

SCTE TECHNICAL JOURNAL

SCTE TECHNICAL JOURNAL

VOLUME 3, NUMBER 3

November 2023

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

© 2023 by the Society of Cable Telecommunications Engineers, Inc. All rights reserved.

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

Table of Contents

4	Foreword
5	The Importance of Performance Testing for Residential Wi-Fi Deployment Leigh Chinitz, Spirent Communications
30	Scaling the DAA Remote PHY CIN Deployment with Cloud Native Routing Hitesh Mali, Juniper Networks, Inc. Bob Gaydos, Comcast Brian Sherwood, Juniper Networks, Inc.
38	Automatic Real-time Detection of Congested Multimedia Networks Mun Gi Choi, Charter Hany Heikal, Charter Manish Jindal, Charter
55	Climate Control Automation for Critical Facilities: Supporting an “Always On” Network Through Automation John Dolan, FTCO Past Chair Arnold Murphy, FCTO Ken Nickel, Quest Controls Tom Hurley, Comcast
68	Automating Greenfield Access Planning Using Intelligent GIS Algorithms: Getting Buildout Cost Estimates in Minutes, Not Months Jared Faulk, First Principles Innovations Luc Absillis, First Principles Innovations
85	Improvements in Upstream Profile Management Algorithms: Delivering Reliable and High Upstream Speeds in Production Networks with OFDMA Technology Marie-Sophie Graftiaux, VodafoneZiggo Ying Wang, VodafoneZiggo Karthik Sundaresan, CableLabs

***SCTE Engineering Committee
Chair and Data Standards
Subcommittee (DSS) Chair:***
David Fellows,
SCTE Member

Senior Editor
Paul Schneider
Paul Schneider Public Relations

Publications Staff
Chris Bastian
SVP & Chief Technology Officer,
SCTE

Dean Stoneback
Senior Director, Engineering &
Standards, SCTE

Natasha Aden
Document Editor, SCTE

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, 140 Philips Road, Exton, PA 19341-1318 or email journals@scte.org.

Submissions: If you have ideas or topics for future journal articles, please let us know. All submissions will be peer reviewed and published at the discretion of SCTE. Electronic submissions are preferred, and should be submitted to SCTE Journals journals@scte.org.

I just finished authoring an article on how cable networks are changing and transforming to exceed customer expectations, and it struck me that that theme aligns perfectly with this moment in SCTE time. Change is happening with the departure of Mark Dzuban from the CEO's seat, and expectations were certainly exceeded by the impressive turnout at SCTE Cable-Tec Expo in Denver.

Expo 2023 was a fitting departure for Mark, whose imprint on SCTE is large and lasting. Energy management and sustainability, business continuity, telehealth and more – just some of the programs launched with Mark at the helm – all were part of the jam-packed agenda, which helped to attract a significant increase in attendees when compared to Expo's last visit to The Mile High City.

With Expo 2024 in Atlanta next September 23-26 firmly in our sights, we're continuing to grow our role as the hub of applied science for the broadband industry. This month's SCTE Journal article lineup is a reflection of how the industry continues to lead the way in maximizing the value of telecommunications networks.

- **"The Importance of Performance Testing for Residential Wi-Fi Deployment,"** by Spirent Communications' Leigh Chinitz.
- **"Scaling the DAA Remote PHY CIN Deployment with the Cloud Native Routing,"** by Juniper Networks' Hitesh Mali and Brian Sherwood and Comcast's Bob Gaydos.
- **"Automatic Real-Time Detection of Congested Multimedia Networks,"** by Charter Communications' Mun Gi Choi, Henny Heikal, and Manish Jindal.
- **"Climate Control Automation for Critical Facilities,"** by Strategic Clean Technology, Inc's Arnold Murphy, Quest Controls' Ken Nickel, and Comcast's Tom Hurley.
- **"Automating Greenfield Access Planning using Intelligent GIS Algorithms,"** by First Principles Innovations' Jared Faulk and Luc Absillis.
- **"Improvements in Upstream Profile Management Algorithms,"** by VodafoneZiggo's Marie-Sophie Graftiaux and Ying Wang and CableLabs' Karthik Sundaresan.

The year ahead will bring new leadership and continued expansion of SCTE's game-changing mission to make our industry better and subscribers more satisfied. I'm thrilled to be together with all of you on that journey and invite you to share your knowledge with SCTE Journal audiences and to be on the lookout for the Call for Papers for Cable-Tec Expo 2024.

Best wishes for the holiday season. See you in 2024!

The Importance of Performance Testing for Residential Wi-Fi Deployment

A Technical Paper prepared for SCTE by

Leigh Chinitz, Senior Director of Solutions Architecture
Spirent Communications
305 Foster St #104,
Littleton, MA. USA, 01460
Leigh.Chinitz@Spirent.com
+1.781.248.3952

Table of Contents

Title	Page Number
Table of Contents	6
1. Introduction	8
1.1. Regulatory Testing	8
1.2. Conformance Testing	8
1.3. Performance Testing	9
2. Performance Testing in Wi-Fi	9
2.1. The Meaning of the Word "Performance"	10
2.2. The Turning Point for Performance Testing	11
3. Survey of the Performance Testing Landscape	11
3.1. The Assumption of Compact, Consistent Testbeds	12
3.2. BBF TR-398	14
3.2.1. TR-398 Issue 1	14
3.2.2. TR-398 Issues 2 and 3	16
3.2.3. Wi-Fi Alliance Device Metrics	16
3.2.4. ETSI Multiple Access Point Performance	17
4. The Future of Wi-Fi Performance Testing	17
4.1. New Wi-Fi Features, and their Area(s) of Application	18
4.2. Multi-User Operation in Wi-Fi 6	20
4.3. Advanced Power Save Operation in Wi-Fi 6	21
4.4. More Features Arriving in Wi-Fi 7	21
5. Challenges Related to Advanced Feature Testing	21
5.1. Automation as a Key Component to Wi-Fi Performance Testing	23
5.1.1. OFDMA Performance Testing	23
5.1.2. Roaming Performance	25
6. Conclusion	27
7. Abbreviations and Definitions	28
7.1. Abbreviations	28
7.2. Definitions	28
8. Bibliography and References	28

List of Figures

Title	Page Number
Figure 1- Wi-Fi by the numbers	8
Figure 2 - Illustrative Data Taken during a Walk Test in a Real Building	10
Figure 3 - BBF Single-Chamber Diagram	12
Figure 4 - BBF Multiple Chamber Testbed Implementation	13
Figure 5 - Wi-Fi Alliance Example Three-Chamber Testbed	13
Figure 6 - Multiple Chamber Testbed for Multiple Access Point Testing	14
Figure 7 - Example of a Range vs. Rate Curve	15
Figure 8 - TR-398 Multiple STA Performance. A Visual Representation.	16
Figure 9 - The World's Most Used Apps by Downstream Traffic	19

Figure 10 - Some New Wi-Fi Features Mapped with Applications	20
Figure 11 - A Testbed for Running Traffic between Two Endpoints	22
Figure 12 - Example Testbed for a 4-Node Mesh with Multiple STAs	22
Figure 13 - OFDM vs. OFDMA channel access	23
Figure 14 - Example Analysis of RU Allocations in an OFDMA System	24
Figure 15 - Example Analysis of Channel Allocation Efficiency in an OFDMA System	24
Figure 16 - OFDMA Communication in Wi-Fi, with Sniffing Indicated.	25
Figure 17 - Multiple Access Points Performance, Two-Stage Networking Test	26
Figure 18 - Wi-Fi Alliance Device Metrics Roaming Topology	27
Figure 19: Maximum Roaming Throughput Achieved over a Set of Runs	27

1. Introduction

Somewhat surprisingly, the concept of standardized performance testing is a relatively new concept in the world of Wi-Fi. When we talk about performance testing, we are distinguishing it from other product/device testing, such as regulatory testing and conformance testing. So, it will be useful to begin by defining those other types of testing that are not the focus of this paper.

1.1. Regulatory Testing

Everywhere in the world, before wireless products can be used in an open-air environment, they must conform to certain regulatory requirements. In the US, for example, Wi-Fi products must conform to the Federal Communications Commission's (FCC's) Part 15 rules for Radio Frequency Devices. These tests are aimed at making sure that devices obey the operating rules for the spectrum they will use, and that they don't interfere with other users in that spectrum or in other spectrum bands. The rules say nothing about what technology should, or should not, operate in the spectrum. This is why, for example, a Wi-Fi device and a garage door opener can operate in the same band. The devices don't interoperate (that is, they don't communicate with each other) but they can share the spectrum. This is the focus of almost all regulatory testing – basic spectrum compatibility.

1.2. Conformance Testing

From a Wi-Fi perspective, this is the type of testing that is most well-known/understood and has contributed to the incredible growth of Wi-Fi over the past 20 years.

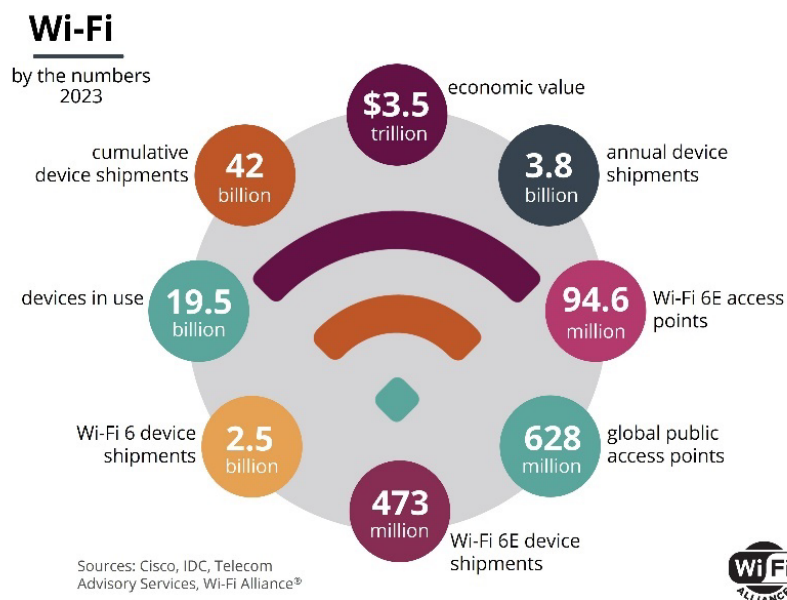


Figure 1- Wi-Fi by the numbers¹

The success and growth of Wi-Fi, as demonstrated in Figure 1, is due in large part to the conformance testing regime created and managed by the Wi-Fi Alliance. Although it is the IEEE's 802.11 group that

creates the specifications that describe the underlying Wi-Fi technology, it is the Wi-Fi Alliance that has imposed a conformance testing regime that awards certification to products declaring that they support some subset of the IEEE specification(s). It is this conformance regime that has allowed users of the technology, both consumers and service providers, to have confidence that certain devices will interoperate with other devices. When those devices bear the “Wi-Fi Certified” label, interoperability is all but guaranteed. People trust that products from one company will work with products from another and that it will be simple to create a multi-vendor deployment. It is partly for this reason that, for example, a survey in the UK declared that “Wi-Fi is the best invention of the past 25 years.”²

1.3. Performance Testing

Given the success that Wi-Fi has seen over the decades, why is even more (and different) testing important? That is, why on top of regulatory and conformance testing is there now a focus on standardized performance testing? And how has it been that, over all this time, performance testing has not been an industry focus? We will answer these questions and get into the details of performance testing in the next sections.

2. Performance Testing in Wi-Fi

The difference between conformance testing and performance testing is the difference between the question “*will* this device interoperate with that device,” and “*how well* will this device perform when interoperating with that device?” It is, in many respects, the success of the Wi-Fi Alliance’s conformance testing regime that has elevated the need for performance testing. If the conformance testing were not a success, there would be very few interoperable products and very few deployment options to consider. The most extreme example of this is for proprietary technologies with no conformance standards, in which all radios in the system must be manufactured by the same vendor.

Given that Wi-Fi products are so interoperable, when a service provider, for example, decides to deploy a Wi-Fi enabled router, they have a large number of options among which to choose, all of which are, at a minimum, Wi-Fi certified. The question then becomes, do they all give the same performance? This is a question that conformance testing is generally not designed to answer, and this is where performance testing comes into the picture.

It would be misleading to say that throughout the history of the development and deployment of Wi-Fi products, no one has done any performance testing. This is clearly not the case. What is the case, however, is that no *standards organizations* have addressed the topic of performance testing until very recently. In the early days of Wi-Fi, much performance testing was done using walk-test methods, in which Wi-Fi devices were deployed into actual buildings, and measurements were taken with test tools to look at how the devices behaved under various conditions. (See Figure 2.)



Figure 2 - Illustrative Data Taken during a Walk Test in a Real Building

This methodology has a number of drawbacks, in that:

- (a) The results are not very repeatable because in open-air environments, the RF environment can be constantly changing, which will affect the results.
- (b) The results are not easy to reproduce across different testing labs, because one real environment will likely be very different from another.
- (c) The testing can take a very long time, since networks need to be configured, and testers have to physically walk around the environment to collect data before it is analyzed.
- (d) Various test conditions can be very difficult to set up in a real environment; for example, many devices send various kinds of traffic under different interference conditions, etc.
- (e) The testing is not amenable to standardization because it's very difficult to standardize a physical environment (like a home) in which devices would be deployed and walk-test data would be collected.

These are among the reasons that performance testing has not been standardized until recently, but there is one more reason that has limited the development of standardized performance testing.

2.1. The Meaning of the Word “Performance”

Standards bodies are made up of participants from many areas, including academia, government, service providers, and equipment vendors. While all of these stakeholders see the need for equipment to conform to regulatory requirements (regulatory testing), and in many cases they all see advantages to products meeting conformance requirements (conformance testing), it has not always been quite as clear that there is a common understanding of performance testing.

Equipment vendors in particular have historically been concerned that once a test defines performance characteristics, certain products will be perceived as “better” than other products, even though there might be a complicated mapping of performance to price point, so that what is “better” for one user might not be “better” for another with different requirements.

2.2. The Turning Point for Performance Testing

Despite the difficulties in creating standardized performance tests for Wi-Fi, the past several years have seen a rapid uptick in the number of such standards, and the momentum for them is continuing. How has this happened?

- (a) The advent of compact Wi-Fi testing tools has removed many of the difficulties listed above. Wi-Fi testing can now be done more quickly, more repeatably, under more complex scenarios, and in ways that can be more easily defined by a standards organization.
- (b) An understanding has emerged across the industry that performance testing does not imply “better;” rather, it looks at whether a product is “fit for purpose.” There is a complicated mapping of a product’s feature support, performance, price, and intended area of operation. Because of this, there is no real meaning to saying that one product is “better” than another, but it is possible to say that a product may be more fit for use in, for example, a small home environment, while a different product is a better fit for a larger, crowded, interference-prone office environment.

3. Survey of the Performance Testing Landscape

The demand for standardized performance testing for Wi-Fi was ultimately driven by service providers (although other stakeholders have been active participants and supporters). The success of conformance testing led to a dramatic rise in the use of Wi-Fi, and as a result, most users experience their broadband connection via Wi-Fi as the final link. According to one study, 92% of US internet households use Wi-Fi at home.³ This being the case, it became vitally important that broadband service providers manage this part of the connection since the entire broadband experience depends on it. With ISPs making Wi-Fi a part of their service offering, they need to be able to compare different Wi-Fi products to see which have the best fit with their offering, so they started to advocate that standardized performance testing was necessary.

The first standards organization to act on this was the Broadband Forum (BBF), which in 2019 issued the first version of its TR-398 test plan, “Wi-Fi In-Premises Performance Testing.” The focus of this testing was residential access points. (Client devices are not a focus of this test plan.) This was the first Wi-Fi test plan to focus on performance by including not only test cases and methodology but also pass/fail criteria to go along with each test. Obviously, this could only be done under the assumption/requirement that largely similar, repeatable testbeds could be created in which these tests could be run.

A few years later and at roughly the same time, both the Wi-Fi Alliance and ETSI also entered the Wi-Fi performance testing space with the Wi-Fi Customer Experience group’s release of the Wi-Fi Device Metrics test plan⁴, and ETSI Broadband Radio Access Network’s (BRAN’s) release of the Multiple Access Points Performance Testing plan⁵, both in mid-2022. As is already somewhat obvious from the names of these test plans, they do not all cover the same thing in terms of “performance.” (This relates to our previous comment about products being “fit for purpose.”) For example, while the BBF TR-398 test plan is explicitly called out as being for in-premises performance, the ETSI test plan is called out as being explicitly about multiple access point performance (so performance of APs when configured in a mesh, or extender, scenario). There are many areas of “performance” that can be examined, and in the next sections, we will summarize what each of these groups has covered so far, and then we will discuss what types of performance testing have yet to be addressed.

3.1. The Assumption of Compact, Consistent Testbeds

One thing that is immediately obvious when reviewing these three specifications is the impact that compact test systems have had on the ability of groups even to consider widespread, standardized performance testing. The BBF, for example, describes how the testing can be done in either single, or multiple-chamber testbeds (Figures 3 and 4). In some more advanced test cases (discussed below) only the multiple chamber testbed option is available.

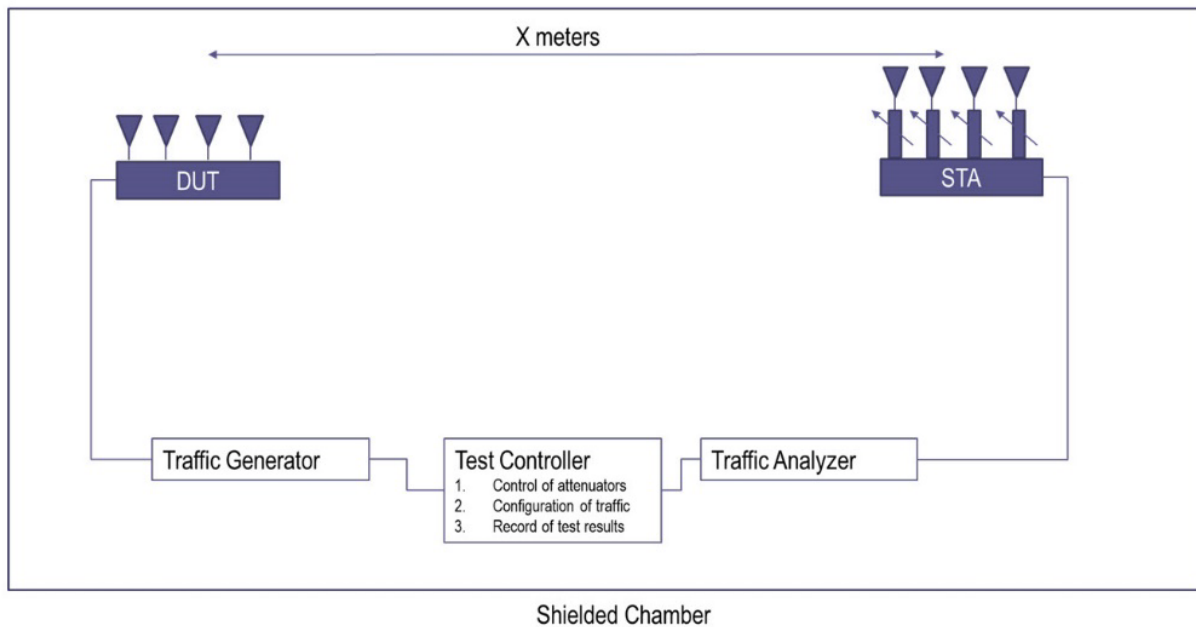


Figure 3 - BBF Single-Chamber Diagram

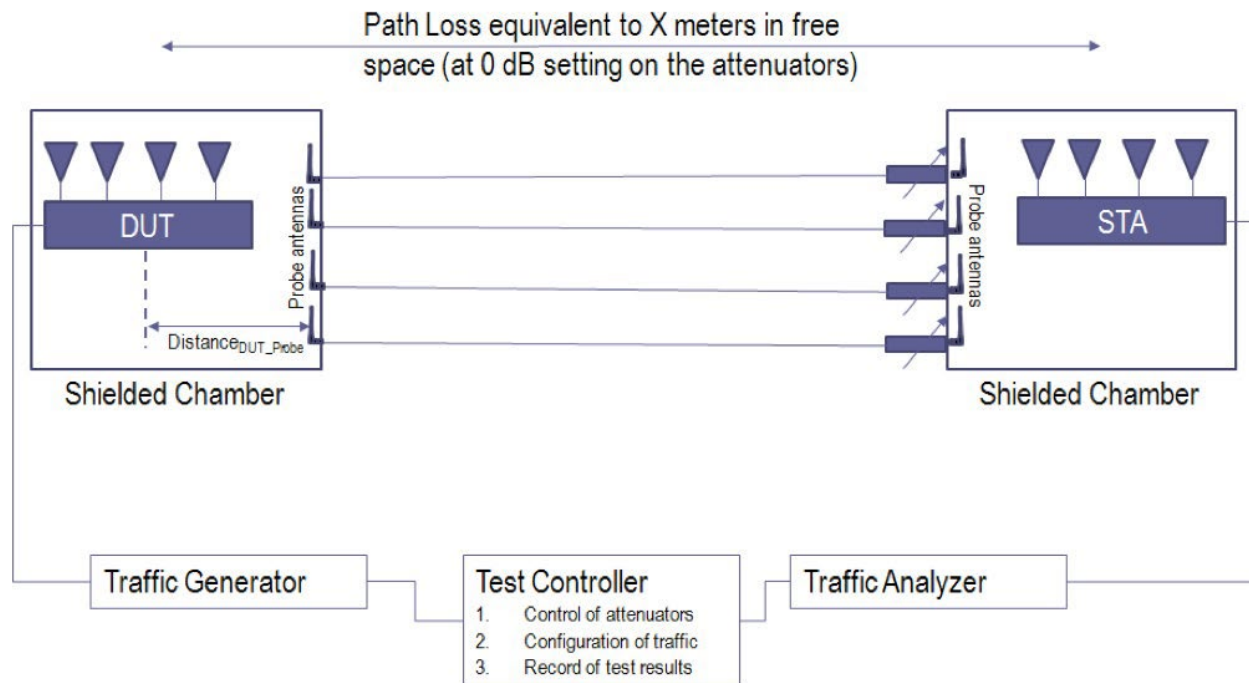


Figure 4 - BBF Multiple Chamber Testbed Implementation

The Wi-Fi Alliance shows, as an example, a multiple chamber implementation that can be used for their testing (Figure 5).

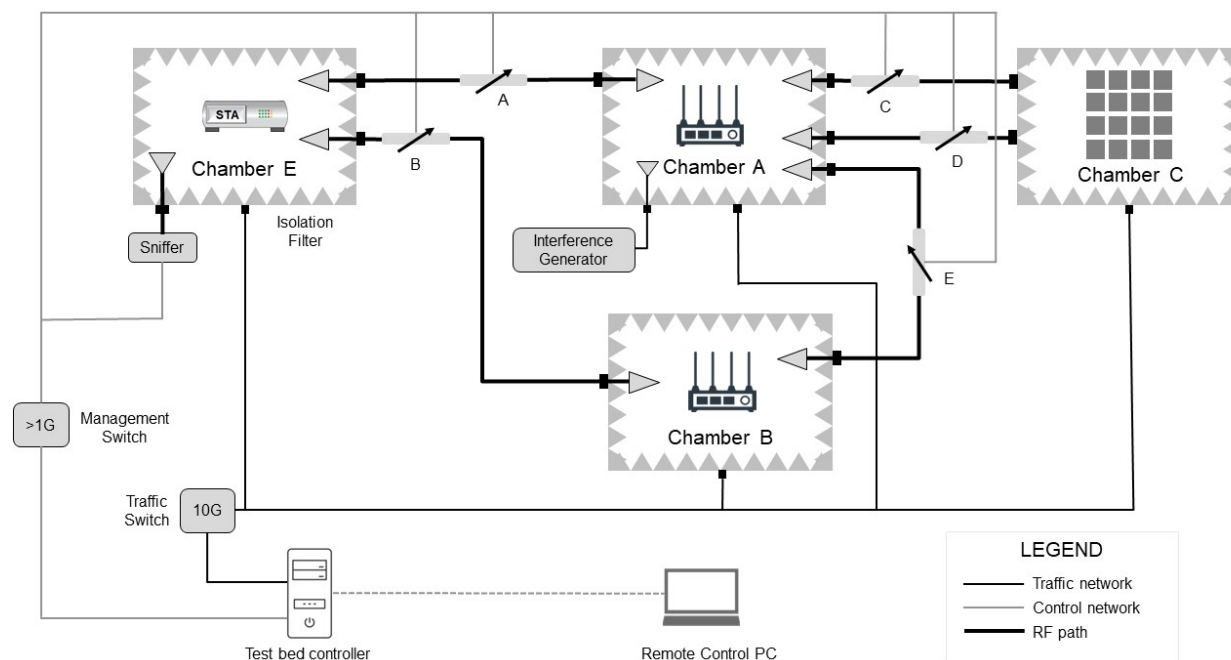


Figure 5 - Wi-Fi Alliance Example Three-Chamber Testbed

And, similarly, the ETSI test plan describes how a multiple chamber testbed can be used to configure the complex, multiple access point scenarios described in that test plan.

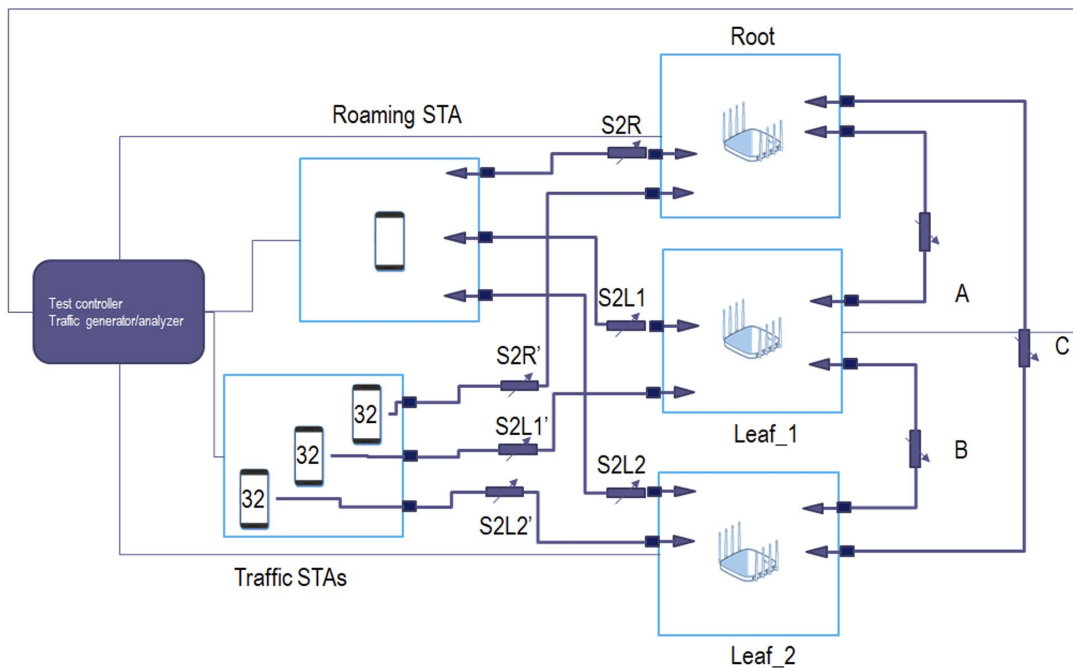


Figure 6 - Multiple Chamber Testbed for Multiple Access Point Testing

There is commonality among all of these standards: they all require the testing be done in a controlled RF environment (not in open air). In addition, they all require (at least for some of the test cases) that multiple, independent, interconnected RF chambers be used to create more complicated testing topologies than would be possible in a simple shield room scenario.

3.2. BBF TR-398

As mentioned above, the first group to enter the Wi-Fi performance test space was the Broadband Forum with TR-398, which is why that group is currently nearing the release of the third revision (or “issue”) of that specification. As these tests were primarily created to respond to the needs of service providers, they are focused on the performance of Wi-Fi access points, or APs, (not clients, or “STAs” in Wi-Fi terminology). In addition, the focus has been primarily (but no longer solely) on single access point deployments, which implies a fairly simple home deployment with a single AP. This is the use case for most home Wi-Fi deployments.

3.2.1. TR-398 Issue 1

As the first mover in the Wi-Fi performance testing space, the BBF group began with a fairly wide-open field of items to test. In Issue 1, they broke their plan down into five sections:

- 1) RF capability
- 2) Baseline performance
- 3) Coverage

- 4) Multiple STA performance
- 5) Stability/Robustness

We will not, in this paper, get into the details of each test, its configuration, etc. (The specification itself is freely available, and all of the details are spelled out there.) Rather, we will focus on how, even in this first set of tests, the complexity of Wi-Fi performance testing can be seen, and the idea of “fit for purpose” becomes ever more apparent.

In some of the tests above, the concept is quite simple, and it might seem obvious what a good result is and what a bad result is. Under “baseline performance”, for example, is a simple test of maximum throughput. Traffic is run from the device under test (DUT), which is an AP in this test plan, to a test instrument, and the throughput is measured. It would seem obvious that in this test, a higher number is better.

But not all tests are quite so obvious. A fairly simple test (in the coverage section) is the “Range vs. Rate” test, in which the test emulates moving a STA away from the AP so we can measure not only the maximum throughput (at the lowest attenuation point) but also the throughput as the attenuation between the endpoints is increased (emulating an increasing spatial separation). However, depending on the expected deployment scenario, the performance of an AP at high attenuation might be more important for some use cases than for others. Nevertheless, the test plan still needs to provide a baseline of performance in order to be able to qualify as a performance test plan.




Figure 7 - Example of a Range vs. Rate Curve

While the Range vs. Rate test might seem to imply an obvious result (more throughput is better), there are even more complicated tests where the ideal result can be less obvious. Take for example, the “multiple STA performance” test. In this test, we emulate sending traffic first to three clients close to the AP. Then we add in another three clients that are a “medium” distance away. And finally, we add in three more clients that are “far” away.




Figure 8 - TR-398 Multiple STA Performance. A Visual Representation.

While the test plan does call out the threshold throughput that must be achieved in each scenario, it is clear that in tests like this, we start to enter subjective territory about what is the best way of handling the situation. The highest throughput can always be achieved by handling only the close STAs and ignoring the far STAs but this is obviously not a good solution. How much should far-away STAs be allowed to impact the overall throughput? What is fair? These are questions that different AP vendors may handle differently and what is fit for one application may not be fit for another.

3.2.2. TR-398 Issues 2 and 3

As the TR-398 specification has progressed through Issue 2 (released March 2021⁶) and Issue 3 (not released as of the writing of this paper, but likely to be finalized in late 2023 or early 2024), more complex scenarios have been addressed. For example, Issue 2 includes:

- Dual-band and bi-directional throughput tests
- Channel auto-selection tests
- A first attempt at “roaming” and “repeater” tests

In addition, since the primary Wi-Fi technology in the market at the time of Issue 1 was 802.11ac (Wi-Fi 5), there were no tests in Issue 1 related to 802.11ax (Wi-Fi 6.) So, Issue 2 added in a few new tests to begin to address Wi-Fi 6.

With Issue 3, the group has begun to address the new 6 GHz spectrum, as well as areas such as quality of service (QoS) and latency, and the accuracy of measurements reported by the AP that may be used for radio resource measurement (RRM) algorithms, such as received channel power indicator (RCPI), noise, and channel utilization.

3.2.3. Wi-Fi Alliance Device Metrics

The Wi-Fi Alliance Device Metrics test plan is similar to that of TR-398 but with a noticeably different focus. The test cases are familiar to TR-398:

- Rate vs. range

- AP latency
- Channel switching [for various reasons, like interference or dynamic frequency selection (DFS)]
- Band steering
- Roaming
- AR/VR device performance

As with TR-398, each test case comes with a description of the test configuration and the methodology. The Device Metrics test plan has a heavy focus on how the collected data should be analyzed and presented to the user. A difference from TR-398 is that the Device Metrics test plan does not provide any pass/fail criteria for these tests. The philosophy behind this test plan is even more along the “fit for purpose” concept, the thought being that if people are presented with the results of these tests in a clear and consistent way, they can decide for themselves (based on their own understanding of their market/customer requirements) whether any given device satisfies their requirements. It is the focus on the statistical analysis and consistent data presentation that distinguishes the Device Metrics test plan.

3.2.4. ETSI Multiple Access Point Performance

Another example of a standards body addressing Wi-Fi performance is the ETSI specification (TS 103 754) called “Multiple Access Points Performance Testing.” This specification is an excellent example of the “fit for purpose” concept we have been discussing. This specification, as its name implies, is only related to multiple-AP scenarios (mesh/repeater installations). For a user/operator that is not interested in a multiple AP scenario, there would be no reason to use this test plan. That seems to be a good description of the direction that Wi-Fi performance testing is moving in general: “Performance” will be defined by whatever the group feels is important to the members of that group (and what they perceive to be important to the broader community of interest). Tests will then be defined to give consistent, repeatable, quantitative insights to tests of interest to address those requirements.

For the Multiple Access Points group, those tests include:

- Roaming time and throughput
- One-hop and two-hop throughputs
- Network configuration and self-healing
- Band steering

4. The Future of Wi-Fi Performance Testing

Based on the above description, the reader would be forgiven for concluding that performance testing for Wi-Fi is a solved problem. There are at least three major standards bodies addressing it (Broadband Forum, Wi-Fi Alliance, ETSI), and each group is tackling the problem from a somewhat different direction.

However, Wi-Fi performance testing is anything but a solved problem. In fact, the efforts described above are a description of the early stages of grappling with how to define Wi-Fi performance. Some of the most interesting features of Wi-Fi have not been dealt with by any of these specifications, and more features are being developed all the time. We describe just some of these advanced features below.

4.1. New Wi-Fi Features, and their Area(s) of Application

Before we get into a specific discussion about some of the newer (and more complicated to test) features of Wi-Fi, it is useful to describe *why* these new features were developed in the first place.

When first publicly introduced in 1997, Wi-Fi was considered mostly a “nice-to-have” feature for broadband users, allowing them to connect to the Internet wirelessly “when possible.” Coverage was often limited to specific locations (the living room, the study), and speeds were not high. At that time, the main applications were web browsing, email access, etc. Video streaming, IoT, voice-over-Wi-Fi, and the like were virtually unheard of.

The world has changed dramatically since that time. First of all, Wi-Fi is no longer a “nice-to-have” feature of a broadband service plan. As mentioned above, the vast majority of users are interacting with their broadband service over Wi-Fi as the last link in the connection. And the applications are no longer what they were a quarter-century ago. According to the graphic shown in Figure 9⁷, nearly half of all mobile traffic is streaming video. And many of the other applications, like social networking, much of the online gaming industry, etc., didn’t even exist in 1997.



Figure 9 - The World's Most Used Apps by Downstream Traffic

Wi-Fi has moved from the periphery of the offering to being a critical component, and it needs to be able to deliver a lot more in a variety of areas, like higher throughput, lower latency, better overall coverage, and higher overall capacity. Take, for example, the issue of coverage. In the early days of Wi-Fi, having some wireless connectivity was enough. Now, however, the Internet of Things has become widespread, and many home devices (stoves, refrigerators, thermostats, alarm systems, water shutoff systems, sprinkler systems, electric vehicle charging systems, internet-connected photo frames, etc.) rely on available Wi-Fi coverage, so Wi-Fi has to be available “everywhere” in the home.

To address the flood of new requirements, the Wi-Fi standards bodies have added a host of new features to Wi-Fi to improve performance in these and other areas. In Figure 10, we show just some of these features (in grey) mapped along the four axes of throughput / latency / coverage / capacity, along with some relevant new technologies in the Wi-Fi standards (in yellow.) For example, while video calling is sensitive to both throughput and latency [and, therefore, may take advantage of the new higher data rates as well as the orthogonal frequency-division multiple access (OFDMA) functionality of Wi-Fi 6], something like residential IoT is sensitive mostly to coverage and therefore will be helped most by the

newer mesh capabilities as well as the power save features such as targeted wait time (TWT). This figure is not intended to be exhaustive, and it's certainly possible to argue that some of these features may be useful in other parts of the requirement space. The point of the figure is just to show that with exploding requirements for Wi-Fi networks has come an exploding list of new features to be tested.

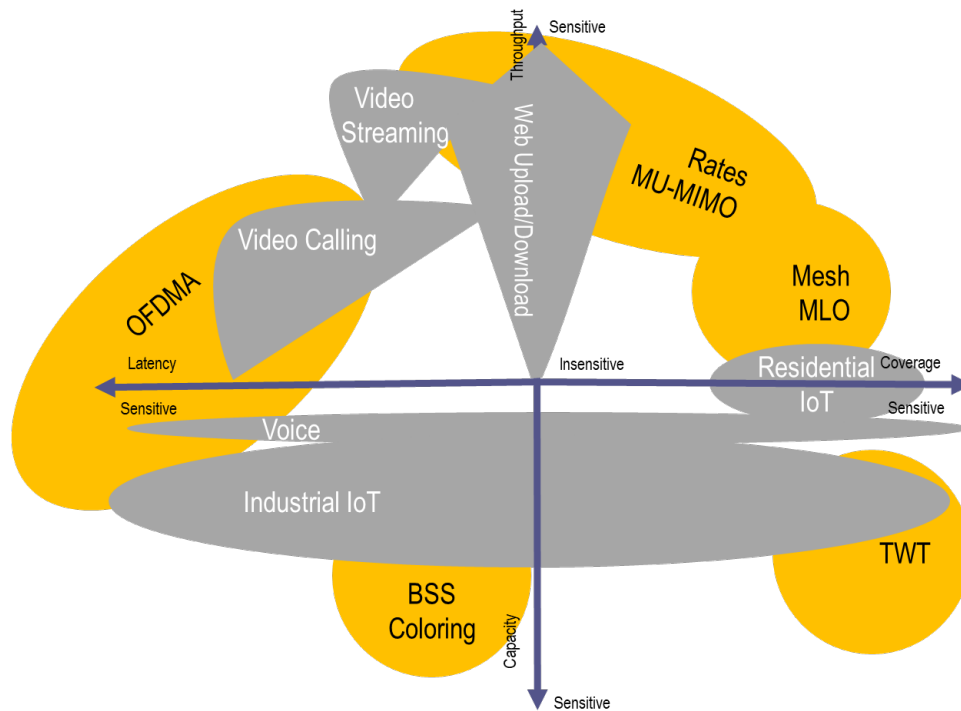


Figure 10 - Some New Wi-Fi Features Mapped with Applications

In the following sections we will discuss just a few of these new features in more detail.

4.2. Multi-User Operation in Wi-Fi 6

The first multi-user Wi-Fi capability was introduced in Wi-Fi 4 with multi-user MIMO, and indeed some of the test plans described above do probe that functionality (e.g., TR-398). Even at that, Wi-Fi supports both downlink and uplink MU-MIMO, and so far, test plans only address the downlink capability.

In addition, one of the most talked-about new elements of Wi-Fi 6 is the addition of OFDMA as a channel access capability. OFDMA is also a multi-user feature, capable of allowing an access point to address and receive from multiple STAs simultaneously. One of the ways that this functionality can be useful is in terms of limiting per-STA latency since STAs would no longer need to wait for the AP to address each STA separately; rather, an AP could address many or all STAs at the same time.

Although, as discussed above, the BBF's TR-398 added in some Wi-Fi 6 testing functionality as part of Issue 2, that functionality was focused solely on the higher throughputs that Wi-Fi 6 offers. Looking at the performance of OFDMA has not been broached, and coming up with an implementable, repeatable scenario that can provide a quantitative measurement of an AP's OFDMA performance is a challenge that is, so far, for further study.

4.3. Advanced Power Save Operation in Wi-Fi 6

Another highly touted feature of Wi-Fi 6 is what is known as Targeted Wake Time (TWT). This feature is an advanced version of Wi-Fi's power save functionality that allows for very flexible sleep times for Wi-Fi devices. None of the performance testing has yet dealt with this functionality; indeed, none of the testing has yet addressed any of the power save functionality of Wi-Fi.

4.4. More Features Arriving in Wi-Fi 7

Wi-Fi 7 is right around the corner⁸, and the number of new features in that specification will require whole new rounds of performance testing updates. For example:

- There are updates to basic modulations and bandwidths, so, higher throughputs.
- OFDMA has been upgraded to support more flexible spectrum allocations, so even though basic OFDMA functionality has not yet been addressed by these performance standards, Wi-Fi 7 will add more functionality to be addressed.
- TWT has been upgraded to provide restricted access to the channel for certain users and, so as with OFDMA, the TWT functionality has been upgraded before the performance testing specifications have even addressed the initial feature.
- One of the most highly anticipated features of Wi-Fi 7 is Multi-Link Operation (MLO), in which an AP/STA pair can communicate using a combination of channels across bands. There are a number of ways in which this functionality is likely to be useful, from increased throughput to interference robustness to low-impact band steering. Defining use cases, test methodologies, and metrics for MLO are going to be major areas of activity for at least some of the performance test specifications.

5. Challenges Related to Advanced Feature Testing

As should be evident from the above discussion, the difficulty with performance testing comes from the definition of "performance." There are some aspects of performance that can be understood fairly simply, and there are some that are much more complex. Compare the questions:

- (1) At high signal strength, what is the maximum throughput I can achieve between this AP and this STA?
- (2) How much more efficient is this product in Wi-Fi 6 mode than it is in Wi-Fi 5 mode?
- (3) What kind of experience does a client have, on a mesh system, while the mesh is supporting some set of other users, under specific traffic loads, with different interference conditions when roaming in a specific way?

The differences between these questions come down to the complexity of the question itself. The first question is fairly simple, and so leads to a fairly simple test: throughput is run between two points and measured. The second question is somewhat more complicated. What do we mean by "efficient?" Wi-Fi 6 is based on the IEEE 802.11ax specification, known as the "high efficiency" enhancement⁹, but what kind of test should be used to define efficiency? What metric should be evaluated for that case?

The final question indicates the level of complexity that these performance tests can elicit. While it is exactly the kind of question that a service provider would care about ("how well will this mesh system perform if I deploy it in the homes of my customers?"), the level of complexity is quite high. What do we

mean by “well?” Is that throughput? Latency? A combination? Is it different for different applications? How do I create a mesh system in a compact testbed? How can I be sure that my emulated mesh system is a reasonable representation of what my customer will experience?

With the complexity of the question comes the complexity of the test itself. Compare the difference in testbed complexity between a testbed required to satisfy the first test...




Figure 11 - A Testbed for Running Traffic between Two Endpoints

... and a testbed that would be needed to answer the question about a mesh system with roaming stations (and possible interference conditions.)




Figure 12 - Example Testbed for a 4-Node Mesh with Multiple STAs

Also, increasing testbed complexity is not the only way in which the performance tests can become more and more complicated. The tests themselves get much more complicated and, as a result, rely more and more heavily on automation. To be clear, the assumption is that all of the tests can be automated (and usually are), but a simple throughput test (for example) can also fairly easily be run manually and therefore require fairly simple automation. The more complex tests, however, move beyond the stage for which manual testing can be reasonably expected to work.

5.1. Automation as a Key Component to Wi-Fi Performance Testing

To describe what we mean by testing complexity, and to illustrate how important automation becomes to enable these tests, let's consider the two more complicated tests described above: the OFDMA test and the mesh test.

5.1.1. OFDMA Performance Testing

The difference between orthogonal frequency division multiplexing (OFDM) and OFDMA is much discussed, but briefly, it is the difference between an AP being able to service only one user at a time in a given channel or multiple users. Figure 13 shows an example:

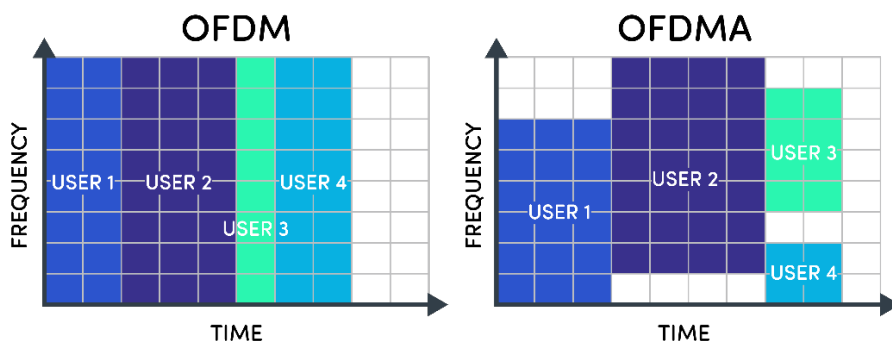


Figure 13 - OFDM vs. OFDMA channel access

Notice that at any given time, in the OFDM case on the left, only a single user has access to any (actually to all) of the frequencies shown on the vertical axis. But in the OFDMA case on the right, there are times when both User 3 and User 4 are being served, with User 3 getting access to some of the frequencies and User 4 getting access to different frequencies.

These combinations of frequencies and times (represented by the small squares in the plots) are known as resource units (RUs) in Wi-Fi. To address the question asked above ("how efficient is the Wi-Fi 6 operation?"), it might be very useful to look at the RU allocations to see how much of the channel is actually occupied. So, for example, we may want to look at how the RUs have been allocated to the STAs.

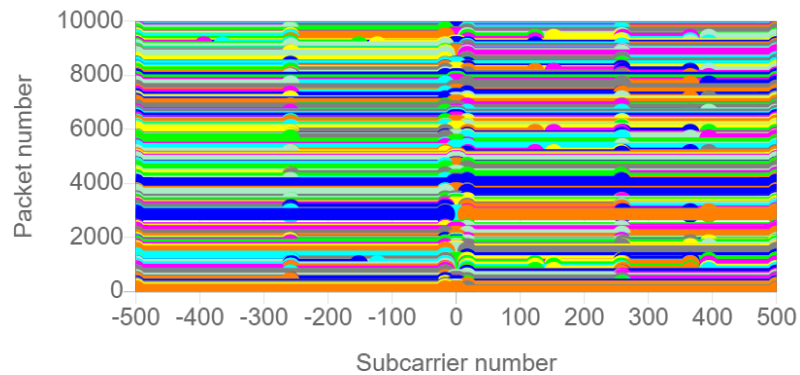


Figure 14 - Example Analysis of RU Allocations in an OFDMA System

We may also care about how efficiently the channel was allocated to the different users. That is, how much of the channel is actually occupied.

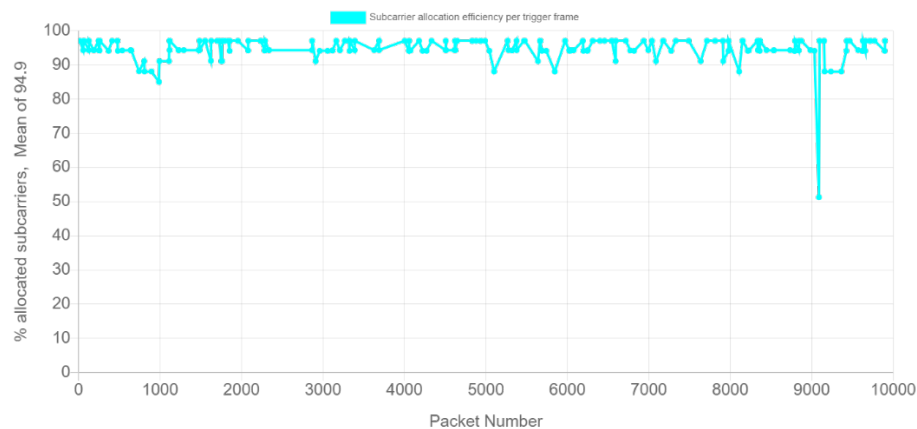


Figure 15 - Example Analysis of Channel Allocation Efficiency in an OFDMA System

Doing OFDMA analyses can be much more complicated than doing throughput measurements. RU allocations are only visible “under the hood” by looking at the actual packet stream and decoding the OFDMA information. More than that, not all OFDMA information is visible to a standard Wi-Fi sniffer. A simple OFDMA setup is shown in Figure 16.




Figure 16 - OFDMA Communication in Wi-Fi, with Sniffing Indicated.

In this figure, we show an AP communicating with four STAs using OFDMA. In this process, the STAs have received an association ID (AID) at association time, and this is used to indicate the RU that it should use on a packet-by-packet basis. That AID is buried in the signaling information and needs to be pulled out to do this kind of analysis. And the OFDMA traffic flowing between the two endpoints will not even be visible to a Wi-Fi sniffer unless the sniffer knows, and has been provisioned with, the AID being used. So, if we want to watch the traffic, we can no longer use a single sniffer. We will need to use four separate sniffers, each with a separate provisioned AID.

The complexity of this test should be obvious:

- Set up the OFDMA traffic
- Figure out the assigned AIDs
- Extract RU information, per AID, on a per-packet basis
- For OFDMA traffic, use separate sniffers with correctly provisioned AIDs to gather that traffic from that specific AID
- Combine all of this information into a coherent result in order to answer the relevant performance question

The sheer level of complexity involved in this test means that without a way of fully automating the test and analysis this kind of test would not be practical or even possible.

5.1.2. Roaming Performance

The roaming performance test is another good example of how the more complicated the test, the more important automation will become.

As discussed above, the ETSI Multiple Access Points test plan focuses on the performance of a system designed for roaming. However, many of those tests are fairly straightforward from a conceptual

perspective. Take, for example, the “two stage networking test.” This test is designed to look at the throughput achieved on a STA/AP link when the STA is separated from the AP by a two-hop mesh network.

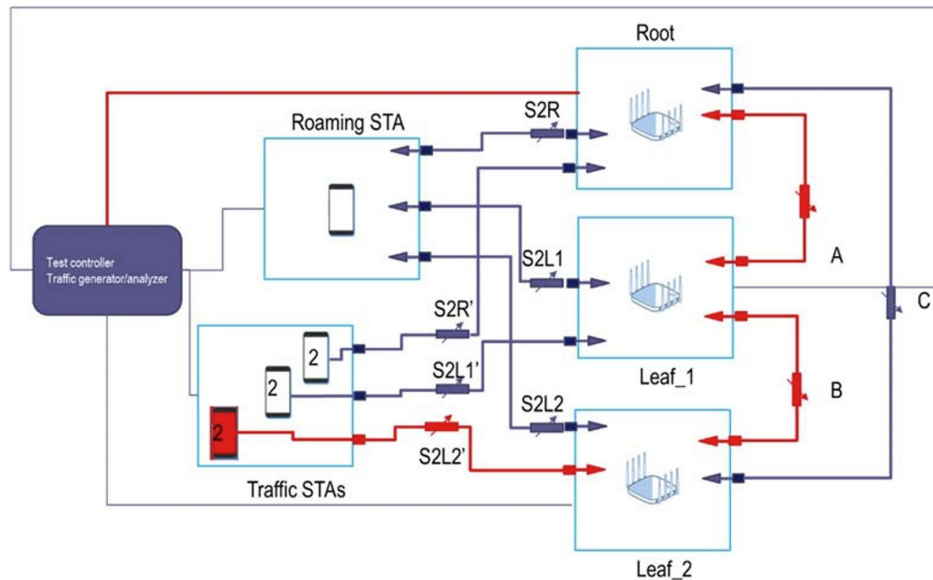


Figure 17 - Multiple Access Points Performance, Two-Stage Networking Test

In this test (see Figure 17), the red elements identify the test configuration. Traffic is run between a STA (the phone icon shown in red) and the root AP, passing first between leaf nodes Leaf_2 and Leaf_1. This is a two-hop mesh throughput test. While this takes some work to configure, it is not difficult to imagine how this could be performed manually.

Now let’s consider the Device Metrics roaming test. This test does not fix the STA statically on one element of the mesh. Rather, it expects the STA to move between the elements of the mesh, thus roaming from node to node. The defined topology is shown in Figure 18 and is simpler than the above:

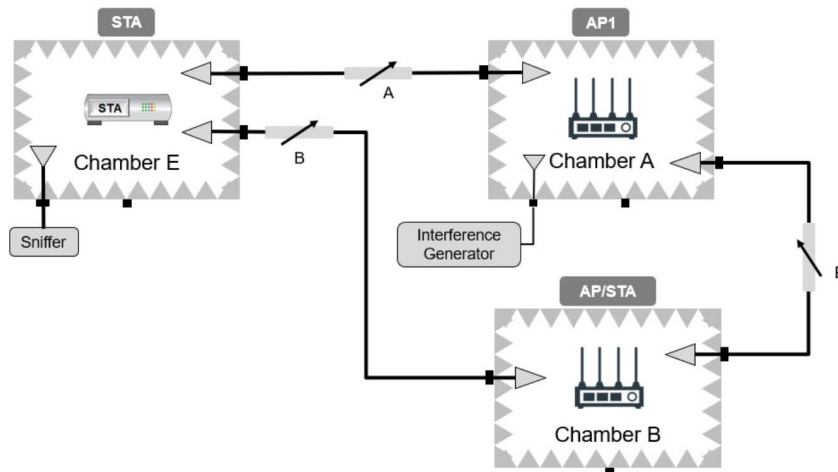


Figure 18 - Wi-Fi Alliance Device Metrics Roaming Topology

However, the test itself is much more complicated. The STA is caused to roam between Chamber A and Chamber B by modifying the programmable attenuators of A and B. As described in the test case, the attenuations used come from real-world measured data, and the recording is then used to play back the walk through the attenuators creating a highly repeatable, lifelike test. Already, a test like that is virtually impossible manually because of the requirement that the attenuators be modified in this highly specific way. The test then also looks at the throughput achieved over multiple runs as a function of time. This data can be different from run to run based on the roaming algorithm and the probabilistic nature of the decisions it makes. An example from the test plan is shown in Figure 19.

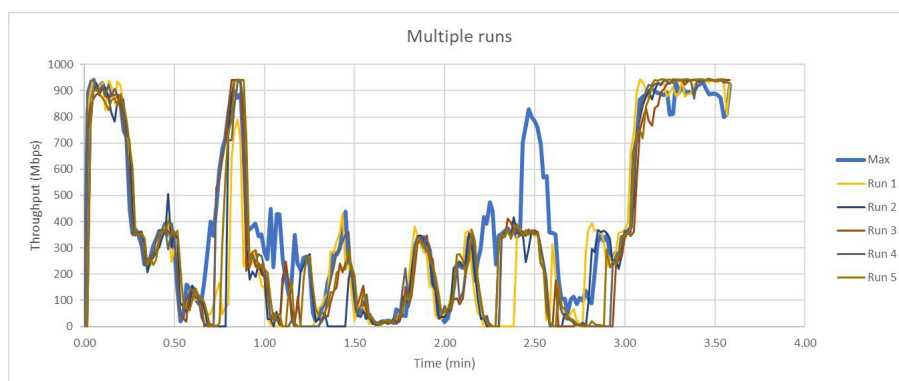


Figure 19: Maximum Roaming Throughput Achieved over a Set of Runs

Adding on the statistical nature of these tests and the resulting analysis required, it becomes clear that without automation, the more useful Wi-Fi performance tests would not be possible.

6. Conclusion

The Wi-Fi community has finally embraced performance testing. Since 2019 three of the most important standards organizations (BBF, Wi-Fi Alliance, ETSI) have developed and published specifications

targeting Wi-Fi performance. It remains, however, early days for Wi-Fi performance tests. Many of the existing tests are simple speed tests, albeit sometimes under more complex conditions.

But Wi-Fi is used in challenging environments under challenging conditions. Almost always used in shared spectrum, Wi-Fi must perform well even when interference exists from other sources (either Wi-Fi or not.) Wi-Fi is the main access network by which people experience their broadband connections, so Wi-Fi must be adept at handling high throughput connections for movies and downloads while simultaneously being able to handle low latency connections for video calls and gaming. Creating test cases to capture these conditions accurately while identifying and measuring the relevant metrics will be the work of these groups, and undoubtedly others, as the need to understand Wi-Fi performance continues.

7. Abbreviations and Definitions

7.1. Abbreviations

AID	association ID
AP	access point
BBF	Broadband Forum
BRAN	broadband radio access network
DFS	dynamic frequency selection
DUT	device under test
ETSI	European Telecommunications Standards Institute
IEEE	Institute of Electrical and Electronics Engineers
IoT	internet of things
MLO	multi-link operation
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
QoS	quality of service
RCPI	received channel power indicator
RRM	radio resource measurement
RU	resource unit
SCTE	Society of Cable Telecommunications Engineers
STA	“station,” a non-AP device in a Wi-Fi network
TWT	target wake time

7.2. Definitions

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

8. Bibliography and References

TR-398 Wi-Fi In-Premises Performance Testing, Issue 1, February 2019
 TR-398 Wi-Fi Residential & SOHO Performance Testing, Issue 2, March 2021
 Wi-Fi Device Metrics, v1.0, 2022
 Broadband Radio Access Networks (BRAN); Multiple Access Points Performance Testing, ETSI TS 103 754 V1.1.1 (2022-06)

-
- ¹ https://www.wi-fi.org/system/files/Wi-Fi_by_the_numbers_2023.jpg
 - ² <https://wifinowglobal.com/news-and-blog/wi-fi-is-the-best-invention-of-the-past-25-years-says-uk-study/>
 - ³ <https://www.mediaplaynews.com/parks-92-percent-of-us-internet-households-use-wi-fi-at-home/>
 - ⁴ https://www.wi-fi.org/system/files/Wi-Fi_Device_Metrics_Highlights_202209.pdf
 - ⁵ https://www.etsi.org/deliver/etsi_ts/103700_103799/103754/01.01.01_60/ts_103754v010101p.pdf
 - ⁶ https://www.broadband-forum.org/technical/download/TR-398_Issue-2.pdf
 - ⁷ <https://www.visualcapitalist.com/the-worlds-most-used-apps-by-downstream-traffic/>
 - ⁸ <https://www.wi-fi.org/who-we-are/current-work-areas#Wi-Fi%207>
 - ⁹ <https://standards.ieee.org/ieee/802.11ax/7180/>

Scaling the DAA Remote PHY CIN Deployment with Cloud Native Routing

A Technical Paper prepared for SCTE by

Hitesh Mali, Architect, Juniper Networks Inc
1133 Innovation Way
Sunnyvale, California, 94089 USA
hmali@juniper.net
(267) 495-6851

Bob Gaydos, Comcast VP Fellow, Comcast
Comcast Technology Center
1800 Arch Street
Philadelphia, PA 19103
Robert_Gaydos@cable.comcast.com
(215) 286-8737

Brian Sherwood, Sr. Systems Engineer, Juniper Networks Inc
1133 Innovation Way
Sunnyvale, California, 94089 USA
bsherwood@juniper.net
(215) 499-3135

Table of Contents

Title	Page Number
Table of Contents	31
1. Introduction	32
2. Overview	32
3. Cloud Native Routing Design	33
4. Benefits of Cloud Native Routing in DAA RPHY CIN Design	36
5. Conclusions	36
6. Abbreviations	36
7. Bibliography and References	37

List of Figures

Title	Page Number
Figure 1 - The Cloud Native vCMTS Remote PHY Architecture - vCMTS Microservices Hosted on Kubernetes Compute and Spine and Leaf CIN Provide the IP Connectivity between vCMTS and RPDs	33
Figure 2 - The Cloud Native Routing Design for vCMTS Microservices Remote PHY Architecture.	34
Figure 3 - Route BGP Update Flow in Cable Modem Online and Offline Events	35
Figure 4 - ARP and NDP Flow	36

1. Introduction

As cable operators continue to seek ways to keep pace with the growing demand for faster internet speeds and higher bandwidth, the deployment of distributed access architectures (DAA) has emerged as a popular solution. Remote PHY (RPHY) is one such DAA architecture that has gained significant traction in recent years. Comcast's vCMTS RPHY solution is deployed as a containerization microservices to improve agility and scalability. The DAA RPHY vCMTS deployment at Comcast is made up of a collection of servers that use Kubernetes to orchestrate the containerized vCMTS microservices. These microservices communicate with a remote physical device (RPD) via a networking layer referred to as the converged interconnect network (CIN), a leaf-spine architecture (see Figure 1).

The converged interconnect network is a vital component of the DAA RPHY architecture that enables the connectivity between the vCMTS microservices and the remote physical device (RPD) in remote PHY architecture. The deployment of a high scale CIN for cloud native vCMTS deployment presents a unique set of challenges, including the need for efficient routing of traffic between the cable modems (CMs) and the vCMTS. To address these challenges, cloud-native routing can be leveraged to provide a scalable, agile, and efficient solution for the cloud-native vCMTS deployment at scale.

This white paper explores the benefits of scaling remote PHY CIN deployment with cloud-native routing, highlighting how it can reduce the routing and switching complexity in CIN and make it more scalable and efficient. We will discuss how cloud-native routing can be leveraged to enhance the performance of the remote PHY CIN, and how it can be used to simplify deployment, reduce operational expenses, and improve network agility.

2. Overview

In cloud native DAA remote PHY deployment the leaf switch pair that connects to the commute server hosting the vCMTS plays a key role in handling traffic between the vCMTS microservice and the external network. A vCMTS container supports a certain number of cable modems. A head end may have many servers, hosting hundreds of vCMTS microservices; these vCMTS microservices host tens of thousands of cable modems. For the leaf to route the traffic to the appropriate vCMTS it needs to be aware of which particular vCMTS microservice is hosting the cable modem (CM).

In a typical deployment the L2 (layer 2) domain is extended to the leaf. This enables the leaf to learn the CMs ARP/NDP (Address Resolution Protocol/ Neighbor Discovery Protocol) as it comes online or goes offline. Address Resolution Protocol is a link layer protocol used for discovering the MAC address CM's for IPv4; Neighbor Discovery Protocol performs functions that are similar to those addressed by the ARP for IPv6. A large L2 domain can create issues, such as broadcast storms, as tens of thousands of devices refresh their ARP tables regularly. This can also burden the leaf switch, as it needs to keep track of tens of thousands of L2 MAC addresses, which can also exceed the capacity of lower-end commodity switches and create a significant amount of L2 traffic.

Also, there are limits on the number of ARP/NDP updates that can be learned on given leaf switches and the number DHCP (Dynamic Host Configuration Protocol) transitions the leaf switches in the converged interconnect network can handle at a given point. Especially in scenarios such as power outages, when all cable modems try to come online simultaneously, the ARP/NDP and the DHCP transitions supported by the leaf switches in the CIN can become a bottleneck. This can set up the limit on the number of subscribers that can be supported by cloud native vCMTS RPHY deployment.

One of the major advantages of cloud native vCMTS RPHY architecture deployment with Kubernetes is it allows us to move the vCMTS containers to different compute hosts without impacting the vCMTS subscriber states, which means there are no impacts to ongoing subscriber traffic flow. This also means that in the event that there is failure of an existing compute host, or there are capacity constraints, the vCMTS container can be moved to a different host without impacting the subscriber. However, with an L2 design the leaf switch would have to relearn the new routing and switching state of the vCMTS and CM to route the traffic to the new host, which can cause disruptions to subscriber traffic flows, including interruptions to sensitive traffic such as voice calls.

In the Comcast deployment this problem is addressed by using cloud native layer 3 control plane. In this approach the CM states are advertised as a layer 3 BGP (border gateway protocol) route to the leaf switch in the CIN using an external cloud native routing container.

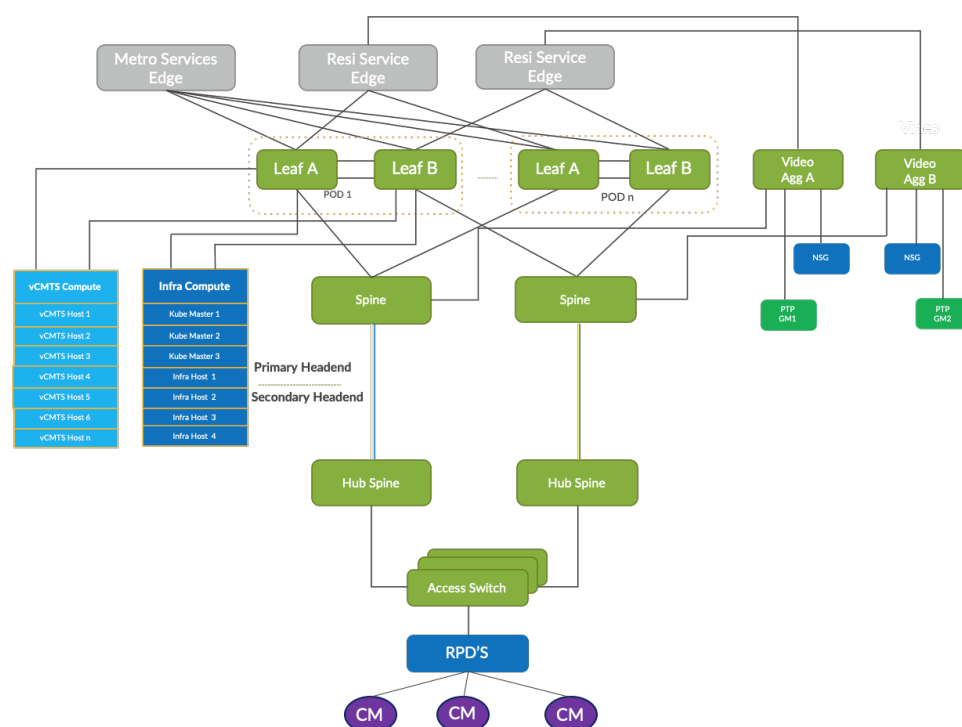


Figure 1 - The Cloud Native vCMTS Remote PHY Architecture - vCMTS Microservices Hosted on Kubernetes Compute and Spine and Leaf CIN Provide the IP Connectivity between vCMTS and RPDs

3. Cloud Native Routing Design

The Fibber application developed by Comcast for the cloud native routing plays a key role within the Comcast cloud native routing design for the distributed access architecture (DAA). It encompasses the cloud native routing stack and API broker, serving as an entry point for external clients to access internal routing APIs and configurations. These APIs facilitate communication between the Fibber application and the DHCP servers responsible for assigning IP addresses to the Cable Modems (CMs). To enhance

control plane scalability and resilience, this architecture uses two BGP route reflectors. In this solution Comcast uses the Juniper containerized routing domain (cRPD) as a BGP route reflector. These cRPDs significantly reduce the number of required BGP sessions on the CIN leaf switches, subsequently simplifying the control plane complexity. The cRPD route reflector BGP peers with the leaf switches and the Fibber application. The Fibber application communicates changes in CM states to the cRPD through BGP updates. Figure 2 illustrates the elements comprising the cloud native routing design for the vCMTS microservices remote PHY architecture.

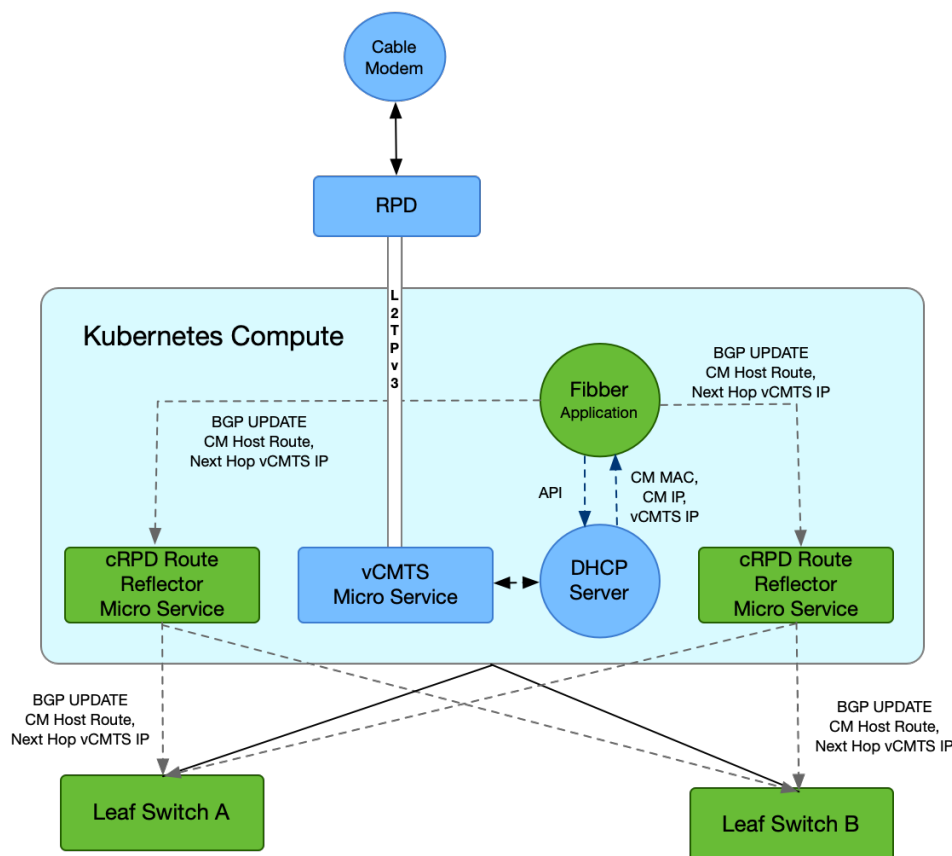


Figure 2 - The Cloud Native Routing Design for vCMTS Microservices Remote PHY Architecture.

As the cable modem is turned on it will use DHCP to obtain an IP address. The DHCP request is sent from the cable modem to the RPD over the DOCSIS network. The RPD encapsulates the DOCSIS request in the L2TPv3 and sends it to the vCMTS microservice. The vCMTS de-encapsulates the packet and relays the request to the DHCP server. Once the DHCP server successfully assigns an IP address to the cable modem, it also sends the notification to the Fibber application using APIs. This notification includes the MAC address of the cable modem, the IP address assigned to the cable modem, and the IP address of the vCMTS that hosts the particular cable modem. The Fibber application then sends the IP address of the CM and the IP address of the vCMTS information as a BGP update, with the next-hop as a vCMTS container. This update is relayed as a BGP route to the cRPD Route Reflector, which further relays it to the leaf switches as a BGP route. The leaf now has correct routing state in the routing information base

(RIB) to route the traffic to the vCMTS microservice for the CM that it is serving to. Figure 3 outlines the BGP routing process for cable modems in both online and offline scenarios.

When the vCMTS container shifts from one host to another — whether due to the failure of the compute host or for capacity reasons — the cloud-native control plane design ensures that massive routing/switching updates are not necessary. In this scenario, the change will be accommodated by a single route update of the vCMTS IP address to a different MAC address. This streamlined approach simplifies the transition process, minimizing disruption and complexity.

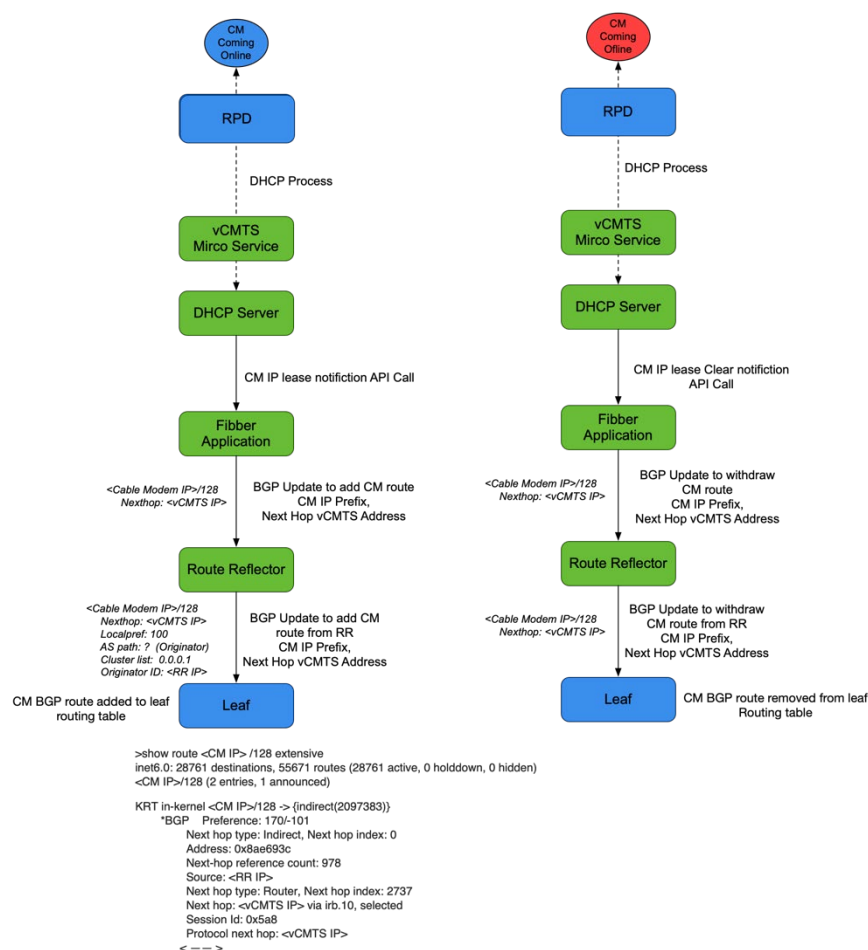


Figure 3 - Route BGP Update Flow in Cable Modem Online and Offline Events

For data plane traffic the vCMTS container is a gateway, but the vCMTS container also works as a proxy gateway for the CM ARP/NDP requests, which also are answered by the vCMTS. When the CM sends ARP to its gateway the vCMTS intercepts the CM ARP and sends its own MAC address in response since it is the upstream gateway. For a visual representation of this process, please refer to Figure 4, which outlines the sequence of ARP interactions between the cable modem and the vCMTS.




Figure 4 - ARP and NDP Flow

4. Benefits of Cloud Native Routing in DAA RPHY CIN Design

The deployment of DAA CIN with cloud native routing provides several benefits to cable operators, including:

- **Improved network performance:** DAA CIN deployment with cloud native routing improves network performance by reducing switching/routing complexity in the converged interconnect network
- **Increased Scalability:** Increased scalability from DAA CIN implementation with cloud native routing enables more subscribers to be supported in cloud native vCMTS deployments.
- **Cost Savings:** DAA CIN deployment with cloud native routing reduces the need for high end leaf routing equipment for CIN, resulting in cost savings for cable operators.

5. Conclusions

The deployment of vCMTS microservices with cloud-native routing through advanced networking architecture provides several benefits: it improves network performance, simplifies network operations, provides increased scalability, and reduces costs. The scalability of the vCMTS remote PHY deployment is not limited by the layer 2 ARP/NDP scale of the CIN leaf switches thanks to the CIN design's incorporation of cloud native routing at the layer 3 routing level.

6. Abbreviations

vCMTS	virtual cable modem termination system
DAA	distributed access architecture
cRPD	containerized routing protocol daemon
CIN	converged interconnected network
RR	route reflector
CNR	cloud native routing
CM	cable modem
L2TPv3	layer 2 tunnelling protocol version 3
API	application programming interface
DHCP	dynamic host configuration protocol
APR	address resolution protocol

NDP	neighbor discovery protocol
BGP	border gateway protocol
Fibber	Fibber Routing Application developed by Comcast for cloud native routing.

7. Bibliography and References

Cable Labs Remote PHY Specification CM-SP-R-PHY-I17-220531

BGP Route Reflection RFC4456 <https://datatracker.ietf.org/doc/rfc4456/>

The Kubernetes <https://kubernetes.io/>

Juniper cRPD datasheet

Automatic Real-time Detection of Congested Multimedia Networks

A Technical Paper prepared for SCTE by

Mun Gi Choi, Principal Engineer, Charter Communication Inc, SCTE Member
6399 S Fiddlers Green Cir,
Greenwood Village, CO 80111
mun.choi@charter.com
949-510-4796

Hany Heikal, Senior Director, Charter Communication Inc, SCTE Member
6399 S Fiddlers Green Cir,
Greenwood Village, CO 80111
hany.heikal@charter.com

Manish Jindal, General Vice President, Charter Communication Inc, SCTE Member
6399 S Fiddlers Green Cir,
Greenwood Village, CO 80111
manish.jindal@charter.com

Table of Contents

Title	Page Number
Table of Contents	39
1. Introduction	41
2. Visual Quality Analysis of Video Communications in Congested DOCSIS Networks	41
2.1. Congestion Simulation Network of Video Conferencing Application Metrics	41
2.2. Congestion Simulations	42
3. Detection Algorithm of Frame Repeats	44
3.1. Average of Absolute Frame Difference	44
3.2. Frame Repeat Index Estimation	45
4. Temporal Congestion Percentage Estimation	46
4.1. Temporal Congestion Detection Window Setup	46
4.2. Temporal Congestion Percentage Estimation	46
5. VQM Estimation	46
5.1. Video Quality Evaluation Guideline	46
5.2. VQM algorithms on Congested Network	47
5.3. VQM Test Results on Congested Network	48
6. Motion Adaptive VQM	51
6.1. Motion Adaptive VQM algorithm on Congested Network	51
6.2. Motion Adaptive VQM Test Results	51
7. Network Traffic Monitoring	52
7.1. Bottom-Up Network Traffic Monitoring and Notifications	52
7.2. Root Cause Analysis	53
8. Conclusions	53
9. Abbreviations and Definitions	54
9.1. Abbreviations	54
9.2. Definitions	54
10. Bibliography and References	54

List of Figures

Title	Page Number
Figure 1 - Congestion Simulation of Video Conferencing Application over DOCSIS network	5
Figure 2 - Received Frames of 80% Network Congestion (Frame Detection = 1: Received Frame, Frame Detection = 0: Frame Drop)	6
Figure 3 - Received Frames of 90% Network Congestion (Frame Detection = 1: Received Frame, Frame Detection = 0: Frame Drop)	6
Figure 4 - Received Frame Rates of Network Congestion Percentage Levels	7
Figure 5 - Frame Repeat Index Estimation Scheme with Binary Format	8
Figure 6 - Subjective MOS vs. Objective VQM scores	10
Figure 7 - Plots of VIF Scores of the Congestion Percentage Levels	13

Figure 8 - Plots of VMAF Scores of the Congestion Percentage Levels _____	13
Figure 9 - Motion Adaptive VIF Scores of Different Congestion Percentage Levels _____	15
Figure 10 - Bottom-Up Network Traffic Monitoring over DOCSIS Network _____	16

List of Tables

Title	Page Number
Table 1 - Bad, Marginal, Good Quality Guideline of VQM Metric Scores _____	10
Table 2 – VQM Test Results of Webex Video Conferencing _____	12

1. Introduction

Network traffic monitoring is essential to optimizing network performance. While congestion monitoring tools such as SNMP, NetFlow, IPFIX, jFlow, sFlow, and NetStream can be useful, it's important to consider the differences to select the best option because these tools each have their own pros and cons. For example, flow-based monitoring relies on flow data received from switches, routers, and hosts during a specific period, and it uses more bandwidth. A flow analyzer evaluates the flow records received and stored by a flow collector, and it provides bandwidth usage, traffic patterns, and other performance metrics. However, it is difficult for these heavy monitoring systems to monitor/analyze time-sensitive media traffic such as video conferencing, video calling, online gaming, cloud gaming, and live video streaming for individual clients in a real-time manner. This paper proposes new lightweight and real-time congestion monitoring frameworks.

When network congestion occurs because of insufficient bandwidth, some video packets can be lost, and they cause the video frame drops at the decoder. The current well-known video conferencing applications, such as Webex, Zoom, Skype, etc., repeat the previously decoded frames to compensate for the dropped frames to preserve the spatial video quality. In these scenarios, the temporal video quality can be degraded significantly although the spatial video quality is maintained similarly.

In this paper a new detection algorithm of frame repeats in the network congestion scenarios is proposed, and the percentage of temporal congestion is estimated to find the network congestion level. In addition, the adaptive temporal video quality measurement algorithm is also proposed to find the degree of degradation of temporal video quality because the static video quality measurement cannot reflect the real video quality of video conferencing applications.

2. Visual Quality Analysis of Video Communications in Congested DOCSIS Networks

When the linked networks of video communication applications have congestion caused by heavy traffic, some video packets can be discarded and the video frames that include the lost packets are dropped. The lost frames at the decoder have a greater impact on the quality of the video streams.

2.1. Congestion Simulation Network of Video Conferencing Application Metrics

For this paper, channel impairment tests were performed over DOCSIS 3.1 network to find the cable modem and CMTS saturation effects on the video conferencing applications. The interactive video communication using Cisco Webex video conferencing was held from one video conferencing client to another client. Since the reference frame capture for VQM is not possible in Webex video conferencing, the content share option of Webex was used to send the video from one client to another client. Webex Content Sharing extension is an enhancement for Webex video conferencing. The shared 1080P video content was transcoded by using H.264 at the Webex server.

For the network congestion effect in the video communication applications, Webex transcoded the 1080p resolution and 10fps frame rate although the input resolution and frame rate were 1080p and 30fps, respectively. The simulation network is described in Figure 1.




Figure 1 - Congestion Simulation of Video Conferencing Application over DOCSIS Network

2.2. Congestion Simulations


During the video conference, the maximum available up-stream bandwidth was set at 10Mbps, and an iPerf tool was used to control the network congestion levels such as 1%, 50%, 60%, 70%, 80%, 90% and 100%. This means that at the approximately available bandwidths of 1%, 50%, 60%, 70%, 80%, 90% and 100% the congestion levels were 10Mbps, 5Mbps, 4Mbps, 3Mbps, 2Mbps, 1Mbps and 0.1Mbps, respectively.

Our tests showed that the received frame rates are varied and they are dependent on the congestion levels of up-stream. As the congestion level increases, the received frame rate decreases. The channel congestion significantly reduces the received frame rates at the 80%, 90% and 100% congestions. This causes serious motion jerkiness at the decoder of Webex client and lip-sync problems if the audio channel is added.


The Webex server receives the clients' encoded bit streams and transcodes the received encoded bit streams by using H.264 Scalable Video Coding (SVC) [8]. The received video is the enhancement layer video, and the participant image is the base layer in our test cases. SVC-based conferencing server transcoding uses the spatial scalability and temporal scalability of SVC:

- Spatial scalability: base layer and enhancement layer.
- Temporal scalability: base layer (variable frame rate) and enhancement layer (constant frame rate).

For the 600 input frames, the video conference participant should receive 200 frames if there are no frame drops during the video conference because the video conferencing application transcodes the frame rate from 30fps to 10fps. However, the received frames at the participant are much less than 200 frames because of the network congestion. Figures 2 and 3 show the received frames for 80% and 90% network congest levels, respectively. Furthermore, the frame rate of received video is proportionally dropped as the network congestion level increases and the available bandwidth is insufficient to deliver good quality of video. Figure 4 shows the received frame rates of network congestion percentage levels for the video conference.



**Figure 2 - Received Frames of 80% Network Congestion (Frame Detection = 1:
Received Frame, Frame Detection = 0: Frame Drop)**



**Figure 3 - Received Frames of 90% Network Congestion (Frame Detection = 1:
Received Frame, Frame Detection = 0: Frame Drop)**




Figure 4 - Received Frame Rates of Network Congestion Percentage Levels

3. Detection Algorithm of Frame Repeats

Due to the increase of network latency, jitter also increases at the video decoder. If jitter is bigger than the threshold of jitter, then the video decoder drops the frame. This is an especially inevitable problem in interactive communications such as video conferencing. Thus, lost frame detection is one of essential tasks for the VQM analysis.

3.1. Average of Absolute Frame Difference

Let's define Webex decoded frames $\{f_1, f_2, \dots, f_L\}$ and the resolution of $f_k, k = 1, 2, \dots, L$ as equal to $H \times W$. The average of absolute frame difference of the k^{th} frame, $AAFD_k$, is defined as

$$AAFD_{k-1,k} = \frac{1}{HW} \sum_{y=1}^H \sum_{x=1}^W |f_k(x, y) - f_{k-1}(x, y)|$$

Between frame k and frame $k - 1$, if $AAFD_{k-1,k}$ is less than or equal to the frame repeat threshold (TH_{fr}) that is much smaller than normal absolute frame difference or the scene change threshold, the k^{th} frame is defined as a repeated frame.

If $AAFD_{k-1,k} \leq TH_{fr}$, then the frame repeat is occurred at k^{th} frame.

$$\{f_1, f_2, \dots, f_k, f_k, \dots, f_k, f_n, f_{n+1}, \dots, f_{L-1}, f_L\}$$

$$\text{Total repeated frames} = n - k$$

Else No frame repeat

This frame repeat can randomly occur at one time or consecutively, thus the total number of frame repeats should be estimated when the frame repeats occur. These are key findings and we propose a new binary based estimation algorithm of frame repeats as consecutive frame repeats are more objectionable to the end user.

3.2. Frame Repeat Index Estimation

Let's define the Frame Repeat Index as $FRI_{k-1,k}$ between frame k and frame $k - 1$, and $FRI_{k-1,k}$ estimation algorithm is described as follows:

If $AAFD_{k-1,k} \leq TH_{fr}$, then $FRI_{k-1,k} = 1$

Else $FRI_{k-1,k} = 0$

This algorithm is depicted in Figure 5:




Figure 5 - Frame Repeat Index Estimation Scheme with Binary Format

The pseudocode of the Frame Repeat Index estimation algorithm is described in the following:

$SFR = 0$; //sum of frame repeats

While ($k > 1$ && $k < L$)

 If $FRI_{k-1,k} == 0$;

 return //The media client does not need to send any data to the regional data center.

 Else if $FRI_{k-1,k} == 1$;

$SFR += 1$;

$k += 1$;

The frame repeats occur randomly, and this event is transmitted with the sum of frame repeat to the network monitoring server such as regional data center.

4. Temporal Congestion Percentage Estimation

Monitoring packet loss can help identify areas of congestion and troubleshoot issues with network performance. In addition, latency measures the time it takes for data to travel from one Webex client to another on the network. High latency can indicate congestion or other issues that are causing delays in data transmission. Due to the network congestion, the frame drops caused by packet losses occur. By using the estimated number of dropped frames in the congestion detection window, the congestion percentage is estimated.

4.1. Temporal Congestion Detection Window Setup

The frame rate (i.e., frames per second) N is used for real-time congestion detection to set up the congestion detection window. For example, the congestion detection window size for the video conferencing applications is set at 30 frames because the frame rate of most video conferencing applications is 30 fps. Frame rates for gaming applications are set at 60 frames or 120. In real applications, the detection window size can be adaptively changed.

4.2. Temporal Congestion Percentage Estimation

Let's define N frames as the congestion detection window size, and this N frames depend on the frame rate of media applications. The Temporal Congestion Percentage (TCP) is estimated by the following formula:

$$TCP = \frac{\text{Repeated (or dropped) frame numbers in congestion detection window}}{N} \times 100$$

5. VQM Estimation

Current video conferencing servers and client apps such as Webex, Zoom, Skype, etc., repeat the previous frame when the current frame is lost during network congestion. Then the full-reference objective VQM [6] metrics at video conference clients don't match with the actual video quality because of the frame repeats. This paper proposes the new methodology and the supporting algorithms to resolve these problems.

5.1. Video Quality Evaluation Guideline

Subjective ratings were based on a five-point impairment scale with respondents' opinions mapped to the values 1 to 5, where 5 (Excellent) means impairments are imperceptible, 4 (Good) means impairments are noticeable, but not annoying, 3 (Fair) means impairments are slightly annoying, 2 (Poor) means impairments are annoying, and 1 (Bad) means impairments are very annoying. Then subjective ratings for each test video were averaged into the mean opinion score (MOS) [5].

The relationships between MOS and the full-referenced VQM, Video Multimethod Assessment Fusion (VMAF) [2], are presented in Figure 6. The VQM metrics of VMAF, Visual Information Fidelity (VIF) [1], Structural Similarity (SSIM) [3] and Peak Signal-to-Noise Ratio (PSNR) are estimated by using the

reference video and the captured video. These VQM metrics are used to evaluate the congestion effects on the Webex video conferencing.

Three quality levels such as Bad Quality, Marginal Quality and Good Quality are defined by the ranges of the estimated VQM metric values. Table 1 shows the ranges of Bad Quality, Marginal Quality and Good Quality for VMAF, VIF, SSIM and PSNR [5].

Table 1 - Bad, Marginal, Good quality guideline of VQM metric scores

VQM Metrics	Bad Quality	Marginal Quality	Good Quality
VMAF	$VMAF \leq 50$	$50 < VMAF \leq 80$	$VMAF > 80$
VIF	$VIF \leq 0.4$	$0.4 < VIF \leq 0.6$	$VIF > 0.6$
SSIM	$SSIM \leq 78$	$78 < SSIM \leq 88$	$SSIM > 88$
PSNR	$PSNR \leq 30$	$30 < PSNR \leq 35$	$PSNR > 35$




Figure 6 - Subjective MOS vs. Objective VQM Scores

5.2. VQM algorithms on Congested Network

Specifically, for a received video stream, assume that $f = (f_1, f_2, \dots, f_L)$ is the video input sequence that has L frames and these L frames are the reference frames of full reference video quality measurement. Let f_i define the i^{th} frame and L frames are sent to the video conferencing server, but the receiver at one participant receives only M frames because of the frame drops (i.e., $L > M$) caused by the video conferencing application's network congestions, and the received video is defined by $g = (g_1, g_2, \dots, g_M)$.

Since the frame drops are randomly occurring, the reference frame f_i is not matched with the reference frame g_i . Therefore, each g_i finds the best matched reference frame \hat{f}_i among (f_1, f_2, \dots, f_L) using the full search method. The best matched reference video is defined by $\hat{f} = (\hat{f}_1, \hat{f}_2, \dots, \hat{f}_M)$. Let VQM between \hat{f}_i and g_i define $VQM_i = VQM(\hat{f}_i, g_i)$. The average of VQM value is given by:

$$VQM(\hat{f}, g) = \frac{1}{M} \sum_{i=1}^M VQM(\hat{f}_i, g_i)$$

When network congestion occurs for any participants, the clients could save bandwidth by receiving fewer enhancement layers, gradually degrading the video quality while maintaining the best video experience and dynamically adjusting the quality to changing conditions of the network or the participant's computer. By decoding the base layer and only the subsequent enhancement layers required, a decoder can produce a video stream with certain desired characteristics.

Webex transcoding uses the spatial scalability and temporal scalability of SVC:

- Spatial scalability:
 - Base layer: 720P
 - Enhancement layer: 1080P
- Temporal scalability:
 - Base layer: variable frame rate (for example of 0% congestion; min: 0.853fps, max: 62.5fps)
 - Enhancement layer: constant frame rate (~10fps) and this frame rate is decreased by the network congestion levels.

5.3. VQM Test Results on Congested Network

In the test, congestion simulations were performed over the DOCSIS network, and the received frame rate, VMAF score, VIF score and SSIM score were estimated. VQM metric estimated results of Webex Video Conferencing for each temporal congestion percentage are shown in Table 2. All metrics show that the video quality of received Webex conferencing video is in the marginal quality range (acceptable video quality in general) based on Netflix and Elecard video quality evaluation criteria [2][5]. The estimated VIF, VMAF, SSIM, PSNR metrics in Table 2 don't reflect the network congestion levels at all.

Table 2 – VQM test results of Webex video conferencing

Temporal Congestion Percentage (%)	Received Resolution	Received Frame Rate (fps)	VMAF Score	VIF Score	SSIM Score
0	1920x1080	9.909	68.1023	0.52712	0.92116
50	1920x1080	9.804	67.8556	0.524477	0.915945
60	1920x1080	9.673	67.6664	0.52302	0.914949
70	1920x1080	9.418	68.0504	0.525282	0.916108
80	1920x1080	6.581	68.2771	0.525541	0.916764
90	1920x1080	2.474	71.2005	0.526827	0.915345
100	1920x1080	1.641	71.6933	0.520621	0.917927

Since the differences of VMAF scores are negligible as the temporal congestion percentage levels increase, and the differences of VIF scores and SSIM scores are similar with VMAF, the spatial VQM scores can't be matched with the actual visual qualities of video conferencing applications in the congested networks although the spatial VQM scores are used in the non-congested networks.

VIF scores and VMAF scores with temporal congestion percentage levels are shown in Figure 7 and Figure 8, respectively. The differences of VIF scores are negligible, and the differences of VIF scores are also negligible. These show that the spatial VQM scores can't reflect the actual visual qualities. In the real time video quality evaluation, there is no motion jerk in the 0% congestion, but there is serious motion jerk in the 100%. Therefore, the temporal VQM methodology should be applied to the actual visual quality measurements in the congested networks, and this paper proposes a new VQM formula with the reflection of motion jerk due to the frame drops.




Figure 7 - Plots of VIF Scores of the Congestion Percentage Levels




Figure 8 - Plots of VMAF Scores of the Congestion Percentage Levels

6. Motion Adaptive VQM

A new proposed VQM formula is based on the motion adaptivity to reduce the motion jerk due to the frame drops.

6.1. Motion Adaptive VQM algorithm on Congested Network

In a full-reference VQM [6], an original undistorted high-quality video is compared to a degraded version of the same video, for example, pixel by pixel. During interactive communication such as chatting, the full-reference VQM is not possible because the reference video can't be captured. However, video content sharing can achieve the full-reference VQM of video conferencing.

Assuming that frame drops are randomly occurring during video stream transmission, a spatial-temporal VQM associated with a video conference video stream can be determined as follows:

$$VQM_{MA}(\hat{f}, g) = \sqrt{\frac{\sum_{i=1}^M VQM_i^2(\hat{f}_i, g_i)}{L}}$$

where the reference video \hat{f} , the decode video g , and VQM_i is the VQM value of i^{th} frame, and M is the actual received frame numbers at the video conference participant and L is the frame numbers after video conference server transcoding. The motion adaptive VQM methodology provides a mechanism for determining a spatial-temporal VQM associated with a video stream. It is noted that a spatial-temporal VQM according to various embodiments may be generated using other image quality assessment mechanisms alone or in any combination, as VIF, VMAF, SSIM, PSNR and the like.

6.2. Motion Adaptive VQM Test Results

Using the motion adaptive VIF formula, the motion adaptive VIF scores are plotted in Figure 9. If the motion adaptive VIF scores in Figure 9 are compared with the received frame rates in Figure 4, the motion adaptive VIF scores reflect properly the frame drops caused by the network congestion levels. In addition, the motion adaptive VIF score variations are similar with the variations of received frame rates. However, if the original VIF scores in Figure 7 are compared with the received frame rates in Figure 4, the VIF scores are not changed although there are lots of frame drops.

The motion adaptive VQM information for temporal video quality along with congestion detection, percentage of congestion is transmitted from the media application devices to the regional data center for network monitoring and customer care.




Figure 9 - Motion Adaptive VIF Scores of Different Congestion Percentage Levels

7. Network Traffic Monitoring

Unlike top-bottom network traffic monitoring approaches, this paper proposes the bottom-up network traffic monitoring approach.

7.1. Bottom-Up Network Traffic Monitoring and Notifications

Bottom-up network traffic monitoring over DOCSIS network is described in Figure 10. The information of congestion detection, percentage of congestion, and temporal video quality is sent to the network monitoring system when the network congestion is detected. Unlike flow-based network monitoring, those are much less bandwidth burdens. Whenever frame drops occur at media devices, the congestion information is automatically detected and estimated. Then it is sent to the regional data center for real-time customer care such as root cause analysis and notification.

Network administrators can notify the media application users through SMS and/or email whenever an alarm occurs. Administrators can also configure the operations manager to run external programs or network monitoring scripts automatically when an alarm is triggered. Operations managers not only perform intelligent event processing, but they also correlates raw network events and present only meaningful alarms to the user.

The regional data center also sends the alarm events to the national data center. All collected malfunction events at the national data center can be used for long term care solutions to media clients.

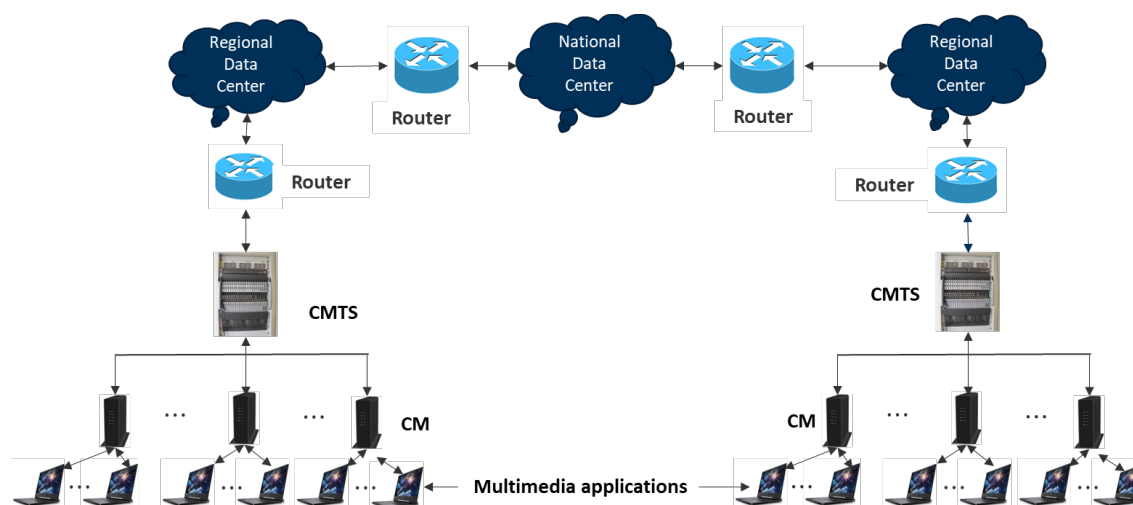


Figure 10 - Bottom-up network traffic monitoring over DOCSIS network

7.2. Root Cause Analysis

Root cause analysis (RCA) leverages the auto-discovered relationships among our monitored resources, as estimated by the proposed algorithms, to determine the root cause of an incident that is impacting dependent resources. Rapid RCA dramatically condenses the time it takes to resolve incidents/outages. Network administrators at regional data centers receive information of congestion detection, percentage of congestion, and temporal video quality and get end-user perspectives of network and application performance, which helps them detect easily and resolve issues as quickly as possible.

RCA allows admins to test how a network performs under a certain load so network administrators can better understand their network's events and see whether they need to better allocate resources or not. RCA aids admins identifying sudden spikes in traffic levels to reduce device downtime.

8. Conclusions

Bottom-up network traffic monitoring methodology over DOCSIS networks is described in this paper. A new detection algorithm is proposed to detect the repeated frames, and it can calculate the total number of repeated frames used to estimate the percentage of congestion. The VQM analysis of congested networks found that the spatial VQM algorithms cannot be directly applied to evaluate the congested real-time media networks. Thus a motion adaptive temporal VQM algorithm is proposed to find the degree of degradation of temporal video quality.

The information of congestion detection, percentage of congestion, and temporal video quality is sent to the network monitoring system when network congestion is detected. Unlike flow-based network monitoring, those are much less burdensome on bandwidth. The congestion information transactions and the roles of Regional Data Center and National Data Center are described for whenever frame drops occur at media devices.

9. Abbreviations and Definitions

9.1. Abbreviations

CM	cable modem
CMTS	cable modem termination system
DOCSIS	Data Over Cable Service Interface Specification
fps	frame per second
Mbps	megabits per second
PSNR	peak Signal-to-Noise Ratio
RCA	root cause analysis
SCTE	Society of Cable Telecommunications Engineers
SNMP	simple network management protocol
SSIM	structural similarity index measure
SVC	scalable video coding
VIF	visual information fidelity
VQM	video quality metric
VMAF	Video Multimethod Assessment Fusion

9.2. Definitions

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

10. Bibliography and References

- [1] H.R. Sheikh and A.C. Bovik, "Image information and visual quality," IEEE Transactions on Image Processing, vol. 15, no. 2, pp. 430-444, February 2006.
- [2] Zhi Li, Christos Bampis, Julie Novak, Anne Aaron, Kyle Swanson, Anush Moorthy and Jan De Cock, "VMAF: The Journey Continues", Oct. 25, 2018.
- [3] Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," IEEE Transactions on Image Processing, vol. 13, no. 4, pp. 600–612, April 2004.
- [4] Cisco Webex, "Cisco WebEx Network Bandwidth", Cisco Webex White Paper, 2015.
- [5] www.elecard.com, "Interpretation of objective video quality metrics", October 26, 2022.
- [6] Oleg I. Ieremeieva, Vladimir V. Lukina, Nikolay N. Ponomarenkoa, Karen O. Egiazarianb, Jaakko Astolab, "Combined full-reference image visual quality metrics", 2016 Society for Imaging Science and Technology, DOI: 10.2352/ISSN.2470-1173.2016.15.IPAS-180.
- [7] N. Ponomarenko, F. Silvestri, K. Egiazarian, M. Carli, J. Astola, and V. Lukin, "On between-coefficient contrast masking of DCT basis functions," in Proceedings of the Third International Workshop on Video Processing and Quality Metrics for Consumer Electronics, January 2007.
- [8] "ITU-T Recommendation H.264 (10/2016)", Annex G. ITU. October 14, 2016. Retrieved June 14, 2017

Climate Control Automation for Critical Facilities

Supporting an “Always On” Network Through Automation

A Technical Paper prepared for SCTE by

John Dolan, FCTO – Past chair

Arnold Murphy, President, Strategic Clean Technology Inc (SCTi)

FCTO – Vice-Chair

1057 Frei Street

Cobourg, ON

K9A 5G5

a.murphy@sct-inc.com

613-558-4415

Ken Nickel, EVP -Quest Controls - SCTE Member

208 9th Street Dr. West

Palmetto, FL 34221

knickel@questcontrols.com

775-409-4312

Tom Hurley, Principal Mechanical Engineer, Comcast – SCTE Member

55 Executive Drive

Hudson, NH 03052

thomas.hurley@comcast.com

603-481-0909

Table of Contents

Title	Page Number
Table of Contents	56
1. Introduction	57
2. The Role Automation Can Play in Critical Facilities	58
2.1. Resilience	58
2.2. Efficiency	58
2.3. Capacity Management	59
3. Benefits of Automation	60
3.1. Site Uptime	60
3.2. Site Operation Efficiency and Reporting	60
3.3. Reduce Costs	60
3.4. Better Use of Resources	61
4. What Levels of Automation Should be Deployed?	61
4.1. Level 1 Automation – Onsite Standalone	62
4.2. Level 2 – Onsite Complex	62
4.3. Level 3 - Central Monitoring and Control Automation	62
5. Best Practices for Choosing and Implementing Automation	63
5.1. Define Clear Objectives	63
5.2. Start with a Proof of Concept (PoC)	63
5.3. Invest in Training and Skill Development	63
5.4. Choose the Right Automation Tools and Platforms	63
5.5. Implement Robust Security Measures	63
5.6. Continuously Monitor and Optimize	63
6. Automation in Practice	63
6.1. Level 1 Use Case: Rack Inlet Temperature Data Collection	63
6.2. Level 2 Use Case: HVAC Device Data and Control	64
7. The Path Forward	65
8. Abbreviations and Definitions	66
8.1. Abbreviations	66
8.2. Definitions	66

List of Figures

Title	Page Number
Figure 1 - SCTE Site Classification	62

1. Introduction

Three major events related to data center operations are taking shape. First is the world's increasing appetite for data, with the result that everyone expects an "always on" state of communications to retrieve data, run their businesses, respond to social media, enable remote working, etc. The evolving and ever-expanding realm of IoT (Internet of Things) will drive greater dependence on the availability of digital communications services.

The second world event that has affecting data center cooling operations in many areas is the warming climate trends being experienced in many parts the world, including North America. For critical facilities the biggest impacts are the unusually high temperatures being experienced for a few days or weeks. Summer of 2022 saw Amazon Web Services experience a major power outage due to a thermal event in the United Kingdom. Within a week of this Google and Oracle services were dropping off-line due to cooling issues. In Sacramento, Twitter experienced an outage due to an unprecedented heat wave.

1.1 Why Critical Facility Automation?

As reliance on digital communications increases, new methods of managing critical facilities are required to better respond to outside influences. Keeping a critical facility always available requires a reliable power supply and adequate cooling.

When a thermal event arises, such as an outside climate issue or an internal mechanical problem that could impact operations, an automation system can notify and respond accordingly to prevent or lessen the impact of a crisis condition. Not having an effective system can result in network downtime and extended service outages.

Operational data must be delivered quickly and accurately to the appropriate groups to prevent outages. Advances in automation with machine learning (ML) and artificial intelligence (AI) are enabling automation systems to go beyond just reporting data. AI and ML based systems can provide recommended actions and even take action to overcome an issue based on thousands, if not millions, of data points. This can eliminate an inappropriate manual intervention which may actually worsen the situation, especially if the person uses poor judgement or uses outdated documentation. A recent article from the Uptime Institute found that "more than 70% of all data center outages are caused by human error and not by a fault in the infrastructure design." Not having the information to take action or improper management of the information can be considered a failure due to lack of human intervention.

This article will focus on automation of critical infrastructure and the use of computer systems and software to develop actionable data for monitoring a facility to minimize human interaction, provide real time response and optimize operations.

1.2 Automation - A Cost or an Investment?

Automation systems are not inexpensive to buy and install. Many organizations encounter the hurdle of procurement requiring the investment to show a "simple payback" based on cost savings. A key value that an investment in automation brings is ensuring a greater degree of network uptime can be achieved. The cost of a network outage to businesses was estimated at \$5,600 per minute by Gartner in 2014. In 2022 the Ponemon Institute raised that average to \$9,600 per minute. Obviously, this varies depending on the industry vertical; however, given that digital services are the revenue generator for major multiple system operators (MSOs), the average per minute cost of an outage is undoubtedly much higher.

Business customers rely on internet access and any service outage can severely impact their revenue stream. In the case of an outage they will be looking to their service provider for some form of compensation. Work from home, IoT applications, security monitoring, emergency response, etc. all rely on having data available instantly so an outage can have serious financial repercussions.

Internal to MSOs, automated systems can lower operations costs by requiring less time for mundane tasks such as tracking temperature values or cooling system operation. Automation will enable operations to provide better information on network performance. Reports on energy use and greenhouse gas (GHG) values can be generated easily. Decisions on major capital costs, such as additional cooling systems, can be better managed with more in-depth information. The total value derived from implementing an automation system is broad in scope. The value and investment of such systems should not be reduced to a simple payback calculation.

Return on investment (ROI) is a more realistic view of the value automation brings to a site. Cost savings such as reduced truck rolls, staff overtime, cost of network downtime in lost business, and the reduced time to resolve failure issues should all be taken into account in determining the value of automation.

2. The Role Automation Can Play in Critical Facilities

Automation addresses the need to reduce outages and customer service impacts to an “always on” network facility by providing better information in near real time to identify potential issues before they become problems. There are three functions where automation benefits MSOs.

2.1. Resilience

Job One of the automation system related to cooling is to maintain the operating thermal environment of the critical facility. All control decisions and alarm notifications must be configured to properly control the heating ventilation and air conditioning (HVAC) systems to maintain optimal operating conditions and to react to an issue impacting operations. For example, if a particular HVAC unit begins to falter, the automation system can notify appropriate groups and also can react by switching to standby HVAC units or by ramping up cooling operation of other units.

2.2. Efficiency

Automation systems save energy by monitoring and controlling excess cooling capacity. Systems using ML and AI can automatically control cooling systems to respond quickly to changes in the environment and take corrective actions to ensure the desired thermal conditions are maintained. Coordinating and determining what HVAC equipment should be operating to maintain the environment is key to reducing kWh (kilowatt hour) consumption. With proper monitoring of room and rack inlet temperatures, the set points used for cooling can be increased on sites that are presently over cooled to reduce the run time required for the HVAC compressors, reducing the kWh used for cooling.

The ability to use auto responding strategies allows for a self-correcting capability and reduces maintenance call outs for non-actionable preventative maintenance (PM), e.g. low humidity. What used to be an emergency call out can now be properly identified and scheduled during normal operating hours.

Additional efficiencies can be realized with HVAC systems that have economizers to use outside air for cooling instead of running the compressors. Advanced systems learn what outside air temperature is

required to properly cool and will adjust upwards to take advantage of additional time to use outside air for cooling.

2.3. Capacity Management

Automation can identify rack locations with available cooling and air flow capacity to accommodate additional equipment, preventing hot spot issues and helping to eliminate the perception of additional cooling capacity being required due to poor equipment placement. It can also identify if equipment has been improperly placed, i.e. exhausting into cold aisle and contaminating air flow to nearby equipment. Having real time data that can identify where cooling is available will substantially reduce equipment engineering time.

In many facilities it is assumed N+1 cooling exists based on a paper calculation. Using an automated system, operations can identify high heat load areas which may experience insufficient cooling and air flow capacity.

Automation Applications

What should be automated? This depends on your corporate goals and initiatives. The biggest challenge today is how to ensure that your “always on” network is functioning properly and that your facilities are operating efficiently. Below are some use cases for automation that should be considered when implementing your automation plans.

- **Rack/room temperature monitoring** – Real-time monitoring of the ambient temperature to which the network equipment is being subjected. Identification of hotspots and adjustment of airflow based on temperature feedback.
- **HVAC performance analysis** – Monitoring the performance of your HVAC systems to indicate which units are operating properly and which are operating poorly.
- **Better/consistent operating environment** – Automation systems can respond quickly to changes in the environment and take corrective actions to ensure that desired thermal conditions are maintained. Truck rolls for non-actionable alarms can be reduced. Analysis of trends in historical data can identify potential service-affecting events in advance.
- **Redundancy and resilience in your facilities** – Testing and confirmation of true cooling and air flow redundancy can be tested and tracked.
- **Predictive maintenance of HVAC equipment** – Central monitoring software can provide operators with stored historical trend information that can be used to help predict the operation of equipment. In addition, early warnings of impending failures which may be related to weather conditions can be acted upon to lessen or avoid failure. Runtime on fans, motors or pumps can be monitored to match mean time between failure (MTBF), for proactive replacement.
- **Intelligent data focused decisions for operation and maintenance** – The ability to create actionable data, without spending hours manually combing through data records, can further optimize operations to reduce energy consumption of equipment and verify carbon emissions reduction.
- **Automated compliance management** – Automated systems can monitor the situation of business processes and data to track changes in compliance regulations related to critical facility protection.

- **Reduce carbon emissions** – With the advancements of control and monitoring systems, organizations are able to realize significant reductions in direct energy usage and carbon footprint due to gained operational efficiency.
- **Energy consumption reduction** – Automation systems save energy by monitoring and controlling excess capacity, coordinating and determining what HVAC equipment should be operating to maintain the environment.
- **Operations expense** – Unnecessary truck rolls, staff call outs and overtime will be reduced with an automated system that clearly identifies the criticality of the alarm or failure, enabling operations to determine if a next day response is adequate.

3. Benefits of Automation

Automation offers a number of significant benefits that will improve site operation, reduce costs and enable better use of the facilities resources.

3.1. Site Uptime

- The primary benefit of a properly implemented automation system is the increased network uptime and reduced time to resolve failures or identify issues before they cause an outage, greatly reducing customer impacts.

3.2. Site Operation Efficiency and Reporting

- Improving response time to issues by eliminating many common tasks and recording all current and historical state information and previous issue responses, enabling intelligent, visual decision-making and report generation.
- Reduce errors from human manual processing by providing current information on site conditions.
- Reduce or eliminate routine, mundane, and potentially dangerous tasks, thereby allowing staff to focus on more strategic, proactive, and predictive areas.
- Track energy consumption of the facility and the operation of HVAC units as well as whether power usage effectiveness (PUE), kWh used for cooling per HVAC unit, is increasing or decreasing. This data can also be used to support corporate environmental, social and governance (ESG) initiatives and reporting.
- Facilitates reporting and feedback on carbon reduction initiatives such as refrigerant leaks, energy reduction, reduced truck rolls, etc.
- Enables increasing temperature set points to maximize cooling capacity.

3.3. Reduce Costs

- Automation can eliminate many of the redundant administrative tasks that are performed to track and collect data and provide management update reports.
- Routine site visits and need for travel can be reduced or eliminated by being able to access detailed data remotely, by having eyes on the site conditions and by having better classification of alarms.
- Track operational times for each HVAC unit and provide notification of HVAC units that may be short cycling (cycling on/off) or are continuously running but not satisfying the zone temperature.

- Monitor control setpoints for the HVAC units and rack inlets to identify hot spots or excess cooling.
- Ensure all HVAC units are operating in unison rather than in conflict.
- Capital reduction opportunities due to better site operation related to:
 - Right sizing of cooling to site loads.
 - Increased economizer operation through ambient monitoring.

3.4. Better Use of Resources

- Make interventions more effective by having access to tools and information based on data rather than best guesses.
- Facilitate management of carbon reduction/sustainability targets and improvement.
- Reduce reliance on an individual's site knowledge to help address loss of expertise through retirement or attrition.
- Use capacity calculations to reduce time in planning new equipment placement.
- Gain access to data that can improve predictive maintenance and alarming for priority events, increase energy efficiency, and maintain consistent thermal conditions.
- Have access to real data that enables operators to meet corporate directives.

4. What Levels of Automation Should be Deployed?

Deciding what to automate and to what depth is challenging. The benefits obtained from automation increase as more and more tasks and processes are included. Implementing automation in stages based upon corporate goals and finance availability is one strategy that can be successful and ensures that staff become familiar with the tools operation without being overwhelmed.

Before any decision is made on automation a clear requirements document should be created to ensure that the automated system addresses the key requirements of the operations team and management. A comprehensive checklist of requirements should be developed for evaluation of a system. Ease of expansion of functionality and coverage are important considerations to ensure the system can grow as operational and corporate needs grow.

Not all sites need to be automated to the same extent. Using the SCTE Site Classification (SCTE - 226) as an example, Class D smaller sites may be adequately served by basic temperature, humidity and power monitoring, while Class C sites may require additional information for capacity planning. Depending on the organizational structure the data from a subset of Class C & D sites may be amalgamated to provide multi-site reports.

In the larger Class A and B sites additional tools and functionality to facilitate equipment placement, capacity planning and automated control of cooling system operation may be applicable due to the size of the site and movement of information technology equipment (ITE). Information on these sites and all relevant sites in the hierarchy could be amalgamated to provide regional management reports.













CLASSIFICATION	CLASS A	CLASS B	CLASS C	CLASS D	CLASS E
GEOGRAPHIC AREA	ENTERPRISE/NATIONAL	MARKET	REGION	EDGE	OUTSIDE PLANT/ ACCESS NETWORK
PRIMARY FUNCTION	 	  	   		 

Figure 1 - SCTE Site Classification

4.1. Level 1 Automation – Onsite Standalone

Simple repetitive tasks such as real time monitoring of rack inlet temperatures and HVAC supply and return temperatures are ideal uses of Level 1 automation. Making this data available at the site level in near real-time and using it for control and monitoring decisions can prevent high temperature alarms and unnecessary call outs in the off hours.

Today's HVAC systems have become smarter by incorporating automation through the use of programmable logic controllers (PLCs) or other types of proprietary control systems built into the HVAC systems. With these controls, HVAC units can better match the heat load by using variable speed fans and variable capacity cooling, whether that be multiple stages of cooling or variable speed compressors. While the HVAC units are individually very efficient, there is not much sharing of data to coordinate the overall performance of a facility. By adding networking capability with automated intelligence, the HVAC systems can be more responsive to site cooling needs.

4.2. Level 2 – Onsite Complex

The use of ML and AI extends the capability of automation from one of data collection to being able to respond intelligently to an impending issue. By using the historical operational data, an intelligent automation system can identify if a cooling system is beginning to perform abnormally, which may be the result of a dirty condenser coil, slipping fan belt or low refrigerant charge

By having access to more sensor data, the systems can make intelligent decisions about which HVAC or combination of units needs to run in order to satisfy the heat load in a facility. This coordination also allows for lead/lag control of multiple HVAC units to ensure even run times and prevent units from fighting each other by operating at different setpoints or in different modes of operation. Multiple pieces of equipment can be trended over time and run holistically to optimize overall facility conditions and efficiency.

4.3. Level 3 - Central Monitoring and Control Automation

Central monitoring builds on the onsite comprehensive automation by providing access to data from multiple facilities. Central monitoring allows for gathering operational status of all HVAC units in every facility being monitored whether the data is gathered directly from smart HVAC units or through the onsite building automation system (BAS). Comparisons can be made across the portfolio to determine

outliers and where remediation is required. Long term monitoring can also spot changes in the operating conditions of HVAC units. This can help predict when maintenance will be required in order to return a unit to proper operational condition.

5. Best Practices for Choosing and Implementing Automation

5.1. Define Clear Objectives

Start by identifying the specific goals and outcomes you want to achieve with automation. Having a clear vision will guide your automation strategy.

5.2. Start with a Proof of Concept (PoC)

Before implementing automation at scale, conduct a proof of concept (PoC) to test your chosen automation tools and workflows. This allows you to identify and address potential issues early. If possible PoC should be conducted in a lab environment to verify/validate operation and impact on systems rather than in a live customer setting.

5.3. Invest in Training and Skill Development

Provide training and resources to upskill your staff in automation technologies and best practices. A knowledgeable team is essential for successful implementation.

5.4. Choose the Right Automation Tools and Platforms

Select automation tools and platforms that align with your needs and objectives. Consider factors such as scalability both in size and functionality, compatibility with existing in-house systems, and long-term vendor support.

5.5. Implement Robust Security Measures

Prioritize security throughout the automation process. Use encryption, access controls, and regular security audits to protect your critical infrastructure from threats. Keep software up to date, applying all pertinent patches. Ensure all vendors and suppliers are compliant with company security requirements.

5.6. Continuously Monitor and Optimize

Regularly review and refine your automation workflows. Monitor performance, gather feedback, and make adjustments to improve efficiency and reliability. Perform regular back-ups to facilitate network recovery in event of an outage.

6. Automation in Practice

To illustrate implementing automation the following are real examples of how MSOs are using automation to support their “always on” networks.

6.1. Level 1 Use Case: Rack Inlet Temperature Data Collection

Pre-Automation: HVAC monitoring and alarming

A major cable operator required staff to manually collect and record room, row and rack inlet temperatures and HVAC supply and return air temperatures twice daily, requiring four hours of staff resources per day. Averaging and review of the data collected was performed manually. Hotspots were determined to be any rack or row temperature above 80°F (26°C) were flagged and reported. Paper records were kept and archived in boxes and weren't readily available for troubleshooting when problems arose.

Post-Automation Implementation

A monitoring system was deployed with temperature sensors installed to collect rack inlet temperatures and HVAC supply and return. The monitoring system aggregated the data and provided a real-time view of temperatures along with a historical log.

The results of automating temperature monitoring include:

- Staff time savings – no longer required to manually take temperature readings.
- Real time reliable data – accurate readings available 24/7
- Non-intrusive to headend – no impact to operation
- Hotspot identification for prompt remediation – quick response time means problems are resolved faster without impacting the headend.
- Self-reporting sensor failure – real time notification of issues.
- Auto-alarming to engineering and monitoring teams. – immediately know when an impending problem is occurring before it becomes a critical situation.
- Baseline to room heatmap – a simple but effective method to heatmap the facility to correct airflow issues.

6.2. Level 2 Use Case: HVAC Device Data and Control

Pre-Automation: HVAC operation, monitoring, and alarming

Stand-alone legacy HVAC performance monitored did not provide any data regarding the operating performance and did not coordinate with other HVAC systems to efficiently cool the headend. Because there was no awareness of how the system operated, technicians were required to be dispatched whenever they received an HVAC alarm from a summary alarm contact. These could trigger for multiple reasons, most of which don't require immediate dispatch. When technicians were on site, there was no method for gathering data for troubleshooting, resulting in extended service calls and possibly multiple calls since the technician might not have had all the necessary supplies to solve the problem. Additionally, since it was unknown what the issue causing the summary alarm was, it required emergency dispatch, possibly after hours, to respond to something that could have been resolved during standard operating hours.

Post-Automation HVAC monitoring and control:

HVAC systems being deployed were required to be smart systems that share performance data with the automation system. Depending on the HVAC system, various telemetry points are available for trending and analysis such as:

- Supply, return and zone temperatures;
- Outdoor ambient temperatures;
- Indoor and outdoor relative humidity;
- Compressor and fan speeds;

- Economizer operation and engagement;
- Remote setpoint adjustments and control;
- Duct and coil pressure differential readings;
- High visibility to operational status;
- Cooling, humidify, dehumidify, reheat;
- Multiple specific alarm points;
- Communication loss to device;
- Energy consumption recorded.

The additional sensing of the HVAC systems provides thorough operational performance and specific alarms for precise response, reducing nuisance truck rolls for non-serious events or conditions, saving money.

Tracking and archiving of data can reveal trends and allow for predictive analysis of the HVAC performance. In addition, HVAC systems are now coordinated and can be controlled using lead/lag or other sequencing strategies to select the correct level of cooling and air flow required to optimally match the heat load and automatically adjusting as changes in the facility occurs.

7. The Path Forward

Automation is a key operational component in critical facilities, and it is not going away. The longer major cable operators stall in gaining the experience of using automated system the less knowledge and opportunity they will have to influence the design of systems that meet their specific requirements.

Automation is advancing rapidly as equipment is being deployed with telemetry capabilities that were not available just a few years ago. The ability to monitor multiple data points within seconds and replay critical information to operators allows critical facilities to take advantage of Machine Learning (ML) to respond and adapt to changes in the facility to further optimize operations and prevent downtime.

Automation systems can learn and optimize HVAC operation to optimally maintain the network equipment inlet temperatures. By taking into account such influences as outside air temperature and changes in equipment heat load, HVAC systems can learn to adjust cooling operation and supply fans to adjust airflow to match cooling needs and reduce hot spots.

The ability to deploy multiple sensors to capture thousands or millions of data points using advanced telemetry from HVAC and other equipment will allow operators to progress from simple machine learning functions to more advanced capabilities of artificial intelligence.

Artificial intelligence-based applications will improve energy efficiency, resiliency and indoor environmental quality while reducing an operator's carbon footprint. The ability for a system to interpret multiple inputs including site cooling and power loads data, weather forecasts, utility data, and historical performance will further aid in optimizing operations.

A few additional examples of AI potential include:

- Intelligent decision making on dispatching of service technicians reducing carbon emissions associated with service vehicles.
- Ability to answer the “what ifs,” i.e., what if changes are made to the operation such as raising rack inlet/room temperatures? How will that impact the site operation?

- Ability to understand the current cooling and power requirements of a facility and to determine available capacity for network equipment expansion.
- Predictive failure analysis, or the ability to identify trends and subtle differences in equipment operation to determine if a failure condition is imminent.
- Ability to incorporate other systems and operation workflows into the overall operation of the critical facility, allowing operators to prioritize precious resources to ensure network integrity.

Advances in ML and AI will make automation a necessary component in site operation. As more tasks are automated, a greater database of information will be available to make decisions and predictive analysis which will help to identify operational issues, improve resiliency resulting in reduced outages. Critical facilities will be more energy efficient with reduced carbon emissions and GHGs.

8. Abbreviations and Definitions

8.1. Abbreviations

AI	artificial intelligence
AP	access point
AWS	Amazon Web Services
BAS	building automation system
GHG	greenhouse gases
HVAC	heating ventilation and air conditioning
IoT	Internet of things
ITE	information technology equipment
kWh	kilowatt hour
ML	machine learning
MTBF	mean time between failure
PLC	programmable logic controller
PM	preventive maintenance (see ISO/IEC 30134-2:2016)
PoC	proof of concept
PUE	power usage effectiveness
SCTE	Society of Cable Telecommunications Engineers

8.2. Definitions

Always On	Resilient network with zero unavailability, no customer downtime
Artificial Intelligence	An evolving term first coined by Stanford Professor John McCarthy in 1955: “the science and engineering of making intelligent machines”.
Downtime	Measure of a critical facility unavailability
kWh	kilowatt hour, standard measure of energy use
Machine Learning	A subset of AI involving computer agents improving perception and actions based on experience or data

Bibliography and References

Uptime Institute - <https://uptimeinstitute.com/component/content/article/26-news-external/458-uptime-institute-70-of-dc-outages-due-to-human-error>

Automating Greenfield Access Planning Using Intelligent GIS Algorithms

Getting Buildout Cost Estimates in Minutes, Not Months

A Technical Paper prepared for SCTE by

Jared Faulk, Software Engineer, First Principles Innovations, SCTE Member
111 Fieldbrook Ct.
Cary, NC 27519
Jared.faulk@fpinno.com
+1-704-775-5926

Luc Absillis, COO, First Principles Innovations, SCTE Member
111 Fieldbrook Ct.
Cary, NC 27519
luc@fpinno.com
+1-919-345-2655

Table of Contents

Title	Page Number
Table of Contents	69
1. Introduction	70
2. Piecing the Puzzle Together	70
2.1. Inputs to the system	70
2.1.1. HHP and Building Data	71
2.1.2. No Perfect Data Source Exists	73
2.1.3. Multi-Source Merging and Verification	74
2.1.4. Gathering Roads	75
2.2. Outputting a Full Network	75
2.2.1. All the Ingredients Together	75
2.2.2. Estimating Node Placements	77
3. Conclusions	82
4. Abbreviations and Definitions	83
4.1. Abbreviations	83
4.2. Definitions	83
5. Bibliography and References	84

List of Figures

Title	Page Number
Figure 1 - A House that Appears to be a Duplex, with Two Doors and Two Address Numbers, is a Quadplex.	71
Figure 2 - "Street View" Not Available Everywhere	72
Figure 3 - A View of Multiple Building Sources Overlaid on Top of eOne Another	73
Figure 4 - An Example of Varying Views for the Same Building	74
Figure 5 - The Total View of All Buildings and Roads Collected	76
Figure 6 - A Closer View of Collected Buildings and Roads	76
Figure 7 - One of Approximately 27,000 buildings with Calculated Drop Lengths	77
Figure 8 - High Level View of Entire Calculation	78
Figure 9 - A Closer View of the Calculation	79
Figure 10 - A Much Closer Look at the Nodes and Buildings	79
Figure 11 - A View of Metadata at the Node Level	80
Figure 12 - A Map View of Calculated Feeder Miles	81
Figure 13 - A Summary of Calculation Results	81

1. Introduction

Expanding one's network is no small task— from large government grants like Rural Digital Opportunity Fund (RDOF) or the Broadband Equity, Access, and Deployment program (BEAD), to continual expansion of the population, efficient and quick greenfield planning is crucial. However, sending a person to walk the streets of every potential area, or manually scoping areas out in a geographic information system (GIS) tool like Google Earth can be time consuming, prone to error, and algorithmically limiting. Furthermore, only a small percentage of buildouts can and will come to fruition, making time efficiency ever more critical.

Building quick network expansion plans are difficult and complex because of:

- *Inaccurate information:* Getting a single source of truth to calculate accurate plans, thus cost and resource forecasting, can be next to impossible.
- *Too many unknowns:* Including unknown node locations, distribution miles, and homes passed information.
- *Heavy time restrictions:* Getting a cost estimate from detailed design takes too long, and if started and abandoned, can lead to huge, wasted efforts in both time and investments.
- *Difficult what-if analysis:* Viewing the deployment strategy from different angles can add an additional level of complexity – resulting in suboptimal deployments or worse, lost opportunities.

From a high-level view, operators need to know what it is going to cost to build out an area accurately and quickly, without the cost and risk of diving into detailed design. Instead, with the right algorithmic intelligence, coupled with readily available GIS data— a full network plan with detailed costs, resources, and topology is achievable.

In this technical paper, we explore available GIS data (like roads, buildings, and area information) and their current inadequacies. We propose novel algorithms in automating node placements. We discuss how to resolve the difficulties of creating a viable network expansion algorithm that strictly adheres to different technology constraints, which, in the end, should allow for hypothetical buildouts using every array of technology available. Furthermore, we detail how the proposed algorithms can solve node location placement, distribution miles, cost and resource requirements, as well as generate a full network topology, all in minutes— not months.

2. Piecing the Puzzle Together

2.1. Inputs to the system

To create a build out plan in a new area, we need a few core inputs: buildings (latitude and longitude), roads, and households passed (HHP). HHP are different than buildings, because there can be multiple HHP in a single building. We need buildings (with location data) and HHP to calculate several outputs such as node placements and drop lengths. We need roads to calculate distribution miles and feeder miles. To refine a buildout will obviously take more input (i.e., construction cost and resources, technology costs, etc.), but those details will be specific to each use case and without these three core inputs (buildings, HHP, and roads), no automated greenfield plan can be created.

2.1.1. HHP and Building Data

Manual gathering of these inputs is the traditional approach to solve this issue – however, manual data gathering is very time consuming, prone to error, and usually relies on a single data source. For example, physically walking the streets is going to be the most accurate way to gather information on HHP, though also the most time consuming, expensive, and hard to scale. It requires sending a field worker to scout the area manually, potentially driving long distances just to reach the area. Manually scouting a single block in a city is easy, but what about manually scouting 100 blocks in four or five different cities scattered over the state? The amount of time and effort to get accurate data at scale then grows considerably, and may become outright impossible.

What about Google Earth? With modern technology we can now see anywhere in the world with quite amazing detail. Instead of physically having to “walk the streets,” operators virtually do so, via Google Earth and other related GIS software tooling. This can be a significant time saver, preventing one from having to physically scout neighborhoods. However, it is still a manual process and requires considerable hours from a person, or a team of personnel to scout out even a relatively small area. Not to mention, the data gleaned may vary in accuracy—for example:




Figure 1 - A House that Appears to be a Duplex, with Two Doors and Two Address Numbers, is a Quadplex.

Figure 1 shows a house that appears to be a duplex, having two doors, and two address numbers. However, this building is a quadplex. This was verified via contacting the residing renter there – a good friend of the author. No amount of walking the streets (physically or virtually) could amass this answer, short of entering the building. Exceptions like this can occur often, which makes HHP gathering very difficult. By the way, the difficulty of objects blocking houses, like the tree in this example, is quite common in Google Earth.




Figure 2 - "Street View" Not Available Everywhere

Another example where these software tools can fall short is when streets are not available to “walk.” Figure 2 shows an entire street missing from Google Earth’s “Street View” functionality. This can lead to unverified buildings, which in turn leads to either miscounted HHP or an excessive amount of time spent trying to verify individual buildings. Is the above picture 2, 3, 4, or 5 HHP? For such a small area, the answer is not as straightforward and as easily attained as one would hope. Exceptions exist like this all over the map; these add up overtime for the many more man hours needed to accurately glean just the input of HHP/buildings.

If manually gathering data is expensive and time consuming, what other alternatives exist?

Enter the API (Application Programming Interface). An API is a link from one piece of software to another piece of software or database—it is like a gateway that allows different tools to interact with one another.

We can use APIs to link into available databases with building data. Through APIs one can access a database without having to download all the data in the database and process/store it locally. APIs thus encapsulate the complexity and size of the data required only at the time of request, allowing for efficient scaling. In a perfect world, there would be a single database, with open and free access to APIs that provide accurate buildings and unit data. Unfortunately, this database does not exist, and likely never will. Even the USPS, with a database of over 165 million house addresses, does not open their address database to the public for access, not even for a commercial price [9].

2.1.2. No Perfect Data Source Exists

Since no perfect data source exists, are operators consigned to manually gather HHP forever? We propose not. Though no single database exists with every building all over the world, there exist many databases and data sources with localized, highly accurate levels of building data. These data sources can vary in scale from across the country to the county level to the census block level. A few examples of free, publicly available building data sources are the United States FEMA (Federal Emergency Management Agency) USA Structures database that documents all buildings in the United States with an area greater than 450 square feet, Open Street Maps building data available via their Overpass API, and the Department of Transportation's National Address Database (NAD) [3][6][10]. Many of these data sources are available through Esri servers as well [1][3][7]. Furthermore, many counties in the United States collect and document all addresses, parcels, and building types with latitude and longitude information, freely available for download. If you mapped a couple of these different sources on top of each other, you might get something like this:

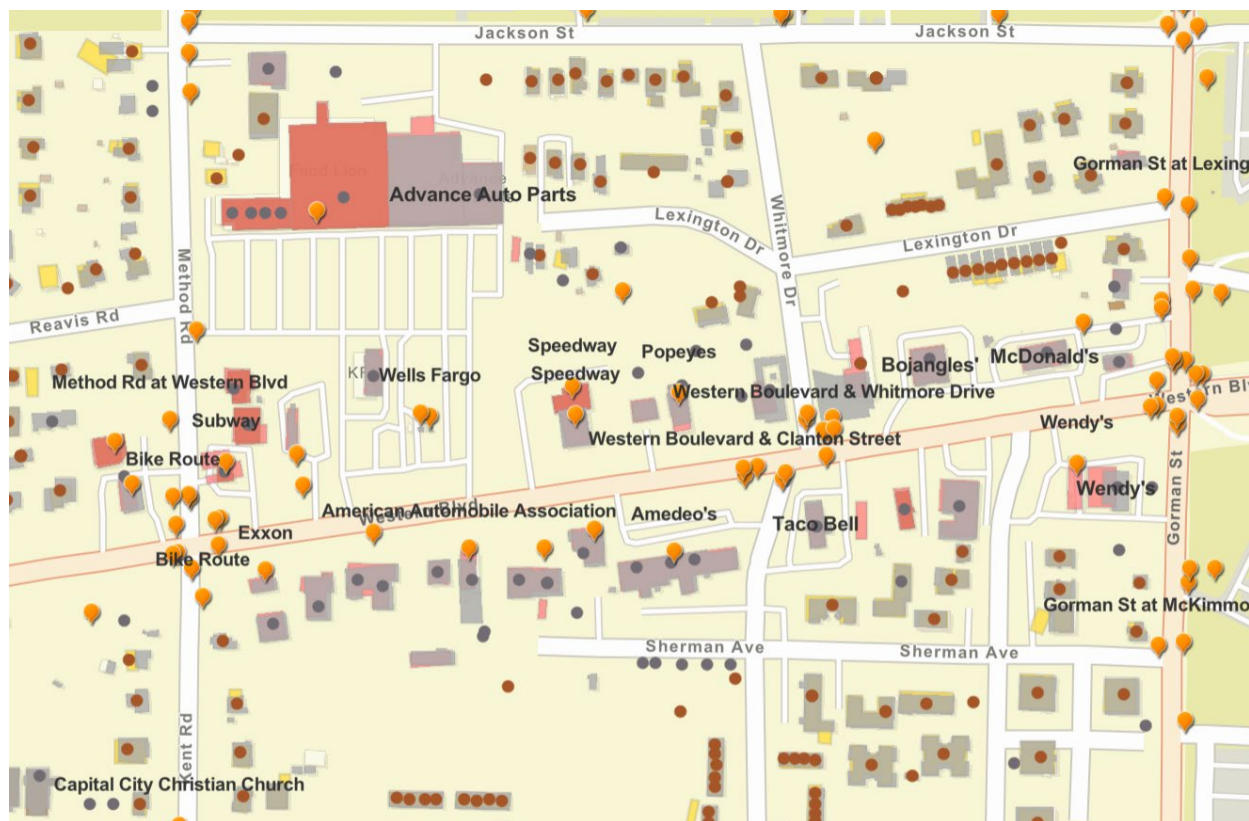


Figure 3 - A View of Multiple Building Sources Overlayed on Top of eOne Another

For any given area in the United States, we can merge multiple data sources – each with metadata, say, classifying businesses from residential buildings etc. – to glean a highly granular view of HHP and building locations, like in the picture above. By leading with programmatic building gathering, it is possible to get the total HHP in the above example in less than 10 seconds. No manual street surveying, no virtual walking necessary.

2.1.3. Multi-Source Merging and Verification

“Merging” as referenced above, was partially glanced over. The reason was intentional, because simply merging two, three, or four data sources – as anyone who has tried to do will attest – is not so simple and requires its own section. There is, however, enormous benefit to be had from using multiple data sources for building and HHP gathering: as shown above, one can achieve a high level of coverage for all buildings in a single area extremely fast; one can also glean HHP information from more areas than any single source could glean (because no perfect data source with perfect coverage exists). In addition, building/unit verification becomes stronger as more sources confirm one another – i.e., if Source A gives a point with one business, it may or may not be correct, but if Source B and Source C also give a point at the same location with one business, then the level of confidence that this point is indeed a building and a business, grows substantially.

The biggest issue with multi source merging is the varying level of degrees of detail, location data, and metadata between differing sources. For example, Source A may have all building polygons with metadata classifying it as either business, single family unit (SFU), or multi-dwelling unit (MDU); however, it may not have the number of units of each building. Source B may have the number of units of each building, but not the type of building. Furthermore, these locations, coming from different data sources, may have slight variances, leading to difficulties merging the two sources referencing the same building. An example:




Figure 4 - An Example of Varying Views for the Same Building

Figure 4 shows three different data sources mapped on top of one another. There are four buildings, and each data source recognizes those four buildings, but each data source has its own idea of where those buildings are. Two of the data sources intersect, but the third one does not. Assuming these are SFUs we only want to count four HHP here, not eight.

The other complexity with using multiple data sources is the varying degree of availability depending on the area. Source A may be very good in North Carolina, but non-existent in Alabama, whereas Source B may be available in both, but not as highly detailed as Source A. Choosing the best data source combination, based on the level of availability, for any given area in the United States, is simply impossible to do by hand, but only achievable through a programmatic approach to data collection.

2.1.4. *Gathering Roads*

To automate a buildout, we need roads to calculate node placements (nodes are generally placed along a road), reachability of houses from the node (distribution miles), drop lengths (how far the house is from the road), and feeder miles (how much feeder fiber is needed to service the nodes).

Roads, thankfully, are not as complex to gather as HHP, at least in the United States. Because routing has become so mainstream in the 21st century, there have been a lot of efforts to provide accessible road data via APIs and downloads with broad coverage for entire countries. For example, the United States Census Bureau publishes yearly all shapefiles (a file that merges metadata with a geometry, in this case a line representing a road) for all roads covering the United States [2]. Many companies take this data, clean it, and expand it with their databases. One such example is Esri's USA "Transportation" server that consolidates and cleans all of the Census Bureau's road shapefiles into a single server with APIs [7]. Open Street Maps is another example of a free database of all roads inside of the United States with open access via APIs [5][6].

2.2. Outputting a Full Network

2.2.1. *All the Ingredients Together*

Once we have the necessary inputs— buildings, HHP, and road data— the possibility for an automated network deployment, and thus an extremely fast and accurate cost estimate, becomes real. Manually gathering these inputs for any area can take a considerable amount of time, but at this point, gathering all HHP, buildings, and roads for a section of a city, say with 35,000 HHP via programmatic data gathering, can be achieved in about one minute. A map of this data might look something like this:




Figure 5 - The Total View of All Buildings and Roads Collected

If we zoom in, we can see the details a bit more:

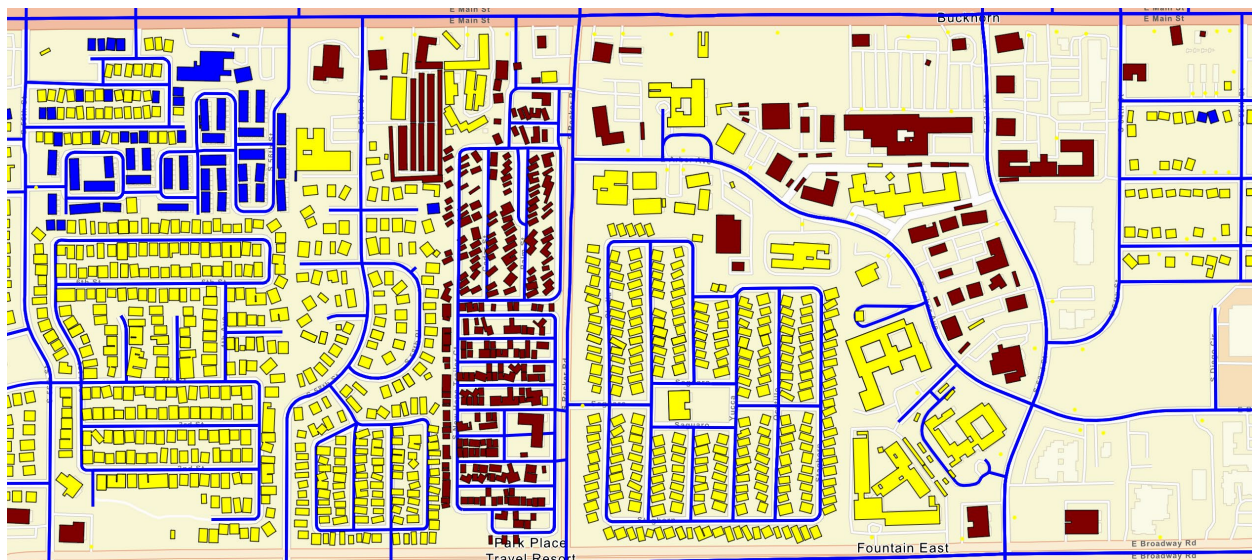


Figure 6 - A Closer View of Collected Buildings and Roads

Note in Figure 6, red signifies a business, yellow signifies a single-family unit (SFU), and blue signifies a multi dwelling unit (MDU).

The picture above shows a subsection of an area queried, with all buildings, HHP, and road data available. For 37,282 HHP, gathering roads and buildings, and running a merging strategy took only one minute and seven seconds. This is a feat clearly impossible by manual data gathering. Furthermore, now

that we have this data, there are countless possibilities to utilize intelligent algorithms to automate a greenfield buildout.

An easy calculation to achieve, given we have all building data and road data in an area, is calculating approximate drop length for each building.



Figure 7 - One of Approximately 27,000 buildings with Calculated Drop Lengths

Figure 7 shows an example of one building, out of 27,000 with a calculated drop length.

2.2.2. Estimating Node Placements

The more difficult problem than gathering HHP, roads, and calculating drop lengths in an efficient manner, is estimating node placements such that nodes can physically reach the homes they are serving, given the total buildout picture. Each node slots into an area that is optimized to service the most amount of HHP possible with the minimal amount of distribution miles. In computer science, this specific problem is something we would call an NP Hard (non-deterministic polynomial-time hard) problem, meaning a perfect mathematical solution cannot be achieved in polynomial time, that is, with a time complexity of Big O (n^c). An NP Hard problem may have a solution, but that solution cannot be run in a reasonable time on modern computer architecture, essentially.

Thankfully, calculating a mathematically perfect solution is not necessary for most problems, including this one. Where operators place nodes is never mathematically perfect because in the real world there are other constraints such as permitting, terrain, etc. By approximating, just slightly, we can come up with algorithms that do a very good job at placing nodes in a reasonable spot in the buildout area, respecting each other's node HHP, and aiming to minimize distribution miles per node as much as possible. One such algorithm that is possible may look like this:




Figure 8 - High Level View of Entire Calculation




Figure 9 - A Closer View of the Calculation

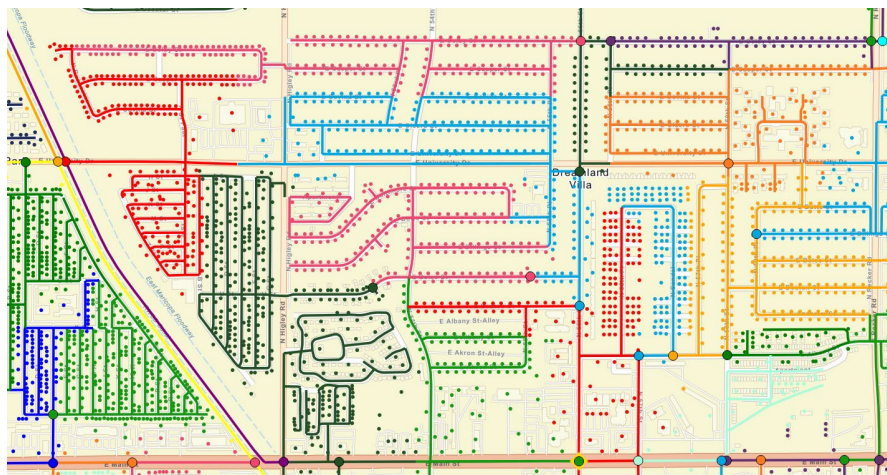


Figure 10 - A Much Closer Look at the Nodes and Buildings

Note: each node has its own color, and each building that is services shares that same color. Some nodes may have the same color but are located on different parts of the map.

The picture in figures 8-10 above display an algorithm that services all HHP in the chosen area by placing nodes such that a household is reachable by the node and the node has capacity for that household. This calculation was done using 10-Gigabit-capable passive optical network (XGS PON) node technology with a limit of 256 HHP per node. Given we then place nodes such that each house is reachable by the node, we can then calculate all the roads needed to get from the node to those houses. In essence, we can

calculate distribution miles. In the end we can accumulate very detailed data about each node, and a total picture of the buildout area. For example, some metadata for each node might look like:

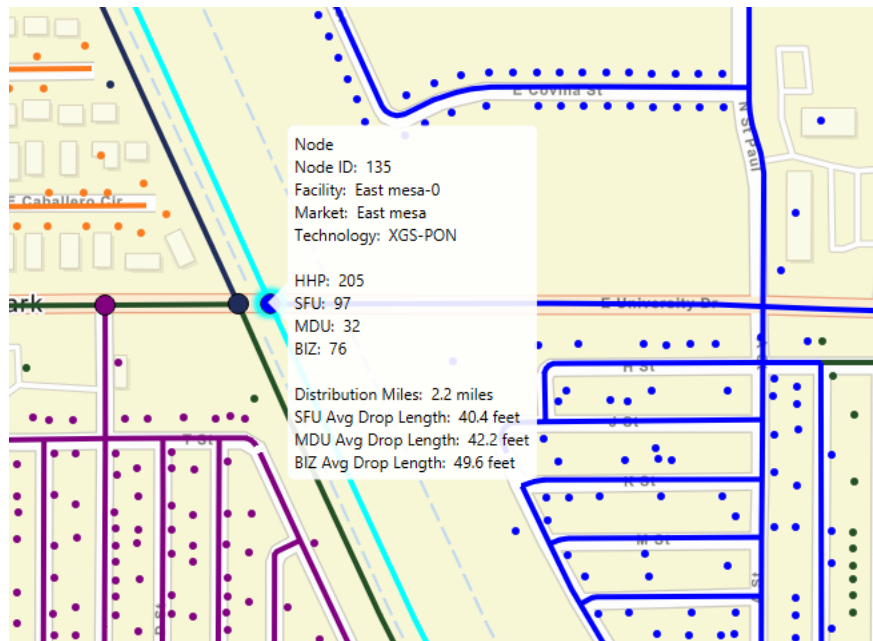


Figure 11 - A View of Metadata at the Node Level

Figure 11 shows some of the possible metadata that can be gleaned from this calculation. Also, since we have all nodes in allotted placements along roads, we can then calculate feeder miles as well. It may look something like this:

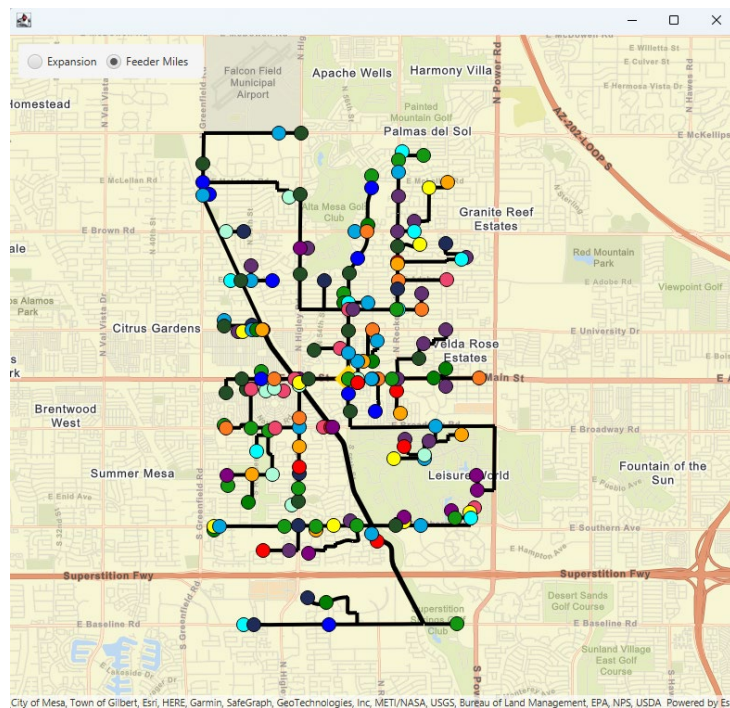


Figure 12 - A Map View of Calculated Feeder Miles

For this selected area in Phoenix, Arizona, calculating everything from building and road gathering to adding node placements, and calculating drop length, distribution miles, and feeder miles took only two minutes and four seconds. Performance, accuracy, speed, and scalability simply cannot be achieved via manual data gathering and calculation. Furthermore, with a full greenfield buildout calculated, if cost and resources are added, one can then come up with a very accurate estimate as to what it is going to cost to implement such a buildout. These algorithms work anywhere in the United States. An example using this same scenario in Phoenix, Arizona:

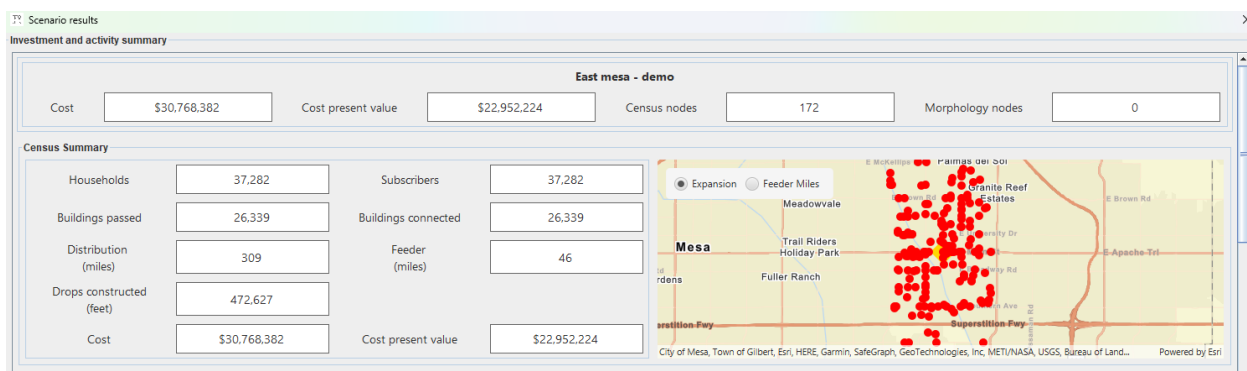


Figure 13 - A Summary of Calculation Results

The above picture shows the estimated cost for total buildout, the cost present value, the number of calculated nodes, the number of buildings passed, total distribution miles, total drop lengths, total number of subscribers (in this scenario, we used a 100% take rate), total number of buildings connected, and total feeder miles.

3. Conclusions

When a programmatic approach is implemented, extremely hard and cumbersome tasks such as estimating a greenfield build out can be done in just a couple of minutes with data down to the building level. Manually approaching the fast-moving world of government grants, increasing populations, and company competition, operators can no longer afford the large wait time that traditional methods of greenfield access planning required, due to the insurmountable scalability, speed, and detail issues of manual planning.

In this paper, we have proposed a programmatic approach to tackle the issue of fast, efficient, and accurate buildout cost estimates. Using this approach can help operators and overbuilders more accurately plan for building out areas with true confidence in cost and resources. Government grant bidding can also benefit greatly from such an approach, because estimating an area's requirements can allow for quick turnarounds needed to keep up with competition. Lastly, using a programmatic approach allows for extreme flexibility in adjusting and constructing algorithms to suit one's specific use case. Something that may take months for a team of manual inputting can be achieved with reasonable accuracy, in a couple of minutes.

4. Abbreviations and Definitions

4.1. Abbreviations

API	application programming interface
BEAD	broadband equity, access, and deployment
GIS	geographic information system
HHP	households passed
MDU	multi-family dwelling
NP-Hard	non-deterministic polynomial-time hard
RDOF	rural digital opportunity fund
SFU	single family dwelling
XGS-PON	10-Gigabit-capable passive optical network

4.2. Definitions

NP-Hard	A defining class of problems in computational complexity theory, which defines a problem as “non-deterministic polynomial time” if there exists no such algorithm that can solve the problem in Big O (n^c).
---------	--

5. Bibliography and References

1. Address_points_from_national_address_database (FeatureServer),
services.arcgis.com/xOi1kZaI0eWDREZv/ArcGIS/rest/services/Address_Points_from_National_Address_Database/FeatureServer. Accessed 2 Nov. 2023.
2. Bureau, US Census. "Tiger/Line Shapefiles." Census.Gov, 5 Dec. 2022,
www.census.gov/geographies/mapping-files/time-series/geo/tiger-line-file.html.
3. "National Address Database." U.S. Department of Transportation,
www.transportation.gov/gis/national-address-database. Accessed 2 Nov. 2023.
4. OSM_NA_Buildings (FeatureServer),
services6.arcgis.com/Do88DoK2xjTUCXd1/ArcGIS/rest/services/OSM_NA_Buildings/FeatureServer. Accessed 2 Nov. 2023.
5. OSM_NA_Highways (FeatureServer),
services6.arcgis.com/Do88DoK2xjTUCXd1/arcgis/rest/services/OSM_NA_Highways/FeatureServer. Accessed 2 Nov. 2023.
6. "Overpass API." Overpass API - OpenStreetMap Wiki,
wiki.openstreetmap.org/wiki/Overpass_API. Accessed 2 Nov. 2023.
7. Transportation_v1 (FeatureServer),
services2.arcgis.com/FiaPA4ga0iQKduv3/arcgis/rest/services/Transportation_v1/FeatureServer. Accessed 2 Nov. 2023.
8. USA_Structures_View (FeatureServer),
services2.arcgis.com/FiaPA4ga0iQKduv3/ArcGIS/rest/services/USA_Structures_View/FeatureServer. Accessed 2 Nov. 2023.
9. "USPS by the Numbers 2023: Key Data for Congress - about.Usps.Com." USPS,
about.usps.com/who/government-relations/assets/usps-by-the-numbers-2023.pdf. Accessed 2 Nov. 2023.
10. "Welcome to the Disasters Public Files." Disasters Public Files - USA_Structures,
disasters.geoplatform.gov/publicdata/Partners/ORNL/USA_Structures/. Accessed 2 Nov. 2023.

Improvements in Upstream Profile Management Algorithms

Delivering Reliable and High Upstream Speeds in Production Networks with OFDMA Technology

A Technical Paper prepared for SCTE by

Marie-Sophie Graftiaux

Technical Improvement Manager

VodafoneZiggo

Boven Vredenburgpassage 128, 3511 WR Utrecht, Nederland

marie-sophie.graftiaux@vodafoneziggo.com

Ying Wang

Sr. Data Scientist

VodafoneZiggo

Boven Vredenburgpassage 128, 3511 WR Utrecht, Nederland

ying.wang@vodafoneziggo.com

Karthik Sundaresan

Distinguished Technologist & Director of HFC solutions

CableLabs

858 Coal Creek Circle, Louisville, CO 80027 USA

k.sundaresan@cablelabs.com

Table of Contents

Title	Page Number
1. Introduction.....	88
1.1. Background of DOCSIS 3.1 Profiles and the Need for PMA	88
2. Profile Calculation Methodology.....	90
2.1. Overview and Evolution of PMA Algorithms	90
2.2. Input Data Samples to the PMA Algorithm	93
2.3. Optimization Metrics.....	93
2.3.1. J-value.....	93
2.3.2. A New Simplified Optimization Metric: Capacity Impact.....	94
2.4. Summary of Current PMA Algorithms	94
2.4.1. An Optimized Profile-Coalescation Algorithm	95
2.4.2. An Upstream PMA Algorithm 2C	96
2.4.3. A Two-Stage Upstream PMA Algorithm	96
2.5. Introducing the New Two-Stage PMA Algorithm with Pre-Grouping	96
2.5.1. Pre-Grouping of Samples	96
2.5.2. PMA Stage 1	97
2.5.3. PMA Stage 2	97
2.5.4. Different Configurations in Two Stages	97
3. Field Trial Results on VodafoneZiggo Network.....	98
3.1. Implementation of the New Two-Stage US Algorithm with Pre-Grouping	98
3.1.1. Dataset.....	98
3.1.2. Pre-Grouped Samples	98
3.1.3. Two-Stage PMA Results	99
3.1.4. Computation Time.....	100
3.1.5. Capacity Gains.....	101
3.2. Impact of the PMA Algorithm Implementation on the Network Performance	101
3.2.1. VodafoneZiggo US Spectrum	102
3.2.2. Building Blocks of PMA End-to-End Pipeline.....	102
3.2.3. Field Trial Setup	103
3.2.4. Network Performance	104
4. Conclusion and Future Work.....	110
5. Appendix	111
5.1. Configurations for Two-Stage PMA.....	111
5.2. Detailed Example of a Step within the Profile Coalescation Algorithm.....	111
5.3. PCDA	113
Abbreviations	114
Bibliography & References.....	114

List of Figures

Title	Page Number
Figure 1 - Upstream Profile Example, Baseline Modulation with Exceptions	89
Figure 2 - US OFDMA Profile Switching Example	90
Figure 3 - Downstream OFDM RxMER from One CM, over One Week.....	91
Figure 4 - Upstream OFDMA RxMER from One CM, over One Week.....	91
Figure 5 - Upstream OFDMA RxMER from Multiple CMs over One Week	92
Figure 6 - Downstream OFDM Example with Two Distinct RxMER Signature Clusters	93

Figure 7 - Profile Coalescation Algorithm Example	95
Figure 8 - Recomputing Impact only for New Clusters	95
Figure 9 - RxMER Values of the First Group (out of 16 Groups), of Interface 4.....	99
Figure 10 - RxMER Values of the Second Group (out of 16 Groups), of Interface 4	99
Figure 11 - RxMER Values of 15 th & 16 th Groups (out of 16 Groups), of Interface 4	99
Figure 12 - Generated Profiles with OPCA	100
Figure 13 - Generated Profiles with the New Two-Stage Upstream Algorithm with Pre-Grouping	100
Figure 14 - Generated Profiles with the Upstream Algorithm 2c	100
Figure 15 - J-values of all Interfaces, using the Two-Stage US Algorithm w. Pre-Grouping	101
Figure 16 - US Spectrum in VodafoneZiggo Network.....	102
Figure 17 - IUC Usage and Average Channel Efficiency, Scenarios (1) & (3)	107
Figure 18 - IUC Usage and Average Channel Efficiency, Scenarios (2a) and (4).....	107
Figure 19 - IUC Usage and Average Channel Efficiency, Scenarios (2b) and (5).....	108
Figure 20 - Profile Coalescation: Combine Clusters using Capacity Impact	112
Figure 21 - Profile Change Decision Algorithm (PCDA) Decision Tree	113

List of Tables

Title	Page Number
Table 1 - CMTS interfaces used to poll US RxMER data.	98
Table 2 - Computation time	101
Table 3 - Conversion Table, Modulation to Mbps/MHz.....	104
Table 4 - Summary of Deployment Scenarios	105
Table 5 - Static Exceptions	106
Table 6 - Average Capacity in Mbps/MHz per IUC for 5 scenarios	106
Table 7 - Configurations for Two Stage PMA	111

1. Introduction

DOCSIS 3.1 technology is now widely deployed and delivering gigabit services to customers. Operators are making significant strides in understanding how to make upstream orthogonal frequency division multiple access (OFDMA) technology work robustly and deliver the high upstream speed tiers. VodafoneZiggo, a cable operator in the Netherlands, and CableLabs have been collaborating on implementing an upstream profile management application (PMA) solution in the field. By tailoring profiles to groups of cable modems (CMs) on the OFDMA channel, the aim is to increase the stability of the channel and the total capacity realized without increasing the costs. The upstream channels are operating in a very low and challenging part of the spectrum; implementing PMA within that spectrum can create a lot of added value.

This paper will give an overview of the main algorithms for both downstream and upstream, define updated/refined optimal metrics for use in calculating profiles, and share the algorithms and data analysis methods that produce the optimum profile definitions for an OFDMA channel. This paper will deliver a solid understanding of the latest in PMA technology and some of the latest methodologies that have been developed.

We will compare results when using the different algorithms on upstream OFDMA channels and using those profiles in a live network. The field trial consists of a fully closed loop solution, including upstream data collection, configuring the profiles / interval usage codes (IUC) generated by PMA on a live plant, and then tracking the impact i.e., behavior of the CMs and the performance of the upstream.

The experiments and trials during the upstream PMA deployment at VodafoneZiggo have produced very interesting results. When implementing PMA algorithms in the network we see an increase of 60% in capacity compared to a static two-profile solution. In a network configured with static exceptions and more profiles, the gain is reduced to 30%. In addition, we see substantial improvements in stability.

1.1. Background of DOCSIS 3.1 Profiles and the Need for PMA

DOCSIS 3.1 introduced features using the orthogonal frequency division multiplexing (OFDM) based PHY layer, including variable bit loading and defining multiple modulation profiles for downstream and upstream channels. Additional new functionalities include measuring downstream OFDM channel quality, measuring upstream OFDMA signal quality and testing profiles and IUCs. These features are managed on a per CM basis through DOCSIS media access control (MAC) management messages and are also supported by an expanded set of proactive network maintenance (PNM) capabilities. These features measure the physical layer metrics and make them available for operators to analyze in an external application and use for various optimizations or troubleshooting.

DOCSIS 3.1 OFDM and OFDMA Profiles offer a diverse selection of modulation options. Leveraging the OFDM/A physical layer's benefits involves using different modulation orders for various subcarriers. Profiles can be strategically employed to optimize the cable modem termination system (CMTS) downstream and CM upstream transmissions, ensuring peak performance in varying network conditions. Well-crafted, optimized modulation profiles enhance a channel's ability to withstand interference and boost overall user and system throughput.

To optimize functions like downstream profile creation and upstream profile creation, an operator can run a profile management application (PMA) external to a CMTS. The PMA communicates with a data collector which collects data from the CMs and the CMTS; typically, the data collected is stored on a data lake for multiple applications to use. The PMA processes the input receive modulation error ratio

(RxMER) data and topology information to create optimized and intelligent profile definitions. It then passes on these profiles to the CMTS to configure the CMs to use those profiles.

A D3.1 CM supports two or more OFDM channels, each up to 192 MHz in the downstream spectrum. A modulation profile consists of bit-loading values for active subcarriers, with modulation orders from 16-QAM to 4096-QAM. CMs can support up to four profiles, including Profile A, which is a baseline profile that all CMs use for DOCSIS MAC Management messages and for data traffic. The CMTS can support up to 16 profiles. Currently, most CMTSs support three or four profiles per channel, assigning all supported profiles to all modems.

On the upstream spectrum (US), a CM supports 2 US profiles at a time. US profiles are also known as IUCs in the DOCSIS specification, and these terms are used interchangeably. A CMTS supports seven IUCs (IUC 13, 12, 11, 10, 9, 6, 5). See [D40MULPI]. The CMTS assigns one or two OFDMA upstream data profiles (IUCs) to a CM. The CMTS grants OFDMA bandwidth for data transmission to a CM using one of the CM's assigned IUCs.

A CM starts on the OFDMA channel with IUC 13 (the baseline US profile, similar to Profile A on the downstream). At a later point, the CM is assigned an additional IUC. When CMTS sees US forward error correction (FEC) errors on the secondary profile (or an US probe from the CMTS shows the RxMER to be below a defined threshold) it can remove that profile on a CM and assign it a new profile. This switching mechanism is depicted in an example in Figure 2.

Figure 1 below shows an example US IUC definition. In this case, the IUC is using a “base modulation” of 1024-QAM (meaning most of the minislots use that modulation) and then it has various modulation exceptions, which are different from that baseline level. In simple terms, the PMA algorithm decides on a channel level the frequency zone and modulation rate of exceptions to the baseline. With PMA the baseline of a certain IUC is no longer fixed and will vary based on the actual physical channel condition.




Figure 1 - Upstream Profile Example, Baseline Modulation with Exceptions

The CMTS supports seven IUCs. In the example shown below in Figure 2 the profiles indicate the base modulation used for that profile. In blue are the provisioned but unused IUCs on a CMTS. In orange are the IUCs that are assigned and used, based on a probe measurement). In red is the baseline profile (IUC 13). These profiles are only examples and computed modulations could be higher or lower on a given plant.




Figure 2 - US OFDMA Profile Switching Example

Figure 2 depicts the process that a CMTS follows when it sees US FEC errors. The IUC switch sequence is as follows: Let's assume the CM has stable upstream operation on IUC 6, and at some point in time code word errors (CER) pop up on IUC 6. By default, data communication will switch to IUC 13, as all CMs are assigned IUC 13. The steps in the switching mechanism below correspond to the arrows in Figure 2.

1. Check if the use of IUC 9 is without CER; if no errors use IUC 9.
2. If IUC 9 has errors, continue data communication with IUC13.
3. Check if the use of IUC 10 is without CER; if no errors use IUC 10.
4. If IUC 9 has errors, continue data communication with IUC13.
5. Check if the use of IUC 11 is without CER; if no errors use IUC 11.

The modem sends data on IUC 11 until errors arise or the promotion mechanism is started. Every 60 - 70 minutes (configurable) a promotion mechanism is initiated to see if a switch to a higher IUC can be done without CER within the list of provisioned IUCs. The decision is always based on a probe of the average RxMER of the channel.

2. Profile Calculation Methodology

The capability of a CM to use a particular profile (downstream or upstream) is determined by the RF channel characteristics of that CM. The PMA problem statement is to create an optimal set of profiles given the set of RF channel characteristics (RxMER data). In this paper, the focus is on upstream PMA algorithms. All the upstream algorithms are derived from and built on the downstream algorithms. Hence, this paper reviews both downstream and upstream algorithms in for continuity and as a reference.

2.1. Overview and Evolution of PMA Algorithms

The PMA algorithm analyzes the variances in RxMER signatures among the CMs. The basic idea is to organize the CMs into groups based on similar RxMER signatures (across the whole OFDM/OFDMA channel) and then allocate distinct modulation profiles to each of these CM groups.

There are a few different methods to calculate profiles for a given set of CMs on a channel, introduced in [D31PMA-INTX16]. These include:

- *Profile coalescence algorithm (PCA)*,
- *K-means clustering*,
- *K-means coalescence algorithm (KCA)*.

These methods for designing OFDM/A profiles and choosing appropriate modulation orders help determine the right profiles for CMs and the optimal set of profiles for an OFDM/A channel. The best algorithm recommended at that time was the K-means coalescence algorithm, as the best performing algorithm in terms of execution time and gave results that matched the much slower but more optimal profile coalescence algorithm.

Since then, a newer version of the profile coalescence algorithm has been developed that optimizes the needed calculations, this is described in section 2.4.1, and is referred to as the:

- *Optimized profile coalescence algorithm (OPCA)*.

The cable industry when deploying DOCSIS 3.1 technology, started off with downstream OFDM channels. The PMA algorithms were initially designed for the Downstream OFDMA channels, as only downstream CM RxMER data was available (see Figure 3).




Figure 3 - Downstream OFDM RxMER from One CM, over One Week

As the industry started deploying upstream OFDMA technology, OFDMA US RxMER data started becoming available at the CMTS. On the upstream side, due to noise funneling, the RxMER data for a single CM, (collected at the CMTS receiver) often shows a much more noticeable variance over an extended timeframe (see Figure 4). Many of the figures of US RxMER show some subcarriers to have high values (0xFF); these usually indicate inactive subcarriers within the channel.




Figure 4 - Upstream OFDMA RxMER from One CM, over One Week

On the upstream side, the RxMER data from every CM of the plant often displays similar patterns when captured at about the same time.




Figure 5 - Upstream OFDMA RxMER from Multiple CMs over One Week

The upstream PMA algorithm philosophy then changes slightly from the downstream. Instead of looking at one representative sample from each CM, the idea is to assess every sample data point from every CM over a longer period of time. The basic idea is to organize all of the CM US RxMER samples into groups based on similar RxMER signatures (across the whole OFDMA Channel) and allocate distinct modulation profiles/IUCs to each of these sample groups. [USPMA-SCTE2020] describes some of these algorithms.

The current upstream algorithm uses the *OPCA* algorithm and additionally adds a framework that preprocesses and cleans up some of the data, as summarized in Section 2.4.2 of this paper. This is referred to as:

- *Upstream algorithm 2C*

Based on a number of upstream samples in a service group/node, algorithm 2C was taking a long time to execute, so for the upstream, we devised a new process that uses two steps in a hierarchical fashion to calculate profiles, known as:

- *Two-stage upstream (US) algorithm*

This algorithm uses upstream algorithm 2C in each of the two stages and the two-stage process described in Section 2.4.3.

In this paper we further develop the two-stage upstream algorithm and improve it using pre-grouping of data. This new algorithm is described in Section 2.5 and is known as:

- *Two-stage upstream (US) algorithm with pre-grouping*

2.2. Input Data Samples to the PMA Algorithm

For the downstream direction, the PMA algorithm has the flexibility to utilize either a one-time data snapshot from each of the CMs or pre-processed data, such as averages, minimum values, and percentiles, obtained from multiple historical data captures of each CM. If there are 100 CMs and an operator collects six samples of downstream RxMER per day for each CM, one method could be simply to calculate a representative sample, i.e., an average (or lowest) per sub-carrier RxMER of the last few days of samples, to use for each CM. This means for this OFDM channel, there will be a total of 100 RxMER samples (one for each CM) to analyze. These samples will cluster into multiple groups to create the needed number of profiles (say four or 10). The figure below shows multiple samples from different modems, with two clear clusters that can be seen.




Figure 6 - Downstream OFDM Example with Two Distinct RxMER Signature Clusters

For the upstream direction, the idea as described previously is to use all the RxMER samples over time (i.e., over two or three weeks) from all the CMs. If there are a 100 CMs, and the operator collects 18 samples of upstream RxMER per day for each CM, over 2 weeks, there will be $14 \times 18 \times 100 = 25200$ samples to analyze. These samples will be clustered and reduced into groups to create the needed number of profiles (seven IUCs supported in a DOCSIS 3.1 CMTS). As seen in Figure 5, which shows multiple US RxMER samples from different modems, the best method to create profiles is for the algorithm to pick out the groups and to design the best clusters.

2.3. Optimization Metrics

The objective function for the profile management application algorithm focuses on creating the best set of profiles based on the given set of CM RxMER. The goal is to maximize the network capacity delivered with a given set of profiles, this metric is best captured by the J-value.

2.3.1. J-value

The J-value is a capacity gain metric to evaluate a set of profiles with respect to the lowest profile. See [D31PMA-INTX16] which introduced the concept of the J-value. The J-value denotes the capacity gain a set of profiles makes instead of running all CMs at the lowest profile A.

$$J_{P,A} = \frac{1}{K_A \cdot \sum_{x \in P} \frac{\Phi_x}{K_x}}$$

- $\Phi_x = N_x/N$ is the fraction of users assigned to profile x.
- K_x is the total efficiency of profile x (i.e., sum of the bit-loading values of all of the subcarriers).
- P is the set of profiles and A denotes Profile A.

A more practical way to compare the J-value across different downstream channels is to compare a set of profiles to a flat 256-QAM profile, (common for SC-QAM downstream Channels). So, we have:

$$J_{P,256-QAM} = \frac{1}{K_{256-QAM} \cdot \sum_{\forall x \in P} \frac{\Phi_x}{K_x}}.$$

Similarly, the practical way to compare the J-value across different upstream channels is to compare a set of profiles to a flat 64-QAM profile, (common for SC-QAM upstream channels). So, we have:

$$J_{P,64-QAM} = \frac{1}{K_{64-QAM} \cdot \sum_{\forall x \in P} \frac{\Phi_x}{K_x}}.$$

The objective of profile creation is to find the set of profiles P that maximizes the metric $J_{P, 256/64-QAM}$ for the channel.

The J-value can be useful for operators to understand the capacity of the system of all the profiles relative to 256-QAM or 64-QAM and hence compare sets of profiles across different channels. The J-value indicates the improvement or decline in capacity for the whole system (i.e., a system using a set of profiles vs. a system running at 256-QAM)

2.3.2. A New Simplified Optimization Metric: Capacity Impact

The J-value computation can be somewhat expensive to compute, given processing cycles or time to execute. If we use the J-value calculation on every iteration of the PMA algorithms, the processing time is vastly increased. This leads to the definition of a new metric for optimization, that helps speed up the intermediate calculations, and introduces the concept of “capacity impact.” The capacity of a profile can be thought of as the product of the sum of the bit loading of the profile and the number of CMs in the cluster using that profile. So, the capacity impact of combining two groups would be the difference between the capacity of the two profile groups from the capacity of the combined group.

$$\text{Capacity Impact} = (K_x * N_x + K_y * N_y) - (K_{x+y} * N_{x+y})$$

This capacity impact is a temporary, intermediate computation used only within the PMA algorithm. It is a pseudo capacity number, purely used as an optimization metric internal to the algorithm. See Figure 20 in the Appendix, which shows an example of how this is computed. Figure 20 shows an example with three clusters being combined into two, using the capacity impact of each cluster along with the J-value calculation for each.

While the capacity impact is used within the algorithm iterations, the ultimate definition of capacity gain of a set of profiles is the J-value.

2.4. Summary of Current PMA Algorithms

This section is a summary of the current recommended algorithms for use with downstream and upstream PMA calculations. The core downstream PMA algorithm, OPCA, is also the basis for the upstream PMA algorithms, which add additional steps of data clean up and split up the process into multiple stages.

2.4.1. An Optimized Profile-Coalescation Algorithm

Over time with more mature implementations and refinements, advancements have been made to the profile coalescation algorithm [D31PMA-INTX16]. The optimized profile coalescation algorithm (OPCA) is a new refinement that takes the original profile coalescation algorithm (PCA) and stores the computations from one iteration for the next, dramatically reducing the computations needed in each iteration. The figure below shows the profile coalescation algorithm in action.

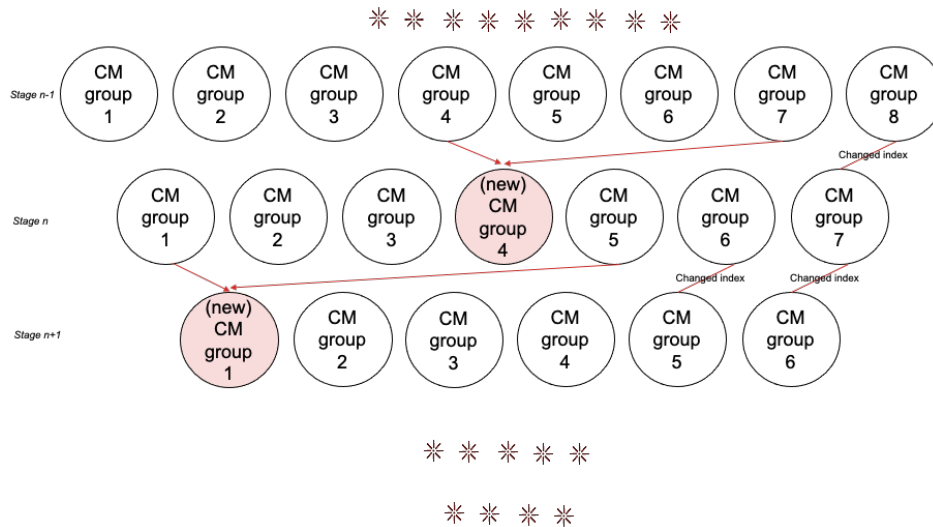


Figure 7 - Profile Coalescation Algorithm Example

In Stage $n-1$ above, all pairs of combinations of the CM Groups 1, 2, 3, 4, 5, 6, 7, and 8 have been computed. For Stage n , CM Groups 4 & 7 get combined into the new CM Group 4, and CM Group 8 gets reindexed as the new CM Group 7. For this Stage n , the computations for all of the pairs of CM Groups except the new group 4 can be reused from the previous Stage $n-1$. All that needs to be newly computed at this Stage n is the capacity impact of the combination of the “new” group 4 with each of the other groups. This reduces the computations significantly.

Stage n-1								
	Group 1	Group 2	Group 3	Group 4	Group 5	Group 6	Group 7	Group 8
Group 1		X	X	X	X	X	X	X
Group 2			X	X	X	X	X	X
Group 3				X	X	X	X	X
Group 4					X	X	X	X
Group 5						X	X	X
Group 6							X	X
Group 7								X
Group 8								

→

Stage n							
	Group 1	Group 2	Group 3	Group 4 (4+old 8)	Group 5	Group 6 (old 7)	Group 7 (old 8)
Group 1		X	X	X	X	X	X
Group 2			X	X	X	X	X
Group 3				X	X	X	X
Group 4					X	X	X
Group 5						X	X
Group 6							X
Group 7							

New Computation needed

Old Computation reused

Figure 8 - Recomputing Impact only for New Clusters

2.4.2. An Upstream PMA Algorithm 2C

Here is a summary of the main algorithm for upstream PMA, dubbed algorithm 2C:

Input Samples: Use all samples of RxMER for all of the CMs (i.e., over the last two weeks)

Threshold: Create an artificial threshold sample using a percentile value (e.g., 1%) of the set of RxMER values. The threshold RxMER value is chosen using the 1 percentile value for each subcarrier across all samples.

Apply the threshold: For every sample available, if any individual subcarrier RxMER within the sample set is below the threshold sample, then bump up those RxMER values (for those subcarriers) to the threshold RxMER level using a ceiling function. This allows the sample to keep its weight in other areas of the channel when applying the threshold. This works better than discarding the sample.

Time Clustering: Find time clusters of RxMER values from all of these samples. Use the clustering algorithm to find groupings across all RxMER values e.g., reuse the optimized profile coalescence algorithm, to find the required number of clusters.

2.4.3. A Two-Stage Upstream PMA Algorithm

Running the PMA clustering algorithm (e.g., OPCA) is not scalable when analyzing large numbers of samples (e.g., 25200 from the example above).

So, the first idea is to modify the upstream PMA algorithm 2C into a two-stage process.

PMA Stage 1: Here the idea is to split up the large data set of samples into manageable chunks of input through which the PMA algorithm can quickly work. As an example, one can split the 25200 samples into 26 groups of 1000 samples each (with the last group containing 200). Then the idea is to run the algorithm (OPCA) on each of these subgroups (of 1000 samples each) and reduce them to 50 clusters.

PMA Stage 2: The clusters from each of the groups from Stage 1 will be used as input for Stage 2. In the above example, there are 50 clusters from each of the 26 subgroups which are used as input for the second stage of processing. This input to the second stage will be $26 \times 50 = 1300$ samples, on which one can run the OPCA algorithm again to reduce it down to the required final seven clusters or IUCs.

The results of this two-stage process are somewhat different based on the order in which these samples are chosen. Further optimization prior to Stage 1 enables better clustering of data, so that not as much information is lost in the two-stage PMA process. We start with pre-grouping the data prior to Stage 1. This new two-stage algorithm with pre-grouping is described next.

2.5. Introducing the New Two-Stage PMA Algorithm with Pre-Grouping

Our new proposed method consists of the following steps. It introduces pre-grouping of data as a preparatory step prior to Stage 1 and Stage 2.

2.5.1. Pre-Grouping of Samples

Pre-grouping organizes all the data in increasing order of average RxMER of the sample, before splitting the whole data set into sub-groups. This puts a lot of similar samples in the same sub-groups leading to better clustering results from the two-stage process.

For splitting a large set of data into subgroups many different types of methods can be used. Arranging samples by the average RxMER is one way that gives a good grouping of similar samples. Other options such as arranging according to time of the day (all samples from the same hour get into one group, etc.) or arranging according to the greatest variation across each sample along with absolute values, etc., or other useful categorization of the data could also be used as different ways to pre-group the data for Stage 1.

We calculate the average RxMER over subcarriers for each sample,

$$r_s = 1/K \sum_k r_{s,k},$$

where $r_{s,k}$ is the RxMER at the s -th sample, $s = 1, 2, \dots, N_s$, and the k -th subcarrier, $k = 1, 2, \dots, K$, with N_s being the total number of samples and K being the total number of frequency subcarriers.

Then we order samples according to sorted r_s ascendingly, and split samples into equal sized groups with G samples per group, and N_g groups.

As described above the two-stage evolution of the original method “upstream algorithm 2C” employed a simpler grouping approach by collecting G samples consecutive in time. This method could lead to the outlier (lower RxMER) samples being scattered over many groups, which would drag down the profile performance of each of those groups in Stage 1, ultimately leading to very low RxMER profiles. Our new grouping approach makes sure that we group similar samples together, i.e., lower-valued RxMER samples in one group and higher-value samples in another group, so that we can minimize the information loss in Stage 1.

2.5.2. PMA Stage 1

Within each group, we apply the optimized profile coalescation algorithm (OPCA) search and generate P profiles, with $b_{p,k}$ denoting the bit loading at the p -th profile and the k -th subcarrier, and w_p denoting the number of original RxMER samples assigned to the p -th profile. Then we convert each bit loading to an equivalent RxMER,

$$\hat{r}_{p,k} = \text{bit_to_rxmer}(b_{p,k})$$

based on a fixed mapping between bit loading and Rx MER.

2.5.3. PMA Stage 2

We collect all $\hat{r}_{p,k}$ together with the associated w_p from stage 1 and start a second stage of profile computations, to perform the OPCA search over $N_g * P$ samples and generate finally Q profiles. In Stage 2, each input equivalent RxMER sample (from Stage 1) is weighted by the number of original samples from Stage 1, so that the PMA optimization procedure keeps track of the system’s capacity gain properly. This is a very important property that a developer needs to make sure to implement.

2.5.4. Different Configurations in Two Stages

We use different PMA configuration parameters for PMA Stage 1 and PMA Stage 2. In Stage 1, we do not impose any practical limitations on modulation orders or number of frequency segments, and we allow a lot more profiles than the final target, such that P can be a lot larger than Q . This is to ensure that we do not lose too much information in Stage 1. Meanwhile, P cannot be too large because we want to keep $N_g * P$ at a reasonable value and to avoid an expensive profile coalescation search in Stage 2.

3. Field Trial Results on VodafoneZiggo Network

In this section, we describe the results from the implementation of the improved two-stage PMA algorithm within the VodafoneZiggo network. First, we first look at the improvements achieved by running the improved PMA algorithm on a dataset of two weeks in Section 3.1. In Section 3.2 we look at the impact and compare the network performance after the integration of two PMA algorithms in the network with parts of the network running on static profiles.

3.1. Implementation of the New Two-Stage US Algorithm with Pre-Grouping

In this first section, we look at the results achieved by running the proposed new two-stage US algorithm with pre-grouping on a dataset from the VodafoneZiggo network. In the next four paragraphs, we dive into statistical data and show how the proposed two-stage PMA performs better in profile generation and J-value.

3.1.1. Dataset

The dataset consists of 14 days of RxMER data polled from 10 interfaces. In Table 1, we show the amount of 3.1 CMs behind each interface. In total, these 10 interfaces have 478 (3.1) CMs. We intend on running the PMA algorithm once every two weeks in the network, using RxMER collected over the past two weeks. This is the reason we choose 14 days of data polling as input for the results in the following sections.

Table 1 - CMTS interfaces used to poll US RxMER data.

Interface ID	Interface Name	Number of CMs
1	34218185	88
2	84553945	82
3	34226377	60
4	84549849	57
5	67780809	42
6	84558041	28
7	67801289	46
8	34230473	24
9	101363913	15
10	84545753	20

3.1.2. Pre-Grouped Samples

In the next two sections, we show results using data collected from one interface, interface 4 (“84549849”), containing 57 modems. Each modem reports RxMER per subcarrier roughly every 30 minutes. In this run we have $N_s = 7708$ RxMER data samples. This number is lower than the expected $57 * 48 * 14$ due to a faulty script that caused a loss of data during the polling.

As described in Section 2.5.1 after sorting samples based on average RxMER, we split them into $N_g = 16$ groups, with $G = 500$ samples per group. We show the grouped RxMER samples in the figures below, using the first two groups and the last two groups to demonstrate an increasing trend of RxMER from Group 1 to Group 16. Each line represents one sample, collected at one point in time from one CM. We can see that the first group contains the worst RxMER samples, with values fluctuating between 0 dB and

20 dB. From the second group onwards to the last group, the RxMER samples are getting better and better, with higher values and fewer fluctuations.




Figure 9 - RxMER Values of the First Group (out of 16 Groups), of Interface 4




Figure 10 - RxMER Values of the Second Group (out of 16 Groups), of Interface 4




Figure 11 - RxMER Values of 15th & 16th Groups (out of 16 Groups), of Interface 4

3.1.3. Two-Stage PMA Results

Given pre-grouped data samples, we apply the two-stage US algorithm with pre-grouping as described in Sections 2.5.2 and 2.5.3, with (target number of profiles) $P = 50$ profiles in stage 1 and $Q = 5$ profiles in Stage 2. We show the generated profiles of three different methods in the figures below, 1) the optimized profile coalescence algorithm (OPCA), 2) the two-stage US algorithm with pre-grouping, and 3) the US algorithm 2c. The OPCA performs an exhaustive search over all samples, which is a global optimization

and guarantees to generate the best possible profiles. We use OPCA profiles as a reference for evaluating different methods. The OPCA in Figure 12 produces optimal profiles, at the cost of an expensive (in time) search that, for this interface, takes 6 minutes. The two-stage upstream algorithm with pre-grouping in Figure 13 generates profiles very close to the optimal ones, at a much lower computation cost of 0.1 minutes. The US algorithm 2c in Figure 14 takes also about 0.1 minutes in computation, but the generated profiles are less optimal.




Figure 12 - Generated Profiles with OPCA




Figure 13 - Generated Profiles with the New Two-Stage Upstream Algorithm with Pre-Grouping




Figure 14 - Generated Profiles with the Upstream Algorithm 2c

3.1.4. Computation Time


We compare the computation time between the proposed two-stage upstream algorithm and the OPCA in Table 2. We can see that the OPCA is costly in computation and can take almost an hour for a large interface with 82 modems. The two-stage upstream algorithm is much faster and takes far less than a minute for any interface.

Table 2 - Computation time

Interface	Number of samples	Number of modems	"Two-stage PMA" (minutes)	"Profile Coalescation PMA" (minutes)
34218185	15403	88	0.22	43.79
84553945	14646	82	0.21	52.88
34226377	9411	60	0.11	10.5
84549849	7708	57	0.09	5.91
67780809	7407	42	0.08	5.3
84558041	5909	28	0.06	2.81
67801289	5691	46	0.06	3.65
34230473	5368	24	0.06	2.13
101363913	3565	15	0.04	0.68
84545753	2272	20	0.02	0.24

3.1.5. Capacity Gains

We calculate the J-value of the two-stage upstream algorithm method using 64-QAM as a baseline. The $J_{P,64-QAM}$ indicates a theoretical system capacity gain of the proposed method over the baseline profile where all modems use 64-QAM. We illustrate the J-values of all interfaces in the figure below. The vertical line at a J-value of 1.0 indicates a performance equal to 64-QAM. The two lines at 1.25 and 1.5 indicate capacity gains of 25% and 50% respectively, over the baseline of 64-QAM. On average the two-stage PMA algorithm achieves a capacity gain of 42% compared to the 64-QAM baseline.

**Figure 15 - J-values of all Interfaces, using the Two-Stage US Algorithm w. Pre-Grouping**

3.2. Impact of the PMA Algorithm Implementation on the Network Performance

In this section, we look at the VodafoneZiggo upstream (US) spectrum and how the PMA algorithm improves its network performance. First, we look at how we utilize the US spectrum at VodafoneZiggo in Section 3.2.1, next we show how we implement the PMA algorithm in an end-to-end automated network configuration setup in Section 3.2.2. We then explain how we set up a field trial in Section 3.2.3 and finally, we analyze the impact on network performance metrics with different scenarios in our network in

Section 3.2.4. In this last section, we compare different PMA algorithms and static exception setups and their impact on the network.

3.2.1. *VodafoneZiggo US Spectrum*




Figure 16 - US Spectrum in VodafoneZiggo Network

The VodafoneZiggo US spectrum is currently set up to be used only up until 65 MHz across the whole country. The frequencies between 19-34MHz are set up to use a single DOCSIS 3.1 channel. The PMA implementation supports extension to other parts of the spectrum in the future. A second channel is set to go until 85MHz (Phase 2). In the future, the spectrum will be increased again with several channels going until 204MHz (Phase 3). The FM part of the spectrum shows ingress that will tremendously benefit from PMA. The lowest part of the spectrum is especially interesting to run with PMA as it is plagued by a lot of ingress. We will discuss the results of our network only in this lower part of the spectrum in this paper. Phase 2 will also run on PMA, but we expect the gains to be lower.

At the moment the PMA algorithm setup runs on an integrated CMTS platform and will be initiated onto a CCAP and RPD setup in the future.

3.2.2. *Building Blocks of PMA End-to-End Pipeline*

In our network implementation, the PMA algorithm represents a pivotal component within the broader context of complete end-to-end automation. We believe that PMA is most effective when it seamlessly integrates into a fully automated network configuration loop. Hence, before we implement the PMA algorithm, we carefully outline the various steps necessary to create a fully automated network configuration loop.

3.2.2.1. *Data Collection*

Data collection is initiated by extracting PNM US RxMER data from the CMTS platform using data polling software. This data serves as the foundation for subsequent analysis. We use 30-minute polling cycles. Due to limitations in the CMTS configurations at the moment a shorter polling cycle is simply not feasible. The software that we use is a CableLabs US data collection script, which polls each CMTS on each OFDMA channel and US RxMER data for all the CMs on it.

3.2.2.2. *Data Storage*

The collected data is systematically stored as a table in a database until the PMA algorithm is run on the data. The data is then stored in an aggregated form for subsequent statistical analysis. Storing the raw near-real-time data for longer than needed should be managed as appropriate. The raw RxMER data is ~300 Mb for the 10 interfaces in this trial over two weeks, so not much of a burden, though if not managed it could cause surprises in the long term.

3.2.2.3. *Profile Calculation*

The PMA algorithm is then run on the data stored in the database every two weeks. The output of the algorithm is stored again in our databases ready for post-processing.

3.2.2.4. Profile Change Decision Algorithm (PCDA)

To evaluate the suitability of the profiles proposed by the PMA algorithm in comparison to existing profiles, an additional algorithm is introduced at this stage. This step, which we will call the "profile change decision algorithm (PCDA)," assesses the necessity of proposed changes within our closed automated loop. Find the pseudo algorithm in Figure 21 in the Appendix.

The logic behind this algorithm is that if the network presents good performance data, we do not want to implement (useless) changes in the network. Performance is measured by channel efficiency and channel impairments and will be explained further in Section 3.2.4.1.

3.2.2.5. Post-Processing

Upon approval of the new profiles by the PCDA, the relevant data is processed and formatted appropriately to ensure compatibility with the CMTS. During this step, we also enforce two network-specific changes to the results of the algorithm.

1. We enforce that IUC 13 remains lower than equal to 64-QAM for the sake of stability. This entails that if the PMA algorithm recommends higher modulations, they are adjusted to 64-QAM. Given that IUC 13 serves as the fallback IUC for the CMs in case of issues, the objective is to preempt channel impairments in the event of ingress.
2. We configure the last three minislots of all profiles to utilize 64-QAM (or a lower) with the intention of preventing interference with adjacent SQ-QAM (DOCSIS 3.0) channel.

3.2.2.6. Reconfiguration of the Network

In the final step, network configurations on the CMTS are adjusted to implement the newly proposed profiles, thereby realizing the intended network optimization. This is done by a virtual machine that has access to specifically change the configurations.

3.2.3. Field Trial Setup

We start by implementing a Proof of Concept (PoC) version of these steps. The PoC involves the manual execution of the aforementioned steps on a limited scale. The PoC serves as a valuable precursor to full automation and allows us to refine our procedures and ensure the viability of the closed automated loop.

In our initial implementation, we manually executed the following steps:

- Initiation of data collection using custom scripts;
- Manual triggering of pre-/post-processing scripts and PMA algorithm execution;
- Profile decision algorithm executed manually following the logic in Figure 21; and
- Manual execution of network configuration changes by engineers from the operations team.

In the period from August 2022 until November 2023, a field trial has been running in the VodafoneZiggo network to trial the PMA algorithm on the network in the aforementioned end-to-end setup. This trial started with a manual PoC and culminated in the current automated setup. The trial area encompasses 10 segments/interfaces of a CMTS located in the Utrecht area. These 10 segments are provisioned with a total of around 455 DOCSIS 3.1 CMs. Due to the changing nature of the network

(rerouting of modems, customers taking or canceling subscriptions, and the introduction of more 3.1 CMs), this number fluctuates over time.

Until now the field trial made use of upstream algorithm 2c. The improved two-stage upstream algorithm with pre-grouping proposed in this paper has just been introduced into the network and only a week of results will be discussed in this paper. In the next section (3.2.4), we will compare the network performance of these 10 interfaces running on several algorithms with the average performance of the rest of the network.

3.2.4. Network Performance

In this section, we introduce the metrics used to evaluate the network performance used to quantify the results achieved by implementing PMA in the VodafoneZiggo network. We then compare these metrics on a network running with and without PMA.

3.2.4.1. Metrics

To measure network performance in our network we look at the following measurements:

- **IUC Usage and Impairments:** IUC usage corresponds to the percentage of time a CM uses a certain IUC to send data over the network. Impairment is the percentage of time a channel is *impaired*. An impaired channel means that the quality of the network is so bad that the CM cannot send data, even over the fallback IUC (IUC13). This event is known in the DOCSIS specifications as a partial service mode of operation. In this case, the CM only uses the DOCSIS 3.0 SC-QAM channels to send data. In our network, the CM tries again sending data on the OFDMA channel after a cool-down period of 70 minutes.
- **(Channel) Efficiency:** Efficiency corresponds to the speed data is sent over the network normalized by channel bandwidth and is expressed in megabits per second per megahertz (Mbps/MHz). The speed is calculated using the conversion in Table 3 and is dependent on the modulation. For completeness, we have also included the RxMER threshold values to support a certain modulation and the bit-loading.

Table 3 - Conversion Table, Modulation to Mbps/MHz

RxMER	Modulation	Bit Loading	Mbps/MHz
43	4096-QAM	12	Not tested
39	2048-QAM	11	Not tested
35.5	1024-QAM	10	9.1
32.5	512-QAM	9	8.2
29.0	256-QAM	8	7.2
26.0	128-QAM	7	6.3
23.0	64-QAM	6	5.3
20.0	32-QAM	5	4.4
17.0	16-QAM	4	3.6
11.0	QPSK	2	1.5

3.2.4.2. IUC Scenarios

Over the course of OFDMA technology deployment in our network, we developed five different scenarios:

- (1) **Two IUCs without Exceptions (implemented network-wide until August)**
This scenario encompasses the network-wide configuration of two IUCs without any exceptions. The implementation of this scenario was prevalent on the entire network until August, and data measurements used in this section were collected in June 2023, which we will call Time Period 1.
- (2) **Five IUCs with Constant Exceptions (implemented network-wide as of September)**
This scenario encompasses the network-wide implementation of five IUCs with constant exceptions (see Table 5). A major part of the network finished transitioning to this configuration in September, and data measurements used in this section were collected in September 2023 (2a) and November (2b), which we will call Time Period 2 and Time Period 3, respectively.
- (3) **Four IUCs with PMA Profiles (implemented in field trial area until August)**
In this scenario, the network is configured with four IUCs determined by the US algorithm 2c. This configuration was confined to the field trial area until September and measurements in this section were collected in June 2023.
- (4) **Five IUCs with PMA Profiles (implemented in field trial area as of September)**
Similar to the preceding scenario, the network is configured with five IUCs determined by the US algorithm 2c. This configuration is confined to the field trial area, commencing in September. Measurements used in this section were collected in September 2023.
- (5) **Five IUCs with PMA Profiles and Two-Stage Improved Algorithm (implemented in field trial area as of October)**
In this scenario, we use the same setup as in (4) but the IUCs are determined with the two-stage US algorithm with pre-grouping. This advanced configuration was implemented in the field trial area commencing in October and corresponding measurements used in this section were collected in November 2023.

Table 4 - Summary of Deployment Scenarios

Scenarios	(1)	(2a)	(2b)	(3)	(4)	(5)
Number of interfaces	13	13	13	10	10	10
Number of configured IUCs	2	5	5	4	5	5
PMA used	No	No (static exceptions)	No (static exceptions)	Yes (US Algorithm 2c)	Yes (US Algorithm 2c)	Yes (Two-Stage US Algorithm with Pre-Grouping)
Time periods	May'23 <i>Time period 1</i>	Sept '23 <i>Time period 2</i>	Nov'23 <i>Time period 3</i>	May'23 <i>Time period 1</i>	Sept'23 <i>Time period 2</i>	Nov'23 <i>Time period 3</i>

Since the start of the end-to-end closed-loop network automation field trial using PMA, we have had three time periods running at different scenarios simultaneously: Until August (1) ran on the whole network and (3) in the field trial area. Since September (2) has been implemented network-wide and (4) runs in the field trial area. Since November (5) runs in the field trial alongside Scenario (2b).

The reason to introduce (2) network-wide is that we realized that static exceptions might be able to reduce most network impairments while increasing capacity. This part of the US spectrum is especially hit by US ingress in the lower frequencies up until around 28MHz. The implemented static exceptions can be found in Table 5.

From the 10 segments on which the field trial has been running, six have been chosen for their extremely bad channel conditions. The reason for adding these interfaces was to test the limits of the PMA algorithm in the worst-case scenarios. The state of the interfaces changes over time and is not the same across different scenarios due to maintenance in the network and a number of CMs added.

Table 5 - Static Exceptions

	Frequency (MHz)		
	19.05-28.2	28.2-33	33-34.6
IUC9	64	1024	64
IUC10	32	512	64
IUC11	16	256	64
IUC12	8	64	64
IUC13	QPSK	16	16

In the next sections, we will always compare the average performance of the field trial (running on PMA; Scenarios 3-5) with the average performance of the rest of the network (running without PMA). We always compare the two (field trial with rest) in the same timeframe. This means that we will look at time period 1: (1) with (3), time period 2: (2a) with (4) and time period 3: (2b) with (5).

3.2.4.3. Results Metric: IUC Usage & Impairments

In Figure 17, Figure 18, and Figure 19, we show the two proposed network measurements for all time periods and the corresponding scenarios (1) through (5). The figures indicate the time spent by CMs in a particular IUC. Note that the IUC definitions are different for each scenario, even though they have the same color code. The figures also indicate the average Mbps/MHz over all IUCs of each scenario. To be able to read Figure 17, Figure 18, and Figure 19 correctly we also need to look at the values in Table 6.

Table 6 - Average Capacity in Mbps/MHz per IUC for 5 scenarios

Scenarios	Without PMA		With PMA		
	(1)	(2)	(3)	(4)	(5)
IUC9	-	6.2	-	7.9	8.1
IUC10	-	5.4	7.1	6.8	7.7
IUC11	-	4.6	6.3	5.7	7.2
IUC12	4.6	3.6	5.1	4.1	6.2
IUC13	3.0	2.4	1.3	1.9	2.2

We start by looking at the IUC usage. In general, we see that most CMs run on the highest available IUC for most of the time. Based on our quality of experience thresholds an average capacity of about 4.5 Mbps/MHz starts becoming critical in terms of service to the customer. We compare the time spent by CMs in an impaired mode or using an IUC below 4.5 Mbps/MHz. We see that without PMA (Scenarios (1) and (2)), about 13% of the time the channels in the network are using IUCs that have an average efficiency of 4.6 Mbps/MHz or less. With PMA (Scenarios (3), (4) and (5)) this percentage reduces to about 3.5%. With PMA implemented in the network, we utilize the good-performing parts to the maximum while limiting the loss of efficiency by the worst-performing part of the network.




Figure 17 - IUC Usage and Average Channel Efficiency, Scenarios (1) & (3)




Figure 18 - IUC Usage and Average Channel Efficiency, Scenarios (2a) and (4)

Next, we look at the percentage of time the channels are impaired. We observe that in the situation where the network is running on two IUCs, scenario (1), the channels in the network are impaired on average 3.1% of the time. For the channels running on PMA (3) we have observed a much lower percentage of impairments of 0.3%. This corresponds to a reduction of 90% in impairments due to PMA. The reason we can reduce impairments further is that we allow PMA to set the modulation in the network to lower modulations (QPSK) for certain minislots/subcarriers.




Figure 19 - IUC Usage and Average Channel Efficiency, Scenarios (2b) and (5)

In the network-wide scenario (2), where we have introduced static exceptions and three more IUCs, the observed network channel impairments have reduced to 0.8%. Compared to the situation in (1) with only 2 IUCs, this corresponds to a reduction of 75%. If we look at the part of the network running on PMA-generated profiles during that same period (4), we witness a channel impairment of 0.3% on average which is a 50% reduction in impairments of the situation in (2). The comparison between the most recent algorithm implementation (5) in time period 3, with the averaged rest of the network (2b) results in a similar conclusion. The average impairments reduced from 0.7% to 0.4%, a 40% decrease. We must add at this point that average values under 1% are very good. The PMA algorithm certainly provides a measurable better stability but is not as significant anymore compared to a time period running on IUCs without static exceptions.

What should be noted is that a reduction in impairments is achieved by the usage of QPSK modulation. This modulation level is very low and provides only very low Mbps/MHz rates of about 1.9. This can cause considerably slow internet connections.

The exceptional results of PMA lie in the combination of very robust impairments combined with good efficiency results. We dive deeper into those in the next section.

3.2.4.4. Results Metric: Efficiency

In terms of efficiency, we observe the blue line in Figure 17, Figure 18, and Figure 19, as the average Mbps/MHz achieved per scenario. Looking at Time Period 1 we observe the average Mbps/MHz for scenario (1) to be 4.3. The average Mbps/MHz for scenario (3) is 6.9. This corresponds to an increase of about 60% in efficiency with the use of PMA-generated profiles. In this case, the generated profiles use US algorithm 2c.

If we look at the second time period with a scenario using five IUCs and static exceptions (2) we witness an average Mbps/MHz of 5.7. This is further improved by implementing PMA-generated IUCs (4) to an average of 7.6. Here again, the profiles are generated with US algorithm 2c. This entails an increase in efficiency of about 33%.

In the last time period with the same scenario in a new time period (2b) and the profiles calculated with the two-stage US algorithm with pre-grouping (5) we see an average Mbps/MHz of 5.8 for (2b) and 7.4 for (5). This is still a considerable increase of 27% in efficiency. Nevertheless, we explain why this is not higher than scenario (4).

There are a few challenges we encountered in this first implementation of scenario (5). First of all, the polling mechanism did not poll data as often as usual (on average data was polled every two hours vs. every 30 minutes). Due to the changing nature of the US spectrum less polling also means a lesser chance to poll a short-lived upstream ingress. This can in turn lead to less conservative profiles. Indeed, we see in Table 6 that the average efficiency is way higher in scenario (5) than in for example scenario (4). Higher modulation profiles are only useful if the respective CMs can actually attain them. The slight increase in impairments witnessed between scenarios (4) and (5) also shows that the higher modulation in IUC 13 will lead to more impairments. Ideally, you want a set of profiles that are conservative enough to prevent errors and efficient enough to allow for high-capacity throughput.

A second issue we face with the first implementation of scenario (5) is that we did not run the PCDA algorithm. This algorithm is an addition to the original PMA algorithms and adds a conservative dimension to the output of the algorithm. Running this algorithm several implementations in a row will provide better results. We can therefore safely conclude that profiles generated with the new Two-Stage US Algorithm with Pre-Grouping are currently performing similarly to the ones generated by the US Algorithm 2c with the opportunity to perform better in the future if the polling cycles go back to the usual levels.

4. Conclusion and Future Work

We have proposed a new two-stage PMA algorithm consisting of two major improvements. First, we identified an improved grouping of RxMER samples that avoids a few bad samples deteriorating the profiles of an entire group. Next, we propose a two-stage PMA that generates profiles close to the optimal ones by the Profile Coalescation method, while at a much lower computation cost that runs at least 100 times faster than the Profile Coalescation method. For upstream PMA our proposed method achieves a theoretical system capacity that is typically 20% to 50% higher than the baseline of 64-QAM.

In the VodafoneZiggo network, we implement the PMA algorithm as part of an automated closed-loop architecture. The results in a limited field trial with 10 interfaces show an increased capacity of about 60% compared to a baseline running on two IUCs and of about 30% compared to a network configured with five IUCs and static exceptions. Stability calculated in terms of channel impairments in time also shows benefits from the PMA implementation. In a network configured with five IUCs and static exceptions, PMA can still reduce the amount of time spent by CMs in an impaired mode by 50%.

The new improved two-stage upstream algorithm with pre-grouping shows considerably better results in terms of J-value and profile generation. In the network, the algorithm does not yet show improved results, though as of the time of writing, we expect that it will improve with fixed data polling mechanism. We for now conclude that it is at least as good as the previously implemented upstream algorithm 2c.

For the next steps, we will deploy the new two-stage PMA algorithm in the field trial via a fully automated pipeline running on AWS, and then scale up the deployment to all CMTSs in the country. We will also monitor several practical metrics (usage percentage per IUC, channel impairment) reported from the field, based on which we can further improve the algorithm and finetune the PMA profile change decision and deployment schedule.

5. Appendix

5.1. Configurations for Two-Stage PMA

Table 7 - Configurations for Two Stage PMA

Configurations for PMA Stage 1	Configurations for PMA Stage 1
<pre>{ "num_profiles": 50, "num_segments": 50, "max_num_profiles": 7, "available_modulation_orders": ["QPSK", "QAM8", "QAM16", "QAM32", "QAM64", "QAM128", "QAM256", "QAM512", "QAM1024", "QAM2048"], "profile_A_for_data": true }</pre>	<pre>{ "num_profiles": 4, "num_segments": 4, "max_num_profiles": 7, "available_modulation_orders": ["QPSK", "QAM16", "QAM64", "QAM256", "QAM512", "QAM1024"], "profile_A_for_data": true }</pre>

5.2. Detailed Example of a Step within the Profile Coalescation Algorithm

The figure below in the Appendix shows three clusters of modems in green blue and orange. The number of modems and the profile definition of those clusters of modems are shown pictorially. Combining the first two clusters leads to a capacity impact of 440 while combining the second two clusters leads to a capacity impact of 240.

5.3. PCDA

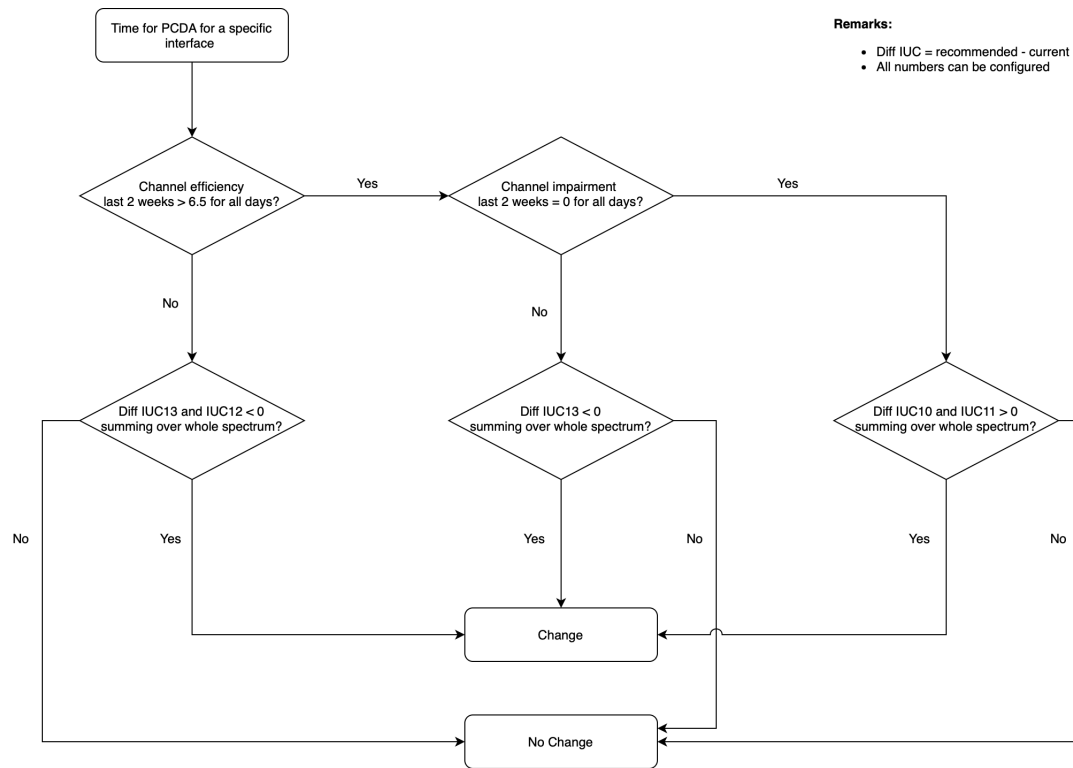


Figure 21 - Profile Change Decision Algorithm (PCDA) Decision Tree

In this example the network is configured with four IUCs (10, 11, 12 and 13). IUC 13 is the default IUC, IUC 10 is the highest data IUC.

Abbreviations

bps	bits per second
Hz	hertz
FEC	forward error correction
CMTS	cable modem termination system
CM	cable modem
CER	codeword error rate
DS	downstream
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
PMA	profile management application
PNM	proactive network maintenance
RxMER	receive modulation error ratio
US	upstream
IUC	interval usage code (alternate name for upstream profile)

Bibliography & References

[PMA-TR-CL] DOCSIS 3.1 Profile Management Application Technical Report CM-TR-PMA-V01-180530

[D31PMA-INTX16] DOCSIS 3.1 Profile Management Application and Algorithms (2016), G. White, K. Sundaresan, CableLabs. <https://www.nctatechnicalpapers.com/Paper/2016/2016-docsis-3-1-profile-management-application-and-algorithms>

[USPMA-SCTE2020] Field Experiences with US OFDMA and using US Profile Management, K. Sundaresan, J. Zhu (CableLabs) J. Fernandes (NOS)

[D40MULPI] *DOCSIS 4.0 MAC and Upper Layer Protocols Interface Specification, SP-CM-MULPIv4.0-I07-230503, 2023.*

[D40PHY] *DOCSIS 4.0 Physical Layer Specification, SP-CM-PHYv4.0-I06-221019, 2022.*

