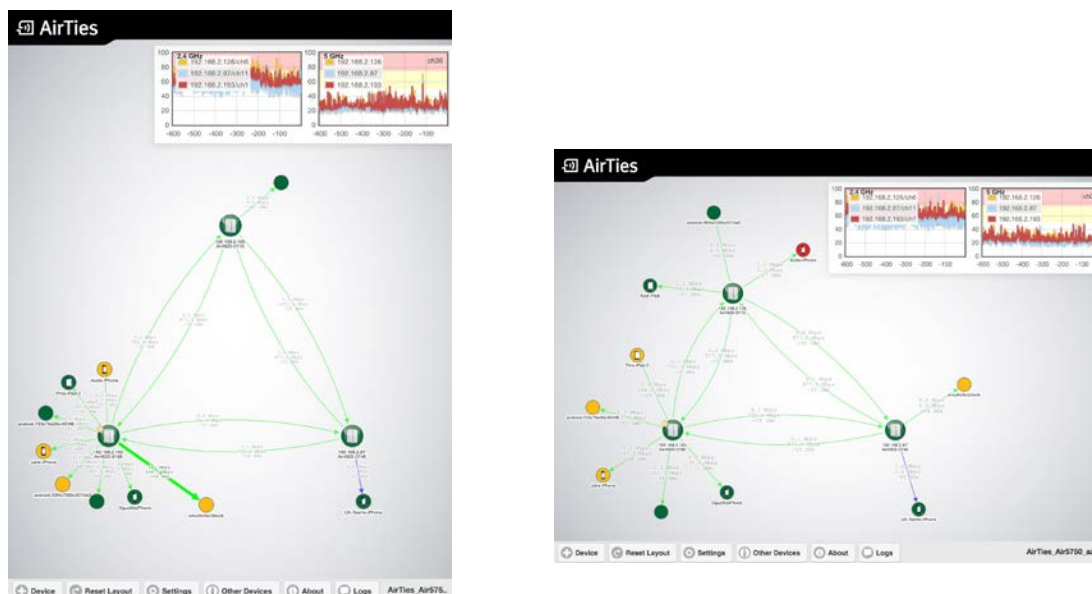
In Figure 3, we see an ANV screen shot of a three-AP Mesh network. The main activity in this image is the 2 megabits per second (Mbps) data consumption by the “book” that draws from the top AP, and by extension the lower-right AP connected to the Internet gateway. The map also reveals a low receive signal strength indicator (RSSI) value of -75 dBm on the bottom link, and the top-left AirTime Graph indicates that 2.4 GHz channels have 50% to 80% interference whereas 5 GHz channel is utilized 20-40% including this home’s network activity. Automatic channel selection enables selection of least busy channel at both 2.4 GHz and 5 GHz and band steering algorithm makes sure that dual band wireless clients are connected to “cleaner” 5 GHz channels. If necessary, and as for the low signal between the bottom-two APs, the described application can flag that deficit by changing the nodes’ background colors. This highlights the need to place one AP closer to the other.





**Figure 3 - Three AP, Nine-Device Wi-Fi Mesh Network**

Figure 4 shows two images of the same network. In this case, three of the devices linked to the lower-left AP moved: two to the top AP, while another to the bottom right. Everything took place due to the client-steering functionality of the Mesh networking platform, which moves clients to different APs when its algorithms determine a better fit. In real time and within the app, a steering occurrence generates a wind sound and dashed-line indicators.



**Figure 4 - Three-AP, 10-Device Wi-Fi Mesh Networks**

As a final example of the kind of visibility and management capability this Wi-Fi Mesh network enables, we return to the “Remote View” platform and device/client-level detailed data related to steering. In Figure 5, we see that an iPhone 6 supporting 802.11ac with 1x1 MIMO was steered 25 times (23 being

2.4 GHz to 5 GHz) over the course of a week. As background, this network was located in a two-story duplex with steel-reinforced concrete floors and walls. The result is that for 99 percent of this time, the iPhone was connected to the closest 5 GHz AP, with an average PHY rate of 410 Mbps, or eight times the PHY rate for the 2.4 GHz AP, to which it was connected only 1 percent of the time.




Figure 5 - Device/Client-level Detailed Information

## 4. Conclusions

Consumers who are using more Wi-Fi devices than ever have a reasonable expectation that wireless performance within the home should be good as the broadband service delivered to the home. When that proves not the case, many contact their service providers or opt to buy and install new Wi-Fi routers, repeaters or extenders. These do-it-yourself projects, however, often run into trouble. Such situations can lead to more trouble calls, forcing broadband service to roll trucks because they have no visibility into the home beyond the managed router/gateway.

Already having invested significantly to improve broadband services, cable operators have a choice. They can sit on the sidelines and be on call for repairs as consumers undertake their own in-home upgrades. Or they can become more proactive. One option is to offer their subscribers smart Wi-Fi Mesh technology that not only boosts capacity and resilience but also gives operators access to heretofore-unavailable information on connections, traffic and throughput. These data can in turn drive a Wi-Fi performance monitoring and monitoring system set up to ensure QoS across both the managed and increasing number of unmanaged devices within consumer homes.

## 5. Abbreviations

ANV	AirTies Network Visualizer
AP	access point
bps	bits per second
CCE	clear channel assessment
CE	consumer electronics
CS	carrier sense
dBm	decibel milliwatt

DVR	digital video recorder
GHz	gigahertz
HFC	hybrid fiber-coax
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of things
IP	Internet protocol
Mbps	megabits per second
MoCA	Multimedia over Coax Alliance
MIMO	multiple-input, multiple-output
MSO	multiple system operator
OFDM	orthogonal frequency division multiplexing
OTT	over-the-top
PC	personal computer
PHY	physical layer
PLC	power line communication
QoS	quality of service
RSSI	receive signal strength indicator
SCTE	Society of Cable Telecommunications Engineers
SNMP	simple network management protocol
SSID	service set identifier
TR	technical report
UI	user interface
WLAN	wireless local area networking

## 6. Bibliography and References

“Are Wi-Fi Repeaters and Extenders Beneficial?” Vikas Sarawat; CableLabs, Feb. 23, 2016

“The ARRIS 2015 Consumer Entertainment Index – Global Results”

“Global Wi-Fi Market by Business, Product, Service, Vertical, Regional- Global Forecast to 2020,” Market and Markets; July 2015

“Maintaining Customer Experience in the Face of Increasing Wi-Fi Congestion,” Christopher Green; SCTE CableTec Expo, 2015

“Modernizing In-Home W-Fi for Video: Multiple Access Points, Intelligence, and Mesh,” Bulent Celebi; NCTA, 2016 Spring Technical Forum

# Orchestrating the Software-Enabled Operator

## An Overview of Key Concepts in Network Function Virtualization Orchestration as Applied to the Cable Network

A Letter to the Editor prepared for SCTE/ISBE by

Alon Bernstein, Distinguished Engineer, Cisco Systems, SCTE/ISBE Member  
375 East Tasman St.  
San Jose, CA 95134  
alonb@cisco.com  
408-525-5970

## 1. Introduction

The term “orchestration” is commonly used in the context of the European Telecommunications Standards Institute (ETSI) standard for network functions virtualization (NFV) or ETSI-NFV. This presentation outlines what orchestration is, what are the pieces that in the author’s view are missing from the ETSI-NFV definition and how it all relates to cable networks.

## 2. Orchstration Overview

### 2.1. Definition

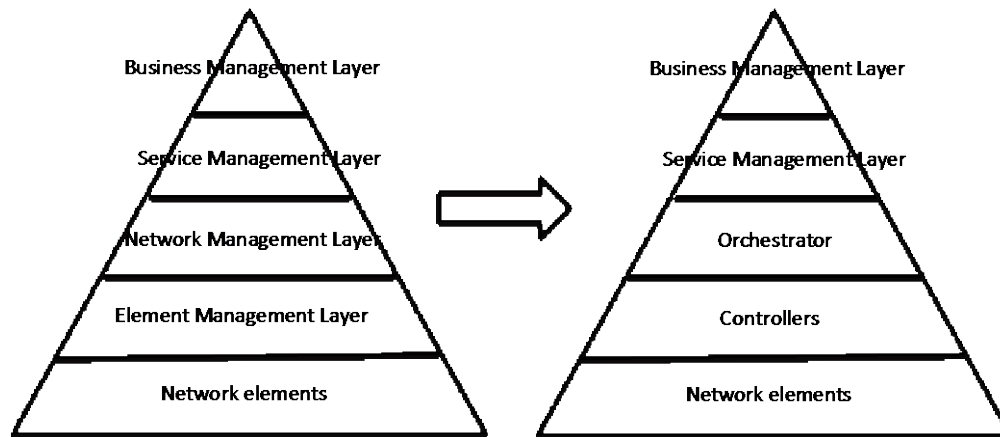
Orchestration is about managing a workflow: Certain operations have to be performed in a certain order and if a failure occurs recovery has to be performed in a certain order as well. An Orchestrator helps manage these workflows. However, since workflows can be defined as a sequence of instructions, why do we need an “Orchestrator”? Isn’t it sufficient to write a script to take care of the sequence of events as well as recovery scenarios? The answer is that an Orchestrator is a platform that makes creating and maintaining these workflows easier to implement. For example, many Orchestrators can accept as an input an end-state and derive a workflow from the end state. Another example: An Orchestrator configures/creates multiple elements (hence the “orchestra” analogy) in what can be viewed as a “distributed transaction” because multiple elements have to be in a well-known state once the transaction is done. A good orchestration platform can help manage this transaction without the need to write a transaction manager for each use case separately.

A common issue in orchestration concerns the application of web deployment tools such as Vagrant, Puppet/Chef as an orchestration solution. The general problem with the current plethora of tools available for managing the cloud is that they do not appear to solve all problems, but instead are optimized for specific use cases (e.g., managing the life cycle of a virtual machine (VM), pushing configurations, and so on). What is needed is a more complete and efficient solution for NFV orchestration.

Orchestration is critical for the development-operations (DevOps) model. Without an automated way to load new software modules into a system the DevOps model cannot be realized.

Orchestration can also be viewed through the lens of traditional network management as depicted in Figure 1.



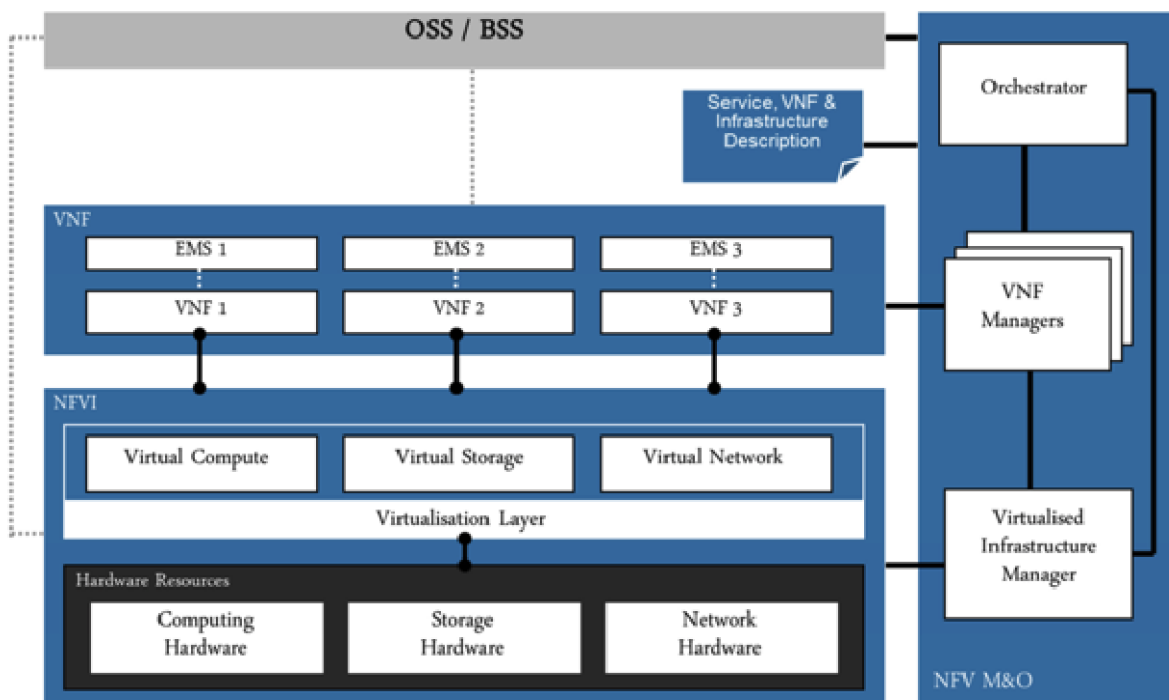


**Figure 1 - Orchestration as network management**

As can be seen, the software-defined network (SDN) controllers are functionally at the element management layer while the Orchestrator is at the network management layer. Unfortunately, the “network management” term became synonymous with “network monitoring.” However, that was not the original intent of the model depicted here. In general, the term orchestration refers only to the part of creation/deletion and configuration, not to data collection or analytics, in other words orchestration is part of a complete network manageability solution.

## **2.2. ETSI NFV MANO**

The ETSI NFV management and orchestration (MANO) framework [1] defines the general framework for network function virtualization as well as the role the Orchestrator plays in it as defined in Figure 2.



**Figure 2 - ETSI NFV MANO**

The reader is encouraged to read the ETSI document in order to get familiar with the next level of detail. Here, we will only mention the areas where more work is needed for the aforementioned NFV model:

1. The NFV model is very data center-focused. But as we explore later, in cable networks we need to orchestrate both physical and virtual appliances.
2. An Orchestrator needs to coordinate several domains in order to provide end-to-end solutions. In a cable network “domains” can be the data center, the wide area network and the access network. This means that above the ETSI-MANO Orchestrator there needs to be a higher-level Orchestrator to coordinate an end-to-end solution across all these domains.
3. The diagram in Figure 2 presents the MANO as a self-contained module but in reality the Orchestrator is a platform and there are “orchestration applications” that are written above the platform orchestration layer.

In summary the ETSI NFV model for orchestration covers only one domain (data center) of the overall solution.

Very recently AT&T has published an open source implementation of their NFV platform (ECOMP - Enhanced Control, Orchestration, Management & Policy) where they extend MANO to include more of the “application level,” including policy framework and meta data description of services [2].

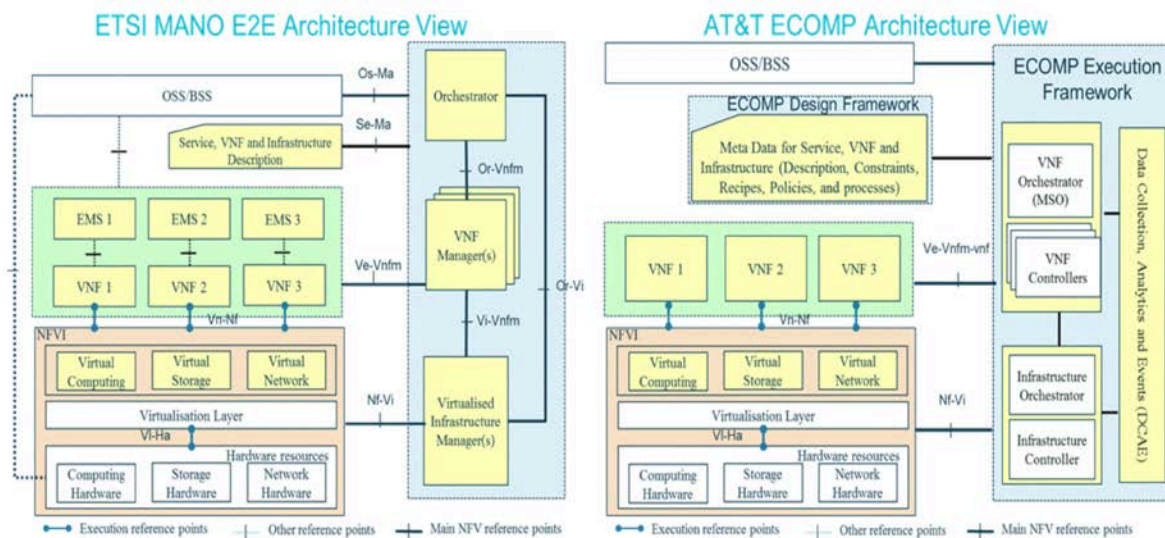


Figure 3 - AT&T ECOMP and ETSI MANO

### 2.3. Orchestration In The Web World

The success of the cloud and the cloud-based tools have inspired ETSI NFV but when it comes to orchestration it is apparent that while orchestration is central to ETSI NFV it is less elemental in the application “cloud native” environment. This is most likely because the cloud native services of today are focused on application deployment and less about networking these applications. Having said that, the NFV world can definitely adopt a couple of key concepts from the application-focused world, specifically when it comes to automatic scaling and fault recovery, to the cable industry.

### 2.4. Customer-Facing and Infrastructure-Facing Orchestration

One area that can be a source of confusion is the fact that infrastructure-facing NFVs, such as virtual cable modem termination system (CMTS) or virtual optical line termination (OLT) can use the same orchestration tools as customer facing services such as a layer 2 virtual private network (L2VPN) or parental controls. This duality is not a problem for the Orchestrator since it operates at a high enough level where a VM is just a VM, regardless of whether it is a customer application (e.g., virtual disk) or a provider-facing application.

### 2.5. Orchestration For Cable Networks

Orchestration for cable networks includes a domain that is unique to cable, the hybrid fiber-coax (HFC) access network. With distributed architecture such as remote physical layer (PHY), there is a network between the CMTS and/or edge quadrature amplitude modulation modulator (EQAM) cores that has to be configured and maintained in addition to other networking domains. The sheer number of remote PHY devices possible means that the network cannot be managed manually and an orchestration framework is essential.

As cable operators need to support more types of access media, such as Wi-Fi, passive optical networks (PONs) and possibly others, a solid orchestration framework is key to unified deployment models and the ability to quickly switch subscribers from one access media to another.

### 3. Conclusions

The data center scale and complexity require automation, and it is not feasible to create services on this infrastructure with Excel spreadsheets and manual configurations, for example. The NFV world requires this automation as well. More specifically, the scale of the remote PHY architectures and the need to support multiple access technologies will drive cable operators to adopt orchestration solutions as well.

### 4. Bibliography and References

1. ETSI NTV MANO : [http://www.etsi.org/deliver/etsi\\_gs/NFV-MAN/001\\_099/001/01.01.01\\_60/gs\\_nfv-man001v010101p.pdf](http://www.etsi.org/deliver/etsi_gs/NFV-MAN/001_099/001/01.01.01_60/gs_nfv-man001v010101p.pdf)
2. AT&T ECOMP : <http://about.att.com/content/dam/snrdocs/ecomp.pdf>

## **PNM Developments**

### **Improving Customer Satisfaction with Better Proactive Network Maintenance**

A Letter to the Editor prepared for SCTE/ISBE by

Tom Williams, Principle Architect, CableLabs  
858 Coal Creek Cir  
Louisville, CO 80027  
t.williams@cablelabs.com  
303-661-3486



## 1. Introduction

This letter describes some efforts underway at CableLabs to improve customer satisfaction with better proactive network maintenance (PNM), specifically in the areas of leakage testing, testing of plant for impedance mismatches, using SCTE-109 in the field to detect common path distortion (CPD), and that the detectable CPD often varies at a 120 Hz rate.


## 2. Recent PNM Innovations

As you extend sources of plant data, you learn more. We currently have a number of DOCSIS<sup>®1</sup> (Data-Over-Cable Services Interface Specifications) indicators that tell us something is wrong, like low receive levels, and forward error correction (FEC) errors, but not exactly what, or where, or for how long (1). For example, ultra high frequency (UHF) leakage testing has led to discovery of plant defects, such as radial cracks and holes in the cable's sheath. These defects reflect signal energy back to a launch amplifier where it is absorbed, while some of the energy is radiated into space. The net effect is subscribers receiving low signal levels and interference with long term evolution (LTE) services. Often these single impedance mismatches don't form an echo tunnel, so there is a flat frequency response in a home's received broadband signal, albeit at reduced RF levels. This makes detection difficult.

CableLabs has been developing a leakage detection method using a phase stable continuous wave (CW) test signal, which is received by software defined radios (SDRs) and analyzed for Doppler shift. The signal samples are captured as a test vehicle moves and are used to create a synthetic phased array (SPA). The SPA reveals a first bearing angle to the direction of the leak. After the vehicle moves, the crossing of the first angle with a second bearing angle shows precisely the origination point of the leak. This is illustrated in Figure 1. Note that GPS data shows where the leak originated, not the location of the truck was when it detected the leak. It is anticipated that the leakage detection system, combined with automatic data collection, will provide information for a leakage overlay on a PNM map, resulting in faster problem resolution. Automated leak detection, for example by a pizza delivery or a garbage truck, promises to provide a large valuable set of data for relatively low cost. The leakage data (latitude, longitude, signal strength) are pushed into a data base for mining by PNM applications.

---

<sup>1</sup> DOCSIS is a trademark of CableLabs



**Figure 1 - Drive path showing vehicle's current location at green arrow. The intersection of bearing angles reveals a leak's source location. Colored dots are leaks discovered earlier on the drive path.**

Another area of investigation involves using the magnitude component of broadband data signals to reveal reflections inside cables. These broadband random data signals are captured with a high impedance probe touching the cable's center conductor and analyzed for standing waves. The standing waves are processed using a math trick (2) to make an impulse response, so a tech can know a distance to a reflection. This test method has advantages of being non-service interrupting, since the actual live digital cable signals are used as a test signal. It also has an advantage of high distance accuracy because of the wideband nature of the data signals, which can be a block of DOCSIS 3.0 signals or a DOCSIS 3.1 signal. Another advantage is revealing exactly where the reflection is located: it is downstream from the probe's location. An echo tunnel formed by two reflections can be analyzed for its length, but *where* it is located must be determined by other means. There are multiple cost-effective ways to get wideband magnitude data, including full band capture (FBC) chips, SDRs, and even old analog spectrum analyzers.

Another discovery that was made relates to upstream noise created by corrosion-induced diodes mixing the downstream signals. This phenomenon is called CPD. As plants are switching from analog to digital signals, the CPD is losing its three beats every 6 MHz appearance and starting to look like an elevated noise floor. But it is not going away like some technicians mistakenly are reporting. We found that the SCTE-109 test method, which is a lab bench test method, also works for outdoor in-service plant. We put up two downstream CW carriers at 800 MHz and 840 MHz on in-service plant and looked for an upstream difference beat at 40MHz. Running field tests show us something surprising: the 2nd order beat product is frequently modulated at a 120 Hz rate by alternating current (AC) on the coaxial line. Perhaps one node in 20 shows elevated levels of static CPD, and 5% of the nodes measured exhibit this dynamic behavior. Furthermore, the CPD test results vary with time and temperature. This phenomenon needs to be investigated further.

### 3. Conclusions

In conclusion, this letter briefly discussed recent innovations in leakage testing, non-service interrupting testing of plant for impedance mismatches, the discovery that CPD is detectable using a field implementation of SCTE-109, and that the detectable CPD often varies at a 120 Hz rate. The

development of new PNM tools is on a fast track. An upcoming challenge for both cable management and the technical workforce will be an assimilation of the new high-value information.

#### 4. Notes and References

- (1) Complaining customers are another vague indicator that there may be a line problem, but often without sufficient details to make a quick repair.
- (2) Inverse Fourier transforms require complex frequency values as an input and return complex time values as an output. The math trick is to (falsely) use zero as the phase value for the frequency values. While this method will not work to reveal linear distortions like group delay, it works very precisely with standing waves.
- (3) Various aspects of these PNM tools are patent pending
- (4) CableLabs publicly-available publications can be found at:  
<http://www.cablelabs.com/resources/publications/>



**Society of Cable  
Telecommunications  
Engineers**



**International  
Society of Broadband  
Experts**

Society of Cable Telecommunications Engineers, Inc.  
International Society of Broadband Experts™  
140 Philips Road, Exton, PA 19341-1318  
T: 800-542-5040 F: 610-884-7237  
[www.scte.org](http://www.scte.org) | [www.isbe.org](http://www.isbe.org)