EPoC Application & MAC Performance

Edward Boyd
Broadcom Corporation
Kevin A. Noll
Time Warner Cable

Abstract

Ethernet Passive Optical Network (EPON) systems have been successfully deployed worldwide for high-speed access networks. EPON uses the 802.3 Ethernet MAC over optical fiber to provide high-speed IP connectivity to the home or business. In November 2011, the IEEE 802.3 formed a study group [3] to study the feasibility of creating a coax cable physical layer (PHY) for the EPON MAC. With the Ethernet-Protocol-over-Coax (EPoC) PHY, cable operators can deploy high speed IP connectivity using the EPON MAC over optical fiber or coaxial cable. Key criteria for selecting and evaluating a PHY layer will be the application in which it is used and the MAC performance over the system.

The MAC layer performance over a Coax PHY layer will be different than an optical fiber PHY layer. Emerging interactive services and higher speed data links will require shorter delays than today's services over low speed links. In this paper, the bandwidth, buffering requirements, and delay over an EPOC network will be predicted for different deployment scenarios and physical layer technologies for the EPOC PHY. The impact of increasing the round trip delay will be considered in a comparison between EPON and EPOC with expected services requirements.

Introduction

EPoC provides a solution for Cable TV operators to provide fiber performance over a coax network or Hybrid Fiber Coax (HFC) network. By re-using the EPON OLT, EPoC promises common head-end or hub site equipment for both fiber and coax customers. There are many architectural choices for EPoC to connect the OLT to the Coax Network Unit (CNU) in the customer's home.

The IEEE 802.3 working group will define a new physical layer to operate on the coax cable. During this process, decisions will be made to achieve reliable performance, high efficiency, and low delay. This paper will explore a set of service requirements for VoIP and Metro Ethernet Forum (MEF) services operating on a potential EPoC implementation. The coax physical layer will require additional functionality that will add delay and increase the round trip time. The efficiency, buffer requirements, frame delay, and frame delay variation (jitter) will be considered for a range of round trip times to understand the impact to the operator. In a point-to-multipoint network like EPON or EPoC, the shared upstream contains the highest frame delay and frame delay variation. The upstream MAC layer differences with DOCSIS and bandwidth requesting mechanisms will be considered. This paper focuses on upstream traffic performance since it is the most challenging.

EPOC Architecture

There are several possible architectures that EPoC could follow. All are rooted at an

EPON OLT and have Coax Network Units (CNU) at the leaves. The variations exist in the outside plant configuration and the implementation of the electrical interface.

Direct Coaxial Connection

One possibility removes optical fiber from the link and attaches the coaxial cable directly to the OLT system. This approach, pictured in Figure 1, mirrors what is implemented with DOCSIS CMTSes today. In DOCSIS, the electrical interface is a coaxial cable secured to the CMTS (or Edge QAM) chassis with an F-connector. It is easy to imagine that an EPOC implementation would have the same electrical interface and F-connector mechanical attachment.

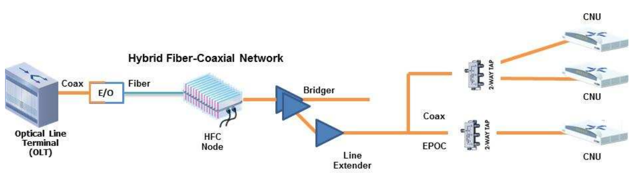

Figure 1: Direct Coaxial Connection

The practical application of this approach suffers from the fact that the bulk of coaxial plant is separated some distance from the hub site and connected via fiber optic cables. This means that the OLT would need to connect to a fiber optic link anyway. The development time and expense to develop a solution of this type is likely to be unproductive.

An alternate approach might carry the RF modulated EPoC signal over analog optics to an HFC node to be converted back to an electrical signal. This approach, however, does not provide the EPoC signal some easily realizable gains in the outside plant characteristics.

Baseband Signaling to Remote CMC

A more preferred architecture is one that uses baseband Ethernet or EPON signaling across the fiber plant. In this scenario, the hub site equipment might be (for example) an Ethernet switch containing WDM baseband optics connected to an OLT that is installed on the strand near an existing HFC node, or even in the HFC node. The OLT in this case could have a direct electrical connect to the coaxial cable and directly implement the EPOC PHY.

This architecture moves in a direction to reduce the use of expensive linear optics in the transmission path to support this type of application. However, the cost and operational complexity of installing an OLT in the outside plant is best avoided in most situations. In addition, for operators that already have EPON OLTs deployed in their hub sites, this approach is not a very effective use of capital.

A similar approach, and the one that is the focus of this analysis, uses the existing OLT and fiber plant to connect to an optical-electrical media converter that is installed in the coaxial plant. A typical configuration is shown in Figure 2.

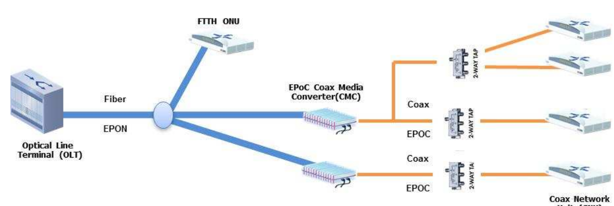

Figure 1: Baseband to Remote CMC

The EPON OLT provides the interface between the PON and external networks (the Internet, for example). It also is responsible for the well-known management functions in an EPON – admission control, station maintenance, scheduling upstream transmission, and other tasks. The role of the OLT in an EPOC network is no different than in an EPON and the CNUs appear to the OLT as if they are ONUs.

From the CMC to the CNU

As mentioned above, this chosen architecture, shown in Figure 2, requires an optical-electrical conversion. The implementation under study refers to this device as the Coaxial Media Converter (CMC). The CMC could be installed in or near an HFC node or somewhere closer to the subscriber. The CMC could be an Ethernet Switch or an Ethernet Repeater. The Ethernet Repeater could be a simpler and lower power device connecting the EPON optical PHY with the EPoC coax PHY. The Ethernet Switch would contain a bridge between an EPON MAC/PHY and an EPoC MAC/PHY.

Operators' Plant Characteristics

A coaxial cable plant, like any other transmission medium, has a set of characteristics that constrain the performance of the communication channel. The typical (but not exhaustive) list of physical-layer metrics includes signal-to-noise ratio or carrier-to-noise ratio, carrier-to-distortion ratios (Composite Triple Beat, Composite Second Order, etc.), carrier-to-interference ratio, group-delay, and micro-reflections. Each of these parameters varies based on operating frequency and bandwidth, so these two parameters must be specified as well.

Fully characterizing a coaxial cable-based network is a nearly intractable problem. Further complicating this is the variation in construction and operating practices from operator to operator and sometimes within a single operator's footprint. This study is focused on the MAC layer performance; therefore this study assumes that physical layer conditions are not a variable and circumstances allow the system to achieve the desired MAC signaling rates.

In addition to the physical-layer, the EPoC system will be expected to adapt to each operator's plant topological design and construction. The primary factors that characterize topology and affect capacity include the distance (which helps define loss characteristics and system timing constraints) from the CMC to the nearest and farthest subscriber, number of active subscribers, and offered subscription tiers (speeds).

Number of CNUs

The number of CNUs to be supported on the EPoC network needs to closely align with the number of active users on an HFC node today. This will help the operator avoid the cost of plant modifications required to deploy the EPoC system.

Today's HFC node-branch typically serves as few as 50 subscribers and as many as 500 subscribers (there are certainly cases where the node serves more or less than this). Based on this it is safe to require that the EPoC system support a similar range.

For the purpose of comparing EPOC performance to EPON performance, we should choose 32 CNUs. For the purpose of analyzing EPOC under conditions similar to today's average density this study will analyze network populations up to 512 CNUs and activity on up to 256 CNUs.

Distances

The propagation delay, that time required to transmit a frame across the coaxial cable plant, can have significant impact on the scheduler in the EPoC implementation. Therefore the distances spanned by fiber and coaxial cable in the network are an important parameter in the MAC performance.

Given the topology chosen for analysis – baseband EPON to a CMC located in or near an HFC node – we must consider two contributors to the distance from OLT to CNU. The first is the fiber from the OLT to the CMC. This distance can range from 0 meters (when the node is located in the hub site) to a typical maximum of 30km.

The second contributor to distance is the coaxial link from the CMC to the CNU. In an N+0 configuration the coaxial distance can range to around 150 meters. In an N+5 configuration with 1000-foot spacing the coaxial plant contributes about 1.7km to the total distance.

Subscription Tiers

Another factor in the system's ability to deliver traffic in a timely fashion is the speed tiers offered to subscribers. The typical Internet access service is a best effort service and ranges widely in offered data rates. A sampling of current offerings across the industry shows offered tiers 3x1Mbps (downstream x upstream bandwidth), 50x5Mbps, 60x6Mbps and as high as 100x10Mbps.

Operator Service Requirements

Operators offer many different services over their networks. Services include Internet access, Voice (VoIP), Video, Cellular Backhaul, Enterprise-class Ethernet circuits and more. Each service has its own set of network service level objectives.

Conveniently, there are two sets of specifications that can be referenced to cover the majority of these services and use cases. These specifications are the PacketCable specifications published by CableLabs and the MEF23 Implementation Agreement published by the Metro Ethernet Forum.

Packet Cable VOIP

MSOs provide packet cable VoIP service to residential and business subscribers. These are often a single line per home but multiple lines are possible, especially for business services customers.

Performance requirements for an access network supporting voice services are widely understood. Requirements specifications include packet loss, latency, and jitter. The major source of jitter in the EPON/EPOC network is scheduling the upstream transmission.

There are several sources of delay in the EPON/EPOC network. These include DSP processing and encryption, packetization, upstream transmission, and forwarding at the OLT.

Impairment	Value
Packetization Delay	20ms
Forwarding and Transmission Delay	< 10ms
Jitter	< 10ms

Table 1: VoIP Requirements

Table 1 summarizes these impairments and gives some typical tolerances in use by various service providers. In this analysis, we will assume that packet loss is trivial.

MEF 23H

The Metro Ethernet Forum defines a set of performance metrics that specify High, Medium and Low parameters that set the expectations for Ethernet services that traverse Metro, Regional, Continental, and Global distances (Performance Tiers). The general description of each performance tier (PT) is given in Table 2. In the context of this study, only the Metro PT is interesting and the EPON/EPoC network segment will generally only be a small portion of any one Ethernet service. The expected contribution of the EPON/EPOC link to the performance budget is expected to be small.

Performance Tier	Distance
PT1 (Metro)	< 250 km
PT2 (Regional)	< 1200 km
PT3 (Continental)	< 7000 km
PT4 (Global)	< 27500 km

Table 2: MEF Performance Tiers

Each PT definition includes a maximum frame delay (FD), mean frame delay (MFD) and a maximum inter-frame delay variation (IFDV).

The MEF 23 high quality service definition (H) is intended to carry delay sensitive traffic such as VoIP and financial trading transactions. These performance metrics for point-to-point delivery are summarized in **Error! Reference source not found.**.

Metric	Value
FD	≤10ms
MFD	≤7ms
IFDV	≤3ms

Table 3: MEF 23H Parameters [2]

<u>MEF 23M</u>

The MEF 23 medium quality service definition (M) is intended to carry traffic like Fax and network control traffic which are delay-sensitive but non-interactive. These performance metrics for point-to-point delivery are summarized in **Error! Reference source not found.**.

Metric	Value
FD	≤20ms
MFD	≤13ms
IFDV	≤8ms

Table 4: MEF 23M Parameters [2]

<u>MEF 23L</u>

The MEF 23 low quality service definition (L) is intended to carry Internet data service for business or residential where delay and jitter are not of any significant concern. These performance metrics for point-to-point delivery are summarized in Table 3.

Metric	Value
FD	≤37ms
MFD	≤28ms
IFDV	Unspecified

Table 3: MEF 23L Parameters [2]

<u>EPoC System for Analysis</u>

<u>EPoC Sources of Delay</u>

EPON Delays

In 1Gbps EPON, a round trip time of 250µs includes the propagation delay and physical layer delay for 20Km of fiber. The fiber propagation delay is about 100µs in each direction and 50µs covers the physical layer and synchronization delays in the OLT and ONU. For the analysis in this paper, the EPON round trip time of 250µs will be used as a baseline for comparison. EPoC bandwidth overhead (same FEC, 64/66) will be used on all RTT values so the difference is limited to the round trip delay.

EPoC Architecture

The MSO network has cable distances longer than the traditional TELCO network. While 20km may cover the entire network in EPON, EPoC will likely need to cover 30 km spans. The extended distance could add another 100µs of propagation delay.

EPoC PHY Functions

The EPoC PHY will require additional functionality to provide reliable performance when faced with burst or narrow band interference. A forward error correction (FEC) and interleaver will be selected to handle 25µs or more of burst error. The interleaver and FEC could add 400µs to 800µs delay to the round trip time.

Long symbol times of 20µs or 100µs will help combat multipath reflections. To gain better granularity, a block of symbols will be transmitted in selected carriers. Depending on the symbol and block size, an additional 400µs could easily be added.

Sharing Upstream & Downstream Frequency

Some operators like the option of using the same frequencies in the upstream and downstream in a Time Division Duplex (TDD) mode. While EPON is a full duplex protocol, half duplex operation to support TDD might be achieved by alternating between upstream and downstream transmissions in a fixed time block. To get reasonable efficiency on the upstream and downstream, a large block of transmission from each direction is needed. The larger block would be more efficient but it would add a significant amount of delay to the upstream and downstream. For example, an EPoC system that gave 1 millisecond of slot time to the upstream and 1 millisecond of slot time to the downstream would add 2 milliseconds of delay to the round trip time. The split between upstream and downstream maybe 50/50 or it might give a larger percentage to the downstream. In either case, the round trip time delay is the sum of the upstream block size and downstream block size. Small upstream block sizes would provide an additional restriction on the per-CNU upstream burst size. This paper will only consider the effect of the round trip time. An EPoC system using TDD would likely add 2 to 4 milliseconds of round trip time.

Switched or Repeated

The EPoC CMC provides a link between the optical fiber to an EPON OLT and the coax cable link to a CNU. The EPoC CMC could be defined as a switch or as a repeater.

An EPoC CMC Switch would contain an EPON ONU MAC layer connected to an EPoC OLT MAC layer through an 802.1D Ethernet Bridge. In this case, the access plant has two networks. The CMC will schedule and aggregate data from the CNUs and the OLT will schedule and aggregate data from the CMCs. The two layers of scheduling and aggregation allow for a more efficient use of the fiber. To determine the service delays, the fiber network frame delay and the coax network frame delay would be added together.

In an EPoC CMC Repeater, the EPON PHY and the EPoC PHY will be connected together in a fixed delay repeater. A single layer of scheduling and aggregation from the OLT handles upstream traffic. This system allows for a much simpler device but doesn't provide the second level of aggregation so it will not get full utilization of the fiber network when multiple CMCs share an OLT port. In networks with large Coax plants, the fiber to the OLT would likely be point-to-point so there is no needed for aggregation on the fiber. When there are very few CNUs connected to each CMC coax segment, data from the CNUs could be aggregated to the fiber as if the CNUs are on the same coax plant. For example, four CMCs with 10 CNUs each could share an OLT port as a single 40 CNU network. The EPoC CMC Repeater does not require QoS buffers, classification, SLAs, or scheduling in the CMC.

For round trip delay analysis, only the CMC Repeater is considered in this paper. The CMC Switch performance can be determined by assuming 300µs less round trip delay on the CMC Repeater RTT time and adding a second system with the EPON delay of 250µs. For example, the FD results for a CMC switch could be determined from the CMC repeater results by the following equation.

FD-Switch(RTT) = FD-Repeater(RTT-250us) + FD-repeater(250us)

The IFDV would follow the same equation since the delay frame variation from the coax scheduling would be added to the fiber network. The total delay budget for the access plant must be shared between the coax aggregation and fiber aggregation to guarantee compliance. In all cases, the CMC Switch will add delay to the access plant because of the two stages.

Delay Summary

The EPoC system could have delays from 1ms to 6ms based on decisions made in the standard and architecture deployed by the operator. In the performance analysis, a selected set of round trip times will be used to analyze the performance impacts. In most cases, the delay would be different for upstream and downstream. To simplify the analysis, the round trip time will be divided evenly between upstream and downstream.

<u>EPoC MAC Layer Performance</u>

<u>EPoC MAC Layer Differences</u>

Packet Fragmentation

Like other Ethernet MAC solutions, EPoC does not support layer 2 fragmentation of packets in multiple flows [1]. Fragmentation in ATM and other networking technologies allow for improved Quality of Service on low speed links along with a large unit of granularity. EPoC will need to support variable packet sizes and burst sizes with a finer granularity. On higher speed links like EPoC, the value of fragmentation and reassembly is questionable for the additional complexity. Since QoS is measured by frame delay variation and maximum frame delay, QoS on cells (fragments of packets) is misleading for packet analysis. The scheduling of cells can increase the worst-case delay and frame delay variation since a packet could span multiple upstream bursts. Even though fragmentation is not supported in EPON and EPOC, this paper will consider the impact of fragmentation on the performance when appropriate.

Stateless REPORT Frame

EPON and EPOC use a REPORT frame to pass queue information from the subscriber side CNU to the operator side OLT. The REPORT frame is not a request for bandwidth. It identifies the depth of the queues at the time of generation [1]. REPORT frame values will only change when data moves in and out of the queue. It is the responsibility of the OLT to track what has been granted in the past. This method is commonly referred to as stateless bandwidth reporting since the CNU doesn't hold state on the status of a bandwidth request. The CNU reports the queue size at the present time without regard to previous report frames.

DOCSIS systems use a stateful bandwidth request. The CM will generate a request for an upstream slot and it will not request for the same packets unless there is a timeout. The CMTS must grant the request or acknowledge it so the CM can update state on the request. A second request will not include the request in progress from an earlier request. The CM and CMTS must track the state of the request for the stateful system.

Stateful bandwidth requests were required for DOCSIS to support multicast bandwidth request slots. The multicast slots would only be used by a cable modem that hadn't already requested a bandwidth request. Stateful bandwidth requests are required for this function. EPON does not support multicast slots since the user count is lower and upstream bandwidth is higher. Performing a worst-case delay analysis is greatly simplified without multicast bandwidth request slots.

The stateless queue reporting of EPoC provides a simplification for a higher bandwidth upstream. It allows the CNU to avoid timers and long timeouts from a lost upstream request frame, downstream bandwidth acknowledge frame, or grant frame. In a stateless system, the polling interval determines the delay penalty for a lost upstream REPORT or gate frame. A timeout is considered in the delay penalty.

The REPORT frame provides a solution for reporting to frame boundaries. Since Ethernet doesn't support fragmentation, grants that aren't at frame boundaries will significantly decrease the efficiency. The REPORT frame contains one or multiple queue sets to define a queue's frame boundary at different thresholds. The queue set allows for the OLT to know a frame boundary at maximum size. A REPORT value for every frame in the queue would make a very large REPORT frame. The number of queue sets and maximum number of bytes can be configured with the SLA. For the analysis in this paper, a 4 queue set REPORT frame will be used. All 4 queue sets will have an equal limit. For example, queue set 1 will REPORT up to 4K bytes and queue set 2 will REPORT up to an additional 4K bytes. With a 4 queue set REPORT frame, the OLT can give 4 grants from a single REPORT frame before receiving the next REPORT frame. A smaller queue set will allow for smaller bursts and shorter delays for the upstream. Larger upstream queue sets will result in more efficient upstream bursts but longer delays.

Contention Slots

The EPoN MAC and EPoC system won't support contention or multicast slots. The lone exception to this rule is the discovery slot where multiple CNUs may respond. After discovery, all grants to an ONU or CNU will be unicast. Only one CNU or ONU will transmit in the slot. While the contention slots are useful in a large user network with many CMs, contention slots will prevent a smaller user network to reach high upstream data performance.

Since contention slots are not used in the EPoC based system, the worst-case delay is easier to determine and guarantee. It is also easier to show stable performance at close to or reaching 100% capacity.

The loss of contention bandwidth request slots also impacts the requirements for SLAs on the subscriber side. In DOCSIS, a cable modem will have an SLA to prevent it from over requesting bandwidth from the CMTS. The stateless REPORT frame of EPoC will only be sent by a CNU when requested by the OLT. The OLT has complete control over the

CNU for bandwidth granting and reporting so there is no need for an SLA on the CNU.

Piggybacking

REPORT frames can be sent in a single frame burst or as a frame in a longer burst with many frames. Since the REPORT frame contains the status of the upstream at the time of generation, it is normally sent as the last frame of the burst to exclude the frames in the burst. The OLT uses the force report indicator in the GATE frame to request a report frame in the burst. While a CNU could decide to send a REPORT frame in any burst, the normal practice is to send a REPORT frame only when requested by the OLT. The GATE frame with the force REPORT bit set is commonly referred to as piggybacking while the burst with only a REPORT frame in it is commonly referred to as a polling grant.

GATEs and MAPs

A MAP in DOCSIS provides a time slot description of the upstream with information for all stations. In EPoC, the GATE frame provides a unicast message to the CNU with a start time and length. In some cases, the MAP frame contains many grants over a significant portion of time. In the case of EPoC, the GATE frame will only contain a single grant to a single CNU. The GATE frame allows for up to four grants to the same CNU. In practice and in this analysis, a GATE will only contain a single grant. A MAP block delay or generation time does not exist for this reason.

Multiple LLIDs and Service Flows

A Cable Operator who provides multiple billed services to a single subscriber uses service flows to allow for different service level agreements. In EPON, the logical link identifier (LLID) provides a virtual point-to-point MAC connection between the OLT and CNU. A CNU with multiple LLIDs acts with multiple EPON MACs. With a MAC for each service, the OLT can monitor, enable, or grant the service independently of the other services on the CNU. By using multiple LLIDs, a cable operator can have multiple service flow like DOCSIS.

Activity based Polling

The large number of LLIDs or service flows on an OLT port will require a significant amount of bandwidth to query for status. Since service flows are often inactive for large residential systems, activity based polling can save bandwidth. Any service flow can have an active and inactive polling rate. The active polling rate would be much higher than the inactive rate. A simple example is a VoIP call where the active rate is used when a call is active and the inactive rate is used when no call is active. Activity can be determined by looking at the presence, rate, or type of frames on a link. The OLT system can determine the rules for activity and inactivity.

<u>EPoC PHY Parameters for Analysis</u>

The IEEE 802.3 will define overheads for the physical layer. Commonly suggested options for FEC and encoding burst overhead will be selected to get an estimate of the overhead. The Ethernet Frames will use the 64/66 encoding of 10G EPON and an 85% efficient LDPC FEC code. With these constant overheads, a fixed 20% overhead would be needed. A 1Gbps Ethernet MAC rate would require 1.2Gbps of Ethernet Line rate.

For bursts, a shortened FEC code word is allowed for end of bursts. A common burst overhead for EPON is 32 time quanta (time quanta are 16ns long) for sync time, 64 time quanta (TQ) for laser ON, and 64 TQ for laser OFF. At 1Gbps, the total burst overhead will be 1536 bits or 192 Bytes. EPoC will use the same burst overhead as EPON so the analysis can focus on performance differences due to round trip time. A larger EPoC burst overhead would reduce the performance and it should be considered in future analysis.

Packet Cable VoIP

Packet Cable VoIP service can be mapped to EPoC in a variety of ways. The most obvious is an unsolicited grant similar to DOCSIS. Another solution is a solicited granting based on polling.

For the analysis below, the G.711 codec will be assumed. Based on this code, a 218-byte packet will be generated every 20ms for each subscriber with an active voice call. A maximum FD and IFDV of 10 milliseconds will be required.

Unsolicited Grant Synchronization (UGS) Performance

In the UGS solution, the EPoC system will establish two LLIDs. One LLID will carry signaling while the other LLID will carry the encoded voice. The encoded voice LLID will use unsolicited granting. Unsolicited granting is based on a timer at the OLT. A fixed size grant is given in a fixed time period. A REPORT frame with a non-zero queue set is not required for the grant generation. The signaling LLID will use solicited granting. Using activity based polling, the LLID will be polled at 17ms when the LLID is active and 100ms when inactive. The unsolicited granting could be enabled or disabled in the OLT based on the state of the voice call from observing the signaling channel.

The UGS slot will be sized large enough to carry a single 218 Byte Ethernet frame. The granting period of the UGS must guarantee a maximum delay of less than 10ms. The UGS slot is not aligned with the arrival time of the packet so the worst case scenario is an upstream frame just after the slot passed. The worst case delay of packet upstream will be the upstream transport delay plus the period of the UGS slot. The downstream delay does not factor into the UGS performance since the GATE is autonomously generated by the OLT. It is assumed that the worst case slot jitter from discovery slots is less than 500µs.

The period of UGS slots to a CNU must decrease with increased upstream delay. The period can be determined by subtracting the fixed delays from the worst case delay of 10ms. The equation below can be used to find the UGS period. As the UGS period decreases, the amount of upstream bandwidth consumed increases.

UGS-Period = MaxDelay – RTT/2 – SlotJitter

In the example scenario, each CNU will have a single VOIP session. The system is assumed to have 512 CNUs. The amount of Ethernet Line bandwidth required is shown for a different numbers of active voice calls and for different round trip times.

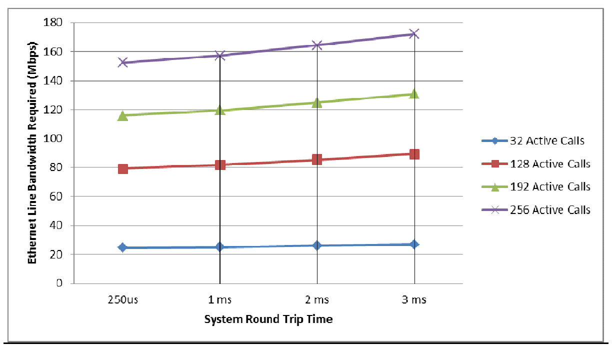

Graph 1: Required UGS Bandwidth

Graph 1 shows the bandwidth required for different round trip times and numbers of active voice calls. The System Round Trip Time of 250µs represents the performance of the all fiber 20Km EPON solution. The 1ms, 2ms, and 3ms show the performance of an

EPoC system with the corresponding total round trip time.

For UGS, the increase in bandwidth required due to longer round trip times is not significant for a small number of active calls. The increased RTT is more significant with 256 active callers.

The UGS efficiency is hurt by the single packet bursts. Additionally, the 20ms arrival time and sub-10ms delay will cause over half of the upstream slots to be empty.

Fragmentation or No Fragmentation

The UGS analysis assumes that EPoC does not allow fragmentation. Would fragmentation improve the capacity or decrease the delay? If packets were fragmented, they would need to wait for an additional UGS slot to be transported upstream. If the packet boundary and slots were miss-aligned, it would take up to 2 UGS slots for a frame to go upstream. In this case, the UGS slot would need to occur twice as often. The payload in the UGS slot could be divided in half. Since the overhead would double for the shorter interval, fragmentation would significantly increase the bandwidth required to transport the UGS flows.

Solicited Granting Performance

The UGS solution provides an adequate solution for transporting VoIP over EPON and EPoC. The UGS has the complexity of detecting the start and end of phone calls. UGS also requires a known packet interval and packet size. Additional phone lines at a CNU require more UGS flows or the complexity of detecting multiple phone calls in a single service flow. UGS is also not easily compatible with compressed voice or video conferencing. Solicited granting would greatly simplify the control and allow for other service options. Solicited is preferred if the performance is similar to UGS.

Solicited granting requires a REPORT frame to transmit upstream, a GATE frame downstream, and a data burst to be received upstream. The transport delay is therefore the downstream delay plus two times the upstream delay. Since data comes in asynchronous to the scheduler, the worst case delay should include the delay from simultaneous upstream slot requests from all active VoIP flows. For this analysis, we assume that VoIP is the highest priority. The polling period is the key parameter for the solicited solution. The following equation can be used to calculate the worst case delay.

Tmax_delay = Tpolling + 2xTup + Tdown + Tall_service

The following equation solves for the polling interval.

Tpolling = Tmax_delay – 2xTup – Tdown – Tall_service

In the case of VoIP, piggybacking will not be used. While piggybacking would decrease the latency for arriving packets, it would not decrease the worst-case latency. In the case of the VoIP example, the packet spacing is larger than the maximum delay so a piggybacking would be useless to detect the next frame.

Graph 2 shows the bandwidth capacity required by the solicited VoIP solution.

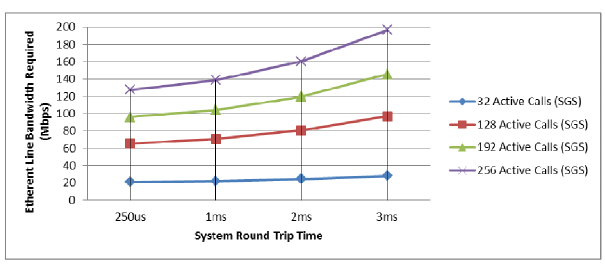

Graph 2: Required Solicited Bandwidth

The UGS bandwidth increase due to increased round trip time was much less than the solicited solution because of extra round trip in the delay equation. At a small number of active calls, the UGS shows little or no difference with a lower or higher round trip time.

The solicited solution is more efficient for the shorter round trip times. The solicited solution benefits from only granting data slots when a frame is present. As the RTT increases, the increasing polling rate to meet the maximum delay consumes more bandwidth than the wasted slot in the UGS solution.

When comparing the 250µs EPON data point, there is less than 10% increase in bandwidth to achieve the same delay performance if the round trip time is in the 1.5ms range. A 3ms round trip adds a 50% bandwidth penalty to achieve the same delays. It is clear that RTT delays beyond 3ms are unusable in the solicited.

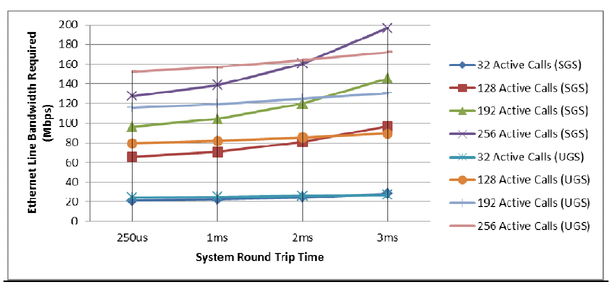

Graph 3: UGS & Solicited Bandwidth

If the UGS and Solicited graphs are overlaid, it shows a cross over point between UGS and solicited around 2ms of round-trip time. A solicited solution is equal performance for a small number of users and it is better performance if the round trip time is less than 2ms. Since the solicited solution is more flexible for video or compressed content and simplifies controls, a lower round trip time that allows for efficient use of soliciting granting is preferred. For DOCSIS systems with many users and long delays, UGS must be used. For EPON systems with fewer users and shorter delays, solicited granting is clearly preferred.

Performance for MEF 23H

Requirement Overview

The MEF 23H service agreement is an example of a higher tier business or residential SLA. For the MEF 23H service, an IFDV of 3 milliseconds and a maximum FD of 8 milliseconds will be used as requirements. For the analysis, a 10Mbps-streaming load will be applied in the upstream direction. The 10Mbps load has a random packet size from 64 bytes to 1518 bytes.

Configuration to reach goals

With an unconstrained packet size, only solicited operation can be used since a UGS would require knowing the packet boundaries. In a system without fragmentation, the unknown packet boundary would be very inefficient. In a system with fragmentation, a packet spanning 2 grant slots would double the delay. In either case, UGS is not the preferred method.

Since the period of polling must be short to meet the IFDV requirement, there is no need to use piggybacking. While piggybacking may lower the average delay in some scenarios, it will not decrease the worst case IFDV or FD. Piggybacking would decrease the efficiency because of the additional REPORT frame in the burst.
The IFDV is the critical constraint in this system. The IFDV in the upstream will be sum of the polling interval and the scheduler delay. A packet arriving just before the

polling slot will have zero delay while a packet arriving just after the polling slot will wait the entire polling interval. The scheduler delay can be zero when only one CNU requests an upstream slot for shortest delay. The longest scheduler delay occurs when all CNUs need a slot at the same time. The maximum scheduler delay is number of CNUs times the maximum slot size.

The best efficiency can be found when the IFDV is equally split between polling and scheduler contention delay. For a 3ms IFDV, the scheduler delay of 1.5ms and a polling delay of 1.5ms are allowed.

Performance Analysis

The maximum delay is defined by the same equation as the VOIP solicited grant example. It should be noted that this equation is the same as the IFDV plus the twice the upstream delay and downstream delay. The delay graph shows the relationship between the round trip time and the FD and IFDV. The bandwidth graph shows the best efficiency is found when the IFDV value is largest. A large IFDV allows for a lower polling rate and larger upstream data bursts which results in higher throughput.

At the EPON round trip time of 250µs, the maximum delay is far below the 8ms maximum. As constant delay is added for the round trip time increases, the IFDV and the efficiency remains the same. When the additional RTT delay causes the maximum delay to be exceeded, the polling period and burst size must be decreased. These decreases cause the bandwidth required to increase dramatically.

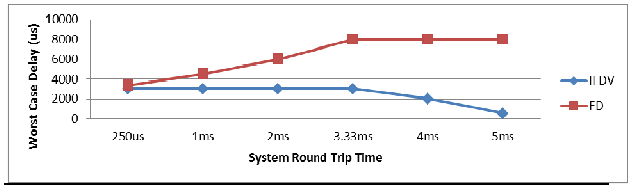

Graph 4: MEF 23H Delay

While the maximum delay increases for the RTT of 250µs to 3.33ms, delay is under the 8ms maximum and the efficiency is constant. If EPoC RTT delay is under 3.33ms, the MEF23H service can be supported without any additional bandwidth. Above 3.33ms, the penalty increases dramatically until the absolute limit of 5ms where the minimum polling period of 250µs is reached. At the 5ms limit, the bandwidth required to meet MEF 23H is more than double EPON at 250us.

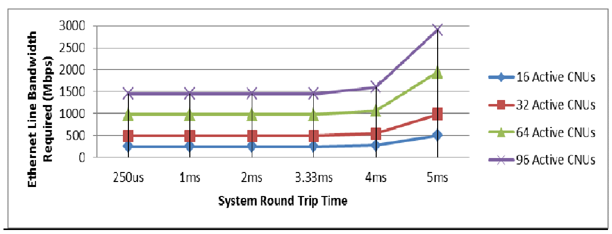

Graph 5: MEF 23H Bandwidth

CNU Buffering Requirements

The additional delay will impact the buffering requirements for a CNU in the upstream direction. For a MEF 23H service, it is assumed that it is a guaranteed bandwidth without best effort data. The MEF 23M and MEF 23L will consider best effort data. With only guaranteed bandwidth to consider, the buffering required can be found by the multiplying the guaranteed rate by the frame delay (FD). Since the buffer is normally store-and-forward, 2000 bytes (the largest 802.3 packet size) is added.

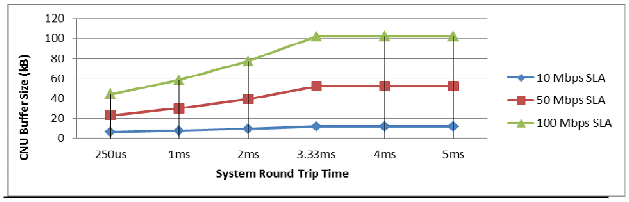

Graph 6: MEF 23H Buffering

The results for MEF 23H buffer size required versus RTT has the same shape as the delay graph and the opposite shape of bandwidth graph. The increase in RTT increases the buffer size until the maximum delay is reached. After the maximum delay, additional RTT increases don't change the buffer size but bandwidth for higher polling rate climbs. Graph 6 shows that while the increase in the EPON fiber RTT delay from 250µs to 3.33ms does not hurt the efficiency, it more than doubles the upstream buffering requirements in the CNU.

<u>Performance for MEF 23M</u>

Requirement Overview

For MEF 23M, a worst case IFDV of 8 milliseconds and FD of 20 milliseconds will be used. For the analysis, a 10 Mbps and 50 Mbps streaming load will be applied in the upstream direction. The load has a random packet size from 64 bytes to 1518 bytes.

Configuration to reach goals

The MEF 23M traffic patterns would not normally be a traffic pattern compatible with a UGS flow. Variable sized bursts of unknown packet sizes are best handled by solicited granting.

The longer FD and IFDV limit allow the use of piggybacking for better efficiency than the polling only solution used in MEF 23H. The polling timer will be reset by the generation of a polling burst or a piggybacked REPORT frame. For a bursting station with a polling period greater than or equal to the scheduler contention delay, no polling bursts will be requested.

There are 2 equations to determine the maximum delay. The first equation is based on a station that has been active but not bursting. Polling will detect the packet in this case. This scenario will be referred to as "burst detection".

Tmax_delay = Tpolling + 2xTup + Tdown + Tall_service

The second equation is based on a CNU that is bursting and not reporting a zero length queue. In this case, the piggybacking will detect the packet arrival. This scenario assumes that the flight delay of 2xTup + Tdown is less than the time to service all stations. The scenario will be referred as "burst continuation".

Tmax_delay = Tup + 2xTall_service

The worst case delay can be determined by taking the longer delay from the burst detection or burst continuation scenarios. For optimum performance, the polling interval should never be less than Tall_service and for best performance, they should be set equal. In this case, the burst continuation equation is not the worst case so the burst detection equation will be used for analysis. The Tpolling interval will be half the result of the maximum delay minus 2xTup + Tdown.

For a system mixed with higher priority services like MEF 23H, the Tall_service should include their burst interruptions. Tall_service should be the sum of all higher and same priority upstream slots. Since the MEF 23H IFDV is 3ms, the MEF 23M will assume no more than a 3ms disruption.

Performance Analysis

Graph 7 shows that as the round trip time increases no increase in the bandwidth required. Since the FD of 20ms is larger than the IFDV of 8ms plus 5ms RTT, there isn't a need to increase the granting rate. The bandwidth increase wouldn't occur in the MEF 23M until a RTT of around 12ms.

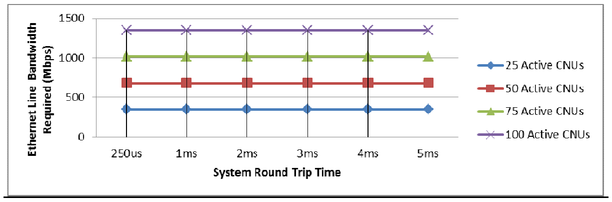

Graph 7: MEF 23M Bandwidth (10Mbps)

Graph 8 shows fewer active CNUs and therefore fewer bursts at a 50Mbps rate each. The charts show that the penalty for extended RTT is larger when there are more users and lower data rates. For a system with many users and higher data rates, the RTT doesn't have significant impact up to 5ms.

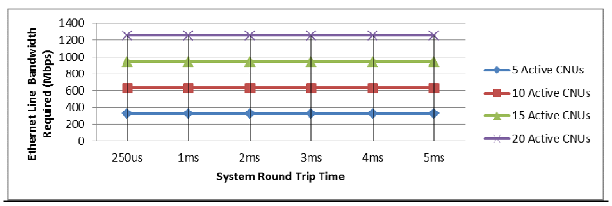

Graph 8: MEF 23M Bandwidth (50Mbps)

Buffering Requirements

While the efficiency of the system for MEF23M is constant from with the increased RTT, the buffering requirements on the CNU are not. The buffer required on a CNU to support the MEF 23M with data rates up to 200 Mbps would need to be more than double the EPON ONU. Graph 10 shows the buffering requirements to support average rates of 50, 100, and 200 Mbps.

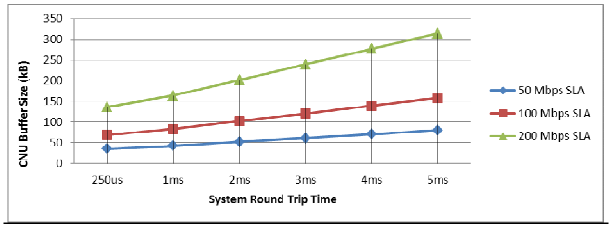

Graph 10: MEF 23M Buffering

<u>MEF 23L</u>

Requirement Overview

The MEF 23L service agreement is an example of a best effort SLA. The MEF 23L specification contains a maximum FD requirement of 37 milliseconds. For the analysis, different data rate streaming load will be applied in the upstream direction. The load has a random packet size from 64 bytes to 1518 bytes.

Configuration to reach goals

For the same reasons as MEF 23M, a solicited granting with piggybacking will be used. A 37ms delay limit is very long for an EPON or EPOC system that is not oversubscribed. The RTT will be a small percentage of 37ms delay limit so it will not have a significant impact on the efficiency like the MEF 23M. The RTT will have a significant impact on the buffering requirements for a CNU to reach high bandwidth. In general, the MEF 23L needs to achieve high efficiency at a high data rate without requiring a large amount of upstream buffering.

The polling rate could be set for MEF 23L to 10ms and meet the FD requirement of 37ms. For the MEF 23L, different polling rates will be considered to balance efficiency with buffering requirements.

The burst detection condition will be considered for the same reason as MEF 23M.

The Tall_service delay is more difficult to determine at this priority level since many higher priority services could be active. The disruption from MEF 23H and MEF 23M services will be limited by the MEF23M IFDV of 8ms. Of course, this analysis assumes a round robin scheduler with guaranteed slots for lower priorities and shaping that streams the higher priority. Without these restrictions to the high priority, the delays to MEF 23L could be unbounded. Tpolling will be equal to Tall_service.

Tmax_delay = Tpolling + 2xTup + Tdown + Tall_service

The polling interval for the MEF 23L service can be determined subtracting the loop time and dividing by 2.
Tpolling = (Tmax_delay - 2xTup – Tdown)/2

To handle the disruption from high priority services, the MEF 23L polling rate shouldn't be set less than the 8ms IFDV of MEF 23M. For a 5ms delay, the maximum polling rate is just under 15ms. The analysis will look at polling rates from 8ms to 15ms.

Performance Analysis

The MEF 23L buffer size is calculated for the different polling rates versus RTT. In Graph 11, the buffer size is considered for a sustained rate of 500 Mbps with 50% of the bandwidth taken by MEF 23M services. The buffer requirements are the maximum delay times the sustained input bandwidth.

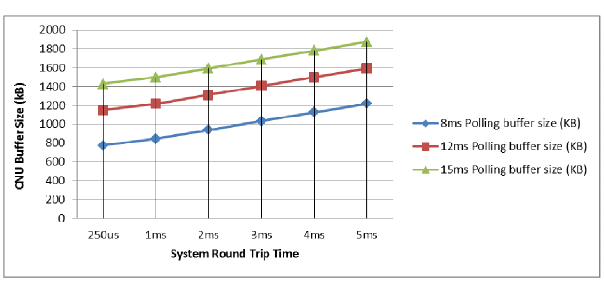

Graph 11: MEF 23L Buffer Size (500 Mbps)

Graph 12 shows the buffer size requirements for an empty system where a single MEF 23L CNU is bursting at 1 Gbps (100%) with no contention delay. Comparing Graph 11 and 12, it is clear that the worst case buffer requirement for MEF 23L is a single user with an SLA to reach maximum bandwidth. The buffer requirement decreases with more users sharing the upstream as the maximum data rate decreases.

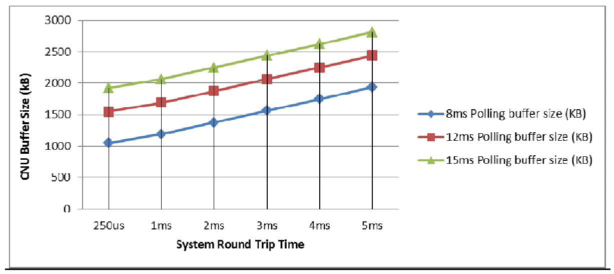

Graph 12: MEF 23L Buffer Size (100% load)

Graph 12 shows that EPoC will require a significant amount of additional buffering (~1 MB) over EPON as the RTT time is increased. From Graph 12, the buffering requirements for EPON and the 5ms RTT are equivalent if the EPON system uses 15ms polling and the EPoC system uses 8ms polling. Since the buffering is a directly related to the delay, the EPON and EPoC would have the same delay as well. For a system with few CNUs, the penalty to compensate for RTT delay with polling will be small but a larger system will require significantly more bandwidth. Graph 13 shows the impact of increasing the polling rate to match the delay and buffer requirements of EPON. In the example for Graph 13, a fixed buffer size of 1.5 MB is used. The 1.5MB buffer represents the 12ms polling, 250μs RTT, and 25ms delay on Graph 12. Graph 13 assumes that the system will carry 1 Gbps of Ethernet traffic split evenly across the stations.

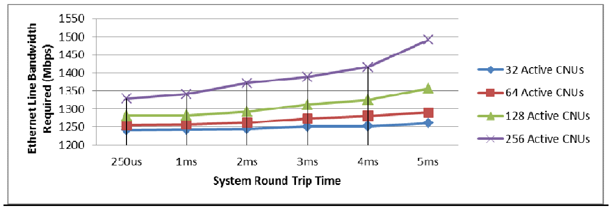

Graph 13: MEF 23L Bandwidth (100% load)

Graph 13 shows that the penalty for increased RTT multiples by the number of users. The system with 256 active users will have around a 15% penalty on bandwidth to match to match the EPON fiber based performance.

Conclusions

An EPoC PHY can be used by a cable operator in multiple network configurations. The EPoC PHY could be placed with an OLT in headend and operate over a traditional HFC network or the EPoC PHY could be placed in a CMC at a remote node and act as a switch or a repeater. The choice of architecture is dependent upon the individual operator's needs and plant design.

EPoC can provide a significant performance improvement over existing cable systems because of small service groups, a fast Ethernet MAC, and a single wide logical pipe. EPoC can provide VoIP, MEF 23H, MEF 23M, and MEF 23L services if the round trip time is low enough. RTT increases will impact the CNU cost dramatically if it requires an EPoC specific chip with more buffering than the standard EPON ONU. Increased polling rates can compensate for larger round trip times to certain limits and still meet MEF 23 requirements but bandwidth efficiency will be reduced. Going over 2ms, forces EPoC from a solicited VoIP into the less flexible UGS VoIP. At RTT's of 3.33ms and 5ms, some MEF 23 services become impossible. Solutions with a shorter round trip time will be more efficient and perform closer to the fiber solutions without additional hardware or bandwidth costs.

The IEEE 802.3 standard should seriously consider the round trip time impacts in selecting the solution. A solution that increases the bandwidth efficiency at the PHY layer by adding significantly delay could hurt the overall system efficiency.

References

[1] Glen Kramer, Ethernet Passive Optical Networks, McGraw-Hill, 2005
[2] MEF Technical Specification MEF 23.1, "Carrier Ethernet Class of Service - Phase 2"
[3] EPON Protocol over a Coax (EPoC) PHY Study Group. http://www.ieee802.org/3/epoc/index.html