CABLE-TEC EXPO'13
OCTOBER 21-24 / ATLANTA, GA

Society of Cable Telecommunications Engineers

DOCSIS Remote PHY

Modular Headend Architecture (MHAv2)

A Technical Paper prepared for the Society of Cable Telecommunications Engineers
By

John T. Chapman
CTO Cable Access BU & Cisco Fellow
Cisco
170 W Tasman Drive
San Jose, CA 92677
408.526.7651
jchapman@cisco.com

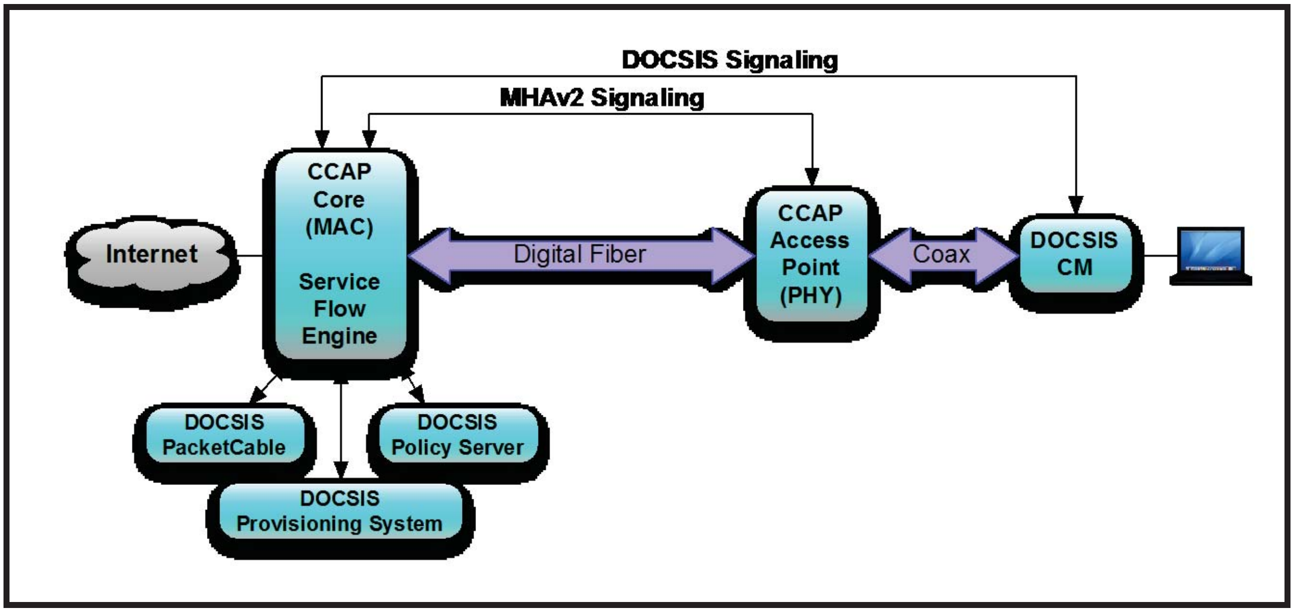

Overview

DOCSIS Remote PHY (RPHY) is known by several names and has a generational history. DOCSIS Remote PHY was originally invented by the author of this paper in 2001. It was brought to the standards process and published in 2004. The name given to the initial suite of protocols was Modular CMTS (M-CMTS). This was to contrast the Integrated CMTS (I-CMTS) that has its PHYs internal to the CMTS.

The initial M-CMTS specification defined two primary specifications. The first was DEPI, (DOCSIS External Downstream Interface) and second was DTI (DOCSIS Timing Interface). There was also the OSSI (Operations Support System Interface), and the ERMI (Edge Resource Management Interface) specifications. ERMI was defined for managing QAMs on the EQAM (Edge QAM) but never got used for DOCSIS.

A few years later, video was more fully defined for the EQAM in some additional specifications; EQAM-PMI (Provisioning and Management Interface) and EQAM-VSI (Video Stream Interface) were added. At this time, the M-CMTS specifications were renamed as MHA, the Modular Headend Architecture.

UEPI (Upstream External PHY Interface) was developed at the same time as DEPI in 2004. At the time, there was no external market application so it was not included in the M-CMTS specifications. In 2007, the ASIC manufacturers adopted UEPI and UEPI became the new MAC-PHY interface in all CMTS upstream designs. There is now some work in the industry to publish UEPI to allow the upstream burst receiver to move to a location external to the CMTS.

The combination of DEPI and UEPI is referred to as DOCSIS Remote PHY. A new timing specification will be developed that will work over longer distances. Another specification is planned to be developed that will define the external QAM channels for MPEG video. This will be Video Remote PHY. The combinations of all these specifications are unofficially referred to as MHAv2.

MHA with DEPI has been adopted by CableLabs. MHAv2 is under evaluation by CableLabs. The MHAv2 with DEPI and UEPI has been adopted in China by SARFT (State Administration of Radio, Film, and Television) for inclusion in the C-DOCSIS (China DOCSIS) specifications. MHAv2 was the recipient of a national innovation award in China.

This white paper will describe the market drivers for RPHY, how RPHY compares to other solutions, and then finally a brief technical description of the primary interfaces, DEPI and UEPI.

The Need for Digital HFC

There are at least two fundamental methods for sending information down and up a fiber – linear and digital. Lets define these as they are referred to today in the context of the Hybrid Fiber Coax (HFC) plant.

> **Linear HFC:** Content is placed on a wavelength by amplitude modulating a laser with a composite RF spectrum. This is the same RF spectrum that exists on the coax segment of the HFC plant, but modulated onto a wavelength.
>
> **Digital HFC:** Content is digitized and placed onto a wavelength. The information is a series of ones and zeros. Digitization could be a baseband sampling or a packet based architecture.

Linear HFC is widely deployed today in the Cable Market. Digital HFC is moderately deployed in the upstream path using a format generally referred to as Baseband Digital Return (BDR). Other newer packet digital formats will be discussed in the following sections.

The advantages of Linear HFC are:

1. It does not “touch the bits”. A linear HFC access network is a simple pipe which does not reformat data and is modulation agnostic.
2. It works. It supports the full service requirement of the cable operator, including native analog video.
3. It is inexpensive
4. When first introduced (1990’s), linear optics provided a higher throughput and a lower cost point than digital optics of the day.

The disadvantages of Linear HFC are:

1. It is distance limited.
2. The linear optics system (node plus amplifiers) requires bi-annual calibration.
3. Its usage is generally restricted to the CATV spectrum and associated applications.
4. It introduces noise that limits achievable SNR.

Linear optics, because it demands a high signal to noise, requires a much higher range of optical powers throughout a link. By contrast, digital optics requires a much lower optical power range for the same capacity transmission. Higher optical powers are, in general, not desired because they make fiber a non-linear medium, and non-linear mediums are always less desirable than linear transmission mediums. This is expressed for instance in the simplicity of digital transmitters, and the complicated electrical drivers of analog transmitters. It is also seen in the relative ease with which digital transmission can do multi-wavelength (80+ lambdas) whereas analog transmission is challenged in doing more than just 16 lambdas.

Despite the great track record of linear HFC, the industry is starting to see the deployment of digital fiber. This has already occurred in North America in the upstream direction (BDR), and in commercial (EPON, GPON), but not in the HFC downstream direction. Internationally, the story is different as there are early cases of Cable Operators turning to digital fiber for the downstream.

The advantages of digital fiber are:

1. In some countries, such as China, Korea, and India, the government is making a regional fiber network available to Cable companies that is digital.

2. Digital fiber can travel longer distances without regeneration.

3. Removing modulation formatting (edge QAM functionality) and demodulation hardware (DOCSIS and QPSK return path tuners) greatly reduces energy and heat loading in Service Provider Critical Spaces such as Head-ends and Hubs.

4. Digital fiber supports more wavelengths per fiber.

5. The CAPEX (capital costs) investment for digital fiber is or will soon be cheaper than linear fiber because of the volume of the 10 and 100 Gbps Ethernet market.

6. The OPEX (operating costs) for digital fiber are lower than linear fiber since calibration of a node is not required.

7. Good scaling properties for deep fiber (N+0, N+1)

8. The optical noise contribution to SNR is eliminated. As a result, a remote QAM modulator that is placed after the fiber link can run at higher orders of modulation when compared to a centralized QAM modulator. Depending upon the HFC plant architecture, the difference can be as high as two orders of modulation. For example, 4096K QAM could be deployed instead of 1024 QAM. (Note: This is a DOCSIS 3.1 example. DOCSIS 3.0 runs at its maximum modulation.)

The disadvantages of digital fiber are:

1. An access network which “touches the bits, unlike the “passive pipe” of current linear HFC networks.

2. Upstream BDR solutions to date are proprietary and do not interoperate

3. Centralized QAM-based services are not supported directly by digital fiber. They have to be converted to digital first and then converted back.

4. The need for a more sophisticated optical node that can do the digital to linear conversion.

5. Native analog video is not supported.

The white paper is focused on one specific problem, and that is the adaptation of DOCSIS to a digital HFC infrastructure. If the HFC plant converts to digital, specifically in the downstream direction, and there is not a well-defined and cost-effective method for adapting DOCSIS to that digital fiber, then DOCSIS technology may get displaced by other digital fiber technologies such as Ethernet or EPON.

The premise of this white paper is that the adaptation of DOCSIS to digital HFC is relatively straight forward and can be accomplished with DOCSIS Remote PHY. The CMTS (CCAP) does not have to be redesigned or discarded. Further, this Remote PHY technology can be applied to all HFC services including MPEG-TS video.

Why Remote PHY?

There are many ways that DOCSIS based services could be sent over digital fiber. These are summarized in Table 1.

Table 1 - Techniques for Digital Fiber

Name	Description	Comments
BDR/BDF	• Remote DAC and ADC. • Digitized bits are framed and sent	• Remote Lower PHY • Remote L0.5
Remote PHY	• PHY chip is remote • Decentralized Software model • Options for remote framer and remote scheduler if needed.	• Remote HW, centralized SW. • Remote L1 to L1.5
Remote MAC	• MAC and PHY chips are remote • Split Software Model. Typically L2 software is remote and L3 software is central.	• Examples are CCAP PASI (defunct) and bridging CMTS. • Remote L2 to L2.5
Remote CMTS	• The entire CMTS HW and SW are remote. • No centralized HW or SW	• Remote L3+

BDR/BDF

The current digital technique in the upstream is to only remote the analog to digital converter (ADC). This technique is known as Baseband Digital Reverse (BDR) and is in wide deployment today. The corresponding technique in the downstream would be to only remote the digital to analog converter (DAC). This would is known as BDF (Baseband Digital Forward). BDF is not utilized in today's HFC networks . Today's HFC plant is linear optics in the downstream and linear or digital optics in the upstream.

BDR/BDF is unique from the next three techniques in that it delivers a digitized spectrum while the other techniques deliver encapsulated data packets.

Some of the advantages of BDR/BDF are:

1. It is very simple.

2. Moving demodulation hardware (DOCSIS and QPSK return path tuners as well as BDR analog signal regeneration at end-of-link) provides a modest but

noticeable energy and heat loading reduction in Service Provider Critical Spaces such as Head-ends and Hubs

3. It is already understood and deployed in the upstream.

4. It will support any current or future service since it does not interpret the service.

Some of the disadvantages of BDR/BDF are:

1. Proprietary and non-interoperable. This could be addressed by standardization, although that would only apply to new products.

2. It is link-optimized rather than network-optimized. The format is TDM based rather than packet based. BDR today does not go over an IP network.

3. Data capacity of the link is inefficient by 3x to 8x when compared to encapsulated data packets techniques such as remote PHY.

4. BDR always runs at continuous peak rate data rather than bursting as needed. This is another 2x to 4x factor.

5. Full spectrum downstream BDF is not a product yet. It will require inexpensive 25 Gbps to 40 Gbps optics.

A simplified example that shows a data capacity comparison is shown in Table 2. The digitization for BDR/BDF assumes 2.5x oversampling with a 12 bit linear codec.

Table 2 - Digital Fiber Throughput

BDR/BDF	Encapsulated Data
Downstream: 1024-QAM, 200 MHz, ~10% Overhead	
• 200 MHz * 2.5 samples/Hz * 12 bits/sample = 6 Gbps	• 200 MHz * 10 bits/Hz * (90%) = 1.8 Gbps
• Ratio = 6 / 1.8 = 3.3	
Upstream: 64-QAM, 6.4 MHz. ~25% Overhead	
• 6.4 MHz * 2.5 samples/Hz * 12 bits/sample = 192 Mbps	• 6.4 MHz * 6 bits/Hz * (75%) = 29 Mbps
• Ratio = 192 / 29 = 6.6	
• Peak rate = average rate	• Peak rate >> average rate

The BDR/BDF performance relative to encapsulated data improves as the modulation increases. Also, a companding codec with 10 bits per sample instead of 12 bits per sample would make an additional improvement.

BDR today is a link technology. It does not use packets to get the information along the link. It uses a TDM structure and has a dependency on clocking that is generated at one end of the link and recovered at the other end of the link. This technology as-is cannot traverse an IP network. It would take a redesign of the protocol and the clocking in order to traverse an Ethernet or IP network.

Further, the BDR/BDF technology is always sending at full data capacity even if there is no actual data payload within the spectrum. For example, if it takes 2.5 Gbps to send an 85 MHz upstream spectrum (typical case), then that 2.5 Gbps data capacity is always in use, even if there is no actual traffic in the spectrum, or if the DOCSIS upstreams only occupy a fraction of the spectrum. The encapsulated techniques only send data when needed.

This is important when network concentration is used. For example, for a BDR/BDF system, if a downstream feeder network was 100 Gbps, and a downstream PHY needed 5 Gbps, then that 100 Gbps network could support 20 PHYs maximum. By contrast, traffic engineering can be used for an encapsulated data system. Traffic engineering can assume that not all the links will be active at the same time and decide that a 2:1 oversubscription will provide the desired service level. Thus, an encapsulated data system like DOCSIS RPHY could support 40 PHYs off of the same feeder network.

DOCSIS Remote PHY

The term Remote PHY refers loosely to locating the DOCSIS PHY chip remotely while keeping the rest of the hardware and software centrally. DOCSIS Remote PHY represents the least amount of hardware and software that can be exported from a CMTS that will yield a system that supports digital HFC.

Remote PHY has a distinct goal of minimization of complexity in the remote chassis or node. This should help keep the cost, power, and size hit of the remote PHY circuitry to a minimum, and keep the reliability as high as possible.

Remote PHY also has a second equally important goal of being as transparent as possible to a centralized CMTS. It should be possible to take a CMTS and have both integrated and remote PHYs on the same CMTS running the same software with the same feature sets.

The protocols that are deployed for DOCSIS Remote PHY also do not interfere with the DOCSIS protocols. This is a key fact. This allows a high level of transparency for software features between I-CMTS and M-CMTS systems. It also allows higher feature velocity as features do not have to be developed separately for each system.

Remote PHY uses an IP pseudowire. This allows the network between the CMTS and the Remote PHY to be any layer 2 or layer 3 network. This also allows the Remote PHY device, such as an optical node, to be used as an IP edge device.

Design calculations show that Remote PHY will work fine with a centralized framer and a centralized upstream scheduler. However, if the distance is to be significantly increased or there are changes with newer DOCSIS protocol, such as DOCSIS 3.1, extension to the protocol could be added that would support these options.

Some of the advantages of Remote PHY are:

1. Preserves the centralized software structure of the CMTS.
2. Least amount of hardware in the remote node
3. Relatively simple.
4. Supports an arbitrary IP network.
5. Can be extended to support a remote US scheduler and/or remote framer if needed.

Some of the disadvantages of Remote PHY are:

1. All services must be identified and supported unless there is an HFC overlay network.

Remote MAC

The term MAC can mean different things to different protocols. In Ethernet, a MAC is basically the layer 2 framer. In DOCSIS, the MAC includes quite a bit more. The DOCSIS MAC has both a complex set of hardware and a very rich messaging set. That messaging set crosses the layer 2 and layer 3 boundary. Although all the messages are formatted in layer 2 frames, Table 3 shows that 40% of the messages are IP aware.

As can be seen from Table 3, the layer 2 and layer 3 aspects of DOCSIS are quite intertwined. This is a historical fact of DOCSIS. DOCSIS was designed as an integrated system and to get the job done. That required layer 3 awareness, even though the MAC management messages were formatted at layer 2. It should be obvious that in order to split layer 2 and layer 3 functionality, the very definition of DOCSIS is changed. That involves changing existing software, which is risky. It also means that for every future change to DOCSIS, those changes would have to be applied separately to I-CMTS software and Remote MAC software code bases.

So the implication of a true DOCSIS Remote MAC system is a system that splits the DOCSIS protocol between the central and remote locations. That split can vary, depending upon the implementation. For the sake of this paper, a Remote MAC system would be any system that did downstream packet classification and rate shaping centrally (a classic Layer 3 function) and had a good portion of the DOCSIS MAC software running remotely.

An example of Remote MAC systems would be the CCAP PASI system. This program was ultimately cancelled before being released due in part to its complexity.

Table 3 - Analysis of DOCSIS MAC messages

DOCSIS MAC Management Messages	L2	L3
SYNC, UCD(3), MAP	5	
RNG-(REQ, RSP, ACK), B-INIT-RNG-REQ	3	
REG-(REQ, RSP, ACK), REG-(REQ, RSP)-MP		3
UCC-(REQ, RSP); DCC-(REQ-RSP)	6	
BPKM-(REQ, RSP)		2
(DSA, DSC, DSD) – (ACK, RSP, ACK)		9
DCI-(REQ, RSP)	2	
UP-DIS, TST-REQ	2	
DCD, MDD		2
DBC-(REQ, RSP, ACK)	2	
DPV-(REQ, RSP)	2	
CM-STATUS, CM-CTRL-(REQ, RSP)	3	
Total	25 (60%)	16 (40%)

It is worth noting that the DOCSIS Remote MAC hardware design may be similar to a DOCSIS Remote PHY with a remote framer, but the software design is quite different.

Some of the advantages of a Remote MAC Design are:

1. Performance offload from the central entity
2. The design of the centralized entity may be simpler.
3. Supports an IP network.

Some of the disadvantages of a Remote MAC design are:

1. It splits the operating software into separate pieces.
2. The design of the Remote MAC entity is more complex than the Remote PHY entity.

Remote CMTS

The Remote CMTS concept is simply to put the entire CMTS into a remote node.

If an entire CMTS is pushed to a remote location, like a fiber node, the software model stays intact. This is a big advantage if re-using an existing CMTS software design. However, there are suddenly 50x to 100x times more CMTSs to manage. This is an operational challenge. Some kind of aggregation system, such as a network management entity, would have to be added to the mix to keep it simple. This negates the CLI manageability.

There is another subtlety. When this device is deployed, the PHY ASICs may have more bandwidth capability than is needed by the network. For example, with DOCSIS 3.1, the PHY chip could have five OFDM channels in the downstream. This is almost a 10 Gbps loading. However, the deployment scenario might only require 1 Gbps. With Remote PHY, 1 Gbps of packet switching equipment can be provisioned at the hub site. With Remote CMTS, however, all 10 Gbps needs to be bought and deployed up front, or it will be necessary to do an upgrade to the node at a later date. This is both a higher upfront CAPEX cost and a longer term OPEX cost.

Then there is a security concern. The CMTS manages link security in terms of BPI (Baseline Privacy Interface) as well as the security around the transfer of the configuration file. It also has a CLI input. If the remote node was broken into, the system could be hacked much easier than if the CMTS was located in a secure hub site.

Some of the advantages of a Remote CMTS design are:

1. Simple.
2. Does not split the software design

Some of the disadvantages of a Remote CMTS design are:

1. Of the three designs, this one may have the highest remote node cost as it has the entire hardware and software design in the node.
2. The Remote CMTS infrastructure must support 100% of the PHY bandwidth from the beginning.
3. The Remote CMTS may be located in an insecure location that would invalidate the security assumption of the DOCSIS design.

Recommendation

The recommendation of this white paper is the DOCSIS Remote PHY design.

DOCSIS Remote PHY allows the best feature compatibility with existing CMTS. This is important to Cable Operators as they will not cut over from a Linear HFC plant to a Digital HFC plant overnight. It will be a transition that takes years. With Remote PHY, it is the easiest for the same CMTS with the same software load to support both styles of HFC network.

Just because the HFC network is changing, it should not be necessary to replace or radically change the CMTS. Rather, the CMTS should evolve to adapt to the new network. Remote PHY permits that minimal impact evolution.

The rest of the paper will provide a high level overview of how DOCSIS Remote PHY works. These techniques can be extended to Video MPEG-TS Remote PHY as well.

Pseudowires

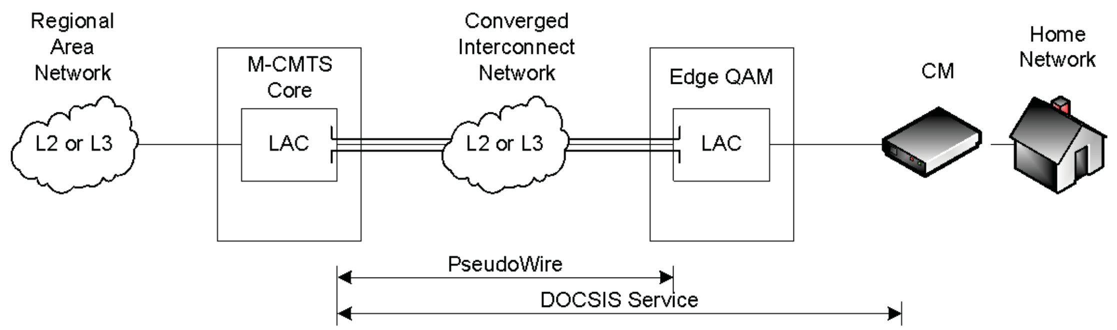

Figure 1 - MHA Pseudowires

To understand DOCSIS Remote PHY, it is first necessary to understand the concept of a pseudowire (PW). A pseudowire was originally defined in "RFC 3955 Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture". The definition has evolved to:

> *A Pseudowire is an emulation of a point-to-point connection over a packet-switching network (PSN).*

A Pseudowire allows an IP network to carry a service without that service having to know the details of the IP network. One transport is encapsulated in another transport.

Pseudowires and tunnels are similar and related concepts. Tunneling something across an IP network using earlier protocols like GRE (generic router encapsulation) came first. As these tunnels became more sophisticated and developed control planes, the concepts of pseudowires evolved. The distinction is that a tunnel may contain one or more pseudowires.

Figure 1 illustrates the use of a pseudowire to transport DOCSIS frames over an interconnecting network.

To build a Pseudowire, you need a service for the pseudowire to connect and a underlying protocol. Examples of underlying protocols are:

1. IP
2. MPLS (Multiprotocol Label Switching)
3. L2TPv3 (Layer Two Tunneling Protocol Version 3)

Simple tunnels that do not require session setup can just use IP. One example is IPv6 over IPv4.

If the network being traversed is an MPLS network, then the natural tunnel/pseudowire technology to use is MPLS for the tunnel and MPLS-TE (MPLS Traffic Engineering) for the session setup.

If the network being traversed is an IP network, and a session setup is required, a convenient choice is L2TPv3. This was the choice for M-CMTS and the DEPI protocol. L2TPv3 offers the following advantages.

1. The underlying network is IP
2. There is a reliable session setup protocol
3. There is support for multiple pseudowires per session
4. The L2TPv3 specification is publicly available for implementation.

DEPI

Intro

DEPI is a pseudowire that connects DOCSIS downstream MAC frames from a CMTS-Core to a DOCSIS downstream PHY located in a separate chassis such as an EQAM. This is shown in Figure 2.

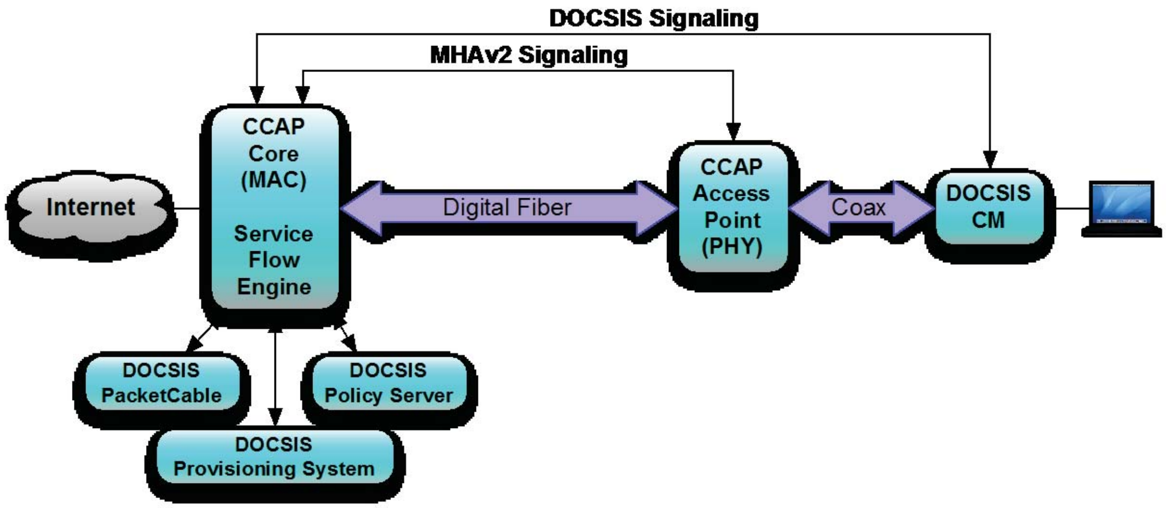

Figure 2 - DEPI Pseudowire

The DEPI signaling is between the CMTS (CCAP) core and the access point. The DOCSIS signaling is between the CMTS (CCAP) and the CM. In general, the Remote PHY access point is transparent to the DOCSIS signaling. One of few exceptions to this rule is that the access point has to regenerate the DOCSIS SYNC message in the downstream.

From a DOCSIS viewpoint, for all intents and purposes the M-CMTS Core and the EQAM act as the equivalent of an I-CMTS (Integrated CMTS). The DOCSIS protocol is the same for M-CMTS and I-CMTS, and the CM is unaware of the difference.

It is worth noting that DEPI is not specifically distance limited. In current M-CMTS deployments, the M-CMTS Core is connected to an EQAM and both of these chassis are co-located because the upstream connections are still on the CMTS-Core.

Both chassis are also connected to a common timing source called DTI (DOCSIS Timing Interface). A single link of DTI is limited to 200 meters. There are methods for having multiple chassis of DTI that can extend the distance, or even multi-site synchronization of DTI chassis, although these extensions were never deployed. The distance limitation of DTI is sometimes thought to limit the distance of DEPI. However, DEPI itself is not distance limited.

The reality is that the M-CMTS scenario does need a clocking solution. So, if a different clocking solution such as IEEE-1588v2 where used, then the solution could be supported over larger links without impacting the definition of DEPI.

To analyze DEPI, we can look at the data plane and the control plane separately.

DEPI Data Plane

The data plane is best described by its choices of encapsulation. DEPI has two encapsulations:

- D-MPT (DOCSIS MPEG Transport): DOCSIS frames over MPEG over L2TPv3
- PSP (Packet Streaming Protocol): Streaming DOCSIS frames over L2TPv3

D-MPT Mode is what is deployed today with M-CMTS and EQAMs in the downstream. PSP did not get deployed in that scenario, but did get used in the upstream direction as will be explained in the next section.

The format of an IP packet that contains L2TPv3 is shown in Figure 3. There are several points worth noting.

- The outer layer 2 encapsulation is Ethernet. This could be any layer 2 encapsulation, such as EPON.
- The next layer is IPv4 or IPv6 as the layer 3 protocol.
- The next layer is L2TPv3. The first field that is always there is the 32-bit session ID. Notice that the L2TPv3 session ID lines up in the same spot that the UDP Source and Destination ports would. This allows existing packet filters that can filter on UDP ports to be able to filter on an L2TPv3 session ID.
- If the L2TPv3 session ID is all zeros, this is a control plane packet. If it is non-zero, this is a data plane packet.
- The L2TPv3 sub-layer header is user definable, although there are some common fields used in various pseudowires that are often included. DEPI and UEPI include those common fields in the second 32-bit field.

For single carrier QAM channels, such as in DOCSIS 3.0 and earlier, there is one DEPI pseudowire for each instance of a QAM Channel and MAC channel pair.

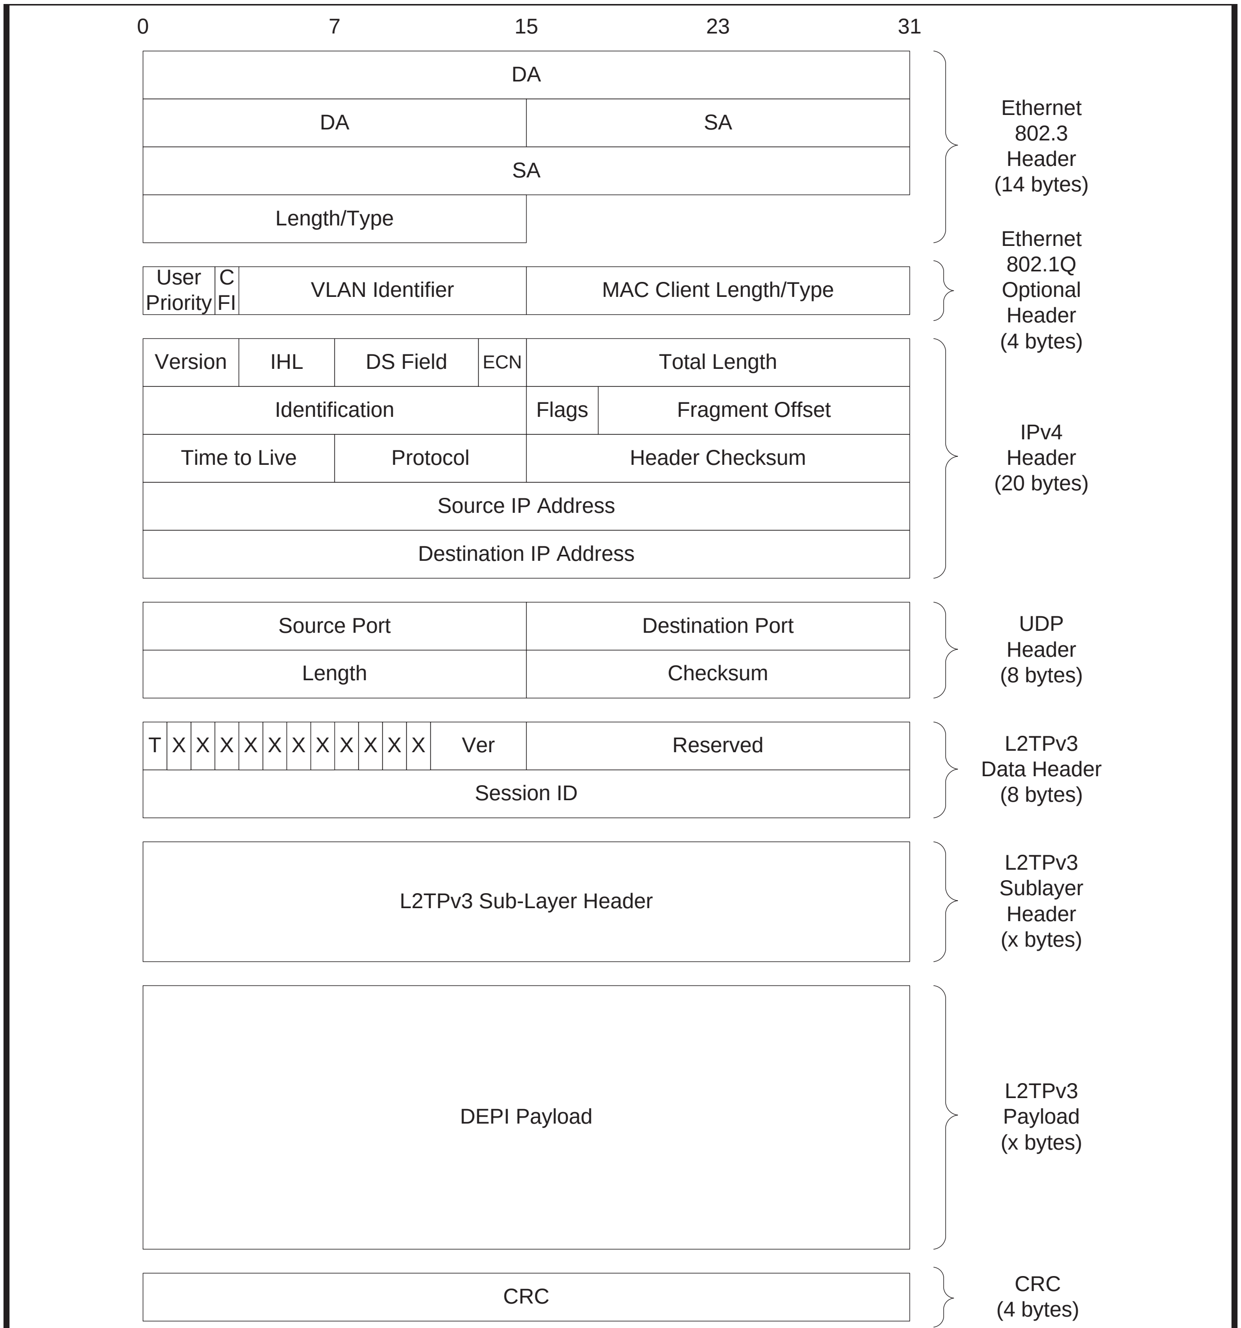

Figure 3 - L2TPv3 Generic IP Packet

The DEPI sub-layer header and payload are shown in Figure 4. The V bit allows the multiplexing of some monitoring protocols but is not used. The sequence number is not needed for resequencing but is often enabled in order for the receiver to detect dropped packets. The Flow ID permits some QoS treatments of flow, although this is not very useful for D-MPT mode as 100% of the DOCSIS MAC frames are all contained in the same DEPI pseudowire.

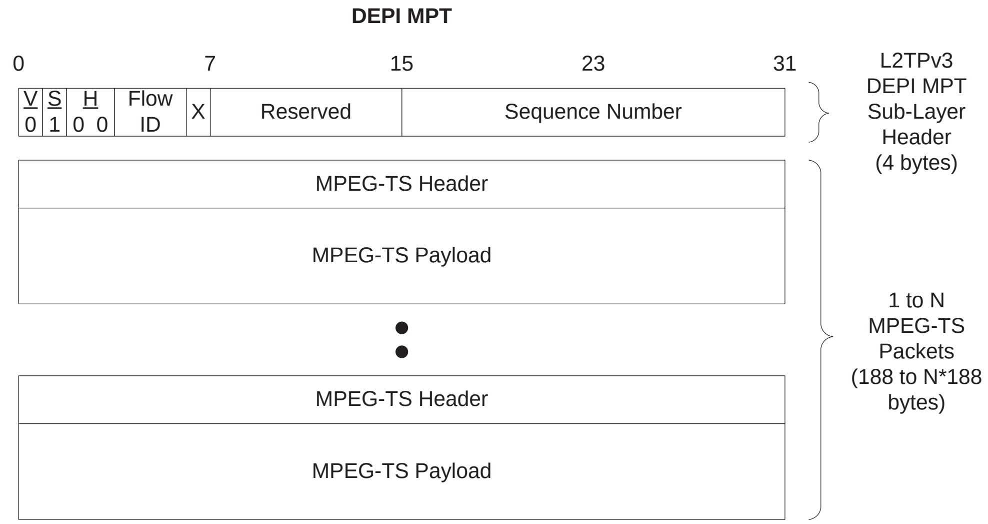

Figure 4 - DEPI D-MPT Mode

The DEPI D-MPT payload contains the exact same MPEG-TS frames that are defined in the DOCSIS specifications. In the payload of those MPEG-TS packets are the DOCSIS frames. With this format, the Remote PHY can take the MPEG-TS packets directly from the DEPI packet and map them into the SC-QAM.

The second encapsulation mode is the Packet Streaming Protocol, or PSP. PSP was the forerunner of the upstream bonding protocol in DOCSIS 3.0. PSP takes a series of packets. These packets can then be broken individually into segments. These segments are then placed into PSP packets. The PSP encapsulation is shown in Figure 5.

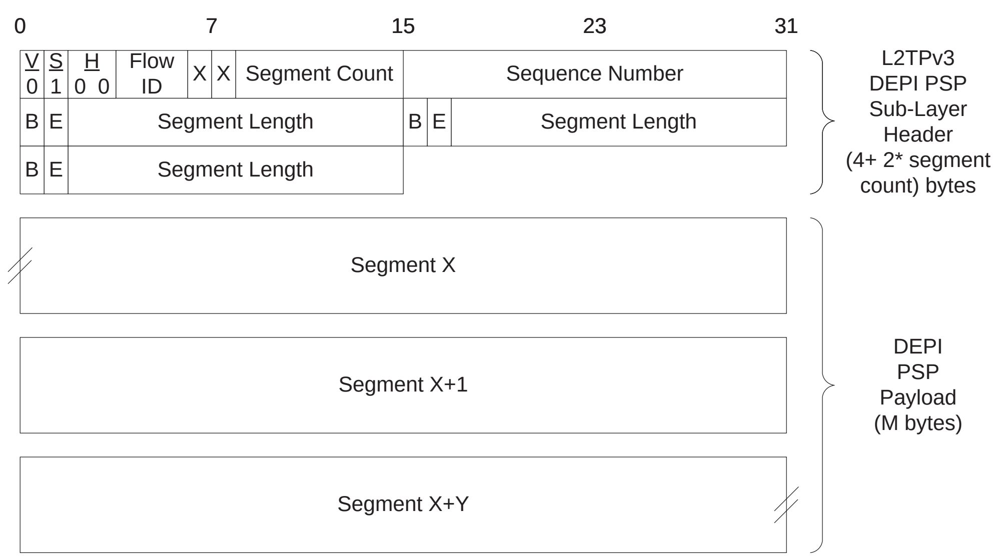

Figure 5 - DEPI PSP Mode

The first 32 bits are similar although there is an additional field that lists how many segments are present. The entire segment table is placed at the header so it can be directly written and read by software without having to traverse the actual segments. Each segment is marked as a beginning, middle, or end segment.

DEPI Control Plane

The L2TPv3 control plane performs the following functions.

1. Session set up and tear down (between two IP endpoints)
2. Pseudowire set up and tear down (per MAC/PHY channel)
3. Reliable message delivery

DEPI has extended the control plane to include:

4. N+1 Pseudowire redundancy.
5. A specified subset of the L2TPv3 AVPs that are used for MHA.
6. New AVPs to describe the D-MPT and PSP pseudowires
7. New AVPs that permit the configuration of a SC-QAM remote PHY.

AVP refers to Attribute Value Pair. It is a data field that can be combined with other data fields in a message. It is similar in concept to a TLV (Type Length Value), although the format is slightly different.

UEPI

UEPI is the Upstream External PHY Interface. There was an earlier MAC chip to PHY chip interface called DMPI and put into the DOCSIS specifications. DMPI stands for DOCSIS MAC PHY Interface. It defined all the necessary information that comes out of a CMTS burst receiver.

UEPI is an extension of DEPI and uses a variation of the PSP encapsulation. UEPI is essentially DMPI mapped into the DEPI PSP encapsulation. The UEPI control plane is the same L2TPv3 control plane that DEPI uses with additional extensions for UEPI.

A UEPI MAC PHY system diagram is shown in Figure 6.

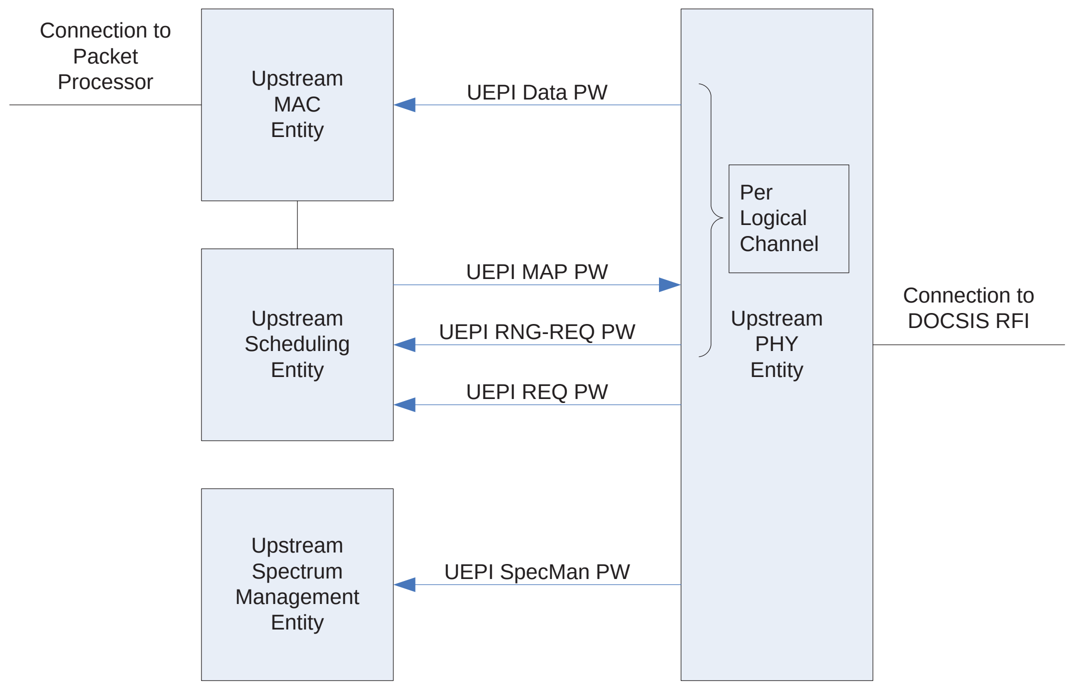

Figure 6 - UEPI System

There are five pseudowires defined for UEPI. The first four are unique per QAM channel. The fifth pseudowire could be per QAM channel or per chip entity.

The UEPI data pseudowire is the DMPI and PSP mixture referred to earlier. Normally, this pseudowire contains 100% of the traffic that is generated by the PHY chip. One of the nice properties of PSP is that the packet flow within a DOCSIS channel can be broken out into separate sub-flows.

This is precisely the idea behind the next three pseudowires. There is a separate pseudowire for the DOCSIS MAP, the DOCSIS RNG-REQ and the DOCSIS REQ messages. The intent is that these messages could be separated from the data traffic, put on different pseudowires, and then those pseudowires could be given higher QoS treatment. In practice, this performance mode is not needed if there is sufficient network bandwidth on the UEPI interface such that congestion never occurs.

The last pseudowire is the spectrum management (SpecMan) pseudowire. At this time, there is no proposed format for the SpecMan pseudowire since the spectrum management designs are often proprietary to a particular product and are not well specified by DOCSIS.

Summary

For MHAv2 to be useful and successful, it has to be driven by some market requirement. That market requirement is not that MHAv2 is a cool way to build a CMTS. It is that there is a need for digital HFC. If there is a business and economic need to put digital fiber into the HFC plant, then MHAv2 becomes the tool that allows it to happen while preserving the rich DOCSIS system architecture.

DOCSIS Remote PHY offers the simplest implementation for a remote node while still being able to traverse an IP network. It does this by minimizing the amount of hardware and software located remotely and also maintains the higher compatibility with existing code at the centralized CMTS.

Abbreviations and Acronyms

AVP	Attribute Value Pair
BDF	Baseband Digital Forward
BDR	Baseband Digital Reverse
C-DOCSIS	China DOCSIS
CAPEX	Capital Expenditures
CIN	Converged Interconnect Network
CLI	Command Line Interface
CM	Cable Modem
CMTS	Cable Modem Termination System
DEPI	DOCSIS External PHY Interface
DOCSIS	Data Over Cable System Interface Specification
DTI	DOCSIS Timing Interface
EQAM	Edge QAM
ERMI	Edge Resource Manager Interface
HFC	Hybrid Fiber Coax
I-CMTS	Integrated CMTS
L2TPv3	Layer 2 Tunneling Protocol Version 3
LAC	L2TP Access Concentrator
M-CMTS	Modular CMTS
MAC	Media Access Control
MHA	Modular Headend Architecture
MHAv2	Modular Headend Architecture version 2
MPLS	Multiprotocol Label Switching
OFDM	Orthogonal Frequency Division Multiplexing
OPEX	Operational Expenses
OSSI	Operation Support System Interface
PASI	Packet Shelf to Access Shelf Interface
PSN	Packet Switched Network
PMI	Provisioning and Management Interface
QAM	Quadrature Amplitude Modulation
RPHY	Remote PHY
SARFT	State Administration of Radio, Film, and Television
TLV	Type Length Value
UEPI	Upstream External PHY Interface
VSI	Video Stream Interface