Distributed Network Architectures for Next-Generation Cable Access

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

Niki Pantelias
Associate Technical Director
Broadband Technology Group
Broadcom Corporation
Duluth, GA 30096 USA
+1-770-232-0018
nikip@broadcom.com

Introduction

In most current HFC deployments, CMTS functions are concentrated within operator buildings such as a local headend or hub. However, the industry continues to explore options for partitioning CMTS functions and distributing them among other locations in the network, particularly the fiber node. Recent technology advances have made some of these options more appealing than they have been in the past. Increasing chip densities, reduction in power consumption per megabit, and the use of direct digital synthesis to generate the downstream RF signal now make it feasible for an entire bidirectional digital channel lineup to be generated/terminated in the node. Increases in CMTS and CCAP port densities are being matched or exceeded by projected future increases in customer demand for faster broadband services, making building space and power an ongoing consideration for some operators. And, the advent of higher modulation orders in DOCSIS® 3.1 offers strong motivation for operators to maximize cable plant SNR by replacing current HFC analog optics with better-performing, lower-cost digital links.

Many possible distributed CMTS architectures have been proposed, each with its own strengths in addressing specific areas of industry concern. Some have already been tried and adopted to varying degrees, and these approaches may offer a foundation upon which to build when considering future proposals.

This paper attempts to shed some light on the various distributed CMTS proposals by comparing them in several important areas. For this discussion, architectures are grouped into three broad classes relevant to the points to be considered.

Classes of Distributed Architectures

Figure 1 illustrates today's Integrated CMTS architecture, showing the elements most central to the discussion of distributed versions. Each of the classes of distributed architecture relocates various elements in various ways.

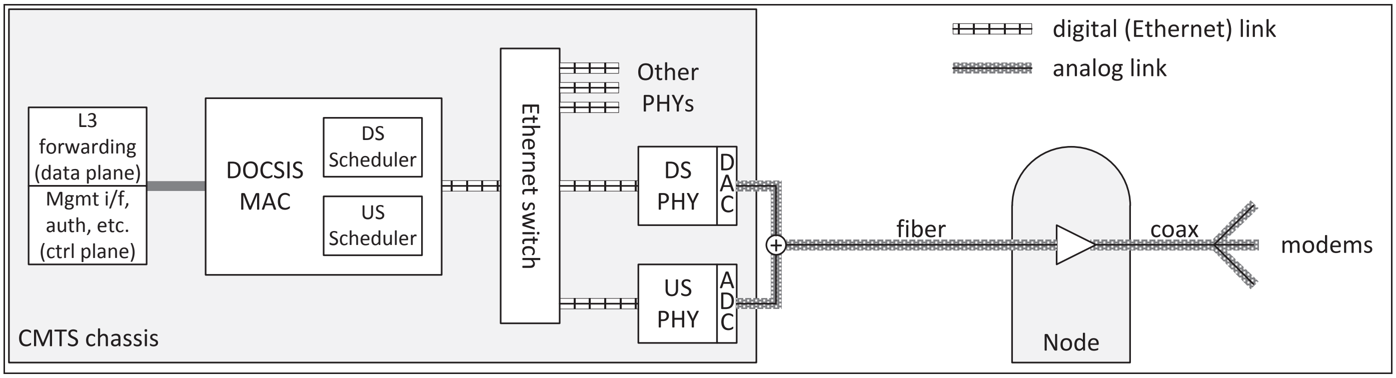

Figure 1 Integrated System Architecture Example

In Figure 1, the functionality of the CMTS chassis located in the headed begins at the layer 3 WAN interface and continues through the DOCSIS® layer 2 MAC, then across an internal digital interface to the downstream and upstream PHYs, including digital synthesis (downstream) or sampling (upstream) of the RF signal to be transmitted/received on the analog fiber link to the node.

A feature of Figure 1 which will be referenced later in this discussion is the interface between the MAC and PHY blocks. In many currently deployed CMTSs, this interface runs over an Ethernet switch fabric internal to the CMTS. Ethernet switching provides a high-rate, low-cost, and easy-to-use many-to-one or many-to-many interconnect between devices, allowing for flexible capacity matching. For example, a high-density MAC device, which must be used to achieve target CMTS port densities, may have enough capacity to support multiple PHY ports; with Ethernet switching, a single high-speed interface on the MAC device can reach as many PHY devices as necessary using readily available components and a standardized protocol.

A key objective of most distributed architectures currently being discussed is eliminate analog lasers over the fiber portion of the HFC. This can give significant improvement in plant SNR by eliminating the contribution of analog laser distortion to the noise and impairments of the plant. It also enables a significant increase in the total data rate delivered to the node, making it possible to provide other types of services (e.g. Metro Ethernet) by using WDN to deliver these services over different wavelengths of the same fiber. In the distributed CMTS architectures considered here, instead of analog

RF signal transmission, some form of digital link is used to connect the local headend or hub location (hereafter simply referred to as the "headend") with the fiber node, where the desired analog RF signal is generated (in the downstream) or terminated (in the upstream). Thus, analog transmission occurs only over the coaxial portion of the HFC plant. All three of the distributed architectures discussed here incorporate this feature, so all can be expected to offer similar RF signal quality.

The most straightforward way to replace the analog fiber link with a digital one, with minimal impact to existing architectures, is to use the architecture class termed "Remote DAC/ADC" in this paper. Figure 2 illustrates an example of such an architecture.

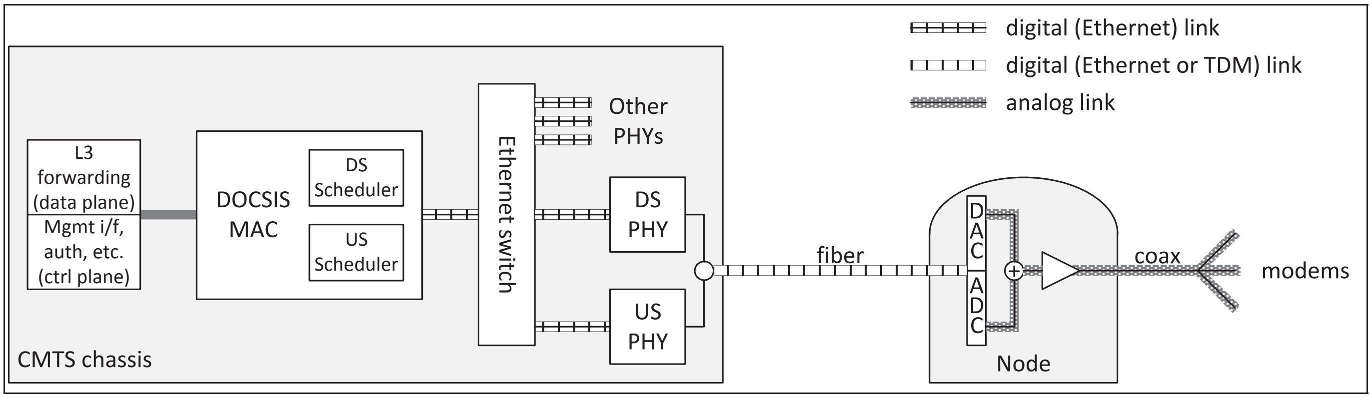

Figure 2 Remote DAC/ADC System Architecture Example

In the Remote DAC/ADC system, only the final step of digital-to-analog conversion (downstream) or analog-to-digital conversion (upstream) is moved to the fiber node; the rest of the CMTS chassis remains unchanged. The analog fiber link of Figure 1 is replaced in Figure 2 by a digital link carrying samples representing the RF signals. This link may use a TDM-over-fiber technology, such as SONET or Fibre Channel, or it may encapsulate samples into Ethernet packets that can be carried using an Ethernet-over-fiber technology.

Currently, a class of products known as "digital return" is enjoying wide acceptance in the industry. These products can be thought of as examples of a Remote DAC/ADC architecture without the DAC. An overview of digital return technology is provided in [1], and the technical aspects of one such implementation are described in [2].

The next class of architectures to be considered relocates somewhat more functionality by bringing the complete downstream and upstream PHY functions into the node. In addition to the DAC and ADC, these functions include FEC encoding/decoding, digital modulation/demodulation and upconversion/downconversion, digital RF combining, timing generation and recovery, and so forth. Figure 3 illustrates an example of a "Remote PHY" system.

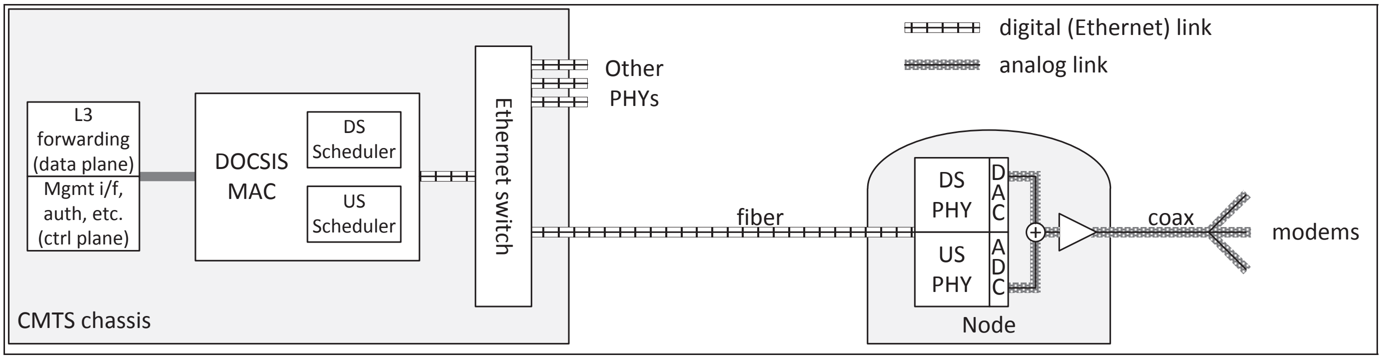

Figure 3 Remote PHY System Architecture Example

The Remote PHY architecture can be thought of as simply “stretching” the Ethernet-based MAC-PHY interface currently in use in many CMTSs so that it now extends across the fiber to the node. Viewed in this way, it exploits a natural “breakpoint” within the CMTS. This is apparent when Figure 3 is compared with Figure 1.

An existing example of a Remote PHY approach is the MHA (Modular Headend Architecture) family of CableLabs® standards [3]. In MHA, the upstream PHY remains a part of the CMTS chassis, so MHA can be thought of as a “Remote Downstream PHY” architecture. The primary objective of MHA was to expand the downstream capacity of a CMTS chassis by allowing it to utilize the downstream ports of an EQAM device to carry DOCSIS® traffic. An Ethernet/IP link connects the CMTS’s MAC function with the PHY function provided by the EQAM, usually within the same building (the EQAM is not normally remoted to a fiber node).

The third class of architectures to be discussed is termed “Remote MAC-PHY” in this paper. There are many possible variants of this architecture, but in all of them, the fiber node contains the DOCSIS® MAC, including the upstream and downstream schedulers and related message processing. (Architectures that place only low-level MAC functions such as encryption or header creation in the node are better described as variants of Remote PHY for purposes of the comparisons done in this paper.) An example of a Remote MAC-PHY architecture is shown in Figure 4.

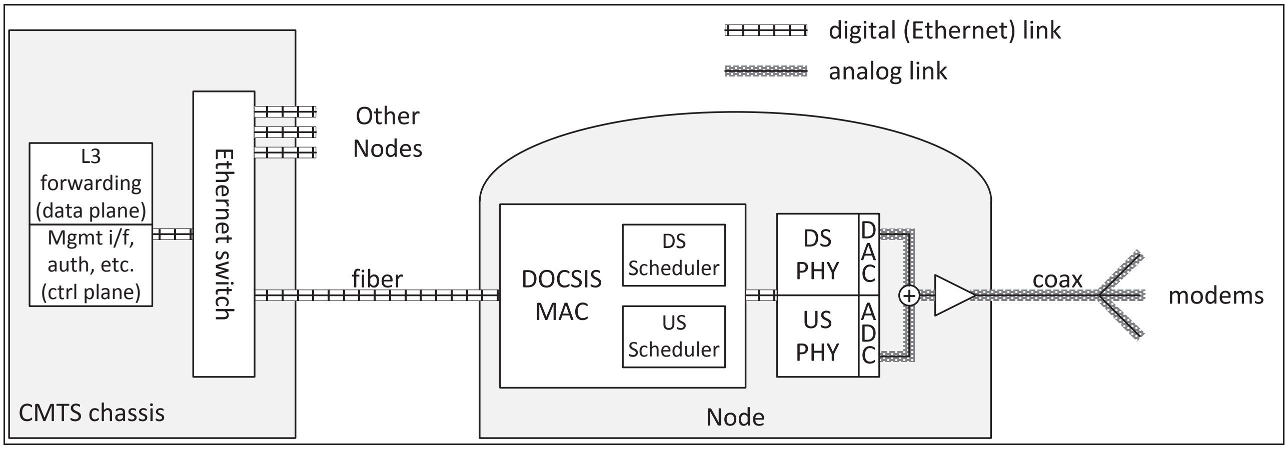

Figure 4 Remote MAC-PHY System Architecture Example

The Remote MAC-PHY approach lends itself to extensive variation depending on exactly how various higher-layer MAC, routing, and control functions are handled. Previous efforts in this area have partitioned DOCSIS® data plane processing, assigned partial handling of control messages to each end of the link, and even attempted to divide the scheduling task among two physically distant devices.

Two examples of Remote MAC-PHY approaches that have been successfully standardized are the C-DOCSIS I and C-DOCSIS II architectures described in the C-DOCSIS specification [4]. The C-DOCSIS II architecture locates DOCSIS® downstream classification and upstream policing in the headend to leverage the capabilities of OLT devices already present there and to reduce the cost of the hardware in the node; the remaining DOCSIS® data plane functions are located in the node. The C-DOCSIS II control plane handles most DOCSIS® MAC Management Messages in the node, with support from an external server for authentication and admission decisions. The external server also provides the management interface to the operator's back-end systems. In C-DOCSIS I, essentially all CMTS functions, including layer 3 forwarding and the management system interface, are located at the node, making a C-DOCSIS I node a kind of "mini-CMTS."

The three classes of architectures will be compared to each other and to the Integrated architecture in the sections that follow.

Points of Comparison

This paper compares the three classes of distributed architectures in the areas of network throughput requirements, impact on DOCSIS® Request-Grant Round-Trip Time, and protocol complexity.

Network Throughput Requirements

A key component of a distributed CMTS system is the digital link between the headend and the node, covering the fiber portion of the HFC network. Digital optics generally cost less than analog optics, and this cost savings is part of the motivation for adopting a distributed architecture. The type and rate of the digital link will have a major impact on the cost of the optical components. Generally, a higher-rate connection will have a higher component cost and a more stringent optical link budget than a lower-rate connection, and a TDM technology (for instance, SONET OC-192) will be more expensive and less flexible than a packet-based technology (for instance, 10 Gigabit Ethernet).

The Remote PHY and Remote MAC-PHY architectures both use an Ethernet link to transfer data packets destined for RF transmission, plus some encapsulation or tagging protocol and a modest amount of control plane traffic. Thus, the data rate required for the Ethernet link is approximately the aggregate RF line rate plus some protocol overhead.

The Remote DAC/ADC architecture is fundamentally different from the other two in that the data packets to be transmitted on the RF link are not sent directly over the digital fiber connection. Instead, the data stream to be transmitted is expressed in the form of digital DAC or ADC samples to be “played out” at the other end of the link to reconstruct the desired RF signal. The samples are generated and consumed at a fixed rate, thus lending themselves to transmission over a TDM link. It is also possible to encapsulate the samples in Ethernet packets to allow use of lower-cost Ethernet transceivers. In either case, the required data rate for the link is close to the product of the sample rate needed to cover the desired RF bandwidth and the ADC or DAC bit depth required to support the target RF fidelity. There is protocol overhead but it is probably somewhat smaller than in the other two architectures.

The discussion and tables below show some sample calculations for required network throughput considering a worst-case system in each direction. Future goals for new frequency plans and higher digital modulation orders, such as might be achieved in a DOCSIS® 3.1 system, are taken into account in describing a worst case.

In the downstream direction, the worst case is assumed to be a downstream RF spectrum of 54 MHz to 1.2 GHz. This spectrum is assumed to accomodate 5 DOCSIS®

3.1 OFDM channels, each occupying 192 MHz of RF bandwidth and supporting 4096 QAM modulation on all subcarriers, *plus* 31 SC-QAM (DOCSIS® 3.0 or earlier) channels using J.83 Annex B parameters with 256 QAM modulation. This is almost certainly an overestimate of the channel capacity, since in practice there would be some unusable areas of spectrum (e.g. due to interference from the FM band or from certain cellular services) and other areas of spectrum which cannot support the highest available modulation orders (e.g. due to impairments such as diplexer rolloff). However, it is useful to examine this case both as an upper bound on required throughput and as a way of comparing the available options.

In the upstream direction, the worst case occurs if the diplexer split is moved to a frequency much higher than the current one. Many operators are discussing future splits in the vicinity of 200 MHz (upstream upper band edge frequency). This would comfortably accomodate 2 DOCSIS® 3.1 OFDMA channels, each potentially occupying 96 MHz of RF bandwidth and supporting 4096 QAM modulation on all subcarriers. (In practice, it is likely that some subcarriers would be forced to use lower modulation orders or be disabled entirely due to impairments, especially in the spectrum below 20-25 MHz and in the FM band.) However, for a true worst case, consideration should be given to a higher split – say, 300 MHz – which allows the two full-width OFDMA channels to be located above a full lineup of 12 legacy SC-QAM (A-TDMA or S-CDMA) channels occupying the 5-85 MHz band. Providing support for up to 300 MHz of bandwidth also provides a degree of “future-proofing” if splits higher than 200 MHz are adopted later.

The tables below show the calculation of required network throughput in the downstream direction for the two cases described above. For the Remote DAC/ADC architecture, it is assumed that a 14-bit DAC and a 12-bit ADC are required to generate RF signals of sufficient fidelity to support 4096 QAM modulation.

Downstream spectrum:	54 - 1200	MHz
OFDM channel parameters:		
Channel bandwidth	192	MHz
Modulation order	12	bits/second-Hz
Max "Raw" bit rate	2304	Mbits/second
FEC coding rate	0.89	(best case)
Max data rate	2051	Mbits/second on each OFDM channel
SC-QAM channel parameters:		
FEC/coding	J.83 Annex B, 256 QAM	
"Raw" bit rate	42.88	Mbits/second
FEC coding rate	0.905	
Data rate	38.8	Mbits/second on each SC-QAM channel
Number of OFDM channels	5	
Total max OFDM data rate	10252.8	Mbits/second
Number of SC-QAM channels	31	
Total SC-QAM data rate	1203	Mbits/second
Total max data rate	11456	Mbits/second
Maximum protocol overhead	up to 10%	
Total rate on digital fiber link	**12601**	Mbits/second

Table 1 Required Digital Fiber Downstream Link Rate for Remote PHY and Remote MAC-PHY

Downstream spectrum:	54 - 1200	MHz
Ratio of sample rate to bandwidth:	2.5	
Sample rate	3000	Msamples/second
Bits/sample	14	bits/sample
Total bit rate of samples	42000	Mbits/second
Protocol overhead	up to 5%	
Total rate on digital fiber link	44100	Mbits/second

Table 2 Required Digital Fiber Downstream Link Rate For Remote DAC/ADC

Upstream spectrum:	5 - 300	MHz
OFDMA channel parameters:		
Channel bandwidth	96	MHz
Modulation order	12	bits/second-Hz
Max "Raw" bit rate	1152	Mbits/second
FEC coding rate	89%	(best case)
Pilot overhead	2%	
Max data rate	1005	Mbits/second on each OFDM channel
SC-QAM channel parameters:		
Channel bandwidth	6.4	MHz
Symbol rate	5.12	Mbaud
Modulation order	6	bits/second-Hz
"Raw" bit rate	30.72	Mbits/second
FEC coding rate	93.33%	
Preamble/guard	3%	
Data rate	27.8	Mbits/second on each SC-QAM channel
Number of OFDMA channels	2	
Total max OFDMA data rate	2010	Mbits/second
Number of SC-QAM channels	12	
Total SC-QAM data rate	334	Mbits/second
Total max data rate	2343	Mbits/second
Maximum protocol overhead	up to 10%	
Total rate on digital fiber link	**2578**	Mbits/second

Table 3 Required Digital Fiber Upstream Link Rate For Remote PHY and Remote MAC-PHY

Upstream spectrum:	5-300	MHz
Ratio of sample rate to bandwidth:	2.5	
Sample rate	750	Msamples/second
Bits/sample	12	bits/sample
Total bit rate of samples	9000	Mbits/second
Protocol overhead	up to 5%	
Total rate on digital fiber link	9450	Mbits/second

Table 4 Required Digital Fiber Upstream Link Rate For Remote DAC/ADC

Table 5 summarizes the comparison of network throughput requirements for the three classes of distributed architectures.

Maximum Possible Downstream Rate:		
Remote DAC/ADC	44100	Mbits/second
Remote PHY and Remote MAC-PHY	12601	Mbits/second
Maximum Possible Upstream Rate:		
Remote DAC/ADC	9450	Mbits/second
Remote PHY and Remote MAC-PHY	2578	Mbits/second

Table 5 Summary Comparison of Required Digital Fiber Link Rates

It can be seen that the Remote DAC/ADC architecture requires significantly higher digital network throughput than the other two approaches. This could result in more expensive optics. However, it is also important for the operator to consider how the link throughput requirements fit into the rates readily available. For example, in the downstream direction, a Remote PHY system may require somewhat more throughput than can be provided over a single 10 GbE link. This may necessitate the use of two such links (most likely multiplexed onto the same fiber using WDM). Or, the operator may conclude that the worst case above does not actually match the specific deployment in question; for example, if the available downstream spectrum is 252 MHz to 1.0 GHz, the maximum downstream rate will be significantly lower and can easily be supported on a single 10 GbE connection. Conversely, in the upstream direction, a

Remote DAC/ADC system can use an OC-192 link with almost full utilization, while a Remote PHY or Remote MAC-PHY deployment would use only about 1/3 of a similarly-sized (but lower cost) 10 GbE link. In this case, the operator may consider "daisy chaining" two or three nodes in the upstream direction so that the link can be more fully utilized, further reducing cost. This approach would necessitate careful network planning as described in the section below.

Impact On DOCSIS® Request-Grant Round-Trip Time

One important contributor to overall performance in a DOCSIS® system is the DOCSIS® Request-Grant Round-Trip Time. This is the time from the moment of transmission of a DOCSIS® Request Frame or piggyback Request by the CM to the moment at which the CM begins transmitting in a Data Grant provided by the upstream scheduler in response to the Request.

Components of the Round-Trip Time include (among other things) the time required for:

- upstream transmission, propogation, and reception of the Request;
- scheduling of a Data Grant by the upstream scheduler, taking into account the interval between MAP messages and other factors;
- downstream transmission, propogation, and reception of the MAP message containing the Data Grant;
- additional "MAP lead time" which ensures that the modem receives the Data Grant far enough in advance of the start of data transmission to complete necessary processing prior to transmission.

Typical values of Round-Trip Time in deployed DOCSIS® 3.0 systems are on the order of 4-6 milliseconds, but may vary depending on plant distances and configurations, network congestion, and other factors.

Round-Trip Time impacts DOCSIS® system performance in a number of ways. Two examples of different but related impacts of Round-Trip Time are as follows:

(1) *As a component of end-to-end latency (Frame Delay).* As a reservation-based point-to-multipoint system, DOCSIS® requires that a cable modem transmit data only in certain time slots reserved for this purpose by the CMTS. For most DOCSIS® scheduling services (except for Unsolicited Grant Services), the CMTS will not normally provide a grant until it first receives a Request from the modem indicating that the modem has data to transmit. Thus, the maximum total latency experienced by a packet arriving at a cable modem's CMCI port for transmission upstream includes the Request-Grant Round-Trip Time as well as the time for transmission of the packet itself. Since the Round-Trip Time includes the full delay of both the upstream and downstream DOCSIS® paths, it is usually the largest component of end-to-end DOCSIS® latency in the upstream direction.

Keeping the Round-Trip Time low is critical for applications such as business services where the SLA includes stringent latency requirements. It is also vital for applications such as gaming in which poor latency performance can result in highly vocal customer complaints which are challenging to resolve, even if a specific latency guarantee is not part of the SLA.

(2) *As a contributor to (or detractor from) TCP protocol performance.* TCP is the layer 4 protocol most commonly used for web browsing and file transfer. It attempts to match transfer speed to available bandwidth by detecting how quickly the network responds to previously transmitted packets. With some TCP algorithms, a delay in transmission and acknowledgement of the first few packets of a transfer may cause the protocol to decide that the link is slower than it actually is, resulting in a delay before the transfer speeds up to full utilization of bandwidth provided. The end result is a transfer that starts slowly and takes longer than necessary to complete. Even modest increases in DOCSIS® Round-Trip Time may trigger this behavior, which affects casual users even when the SLA does not include strict latency bounds.

Because Round-Trip Time is such an important parameter, it is worth comparing the three classes of distributed architectures with each other and with an integrated architecture to understand the effects of each option on DOCSIS® Round-Trip Time.

Figure 5 illustrates the Request-Grant path through an integrated CMTS architecture. The Request originates at the CM and terminates at the upstream scheduler. The upstream scheduler creates a MAP message containing a Data Grant. The MAP message is delivered to the downstream scheduler, which queues it for transmission downstream (generally in the highest-priority category available). The cycle is complete when the CM receives the MAP, processes it, and begins transmission in the data grant.

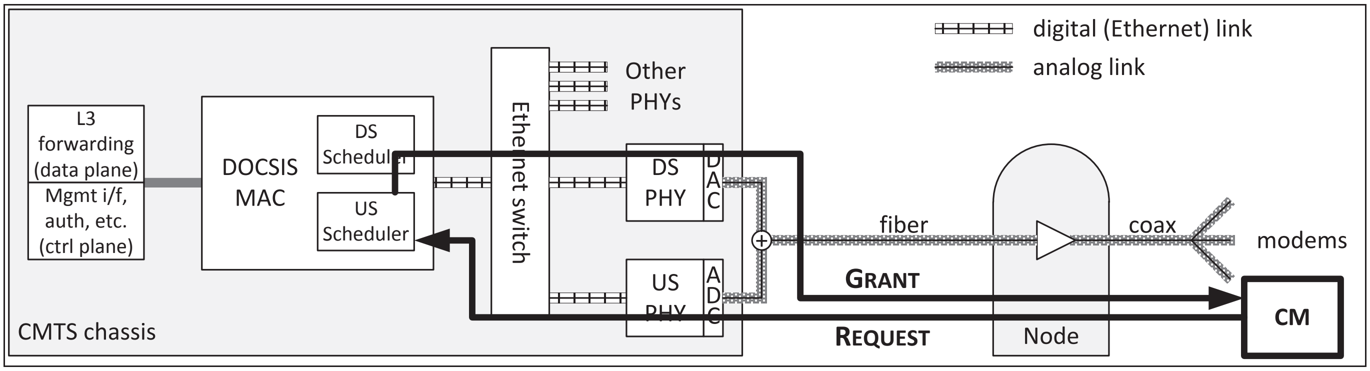

Figure 5 Request-Grant Path For Integrated Architecture

The distributed architecture with the most straightforward Request-Grant path is the Remote MAC-PHY architecture. The Request-Grant path for Remote MAC-PHY is illustrated in Figure 6.

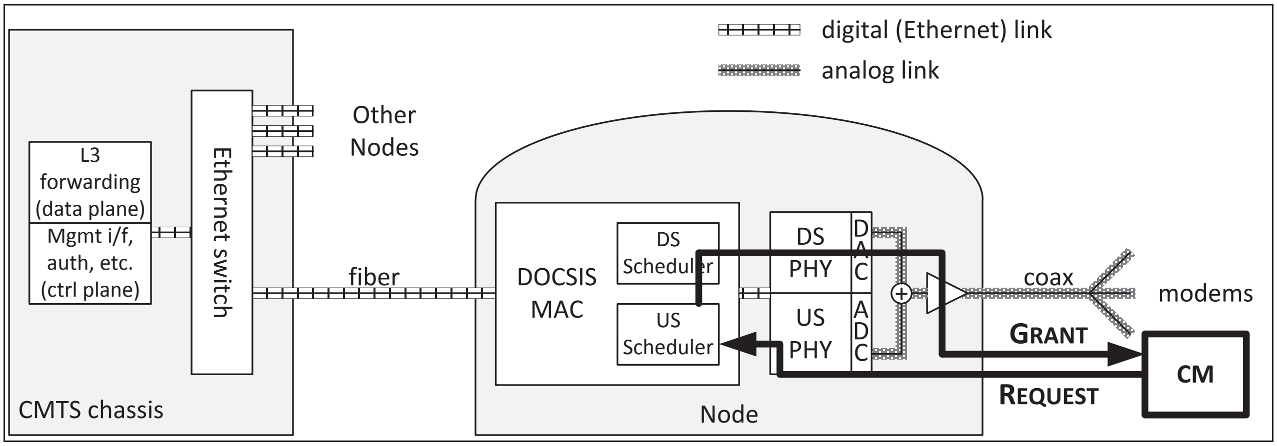

Figure 6 Request-Grant Path For Remote MAC-PHY Architecture

It can be seen in Figure 6 that the Request-Grant path for the Remote MAC-PHY traverses the same system elements as in the integrated case (with the exception of the Ethernet switch which may not be present between MAC and PHY in the Remote MAC-PHY case). The key difference is that the physical extent of the network is shorter; the Request and Grant need only traverse the distance of the coaxial network, not the combined distance of the fiber+coaxial network as in the Integrated case. All other things being equal, the Remote MAC-PHY architecture will offer a Round-Trip Time which is lower than that of the Integrated architecture by the amount of twice the one-way propogation delay of the fiber portion of the network. This may or may not be significant, depending on the extent of the plant. On a plant with 100 miles of fiber (the maximum supported by DOCSIS®), the Round-Trip Time would be reduced by 1.6 msec, potentially a significant improvement. For shorter fiber runs, the savings would be smaller – for instance, 200 microseconds in a plant with 20 km of fiber.

Figure 7 illustrates the same path for the Remote PHY architecture.

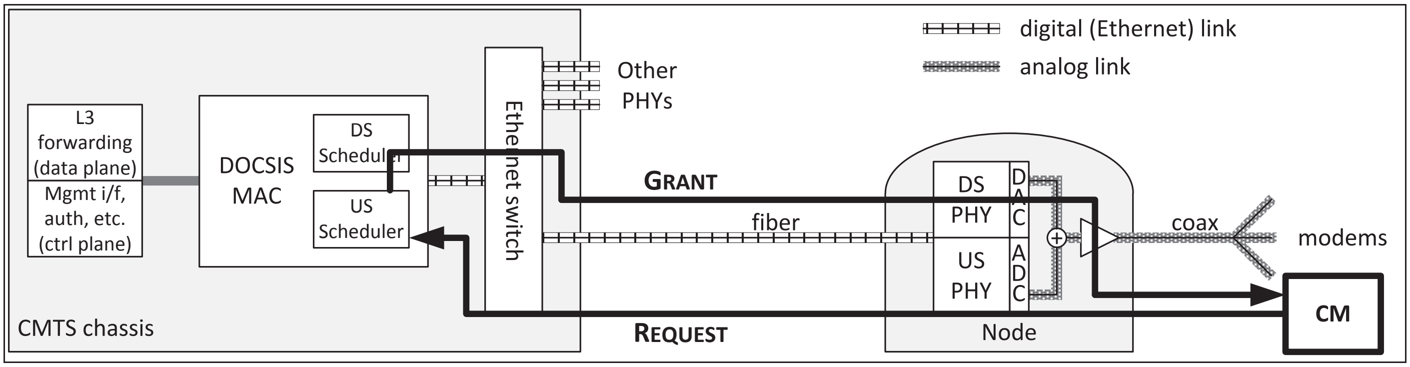

Figure 7 Request-Grant Path for Remote PHY Architecture

It is instructive to compare Figure 7 with Figure 5. In Figure 7, the PHYs have been moved to the node and the Ethernet link between MAC and PHY has been “stretched” to traverse the fiber link, rather than originating and terminating entirely within the CMTS chassis. The digital Ethernet portion of the path is longer, while the analog portion of the path is shorter by the same amount. In both diagrams, the same network elements are traversed in the same order and over the same total physical difference. Thus, in the network as illustrated, the Round-Trip Times of the Integrated and the Remote PHY architectures are essentially identical.

A key assumption behind the preceding statement is that the topologies of the digital Ethernet networks shown in Figure 5 and Figure 7 are in fact identical. Thus, it is essential to examine what happens if this assumption does not hold.

In Figure 5, the Ethernet link between MAC and PHY passes through a single switch device which is almost certainly a line-rate non-blocking device with more than enough throughput on each port to deliver the maximum anticipated data rate required for that port. Congestion within the switch only occurs when multiple packets destined for the same output port arrive on different input ports at the same time. If the output rate is greater than the sum of the input rates, the delay due to this type of congestion is limited to the sum of the line-rate duration of each of the simultaneously received packets. At the data rates in question (usually 10 GbE or higher for a DOCSIS® 3.1 system), the line-rate duration of a packet is on the order of 1 microsecond, and the maximum delay due to congestion within the switch is in the low tens of microseconds – definitely negligible in comparison with a DOCSIS® Round-Trip Time of several milliseconds.

For the Remote PHY architecture of Figure 7, the operator can duplicate this performance simply by providing a point-to-point Ethernet-over-fiber link dedicated to DOCSIS® traffic between a single switch port at the CMTS chassis and a single pair of Downstream/Upstream PHYs at the node. (The downstream and upstream PHYs can be paired since each primarily utilizes the link in only one direction, while the Ethernet link is bidirectional.)

However, there are various reasons why an MSO may prefer not to provide a dedicated DOCSIS®-only link between the headend and the node. As an example, the operator might plan to use only part of the RF spectrum for DOCSIS® channels, with the remainder used for digital video QAM channels. In this case, it might be useful to use the Ethernet link between the headend and the node both for DOCSIS® traffic and for digital “IPQAM” video traffic to be modulated and upconverted at the fiber node. Or, in some plant topologies, the operator may want to have a single Ethernet switch port at the CMTS chassis serve multiple fiber nodes in a “daisy chained” configuration.

Such topology variations can be supported with little or no impact on DOCSIS® Round-Trip Time, provided that the network is carefully planned and managed. All traffic to be sent on the headend-to-node Ethernet link should be carefully characterized and shaped if it is not already. The sum of the anticipated maximum traffic rates from all sources should not exceed the capacity of the link. The aggregate of the maximum burst size/rate from all sources, combined with any queueing/prioritization algorithm that may be in use, will determine the maximum time that a single packet could wait in a queue while other packets are being sent. This queueing delay should be calculated and bounded to ensure that the resulting packet delay and jitter are small enough not to impact system performance. (Total delay and jitter of under 100 microseconds is probably negligible; values of up to several hundred microseconds may be tolerated but should be checked in the context of the overall latency and jitter budgets for the system as whole, taking into account the services to be delivered.)

Most importantly, it is recommended that the operator strictly avoid sharing this network link with uncharacterized or uncontrolled traffic. For example, allowing bursty WAN traffic from the broader internet to share the network between the headend and the node could be very problematic. Though it is possible to make such a scenario work using advanced layer 3 traffic management, the operator may find this to be challenging at best.

Network jitter is of particular concern in a Remote PHY system because, in the downstream direction, jitter on even a single packet can translate into a queueing delay that may persist for an extended period. This issue is described in detail in Appendix I of [5] (especially section I.7). Briefly, it occurs because a packet experiencing higher-than normal delay across the downstream link will block the head of the transmission queue at the downstream PHY, causing packets behind it to be delayed by the same amount. This delay clears only when the downstream RF channel is underutilized for some period of time. Because of this phenomenon, when calculating the maximum latency of the downstream, any jitter on the downstream Ethernet link between the headend and the node should be treated as though it were an additional delay added to many packets. As previously described, careful network planning and traffic characterization should be employed to keep jitter to an absolute minimum.

Finally, Figure 8 shows the Request-Grant path for the Remote DAC/ADC architecture. Again, this Request-Grant path traverses the same physical distance and almost the same network elements as in the Integrated architecture, so the Round-Trip Time should be essentially the same in both cases. The only difference is in the digital fiber link between the headend and the node, which in the case of Remote DAC/ADC may be Ethernet or may be a fiber-based TDM technology.

The Remote DAC/ADC architecture is particularly sensitive to jitter on this link, since any delay variation larger than the sample buffer depth will cause a buffer underrun, resulting in a discontinuity in the reproduced RF signal which could be highly disruptive (as an example, this discontinuity could cause cable modems to lose lock with the downstream signal, temporarily interrupting service). If the link uses a TDM technology, jitter becomes a non-issue; if an Ethernet link is used, the same caveats described for Remote PHY also apply here.

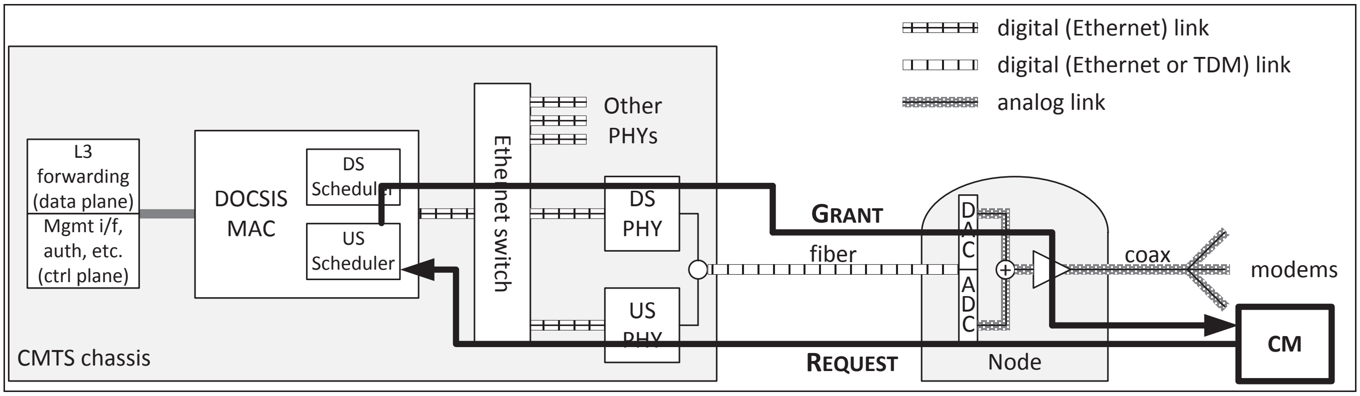

Figure 8 Request-Grant Path For Remote DAC/ADC Architecture

To summarize, the Remote PHY and Remote DAC/ADC architectures can be expected to have DOCSIS® Request-Grant Round-Trip Times very similar to that of the Integrated architecture, provided that the link between the headend and the node is carefully planned, characterized, and managed so that it behaves similarly to the Ethernet switch fabric inside the Integrated architecture's CMTS chassis. The Remote MAC-PHY architecture may have slightly lower Round-Trip Time, and hence slightly better DOCSIS® performance, because of the shorter physical distance traversed by the Request-Grant path.

Protocol Complexity

In order for any distributed architecture to be viable for large-scale deployment, it must be standardized so that multiple vendors can offer fully interoperable products. Thus, the complexity of the protocol used to communicate between the headend and the node is an important factor. A simpler protocol is easier both to standardize and to implement, and multi-vendor interoperability is likely to be demonstrated more quickly, compared with a more complex protocol. A simpler protocol may also be easier for the operator to deploy and maintain. In the data plane, a simple protocol is particularly valuable because data plane processing must be done at a high rate, often requiring application-specific hardware assist to achieve the necessary throughput. On the flip side, a protocol must not be so simple that it lacks features needed to accomplish the task at hand.

None of the architectures under consideration have been fully standardized with DOCSIS® 3.1 requirements incorporated. However, as previously described, implementations do exist for at least some examples of each of the three classes of architectures. To compare protocol complexity, it may be helpful to study these examples and compare the protocols created so far, while recognizing that future standards may differ from current ones, possibly in unforeseen ways.

C-DOCSIS [4] offers two examples of Remote MAC-PHY architectures, termed “C-DOCSIS I” and “C-DOCSIS II.” C-DOCSIS I integrates layer 3 data forwarding and control plane management interfaces into the node, while C-DOCSIS II is more similar to the example Remote MAC-PHY architecture illustrated in this paper (Figure 4 and Figure 6).

In the data plane, C-DOCSIS II uses a protocol called CDT (C-DOCSIS Tagging) to convey Service Flow information between the headend and the node (or vice-versa). CDT is very lightweight, with “encapsulation” consisting only of a single VLAN tag, as shown in Figure 9. C-DOCSIS I actually requires no data plane protocol at all; in this system, the interface on the optical side of the fiber node is analogous to the NSI (Network Side Interface) of a traditional CMTS. The C-DOCSIS I CMC device in the node can receive IP packets directly from the WAN with no special tagging of any kind.

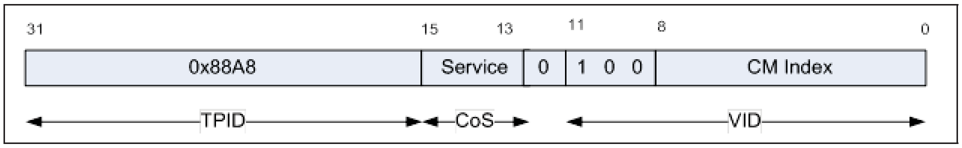

Figure 9 “Encapsulation” (Tag) Format For C-DOCSIS Tagging Protocol

In the control plane, C-DOCSIS II uses a protocol called CDMM (C-DOCSIS Management Messaging) to carry configuration information from an external server to the node, and to allow the external server to gather statistics and monitoring information from the node. The protocol provides a number of messages specific to various types of information to be exchanged and control operations to be performed (such as admitting or deleting modems). The external server is useful if the operator wants to employ a centralized interface point for management of a number of nodes. In such a system, the operator’s management systems interface to the centralized server, which then interfaces to the nodes, solving the potential challenges of having a management system connect directly to a potentially large number of individual nodes. However, this approach is not required; in C-DOCSIS I, each node contains its own complete, independent interface to a management system, and no additional protocol layer is needed.

The most complete existing specification of a Remote PHY system is provided in the CableLabs® MHA (Modular Headend Architecture) family of documents [3]. Although the MHA system remotes only the downstream PHY (the upstream PHY remains a part of the CMTS chassis), the DEPI (Downstream External PHY Interface) protocol between MAC and PHY is fully specified, so it serves as an excellent example of what a Remote PHY protocol might look like. DEPI readily lends itself to extensions for support of DOCSIS® 3.1 data rates and channel types, and similar mechanisms could be used to design a protocol for the upstream direction, unofficially known as “UEPI” (Upstream External PHY Interface).

The DEPI protocol is designed to deliver packets across an arbitrary IP network, and also adds a sublayer of encapsulation using the L2TPv3 protocol to designate constructs specific to DEPI, such as a “PSP flow.” Details of the protocol can be found in [5].

The encapsulation headers for DEPI data plane traffic are shown in Figure 10. These headers are much more complex than the simple tag format of the C-DOCSIS II Remote MAC-PHY architecture. This is partly because the protocol requires additional layers of headers to support tunneling (DEPI packets must carry the addresses of the node and headend devices, not the endpoint addresses already present in the data packets). It is also partly because in a Remote PHY architecture there is more information to be conveyed on a per-packet basis. As an example, a DEPI MPT packet will include MPEG-TS overhead for the DOCSIS® data to be conveyed, while a DEPI PSP packet, though omitting the MPEG-TS headers, must include packet streaming information.

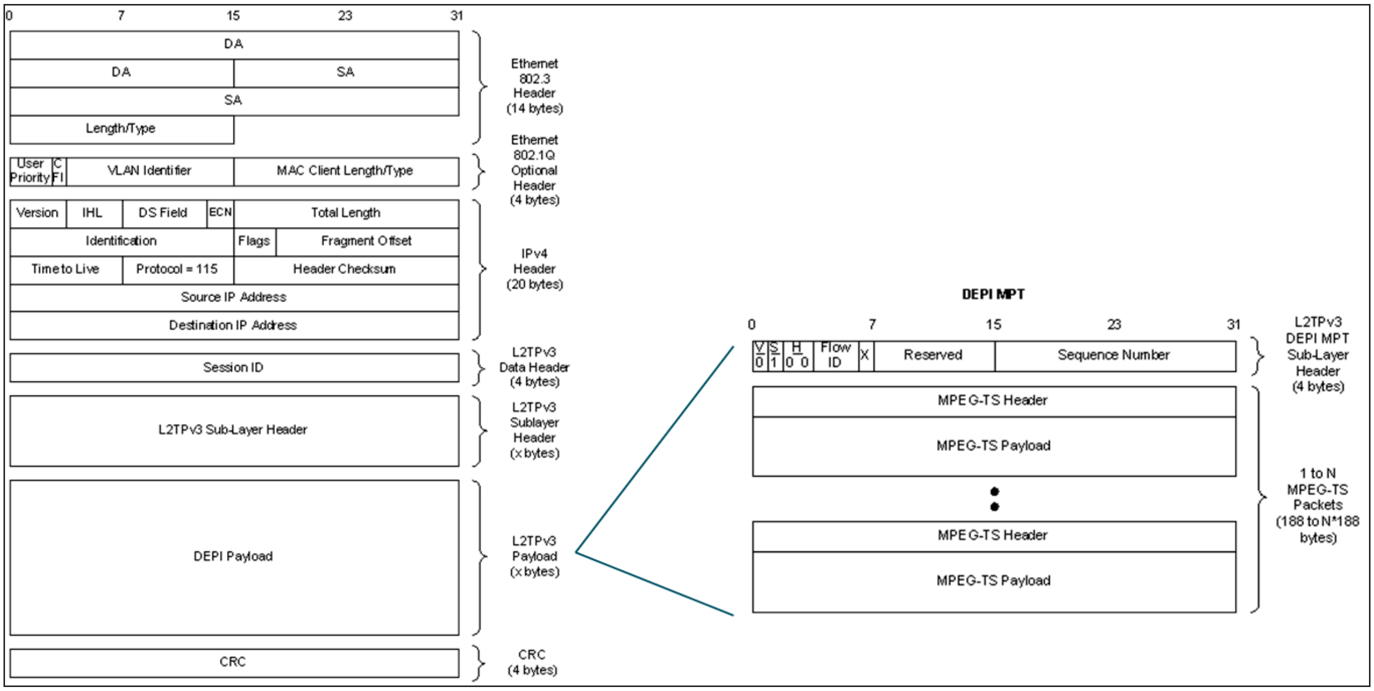

Figure 10 Encapsulation For DEPI Data Plane Protocol (MPT Mode)

In the control plane, DEPI includes messages to set up and tear down control connections, sessions, and flows in order to manage the tunneling protocol. It also contains messages carrying specific types of configuration information and statistics between the two DEPI endpoints. There are fewer such messages and fewer parameters per message in comparison with the C-DOCSIS CDMM protocol, mostly because a Remote PHY node has fewer parameters to configure than a Remote MAC-PHY node.

A number of implementations exist or have existed that exemplify the Remote DAC/ADC class of architectures, though these implementations tend to focus on the upstream path (digital return), leaving the downstream unchanged from the Integrated case. Since none of these implementations have been standardized, no detailed protocol examples are publicly available for study. However, it seems safe to say that these protocols are probably the simplest compared to the other architecture classes. In the data plane, only a minimum of encapsulation is required to package a continuous stream of samples of identical format, and if the link is Ethernet-based, a simple Ethernet/IP/UDP header (or potentially Ethernet only, depending on the network) is all that is necessary to address the packet to the correct destination. In the control plane, this architecture class has the fewest parameters to configure (capture bandwidth, sample rate, sample bit depth, etc.) and the fewest statistics to gather.

Table 6 summarizes the protocol complexity comparison based on the examples studied.

	Data Plane	Control Plane
Simplest		
|	Remote DAC/ADC	Remote DAC/ADC
|	Remote MAC-PHY	Remote PHY
|	Remote PHY	Remote MAC-PHY
\/		
Most Complex		

Table 6 Summary of Comparison of Protocol Complexity

Conclusion

In this paper, three classes of distributed CMTS architectures (Remote DAC/ADC, Remote PHY, and Remote MAC-PHY) have been discussed and compared with each other and with the Integrated architecture in use today. Based on the points of comparison studied, there is no clear winner. The Remote MAC-PHY architecture may offer slightly better Round-Trip Time, but the other architectures will give performance equivalent to the Integrated architecture if the network is carefully planned. The Remote DAC/ADC architecture uses the simplest protocol, but requires a significantly higher-rate network link than either of the other two approaches. All approaches can be expected to offer the same RF performance. All approaches are technically viable, and all ultimately have the potential to serve the industry’s needs in various ways. It is hoped that the comparisons offered here will aid in understanding and evaluating the many possible options for future cable systems.

References

[1] D. Large and J. Farmer, “Digital Return,” in *Broadband Cable Access Networks: The HFC Plant*, 3rd ed. Burlington, MA: Morgan Kaufmann, 2009, sec. 8.5.3, pp. 222-223.

[2] O. Sniezko *et al.*, “Digital Return Technology,” Rev B, Santa Clara, CA: Aurora Networks, 2008.

[3] Cable Television Laboratories, Inc. “DOCSIS® Specifications – Modular Headend Architecture (MHA)” [Online] Available: http://www.cablelabs.com/cablemodem/specifications/mha.html

[4] SARFT, “NGB Broadband Access System – C-DOCSIS” [Online] Available: www.sarft.gov.cn

[5] Cable Television Laboratories, Inc. “Downstream External PHY Interface Specification CM-SP-DEPI-I08-100611,” [Online] Available: http://www.cablelabs.com/specifications/CM-SP-DEPI-I08-100611.pdf

[6] Cable Television Laboratories, Inc. “DOCSIS® - Specifications” [Online] Available: http://www.cablelabs.com/cablemodem/specifications/index.html

[7] “DOCSIS® 3.1 Engineering Symposium,” in *SCTE Cable-Tec Expo.’13*, Atlanta, GA, 2013.

Abbreviations and Acronyms

A-TDMA	Advanced Time Division Multiple Access
ADC	Analog-to-Digital Converter
C-DOCSIS	China DOCSIS
CCAP	Converged Cable Access Platform
CDMM	C-DOCSIS Management Messaging
CDT	C-DOCSIS Tagging
CM	Cable Modem
CMC	Coaxial Media Converter
CMCI	Cable Modem to Customer Premise Equipment Interface
CMTS	Cable Modem Termination System
DAC	Digital-to-Analog Converter
DEPI	Downstream External PHY Interface
DS	Downstream
EQAM	Edge QAM
FEC	Forward Error Correction
FM	Frequency Modulation
GbE	Gigabit Ethernet
GHz	Gigahertz
HFC	Hybrid Fiber-Coaxial
Hz	Hertz
IP	Internet Protocol
IPQAM	IP to QAM
L2TPv3	Layer 2 Tunneling Protocol version 3
MAC	Media Access Controller
MAP	Upstream Bandwidth Allocation Map
Mbits	Megabits
MHA	Modular Headend Architecture
MHz	Megahertz
MPEG-TS	Moving Picture Experts Group Transport Stream
MPT	DEPI MPEG-TS mode
MSO	Multiple System Operator
NSI	Network-Side Interface
OCs	Optical Carrier level x
OFDM	Orthogonal Frequency-Division Multiplexing
OFDMA	Orthogonal Frequency-Division Multiple Access
OLT	Optical Line Terminal
PHY	Physical layer
PSP	Packet Streaming Protocol
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
S-CDMA	Synchronous Code Division Multiple Access
SC-QAM	Single-Carrier QAM

SLA	Service Level Agreement
SONET	Synchronous Optical Networking
SNR	Signal to Noise Ratio
TCP	Transmission Control Protocol
TDM	Time-Division Multiplexing
UDP	User Datagram Protocol
UEPI	Upstream External PHY Interface
US	Upstream
WAN	Wide-Area Network
WDM	Wavelength-Division Multiplexing