Impact of CCAP to CM Distance in a Remote PHY Architecture

John T. Chapman, CTO Cable Access & Cisco Fellow, Cisco, jchapman@cisco.com
Gerry White, Distinguished Engineer, Cisco, gerrwhit@cisco.com
Hang Jin, Distinguished Engineer, Cisco, hangjin@cisco.com

Abstract

The Remote PHY architecture currently being specified by CableLabs and developed by multiple vendors has the potential to change the DOCSIS access network significantly.

The basis of the architecture is to move the PHY and the digital to analog conversion from the CCAP device in the hub to the optical node and replace the analog optical link between CCAP and the node with a digital link such as 10G Ethernet.

This removes the distance limitations imposed by the analog link and could in theory enable the CCAP core to be located in a remote data center thousands of kilometers from the CM. The DOCSIS protocol must continue to operate between CM and CCAP both within the normal 100 mile (160 km) DOCSIS range and over this new extended distance.

This paper looks at how this type of system could be deployed and the impacts of increasing the distance between the CM and the DOCSIS cores on throughput and latency seen by the end user.

The paper is in three parts: theory, simulation and test.

Part One explains the R-PHY approach and discusses some potential deployment scenarios. It looks at the impact on the DOCSIS protocol of moving the CCAP core to a remote data center and using an IP network to connect the core to the R-PHY node.

Part Two describes a simulation model developed to investigate the relationship between CM to CCAP distance and DOCSIS performance. The model has been used to investigate the impact of both latency and jitter in the network connecting the CCAP core to the Remote PHY and the paper will present these results

Part Three describes a series of experiments performed to validate the model and confirm the impact of extended distances in a real world deployment. The paper presents these results and compares them with the theoretical analysis and the simulated data.

This paper also discusses on how a Remote PHY based CMTS can operate at distances much further than an I-CMTS.

TABLE OF CONTENTS

INTRODUCTION

How R-PHY Works

Remote PHY (R-PHY) is the next stage in the evolution of DOCSIS and video service delivery. It is a modification of the current CCAP (Converged Cable Access Platform) architecture in which the PHY component is removed from the CCAP platform and relocated into a separate Remote PHY Device (RPD). To better understand the R-PHY approach, we will compare it to the I-CMTS.

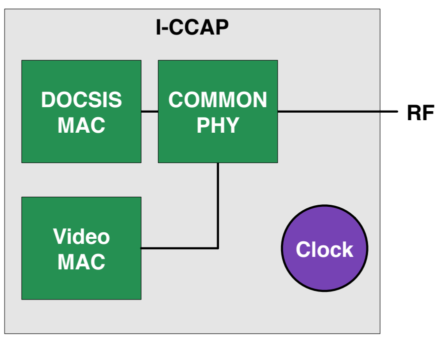

Figure 1 – Integrated CCAP

A very simplified view of a CCAP device is shown in Figure 1. There is a DOCSIS MAC that is a combination of classification, rate shaping and framing for both the upstream and the downstream.

There is also a video MAC that is a combination of encryption, signaling, and MPEG-TS multiplexing. Both the DOCSIS MAC and the video MAC connect into a common PHY. Today this PHY is QAM for DOCSIS 3.0 and video, but will expand to include OFDM for DOCSIS 3.1.

The same components – but arranged differently in Figure 2 – form an R-PHY system. The RPD is connected to the CCAP Core (CCAP minus the RF PHY) by an IP network. The combination of CCAP Core and RPD provide the functional equivalent of the integrated CCAP.

Essentially R-PHY takes the digital interface between the MAC and the PHY components in the CCAP and extends it over the IP network to the RPD using pseudowire technology. That means the I-CMTS and the R-PHY systems are using the same components and the same interfaces.

The resulting system is shown in Figure 3. It is very important to note that the R-PHY Device barely participates in the DOCSIS protocol. The RPD is like a layer 2 gateway device that moves packets from one side the other, providing basic PHY layer conversion

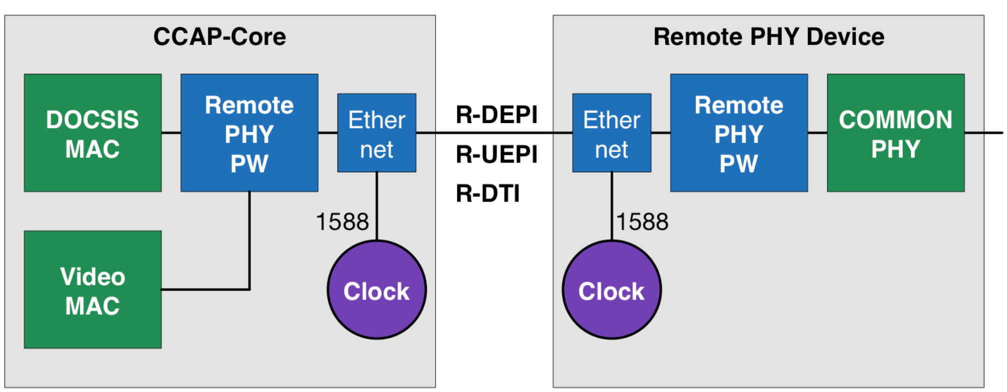

Figure 2 – Remote PHY CCAP

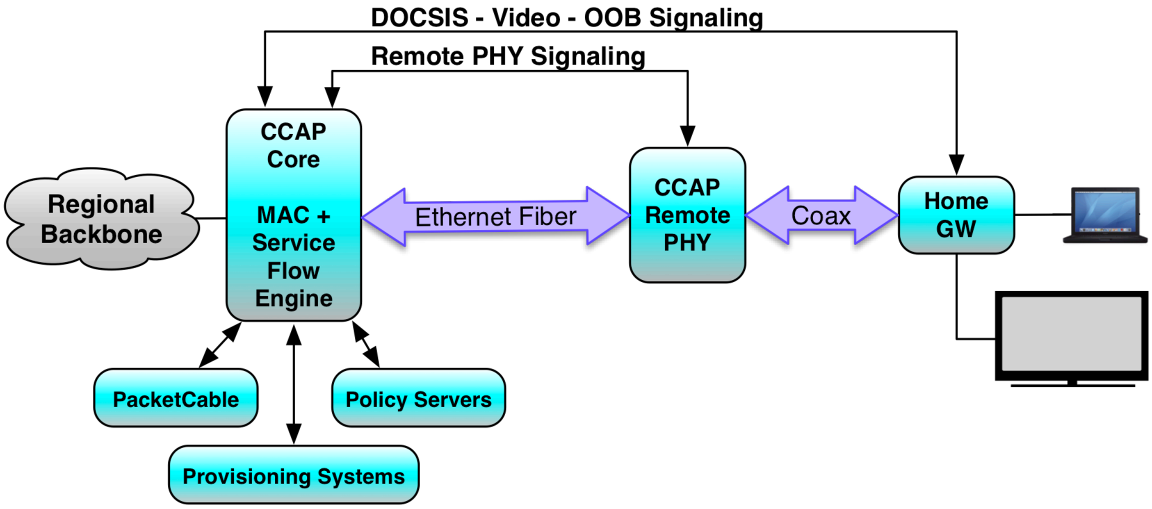

Figure 3 – Remote PHY System

and encapsulation services. There is no scheduler or DOCSIS framing in the RPD. The one exception to this rule is that the RPD is allowed to extract REQ messages out of the upstream so that the REQ messages can be sent with higher QoS than the data.

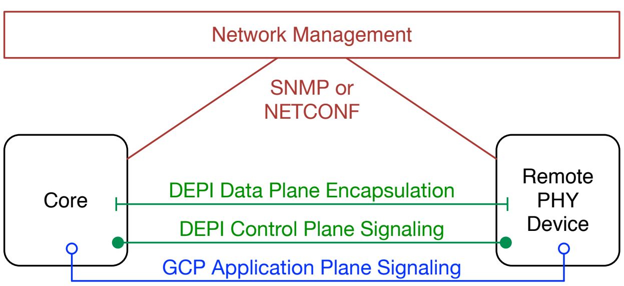

Figure 4 – R-PHY Signaling

The R-PHY architecture uses a set of standardized protocols for its operation as illustrated in Figure 4. The pseudowires (IP tunnels) between the Core and the RPD are managed with a protocol called R-DEPI (Remote Downstream External PHY Interface).

An extension to DEPI is R-UEPI (Remote Upstream External PHY Interface). R-DEPI is based upon the IETF protocol L2TPv3 (Layer 2 Tunneling Protocol version 3). The R-DEPI control plane is used for session establishment and tear down and to manage fail-overs for HA (High Availability). The R-DEPI data plane also includes S-BFD (Seamless Bidirectional Forwarding Detection) to perform data plane diagnostics. R-DEPI supports DOCSIS, Video, and OOB (Out of Band) all using the same techniques and protocols.

To program the ASICs and their corresponding functionality within the RPD, a separate protocol called GCP (Generic Control Plane) protocol is used. GCP is a protocol shim that carries application specific data elements, such as TLVs (type-length-value), over the GCP shim, using a transport such as TCP.

All components are also manageable through SNMP or similar evolved network management protocols such as NETCONF/Yang.

For a more in depth view of R-PHY refer to the References section at the end of this paper.

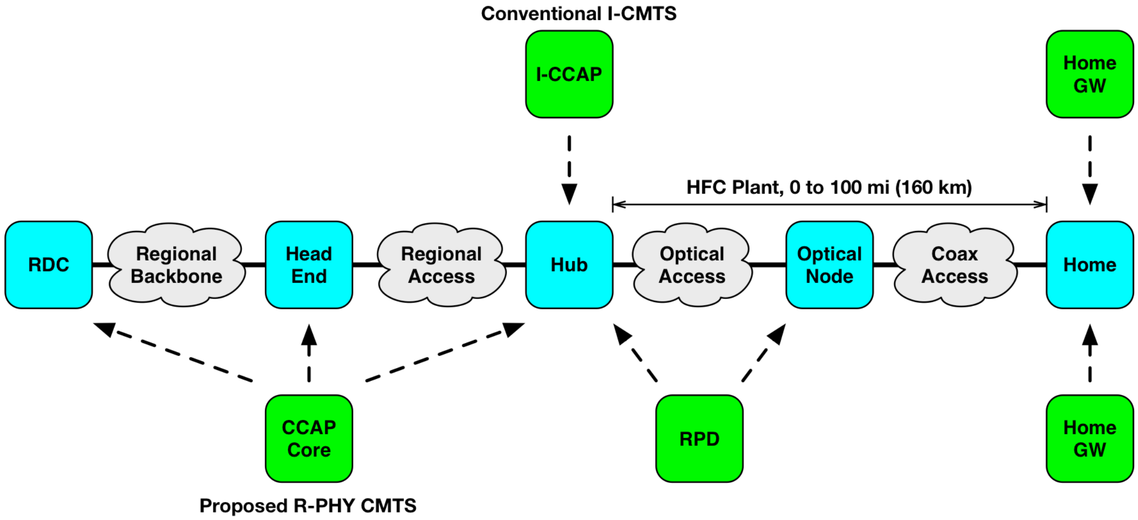

Figure 5 – R-PHY Deployment Model

How R-PHY is Deployed

A reference model for a typical plant of the Cable Operators is shown in Figure 5. A Regional Data Center (RDC) is connected to multiple head ends over a regional backbone network. The head ends in turn are connected to multiple distribution hubs over a regional access network.

Both the regional backbone and regional access networks are typically IP/Ethernet networks running over a digital optical transport. The distribution hubs are connected to optical nodes over the optical access network.

Historically the optical access network has been an analog network but in an R-PHY deployment this can also be IP/Ethernet over digital fiber. From the optical nodes coaxial cable carries an analog signal link to the cable modem at the subscriber premise.

In this reference architecture, the return path fiber from the optical node terminates at the Hub. For this reason, the I-CCAP has been located in the hub. In some systems where the hub has multiplexed and transported the return path fiber to the head end, the I-CCAP can be located at the headend. However, for this white paper, we will refer to the I-CMTS location as being at the hub.

One distinct advantage of the R-PHY architecture is that it removes the distance limitations imposed by the analog link and could in theory enable the CCAP core to be located in a remote data center thousands of kilometers from the CM. Thus R-PHY enables the CCAP-core to be located in the RDC, head end or hub. Similarly the RPD can be located in the Hub or optical node.

The DOCSIS protocol must continue to operate between the CCAP and home gateway over this extended distance. The intent of this paper is to explore the impact of the distance from the Core to RPD on DOCSIS performance and thus to determine which of these deployment scenarios can be used in practice.

Why R-PHY is Cool

An R-PHY System with a CCAP Core and a set of RPDs is functionally equivalent to an I-CCAP. Thus an I-CCAP can still be the best choice if the I-CCAP is located at the hub and feeding a classic analog HFC plant.

The deployment choice of an R-PHY system has to be based on providing value to the customer and solving customer problems that cannot be solved by a traditional I-CCAP. As such, the R-PHY system is positioned to address new and emerging market segments.

Here are a few of the value propositions of a CCAP Remote PHY system.

1. **Enables hub site consolidation**

   Because of the return path fiber limitation, each hub site requires one or more CCAPs. This is not a problem for large hubs where the high density of the CCAP is an advantage but in a small hub may result in underutilized CCAP platforms.

   For operators with lots of small hubs, the total number of partially used CCAP chassis can add up. When the CCAP Core is moved from the hub to the head end in a system with small hubs, better utilization can be achieved and far fewer CCAP chassis may be needed. This can lower the overall system capital costs and operating costs.

2. **Higher bit-rates for DOCSIS 3.1**

   If the DOCSIS 3.1 PHY is located in the optical node, then the transmission impairments from the optical path and the laser and receiver are eliminated.

   That will allow DOCSIS 3.1 to operate at one or two higher levels of modulation. This translates to a 10% to 20% increase in throughput for a given spectrum.

3. **Longer reach**

   Analog fiber cannot reach as far as digital fiber. This distance limitation may prevent deeper fiber deployments.

4. **Enables Ethernet to the Node**

   Replacing the analog optical links to the node with standard Ethernet links enables the operator to leverage the costs, scaling and product development of this much larger market e.g. migration from 1G to 10G to 100G.

5. **Efficient sharing of HFC fiber with FTTH and PON.**

   One of the new major investment areas for Cable Operators is PON (Passive Optical Networks). The current installation plan for PON is to run new fiber all the way from the hub to the home.

   Deployment expense could be lowered on the PON installation if new fiber was installed from the optical node to the home for the PON and backhaul was shared over the same fiber network with R-PHY.

6. **Enables sharing of commercial and residential plants**

   Currently, many cable operators have two plants they own and manage. One is an analog fiber plant for residential and the other is a digital fiber plant for commercial.

If both plants were digital, these could be one common plant. That means that residential or commercial services could be sold and installed anywhere within the combined footprint of both plants. It also means that an investment in one plant is an investment in the other.

7. **Enables HFC fiber to become a full service IP network**.

   Currently the HFC plant is a semi-proprietary network that carries two specific services – DOCSIS and video.

   What if the fiber plant was an IP network? Now that same fiber could carry other services such as wireless backhaul and business services. This would increase the value and usefulness of the fiber plant.

8. **More SG per wavelength**
   In an analog optics system, one wavelength services one node and one SG (service group). In a digital system, the 10GE can service multiple SG since the 10GE can easily multiplex packets.

   If each SG was 5 Gbps, then a 10 Gbps backhaul could support two SG. If there is shared bandwidth across SG due to multicast, then up to four SGs could share one 10GE link. This is one forth the number of forward path lasers that are needed. This provides a 4x scaling difference.

9. **More wavelengths**

   Analog optic systems are typically limited to 16 wavelengths per fiber due to cross talk between the spectrum on each wavelength. Digital optic systems can go as high as 80 wavelengths.

   When deploying deep fiber with more SG, the extra wavelength capacity means that fiber does not have to be pulled all the way back to the hub. This provides a 5x scaling difference.

10. **Better scaling**

    With the PHY separated from the MAC, from the packet engine and from the backhaul, the PHY and MAC can be custom matched to provide a more efficient solution.

    For example, a given backhaul and MAC could support more SG on a DOCSIS broadband only system than on a DOCSIS broadband and Video system. A fixed MAC must be a compromise for all options. This provides a 4x or so scaling difference.

    Note: combining items 8-10 results up to a 80x scaling difference.

11. **Lower plant maintenance costs**

    HFC plants based on analog fiber should be calibrated using a sweep and balance technique twice a year. Digital plants with deep fiber and N+0 passive access do not need the same calibration. Digital optical nodes are simpler to install and operate.

12. **Lower optics costs (10G)**

    With the increase in use of 10GE optics in networking, 10GE optical transmitters and receivers should become less expensive than the analog optical transmitters and receivers.

    The circuit design that supports the

digital optics is also simpler and cheaper than the circuit design that supports the analog optics.

13. **Simpler fiber design rules**

    Analog optics have more complex link budgets as the signal can degrade with distance and the quality of the fiber. Digital systems have simpler design rules. This means that it should be easier and less costly to do field designs.

14. **Replace RF combing with switching**

    In today's hub, there is a large RF combining and splitting network. Ethernet switches could replace all this. This would be more compact and easier to troubleshoot.

    The RF to Fiber to RF conversion would be completely replaced again by Ethernet and Ethernet switches. This will simplify operations, reduce power, cooling and rack space.

15. **Enables CCAP virtualization**

    Virtualization refers to moving the packet forwarding engine and its surrounding services from a dedicated hardware appliance (such as an I-CMTS and moving it into the cloud to run on a generic server architecture.

    An I-CMTS has RF ports and RF ports cannot be virtualized. The R-PHY architecture solves this by moving the RF circuitry to a separate chassis. The remaining circuitry in the CCAP Core can then be virtualized.

How R-PHY Compares

CableLabs has defined an over-all solution umbrella called the Distributed CCAP Architecture (DCA). Within that architecture, there are different approaches to distributing the software and hardware between the core and the remote device. The three recognized approaches are:

- *Remote PHY*: The PHY chip is placed in the remote device. The rest of the hardware and all of the CCAP software remains centralized in a core device.
- *Remote MACPHY*: In a true remote MACPHY, the PHY, the MAC framer, the US scheduler, and half the MAC signaling is placed in the remote device, while the packet classification and switching remains centralized.
- *Remote CCAP*: The entire CCAP is placed in the remote device. There is no core entity other than a standard router.

In addition, a vendor may choose one approach for DOCSIS and a different approach for Video and OOB.

In this white paper, the Remote PHY approach is used consistently for DOCSIS, Video, and OOB. Here are the advantages of the Remote PHY architecture over other architectures.

1. **It's simple**

   The least amount of circuitry is placed in the RPD. This would be the PHY chip as it exists today. None of the CCAP software is placed in the RPD.

   The RPD does of course have a minimum amount of software to operate, but it is not subscriber aware.

2. **It works**

   R-PHY uses a MAC-PHY interface that was initially designed back in DOCSIS 2.0. Back then it was called DMPI, or DOCSIS MAC-PHY Interface. This later led to DEPI and UEPI.

   DEPI was used for M-CMTS (Modular CMTS). UEPI was used for DOCSIS 3.0 chips inside of most if not all I-CMTSs. So, R-PHY is a proven design that is being updated and extended across a longer network.

3. **Complete Interoperability**

   This is really the most important feature. All R-PHY CCAP Core vendors will be able to interoperate with all RPD manufacturers. Any CCAP core will work with any optical node vendor. CableLabs will drive this interoperability.

4. **Complete standardization at CableLabs**

   Only the Remote PHY community has worked with CableLabs to produce seven specifications that will completely describe the Remote PHY operation. Alternate architectures will be described by a technical report, but there will not be any requirements listed. The seven Remote PHY Specifications are:

   - CM-SP-R-PHY: Remote PHY System Specification
   - CM-SP-R-DEPI: Remote Downstream External PHY Interface
   - CM-SP-R-UEPI: Remote Upstream External PHY Interface
   - CM-SP-R-DTI: Remote DOCSIS Timing Interface (based on IEEE-1588)
   - CM-SP-R-OOB: Remote Out Of Band
   - CM-SP-R-OSSI: Remote PHY Operations Support System Interface
   - CM-SP-GCP: Generic Control Plane

5. **Centralized software**

   All CCAP software, including the upstream scheduler, is located centrally and made by the same vendor.

6. **Same feature velocity as I-CCAP**

   The CCAP software for R-PHY is the same software used in the I-CMTS. That means the much larger I-CMTS market will pay for the software for the much smaller and emerging DCA market.

7. **DOCSIS MAC can be upgraded centrally.**

   The DOCSIS MAC contains a huge amount of functionality (the majority in software) – much more than a normal MAC such as an Ethernet MAC – and it is something that changes on an ongoing basis.

   By having a centralized MAC, it can be easily upgraded. With a remote MAC, this is much harder to do, especially if it is a cost reduced hardware based MAC implementation.

8. **MAC and scheduler can be scaled as needed since they are central.**

   It is unknown at this time what the CPU impact will be for the DOCSIS 3.1 scheduler. The DOCSIS 3.0 scheduler also has a CPU impact that is proportional to the amount of data capacity it is managing.

To put the scheduler into the node means designing for the maximum compute resources plus engineering margin now even though those resources may not needed at deployment.

9. **Same consistent approach for DOCSIS, Video, and OOB**

   R-PHY uses the same consistent architecture, signaling, and pseudowire technology for DOCSIS, video and OOB. This provides architectural leverage and operational simplification.

10. **Wi-Fi, EPOC, Cloud-RAN and other access technologies use a similar approach.**

    Many other access technologies arrived at the same conclusion of putting the least amount of electronics and intelligence at the edge. Wi-Fi for example, uses centralized software algorithms for tuning RF performance. C-RAN (Cloud Radio Access Network) goes to the extreme of only putting DAC and ADC with RF at the edge and centralizing the rest. EPOC (EPON over Coax) used a remote PHY in a CMC (Coax Media Converter) with a centralized scheduler.

    The best solutions to difficult problems are usually simple solutions. Centralized solutions historically have been simper than distributed solutions. This is a form of architectural validation.

11. **DOCSIS and Video traffic are already encrypted on the fiber.**

    With a centralized DOCSIS and video MAC, all traffic from the CCAP Core to the CM will be encrypted. If the DOCSIS and video MACs, along with the DOCSIS and video encryption were moved to the RPD, then all the DOCSIS and Video content on the fiber would be in the clear.

    This would lead to needing a new encryption link from the core to the RPD such as IPsec. There would then be a conversion between one encryption scheme to the other at the RPD. This would require two security schemes to protect the data and video, doubling the chances of configuration errors and increasing the vulnerability.

    Why bother with twice the encryption cost, complexity and conversion when you can do it with just one encryption scheme by keeping the payload encryption centralized.

12. **Minimum components in RPD should yield best cost**

    By minimizing the components and the CPU compute requirements in the RPD, the RPD has the opportunity to become more economical. It is also the intent of the R-PHY architecture that the RPD become a single chip device.

13. **Minimum components in the RPD should yield the best availability.**

    Less hardware components means better MTBF. Less software features means less bugs, more stable software and fewer upgrades. All of these increase system availability.

14. **Supported by multiple silicon vendors**

    There is a lot of silicon talent in the industry that can design the PHY chips. However, it requires a different skill set and more investment money and time to get the MAC right.

Recap

- Remote PHY is really a centralized software architecture
- Remote PHY CCAP and RPD systems will be standardized and interoperable.

PART 1: THEORETICAL DISCUSSION

The Case for a Centralized DOCSIS Scheduler

Since the beginning of the R-PHY architecture, there has been a technical debate as to where the upstream scheduler should be. Should it be in the CCAP Core where the rest of the software is or should it be in the optical node with the upstream PHY?

Perhaps we have been looking at this all wrong. Perhaps it is not a technical decision. Perhaps it is a business decision. If an R-PHY system could be built and could work well with the upstream scheduler at either location, are there compelling business reasons to pick one location over the other?

If there was just one really good reason to have a centralized upstream DOCSIS scheduler, that reason would be interoperability. If a centralized scheduler could be made to work, then the manufacturer of the CCAP Core and the RPD could be different.

Any CCAP Core vendor could connect to any RPD vendor. The RPD would not have any knowledge of the upper layers of DOCSIS nor care. The CCAP Core would dictate all system performance.

As a cable customer, do you really want your CCAP purchasing decision determining your optical node manufacturer? In reverse, do you want your purchasing decisions for your optical node to determine your CCAP vendor? Probably not.

If the CCAP Core came from one vendor and the RPD plus a remote scheduler came from another vendor, what is the probability the system would work well together? The upstream scheduler determines all the upstream QoS. That means the downstream QoS would be from one vendor and the upstream QoS from another vendor. Not a good scenario.

And who would you call when something goes wrong or you need a feature added? What kind of system performance, stability and upgrade path would you have with a CCAP Core from one vendor and RPDs from multiple other vendors on that same CCAP core? These are operational considerations that transcend minor differences in performance.

So, will a centralized upstream scheduler work? What is the definition of working well? The big difference in centralized and remote upstream scheduling is the REQ-GNT latency. Lets analyze REQ-GNT latency for that answer.

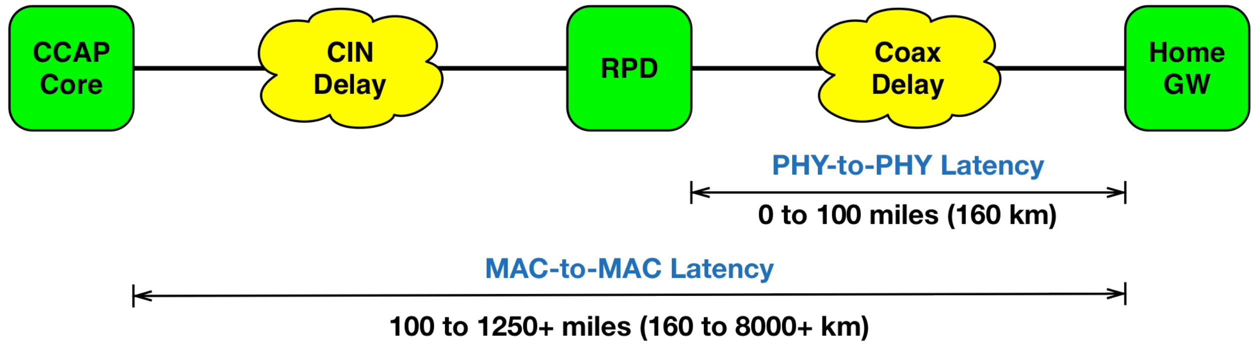

Figure 6 – R-PHY Network Latency Model

Reference Network for Latency Analysis

A simplified, conceptual view of an R-PHY network is shown in Figure 6. This model will be used going forward to discuss the impact of increased CCAP Core to RPD separation. The Converged Interconnect Network (CIN) has been introduced as a single entity to represent the CCAP Core to RPD link.

It is the delay imposed by the CIN that could impact the DOCSIS protocol as the CIN could be anywhere from a single Ethernet link to a multi-hop IP network.

There are two general scenarios for which R-PHY Latency can be evaluated.

1. For distances within the DOCSIS limits of 0 to 100 miles (0 to 160 km).
2. For distances greater than the DOCSIS limit of 100 miles (>160 km)

These two scenarios are very significant. The DOCSIS specification imposes a maximum operating distance from CMTS to CM of 100 miles (160 km) for DOCSIS 3.0. Note that for DOCSIS 3.1, the distance limit has been reduced to 50 miles (80 km).

The first scenario really compares that latency of an R-PHY system to the latency of existing DOCSIS I-CMTS systems.

The second scenario is a new market opportunity that only R-PHY can address. For the R-PHY scenario, the distance limitations that went into the 100 mile limitation are divided out into two distinct categories:

- PHY-to-PHY distance retains the 100 mile distance limitation and pertains to ranging and power levels. Since the CCAP PHY is in the RPD and stays in the hub or node, there is no increase in the PHY-to-PHY distance.
- MAC-to-MAC distance refers to software processes that are dependent upon the delay from the CCAP to the CM. The classic example is the upstream scheduler that is the main subject of this white paper. Other dependencies may include message time-outs.

Before examining if the upstream scheduler will work with R-PHY for each scenario, lets review how the upstream latency impacts the upstream scheduler.

How The REQ-GNT Latency Works

Downstream Latency

The downstream latency seen by DOCSIS packets between arriving at the

CMTS core and being received at the CPE is made up from the following components:

- CMTS packet downstream processing time
- CIN delay due to transmission time on links and switching time in intermediate nodes (refer to Appendix I – Delay Impact of Node to Core Distance)
- RPD downstream processing time
- Transmission delay over the analog network including transmission time, and PHY impacts such as FEC and inter-leaver delays
- CM downstream processing
- Transmission time on the CM to CPE link

Of these components, only the CIN delay varies as the distance between the CMTS Core and the CM change. Downstream DOCSIS operations on a packet are relatively straightforward (classify, apply QoS, add DOCSIS header, transmit) with the impact of CIN delay a linear function.

Thus the measured downstream latency should be equal to the CIN delay plus a fixed processing overhead as shown in Figure 7. The processing delay of 1.7 ms used in this plot was measured for 1518 byte packets using an integrated CMTS.

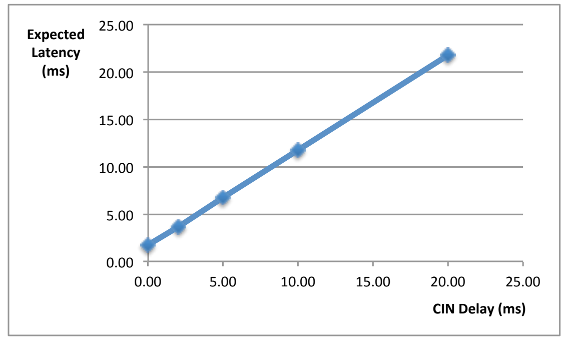

Figure 7 – Expected Downstream Latency vs CIN Delay

The CIN delay for a 100 mile plant is roughly 800 usec. The graph is calibrated in delay that can be later converted to distance.

Note that packets generally always see this delay in Figure 7 since the packets have to travel from the CPE to the RDC on their way to an Internet peering point. So, the significance here is that although the overall network delay is the same, the delay on the regional link is smaller and the delay that the CIN sees is greater as the CCAP Core moves away from the RPD.

Upstream Latency

The upstream latency seen by DOCSIS packets between being transmitted from the CPE and being transmitted into the core network from the CMTS core is much more complicated due to the DOCSIS upstream scheduling.

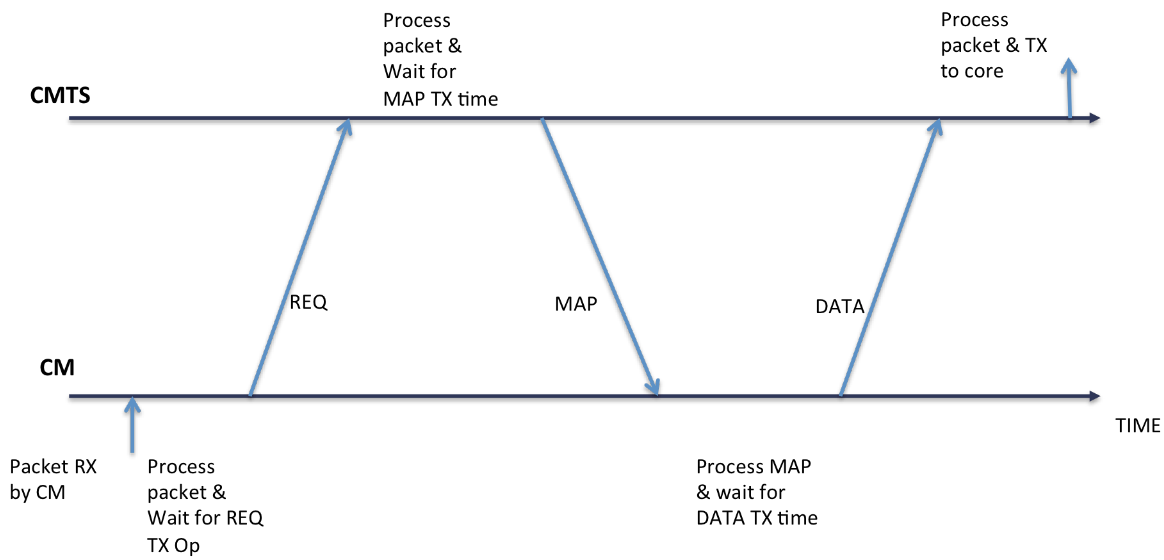

Figure 8 – Upstream DOCSIS Data Transmission

The upstream latency is made up from the following components and shown graphically in Figure 8:

- Transmission time on the CPE to CM link
- CM upstream processing
- CM waits for a transmit opportunity to send a request
- CM sends a bandwidth request to CMTS via RPD and CIN
- CMTS processes the Request and generates MAP with a data grant
- MAP transmitted to CM via CIN and RPD
- CM waits for transmit opportunity to send data
- Data packet transmitted to CMTS core via RPD and CIN
- Packet processed by CMTS and transmitted to core network

From Figure 8 it can be seen that approximately three transmissions traverse the CIN for each packet sent, the bandwidth request, the MAP and the data packet itself. (Approximately because requests are in bytes which could be part of a packet or multiple packets)

Thus the measured upstream latency would be expected to be equal to at least three times the CIN delay plus a fixed processing overhead plus a potentially variable overhead related to the DOCSIS protocol.

Of these components, only the CIN delay and possibly the DOCSIS protocol overhead vary as the distance between the CMTS Core and the CM changes. A first order approximation could assume the DOCSIS protocol delays to be invariant with CIN delay and model the upstream latency as a fixed offset plus three times the applied CIN delay as shown in Figure 9.

This uses 6ms as an approximation for the fixed offset and DOCSIS protocol delay (the 6ms value is based on measurements of upstream latency taken from I-CMTS systems).

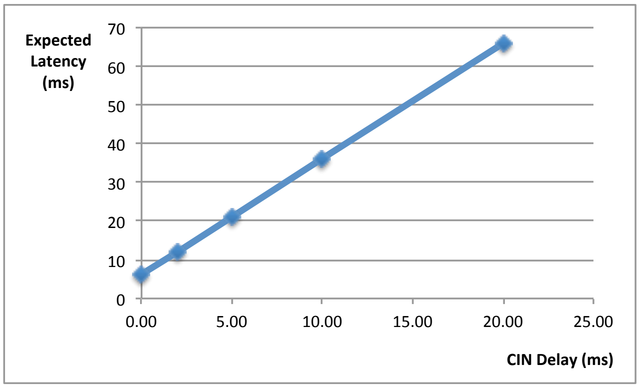

Figure 9 – First Order Approximation of Upstream Latency as a Function of CIN Delay

However the interaction between the CIN delay and the DOCSIS scheduling is potentially more complex than in the downstream case as some of the processing and wait times within the DOCSIS protocol can overlap with CIN transmissions. Thus Figure 9 should be considered to be a rough guide at best.

A Gut Check

Do upstream packets on a DOCSIS system really see a 6 ms delay? Does that make sense? Let's do a gut check by looking at the problem completely differently. Lets look at it from a CM viewpoint and how the CM interacts with MAPs.

Today, for DOCSIS 3.0 systems, a CMTS typically sends a bandwidth MAP to the CM roughly every 2 ms. This is shown in Figure 10. The CMTS also sends the MAP in advance of when it is needed. So at any one time, there is a one MAP in use and one MAP on deck that is about to get used.

Thus the most efficient a DOCSIS system gets is a REQ-GNT cycle that is every second MAP which is approixmately 4 ms. This can vary quite abit because MAPs are rounded off to packet boundaries, so they can be longer or shorter than 2 ms. Also the REQ and GNT can be anywhere in the MAP which adds another +/- 2 ms of uncertaintly.

But, if we stuck with 4 ms as an average, and then allowed for 1 ms for queuing and process time in both the CM and CMTS, we would get the 6 ms of minimum upstream latency for a packet that is seen in the test results. This is a minimum latency as the REQ could see contention and have to be resent which would result in a much larger request latency.

Poisson Distribution of REQ-GNT Times

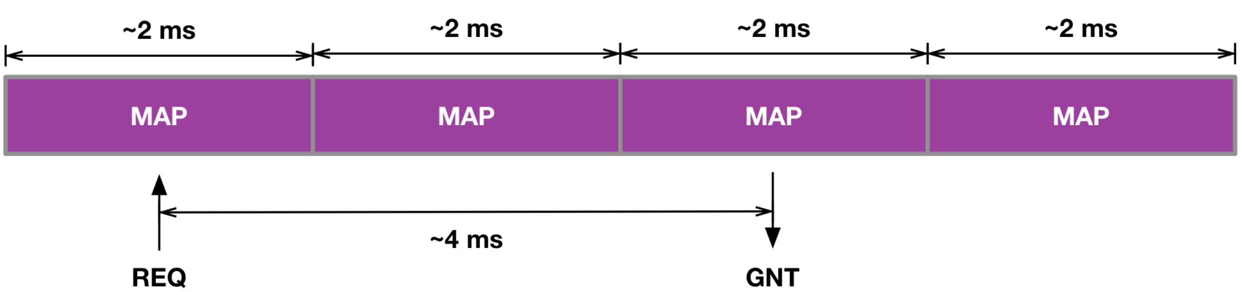

Figure 10 – MAP with one REQ-GNT cycle

In this basic gut check, the REQ-GNT delay is about 6 ms. However, this is really a minimum delay. During a loaded upstream, there will be contention on the REQ slots. This contention will cause the REQ-GNT delay to increase, perhaps well beyond the 6 ms.

The increase in delay might not happen if the requests are sent in a piggyback. However, since Piggy backing only occurs in tight bursts that are within 4 ms of each other, there can still be many restarts of DOCSIS flows.

One could argue that the increase in request-grant delay actually helps to throttle TCP connections during congestion. Delay is a natural part of the networking process. The trick for a protocol such as DOCSIS in a congested situation is to continue to work in a well-defined manner and do so at wire speed.

With 6 ms (or a number such as this) and random occurances of contention with random arrival times of REQs, the REQ-GNT delay most likely will follow an offset Poisson Distribution.

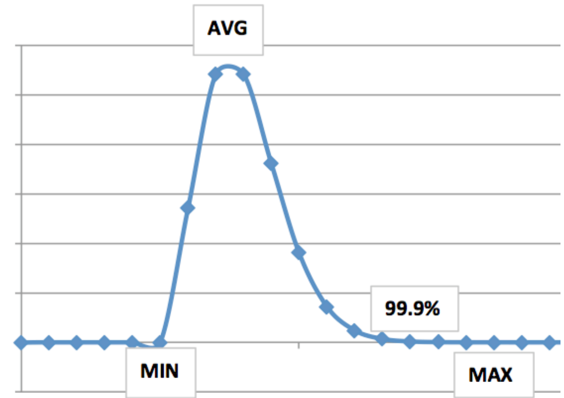

Figure 11 – Offset Poisson Distribution

In this distribution, there would be a minimum delay which is the distribution offset. This might be a value like 6 ms. Then there would be an average value, such as 12 ms. Then there would be a max value, like 50 ms.

The tail of this distribution may need to be modified to be more exponential to mirror the exponential backoff algorithm of the DOCSIS Request algorithms, but that is a refinement for another time. It is more the concept that is important right now.

DOCSIS Pipelining

How does DOCSIS work with a 6 ms or higher upstream packet latency? The answer is the DOCSIS does what all high speed

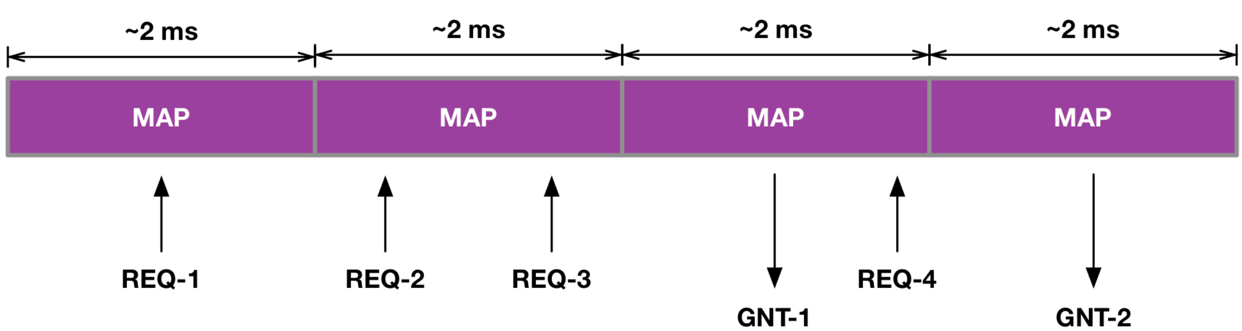

Figure 12 – MAP with pipelining

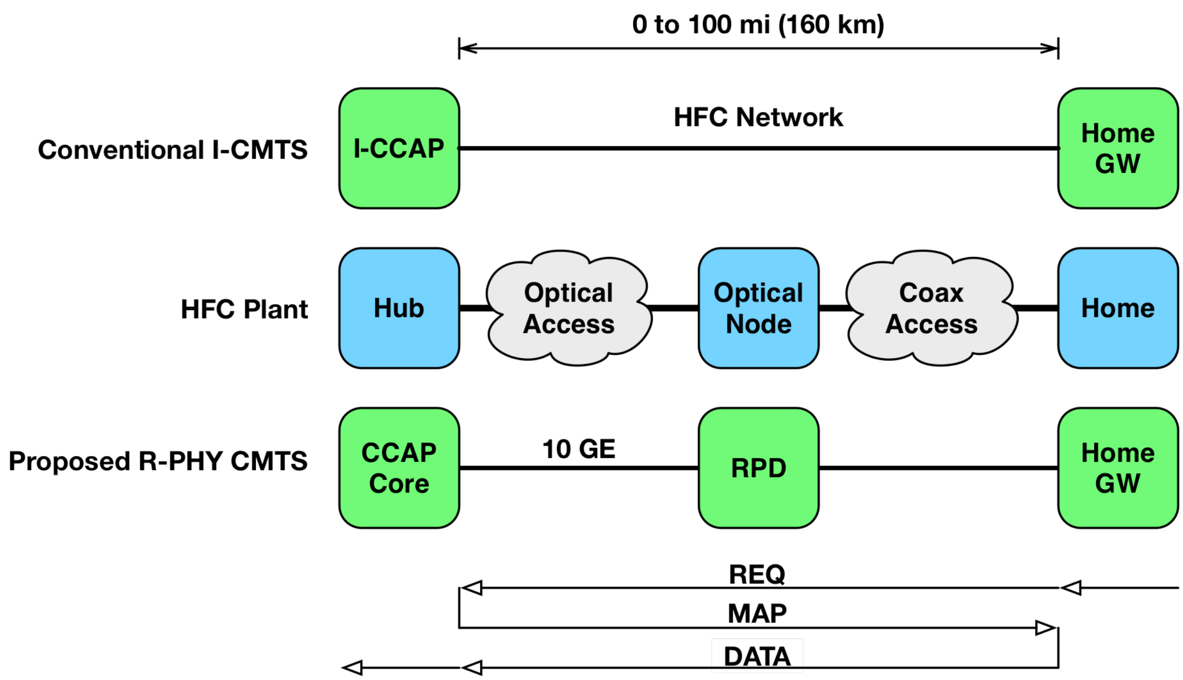

Figure 13 - Latency Scenario #1: 0 to 100 miles

systems with latency do – it pipelines.

DOCSIS allows for multiple outstanding requests. This is pipelining and is shown in Figure 12. The next operation is started before the prior one is complete. Further, the REQ and GNTs are byte based, so they can get larger to accommodate a build up of packets in the CM.

The net result is that once the DOCSIS system gets up and running, it is able to achieve wirespeed. In fact, it can achieve wirespeed for very large values of delay. DOCSIS needs to do this for an I-CMTS because on a busy upstream there can be lots of REQ latency due to contention, but the upstream still can operate at saturation.

Note that given this baseline minimum limit of around 6 ms based upon MAP usage, making the plant considerably smaller, such as putting the scheduler in the optical node, may not result in any signifigant performance improvement.

Even if the length of the MAP interval were made smaller, there is a MAP advance time that would dominate and may just require more MAPs between the REQ-GNT time so that the overall delay would be about the same. If there was any improvement, any real increase in latency would only be on the order of 1-2 ms.

Meanwhile, with or without such an improvement in latency, the throughput for either case would still be wirespeed due to the DOCSIS pipelining.

Scenario 1: 0 to 100 Miles

In scenario #1, the I-CCAP is located at the hub site because this is the point in the

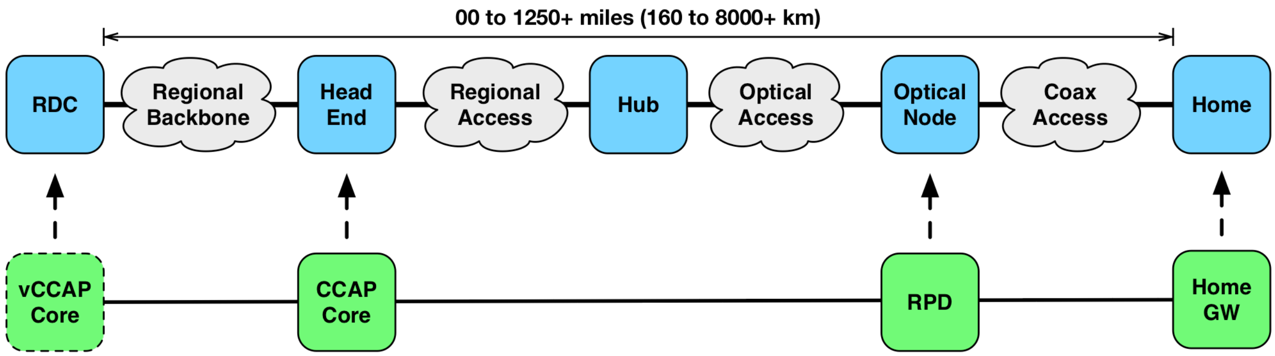

Figure 14 – Latency Scenario #2: Beyond 100 miles

network where the return path fiber terminates. The CCAP Core is also located at the hub site. This is shown in Figure 13.

This means the distance from the CM to the I-CMTS is the same distance as the distance from a CM to a CCAP-Core.

The same PHY is used for an I-CCAP as for the R-CCAP. It is just that PHY is located in a different location. In fact, the 10G MAC-PHY interface is also the same for both the I-CCAP and the R-PHY. This is by design. Both systems use the same chip sets and software protocols.

The REQ-GNT traffic still traverses the same path for both the I-CMTS and the R-PHY system. Despite the location of the PHY (central or remote) the REQ-GNT traffic remains essentially unchanged.

So if the network between the CCAP Core and the RPD were a perfect network with no congestion or jitter, then the round trip REQ-GNT latency should be the same for I-CMTS and R-PHY for DOCSIS networks under 100 miles (160 km). To restate, there is little to no impact due to network latency as that was already built into the DOCSIS network budget.

There is one concern at this point that is raised over network jitter. Network jitter is generally caused by queue build up, which can occur if the network is congested. In the CIN, this can cause an increase in latency.

To allow for that, the designers of DEPI/UEPI allowed the REQ and MAP traffic to be extracted from the data traffic and placed on a separate pseudowire that would allow a higher quality of service to be assigned to the REQ and MAP traffic.

Thus, in a properly configured network, even though the data traffic may see higher latency due to network jitter, the REQ and MAP traffic should not as they will not see that queue build up.

In summary, there is very little differences between and I-CMTS and R-PHY system. What differences that exist are quickly masked by the 6 ms latency of the DOCSIS upstream and corrected for with the DOCSIS pipelining.

Scenario 2: Beyond 100 Miles

So now it gets really interesting.

If DOCSIS has this scalable pipelining built into it that can address the MAC limitations, and the R-PHY architecture has the same the PHY limitations, how high a latency could the REQ-GNT loop tolerate, and how far a distance could that distance be? What would that application be for?

One of the unique features of the R-PHY architecture is the ability to re-locate the CCAP-Core. An earlier example of how this feature can be used is hub site consolidation where the CCAP Core is placed at the head end and the R-PHY devices are placed at the hub or the node.

This allows for fewer CCAPs to be used in a system where there are many smaller hubs. The distance from the head end to the CM, though, may be larger than 100 miles.

Another emerging application for the CCAP Core is to virtualize it. When a CCAP Core is virtualized, then the software that was running on a dedicated hardware appliance such as the CCAP Core hardware is moved to a generic server environment in a data center. That data center may be located even farther way than the head end facility. These two deployment options are shown in Figure 14.

Note that a virtualized CCAP (vCCAP) system by definition must use a server complex for the forwarding plane. It is not sufficient to take a bunch of smaller CCAP devices and connect them together with network management to manage them as one large CCAP. That latter technique is defined by CableLabs as a CCAP Management Abstraction (CMA).

As the latency of the REQ-GNT loop gets excessive, so does the REQ-GNT delay. For example, for a CMTS to CM latency of 10 ms, the REQ-GNT delay could be 32 ms. (three times the delay plus 2 ms of the CMTS and CM). 10 ms is about 1250 miles (2000 km). As you will see in the simulation and test results, DOCSIS with no upstream scheduler modifications still works at these loop lengths and can saturate the upstream link.

However, it is possible to do better. Additional pipelining technology can be used by having a centralized scheduler and measuring the REQ-GNT loop,. Essentially, it becomes possible to proactively send grants to the CM.

This means that the CMTS can use predictive behavior and not wait for each REQ. The net result is that a 10 ms delay from CMTS to CM (2000 km) can result in only a 10 ms or so effective REQ-GNT time. This is incredible.

This means that at extremely long lengths, even in excess of 8000 km, it is possible to operate a centralized scheduler and not incur additional packet latency.

As reference, 8000 km is 1.7x the distance from San Francisco to New York City.

In the following simulation and test sections, we will look at the impact of proactive granting on performance.

Recap

- The upstream latency of a I-CMTS system is about 6 ms
- For the current DOCSIS operating limits of 0 to 100 miles, the additional latency introduced by R-PHY is almost insignificant compared to the existing DOCSIS latency.
- Pipelining techniques built into the DOCSIS protocol make DOCSIS highly tolerant of excessive latency.
- Making the REQ-GNT loop smaller than 100 miles by putting the scheduler

in the optical node may not result in any increase in performance

- R-PHY can actually operate with a central upstream scheduler at very large latency and loop lengths. Increases in REQ-GNT latency beyond normal one-way latency can be managed with advanced scheduler techniques.

PART 2: SIMULATION

To drive beyond theory, a simulation model was established for the DOCSIS US scheduler and simulations were run with typical R-PHY deployment scenarios to assess the impacts of the CIN delay on the latency of US traffic.

The model includes all the essential DOCSIS MAC function blocks of the system relating to the US traffic latency calculation.

Figure 15 is the high-level diagram showing all the function blocks included in the simulation model. A brief description is given for each function block as follows.

Simulation Function Blocks

Ingress Packets

The ingress packets at the network interface arrive randomly in time and sizes. The average throughput is calculated over time, and packet size and arrival interval are adjusted in such a way that the average throughput equals to the pre-configured one for the CM.

Queue Management

In the queue management block, all the ingress packets are time stamped for the purpose of the latency calculation. The packets in the queue are managed as a FIFO. In the simulation, max buffer size is set as 100 KB. Packets will be dropped if the buffer is full. This may occur in the cases where high or close to 100% CMTS loading is configured.

Request Management

The Request Management block manages all the requests: it checks the buffer length and sends the requests with proper requested sizes. It time stamps all the requests that are sent out. Whenever an updated MAP ACK time arrives, the Request management updates the outstanding requests accordingly to ensure no duplicated requests are sent for same packets.

Request Reception

Requests are processed in the request reception block. The CMTS MAP capacity (the total bytes that can be carried in each MAP 2ms interval) is checked, and if the request can be accommodated, the request

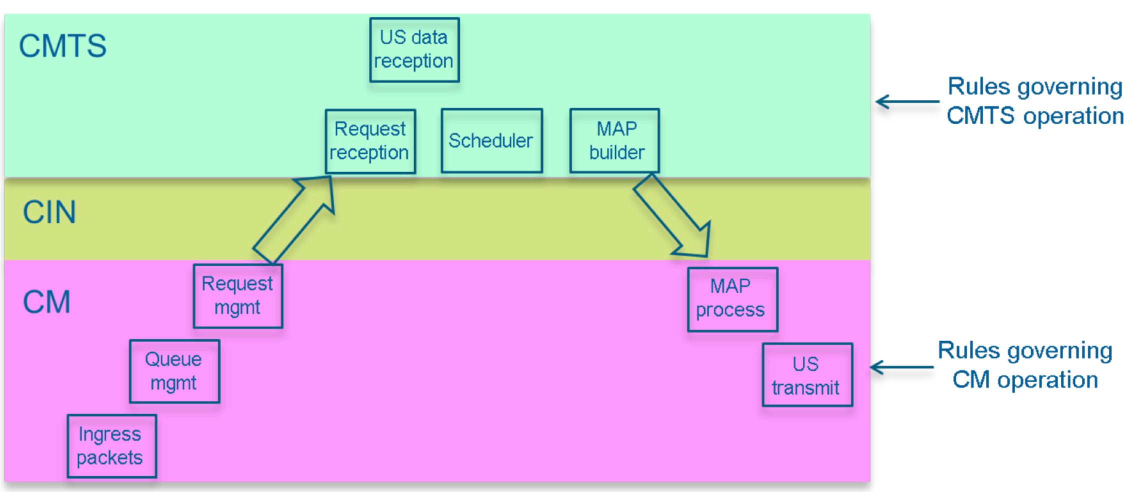

Figure 15 – Simulation Model and Function Blocks

will be granted. The actual grant time will have at least 1 MAP delay to simulate the process delay at the CMTS.

The granting could be partial if there is not enough capacity to fully accommodate the request. Moreover the grant could be a few MAP ahead if the next MAP runs out of capacity.

US Data Reception

All the normal data packets are processed in the Data Reception block. The US latency is computed for each packet by comparing the packet's ingress time stamps and the packet's arrival time at the CMTS.

The US latency is defined as the time difference between the packet ingress time at the network interface of the CM and the packet arrival time at CMTS.

Scheduler

The US scheduling algorithm is running in the scheduler block. For each request received, the scheduler computes the size and the time of a grant if a grant is warranted. The grant is based on both the requested sizes and the MAP capacities. Two scheduler algorithms are simulated:

1. Baseline algorithm

   In this algorithm, a grant is given only when a request is received. This request then grant literately follows DOCSIS MAC US scheduling protocol. In this request-then-grant process, the US latency consists of three parts:

   - the propagation delay of the request traversing the network from CM to CMTS,
   - the delay of the MAP building and the MAP traversing the network from CMTS to CM, and
   - the packet propagation delay from CM to CMTS.

2. Proactive grant algorithm:

   A proactive grant algorithm reduces the latency by activating mechanisms that help establish a pipeline of requests and grants. Depending on the actual mechanisms implemented, the proactive grant scheduling may result in extra BW requests from CMs, or extra grants to CMs, or a combination of both.

   The extra requests that are sent by a CM in addition to the normal requests are denoted as proactive-requests. With the proactive grant algorithm, a CMTS can give out grants to CMs based on received requests, received proactive-requests, or without receiving any requests.

   The grants that are given out by CMTS to CMs in the first case are denoted as reactive-grants; the extra grants that are given out by CMTS to CMs in the two latter cases are denoted as proactive-grants.

The logic behind the proactive grant algorithm is that these proactive grants will allow CMs to bypass the conventional steps of the request-grant mechanism, and send packets right away without going through the normal three stage process (request-grant(MAP)-send), so that US latency can be reduced.

To improve the utilization of the proactive grants, the following rules are observed when the CMTS gives proactive grants:

- Provide more proactive grants to CMs with higher US traffic
- The sizes and intervals of proactive grants are dynamically adjusted based on CM heuristic performance metrics (US utilization ratio, latency, etc.)
- Proactive grants and reactive grants (grants resulting from normal requests) are managed holistically, based on the channel utilization and the total bandwidth (proactive and reactive) assigned to a given CM.

MAP Builder

The MAP builder puts all the grants in place and adds the ACK time. The designated transmit time, which includes the MAP advance Time is added for each grant.

MAP advance time is computed as follows:

- MAP_advance_time = cm_t+cmts_t+CIN_t+ds_phy_t+hfc_t+ margin

Where

- cm_t is the process time required by CM
- cmts_t is the process time (MAP building) required by CMTS
- CIN_t is the propagation delay of CIN (one way)
- ds_phy_t is the processing time required by R-PHY
- hfc_t is the round trip delay of the HFC network
- margin is an extra safety margin

The purpose of the MAP advance time is to allow the CM to receive the MAP ahead of time and have enough time to prepare the US transmission at the designated time. The MAP advance time includes the time budget to allow for the HFC US transmission to arrive at the R-PHY in time.

Unless specified otherwise, all the simulation results are obtained with the following values:

- cm_t=1 ms;
- cmts_t=1 ms;
- ds_phy_t=2ms;
- hfc_t=1 ms
- margin=1ms

In the simulation, we also set us_phy_t = 2 ms, where us_phy_t is the time required for PHY to process US traffics. us_phy_t is required to compute packet US latency.

CM MAP Process

There are three main functions in the MAP Process block:

a) Check the ACK Time in each MAP received and update the outstanding requests accordingly;
b) If there is a grant in the MAP for a CM, the CM prepares the data to transmit at the designated time;
c) Remove the packets from the data queue if they are transmitted.

Simulation Results

The simulation is an event driven process, and runs with an emulated time clock. The time clock runs with a 0.2 ms time increment (it can be set to other time increments, 0.2 ms gives a sufficiently fine time granularity with a reasonable simulation speed).

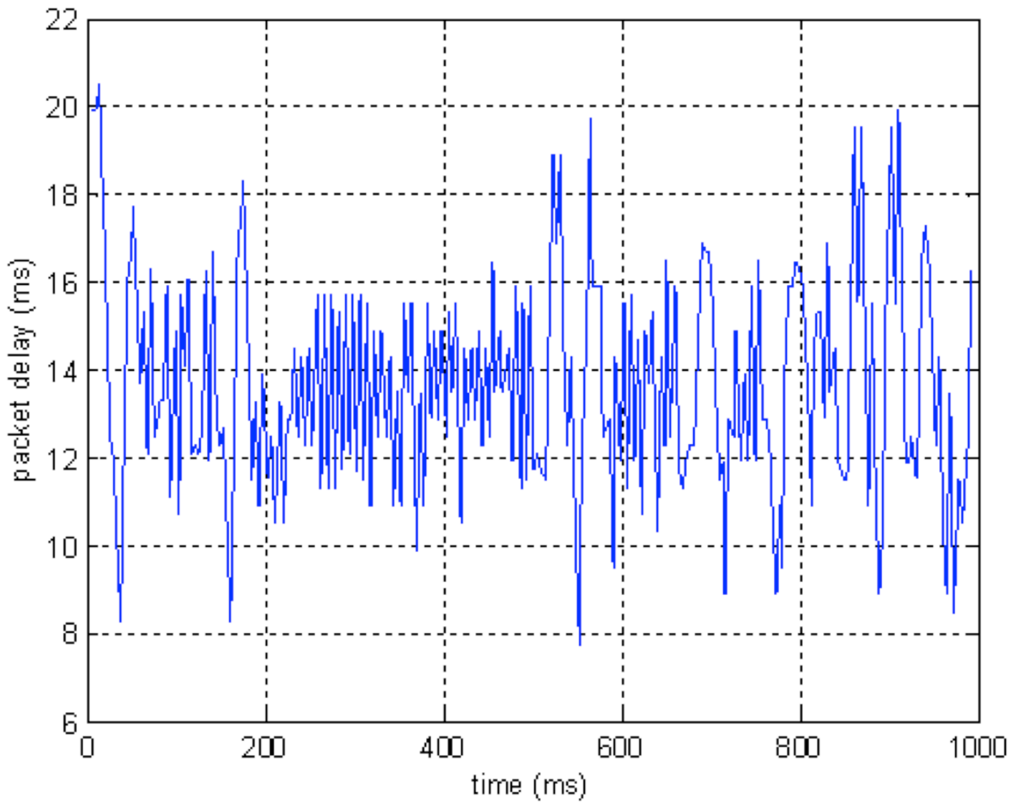

Figure 16 – CM Packet Delay, REQ-GNT scheduler, 2 ms CIN, 30 CM, 90% Loading

All the MAC function blocks execute at each time increment. The simulation runs for a period of one second. However, the US latencies are computed only with packets which arrived at the CMTS after 500 ms to avoid any data abnormalities in the beginning of the simulation that may be due to the simulation initial setup.

A total of 30 CMs are included in the simulation. All the CMs are set to have the same throughput and priority. In the case when the max CMTS capacity is reached (for example, for 90% loading case), the CMTS will throttle the throughput equally for each and every CMs. The bandwidth throttle is done using a first-in-first-serve algorithm, and the sequence of the CMs is selected randomly.

Algorithm 1: Baseline algorithm

There are two scheduler algorithms investigated in the simulation:

1. the baseline algorithm and
2. the proactive grant algorithm.

The simulation results for the baseline algorithm will be presented in this section to establish a baseline for performance (latency) comparisons.

The simulations are first run for the case with 2 ms CIN delay. 2 ms CIN delay corresponds to 400 km in distance and is a good representative value for the majority of deployment scenarios.

The total number of CMs = 30, with each CM having a throughput = 3.64 Mb/s.

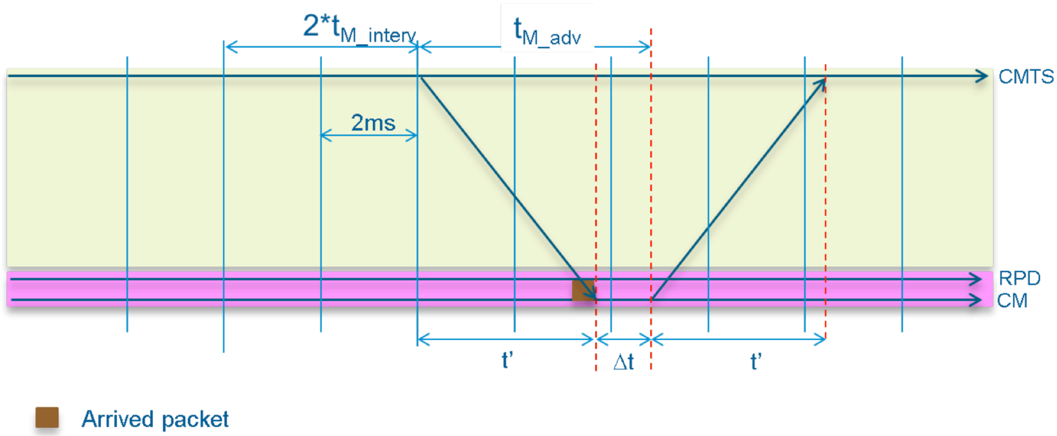

where $t_{M\_interval}$ is the Map interval(2ms), t_{M_adv} is the Map advance time, t' is the DS or US delay (assume symmetric). $\Delta t = t_{M\_adv} - t'$.

Minimal US latency $= \Delta t + t' = t_{M\_adv}$

Figure 17 – The minimal US latency with a pipeline of request and grant

CMTS loading is set to 90%. The simulation results on US packet delay are given in Figure 16.

It is evident from the simulation results in Figure 16 that the minimal delay is 8 ms, and the max delay is about 20 ms. Minimal delay is achieved when the packets arrive immediately before a MAP is received, and the CM is able to send those packets using the grant contained in the MAP. In this case, the US latency is just the MAP advance time, which is illustrated in Figure 17. With 2ms CIN delay, the MAP advance time is 8ms

The maximum latency occurs when the packets go through the normal request-then-grant process and their request arrives at the CMTS right on the MAP time grid (actually just a bit after the MAP grid), causing the grant time postponed by 2 MAP intervals. The initial version of the simulation used does not take into account collisions, which would result in the loss of request packets. Thus the maximum delays will not include retransmission times.

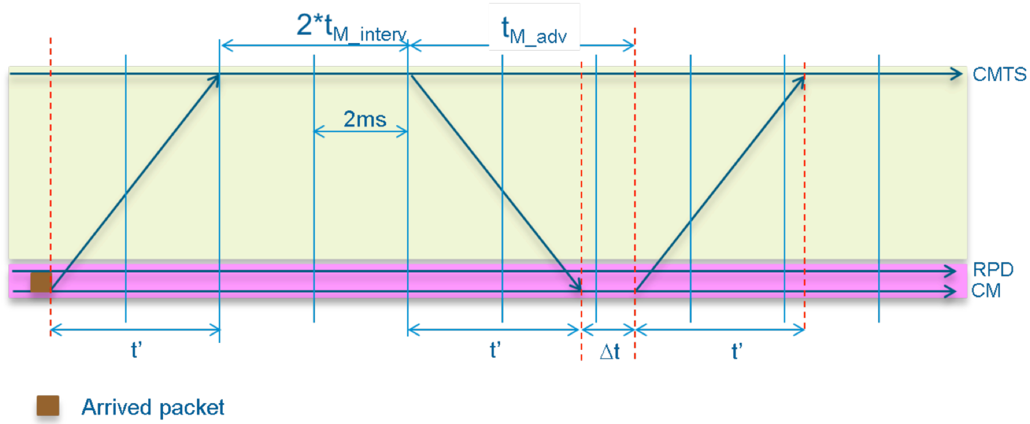

where $t_{M\_interval}$ is the Map interval(2ms), t_{M_adv} is the Map advance time, t' is the DS or US delay (assume symmetric). $\Delta t = t_{M\_adv} - t'$.

Maximal US latency $= t' + 2*t_{M\_interv} + t' + \Delta t + t' = 2*t_{M\_interv} + t_{M\_adv} + 2t'$

Figure 18 – The max US latency in a normal request and grant process

This is illustrated in Figure 18. With the system parameters values configured in the simulation and 2 ms CIN delay, the max US latency is 20 ms.

By averaging the delay over time and over all CMs, one obtains 13.5 ms mean US latency. 13.5 ms is close to the average of the min and max latency ((8+20)/2=14), which means the packets almost have the equal opportunity to be allocated right after it arrives at CM, or have to go through the normal request and grant process with the max latency, or anything in between.

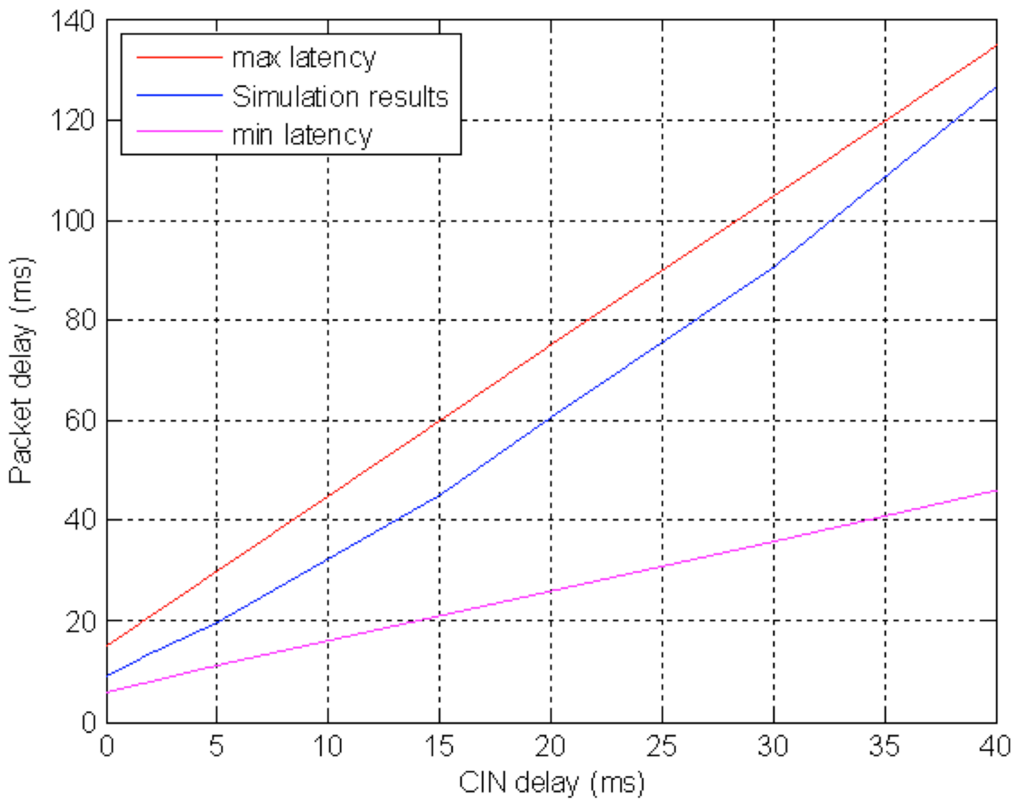

Figure 19 – US Latency vs CIN Delay, REQ-GNT Scheduler, 30 CM, 90% Loading

Simulations were run with other CIN delays. Figure 19 shows the US latency as a function of the CIN delay (CIN_t) for 30 CMs and 90% network loading. The min and max latencies are also plotted on the same figure. The min and max latencies are defined in Figure 17and Figure 18.

Here are the observations:

1. The simulation results falls between the theoretic max and min latencies for all the CIN delays.
2. The simulation results are closer to the max latencies, indicating that, with the baseline request then grant scheduler algorithm, majority of the packets go through the three stage request-grant-send process, and incur the max latency. This is particular obvious with high CIN delay.
3. Some of the packets go through the pipeline of request and grant, incurring less latency than the normal request then grant process. This is obvious because the simulation results are 5~10ms less than the max latency curve.
4. The latency in general follows the max latency curve with a fixed offset

   US latency = fixed offset + Max_delay = Fixed_offset+ 2*t_{M_interv} + t_{M_adv} +2t'

   Where t_{M_interv} is the map interval, and t_{M_adv} is the Map advance time, t' is the DS delay.

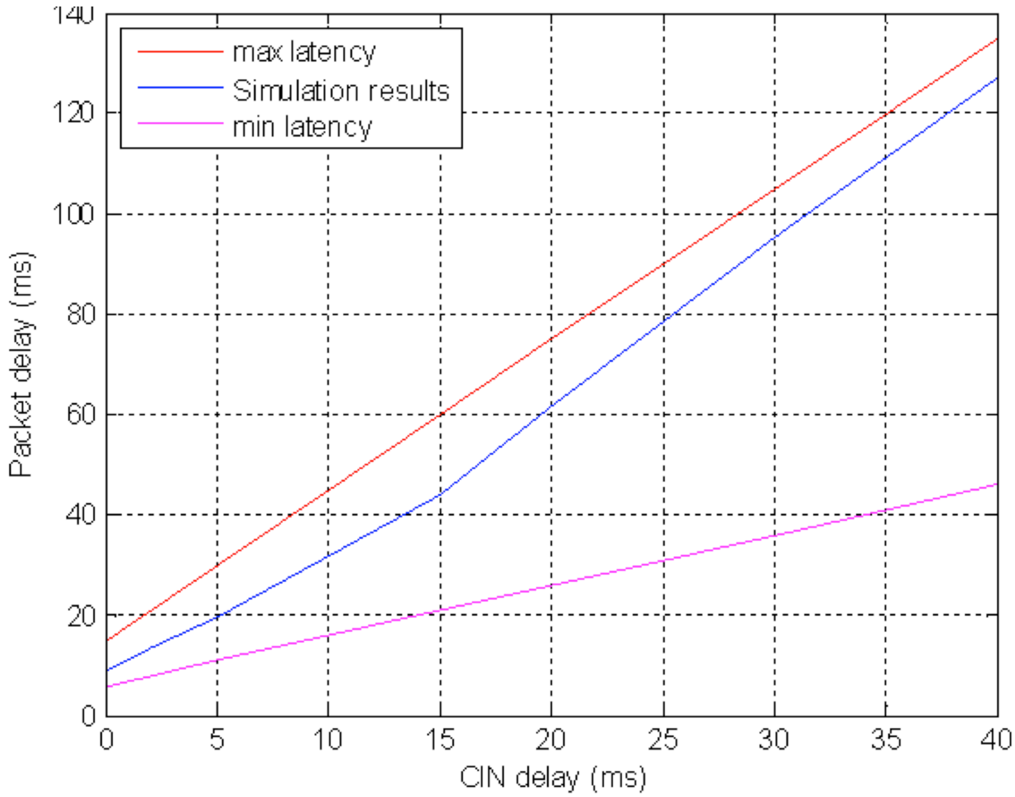

Figure 20 – US Latency vs CIN Delay, REQ-GNT Scheduler, 30 CM, 50% Loading

With the configured system parameters, we have then:

- US latency
  = fixed offset + 4 cm_t + cmts_t
  + ds_phy_t + hfc_t + margin
  + CIN_delay + CIN_delay
  + ds_phy_t + hfc_t/2
  = fixed_offset' + 3*CIN_delay
  ~= 2 + 3*CIN_delay

Three times the CIN_delay is expected, as in the normal request-grant-send three stage process, the packets needs to traverse the CIN three times, incurring 3X CIN delay. The fixed offset counts for other delay (MAP process time, and extra dwelling time in the queue, etc.)

This 3*CIN delay is also observed with different network loading. Figure 20 shows the US latency as the function of the CIN delay (CIN_t) for the case 30 CMs and 50% network loading. It is obvious that all the observations described above for 90% network loading are also valid here. The latency follows the same 3*CIN delay curve.

Algorithm 2: Proactive Grant Algorithm

In this algorithm, the scheduler proactively generates grants to CMs based on received pre-requests or even when no requests are received. Grants that do not result from normal requests are called proactive grants. The proactive granting algorithm needs to follow certain rules to optimize the overall performance (latency, BW utilization, etc.).

The basic idea of the proactive grant is to give extra grants to CMs to increase their

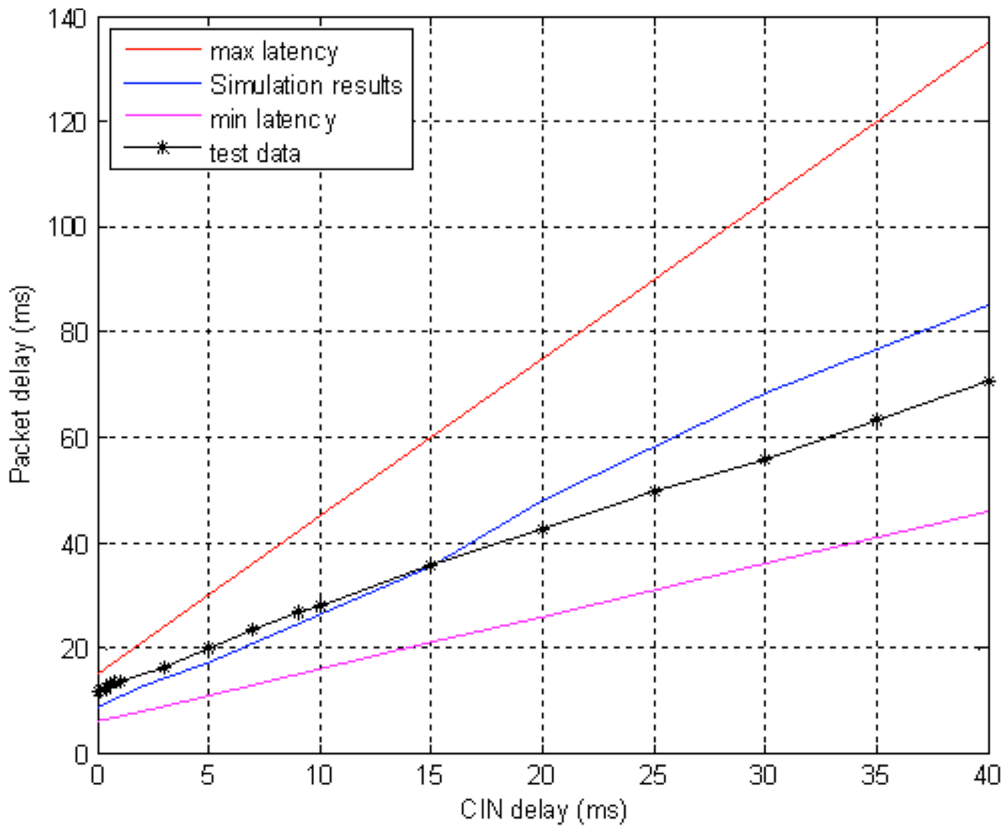

Figure 21 – US Latency vs CIN Delay, Proactive GNT Scheduler, 30 CM, 90% Loading

opportunities to send packets without going through the three stage process, thus the latency is reduced. It is critical to have the right sizes and intervals for the proactive grants so an effective pipeline of request and grant can be established and, at the same time, a high BW utilization is achieved.

Too many proactive grants will definitely help reduce the latency, but may cause some of the proactive grants to be unused or partially used, leading to low BW utilization. So, the sizes and intervals of the proactive grants need be dynamically adjusted with heuristics based on the CM performance metrics and traffic pattern.

Figure 21 shows the simulation results computed with the proactive grant algorithm. The number of CMs included in the simulation is 30, and the network loading is 90%. All other system parameters are the same as before. By comparing with Figure 19 (the request then grant algorithm), it is evident that the proactive grant reduces the US latency significantly.

The mean delay curves lies roughly mid-way between the min and max latency curves, indicating that the packets almost uniformly fall between the two extreme cases of 'send out right away' and 'go through three stage process'.

The simulations show that the pipeline of the request and grant process is established for a considerable number of

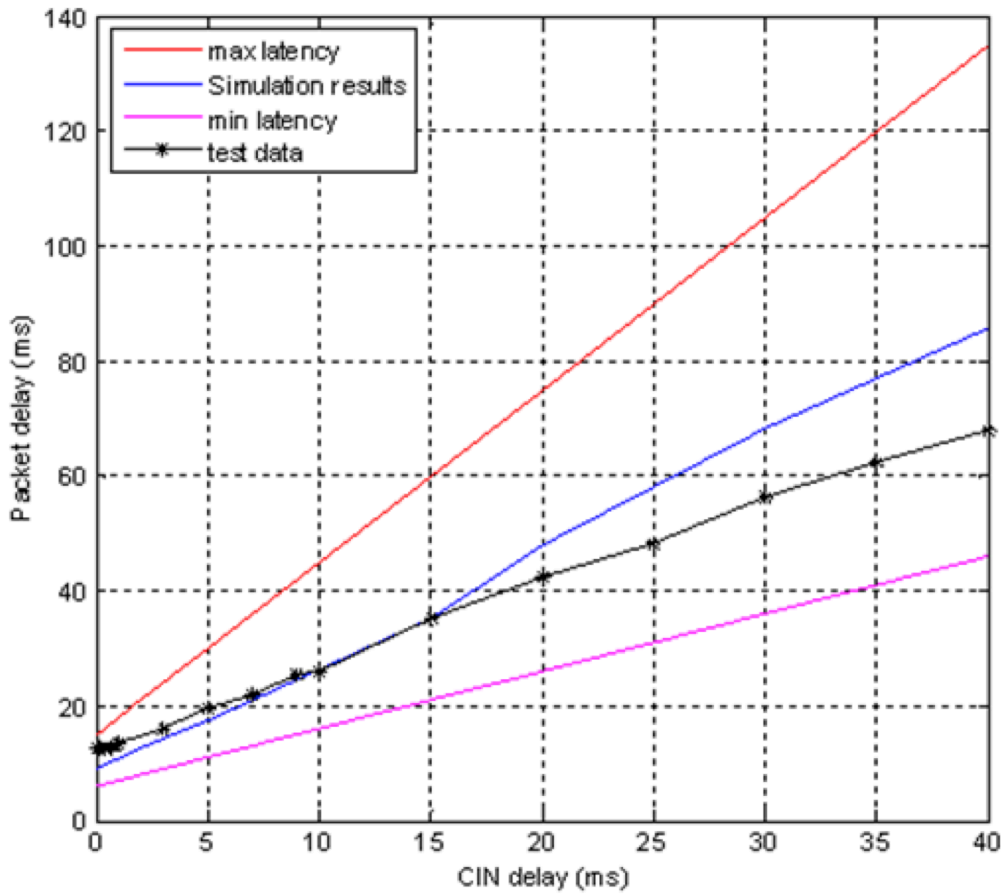

Figure 22 – US Latency vs CIN Delay, Proactive GNT Scheduler, 30 CM, 50% Loading

packets. This results from the extra grants from the proactive grant mechanism.

The test data with the same system configurations are also shown in Figure 21. It is evident that the simulation results agree well with the test data, particularly in the range of 0 ms to 20 ms CIN delay. This range covers the majority of the deployment scenarios, if not all.

Similar performances and agreements are observed with other simulation cases. Figure 22 shows the simulation results and test data for 50% network loading. Both the simulation results and test data show that the proactive grant significantly reduce the US latency.

By examining the simulation results and test data, one can derive the following relationship between the latency and CIN delay.

- US_latency
  = Fixed Delay Offet + 2*CIN_delay
  where the fixed offset is about 9 ms.
- Recall that the US latency follows in general 3*CIN delay curve with the request-then-grant process, the proactive grant scheduler algorithm eliminates 1*CIN delay from the US latency.

In an ideal case where all the requests and grants go through a pipeline, this could result in an US latency of 1*CIN delay (the min delay illustrated in Figure 17). The 1*CIN delay US latency is the theoretical low bound, one could approach to this low bound by designing better scheduler algorithms that fully exploit CM traffic patterns and heuristic performance metrics.

Needless to say, the scheduler algorithm needs to balance out the complexity of the algorithm, BW utilization and actual application QoS requirements.

ReCap

- The max US latency has a slope of three times the CIN delay, and the min US latency has a slope of one times the CIN delay
- The conventional request-then-grant scheduler follows the max US latency curve with a fixed offset.
- The proactive scheduler helps establish a pipeline of requests and grants through proactive granting and can significantly reduces US latency.
- Simulation results show that at least one equivalent CIN delay can be eliminated from the US latency by employing proactive scheduling algorithms.

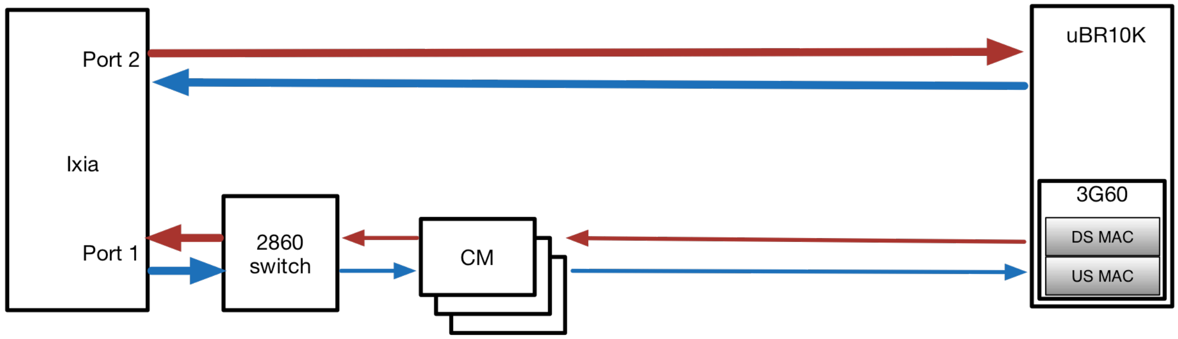

Figure 23 – Test Set Up I-CMTS

PART 3: TEST & MEASUREMENT

Following the simulation experiments we were able to obtain the equipment to generate some experimental data, which could verify (or otherwise) the previous results. To facilitate this a number of tests were run to measure the impact of increasing distance between the cable modem and the CMTS Core.

Test Set Up and Components

The test set up used was as shown in Figure 23 and Figure 24.

CMC

The Cable Media Converter (CMC) used for the testing is a C-DOCSIS (China DOCSIS) type 3 – remote PHY device.

A C-DOCSIS type 3 device such as the CMC contains the DOCSIS PHY with the CMTS Core system providing the higher layer functions including the DOCSIS MAC and all packet classification, QoS and forwarding.

The CMC is connected to the CMTS Core using an IP connection operating over a Gigabit Ethernet link with DEPI / L2TPv3 tunnels to encapsulate the DOCSIS traffic. This architecture is functionally equivalent to a Remote PHY architecture from a DOCSIS 3.0 perspective. It is the first generation of Remote PHY and is a pre-standards version

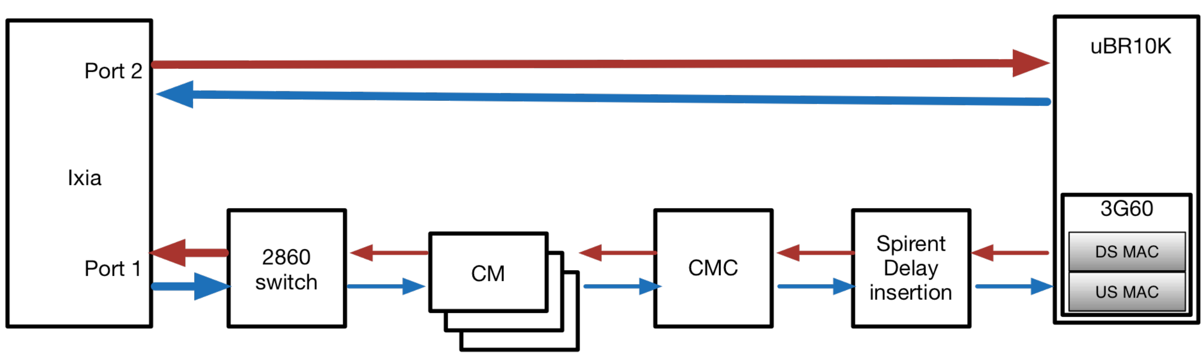

Figure 24 – Test Set Up R-PHY

of the work being done at CableLabs.

This test setup only supports DOCSIS 3.0 and does not support video or DOCSIS 3.1). It does however provide a convenient (and available) platform to investigate the impact of network latency.

CMTS Core

The CMTS core used was a Cisco uBR10012 with an MC3GX60V line card running the CMC controller and DOCSIS MAC software. The physical interface towards the MC was a Gigabit Ethernet. The same CMTS was used for both the Remote PHY and Integrated CMTS tests.

CIN Delay Insertion

A Spirent Aperto test platform was used to inject latency and jitter into the link between the CMTS core and the CMC providing a controlled environment, which was simple to modify during the experiments. Delay can be injected independently to the upstream and downstream using the test equipment but for the data set described in the paper upstream and downstream CIN delays were symmetric

Ixia Test Generator

An Ixia test generator with 10 Gigabit interfaces was used to generate traffic and measure latency. A consistent traffic profile, described in the section Traffic Profile was used throughout the testing.

Cable Modems

For each test there were a total of 30 DOCSIS 3.0 CMs online.

Ethernet Switch

A 2860 Ethernet switch was used to convert between the 10G port on the IXIA and the 1G Ethernet ports on the cable modems. The load on the switch was low enough that it was operating in a non-congested mode. The switch will add some minor amount of latency but is operating at a small fraction of capacity.

HFC

The test bed was a laboratory system with all the components co-located and used a minimal HFC plant with very short cable lengths. Thus transmission delays in the fibers and coaxial cables were essentially zero.

Traffic Profile

The following traffic profile was used for all of the testing (including the baseline):

- Downstream capacity was 16 channels modulated at 256 QAM.
- The downstream traffic was equally divided between each of the 30 DOCSIS 3.0 modems.
- A number of downstream load scenarios were tested varying between 10% and 90% of downstream capacity.
- Upstream capacity was 4 channels of 6.4MHz modulated at 64 QAM.
- The upstream traffic was equally divided between each of the 30 DOCSIS 3.0 modems.
- A number of upstream load scenarios were tested varying between 10% and 90% of upstream capacity.
- The number of frames transmitted in the downstream during each test run varied between 10,000 and 90,000 depending on load.
- The numbers of upstream frames transmitted during each test run varied between 2,100 and 18,900 depending on load.

- All traffic was best effort, 1500 byte packets

Measurements

Referring to Figure 23, upstream latency measurements were made between port 1 and port 2 of the Ixia tester. Downstream latency measurements were made between port 2 and port 1.

For each test the following results were recorded

- Total frames sent, received and lost for each modem
- Minimum, maximum and average latency for each modem.

Note that the measurements do not include the 95% of 99% population points. The maximum number reported is the 100% population point that is at the tail end of the Poisson distribution.

To keep the amount of data contained in this white paper, the reported results are the average seen by each of the 30 CMs. So, the minimum in this paper is actually an average of 30 minimum test results. The average latency reported is actually an average of 30 average latencies. Each of those averages are actually averages of multiple test runs.

So the average results are actually averages of averages.

Latency Added by Test Bed

Any added latency from the test components will be constant as the network delay is increased. As such it is accounted for in the baseline delay measurement.

CIN Delays Tested

The CIN delays used ranged from 0 to 40 ms of delay, representing networks in the range of 0 to 8000 km. The DOCSIS specification allows for delays of up to 0.8 ms between the CMTS and CM so that CIN delays up to 1 ms represent networks roughly comparable to potential I-CMTS deployments.

Non-ideal Testing Conditions

10% versus 90% Loading

Tests were run with a packet loading of 10% to 90%. The most accurate results of system are the results at 90% loading.

This is counter-intuitive because a higher system loading should result in more buffering and more latency.

However, it is the opinion of the authors that the test setup does not exactly model real world traffic. In the test setup, packets are released at an even rate. In the real world, low traffic on the network is when there are fewer CMs generating traffic less often. However, when a single CM generates the traffic, it occurs in a burst.

The net difference is that in the real world, a CM will begin using the DOCSIS piggyback grants even on a lightly loaded system. In this test set up with evenly spaced packets, piggybacking probably did not kick in until the loading was quite high.

	I-CMTS			R-PHY 0 ms			Delta		
Load	**Min**	**Avg**	**Max**	**Min**	**Avg**	**Max**	**Min**	**Avg**	**Max**
10	2.1	2.4	2.7	2.1	2.5	2.8	0.0	0.0	0.0
25	2.0	2.4	12.0	2.1	2.5	2.8	0.1	0.0	-9.2
50	1.8	2.3	12.1	1.8	2.3	2.8	0.0	0.0	-9.3
75	1.7	2.0	2.6	1.8	2.0	2.6	0.1	0.0	0.0
90	1.7	2.0	2.7	1.7	2.0	2.8	0.1	0.0	0.1

Table 1 – I-CMTS and R-PHY Downstream Latency

To illustrate this point, the log file reported the following results:

- In case of 10% usage, there are 2211 BW_REQ, in which, 50 are piggyback request (2%);
- In case of 90% usage, there are total 12205 BW_REQ, in which, 7792 are piggyback request (64%).

1 GE versus 10 GE

Future generations of CMTS Cores and RPDs will use a 10 GE interface. This test setup used a 1 GE interface. Thus, this particular test setup will contain transmission latency that will not exist in a product system.

IEEE 1588 based R-DTI Timing

This test system did not have the benefit of 1588 timing. Thus, extra engineering margin (MAP Advance Time) was used to account for the CIN delay.

With an R-DTI system, one should expect even better performance.

It is worth noting, however, that a R-PHY system could be made to work quite well without R-DTI.

Baseline Results for Latency

To create a baseline, two sets of tests where done.

1. The first was on an Integrated CMTS with 30 CMs.
2. The second was on the same CMTS with an R-PHY Device connected with 0 ms of CIN delay and the same 30 CMs.

The CMTS was the same CMTS in both cases and the software was the same. The line card that drove the R-PHY Device was a modified integrated CMTS line card with the PHY circuitry removed.

In theory, both these systems are similar since the components are the same – they are just packaged differently.

I-CMTS and R-PHY Downstream

The downstream latency for the two reference test setups for various system loadings is shown in Table 1. The differences between the two configurations are also shown.

In comparing the minimum and average downstream latency of an I-CMTS with a R-PHY system with 0 ms CIN delay, there are no appreciable differences.

	I-CMTS			R-PHY 0 ms			Delta		
Load	**Min**	**Avg**	**Max**	**Min**	**Avg**	**Max**	**Min**	**Avg**	**Max**
10	4.8	10.6	43	4.3	13.5	46	-0.5	2.9	3.1
25	4.7	9.5	45	3.4	11.8	44	-1.3	2.3	-0.6
50	4.6	10.4	50	1.7	12.5	48	-2.9	2.1	-1.5
75	4.7	9.8	51	1.7	11.5	50	-3.0	1.6	-0.8
90	4.6	11.3	63	1.7	11.8	57	-2.9	0.5	-5.5

Table 2 – I-CMTS and R-PHY Upstream Latency

The max value is a Poisson tail value. In this test case, it seems that queuing delays on the I-CMTS for two of the loading cases where higher with the I-CMTS than the R-PHY. These are likely test/queuing anomalies and not significant.

I-CMTS and R-PHY Upstream

The upstream latency for the two reference test setups for various system loadings is shown in Table 2. The differences between the two configurations are also shown.

The minimum latency for the I-CMTS upstream is in alignment with the theoretical latency of 4-6 ms.

The minimum latency for the R-PHY system as measured is better than the I-CMTS. This should not actually be the case and is the result of a scheduler enhancement. This is explained in section "Why the Measured Results were Better than Theory"

The average results for both systems show latencies in the 10 to 12 ms range. This is surprising at first, but is a result of buffering in a real world system.

The practical measurement point between the I-CMTS and R-PHY system with 0 ms of CIN delay would be the 90% data point of 0.5 ms. (refer to the Section "10% versus 90% Loading")

The maximum system delays that is a result of contention requesting is surprising high at 50 to 60 ms. It is surprisingly similar in the two systems as well. This can be attributed to the request backoff algorithms for both systems resolving themselves in a similar way.

It is worth noting that these high latencies are natural latencies for DOCSIS and illustrates how well DOCSIS works at these latencies.

CIN Delay and Downstream Latency

Having established the baseline data for the test bed with zero network delay the same tests were run for a range of network delays.

The average, latency for each CM was measured and these values used to calculate the average latency across the CM population as CIN delay was increased.

Average Downstream Latency

CIN delay (ms)	Average DS Delay		
	10% load	50% load	90% load
0	2.46	2.26	1.98
0.2	2.66	2.46	2.18
0.4	2.86	2.66	2.38
0.6	3.06	2.86	2.58
0.8	3.26	3.06	2.78
1	3.46	3.26	2.98
3	5.46	5.26	4.98
5	7.46	7.26	6.98
7	9.46	9.27	8.98
9	11.46	11.26	10.98
10	12.46	12.27	11.98
15	17.46	17.26	16.98
20	22.46	22.27	21.98
25	27.46	27.27	26.99
30	32.46	32.27	31.99
35	37.46	37.27	36.99
40	42.46	42.27	41.99

Table 3 – Downstream Latency with CIN Delay

Table 3 and Figure 25 show the downstream latency measured as the CIN delay was increased for the three load factors.

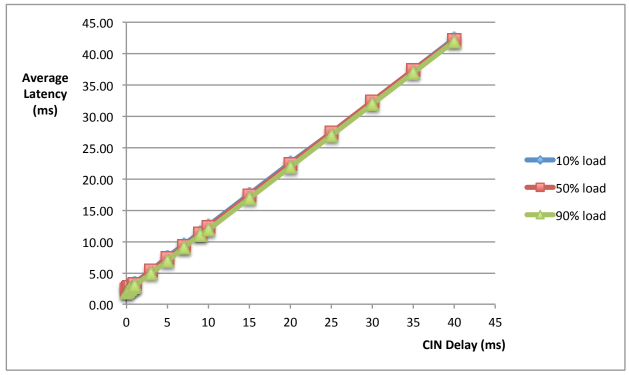

Figure 25 – Average Downstream Latency vs CIN Delay

The measured average latency was seen to be a small, almost constant (~2ms) addition to the imposed CIN delay. In effect it was the baseline delay + the CIN delay.

The downstream performance was seen to be very tolerant of CIN delays remaining linear up to 40ms of imposed CIN delay.

CIN Delay and Upstream Latency

Average Upstream Latency

CIN delay (ms)	Average US Delay		
	10% load	50% load	90% load
0	13.50	12.55	11.82
0.2	13.42	12.55	11.96
0.4	14.42	12.36	12.42
0.6	14.40	12.92	13.24
0.8	15.51	13.18	13.59
1	15.36	13.44	13.48
3	20.93	15.95	16.29
5	26.81	19.47	19.91
7	29.30	21.75	23.38
9	30.48	25.51	26.82
10	30.92	26.09	27.84
15	41.04	35.08	35.66
20	41.20	42.32	42.57
25	54.22	48.12	49.69
30	58.09	56.28	55.69
35	67.09	62.42	63.19
40	69.92	68.00	70.64

Table 4 – Upstream Latency with CIN Delay

Table 4 shows the upstream latency measured as the CIN delay was increased for each load factor. From 0 to 1 ms of CIN delay represents CM to CMTS distances approximately within the DOCSIS specification (0.8 ms is the DOCSIS limit). In this region, the R-PHY system exhibited latencies in the 11 to 15 ms range. For CIN distances beyond the DOCSIS limit the latencies varied from 16 to 70 ms.

Figure 26 shows the average upstream latency plotted against the one way CIN delay for CIN distances in the range of 0 to 200 km. The results show that for each way CIN delays in this range the behavior is relatively linear with a modest increase from the 12 ms baseline to between 13 and 15 ms at the 200 km (1 ms) distance.

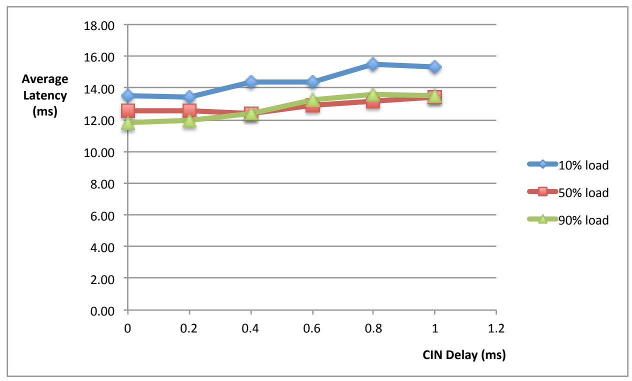

Figure 26 – Average Upstream Latency vs CIN Delay (0 to 200 km)

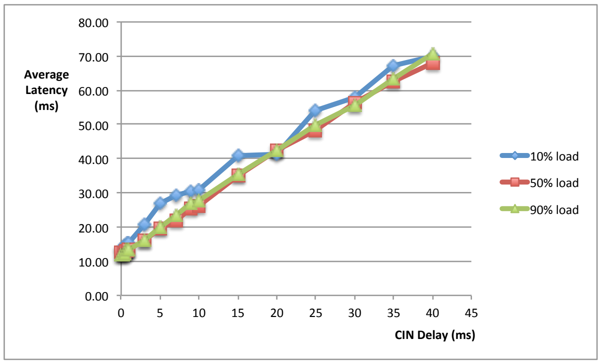

Figure 27 – Average Upstream Latency 0 to 8000 km

Figure 27 shows the latency measured for the entire range of tests. This extended the CIN latency to 40ms representing an 8000 km (each way) network. As seen from the graph the latency increases linearly with the CIN delay. It is interesting to note that the 10% load measurements show higher latencies than the 50% and 90% load at these CIN delays. This is due to the higher level of piggybacking at the higher utilization rate.

Minimum Upstream Latency

The minimum latency was measured at the three load factors and the average across the CM population calculated and shown in Table 5 and Figure 28.

Columns 2, 3 and 4 of Table 5 show the average minimum latency measured for the three load factors. Columns 5, 6 and 7 show the measured latency minus the CIN delay.

CIN delay (ms)	Average Min US delay			Average Min US delay - CIN delay		
	10% load	50% load	90% load	10% load	50% load	90% load
0	4.28	1.73	1.69	4.28	1.73	1.69
0.2	5.29	2.08	1.88	5.09	1.88	1.68
0.4	4.22	2.15	2.08	3.82	1.75	1.68
0.6	6.12	2.30	2.27	5.52	1.70	1.67
0.8	5.06	2.49	2.47	4.26	1.69	1.67
1	5.49	2.69	2.67	4.49	1.69	1.67
3	4.72	4.68	4.67	1.72	1.68	1.67
5	6.74	6.66	6.66	1.74	1.66	1.66
7	8.67	8.66	8.65	1.67	1.66	1.65
9	10.74	10.68	10.64	1.74	1.68	1.64
10	11.71	11.67	11.65	1.71	1.67	1.65
15	16.67	16.66	16.66	1.67	1.66	1.66
20	21.72	21.66	21.66	1.72	1.66	1.66
25	26.73	26.67	26.64	1.73	1.67	1.64
30	31.69	31.67	31.65	1.69	1.67	1.65
35	36.67	36.66	36.66	1.67	1.66	1.66
40	41.73	41.67	41.65	1.73	1.67	1.65

Table 5 – Minimum Upstream Latency vs CIN Delay

For the higher load factors, this is essentially a constant of approximately 1.7ms on top of the CIN delay.

For the 10% loading there is a distinct difference between CIN delays of less than 3 ms and higher delays. Above 3 ms we again see a roughly constant adder of approximately 1.75 ms. Below 3 ms the delta is much higher in the 4 to 5 ms region. Further study is needed to investigate this area.

The measured difference in US latency for a 0 mile plant to a 100 mile plant is 2.47 – 1.69 = 0.78 ms. This matches the 0.80 ms of CIN latency.

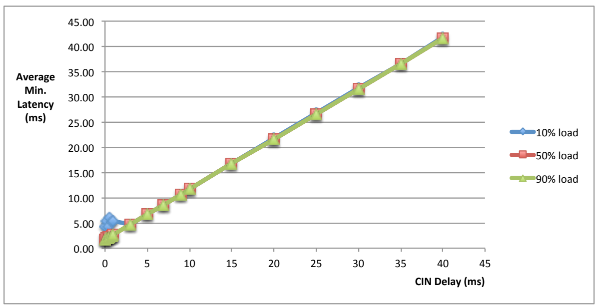

Figure 28 – Average Minimum Upstream latency vs CIN Delay

As can be seen from Figure 28, the minimum latency seen shows very little variation with load.

With the proactive granting strategy described previously the minimum latency occurs when a frame arrives at a CM and is sent "immediately" using an available proactive granted transmission opportunity. Thus the 1.7ms "adder" to the CIN delay represents the processing overhead.

Maximum Upstream Latency

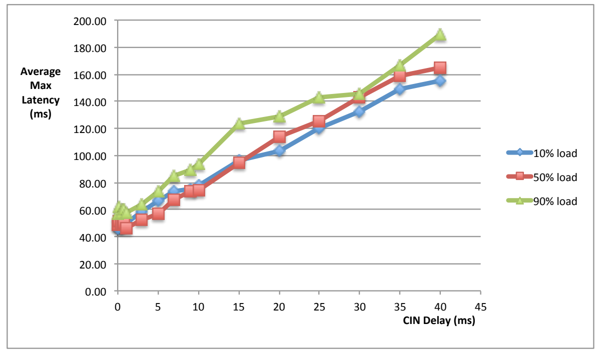

Figure 29 – Average and Maximum US Latency

CIN delay (ms)	Average Max US Delay		
	10% load	50% load	90% load
0	46.30	48.46	57.22
0.2	45.84	51.57	61.67
0.4	47.53	48.91	58.05
0.6	47.43	49.26	59.97
0.8	48.04	49.64	59.92
1	49.53	46.00	57.39
3	58.21	52.83	63.47
5	66.39	56.70	73.46
7	73.11	67.37	85.22
9	75.30	73.44	88.89
10	77.54	73.95	93.46
15	96.02	94.77	123.55
20	103.32	113.77	128.71
25	119.89	125.05	142.94
30	132.63	143.08	145.52
35	148.69	158.69	166.92
40	154.77	165.20	189.63

Table 6 – Average Maximum Upstream Latencies

The maximum latency seen across the CM population is shown in Table 6 and

Figure 29. The maximum latency is high even with a CIN delay of 0. This would be expected from the latency distribution seen in Figure 30 as the maximum latency is measured for those packets at the extreme edge of the Poisson distribution.

These values are included for completeness of analysis but are not statistically significant since they are at the tail end of the distribution.

The impact of CIN delay on maximum latency appears to be more significant than on average or minimum latency and to increase at approximately 3x the CIN delay. This may reflect the scenario in which no proactive grant is available for a packet so that it is subject to the 3x CIN delay (request, grant, data) discussed in the theory section. A more detailed analysis of this is planned as a future work item.

Upstream Latency Distribution for a Single CM

The upstream latency distribution was also measured for a single CM with a CIN delay of 0.6 ms (which is well within the limits of the DOCSIS protocol) and the results shown in Figure 30.

The distribution shows packet latencies varying between 5 and 45 ms even within the DOCSIS operating region. The DOCSIS protocol is clearly very tolerant of varying latencies as the CM was operating perfectly normally when the measurements were taken.

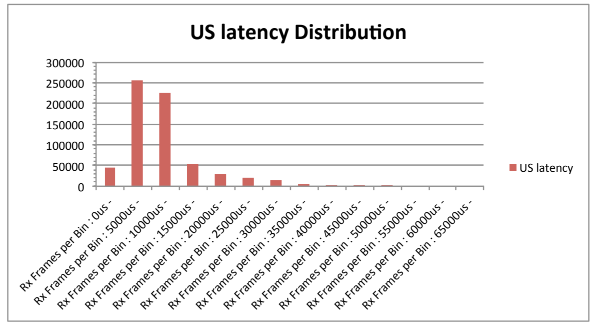

Figure 30 – US Latency Distribution

Thus within the 0 to 200 km operating region the additional 1 to 2 ms latencies due to R-PHY and CIN delay are well within the normal latency distribution seen for DOCSIS traffic.

Why the Measured Results were Better than Theory

Inspection shows that the measured latency was approximately twice the CIN delay. This was in conflict with the theory developed in the previous section on upstream latency, which concluded that a minimum of three times the CIN delay would be incurred as request, grant and data messages must traverse the CIN.

This was the result of pro-active granting. The ACK Time field in the MAP was modified creating a mechanism to generate extra requests and hence grants. This caused the CM to think that requests had been dropped and to resend them. In fact the requests were still in transit to the CMTS so that the effect was to generate additional grants to the CM as shown in Figure 31 – Extra Grants.

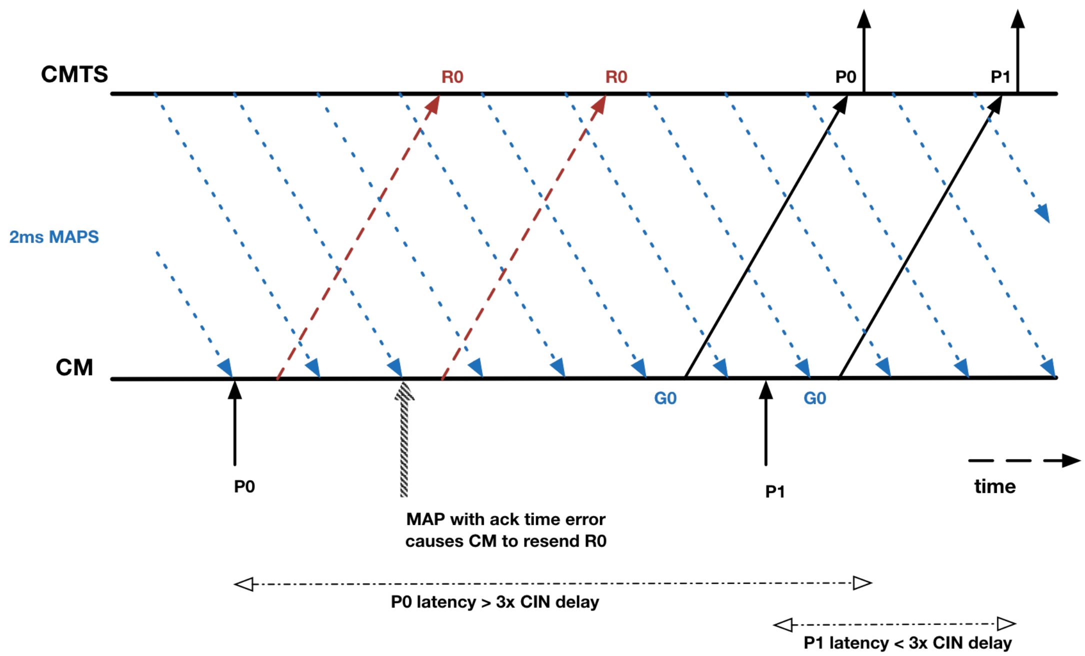

Figure 31 – Extra Grants

The sequence of events is as follows:

- Packet P0 is received at the CM
- Request R0 is generated asking for sufficient upstream bandwidth to send P0.
- A MAP is received with an ack time error. The CM looks at the ack time in the MAP and thinks that R0 has been lost. Thus it resends R0.
- The CMTS receives the first R0 request and responds with a MAP with grant G0
- The CM transmits P0 at the time specified in the map.
- The CMTS receives the second R0 request and responds with a MAP with a second grant G0
- The CM receives packet P1
- The CM transmits P1 at the time specified in the map for G0. Thus P1 has lower latency.

From the above it can be seen that the time to transmit P0 is in fact > 3x the CIN delay as expected (request + grant + data must all traverse CIN sequentially) but due to the additional grant P1 can be transmitted sooner as the request grant cycle is short circuited.

The simple act of generating additional grants to a CM had the effect of significantly reducing the latency impact of the CIN delay for upstream DOCSIS traffic.

The proactive granting applies to all the upstream data measurements in the paper. While the mechanism used was not practical for deployment it was much simpler to implement than changing the scheduler and served to illustrate the effectiveness of even a very crude proactive granting strategy.

Throughput and Frame Loss

This section contains additional test data that was generated and is included for completeness.

Downstream

The downstream tests were run at 10%, 50% and 90% load as described in the section Traffic Profile.

No downstream frame loss at any of the load profiles was seen as the CIN delay was increased suggesting that downstream throughput is not impacted by CIN delay even at the extreme limits of the test. This would be expected from a consideration of CMTS downstream operation as the CMTS core rate shapes the transmission of frames to the RPD to avoid overruns.

Upstream

The upstream tests were also run at 10%, 50% and 90% load as described in the section on Traffic Profile.

Table 7 shows that some frame loss in the upstream was seen, which is to be expected in a DOCSIS system. As seen from Table 7 and from Figure 32, this peaked at 0.04% for the 50% load case and dropped to 0 for the 90% load case.

Frame loss does not appear to increase as a function of CIN delay. The consistently higher frame loss during the 50% load testing appears to be anomalous and may be the result of problems during the test. Unfortunately we did not have time to repeat this test prior to publication but it warrants further investigation.

CIN delay (ms)	Upstream Frame Loss		
	10% load	50% load	90% load
0	0.001%	0.021%	0.000%
0.2	0.000%	0.021%	0.000%
0.4	0.002%	0.024%	0.000%
0.6	0.000%	0.032%	0.000%
0.8	0.001%	0.028%	0.000%
1	0.001%	0.040%	0.000%
3	0.013%	0.036%	0.000%
5	0.000%	0.041%	0.000%
7	0.000%	0.036%	0.000%
9	0.011%	0.036%	0.000%
10	0.004%	0.037%	0.000%
15	0.000%	0.037%	0.000%
20	0.000%	0.035%	0.000%
25	0.003%	0.036%	0.000%
30	0.001%	0.034%	0.000%
35	0.009%	0.037%	0.000%
40	0.009%	0.038%	0.000%

Table 7 – Upstream Frame Loss

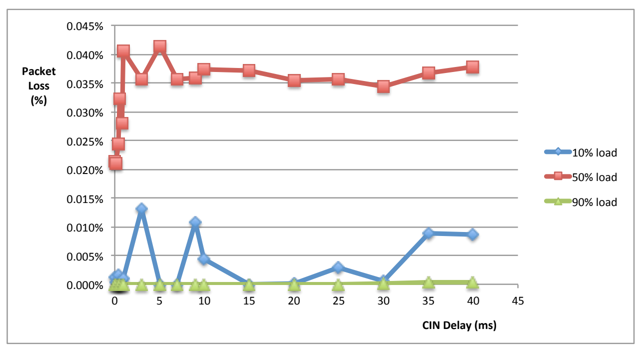

Figure 32 – Upstream Packet Loss vs CIN Delay

Summary of Test Results

The data provides a number of insights relating to a distributed architecture based on R-PHY:

- An R-PHY architecture can support 90% loading on both upstream and downstream links with minimal impact on packet loss for CIN delays of up to 40ms (corresponding to 8000km of fiber).
- Downstream latency increases linearly with CIN delay
- Proactive granting by the CMTS scheduler can reduce the upstream latency for best effort traffic.
    - Upstream minimum latency increases linearly at approximately 1x CIN delay (proactive granting working at its best)
    - Upstream average latency increases linearly at approximately 2x CIN delay (average proactive grant improvement seen during tests)
    - Upstream maximum latency increases linearly at approximately 3x CIN delay (no procative grant available seen during tests)
- Upstream latencies followed a Poisson distribution offset by the minimum delay of the system.
- The latency of the system actually got lower with increased load. This is the opposite of what happens in a normal system. This is probably attributed to DOCSIS "piggy-back" grants kicking in which reduce contention grants. As such, until the test methodology is altered, it is the measured latency at 90% loading that should be the most significant.
- Low traffic density IXIA tests do not accurately represent a true low density network as IXIA testing sends evenly spaced packets rather than burst per CM. This prevents DOCSIS piggybacking from occurring. As such, the test setup should be modified to have a number of traffic bursts rather than evenly spaced packets.
- The R-PHY system for 0 to 100 miles worked very well (note that DOCSIS 3.1 now only requires 50 miles). The added delay for R-PHY versus I-CMTS seem to vary from 0.5 ms (from Table 2) to 2.3 ms (0.5 ms from Table 2 plus 1.7 ms from Table 4) which is close to the predicted one MAP interval time. This is compared to the average US latency of an I-CMTS which is on the order of 10 to 11 ms (Table 2).
- The R-PHY system with a centralized system worked at wirespeed for a 40 ms CIN latency

Recap

- This was a good first pass at testing and can form the basis for further and more in-depth testing.
- The R-PHY test data contained a rudimentary proactive granting scheme. The results without proactive granting should have shown a longer delay but were not available prior to publication. Conversely, a more optimized proactive granting scheme should result in average results that are closer to the shorter delays predicted by theory and simulation than were measured here.

UNIFIED THEORY OF US LATENCY

A rudimentary formula that covers the observed behavior is shown below.

US packet latency = (Q x CIN Delay) + D ms + R ms

Where:

- Q is the impact of the CIN Delay. Q is 3 for reactive granting, and between 1 and 3 for proactive granting, depending upon algorithm efficiency.
- D = DOCSIS processing at CM and CMTS, including REQ backoff.
- R = Additional R-PHY processing.
- Values should all be Min, Avg, or Max

Measured latencies from Table 2 were:

CIN Delay: Min = 0, Max = 0.4 ms (D3.1) or 0.8 ms (D3.0), 40 ms (extended)

D: Min = 5, Avg = 10, Max = 50

R: Min = 0, Avg = 0.5 ms, Max = 2

Based upon these measure values, the CIN latency is insignificant for values below 1 ms and really only becomes significant for value of 2 ms (400 km) or more. It also shows that moving the scheduler to the RPD will have little to no impact.

Note: The measured values are based upon limited experiments and specific test patterns and circumstances and are subject to change.

CONCLUSIONS

One of the biggest features of the Remote PHY architecture is that the software – including the upstream scheduler – is completely centralized. This allows for complete interoperability between a CCAP core and third party RPD vendors. Imagine the challenges if your CMTS Core software came from one vendor and your upstream scheduler came from another vendor. Would it work? Who would you call when it did not work?

The potential technical challenge with centralizing the upstream scheduler has been the REQ-GNT delay. This white paper split the problem into two categories:

1. 0 to 100 miles (DOCSIS 3.0 spec)
2. 100 miles and beyond (beyond spec)

It then compared various R-PHY deployments against an I-CMTS deployment.

The impact of R-PHY on a system 0 to 100 miles is roughly zero to one MAP intervals or about 0 to 2 ms. This was confirmed by theory, simulation, and lab measurements. Meanwhile, the average upstream delay of the I-CMTS and R-PHY at 0 ms was found to be on the order of 10 ms. This is a result of REQ retries, CM and CMTS buffering, processing and reassembly.

The predicted and measured upstream packet latencies in systems less than 100 miles are <u>10x or more</u> than the one-way latency from the CMTS Core to the RPD.

Thus, the latencies in systems less than 100 miles in radius tend to be dominated by MAP duration and packet buffering and reassembly rather than the REQ-GNT delay.

R-PHY can support the same high utilization of DOCSIS channels as an integrated system due to the pipelining inherent in the DOCSIS protocol.

A basic R-PHY system with an unmodified upstream scheduler will provide near-equivalent latencies to an integrated system for distances of up to 160 km (the existing DOCSIS distance limit).

The white paper took the research further to see if the CMTS core could be relocated from the hub site to a more central location.

Upstream packet latencies in systems that are much greater than 100 miles are dominated more by the CMTS Core to the RPD path latency. In fact, the REQ-GNT delay would normally scale by three times the system one-way latency.

However, with the use of a proactive upstream scheduling mechanism, an R-PHY system can provide latencies that are equivalent to an integrated system for distances well beyond the DOCSIS limit.

An R-PHY system was tested at the equivalent of 8000 km (5000 miles) from the CCAP core to the RPD (based upon 40 ms of CIN delay. <u>This is the equivalent of 1.7 times the width of the USA.</u>

The high tolerance of R-PHY to CIN delay will simplify virtual deployments by allowing the use of standard virtual machine environments.

Theory, simulation and measurement align to show that an R-PHY solution can be successfully constructed and deployed.

Further, deployment of a CMTS Core in a regional data center with Remote PHY Devices in fiber nodes with centralized

software is also a very practical architecture
up to significant CIN distances.

APPENDIX I – DELAY IMPACT OF NODE TO CORE DISTANCE

The added delay between the CMTS core and the Remote PHY Device has two components

- The transmission time over the transport links
- The switching time in any intermediate nodes

Transmission Time

Table 8 shows the transmission time for various link lengths based on the speed of light in fiber of $2*10^8$ $ms^{-1.}$

Distance (km)	Accumulated Transmission Time (mS)
100	.5
200	1
300	1.5
400	2
500	2.5
600	3
700	3.5
800	4
900	4.5
1000	5
2000	10
3000	15
4000	20
5000	25
6000	30
7000	35
8000	40

Table 8 – Transmission Time as Function of Distance

Switching Time

The network between the CMTS Core and the Remote PHY Device is expected to be a controlled network that is not oversubscribed. The switching time for current 10G Ethernet switches in the absence of congestion is < 2us for a 1518 byte packet.

Number of Switching Nodes	Accumulated Switching Time (mS)
1	.002
2	.004
3	.006
4	.008
5	.010

Table 9 – Ethernet Switching Time

Table 9 shows the accumulated switching time in a multi-hop network. This shows clearly that switching time can be expected to provide a minimal contribution compared to overall network delay in a well designed and managed network.

Thus the total added latency is approximately 1ms per 200km (ignoring switching time in nodes)

Note: DOCSIS one way latency limit of 0.8ms = 160km

APPENDIX II – CCAP CONFIGURATION

CMTS configuration

```
interface Cable8/0/0
 load-interval 30
 downstream Modular-Cable 8/0/0 rf-
      channel 0-3
 cable mtc-mode
 no cable packet-cache
 cable default-phy-burst 0
 cable map-advance dynamic 1000
      20000
 cable bundle 1
 cable upstream max-ports 4
 cable upstream bonding-group 800
  upstream 0
  upstream 1
  upstream 2
  upstream 3
  attributes A0000000
 cable upstream max-channel-power-
      offset 6
 cable upstream balance-scheduling
 cable upstream 0 connector 0
 cable upstream 0 frequency 15000000
 cable upstream 0 channel-width
      6400000
 cable upstream 0 power-level 2
 cable upstream 0 docsis-mode atdma
 cable upstream 0 minislot-size 1
 cable upstream 0 power-adjust
      continue 6
 cable upstream 0 range-backoff 3 6
 cable upstream 0 modulation-profile
      225
 cable upstream 0 attribute-mask
      20000000
 cable upstream 0 shutdown
 cable upstream 1 connector 0
 cable upstream 1 frequency 21400000
 cable upstream 1 channel-width
      6400000
 cable upstream 1 power-level 2
cable upstream 1 docsis-mode atdma
 cable upstream 1 minislot-size 1
cable upstream 1 power-adjust
      continue 6
 cable upstream 1 range-backoff 3 6
 cable upstream 1 modulation-profile
      225
 cable upstream 1 attribute-mask
      20000000
cable upstream 1 shutdown
cable upstream 2 connector 0
cable upstream 2 frequency 27800000
cable upstream 2 channel-width
     6400000
cable upstream 2 power-level 2
cable upstream 2 docsis-mode
     atdmacable upstream 2
     minislot-size 1
cable upstream 2 power-adjust
     continue 6
cable upstream 2 range-backoff 3 6
cable upstream 2 modulation-profile
     225
cable upstream 2 attribute-mask
     20000000
cable upstream 2 shutdown
cable upstream 3 connector 0
cable upstream 3 frequency 34200000
cable upstream 3 channel-width
     6400000
cable upstream 3 power-level 2
cable upstream 3 docsis-mode atdma
cable upstream 3 minislot-size 1
cable upstream 3 power-adjust
     continue 6
cable upstream 3 range-backoff 3 6
cable upstream 3 modulation-profile
     225
cable upstream 3 attribute-mask
     20000000
cable upstream 3 shutdown
cable sid-cluster-group num-of-
     cluster 2
cable sid-cluster-switching max-
     request 1
```

REFERENCES

Related Papers by the Same Authors

[1] John T. Chapman, "Remote PHY for Converged DOCSIS, Video and OOB", NCTA Spring Technical Forum, June 10, 2014

[2] John T. Chapman, "DOCSIS Remote PHY," SCTE Cable-Tec Expo, October 21, 2013

[3] John T. Chapman, Rakesh Chopra, Laurent Montini, "The DOCSIS Timing Protocol – Generating Precision Timing Services from a DOCSIS System", NCTA Spring Technical Forum, June 14, 2011

[4] John T. Chapman, Jeff Finkelstein, "Maximizing Legacy and DOCSIS 3.0 Bandwidth in the Forward Path using M-CMTS and EdgeQAM Technologies", SCTE Cable-Tec Expo, October 30, 2009 http://www.johntchapman.com/Legacy-and-DOCSIS3-with-M-CMTS-WP-%282009-10-30%29.pdf

[5] John T. Chapman, "Next Generation CMTS – An Architectural Discussion", SCTE Emerging Technologies, January 15, 2008 http://www.johntchapman.com/NG-CMTS-WP-071125.pdf

[6] John T. Chapman, "DOCSIS Timing Interface – Timing for DOCSIS Networks", National Institute of Standards and Technology (NIST) Keynote Address, March 10, 2007 http://www.johntchapman.com/NIST-jchapman-070310.pdf

[7] John T. Chapman, "Modular CMTS Tutorial", CableLabs Summit, December 2006 http://www.johntchapman.com/M-CMTS-Tutorial-Public-070220a.pdf

[8] John T. Chapman, Andrew Page, Carmelo Iaria, Glenn McGilvray, "Transitioning to M-CMTS", NCTA National Show, April 2006

[9] John T. Chapman, "Modular CMTS Architecture", NCTA National Show, April 3, 2005 http://www.johntchapman.com/ncta-2005-modular-cmts-wp.pdf

[10] Gerry White, "A Next Generation Architecture for Business Services Delivery Based on SDN and NFV", SCTE Cable-Tec Expo, 2014

[11] Gerry White, "Transforming Cable Infrastructure into a Cloud Environment", NCTA Spring Technical Forum 2012.

[12] Gerry White, John Ulm, "New Converged Access Architectures for Cable Services", NCTA Spring Technical Forum 2011

Related CableLabs Specifications

[13] "CM-SP-R-PHY: Remote PHY System Specification", Revision I01, CableLabs, 2015

[14] "CM-SP-R-DEPI: Remote Downstream External PHY Interface", Revision I01, CableLabs, 2015

[15] "CM-SP-R-UEPI: Remote Upstream External PHY Interface", Revision I01, CableLabs, 2015

[16] "CM-SP-R-DTI: Remote DOCSIS Timing Interface", Revision I01, CableLabs, 2015

[17] "CM-SP-R-OOB: Remote Out Of Band", Revision I01, CableLabs, 2015

[18] "CM-SP-R-OSSI: Remote PHY Operations Support System Interface", Revision I01, CableLabs, 2015

[19] "CM-SP-GCP: Generic Control Plane", Revision I01, CableLabs, 2015

ACKNOWLEDGEMENTS

The measurements described were part of a series of tests executed by the following Cisco engineers

- Atlanta and RTP teams: especially John Downey, Chris Leghorn, John Skrobko and Abhi Mahendiran.
- China Team: especially Yubin Chen and Yi Zhang

Many other Cisco engineers provided insight into the theoretical and practical subtleties of the DOCSIS protocol, which was invaluable in producing this paper.