Mobile Backhaul over DOCSIS

A Technical Paper prepared for SCTE/ISBE by

John T. Chapman
CTO Cable & Fellow
Cisco
San Jose, CA 95124
408-526-7651
jchapman@cisco.com

Jennifer Andreoli-Fang
Distinguished Technologist
CableLabs
Boulder, CO
303-661-3838
j.fang@cablelabs.com

Table of Contents

List of Figures

List of Tables

Introduction

One of the next big opportunities for cable operators is mobile. For a long time, cable and mobile have been competitors, both competing for that voice or Internet subscriber, or even a video subscriber. Now, they could become the best of friends. To move into small cell deployment, mobile operators need someone with a plant that can backhaul their user traffic. Cable can be that partner. Cable, meanwhile, has already moved in a direction of wireless as the last hop, and LTE is just another access technology for that last hop.

In this white paper, we examine DOCSIS as a viable option for mobile backhaul technology, particularly for small cells. The business case is that when the small cells are deployed deep in the network, DOCSIS will be there and fiber may not be. Should money be spent on installing new fiber, or can LTE use DOCSIS instead? Judging by all the WiFi deployments being backhauled by cable operators, it makes sense to consider DOCSIS as a backhaul for LTE as well.

There is money to be made and money to be lost. As always, it all depends on how much money it takes to build and operate the network. Keeping the costs down is a great motivation and is something DOCSIS has always taken notice of.

The Cable-Mobile Market

1. Market Opportunity

A Mobile Network Operator (MNO) owns their own plant and sells services. A Mobile Virtual Network Operators (MVNO) rents another operator’s plant and sells services. About 50% of CableLabs members are already MNOs or MVNOs. The other 50% are trying to figure out if they should take the plunge as well.

It makes sense. It is difficult to grow revenues significantly in cable when data/voice/video services already have high market penetration and success. Yet, consumers pay similar amounts for cable and mobile services. Most consumers are paying for Internet access from both cable and mobile! By combining cable and mobile services, operators can increase revenue and attract consumers with an attractive combined service at a potentially lower cost.

Today, mobile operators rely on macrocells to provide Long Term Evolution (LTE) services. Macrocells cover areas on the order of one to ten miles in radius. The next step in the evolution of the mobile RF plant is to deploy small cells, and do so in a denser manner. Small cells have an effective radius ranging from 100 meters to 500 meters. Coverage varies and is based upon many factors such as transmitting frequency, antenna height, transmit power and line-of-sight. Small cell footprint also includes both indoor and outdoor coverage.

The movement from macrocell to small cell is analogous to the hybrid fiber-coax (HFC) movement from N+5 (node plus 5 amplifiers) to N+0 (node plus zero amplifiers). Both are segmentations of the physical plant to increase available bandwidth per subscriber and both are happening in the marketplace.

A cable operator could enter the mobile market in a sequence of steps. They rent mobile air time from an incumbent MNO with macrocell-only plant to provide coverage. The cable operator could then begin to build out a small cell footprint. In doing so, they would decrease the macrocell airtime they need to rent. He could then also rent out the small cell footprint back to the same MNO he is renting the macrocell from. This would reduce the operational expenses. Over time, when the small cell network is built out, the rent-back small cell revenue might equal the rent for the macrocells. Now the business case is cash neutral and the cable operator has built a small cell pant with the help of the rent revenue form the MNO. The next step is for the cable operator to move from being an MVNO to being an MNO.

2. Basic Deployment Requirements for Small Cell

There are at least three major items that are required for a small cell deployment:

1. Location
2. Power
3. Backhaul

The cable operators are in an advantageous position when it comes to offering a solution for all three of these requirements.

1. Location

   Cable operators have access to public rights-of-way that can help when placing new infrastructure but that is only true on an overhead plant. Underground plants can be tricky as there is limited antenna access. After that, there is “urban furniture” that works – consumer roof tops, lamp posts, street lights, utility poles, etc.

2. Power

   The HFC plant is already powered and can support a small number of small cells. Current plants may not be able to support a large number of small cells but power distribution over coax is flexible and more power can be added if new small cell deployments require it and a business case is made. The responsibility to power the small cell power might even be borne by the consumer, in a residential or business setting where the consumer is looking add service or improve coverage.

3. Backhaul

   The HFC plant has two resources: fiber and coax.

   Fiber

   Today, signals transmit over fiber are analog-based but fiber communications will soon be upgraded to Ethernet (digital fiber) - a requirement for DOCSIS Remote PHY. Wherever there is fiber, small cells can easily be connected.

   Coax

   Coax is also a very important resource because small cells may not be located near the node. The best location for nodes and small cells don’t always coincide because small cells typically need to be placed above roof tops, tree lines and hills to provide optimal LTE signal coverage. Luckily, coax is nearly everywhere already, especially in residential areas where coax, and thus DOCSIS, is almost always available and fiber may not be.

Imagine a business scenario where the cable team decides it is not cost effective to upgrade a neighborhood from N+5 to N+0 (fiber deep). Then the mobile team from the same operator comes along and deploys fiber to the home to support the small cells. That would be a good example of siloed business economics – where two divisions arrived at opposite business cases instead of a blended business case.

3. DOCSIS vs. Fiber

Throughput

Fiber currently achieves a throughput of 10 Gbps. DOCSIS is competitive with fiber. For example, DOCSIS 3.1 enables throughputs as high as 10 Gbps in the downstream and 5 Gbps in the upstream, using FDX DOCSIS.

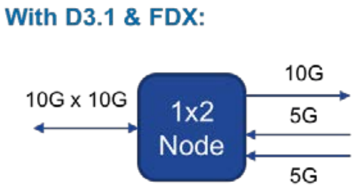

Figure 1 – FDX Optical Node

To achieve a comparable DOCSIS throughput to fiber throughput, a node could be deployed as show in figure 1. Two return paths can be used with FDX DOCSIS on each return path to achieve a bandwidth of 10 Gbps x 10 Gbps - the same bandwidth as fiber!

These DOCSIS throughputs over coax give operators a lot of flexibility. The operators can deploy the small cell anywhere in the HFC footprint and then choose fiber or coax, based on convenience and cost.

Network Timing

Small cells located on the outside plant have the option to receive network timing directly from GPRS. Indoor small cells, for both residential and commercial, don have direct GRPS access and there would need to receive network timing over DOCSIS. Timing for mobile backhaul over DOCSIS is addressed in 0 0.

Latency

In addition to sufficient bandwidth, a backhaul also needs timing and the ability to offer low latency to high priority traffic. In this paper we will address how to achieve low latency backhaul over DOCSIS.

DOCSIS and LTE Working Together

4. Comparing DOCSIS and LTE Latency

If you are familiar with how DOCSIS works, then you will be able to pick up LTE quickly. Both DOCSIS and LTE are trying to solve very similar problems. Both technologies try to manage a point to multipoint network. Both LTE and DOCSIS use a scheduled upstream and both are managing subscriber traffic.

The main difference between LTE and DOCSIS is that LTE manages a wireless network, while DOCSIS manages a wired network and the protocols are different since they were designed by different committees. The wireless channel and mobility aspect of LTE adds challenges that don’t need to be dealt with in DOCSIS but once connected, scheduling and data transfers between customers and the (evolved NodeB) eNodeB and the CMTS follow similar principles.

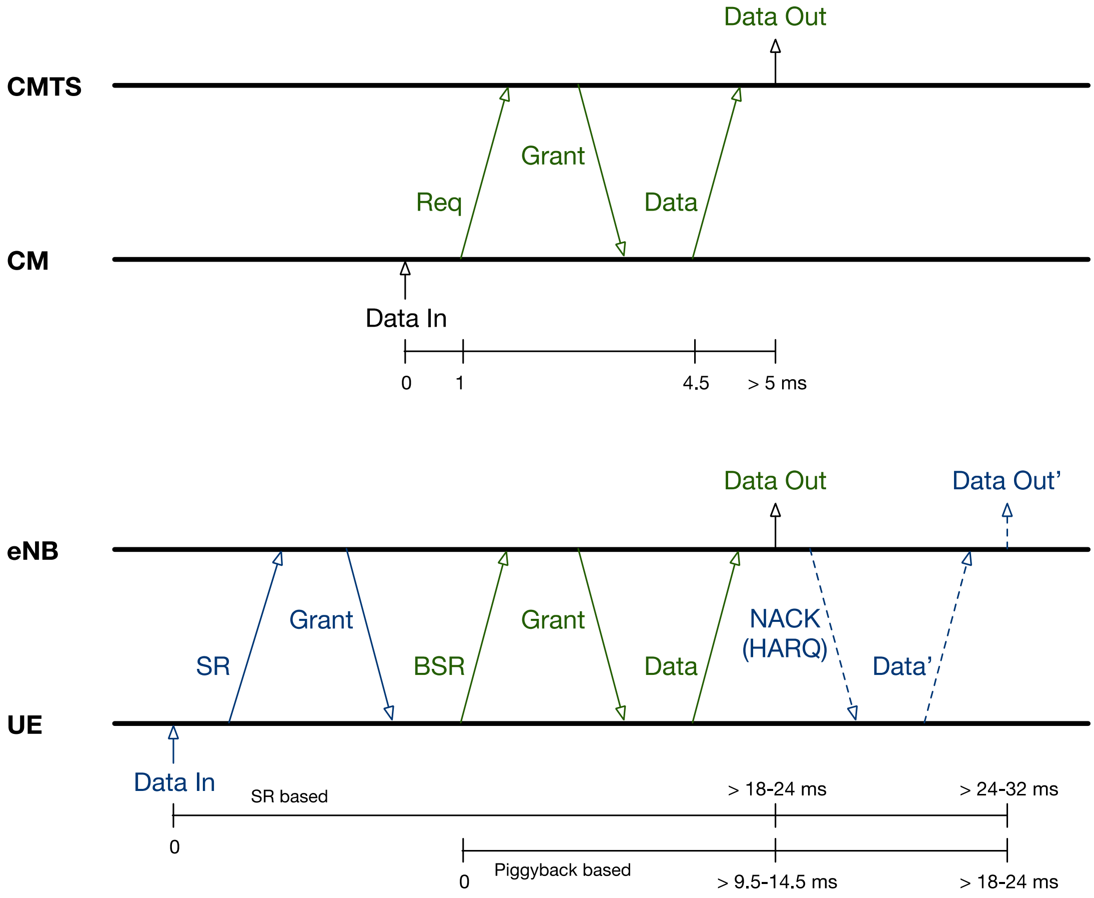

Figure 2 – Comparing DOCSIS and LTE

Duplexing

Another difference between LTE and DOCSIS is that LTE supports frequency division duplex (FDD) and time division duplex (TDD). FDD is similar to HFC in that uplink (UL) & downlink (DL) are segregated by frequency bands. In TDD, essential to massive MIMO antenna operation, the same band is used for UL and DL and the interleaving of the UL and DL transmissions in the same spectrum adds access delay for UL transmissions. Consequently, TDD round trip time (RTT) can be much higher than FDD.

DOCSIS Upstream Scheduling

The DOCSIS upstream uses a request-grant approach as shown in Figure 2. A request message is sent from the cable modem (CM) to the cable modem termination system (CMTS). The CMTS prepares a DOCSIS MAC management message (MMM), called the MAP, and inserts an entry indicating a grant for the CM. The grant entry in the MAP message contains an upstream service identifier (SID) associated with a service flow assigned to the CM, a transmission time signaled as a mini-slot number, and the number of bytes to be transmitted. The CMTS transmits the MAP to the CM, so that the CM can make use of the grant to send its upstream data to the CMTS.

There are many factors that determine the request-grant time in DOCSIS. The DOCSIS specification contains a detailed analysis of the request-grant delay 0. The minimum delay is basically every second MAP time plus some circuit and queuing delay. The DOCSIS 3.0 CMTSs currently use a 2 ms MAP time. One published set of test results for a Cisco uBR10K CMTS showed a 5 ms minimum request-grant response time 0.

The CMTS scheduling algorithm can have a significant impact on request-grant delay. For instance, if the CMTS is using a best effort (BE) scheduling algorithm, requests can be made in contention slots where the requests could fail on the first try. If the CMTS is using the real-time polling service (rtPS) scheduling algorithm, the request will be placed in a dedicated slot which ensures that the requests are always successful. BE can perform better than rtPS when DOCSIS is idle because many request contention slots are available. In that case, BE provides lower latency than rtPS. However, in a busy system where there are fewer contention request slots and BE will need to re-request often, rtPS can provide guaranteed latency. Requests can also be sent as a piggyback message with a data packet. Piggybacking is deterministic within a flow and avoids contention.

LTE Uplink Scheduling

At the heart of the LTE uplink scheduler is a similar request-grant mechanism. The request is sent from the user equipment (UE) to an eNodeB. The UE is typically a cellular phone while the eNodeB is a macrocell or a small cell. When data arrives at the UE, the UE first determines if it already has an LTE uplink (UL) grant. If it does not, the UE waits for an opportunity to transmit a scheduling request (SR), where typical SR opportunities can come along once every 1 to 10 ms depending on the configured periodicity of the SR. The purpose of the SR is to keep the upstream signaling overhead low which is critical in wireless where spectrum is costly.

Upon receiving the SR, the eNodeB schedules an UL grant so that the UE can transmit a buffer status report (BSR) to the eNodeB. Once the eNodeB receives the BSR from the UE, it becomes aware of the outstanding UL data present in the UE's upstream queues. The eNodeB then schedules UL grants and transmits the grants to the UE via a downlink control information format 0 (DCI-0) message transmit in the physical downlink control channel (PDCCH). LTE operates every subframe which is 1 ms. However, due to processing constraints, the eNodeB's BSR-grant delay is typically 4 ms for both the eNB and the UE. In other words, upon receiving the BSR, the eNB performs scheduling and sends DCI-0 to arrive at the UE in 4 subframes, and the UE is scheduled to transmit 4 subframes later. Assuming a very small SR opportunity periodicity (1 ms), the minimum delay before a UE can transmit UL data is 18 ms which is slightly longer than the minimum DOCSIS request-grant latency. Detailed request-grant latency calculations are shown in Table 1.

Table 1 – LTE REQ-GNT Latency

Latency Components	**Latency (ms)**	
	SR	**Piggyback**
Waiting time for SR (assume configured SR period of 5 ms)	0.5 – 5.5	n/a
UE sends SR, eNB decodes SR, eNB generates grant for BSR	4	n/a
eNB sends grant, UE processes grant, UE generates BSR	4	0-4
eNB processes BSR, eNB generates grant for data	4	4
eNB sends grant, UE processes grant, UE sends UL data	4	4
eNB decodes UL data (estimate)	1.5 – 2.5	1.5 – 2.5
Total	**18 – 24**	**9.5-14.5**

It is worth noting that both the DOCSIS and LTE system have design efforts underway to decrease the minimum latency, so the latency figures given in Table 1 are for currently deployed systems. It is also worth noting that in a congested network, both DOCSIS and LTE will have longer latencies.

Re-transmissions

The LTE transmissions over the LTE air interface often occur in a harsh wireless environment (eg. cell-edge). Wireless signals are easily degraded due to path loss and interference and transmission errors are far more likely than in DOCSIS. To overcome the transmission errors, LTE has two retransmission mechanisms:

1. Hybrid Automatic Repeat reQuest (HARQ) which operates at the MAC layer.
2. Automatic Repeat reQuest (ARQ) which operates at the radio link control (RLC) layer.

HARQ

HARQ is intended to quickly re-transmit the LTE transport blocks to recover from most errors in conjunction with a good forward error correcting (FEC) algorithm, while RLC ARQ is a higher layer re-transmission mechanism with which has higher overhead but improves the reliability of the link after HARQ.

RLC ARQ

The use of RLC ARQ is optional in LTE and depends on the type of traffic and quality of service (QoS) parameters (error rate, latency, bit rate, etc) desired. For example, for voice over LTE (VoLTE) packet loss is not that critical, but latency is. The human ear can tolerate small amounts of packet loss but long delays quickly become annoying during a phone conversation. Therefore, for voice traffic, ARQ retransmissions are not required and a LTE EPS bearer for voice can be configured to use the RLC Unacknowledged Mode (UM) entity which does not use ARQ. On the other hand, transmission control protocol (TCP) traffic packet loss triggers slow start, and throughput can suffer. However, longer delays can be tolerated by TCP traffic therefore ARQ is used by configuring the LTE EPS bearer to use the RLC Acknowledged Mode (AM) entity.

The purpose of an additional PHY layer retransmission mechanism is to reduce latency. After transmitting a coded block of bits which is termed a Resource Block (RB), the transmitter, the UE, or the eNB, keeps the RB in the transmission buffer. The receiver PHY layer decodes the RB, and passes the CRC results to the MAC layer. The MAC layer then issues either a HARQ ACK or NACK based on the results. Between the time the RB is received at the receiver and the time a HARQ feedback must be received at the original transmitter of the RB, there is an exact 4 ms of delay.

Overall LTE and DOCSIS Latencies

In summary, if we consider only lightly loaded systems and FDD duplexing for LTE, DOCSIS 3.0 has a minimum request grant-grant delay of about 5 ms while 4G LTE has a minimum request-grant delay of 18 to 24 ms without re-transmission and 26 to 34 ms with one HARQ retransmission. These latency values will increase under higher loads or if TDD is used in LTE.

Applicability of the Work in This Paper

This paper presents a method to reduce DOCSIS backhaul latency for LTE using pipelining and a new inter-system request message called the bandwidth report (BWR). We focus on DOCSIS 3.0 and 4G LTE since both technologies are well documented and widely deployed and the BWR can be used today. The research covers both LTE FDD and TDD, but we only include FDD timing examples for simplicity. The general principles proposed in this white paper will extend to DOCSIS 3.1 and 5G LTE, although some modifications to the protocols may be required.

5. The Bandwidth Report Concept

LTE request-grant time is much longer than DOCSIS, yet the mobile operators often demand backhaul delays to be on the order of a few milliseconds 0. There is a hidden opportunity to this high LTE latency. What if we could use the long LTE latency to start DOCSIS processes in parallel and reduce the overall LTE-DOCSIS system latency? This concept is shown in Figure 3 and is made possible using the BWR message.

Current LTE-DOCSIS systems, which do not have BWR, have cumulative latency as shown in the upper diagram of Figure 3. When the BWR is introduced into the LTE-DOCSIS system, 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 3. This is accomplished by treating the LTE and DOCSIS systems as one pipelined system rather than the two independent systems they are today. For this pipeline mechanism to work, DOCSIS needs some information from LTE in the form of the BWR message.

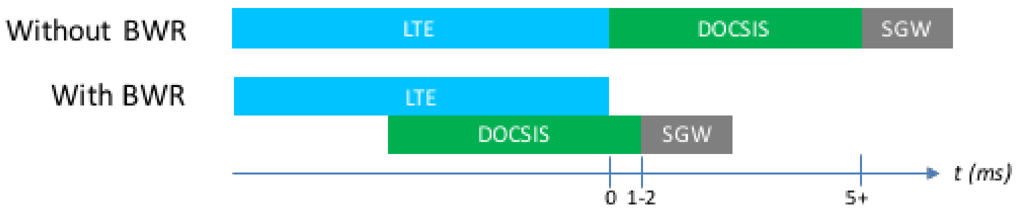

Figure 3 – Reducing DOCSIS Effective Latency Under LTE

How does the BWR work?

The BWR allows an external system, such as a LTE eNodeB, to request bandwidth from a DOCSIS system for some specific time *in the future, before* the arrival of the actual traffic. The eNodeB provides a future traffic profile through the BWR message that allows the CMTS to make QoS and granting decisions earlier than it normally would.

The BWR replaces the CM's internal layer 2 request message for LTE uplink data arriving at the CM. The BWR message itself is an external layer 3 IP based message that is transmitted from the eNodeB to the CMTS, through the CM. The content of the BWR is populated with information describing future data that will arrive at the CM.

The BWR is created by the eNodeB's LTE UL scheduler just after the scheduler finishes granting the UE(s) based on the UE's outstanding UL data buffer sizes. The BWR message created after scheduling contains a description of the amount of UL data that is expected to arrive at the eNodeB and be forwarded to the CM. The BWR message is then transmit from the eNodeB to the CMTS, via the CM, in a dedicated upstream service flow earlier than the CM would normally issue its request.

The result of using the BWR is that DOCSIS is made aware of the LTE scheduler's scheduling decisions using the BWR message. Rather than having two separate and sequential latency-additive LTE and DOCSIS request-grant loops, the BWR starts DOCSIS's request-grant loop early and in parallel to LTE's uplink granting. The CMTS scheduler grants the CM just-in-time for the UE's upstream data to arrive at the CM, which the effect of reducing the upstream latency. This is illustrated in Figure 4.

Figure 4 – Pipelining Requests and Grants

Figure 5 depicts an LTE system backhauled by DOCSIS. The resemblance between the LTE and the DOCSIS systems are striking. In both systems, the UE/CM generates a request when they have data to send, and the scheduler in the eNodeB/CMTS responds with a grant.

When the BWR is used, the LTE system's predictor tries to describe the data flow across the Ethernet interface by estimating the data flow across the air interface using the LTE UL scheduler's grants. Unfortunately, the data flow across the air interface will not exactly match the data flow that will occur later across the Ethernet interface due to segmentation and decoding failures, which will be explained in greater detail in section 6.

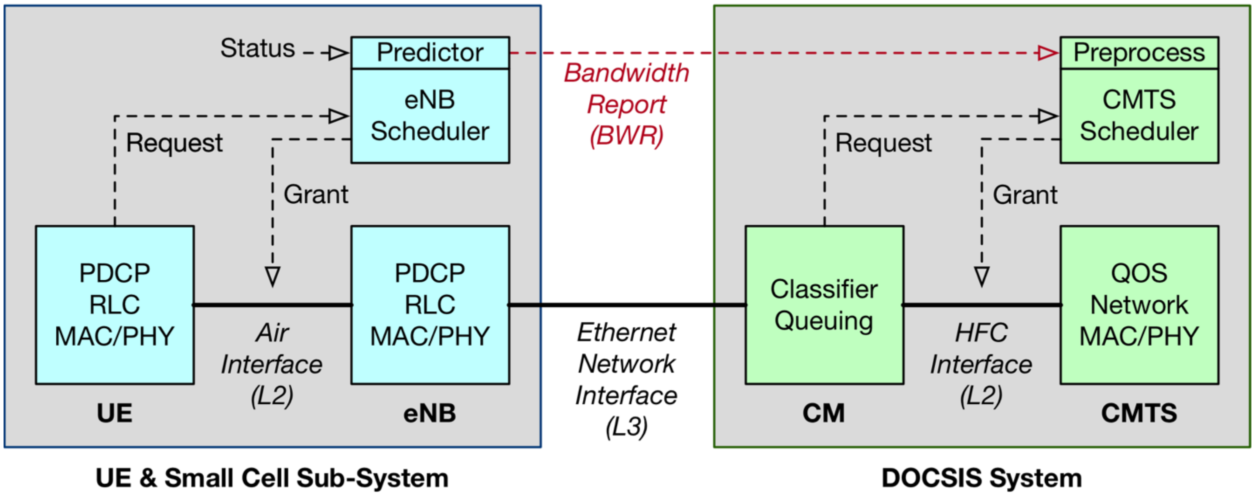

Figure 5 – Linking DOCSIS and LTE Schedulers with BWR

To adapt the BWR message to the traffic flow, a prediction algorithm is used to monitor the operation of the eNodeB scheduler. The predictor looks at the status of the data path to see if there are any changes in behavior and then issues the BWR message. The optimal period for BWR for LTE FDD is 1 ms because the eNodeB scheduler makes scheduling decision once every LTE subframe (every 1 ms). The CMTS uses the information given in the BWR messages, along with other scheduling constraints, and generates grant(s) for the CM.

BWR Example

If a 1000-byte data transfer from the UE will occur in 8 ms and there is an anticipated 1 ms delay in the eNodeB software stack for transport block decoding and packet reassembly before transmission to the CM, a BWR will be transmit to the CMTS to request 1000-bytes in 9 ms. The CMTS could be configured to add an additional 2 ms to the grants generated based on the BWR to ensure they any delayed traffic leaving the eNodeB. Therefore, the CMTS would issue a DOCSIS grant 11 ms after the BWR is created.

That is the basic idea of BWR. Of course, nothing is quite that simple …

6. BWR Additional Design Considerations

6.1. Synchronization and Timing

If the LTE system is scheduling in the future and the DOCSIS system is granting in the future, both systems must be synchronized in time for the grants to line-up. When high accuracy synchronization, such as GPS is not available or practical, another method to achieve good synchronization like the IEEE 1588 protocol is needed.

For DOCSIS to understand the LTE BWR message and issue grants at the correct time, a common time system must be referenced by the BWR. One proposal is to use a time index (TI) field in the BWR that stores the time in a similar way as IEEE 1588 timestamps with a reduced resolution of 1 ms to keep the size of the bit field short. With a 1 ms resolution, a 16 bit field could be used for the TI field, for example, with a defined roll-over procedure.

Another proposal is to communicate the LTE frame and subframe to the CMTS. LTE frames are numbered using 10 bits, from 0 to 1023, with rollover beyond 1023, while subframes are numbered 0 to 9 and can be represented using 4 bits. This scheme also requires a total of 16 bits.

The CMTS would translate the TI value to a mini-slot and issue a grant for the upstream data.

6.2. DOCSIS Mini-slot and LTE Subframe Misalignment

LTE and DOCSIS use different framing mechanisms for their transmissions. For instance, the LTE air interface transmissions are based on frames and subframes. A subframe is 1 ms in duration and a frame is 10 ms in duration. Meanwhile, DOCSIS uses a MAP interval of approximately 2 ms and grants are issued on a mini-slot level of granularity. The two framing systems will not be aligned in time. This is illustrated in Figure 6. In essence, the eNB scheduler and the CMTS scheduler are two asynchronous systems that are now trying to communicate scheduling information.

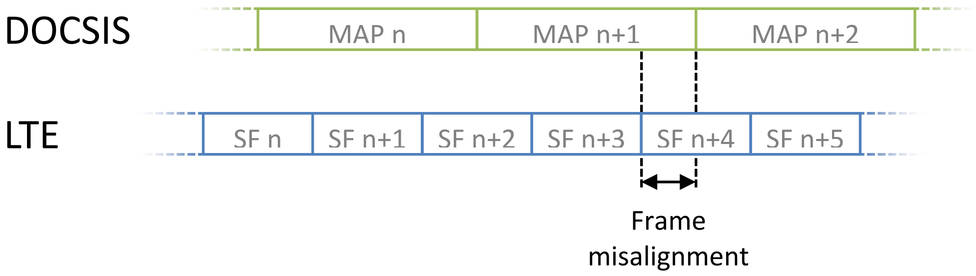

Figure 6 – LTE and DOCSIS Frame Misalignement Despite System Time Synchronization

It is often hard to schedule grants at exact times in a DOCSIS MAP interval because doing so may excessively fragment other flows. Thus, DOCSIS can be configured to allow the CMTS's upstream scheduler some flexibility during scheduling, which results in some grant jitter. This means that the selected mini-slot may not align perfectly with the upstream egress data from the eNodeB.

The CMTS needs to account for all the grant and upstream data egress jitter by adding some engineering margin, otherwise the grant may miss some packets and latency will increase.

6.3. HARQ

If the data from the UE is not received correctly at the eNodeB, the eNodeB will transmit a negative acknowledgement (NACK) to the UE which will trigger data retransmission. This means that the grant requested in the BWR message will not get used for the intended data transfer on the DOCSIS upstream link. The grant can get used by other data, if other data has been delayed, but the grant may also not get used and therefore some DOCSIS bandwidth will be wasted. Refer to 0 for an upper bound on the unused DOCSIS grant.

Since the LTE scheduler knows when a transport block failed to be decoded successfully, the LTE BWR predictor knows when re-transmissions are going to occur. If the LTE transport block is successfully received when it is re-transmitted to the eNodeB, some packets may egress from the eNodeB. The predictor must then account for the delayed and re-transmit data by adding the byte to a BWR corresponding to the subframe that includes the re-retransmission. More details about this operation and its performance can be found in 0.

6.4. RLC Packet Reassembly

The packets that are transmitted from the UE to the eNodeB are segmented (if required) and placed into LTE transport blocks. If a packet fits entirely within a LTE transport block, the UE can transmit it entirely in one subframe. When the eNodeB correctly receives a complete packet, the packet could get passed immediately up the eNodeB's software stack and out the eNodeB's Ethernet interface.

In the case where not all the segments of a packet have been received, the segments will be stored in a buffer at the RLC layer until all the segments are received. Once all the segments are received, the RLC entity re-assembles the segments into a complete packet and the packet can finally be transmitted out the eNodeB's Ethernet interface.

For an example of the effect of segmentation and re-assembly, let us assume that a packet is segmented into 2 pieces across 2 LTE transport blocks and both are received successfully (no decoding failures). The 1st transport block is transmitted from the UE to the eNodeB but it only contains the 1st segment of the packet. Since the packet is incomplete, the RLC entity handling the EPS bearer will store the segment and wait for the next transmission. 1 ms later in the next subframe, the 2nd transport block is received at the eNodeB physical layer and after decoding the RLC entity receives the 2nd segment of the packet. RLC reassembles the complete packet and the packet is sent out the eNodeB's Ethernet interface.

In our example, the result of the packet fragmentation that occurred on the LTE air interface added an additional 1 ms delay to the time when packet was expected to leave the eNodeB. While this doesn't seem like much, consider that for each additional segment an additional 1 ms delay (eg. N_{frag} - 1 milliseconds) is added to the packet egress time and a packet that is late even by a small delay may miss a DOCSIS grant. This is shown in Figure 7.

This packet fragmentation is one of the causes of mismatches between the bandwidth requested by the BWR message and the actual data leaving the eNodeB's upstream Ethernet interface.

Predicting the packet fragmentation and reassembly delay is a difficult problem. The LTE predictor algorithm is predicting the future while the reassembly engine works in the present. If the LTE predictor has packet level visibility when scheduling, then it can address this problem. If it does not then some engineering margin can be added by the predictor to the BWR message time index field to capture segmented packets.

It may appear that the solution is for the BWR predictor to increase the engineering margin in the BWR message. However, there is a tradeoff between increasing the engineering margin to capture more upstream packets and upstream latency gain. The engineering margin that maximizes the upstream latency gain of the BWR feature doesn't necessarily correspond with engineering margin that captures 100% of the delayed packets. So careful BWR predictor and engineering margin selection is needed!

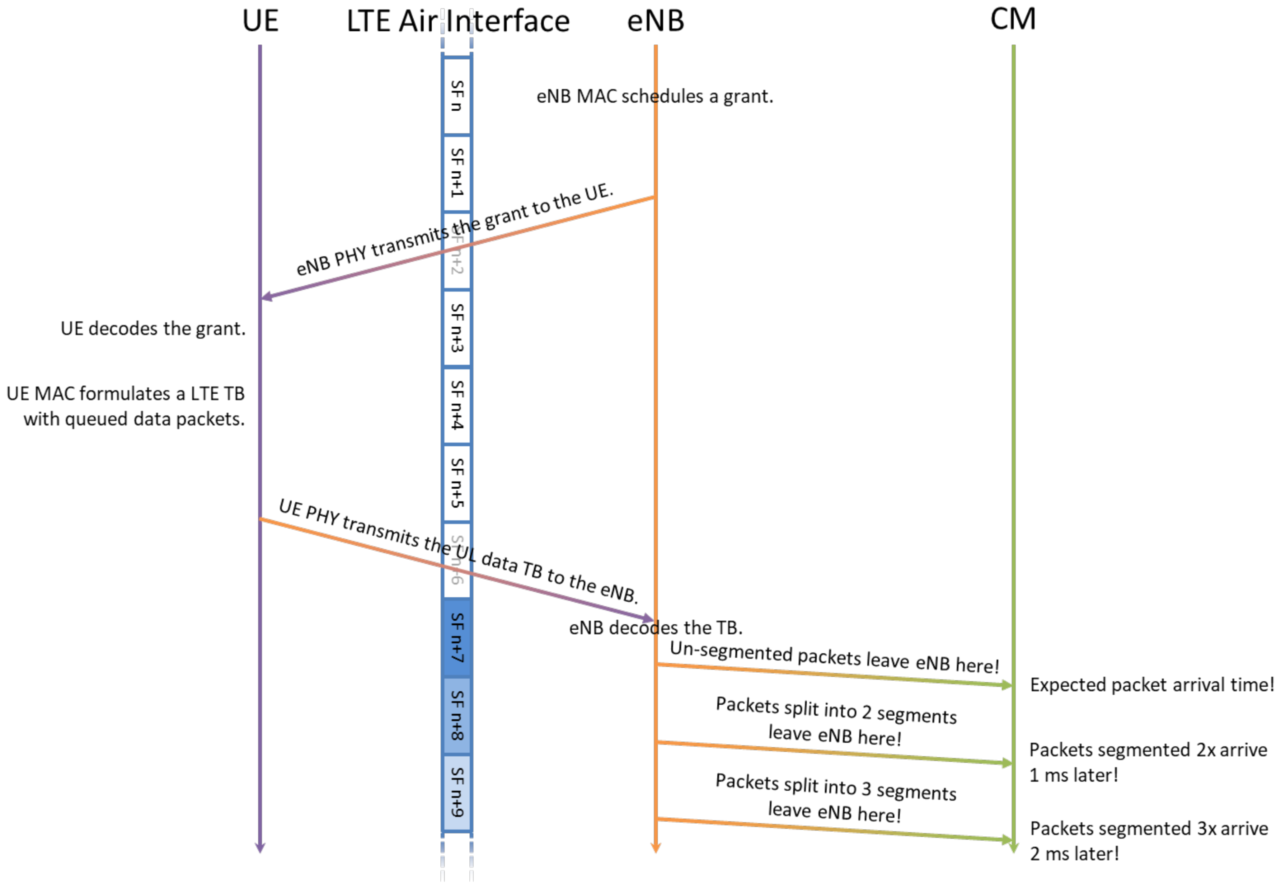

Figure 7 – Packet Delay Due to Packet Segmentation on the LTE Air Interface

6.5. eNodeB Ethernet Queuing and Transmission Delay

All physical devices have speed and bandwidth limitations and queues to help buffer data when the load on the interface is high. In a QoS based system, there will likely be multiple queues serving traffic with different priorities. The queue depth depends on the traffic and the time it takes for a packet to arrive at its destination depends on the time the packet waits in the queue and transmission on the Ethernet interface.

The LTE predictor must allow for the queueing and transmission time. If the predictor knows which QoS level each packet flow is, it can allocate shorter times for the high QoS traffic and longer times for lower QoS traffic to compensate for the time the packet could spend in a queue.

The Ethernet packet transmission time also adds delay. Luckily the transmission delay is small compared to other delays. For example, 1518-byte Ethernet packet on a gigabit interface only takes about 12 microseconds.

6.6. Signaling Traffic

Real-Time Signaling Traffic

LTE signaling could be categorized into the following three categories. This generalization applies to all signaling packets that leave the Ethernet interface of the eNodeB, whether they are intended for another eNodeB or intended for the evolved packet core (EPC).

1. Fully Scheduled Signaling

   UE signaling traffic transmit to the network that is accounted for by the eNodeB scheduler.

2. Partly Scheduled Signaling:

   eNodeB signaling transactions that have one or more transactions to the UE followed by one or more transactions to the Ethernet interface. This signaling is only partially accounted for by the eNodeB scheduler.

3. Unscheduled Signaling

   Signaling traffic originating at the eNodeB, which is not accounted for by the eNodeB scheduler.

The predictor algorithm will have to account for each of these three cases.

Periodic Signaling Traffic

Many signaling messages are time critical and must be transmit immediately. Some signaling messages may be periodic and quite predictable.

Periodic signaling messages are interesting because the grants described by the BWR could be described like LTE describes semi-persistent scheduling (SPS) grants. For example, instead of transmitting 100x 60-byte grants for voice packets arriving in the UE’s transmit queue every 20 ms, LTE issues a single SPS grant to the UE for 100x 60-byte grants spaced out by 20 ms. As long as the signaling traffic is predictable and timely, a SPS version of the BWR could be used by the eNodeB to ask the CMTS for periodic grants of a given size. The benefit of the SPS-BWR would be a reduction in BWR messaging overhead.

6.7. BWR Prediction Error

The predictor algorithm that creates the BWR message at the eNodeB is constantly trying to predict the upstream traffic that will egress from the eNodeB several subframes in the future. Despite the BWR predictor’s best efforts, we explained in section 6.3 and section 6.4 that HARQ failures and packet segmentation can cause differences between the amount of expected data and the actual data the leaves the eNodeB. In addition, timing errors such as poor synchronization, DOCSIS grant jitter and packet egress jitter can cause the packets to miss the grants. The accuracy of the BWR predictor algorithm and error correction algorithm could be differentiating aspects for eNodeB and CMTS vendors.

For instance, on the eNodeB side, an in depth understanding of the eNodeB’s own software processes is a good start to developing a good BWR predictor algorithm that creates accurate BWRs. For example, knowing timing statistics about the LTE UL stack and Ethernet device can be helpful to predicting when packets will egress the eNodeB, if hard real-time code cannot be written to ensure packets egress at an exact time.

On the DOCSIS side, the CMTS can observe the predicted behavior and actual DOCSIS grant usage. The CMTS can then help to manage the error by either increasing the buffer delay in the CM or by doing active buffer management on the CM and sending extra grants when the buffer increases in size.

When errors do occur, the BWR itself can be used to address inaccuracies. For instance, when transport blocks fail to be decoded successfully, the eNodeB re-requests enough bytes for the future re-transmission.

The eNodeB predictor could be even more proactive and monitor the amount of data leaving the Ethernet device. When the actual upstream data differs from the predicted upstream data, the eNodeB could communicate this difference to the CMTS by subtracting the bytes from the TI it originally requested and add the bytes to a new TI.

6.8. BWR Message Flow

The BWR is a signaling message that gets sent to the Ethernet interface. Depending on how the queuing is done, it may need to be assigned a separate QoS level from data. If the BWR goes to a separate queue on the CM with a rtPS or unsolicited grant service (UGS) flow, then the byte count of the BWR message it does may need to be accounted for in the service flows that carry the EPS bearers and thus in the BWR message itself. If BWR is given high priority and placed on a common service flow, then it does have to be accounted for.

The LTE-DOCSIS system may have to account for the network bandwidth used by the BWR message. This will depend upon the CM queuing configuration.

6.9. BWR Message Delay

If the BWR message is predicting 8 ms in the future, then in theory it has 8 ms to complete its journey from the eNodeB to the CMTS and for the CMTS to issue a grant. In practice, the CMTS will need the BWR message before it runs its upstream scheduler and MAP builder routine. That MAP routine is run early depending on the MAP advance time that the CM requires. If the MAP interval is 2 ms and the MAP advance time is 2 ms, and there is some margin added, then the BWR might need to arrive 5 ms early. If the eNodeB takes 1 ms to generate and transmit the BWR, then that only leaves 2 ms of margin.

This means that it is important that when the BWR message arrives at the CM, the BWR is transmit to the CMTS as quickly as possible. Additional latencies, such as contention in a request slot, will cause the BWR to arrive at the CMTS too late. To avoid contention delays, DOCSIS scheduling algorithms such as rtPS or UGS can be used.

Since the BWR can be a variable length message, UGS may not be the best choice unless the size of the BWR can be fixed or the UGS grant is made larger or equal to the length of the longest BWR message. More advanced scheduling mechanisms may be needed at the CMTS such as waterfall granting which is under evaluation by the DOCSIS committee 0.

6.10. BWR Message Loss

One option for sending BWR messages over Ethernet and DOCSIS is UDP encapsulation. While UDP adds little overhead compared to TCP, UDP should not be considered reliable.

Improving BWR Reliability

To make the LTE-DOCSIS system more robust, the BWR could include a sequence number to help the CMTS detect a dropped BWR message. Note that if the eNodeB has no data to send, it may not send a BWR message to same processing time and DOCSIS bandwidth, for example. Therefore, a TI field alone could not be used to detect dropped BWR messages.

Many communications systems use retransmissions to ensure data is received at the receiver. In the case of the BWR, there is not enough time to detect BWR message loss and re-transmit the original BWR message since a margin of only 2 ms is left in the system as described in section 6.9, therefore re-transmission is not a viable solution.

A better solution may be to duplicate the requested bytes per TI and LCG within a BWR message over "N" BWR messages. The number of duplicate messages "N" can easily be configured at the eNodeB but the proposed value for N is three. The duplication of the requested bytes over three messages keeps the total bandwidth requirements low while helping to increase the reliability of the LTE-DOCSIS system.

6.11. DOCSIS Grant Duplication Due to CM Requests

With the introduction of BWR, there are now two request mechanisms – the external layer three BWR request mechanism and the native internal layer two CM request mechanism.

When data packets from the eNodeB arrive at the CM, the CM may generate a request message (REQ) if the grant from the BWR is not immediately available at CM. When the CM BE flow requests its own grant, two grants (REQ + BWR) may be generated for the same data which could lead to wasted bandwidth if those duplicate DOCSIS grants are not used.

The CM BE flow request behavior is dependent on the CM hardware implementations and may be hard to predict. Based on some observations done during Cisco-CableLabs R&D efforts, it appears that the CM may not issue a duplicate request if the grant from the BWR is scheduled and received by the CM within 1-2 ms after upstream data arrives at the CM.

In theory, the BWR message contains all the request information that is needed for the data path so the additional requests from the CM BE flow could be disabled or ignored. In practice, mismatches between the BWR's predicted bandwidth request and the actual upstream data from the eNodeB can occur, and the CM's BE flow's requests can be useful in managing buffer build-up.

The solution is to keep the CM's BE flow's requests enabled and ensure that the BWR is scheduled in a timely manner to prevent the CM from making duplicate requests. This means good synchronization and accurate prediction of when the data will leave the eNodeB and arrive at the CM. In addition, the CMTS may be able to help reduce duplicate messages by detecting the duplicate BE flow requests and ignoring the duplicates, if an intelligent CMTS algorithm can be developed for this purpose.

6.12. Checks and Balances

In order for the BWR concept to work well there needs to be a certain level of trust between the eNodeB and CMTS. For instance, the CMTS doesn't know how much data the eNodeB truly needs to allocate the UE(s) so it must trust the eNodeB - to a certain extent. A rogue eNodeB could intentionally request more bandwidth than it needed if it wanted to try to reduce its latency as much as possible at the expense of wasted DOCSIS grants and potentially other customer's QoS. A poorly designed BWR prediction algorithm could also put an increased burden on the DOCSIS upstream. Whether the BWR inaccuracy is intentional or un-intentional a smart CMTS BWR processing implementation will have checks to ensure the eNodeB is not abusing the BWR.

A simple method to ensure the BWR is not abused is to check if the upstream data flow matches the grants requested. If the eNodeB is using all the grants it asks for there is no reason for concern. On the other hand, a rogue eNodeB that is requesting twice as much bandwidth than it needs would only be using a fraction of the grants it asks for. The CMTS could flag the eNodeB by measuring the eNodeB's grant utilization and enforcing restrictions when poor utilization or other bad behavior is observed. Once an eNodeB is flagged, the CMTS could limit the maximum traffic or scale back the size of the grants dynamically based on the grant utilization of the eNodeB in question.

This puts some responsibility on the shoulders of the eNodeB's BWR prediction algorithm since it must predict the upstream egress data flow with reasonably good accuracy or else it will be bandwidth limited by the CMTS.

7. BWR QoS Considerations

This section will introduce some additional complexities related to the combined data path of the LTE and DOCSIS systems and then discuss how to holistically address these problems.

7.1. A Day in the Life of a Packet

In routing systems, there is a common development exercise called "a day in the life of a packet". The goal is to trace everywhere a packet goes in a system for a given flow. This will help find any data path irregularities that could impact the performance of that data path. Here, we will look at the flow of packets from a source in the UE to a destination just past the CMTS. Figure 8 will be used during this exercise.

The system diagram in Figure 8 depicts an LTE-DOCSIS network. The LTE radio access network (RAN) is connected to the evolved packet core (EPC) using DOCSIS. The LTE system is composed of the UE and the eNodeB which are each composed of their respective sub-systems. The DOCSIS system is composed of the CM and CMTS. For simplicity, this discussion will not go into the detail of the DOCSIS subsystems within the CM and CMTS.

The next thing to notice are the interfaces involved. They are:

- Air
- Ethernet (twisted pair or optical)
- HFC (coax and fiber)

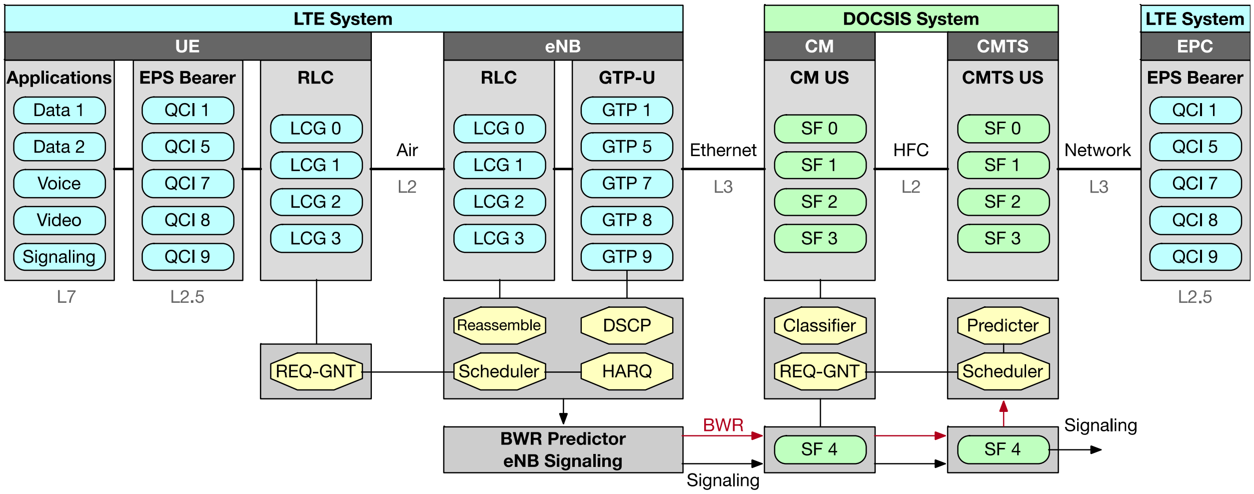

Figure 8 – A Comprehensive Look at the DOCSIS-LTE Interface

The air interface is internal to the LTE system and is a layer 2 interface. The HFC interface is internal to the DOCSIS system and is a layer 2 interface. The Ethernet interface, when combined with IP, is a layer 3 interface. More importantly, the classifiers that the CM uses to classify ingress traffic to a service flow are layer 3 classifiers.

That means that if we have more than one service flow, and hence more than one grant flow, the BWR must describe the request flow *per service flow*. This is a very important point. Service flows are defined by classifiers in the DOCSIS systems. The semantics of the BWR message has to match the semantics of the DOCSIS classifiers.

All this time we have referred to using the scheduling information from the eNB scheduler. That scheduler actually describes flow of segments across a LTE layer 2 construct called logical channel group (LCG). How does this map to something that a DOCSIS classifier can use? Before answering this, let’s review the life of a packet in the LTE and DOCSIS data path.

Starting on the left side of Figure 8, there are a number of applications that will initiate data flows. There can be any number of applications that span the classic realms of data, voice, video and signaling. Each of these generate TCP or UDP flows. The LTE system will map these applications into evolved packet system (EPS) bearers and will tag each of these flows with a QoS class indicator (QCI). EPS bearers are end-to-end bearers between the UE, the serving gateway (SGW), and the packet data network gateway (PGW) and are a collection of radio bearer (RB), S1 bearer and S5/S8 bearers. The S1 bearer is carried in an S1 GTP tunnel. Likewise, the S5/S8 bearer is carried in an S5/S8 GTP tunnel. A radio bearer transports the packets of an EPS bearer between the UE and an eNB. Radio bearers (RBs) are in a one-to-one association with logical channels. The SGW, PGW and the mobile management entity (MME) form the basis of the evolved packet core (EPC).

The logical channels are collectively mapped into LCGs to facilitate efficient use of bandwidth for UE bandwidth requests on the radio link. This is important as there is only a maximum of four LCGs but there can be up to 11 EPS bearers per UE. UE signaling is included in one of these LCGs, typically LCG0. The LTE request-grant exchange is based on LCGs, LCs or EPS bearer.

The segments are received at the eNB and assembled back into packets. The eNB will treat each QCI flow differently. This can result in different delays for different QCI streams. The eNB then maps each flow into a S1 GTP tunnel based upon the value of QCI. GTP stands for GPRS Tunneling Protocol and GPRS stands for General Packet Radio Service. Each GTP tunnel will be assigned a DiffServ Code Point (DSCP) for IP quality of service. *GTP is what the CM sees, not the LCG*. This DSCP will usually be copied from the inner DSCP to the outer DSCP. However, the network reserves the right to overwrite the DSCP if it needs to.

The DOCSIS CM will classify the incoming packets into service flows. Each service flow will receive a grant flow. *It is the data into each service flow that the BWR must accurately predict in order to get the right grant flow*.

On the control plane, we can see the eNB prediction algorithm that gathers information from the LTE scheduler, HARQ, the reassembly engine, and other status. It knows about not only LCGs, but also about GTP and DSCP assignments. But it cannot predict based on GTP or DSCP, because it does not have visibility into the LTE requests on the air interface on a per logical channel granularity.

It can also be seen that there are signaling flows for LTE signaling and BWR signaling. In this example, all signaling is going to a dedicated service flow.

7.2. BWR Flow Types

In the previous section, we described the data path from the point of view of a packet that starts at the UE and eventually goes past the CMTS. That packet traverses multiple layer 2 and layer 3 paths. The ultimate realization is that the BWR must talk in a language or syntax that the DOCSIS system can understand. As such, BWR is really an API into the DOCSIS system for requesting bandwidth in the future.

This section will describe four fundamental ways of describing the data flow from the eNB. They are:

1. Based on Bulk
2. Based on LCG
3. Based on GTP
4. Based on DSCP

The methods are ordered to provide increasingly finer granularity of QoS for different types of traffic. We will look at how each method works, where it would be used, its strength, and its weaknesses.

7.2.1. BWR Using Bulk

In this method, all bytes being transferred for a given future point in time are summed together. This would include both data path and signaling bytes.

A variant of the Bulk Method would be to separate BWR out into its own flow.

The advantage of this method is its simplicity. It is useful when there is ample bandwidth and QoS is not required. Or, if the CMTS can accept all traffic on one Service Flow and sort out QoS on its own, say from IP DSCPs, then this becomes a simple interface. The Bulk Flow could be a good match for the DOCSIS Waterfall Granting proposal 0.

The disadvantage of this method is that there is no specific Quality of Service information in the BWR message. As such, the CMTS may not be able to provide lower latency for higher QoS traffic such as signaling.

7.2.2. BWR Using LCG Data Flows

In this method, BWR would be based on the four flows used on the eNB air interface. These are LCG 0 through LCG 3. BWR would describe how many bytes are required for each flow in each future time interval (say 1 ms).

The CM has at least two choices on how to receive these flows.

The first choice uses five service flows – one to match each of the LCGs and one for BWR. Signaling that is not already accounted for in LCG0 could join BWR in its queue. The mapping of resource blocks (RBs) (or LCs) to an LCG is done at the RB setup time by the eNB based on the corresponding QoS attributes of the RBs such as QCI. This mapping is configured by the mobile operator via a provisioning system. If the LTE provisioning system is able to share provisioning information with the DOCSIS system, then it should be possible to know which IP flows are assigned to each LCG. The provisioning system would program into the DOCSIS system the necessary classifiers to recreate four service flows that exactly match the four logical channel groups.

The second choice is two service flows, one for all the LCGs and one for BWR. This second method simplifies the configuration and classification required on the system. It can work well if the CM and CMTS can manage QoS within a Service Flow 0.

A variant of this approach is for the BWR to request per LCG. This will allow the CMTS to make granting decisions based on packet priorities. However, the CMTS would combine all the grant bytes into one service flow and let the CM use the DOCSIS Waterfall Granting proposal 0 to properly allocate the bytes to queues. The UE and eNB use a similar approach for requesting and granting – the UE BSR requests are per LCG, the eNB DCI-0 is a bulk grant, and the UE assignment of packets to LCG flows can be different at grant time that the original BSR intent.

The advantage of this approach is that it is simple and transparent. The external system can see what is really going on.

The disadvantage of this method it that it can require more system configuration for classifiers.

7.2.3. BWR Using GTP Data Flows

From an IP network viewpoint, this method is connection-oriented. That means that the network contains state information about specific IP flows. Qo S is achieved by assigning bandwidth to these flows.

In this method, the BWR message expresses the number of bytes sent per GTP flow. The CM would be configured with classifiers that would forward GTP flows onto specific service flows. One or more GTP flows could be combined into a single Service flow, but BWR would still report per GTP. Note that the DOCSIS classifier cannot classify on the QCI field of the GTP flow, so the configuration system would have to create a DOCSIS classifier using the IP tuple (Dest IP, Source IP, Dest Port, Source Port). This information could be derived from the LTE Policy and Charging Rules Function (PCRF) policy system.

Note that LTE signaling messages are be sent on a separate protocol such as the stream control transmission protocol (SCTP). BWR would describes the signaling bytes separately.

The advantage of this method is that BWR is describing layer 3 traffic flows instead of internal layer 2 flows. These layer three flows match nicely with DOCSIS classifiers.

The disadvantage of this method is that the provisioning system has to supply precise classifiers. If there is any change in configuration in the LTE system, the DOCSIS system would need updated classifiers. This is achievable, but it does require excellent coordination between the LTE and DOCSIS provisioning systems.

The challenge with this method is that the eNB does not have a proper byte accounting per GTP flow. It only has a byte accounting per LCG flow. Thus, this option may not be implementable in a current LTE system. However, it is included in this proposal in case future mobile protocols or other systems external to a DOCSIS system can use this option.

7.2.4. BWR Using DSCP Data Flow

From an IP network point of view, this method is connectionless. That means that the network does not contain state information about specific IP flows. QoS is achieved by assigning a tag called a DSCP to each packet and using that tag to manage network queuing.

In this method, the BWR message lists the number of bytes per DSCP flow. Usually, each GTP tunnel will have its own DSCP, although one or more GTP tunnels could use the same DSCP. BWR and signaling should have their own DSCP value. This is a very IP centric, connectionless approach, and is a common way of managing QoS on the Internet.

The advantage of this method is that it is very simple. A generic set of classifier rule can be provided to the CMTS at configuration time that is independent of LTE connection information. The LTE system can have any connection it wants. The CMTS will focus on honoring the QoS Policy associated with the DSCP.

The disadvantage of this method is that visibility per GTP flow in the BWR message is lost.

Again, the challenge with this method is that the eNB does not have a proper byte accounting per DSCP flow. It only has a byte accounting per LCG flow. Thus, this option may not be implementable in a current LTE system. However, it is included in this proposal in case future mobile protocols or other systems external to a DOCSIS system can use this option.

7.3. BWR Message Format

This section describes the proposed information elements that will compose a BWR message. These information elements are subject to change as BWR becomes standardized. The actual BWR message format is not specified here as that will be determined by a standards committee.

Table 2 – BWR Message Elements

Information Element	Size	Description
BWR Header		
Version number	3 bits	Version number of BWR message
Sequence number	5 bits	Used for detecting dropped messages
Device Identifier	48 bits	Device unique identifier
1588v2 Reference seconds	48 bits	Value corresponding to the TI reference value.
1588v2 Reference nanoseconds	32 bits	
Time Index (TI) Reference	16 bits	Value corresponding the 1588v2 reference value.
Time Index Increment	32 bits	Total nanoseconds between sequential TI values.
BWR Flow Type	3 bits	Bulk, LCG, GTP, DSCP
BWR Table Size	5 bits	The number of row entries.
BWR Table		
Time Index (TI)	16 bits	Time index that the row refers to
The TI is listed once at the beginning of each row entry. *The following entries are repeated as a block within a row entry.*		
Last Entry	1 bit	This is the last entry in the current row.
Reserved	5 bits	Not defined at this time.
Operation	2 bits	New – This is a new entry for the bytes Subtract – remove bytes in specified TI Add – add back previously bytes to new TI.
BWR Flow ID	8 bits	For Bulk, this value is 0x00 For LCG, this is 0 through 3 For GTP/QCI, this is 0 through 192 For DSCP, this is the six-bit value 0xFF indicates signaling
Bytes Requested	32 bits	These are the bytes per BWR flow per time interval that cross the Ethernet Interface. The bytes are counted based upon a complete Ethernet frame including the CRC.

The BWR message elements are shown in Table 2. There is a BWR header which contains single entries and a body with the BWR Table.

The BWR message starts with a version field. It is reasonable to expect an updated version for 5G. The sequence field is to allow the receiving entity to detect dropped messages. The sequence number also doubles as a transaction ID. There is no expectation of re-ordering of BWR messages. Note that the eNB may not send a BWR message if there are no traffic updates.

The device identifier field is included for convenience. This allows the CMTS to manage and distinguish among multiple eNB clients and to report statistics with an externally known tag such as the Global eNB identifier which is the combination of PLMN (Public Land Mobile Network) + eNB ID within PLMN. The PLMN is 6 digits: 3 integer digits for mobile country code (MCC), 3 integer digits for a network code (MNC). (24 bits). The eNB identifier is defined as equal to the 20 leftmost bits of the Cell Identity Information Element (IE) as standardized by 3GPP of each cell served by the eNB. An alternate device identifier could be a 48 bit MAC address. A value of all zeros is considered as no identifier supplied.

The BWR message will be sent periodically and will describe repetitive units of bandwidth. Typically, the BWR message will be sent every one millisecond and will describe a one millisecond unit of bandwidth that corresponds to an LTE sub-frame. The Time Index field would count LTE subframes and would correspond to a 1588v2 timestamp. This is done to make BWR more readable and to create a transmission window. It can also help the keep BWR message length constrained when the BWR Table gets large. The BWR header contains a cross reference to a 1588v2 timestamp which point to the beginning of the Time Index frame.

In use, the TI at the time the message is created might be X while the TI that bytes are being requested for might be X+8.

There is an entry to indicate what type of flows that BWR is describing – Bulk, LCG, GTP or DSCP. This will dictate what the BWR Flow ID represents. There are some variations on this theme that are possible. For example, bulk and LCG modes could be without much prediction while GTP and DSCP modes would have prediction. Or, prediction could be an option on all modes. This would allow an external prediction mechanism which might be interesting in a cloud implementation.

There is an entry that indicates the number of rows on the BWR Table. Note that the rows can be variable length as flows with nothing to report do not have to be included in the message.

The BWR Table consists of a series of rows. Each row has a time index. If there is no redundancy and no error vectors, there is just the one row entry for the future time index. If the system is configured for redundancy of three, then each row is repeated across three messages.

The BWR operation field allows for New, Subtract, and Add. The New operation is for requesting new bytes for the specified time interval. The Subtract operation is to indicate that the bytes for a predicted TI did not happen. An example of this is a HARQ operation where bytes where supposed to show up but did not. The Add operation is like New except that it is referring to bytes that were previously accounted for. If there are both New bytes and Add bytes for the same flow in the same TI, the bytes would be listed separately. The Subtract and Add operations represent an error vector from the eNB to the CMTS. The CMTS decides what to do with the information.

The BWR Flow ID directly uses the native IDs from LCG, GTP(QCI), and DSCP. A reserved value in this field denotes bandwidth requested for signaling. This would apply for LCG and GTP modes. Bulk and DSCP modes can include signaling within their data flows.

To allow for a variable length row where null flows are not included and there may be separate New and Add entries for a given flow, a last entry marker exists. Then, there is the number of bytes requested for a given flow within a given time interval. A 32-bit field for a one millisecond time interval will support over flows of more than 20 Mbps.

If an eNB uses multiple sectors and/or MIMO channels, each with its own LCG domain, the eNB should collect all byte requests from all sectors and combine it into one BWR message.

A typical BWR message would have an Ethernet/IPv4/UDP overhead of 46 bytes, a BWR header of 24 bytes, a BWR table with three rows (for redundancy) and four flows (78 bytes), for a total of 138 bytes. IPv6 would bring this to 158 bytes.

7.4. Implementation Impacts

If BWR is generated every millisecond, then the eNB has to generate 1000 messages a second that it did not have to do before. This may require additional CPU resources.

To compare actual versus predicted, the eNB would have to monitor its own Ethernet output and timestamp packets in real time. It then could compare off-line the time alignment of the actual packet to the predicted values. This may have some HW impact.

The BWR packet needs an IP destination address (DA) that will cause BWR to be forwarded to the CMTS US Scheduler process. The IP DA could be a well-known IP address, one per CMTS, or one per CMTS scheduler instance. The eNB will need to be configured with this address. This will require some system configuration.

The CMTS has to create a routed path for BWR from an external entity to the scheduler instance.

The CMTS scheduler process will have to receive 1000 BWR message a second *per eNB* that is connected. If there are two eNB per DOCSIS service group (SG) and 32 SG per line card, that would be 64,000 extra messages per second that it is not receiving now. If the CM request flow is disabled for the queues under use, that can help.

Each BWR flow could be on the order of 1.25 Mbps of traffic. If there are a lot of eNBs in a CMTS domain, this bandwidth could be noticeable. There is also the impact of the additional traffic to the control plane in the CMTS.

The CMTS may want to check the percentage of unused grants for a eNB. Unused grants could be a result of improper prediction or over requesting. Results of a grant audit can be combined with the error vector from the eNB to generate statistics.

7.5. Testing out the System

The authors and their respective teams have both simulated and built a system to test out these fundamental theories. Both the MatLab and the physical test environment contain a UE and eNB connected to a CM and CMTS. The test results showed that the desired decrease in latency due to having a BWR interface between the two systems worked.

These test results are being published in a series of IEEE papers 000. The test results showed a decrease in latency from 5 ms minimum to the 2 ms of engineering latency. Also, under heavy load, the overall traffic latency was reduced. Part of that latency reduction under network load can be attributed to BWR not being sent in a contention slot. Figure 9 shows a cumulative distribution function (CDF) of packet delays from the CM to CMTS with and without BWR.

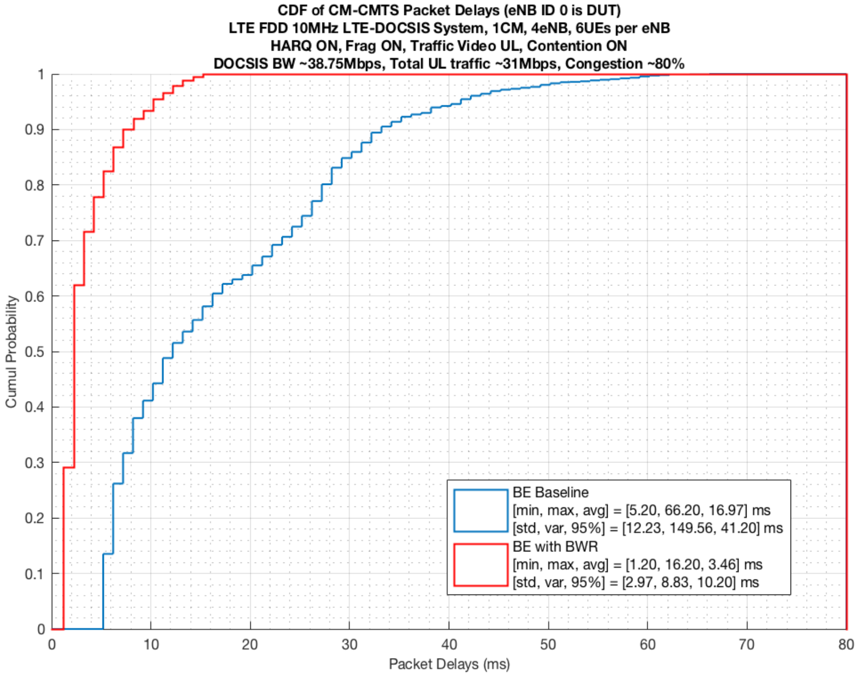

Figure 9 – CDF of CM-CMTS Packet Delays with/without BWR

Extending BWR to Future Platforms

8. Split Small Cell and Cloud CMTS

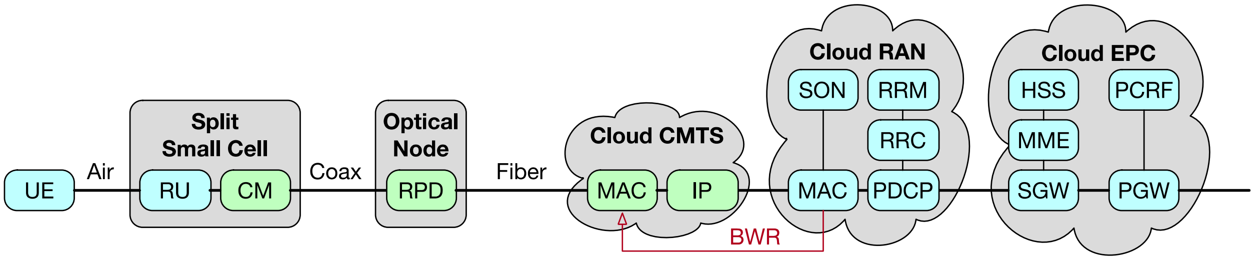

Figure 10 – BWR with vRAN and Cloud CMTS

In the previous sections, the architecture consisted of a physical integrated CMTS and a physical integrated small cell. There are variations of this architecture for which BWR will also work. The CMTS could be an integrated CMTS, a physical CMTS-Core with Remote PHY, or a Cloud CMTS with Remote PHY. The Remote PHY architecture uses a Remote PHY Device (RPD).

The eNB could be integrated or one of eight proposed split configurations 0 0. The split eNB is often referred to as a virtual radio access network (vRAN) or more recently Cloud RAN to reflect updated software architectures. The diagram in Figure 10 illustrates one specific combined scenario known as Network Functional Application Platform Interface (nFAPI) 0, with a Cloud CMTS with a Cloud RAN where the eNB scheduler is located in the cloud. nFAPI is a MAC-PHY split similar in concept to DOCSIS Remote PHY. It can be seen that the BWR message is now passed between software processes within the cloud. The latency for getting BWR from the eNB to the CMTS is greatly reduced now that BWR does not have to traverse up the CMTS access. As a result, the BWR message could be sent less often.

BWR will also not need to include any bandwidth estimation for signaling that would originate and terminate in the cloud. Further, the X2 traffic from eNB to eNB might also be exclusively in the cloud and thus also achieve low latency.

The migration to vRAN does modify the scope of the DOCSIS data path requirements from backhaul to fronthaul. In a backhaul scenario with an integrated eNB, the data packets are fully formed packets and are a tunneled version of the original packets from the original application. In a fronthaul scenario, the content is an interim form that originates from some midpoint in the eNB processing chain.

The ideal fronthaul technology would be a transport that has QoS attributes attached to the data flow. The reason for this is that all transport networks are multi-hop (meaning one or more routers) and bandwidth is typically queued and aggregated. A device like an eNB will have a traffic profile that will contain high priority and lower priority traffic. The task for the network is to maintain that traffic profile through the network.

With a split small cell, where the signaling messages between the UE and the eNB now travel to the cloud, QoS will be required to provide a low latency path for UE signaling. QoS works best with high priority, low bandwidth flows.

From a network QoS viewpoint, a fronthaul like the Common Public Radio Interface (CPRI/eCPRI) 0 where all traffic is equal would not be an optimum choice, whereas a front haul like the network functional application platform interface (nFAPI) 0 would be a good choice.

Conclusion

There were at least three fundamental requirements for DOCSIS to provide mobile backhaul or front haul services: timing, bandwidth and latency. Timing and bandwidth are problems addressed in other white papers. In this white paper, it was described that by coupling the scheduling mechanism from the LTE eNB uplink scheduler to the DOCSIS CMTS upstream scheduler, in effect creating a pipeline that allows for earlier requesting on the DOCSIS system, low latency on the DOCSIS system can be achieved.

The bandwidth report (BWR) message is a proposed API on the DOCSIS system that will allow external systems such as small cells to request bandwidth in advance of when it is needed, thus achieving a low latency transport.

The eNB receives traffic from UEs. In addition, it has its own signaling. The collection of all those bytes and the significance of each byte stream is known as a traffic profile. The eNB needs to place its traffic onto a network. In this case, the network is a DOCSIS system, but even if it was an Ethernet backhaul, that network needs to understand the profile of the traffic from the eNB.

The most popular form of profiling traffic on an IP network today is with IP Differentiated Services (DiffServ). Hence, it is the authors recommendation that the eNB should always properly mark each outgoing IP packet with the appropriate DSCP. The CM classifiers should be setup with DSCP, and thus BWR would be best used by requesting per DSCP. DiffServ techniques also tend to be stateless which means there is less network configuration to worry about. DiffServ is also queuing orientated, and all network ports are queue orientated.

In conclusion, by having LTE and DOCSIS work together, DOCSIS is a viable backhaul mechanism for gigabit-per-second mobile traffic.

Acknowledgements

Many engineers provided software implementation support and technical insights throughout this project. The authors would like to thank Michel Chauvin, Joey Padden, Aaron Quinto, Balkan Kecicioglu, and Vaibhav Singh of CableLabs, and Elias Chavarria Reyes, Zheng Lu, Dantong Liu, Alon Bernstein, and Oliver Bull of Cisco for their contributions.

Abbreviations

ACK	Acknowledgement
AM	Acknowledged Mode
ARQ	Automatic Repeat Request
BE	Best Effort
BSR	Buffer Status Report
BWR	Bandwidth Report
CM	Cable Modem
CMTS	Cable Modem Termination System
CPRI	Common Public Radio Interface
DiffServ	Differentiated Services
DCI-0	Downlink Control Information Format 0
DL	Downlink
DOCSIS	Data over Cable System Interface Specification
DSCP	DiffServ Code Point
eNB	Evolved Node B
EPS	Evolved Packet System
EPC	Evolved Packet Core
FDD	Frequency Division Duplex
FEC	Forward Error Correction
GPRS	General Packet Radio Service
GTP	GPRS Tunneling Protocol
HARQ	Hybrid Automatic Repeat Request
HFC	Hybrid Fiber-Coax
HSS	Home Subscriber Server
IP	Internet Protocol
LCG	Logical Channel Group
LTE	Long Term Evolution
MAC	Media Access Control
MCC	Mobile Country Code
MIMO	Multiple Input Multiple Output
MME	Mobility Management Entity
MMM	MAC Management Message
MNO	Mobile Network Operator
MVNO	Mobile Virtual Network Operator
NACK	Negative Acknowledgement
nFAPI	Network Functional Application Platform Interface
PCRF	Policy and Charging Rules Function
PDCCH	Physical Downlink Control Channel
PDCP	Packet Data Convergence Protocol
PDN	Packet Data Network
PGW	PDN Gateway
PHY	Physical Layer
PLMN	Public Land Mobile Network
QCI	QoS Class Indicator

QoS	Quality of Service
RAN	Radio Access Network
RB	Resource Blocks
RLC	Radio Link Control
RPD	Remote PHY Device
RRC	Radio Resource Control
RRM	Radio Resource Management
rtPS	Real Time Polling Service
RU	Radio Unit
SID	Service Identifier
SF	Service Flow
SGW	Serving Gateway
SON	Self-Organizing Network
SPS	Semi-Persistent Scheduling
SR	Scheduling Request
TCP	Transport Control Protocol
TDD	Time Division Duplex
TI	Time Index
UDP	User Datagram Protocol
UE	User Equipment
UGS	Unsolicited Grant Service
UL	Uplink
UM	Unacknowledged Mode
VoLTE	Voice over LTE
vRAN	Virtual Radio Access Network

Bibliography & References

Jennifer Andreoli-Fang, John T Chapman, "Synchronization for Mobile Backhaul over DOCSIS," Proceedings of SCTE Fall Technical Forum, October, 2017

John T. Chapman, et. al., "The DOCSIS Timing Protocol (DTP), Generating precision timing services from a DOCSIS system," Proceedings of INTX/SCTE Spring Technical Forum, 2011.

"Data-over-Cable Service Interface Scable service interface pecifications, DOCSIS® 3.1, MAC and upper layer protocols interface specification," I-11, issued 2017-05-10, CableLabs. < MULPI Specification, Figure on req-gnt delay>.

John T. Chapman, G. White, H. Jin, "Impact of CCAP to CM Distance in a Remote PHY Architecture," Proceedings of the INTX/NCTA Technical Forum, April 2015.

John T. Chapman, Tong Liu, "Waterfall Granting," Submitted to the CableLabs DOCSIS MAC committee, June 2017.

Jennifer Andreoli-Fang, John T. Chapman, "Latency Reduction for Mobile Backhaul by Pipelining LTE and DOCSIS," to appear in *Proc. of IEEE Globecom*, December 2017, Singapore.

Jennifer Andreoli-Fang, John T. Chapman, "Mobile Aware Scheduling for Low Latency Backhaul over DOCSIS," to appear in *Proc. of IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, October 2017, Montréal, Québec.

Jennifer Andreoli-Fang, John T. Chapman, et. al "Real-Time Experimentaiton with Mobile-Aware Scheduling for Low Latency Backhaul over DOCSIS", submitted to IEEE International Conference on Computer Communications (Infocom) 2018.

R3-161687, Draft TR 38.801 (v030) Study on New Radio Access Technology: Radio Access Architecture and Interfaces, NTT DOCOMO, INC (Rapporteur), 3GPP TSG RAN3, August 2016.

Small Cell Forum Document 159.07.02, Small cell virtualization functional splits and use cases, Small Cell Forum, January 2016.

"Common Public Radio Interface (CPRI); Interface Specification", Release V7.0, www.CPRI.info, Oct 9, 2015.

"FAPI and nFAPI Specifications", Small Cell Forum, Release 9.0, Document 082.09.05 , May 2017.