Technical Brief

May 2018

DOCSIS® Technologies for Mobile Backhaul

Belal Hamzeh, Ph.D., VP, Wireless | b.hamzeh@cablelabs.com
Jennifer Andreoli-Fang, Ph.D., Distinguished Technologist | j.fang@cablelabs.com

Executive Summary

The growth in mobile data consumption has required Mobile Network Operators (MNOs) to deploy small cell networks in order to meet the demand. Although small cells provide a robust solution for increasing capacity, they also increase deployment complexity due to the need for a pervasive backhaul network that can support the mobile network requirements.

The hybrid fiber coaxial (HFC) networks have been traditionally used for video services and high speed broadband services using DOCSIS technologies. With the emerging backhaul needs for small cells, and benefiting from the availability and active nature of the HFC plant, cable networks are becoming a prime candidate for small cell backhaul.

There are three fundamental requirements needed from any network to provide mobile backhaul or fronthaul services: bandwidth, latency, and timing. This document provides an overview on advances in the DOCSIS technology now and in the near future that address these requirements.

Mobile Backhaul Requirements

Bandwidth requirements for each small cell is dependent on the service level agreement (SLA) defined by each mobile network operator (MNO), which is based on the maximum capacity the small cell is expected to carry. Counterintuitively, a cell requires peak capacity when it is serving a single user with good signal and low interference from neighboring cells. This user can thus achieve the highest modulation and coding schemes (MCS), and corresponds to the peak capacity the cell is required to support. In comparison, a cell's capacity during busy hours is significantly lower, when the resources are divided among the likely case of multiple simultaneous users with varying and non-optimum channel conditions.

Typical SLAs for small cells require 50-100 Mbps on the downlink and 10-20 Mbps on the uplink, accounting for the asymmetrical nature of the traffic.

Latency is a critical determining factor of user experience. High user plane roundtrip latency incurred end-to-end translates to lower QoE when the user accesses an application. Additionally, to achieve higher user throughput, capabilities such as dual connectivity require low user plane and control plane latency. While these numbers are vendor-dependent, they may range from 5 to 15 to 20 ms on the air interface.

To increase capacity for cell edge users, advanced interference coordination techniques are needed to coordinate the transmissions of neighboring cells. Low control plane latency is critical to ensure the effectiveness of the techniques. In Table 1, we list several LTE-A techniques with their corresponding latency requirements.

Since LTE is a synchronous technology, the cells need to be frequency synchronized for LTE FDD and phase synchronized for LTE TDD. Table 1 summarizes synchronization and latency requirements.

Table 1: Mobile Backhaul Requirements for Synchronization and Latency

	Frequency Sync	Phase Sync	Latency
LTE FDD	± 50 ppb	-	User plane, varies
LTE TDD	±50 ppb (wide area) ±100 ppb (local area) ±250 ppb (home)	10 µs (wide: cell radius >3km) 3 µs (local: cell radius <3km) 1.33 µs + Tprop (home eNB radius >500m) 3 µs (home eNB radius <500m)	User plane, varies
CoMP	-	± 1.5 µs	< 5 ms
eICIC	-	± 1.5 - 5 µs	< 5 ms
5G TDD	± 50 ppb (wide area) ±100 ppb (local area)	≤ ± 1.5µs	User plane, varies

DOCSIS Network Capacity

In a DOCSIS-based broadband network, the overall network capacity and the number of active CMs connected to the network define what service tiers can be offered by the network, subsequently determining the upper bound on the SLAs and the user experience the mobile network can deliver to the end user.

Network capacity is a function of the deployed DOCSIS technology and the amount of spectrum allocated to DOCSIS channels. The number of active CMs connected to the network is defined by the subscriber attach rate and the number of households spanned by a leg of the HFC network.

Downstream Capacity

DOCSIS 3.0: DOCSIS 3.0 technology implements Single Carrier Quadrature Amplitude Modulation (SC-QAM) on downstream (DS) channels, supporting both 64-QAM and 256-QAM modulation orders. Today, 256-QAM is the modulation order deployed on DOCSIS 3.0 DS channels, delivering a throughput of approximately 42.88 Mbps of PHY rate per 6 MHz channel. The DOCSIS 3.0 specifications defines 762 MHz (108 MHz to 870 MHz) of usable DS spectrum for DOCSIS 3.0 channels thus enabling a potential network capacity in excess of 5 Gbps. With the latest cable modems (CMs) supporting 32 SC-QAM channels (192 MHz of used spectrum), link rates in excess of 1.2 Gbps are achievable to a CM.

DOCSIS 3.1: The development of DOCSIS 3.1 technology delivers higher spectral efficiencies and also significantly increased the available spectrum for DOCSIS channels. The use of OFDM modulation, advanced Forward Error Correction techniques and higher modulation orders such as 4096-QAM significantly increases the spectral efficiency of DOCSIS networks. Field testing has shown that

4096-QAM, 2048-QAM and 1024-QAM modulation orders can be deployed across most of the network footprint, thus increasing spectral efficiency by nearly 50% in comparison to DOCSIS 3.0 networks. For example, 192 MHz of DOCSIS 3.1 channels can potentially deliver a throughput of nearly 1.9 Gbps. Additionally, the DOCSIS 3.1 specification also increases the usable DS spectrum to 1110 MHz (108 MHz to 1218 MHz) and optionally by up to 1686 MHz (108 MHz to 1794 MHz). A DOCSIS 3.1-compliant CM supports 2 DOCSIS 3.1 channels (192 MHz each) and 32 DOCSIS 3.0 channels, thus supporting a link rate in excess of 5 Gbps.

With the higher spectral efficiency and increased supported spectrum for DS channels, a DOCSIS 3.1 system can potentially provide 2x the capacity of a DOCSIS 3.0 network (3x the capacity if using the optional spectrum up to 1794 MHz).

Upstream Capacity

DOCSIS 3.0: DOCSIS 3.0 specification implements Time Division Multiple Access (TDMA) on the upstream channels, supporting modulation orders up to 64-QAM, delivering a throughput of approximately 30.72 Mbps of PHY rate per 6.4 MHz channel. The DOCSIS 3.0 specifications also support lower modulation orders and narrower channel widths to maximize the use of the available upstream (US) spectrum. DOCSIS 3.0 technology enables two upstream bands, 5 to 42 MHz and 5 to 85 MHz, thus supporting upstream bandwidths of 37 MHz or 80 MHz. Although the upstream bands start at 5 MHz, networks typically deploy upstream channels starting at around 15 MHz to avoid ingress noise in the lower part of the band. Subsequently, the available upstream capacity is approximately 135 Mbps (for 42 MHz networks) and 340 Mbps (for 85 MHz networks). The upstream capacity can be further increased by using the spectrum below 14 MHz and using the smaller channelization and lower modulation orders. With DOCSIS 3.1 CMs supporting up to 8 upstream channels, link rates of approximately 245 Mbps are achievable to a CM.

DOCSIS 3.1: Similar to the downstream (DS) channels, DOCSIS 3.1 technology delivers higher spectral efficiencies and also significantly increases the available spectrum for US DOCSIS channels.

The US orthogonal frequency-division multiple access (OFDMA) DOCSIS 3.1 channels support higher modulation orders such as 1024-QAM, and also increases the available US spectrum by enabling additional US spectrum, supporting 5 to 108 MHz and 5 to 204 MHz bands in addition to the legacy US bands.

The higher efficiencies of DOCSIS 3.1 networks significantly increases upstream capacities, enabling capacities up to 700 Mbps in a 5 to 85 MHz network and up to 1.5 Gbps in a 5 to 204 MHz network.

Full Duplex DOCSIS

Full Duplex (FDX) DOCSIS 3.1 technology extends the upstream capacity of DOCSIS 3.1 networks by enabling FDX communications in the 108 to 684 MHz spectrum. Thus, the spectrum that was traditionally used for DS communication only can now be simultaneously used for US communication as well. By enabling FDX communication in this band, the usable US spectrum is increased by 576 MHz, which previously has been limited to the 5 to 42 MHz, 5 to 85 MHz, 5 to 108 MHz or 5 to 204 MHz bands. With most networks using the 5 to 42 MHz or 5 to 85 MHz bands for US communication, this increases the available upstream spectrum by 16x (in the 5 to 42 MHz band case) and 8x (in the 5 to 108 MHz band case). The FDX spectrum is dedicated to DOCSIS 3.1 channels only, thus leveraging the higher efficiencies of the DOCSIS 3.1 technology, and enabling network upstream capacities of approximately 5 Gbps, which in turn enables the delivery of symmetric multi-Gbps service tiers over the HFC network.

Network Migration towards Fiber Deep Architecture

The HFC network architecture is typically described by the number of amplifiers that live between the node and the last house on the coaxial span. For example, if there are 4 amplifiers on the span, then the network is referred to as an N+4 network. A passive network architecture (no amplifiers) is referred to as an N+0 network.

Driving the fiber deeper into the network has many advantages. By eliminating amplifiers, another source of noise is removed from the network and thus the RF signal quality is improved, enabling the network to deliver higher modulation orders to a larger footprint and increasing the overall network capacity. Additionally, the number of houses connected to the same coaxial span is reduced, and thus there are fewer CMs competing for the shared medium and having access to a larger portion of the shared spectrum. As an example, an N+4 network would have 5 coaxial spans, and if we assume each coaxial span has 50 connected homes, then there are 250 homes sharing the medium. By driving fiber deeper, for example to an N+1 architecture, then there are only 100 homes sharing the medium, and thus now each CM has access to 2.5x the capacity of a CM connected to a N+4 network. This is not taking into account the potential increase of the overall network capacity due to the use of higher modulation orders.

DOCSIS Network Latency

Advanced interference coordination techniques in mobile technologies such as coordinated multipoint (CoMP) require 5 ms of X2 interface (eNB-eNB interface) latency to realize significant gain. Furthermore, most of 5G applications require 10 ms user-to-core latency, with some ultra-low latency applications requiring 1 ms latency. These requirements may be difficult to meet with DOCSIS networks today that have not been latency optimized.

DOCSIS Today

In the downstream, DOCSIS can generally achieve 1 ms of latency. On the upstream, DOCSIS specifies several types of scheduling services including best effort (BE), real-time polling service (RTPS), and unsolicited grant service (UGS). Most upstream data is transmitted with BE service. Best effort scheduling follows a request-grant-data loop as shown in Fig. 1. The requests are sent in the contention regions which the CMTS schedules regularly.

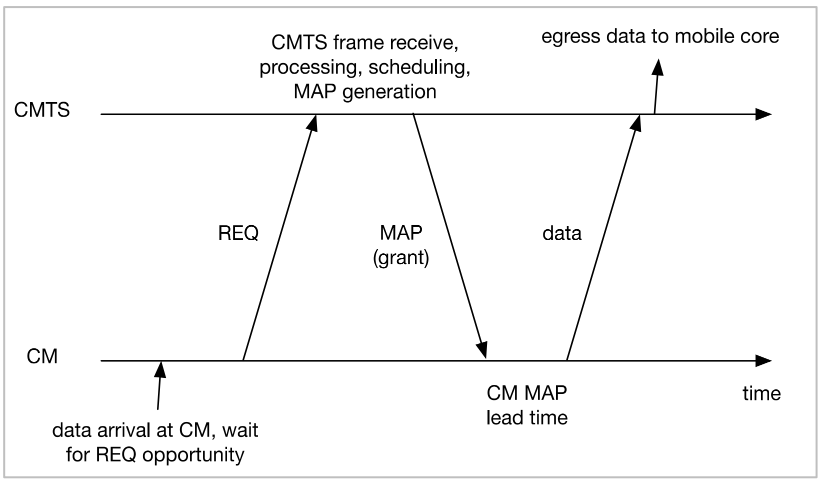

Figure 1 – DOCSIS REQ-GNT-data loop for BE or RTPS scheduling service

RTPS was designed to support real-time data flows that generate variable size packets periodically where the CMTS provides unicast request opportunities periodically.

As shown in Figure 1, after the CM detects data arrival and formulates a bandwidth request (REQ), it waits either for a contention region or for a polling opportunity to transmit the REQ, depending on the scheduling service for which the traffic is configured.

The CMTS scheduler typically processes REQs every 2 ms and generates a MAP that describes 2 ms worth of grant allocations. Since the CMTS sends at least one MAP in advance of that MAP's allocation start time, the shortest REQ-GNT cycle on DOCSIS is theoretically 4 ms. Additionally, the CM and the CMTS each need processing and lead time, typically 0.5 ms on each device. So, the practical minimum DOCSIS upstream latency is about 5 ms.

The average latency is expected to be higher due to the random arrivals of data and therefore, REQs, as well as the scheduling of contention regions. Higher network loading may also result in collision of the REQs, which triggers a truncated exponential backoff that will further increase latency. The average DOCSIS upstream latency for best effort traffic has been measured to be 11 - 15 ms with the potential of a significantly higher maximum latency up to 50 ms under medium to heavy channel utilization.

Lastly, UGS was designed to support real-time data flows, such as VoIP, that periodically generate fixed size packets. The CMTS provides fixed-size grants of bandwidth on a periodic basis. The CM utilizes the periodic grants to transmit data directly without sending REQs. UGS cannot be used to efficiently backhaul bursty mobile traffic.

Low Latency Mobile Backhaul

Current LTE-DOCSIS systems have cumulative latency as shown in the upper diagram of Figure 2. Since LTE uplink channel access works in a similar way as DOCSIS access, the operations of the two systems may be pipelined rather than working in serial. With pipelining, the DOCSIS request-schedule-grant loop can be started earlier and in parallel with the LTE request-schedule-grant loop, leading to much lower latency, as shown in the lower diagram of Figure 2. This is accomplished by treating the LTE and DOCSIS systems as one pipelined system rather than the two independent systems they are today. A “bandwidth report (BWR)” message is sent from the LTE scheduler to the DOCSIS scheduler that includes the amount of bytes the LTE schedule expects at a precise time in the future.

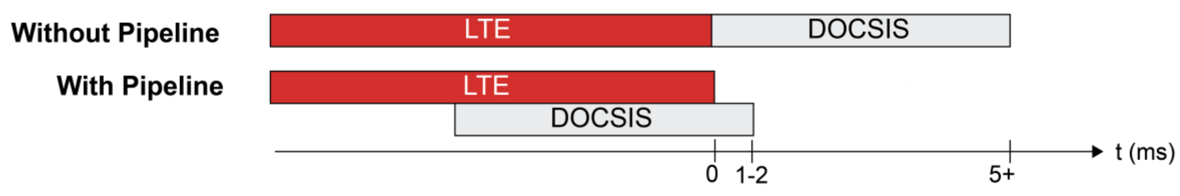

Figure 2 – Reducing DOCSIS effective latency when backhauling LTE

With pipelining, DOCSIS technologies can theoretically achieve an effective latency of 0 ms. In practice, engineering margin needs to be built in to allow for small errors resulting from synchronization, etc. Even so, the DOCSIS latency has been measured to be in the 1 - 2 ms range with pipelining (see Figure 3).

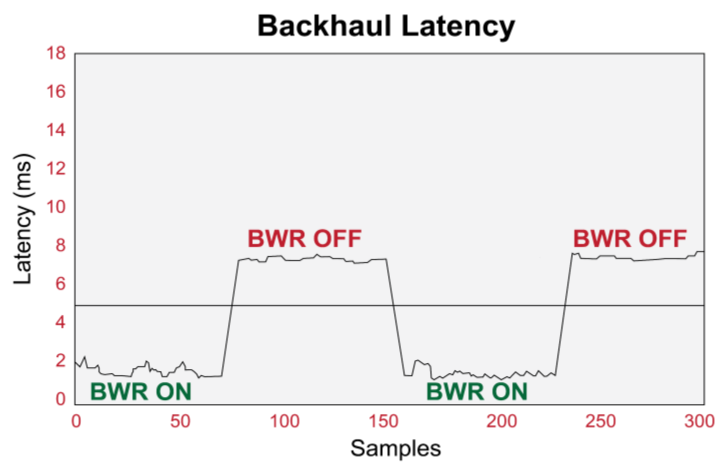

Figure 3 – DOCSIS latency with and without pipelining

As 5G New Radio (NR) continues to use dynamic granting for majority of channel access, the pipelining method is extensible and futureproof for majority of 5G applications. Pipelining can enable 5G low latency applications such as ultra-reliable low latency communications (URLLC) by serving as a real-time error connection mechanism when combined with DOCSIS predictive granting.

Timing and Synchronization

The ITU-T standardizes a set of recommendations on carrying IEEE-1588 timing from a grandmaster clock to a mobile base station over a chain of Ethernet equipment. With DOCSIS performing backhaul for the base station, the DOCSIS Technologies for Mobile Backhaul Specification defines a DOCSIS interworking function (IWF) that translates Precision Time Protocol (PTP) timing to DOCSIS timing and vice versa. This allows the DOCSIS-based timing distribution chain to fit into the standard ITU-T synchronization framework.

DOCSIS IWF supports 3 synchronization approaches for mobile backhaul:

- Full timing support networks distributing phase synchronization: the network is comprised exclusively of network elements that support IEEE-1588 protocol operation. See Figure 4 for an end-to-end view.

- Partial timing support network distributing phase synchronization: the network may include network elements that are not IEEE-1588 aware.

- Physical layer timing support network distributing frequency synchronization: the network supports Synchronous Ethernet (SyncE), and the end application only requires a layer 1 clock.

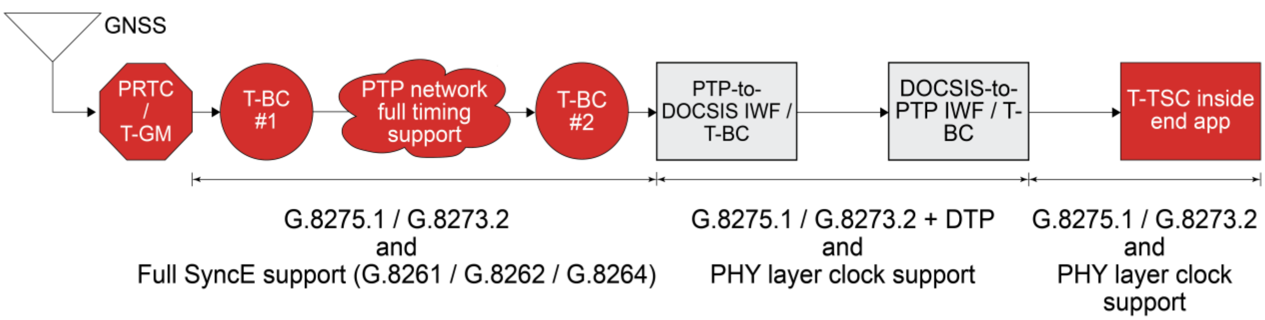

Figure 4 – Full timing support network for phase sync

The asymmetry of the HFC plant poses challenges to carry precision frequency and timing services over DOCSIS networks for mobile backhaul. The DOCSIS Timing Protocol (DTP) was introduced in the DOCSIS 3.1 specifications and allows for the passing of IEEE-1588 timing over the DOCSIS network with no jitter from network buffering. DTP is used to work out network asymmetry by providing two basic services:

1. A hardware path for a timestamp
2. A signaling path that determines the downstream timing offset which is used as a correction factor for PTP

To support the addition of a DOCSIS network in the timing distribution chain, a portion of the time error budget from the Ethernet portion of the overall timing distribution network is reallocated to the DOCSIS part of the network. The DOCSIS Technologies for Mobile Backhaul Specification specifies a set of time error budgeting requirements on the DOCSIS equipment, such that the ± 1.5 µs of grandmaster-to-end applications is met.

Other Considerations

In addition to the technical capabilities of DOCSIS technologies in meeting the mobile backhaul requirements, other factors are also critical: namely reliability, location and power.

DOCSIS technologies have built-in Proactive Network Maintenance (PNM) capabilities that enable a cable operator to remotely monitor the HFC network. With these capabilities, a cable network operator is able to detect, identify and locate network impairments before they become service disrupting. For example, every DOCSIS 3.1 Cable Modem can provide network visibility equivalent to what can be obtained through a multitude of test equipment such as a spectrum analyzer, a network analyzer and a vector signal analyzer.

The HFC network has a large footprint delivering services to large populations. For example, in the USA, the HFC network passes 93% of the households in addition to providing services to a large number of enterprises, thus making the HFC network available where mobile network connectivity is mostly needed. Additionally, the HFC network has the capability to supply power, thus eliminating the need to connect to external power sources and reducing deployment complexity.

Conclusion

There are three fundamental requirements for DOCSIS networks to provide mobile backhaul services: capacity, latency, and timing and synchronization. Depending on the MNO’s defined SLA, DOCSIS 3.0 networks today can already support mobile backhaul. The DOCSIS downstream capacity is significantly improved with DOCSIS 3.1 networks, and Full-Duplex DOCSIS technology will significantly increase upstream capacity.

DOCSIS 3.1 technology added the support for distributing IEEE-1588 timing over the DOCSIS networks. With recent CableLabs work on a Mobile Backhaul specification, DOCSIS 3.1 specifications will be able to achieve the stringent phase precision required for deploying LTE TDD networks. Additionally, control and user plane latency is expected to improve significantly, achieving 1 - 2 ms latency with pipelining across DOCSIS technologies and mobile technologies.