The Power of Distributed Access Architectures (DAA)

Benefits of Digital Fiber Along with Remote-PHY

A Technical Paper prepared for SCTE•ISBE by

John J. Downey
Sr. CMTS Technical Leader
Cisco Systems
RTP, NC
(919)931-9453
jdowney@cisco.com

Table of Contents

Title **Page Number**

1. Introduction ... 3
2. Background ... 3
3. Deployment Scenarios ... 4
4. Remote PHY High Level Benefits ... 5
5. Digital Optics Benefits ... 5
6. Power and Space Savings ... 6
7. The Power of Digital and IP ... 7
8. Automation and Operational Support ... 8
9. “Real Life” Testing and Results ... 9
10. Future Gazing ... 11
11. Don’t Forget/Neglect the “Little Things” ... 12
12. Conclusion & Summary ... 13

Abbreviations ... 14

List of Figures

Title **Page Number**

Figure 1 - Typical Analog Deployment ... 4
Figure 2 - Digital Fiber (DAA) Node Deployment ... 4
Figure 3 - DAA Shelf Deployment ... 5
Figure 4 - Typical Analog HFC Deployment ... 7
Figure 5 - Node Splitting in RF Domain ... 8
Figure 6 - Digital Combining/Splitting ... 8
Figure 7 - Lab Testing Diagram ... 9
Figure 8 - Real Network Distance Testing Diagram ... 9
Figure 9 - DAA for MDU Business Case ... 14

1. Introduction

The industry's latest technology displacement comes in the form of new system architectures. Shifting service provider networks from analog optics to pure digital optics creates a distributed access architecture (DAA). The fiber network now known as a converged interconnect network (CIN) allows for many advantages for future features and performance.

These advantages include features like low-latency functionality for gaming and Mobile (5G) backhaul, better performance to optimize DOCSIS 3.1 for more capacity/speed, and software features for operational simplicity and support.

High speed access via cable or telco has been a long battle that benefits the consumer in the end in the form of competition for better services and/or pricing. The cable industry has re-invented itself many times over with DOCSIS now with version 3.1 deployed and D4.0 on the horizon. We can easily offer > 1 Gbps DS and potentially 100 Mbps or higher US speeds.

The latest "trick up their sleeves" is a new architecture to exploit this even further. By displacing the modulation/demodulation functionality (the physical layer - PHY) remotely, cable operators can offer better reliability, higher quality and in essence, higher speeds.

So, the question, "why R-PHY or DAA in General", will be addressed to provide answers to questions the attendees may not have even known to ask. We look at Remote PHY technology and implementation concerns for CMTS platforms along with benefits of DAA. A new phase is occurring now for outside plant upgrades that convert analog optics to digital optics. There are lessons learned and best practices to make this transition easier, less cumbersome with less customer-generated support cases.

2. Background

Current remote phy devices (RPD) come in node module or shelf designs and support the DOCSIS 3.1 spectrum requirements. The downstream (DS) upper bandedge is 1.218 GHz with an upstream (US) upper bandedge of 204 MHz. The DS lower bandedge will be dictated by the diplex filter and could be hardware or software controlled. Typical US/DS splits will be 42/54, 85/105 (this may actually be 102 to satisfy legacy out-of-band (OOB) settop box DS signaling around 104 MHz), and 204/258 MHz.

DOCSIS 4.0 full duplex (FDX) allows the spectrum band between 108 and 684 MHz to have US and DS overlap, but only for non-primary DS OFDM with US OFDMA signals. D4.0 frequency domain division (FDD), previously known as Extended Spectrum DOCSIS (ESD), is associated with a DS bandedge extension to 1.8 GHz. This also increases US to 684 MHz with an 834 MHz DS start. The US options come in 96 MHz blocks from the D3.1 204 MHz bandedge, but recently one block edge was removed from the specification; 588 MHz.

The physical layer devices (RF signals) from the CMTS will be remotely located from the headend (HE) to the field, whether in a node or some type of shelf design. Also, the fiber will be converted from analog to digital causing all HE signals needing to be created at the remote location. These devices must now generate multiple simultaneous DS signals such as:

- CW carriers for:
  - Leakage test signals for Arcom, Effigis, Comsonics, and Trilithic/Viavi typically at 138 MHz and 612 MHz. Four total carriers where most vendors use two. Trilithic may use three.
  - AGC pilots and alignment tones. Ability for placement on the visual carrier frequency and levels ~6 dB > SC-QAMs.
- DOCSIS SC-QAMs (1.1, 2.0, 3.0) & D3.1 OFDM (multiple 192 MHz blocks)
- MPEG and DVB video
  - Out-of-band (OOB) signaling for legacy STB; 55-1 = Motorola; 55-2 = SA/Cisco

Upstream signals must now be demodulated in this remote location and the US RF spectrum remotely viewed in some fashion. RPDs and D3.1 CMs on the market today support:

- 8 SC-QAMs (ATDMA/TDMA)
- 2 OFDMA 96 MHz blocks

3. Deployment Scenarios

Figure 1 below depicts a typical legacy analog deployment for hybrid fiber/coax (HFC) networks. I-CCAP stands for integrated converged cable access platform where all the RF signals are integrated in the CMTS. Converged means the CMTS can provide DOCSIS and MPEG-2 video signals providing much more savings in the reduction or elimination of video edge-qams.

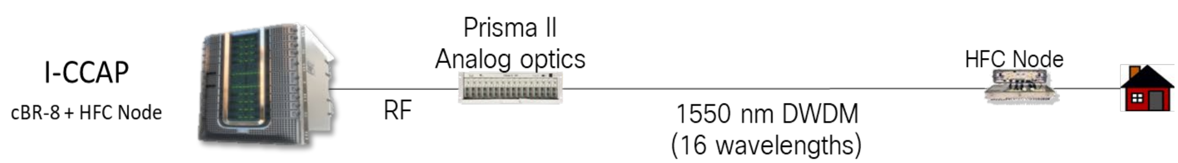

Figure 1 - Typical Analog Deployment

Figure 2 below represents the conversion of a legacy analog HFC plant into a distributed access architecture (DAA) and specifically converting the analog optics into pure digital optics.

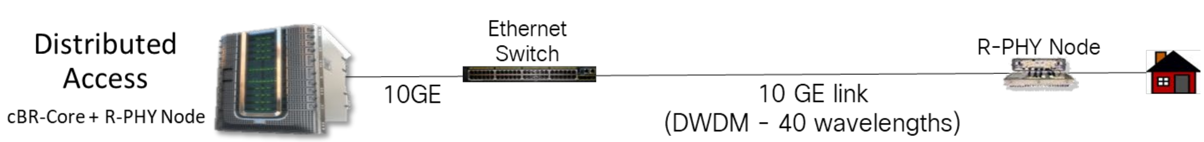

Figure 2 - Digital Fiber (DAA) Node Deployment

Figure 3 below shows another scenario where a DAA can be implemented. This could be a hub site consolidation or even a small HE merged into an existing bigger HE. Digital optics affords us the distance and performance advantages to do this. This scenario shows a hybrid approach where analog HFC is still present, but an entire hub or HE has been eliminated and replaced with signals generated farther away.

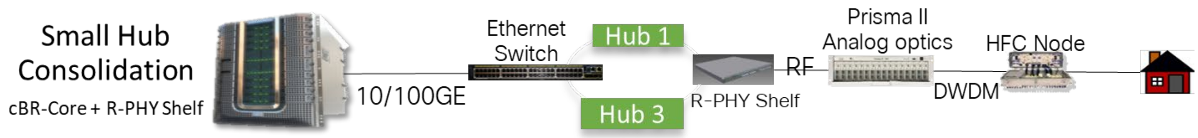

Figure 3 - DAA Shelf Deployment

A fourth idea would be RF port or service group (SG) expansion in the HE. If a current CMTS chassis has 56 DS RF connectors, that in essence is 56 SGs. If a node split or addition is planned, then adding an entire new CMTS is tough to justify, more rack space, power, cost, HVAC, etc. Another option could be to covert one of the RF cards to a digital card and attach via fiber to a few RPD shelves creating more SGs and RF connectors.

4. Remote PHY High Level Benefits

There are many benefits to migrating to digital fiber, but we'll start with some high-level ones first.

R-PHY solves power and real estate problems associated with scale and growth. We can scale out vs scaling up. R-PHY moves spectrum creation from the hub to a node and turns a node into mini-hub. Spectrum represents 100% of services delivered and now this spectrum allocation is targeted closer to the customers. All RF coax, combining and splitting is replaced with fiber, switching and routing.

R-PHY changes HFC from analog fiber to digital fiber and into a 10/100 Gbps Ethernet. The CCAP core can move out of a hub to the headend and turn the hub into an Ethernet switching complex and/or allow hub site consolidation. This enables sharing of commercial and residential plants and provides lower power per SG in the hub. The plant can now become a full- service IP network with simpler fiber design rules and lower plant maintenance costs.

Digital optics also have lower optics costs (10G), more SGs per wavelength and more wavelengths per fiber; 40 vs 16 for analog.

All the core assets such as Video, CMTS, EQAM can be virtualized and allows a cloud-native networking approach and true virtualization.

5. Digital Optics Benefits

There are a plethora of specific digital fiber benefits and the focus of this presentation and paper.

Analog fiber is typically referred to as the "Achilles' Heel" of HFC. Link distance (budget) dictates the optical link performance and usually the entire end-to-end performance. Digital fiber provides much better modulation error ratio (MER), which in turn allows higher D3.1 modulation schemes to be applied with higher speeds.

One of the biggest pitfalls of analog fiber is overdriving the laser also known as clipping. US laser clipping is one fallout of this and well known because of the nature of US noise funneling and ingress in the lower spectrum. DS laser clipping is also a concern when we start adding more channels, more D3.1 OFDM spectrum, emergency alert system (EAS) kicks in, and ingress as well. When we convert to digital fiber, all this concern goes away. There is no laser clipping! Granted, we still need to be aware of analog to digital (A/D) compression, but it's much less systematic than analog laser clipping.

One of the driving forces for capacity is the US and the limited spectrum (typically 42 MHz in North America) we have to allocate for DOCSIS channels. D3.1 supports a 204 MHz upper US bandedge and many systems are looking at this. With the level of US laser clipping we get with 42, 65, or 85 MHz systems, 204 MHz could be a non-starter. Supporting an US of 204 MHz and higher will most likely require digital fiber, in one form or another.

Another great benefit of digital fiber is much longer distances are supported. DOCSIS originally stated .8 msec of delay, which equates to ~100 miles (160 km) of fiber from CMTS to CM. This is possible with EDFA optical amplification, but at the expense of more delay and a performance hit. D3.1 actually dropped the delay to .4 msec or 50 miles end-to-end since almost 99% of fiber nodes are within that realistic distance. Digital fiber transmissions do not degrade in performance like analog (intensity modulated light). It is in essence on/off; 1/0. It's much easier to regenerate a 1! These signals are sent in time vs frequency. More capacity for analog means more spectrum. More capacity for digital means more time/faster links. It's not unfathomable to envision digital links across the country and networks evolving into data/server farms centrally located and feeding remote phy devices across the country thousands of kilometers away. This helps consolidate real estate (hubs).

Another very big advantage/benefit that is often overlooked is cost per wavelength. Digital fiber allows dense wavelength division multiplexing (DWDM) of 40 wavelengths vs only 16 for analog.

Note: This all assumes analog video is retired or some type of overlay is used, which isn't optimum either. Also, US RF testing and spectrum analysis are now required at the RPD, so any test equipment in the HE that requires RF input will be obsolete unless upgraded.

6. Power and Space Savings

A huge benefit of a converged platform along with digital fiber is power and space savings. Converging DOCIS with video delivery through the same chassis can save lots of power and rack space from video e-qam removal. Removing analog video and migrating to digital fiber eliminates the analog optic transport saving power and rack space as well. Coax goes away, RF combining, and splitting is removed, but optical splitting/combining needs to happen. Some rack space saved is needed now for switches, routers, and possibly timing servers.

Typically, two 13 RU CMTS chassis can be collapsed into one chassis. Less CMTSs = less:

- Power
- HVAC
- Rack space

Below is a Cisco command to view power requirements for different components.

- `cmts#show environment power`

```
  ==============================
```


R0	FRU Power	709 W	
2	FRU Power	390 W	<<< Regular RF linecard with RF PIC
9	FRU Power	150 W	<<< RPHY specific linecard
9/1	FRU Power	18 W	<<< DPIC with 8 SFP+ (SR)

The linecard in slot 2 is a regular integrated linecard that provides 8 DS RF connectors for 8 SGs and uses 390 W. The linecard in slot 9 has the US and DS phy modules removed to support remote-phy and lists 150 W. The RF physical interface card (PIC) must be exchanged with a digital PIC which could have up to 8 optical transceivers. That PIC adds 18 W for a total savings of 390-(150+18) = 222 W. A typical fully loaded chassis with 56 SGs drops from ~5 kW to 4 kW.

Note: Even though the PHY is removed, it still has to be accounted for since the service provider pays for outside plant powering. This Phy is moved to the node causing an increase from ~120 W to 160 W.

If Video is incorporated for a converged solution, there will be much more power and rack savings!

7. The Power of Digital and IP

Another benefit of digital fiber is the conundrum we face with splitting/combining with RF. We cannot combine multiple RF upstream legs without interfering and adding noise/ingress. Node splits equal more SGs, which equals more RF connections equaling more CMTSs. This means more RF cabling, rack space, and powering besides cost. Figure 4 below shows a typical analog deployment with the CMTS feeding US and DS independent connections to analog fiber optic transport.

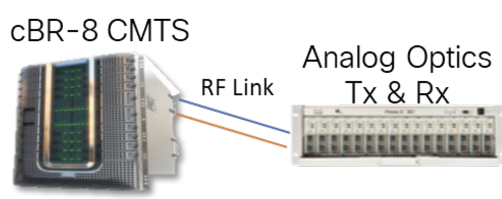

Figure 4 - Typical Analog HFC Deployment

Figure 5 depicts what happens when we do node splits and must add more CMTS chassis to accommodate the extra US connections from the added optical receivers. One could combine the US Rx outputs and share on an existing CMTS US port, but that's like taking 2 steps forward and 1 step back. Why combine noise and ingress?

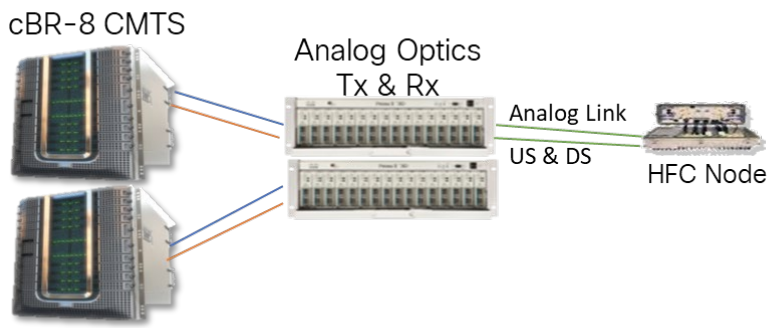

Figure 5 - Node Splitting in RF Domain

Figure 6 below depicts replacing analog optics with digital optics and removing the analog Tx and Rx equipment. Digital optics allows splitting/combining in the time domain with faster digital links. RF does not allow that. If more SGs are needed, they can be deployed in the field without worry of how many RF connections are left on a CMTS in the HE.

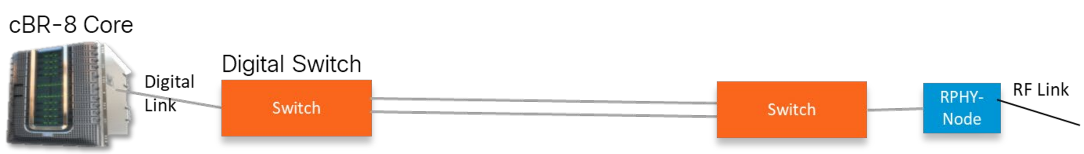

Figure 6 - Digital Combining/Splitting

8. Automation and Operational Support

In a large scale RPD deployment, automation is a must. There will be dozens of steps per RPD with anywhere from 50 to 500 RPDs per CMTS core and 100s of cores in typical network or region. There exist some key steps for RPD deployment automation.

1. Initial RPD discovery
2. RPD to MAC resources mapping
3. Config generation and application to CMTS
4. RPD deployment validation
5. Ongoing health monitoring

DAA is the first, critical step in the path to virtualization. Each step works, has value, and is a good investment.

1. Integrated CCAP
2. R-PHY with physical core and physical orchestration
3. vCCAP stand-alone core and R-PHY shelf or node
4. vCCAP in the data center
5. vCCAP with full orchestration
6. vCCAP core with containers and micro-services

9. "Real Life" Testing and Results

DAA with digital fiber links can traverse much longer distances than analog fiber. To prove that a CM could properly connect and provide HSD along with video services, we put together the design shown in Figure 7 below. This was the control test with very short distance.

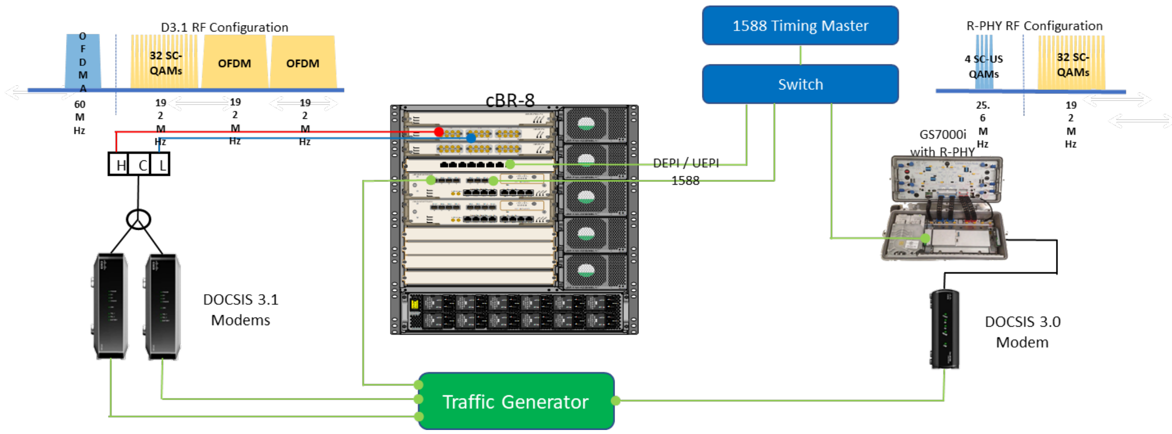

Figure 7 - Lab Testing Diagram

The "real life" distance testing consisted of connecting a CMTS core in RTP, NC with an RPD in Lawrenceville, GA. This is depicted in Figure 8 below. The distance is ~ 350 miles (565 km) at least "as the crow flies".

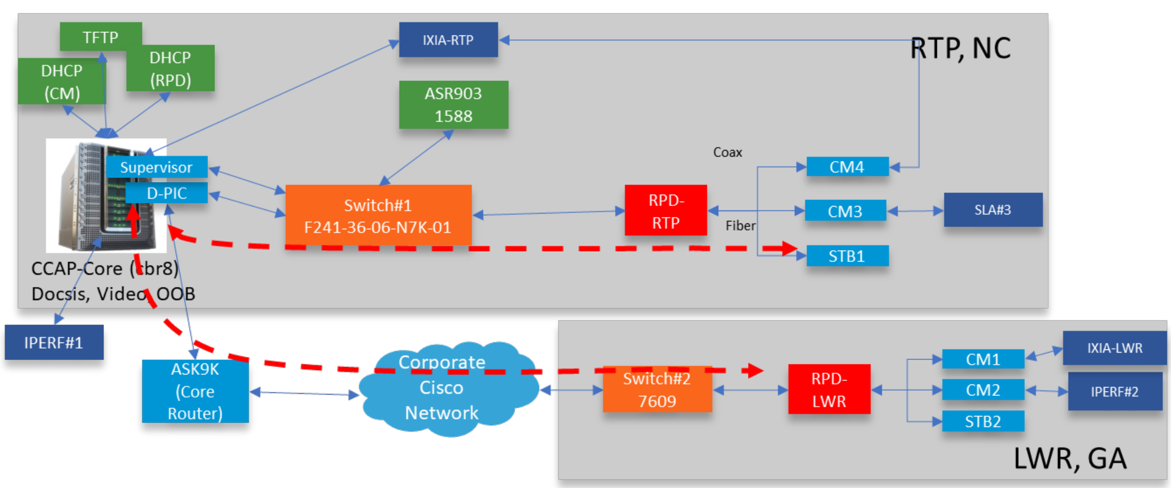

Figure 8 - Real Network Distance Testing Diagram

Running some numbers for velocity of propagation of fiber along with speed of light in a vacuum, we can surmise 2000 km of fiber = ~ 10 ms. This means our fiber distance and delay should only be ~ 3 ms. After some troubleshooting and configuration changes, some observations and major points were made.

- Roundtrip time delay between our CMTS in RTP and the RPD in LWC was 18 ms (avg).

  - To our surprise the corporate network required 13 hops to get to our end location. We quickly determined that router and switch delay could be very high, unstable, and unpredictable!
- Using a Linux box for traffic testing required our MTU packet set to 1434 bytes. The MTU across the IT link was set to 1500.
  - When BPI+ was enabled for security, the MTU on Unix boxes changed to 1428 in order to accommodate the extra bytes used in the DOCSIS Extended Header.
- Precision timing protocol (PTP) is used from proper timing between all the components. Stability and reliability are of utmost importance. Properly designing this part of the equation is critical.
- Certain devices in between the CMTS and RPD may require, or be set for, an MTU of 2000 B. Depending on the RPD vendor or firmware, it may be necessary to turn DEPI fragmentation off:
  - `cable depi fragment off`
- Because our distance was > 100 miles and more specifically, > 500 υs of delay, we had to activate DEPI Latency Measurement (DLM) to properly get the CIN delay added to the Map Advance. If not properly set, the RPD may be fine, but CMs will lose station maintenance and drop offline or not register at all. It is suggested to use DLM with the `measure-only` syntax to track stability of the measurements and if >500 υs, remove the `measure-only` option.
  - ```
    cmts(config)#cable rpd xxx
          (config-rpd)#core-interface Tex/1/x
          (config-rpd-core)#network-delay dlm 1 "measure-only"
    ```
- Because US scheduling is still done at the CMTS for an R-PHY solution, the US request/grant cycle from the DOCSIS protocol can experience delay/latency and affect per-CM US speed. D3.0 with mtc-mode and continuous concatenation and fragmentation (CCF) can get around this. The concern is for D1.x and D2.0 CMs if trying to offer > 5 Mbps on the US over such long distances. You could even have a D3.0 CM register in D2.0 mode (US and/or DS) with the same consequences. One option around this is activating a Cisco proprietary feature known as DOCSIS Predictive Scheduler (DPS).
  - ```
    cmts(config-if)#cab upstream ?
          dps    docsis predictive scheduling on one mac-domain
    ```

  This feature can even help "speed up" US acks that are required for DS OTT ABR video such as Netflix and Hulu TV.
- Another potential cause for concern with Remote-PHY is the CIN traffic. This part of the network (digital fiber) must not be over utilized or oversubscribed. Traffic engineering is critical and proper quality of service given to critical signaling. Also, what if the CIN is a Metro-Ethernet link not owned by the service provider? All the overhead is being billed with no return on investment. Even if the CIN is owned, overhead is a concern for traffic engineering of the CIN. Some of this overhead traffic is created by DS Map traffic, DS video, NDF/NDR and US triggered spectrum captures.

- SC-QAM primary DSs contain MAPs (500/sec) and that will be related to how many are primary and how many USs are in the mac domain. The mac domain could be a 1x2 architecture causing more USs.
  - **Note**: If no data traffic is present during testing, the MAPs are inflated since SC-QAM is MPEG-2 encapsulated leading to 188 B MAPs. Once DOCSIS traffic fills in the 188 B frames, MAPs are closer to 100 B. Also, of note, OFDM primary is not MPEG encapsulated, so 100 B for MAPs is a good number to use. MAPs in DEPI cause UEPI as well, so the CIN MAP overhead will be slightly increased.

Here is an example of potential CIN overhead for a 10Gig link based on a typical deployment.

- 500 Maps/s*100 B/Map*8b/B*8 USs/SG*(24 DSs+1 UEPI)*2 RPDs = 160 Mbps.

It's possible to get this down to 52 Mbps by implementing D3.1 OFDMA US and removing some US SC-QAM chs along with decreasing the DS primary from 24 to 12.

DOCSIS rate is 38.8 Mbps for 256-QAM Annex B since all layer 1 overhead (FEC and Trellis coding) is done at the phy, so not transferred on the CIN. A raw rate of 42.88 is wrong to use for CIN calculations. That 38.8 may include MAPs as well and will be calculated separately below. The CMTS reports 37.5 Mbps, but actual user-rate is closer to 36 Mbps after subtracting MAPs. One could use 37.5 for the CIN calculation, but 38.8 is probably safer. This is close to 50 Mbps for Annex A Euro-DOCSIS.

Video needs 38.8 Mbps and VoD sends null bytes even if no one is watching. So that traffic has to be accounted for always.

The Map overhead will be worse with 1 ms Maps for LLD, more USs and primary DSs, and multiple RPDs per CIN link. Another overlooked aspect is in European markets that are required, or plan, to keep transmitting FM radio over their cable plant. NDF mode 7 is specified for the digitization of the whole FM band and then sent through the CIN link causing 512 Mbps of overhead! This could be a great time to remove this carriage. We also need to be concerned with US Triggered Spectrum Captures (UTSC) as each could be 50 Mbps. Gathering this information from many RPDs, many ports and continuously could have dire consequences.

10. Future Gazing

"Those who do not learn history are doomed to repeat it". So, how did we get here? DAA can be many variants such as Remote-PHY, Remote MAC-PHY, Remote Distributed CMTS, Flexible Mac Architecture (FMA) all using a CMTS core (hardware) or virtual /cloud using servers.

There were concerns with Remote MAC-PHY because of cost, complexity, and power requirements. The cable industry settled on a max wattage of 160 W for the node. More complexity in the field (remote) and in many devices brings a level of unease and unknown.

Remote-PHY gained traction and came first because it was simple. There's something to be said for the "KISS" principle. With complex components in a central location and the PHY in field, it is a simplified design with less chance of issues in the field. Even firmware/software upgrades, bug fixes, and feature additions are easier and more stable.

Knowing that the first iteration of analog to digital migration would probably be a simple node lid changeout with no respacing and 95% of node upgrades would be much less than 50 miles, meant no concern for latency. It's really business as usual (BAU).

With that said, one benefit we listed was hub consolidation, which leads to CIN > 100 miles. Now latency is a concern and it will be necessary to address. This is why we have DLM, DPS and low latency DOCSIS (LLD) as part of the Cablelabs' specification. Part of LLD is the implementation of proactive grant service (PGS) which will provide lower latency for applications that require it like gaming or 5G mobile backhaul. There is even a spec called R-PHY 2.0 that allows for US scheduling in the R-PHY device. So, it is not a full Remote MAC-PHY, but a variant of it.

All this finally brings us to what lies ahead; US scheduling in the RPD explained in the R-PHY 2.0 specification for LLR. Because MAPs are generated in the RPD and the CM US Request and DS Grant are processed in the RPD, there is no CIN contribution to latency. Only the RF portion will contribute and that is typically < 1 mile of coax. This is better for gaming, quicker ping responses and faster US acks for better/smoother DS TCP, which could be OTT ABR video like Netflix and Hulu TV. This also means that DLM would not be needed.

In addition to less latency, LLR also affords us little to no MAP traffic overhead on the CIN. Having MAPs and the US scheduling in the RPD eliminates reliance on PTP sensitivity or PTP issues in general. These PTP issues could be self-induced, 3rd party devices, link redundancy failovers, etc. Granted, we would still need PTP for Mobile Xhaul and other features, but simple DOCSIS CM stability, station maintenance and T3/T4 timers would be solid.

A few other positives being investigated; possibly implement the US Scheduler in existing deployed RPD hardware with no additional power draw and maybe faster RPD reboots ☺

11. Don't Forget/Neglect the "Little Things"

I would be remiss to not mention the little things that are forgotten or neglected. Always start with Layer 1 (Physical Layer) when troubleshooting and validating connectivity. Just because we are using digital transmissions doesn't mean there aren't requirements for optical power Tx and Rx. We must verify correct fiber (single mode vs multimode) and SFPs. Care must be taken to not crisscross the Tx and Rx. Singe mode(yellow jacket) vs multi-mode (orange or blue) is usually easy to identify by color.

SFPs come in different Tx/Rx power. Short reach (SR) are specified for 20 km, while long reach (LR) and ZR are used for 40 & 80 km, respectively. Using ZR SFP optics on a multimode short distance fiber jumper is not going to work properly if at all. Using optical attenuators may be a work-around, but it's still not suggested.

Note: The different values are really based on the Rx sensitivity and the Tx are all the same around 3 dBm. The distance quoted is also based on pure link budget, but in reality, you will lose power from connections, splitters, splices, etc.

Other fiber optic concerns are the actual connections. Does it require LC or SC connectors (node vs optical switch vs RPD, etc.)? Don't forget the patch panels and bulk adapters. Are they angled or flat connections (UPC vs APC)? These are also color coded as black/red for flat and green for angled. Even if a bulk adapter is green, don't assume the connector on the other side of the patch panel is green (angled) as well. Talking about patch panels, how do we justify the cost of many breakout fiber connections with short jumpers and the associated cost of all those SFPs? Here's where copper/10G twinax has been tried or better yet, active optical cable (AOC). Know your options. In regard to breaking out a 100G link to multiple 10G links, QSFPs are available. Newer technology may provide 25G or 40G.

One of the benefits of digital fiber explained earlier was the capability to support DWDM of 40 wavelengths of light vs 16 for analog. This is made possible by optimizing different wavelengths defined in the ITU specification. We must make sure the proper ITU is matched up as a Tx/Rx pair if hard-set, but newer technology may allow adjustable; manual or automatic, at a cost of course.

We may be moving RF out of the HE and getting rid of coax and splitters/combiners, but at a price. We will be replacing RF combining/splitting with routing and switching and fiber connections.

The CIN must be symmetrical for the DS and US (Tx and Rx) which could also be a concern when designing redundancy. An Ethernet ring translates to two different length redundant paths making the RPD DLM and Map Advance adjustment/update critical. The CIN design could also incorporate daisy-chaining to take advantage of a 10G link but only needing maybe 3G at each RPD.

12. Conclusion & Summary

Upgrading or migrating from analog fiber to digital fiber should be an easy sell, but many variables must be taken into consideration in addition to cost. If we look at Remote-PHY in particular, let's separate facts from fiction.

- It's simple and it works
- N+0 **not** required
- Minimum components in RPD yields:
    - Best cost
    - Lowest node and plant power
    - Max SG density for given power budget
    - Best availability
- Same consistent approach for DOCSIS, Video, & OOB
- DAA needed for:
    - US expansion to 204 MHz & beyond
    - Virtualization
    - FDX
- MAC and Scheduler can be scaled as needed since they are central
- Much better visibility of node with RPD vs old transponder technology/nothing
- DOCSIS & video traffic encrypted on fiber

- Centralized software
- Security: CMTS software kept in secure location
- Interoperability
  - Supported by multiple silicon vendors
  - OpenRPD Forum

Remote-PHY can be trialed in targeted areas and doesn't require a complete overhaul of the plant. A great business case could be made for multiple dwelling units (MDU). This could be digital fiber to the MDU with a 1RU RPD shelf feeding the existing HFC plant of the MDU as shown in Figure 9 below. Maybe LLR could be exploited as well for low latency applications.

The shelf can run on 110 Vac or -48 Vdc, so no need for a special quasi-square wave outside plant powering required for a node solution. This could provide 3-6 SGs and provide DRFI DS output levels. The US and DS are independent with support for 204 MHz US and 1.2 GHz DS. Each SG (US and DS pair) could feed a riser with no concern for temperature effects on US attenuation. We could take this a step further and implement Technetix diplexer-less booster amps to "future-proof" the US split. When D4.0 CPE become available, we could upgrade the RPD to higher splits without the need for a plant upgrade assuming taps and passives are specified out to 1.8 GHz.

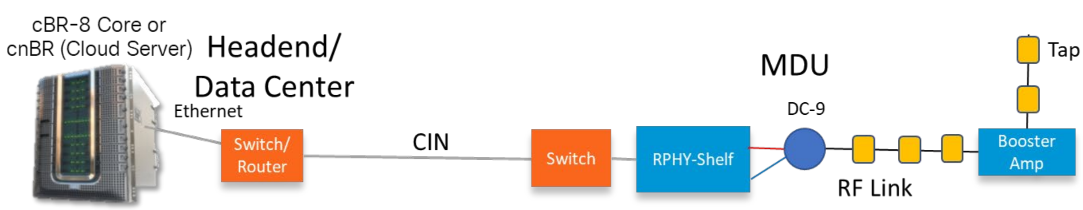

Figure 9 - DAA for MDU Business Case

Converting analog fiber to digital with all its benefits is a foregone conclusion. Whether that is FTTH with EPON/GPON or DAA with R-PHY, MAC-PHY or some variant, depends on other variables such as cost, time to implement, and comfort level. Either way, the benefits of digital fiber are indisputable and the way of the future.

Abbreviations

ABR	adaptive bit rate
APC	angled physical contact
bps	bits per second
CCAP	converged cable access platform
CM	cable modem
CIN	converged interconnect network
CMTS	cable modem termination system
cnBR	cloud-native broadband router
CPE	customer premise equipment
DAA	distributed access architecture
DEPI	DOCSIS external phy interface

DLM	DEPI latency measurement
DOCSIS	data over cable service interface specification
DPS	DOCSIS predictive scheduling
DRFI	DOCSIS radio frequency interface
DS	downstream
DWDM	dense wavelength division multiplexing
FDX	Full duplex DOCSIS
FEC	forward error correction
FM	frequency modulation
FMA	Flexible MAC-PHY
GHz	gigahertz = 1 billion hertz
HE	headend
HFC	hybrid fiber-coax
Hz	hertz
I-CCAP	integrated converged cable access platform
ISBE	International Society of Broadband Experts
LLD	low latency DOCSIS
LLR	low latency remote phy
MAC	media access control
MDU	multiple dwelling unit
MHz	megahertz = 1 million hertz
MPEG	motion pictures expert group
NDF	narrowband digital forward
NDR	narrowband digital return
OOB	out of band
OTT	over-the-top
PGS	proactive grant service
PHY	physical layer
PTP	precision timing protocol
QAM	quadrature amplitude modulation
RF	radio frequency
RPD	remote phy device
R-PHY	remote phy
RU	rack unit = 1.75 inches
Rx	receive
SCTE	Society of Cable Telecommunications Engineers
SC-QAM	single carrier quadrature amplitude modulation
SFP	shared form factor pluggable
SG	service group
Tx	transmit
UGS	unsolicited grant service
US	upstream
VoD	video on demand

