CableLabs®

CableLabs File No. 60229
(assigned by Legal dept.)

INVENTION DISCLOSURE

1. **Invention Title. Give a short** descriptive title of the invention in the box below (in 10 words or less).

System and Methodology to provide Multi-Gigabit Broadband Network Services

2. **Invention Summary.** Give a concise (30 words or less) summary of the invention.
This medium access control and physical layer system and methodology provides multi-gigabit service capabilities over HFC and other networks through a novel scheme that results in low complexity processes.

3. **Invention Description.**

a. Describe the invention in detail and/or attach a description, drawing(s) and/or diagram(s), if available. **Please include flow charts for descriptions of software processes, and block diagrams for descriptions of hardware systems**. Include the description/attachments in electronic form if possible.
This invention describes a broadband networking system that provides multi-gigabit per second services over infrastructures such as the hybrid fiber coax (HFC) networks used in CATV environments. This full-duplex system transmits and receives using one or more frequency blocks that can be configured to operate over multiple frequency allocation scenarios. In this system, the MAC layer is unaware of the PHY layer configuration and operation details which allow the selection of different PHY layer designs. The PHY layer mechanism is designed with the necessary robustness to present a constant and stable amount of resources to the MAC layer. This approach limits the number of PHY layer configuration options which leads to low complexity, low cost and easy to scale implementations. The home network is decoupled from the access network. This reduces the number of required end-stations and improves the physical layer environment. A single scheduling algorithm, coupled with a bandwidth request mechanism, provides fast network access and ensures support for quality of service while keeping implementation complexity low.

The documents entitled "PHY Layer Design for Efficient Multi-Gigabit Transport over CATV Networks" and "MAC Layer Design for Efficient Multi-Gigabit Transport over CATV Networks " describe this system in detail using an OFDM physical layer example that meets the system attributes described here (Documents attached).

b. Why was the invention developed? What problem(s) does the invention solve? How is it better?

The invention was developed to be able to extend the life of the HFC network while using their resources more effectively in order to provide multi-gigabit per second services in a cost effective manner. There may be a limitation to what DOCSIS systems can scale to in a cost effective manner. This approach is an alternative that may achieve higher peak rates at lower cost per bit compared to DOCSIS.

c. Briefly outline the potential commercial value and customers of the invention.

This invention could result in its use for transport of all services and applications across all CATV networks around the world. Interested parties include cable operators, data networking equipment manufacturers, consumer electronics manufacturers, fiber-optic network equipment vendors and many others.

4. HOW is your invention different from existing products, processes, systems?

Please list the closest publication(s), product(s), method(s), patent(s), etc. to your invention. For each item, how is your invention different?

We are not aware of similar systems, patents or products.

HFC Network Impairment Characterization

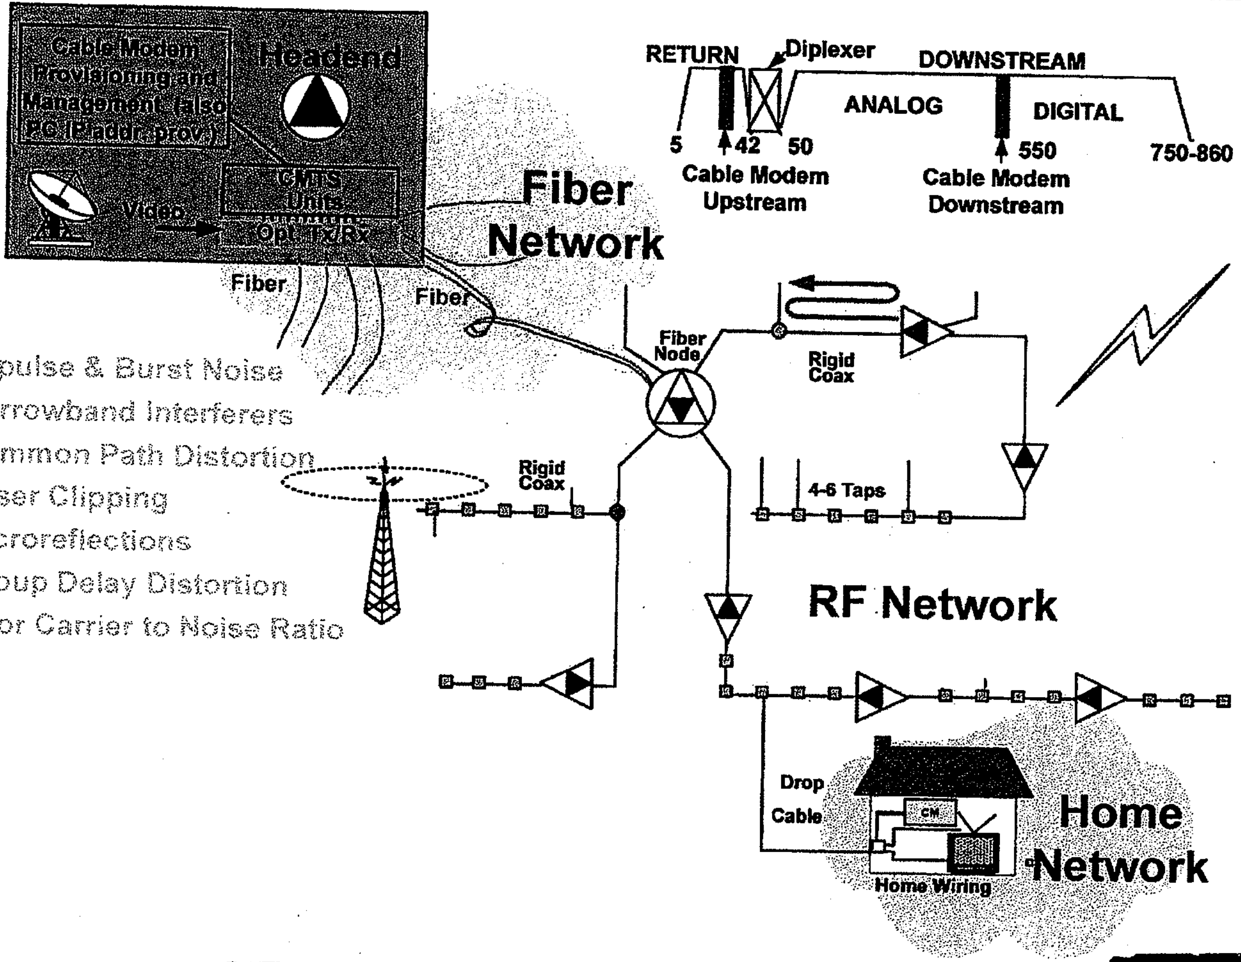

- Impulse & Burst Noise
- Narrowband Interferers
- Common Path Distortion
- Laser Clipping
- Microreflections
- Group Delay Distortion
- Poor Carrier to Noise Ratio

[redacted] [redacted] [redacted]

HFC Network Impairment Characterization

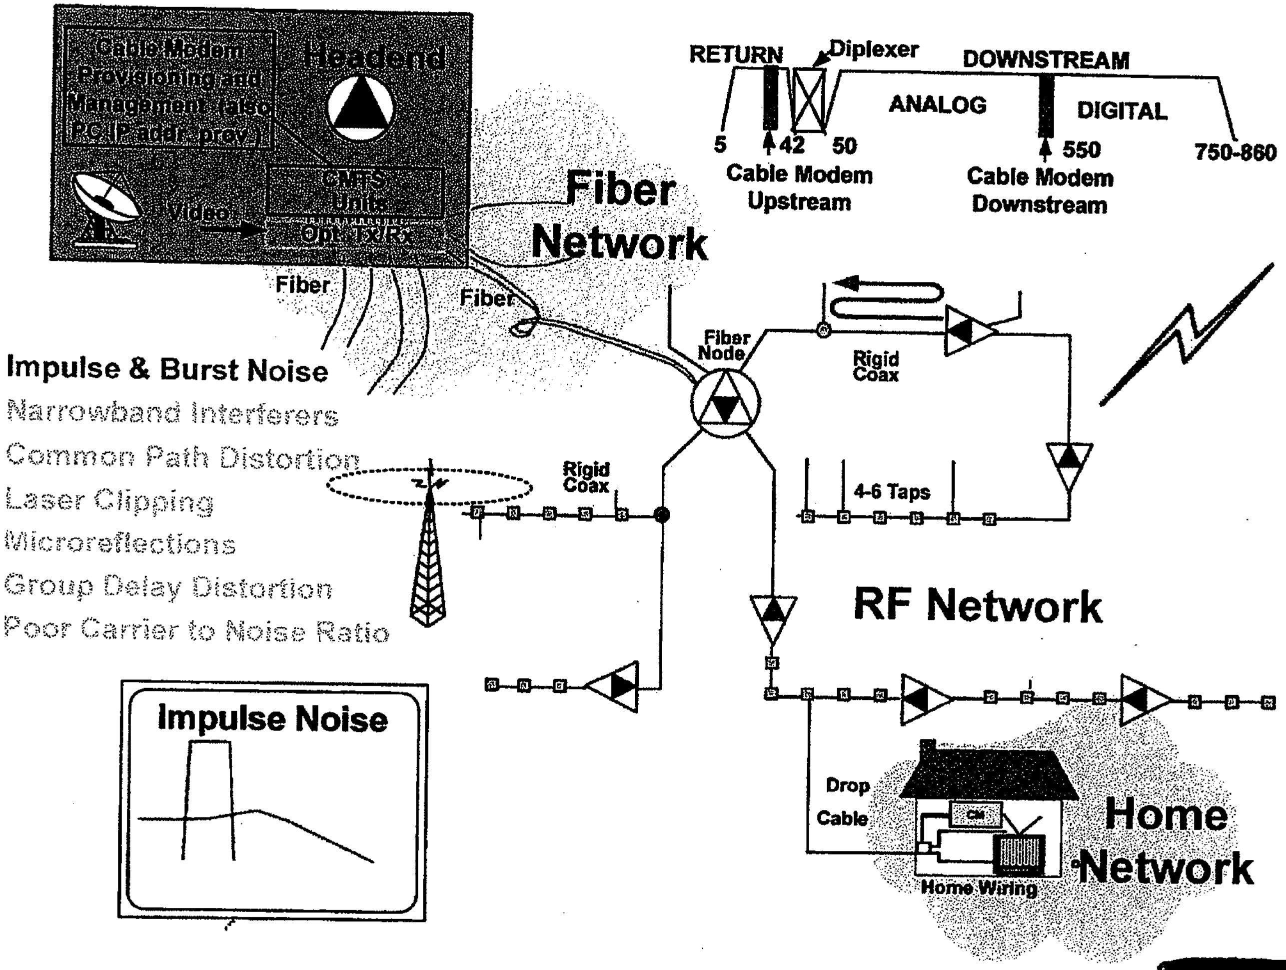

Impulse & Burst Noise
- Narrowband Interferers
- Common Path Distortion
- Laser Clipping
- Microreflections
- Group Delay Distortion
- Poor Carrier to Noise Ratio

[redacted]

[redacted]

HFC Network Impairment Characterization

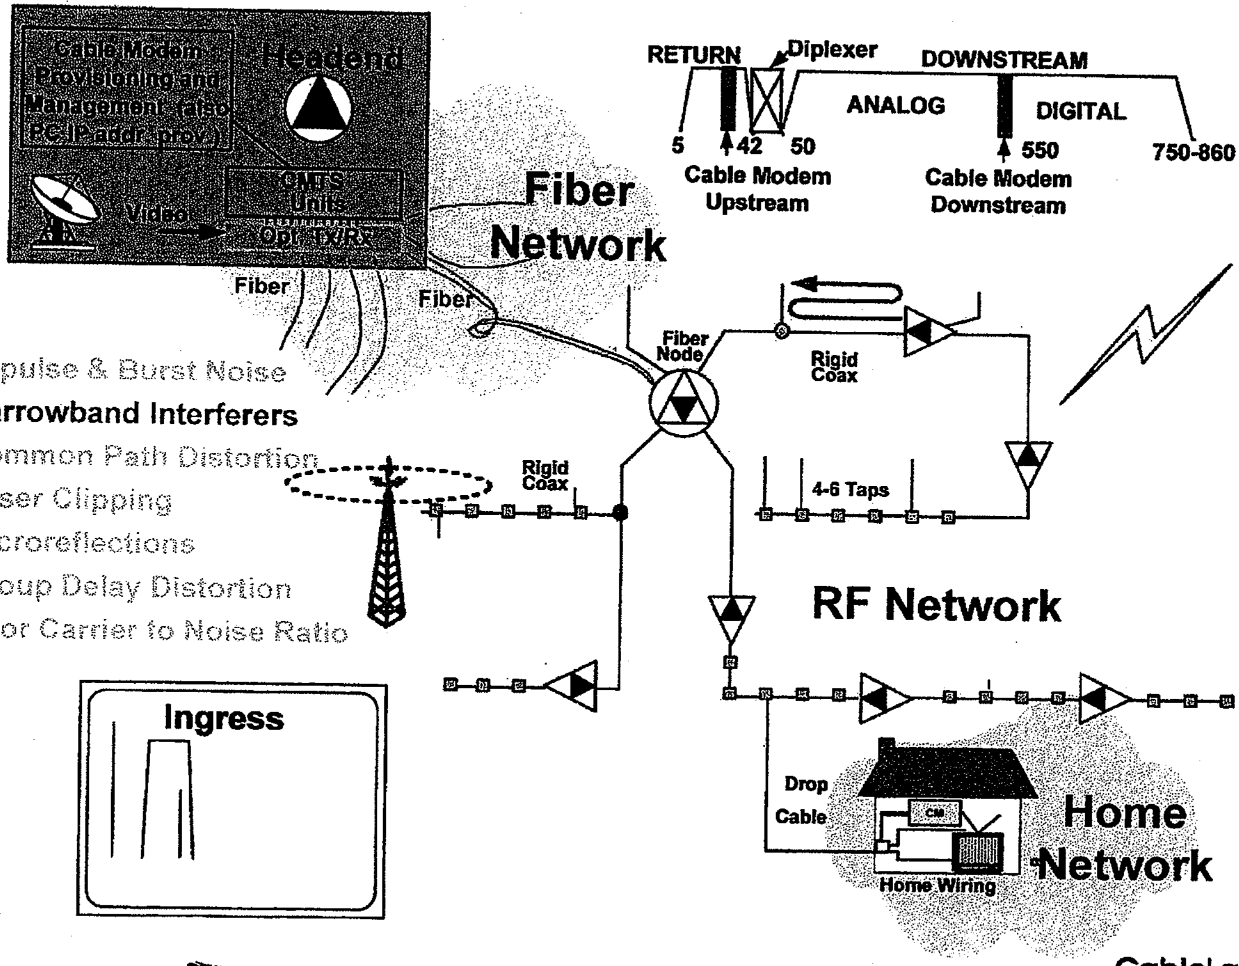

Impulse & Burst Noise

Narrowband Interferers

Common Path Distortion

Laser Clipping

Microreflections

Group Delay Distortion

Poor Carrier to Noise Ratio

[redacted]

CableLabs®
...Revolutionizing Cable Technology®

HFC Network Impairment Characterization

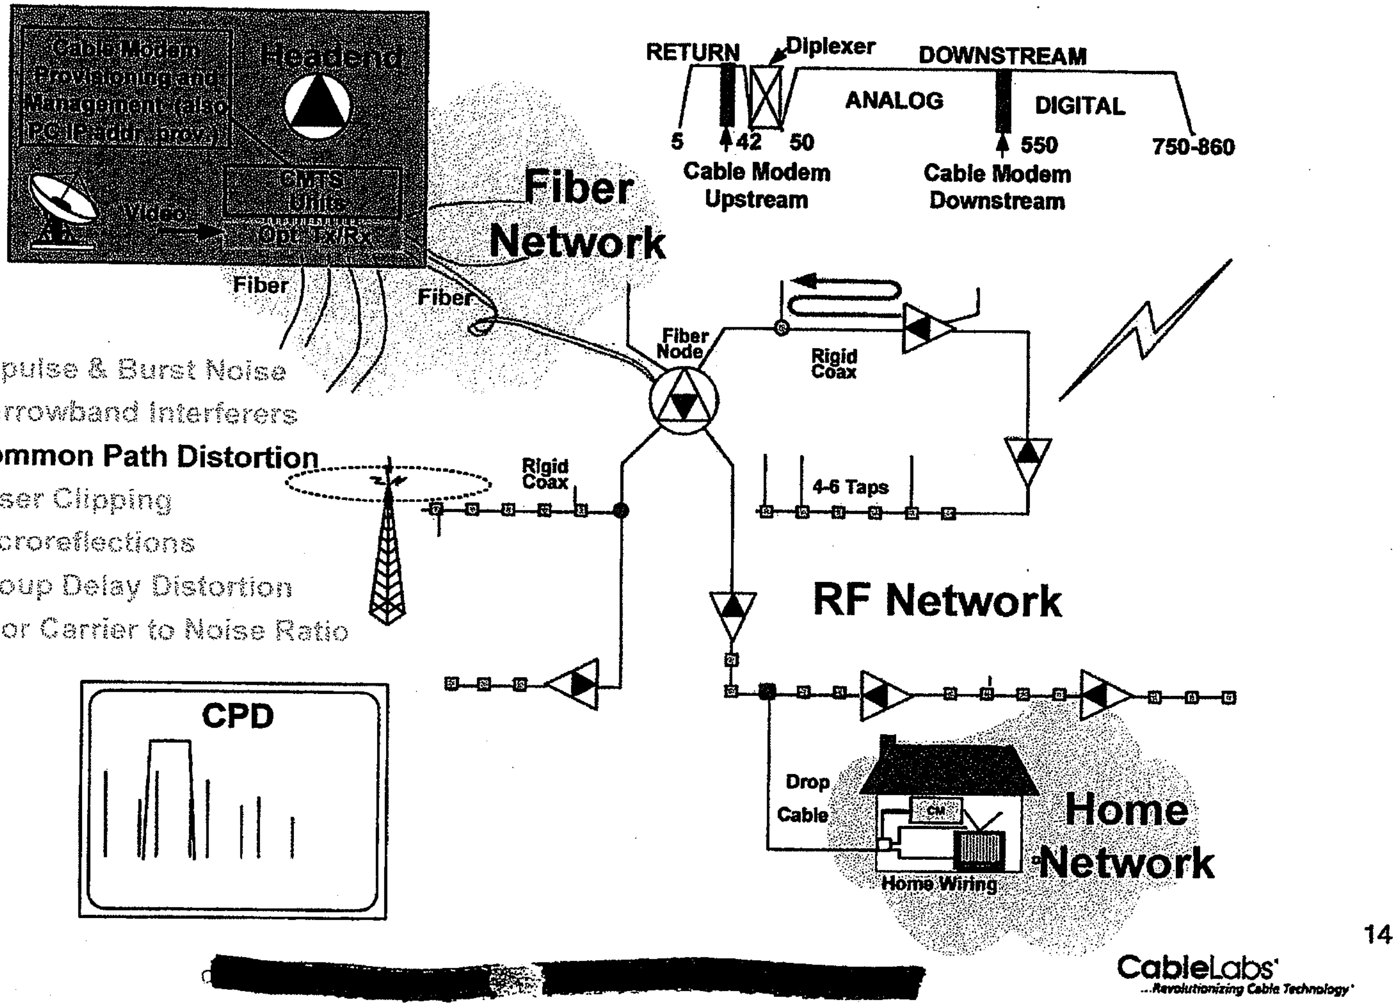

- Impulse & Burst Noise
- Narrowband Interferers
- **Common Path Distortion**
- Laser Clipping
- Microreflections
- Group Delay Distortion
- Poor Carrier to Noise Ratio

HFC Network Impairment Characterization

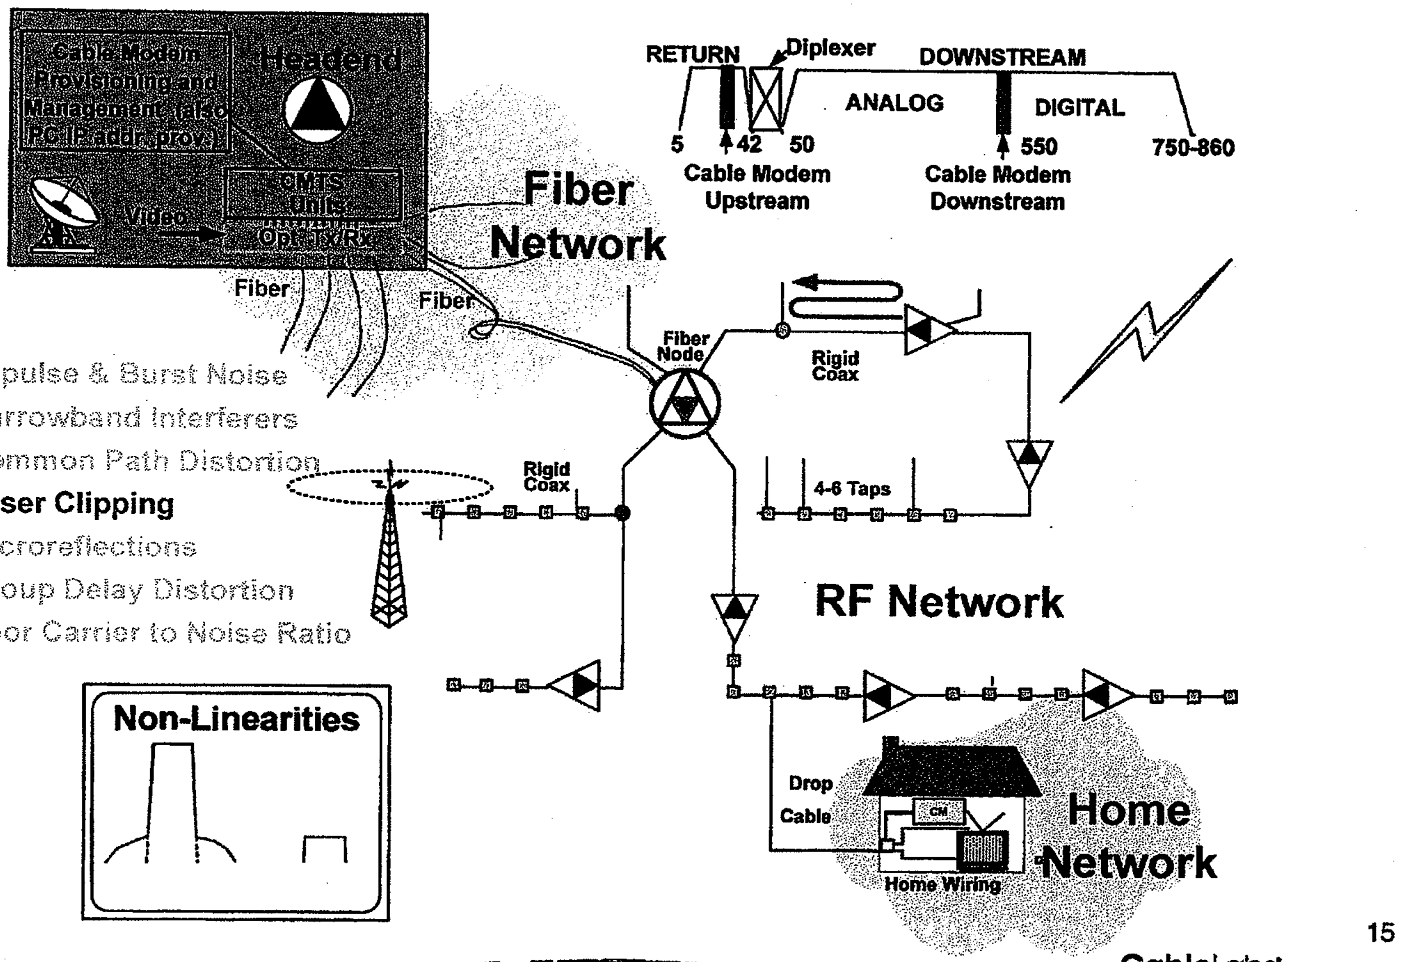

- Impulse & Burst Noise
- Narrowband Interferers
- Common Path Distortion
- **Laser Clipping**
- Microreflections
- Group Delay Distortion
- Poor Carrier to Noise Ratio

[redacted]

CableLabs®
...Revolutionizing Cable Technology®

HFC Network Impairment Characterization

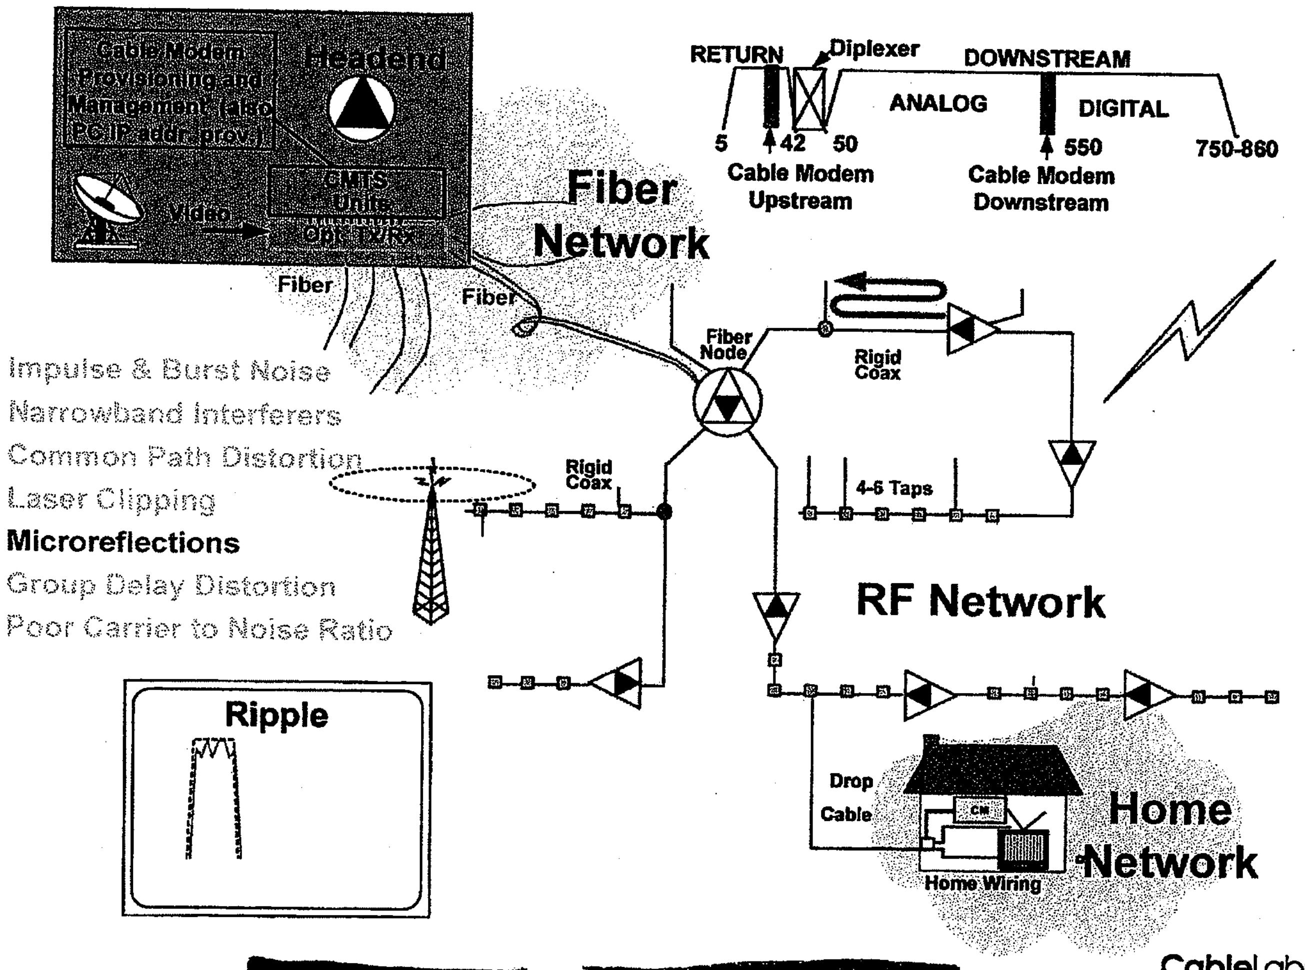

Impulse & Burst Noise
Narrowband Interferers
Common Path Distortion
Laser Clipping
Microreflections
Group Delay Distortion
Poor Carrier to Noise Ratio

[redacted] [redacted]

HFC Network Impairment Characterization

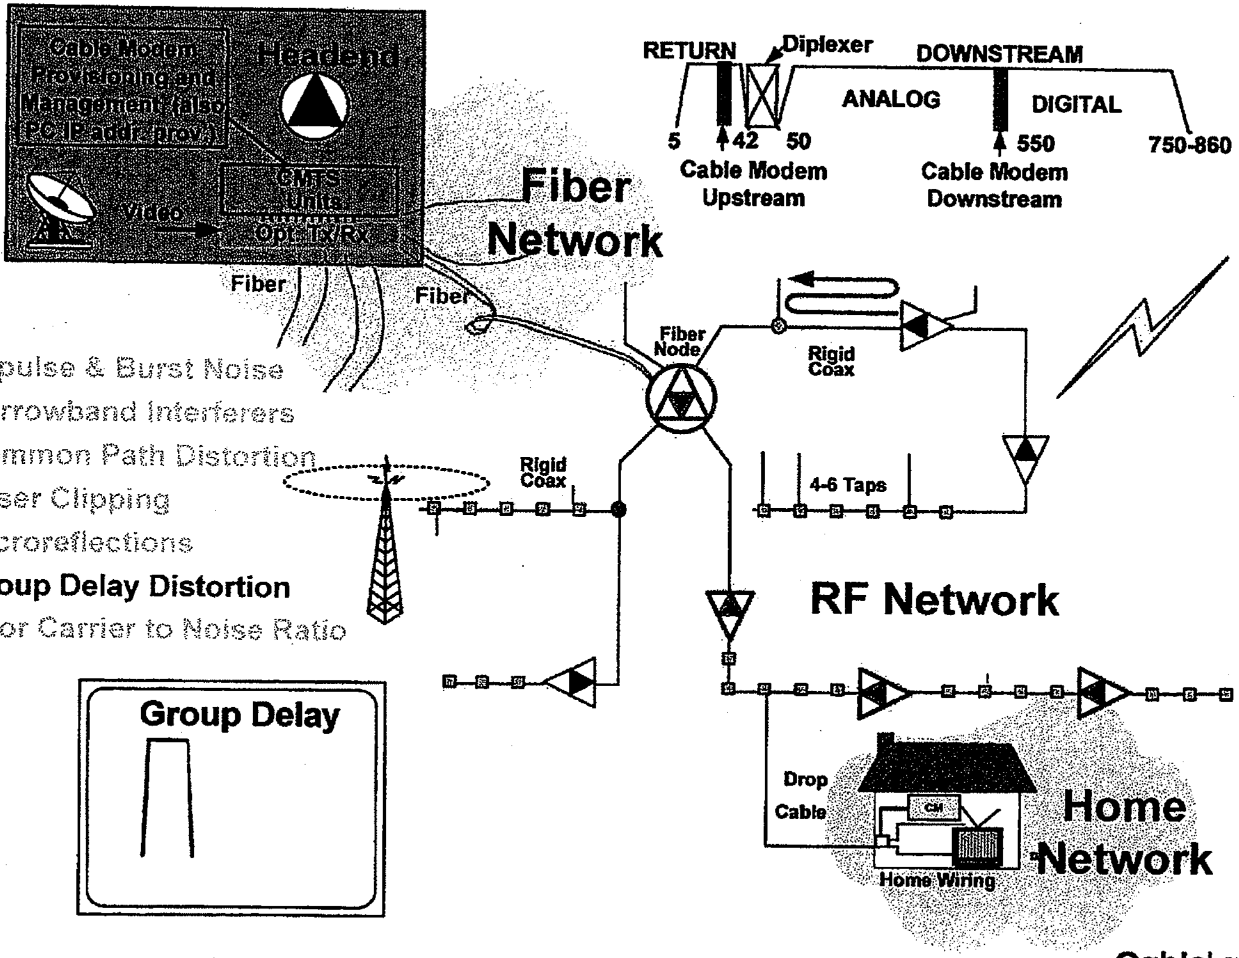

Impulse & Burst Noise

Narrowband Interferers

Common Path Distortion

Laser Clipping

Microreflections

Group Delay Distortion

Poor Carrier to Noise Ratio

[redacted]

HFC Network Impairment Characterization

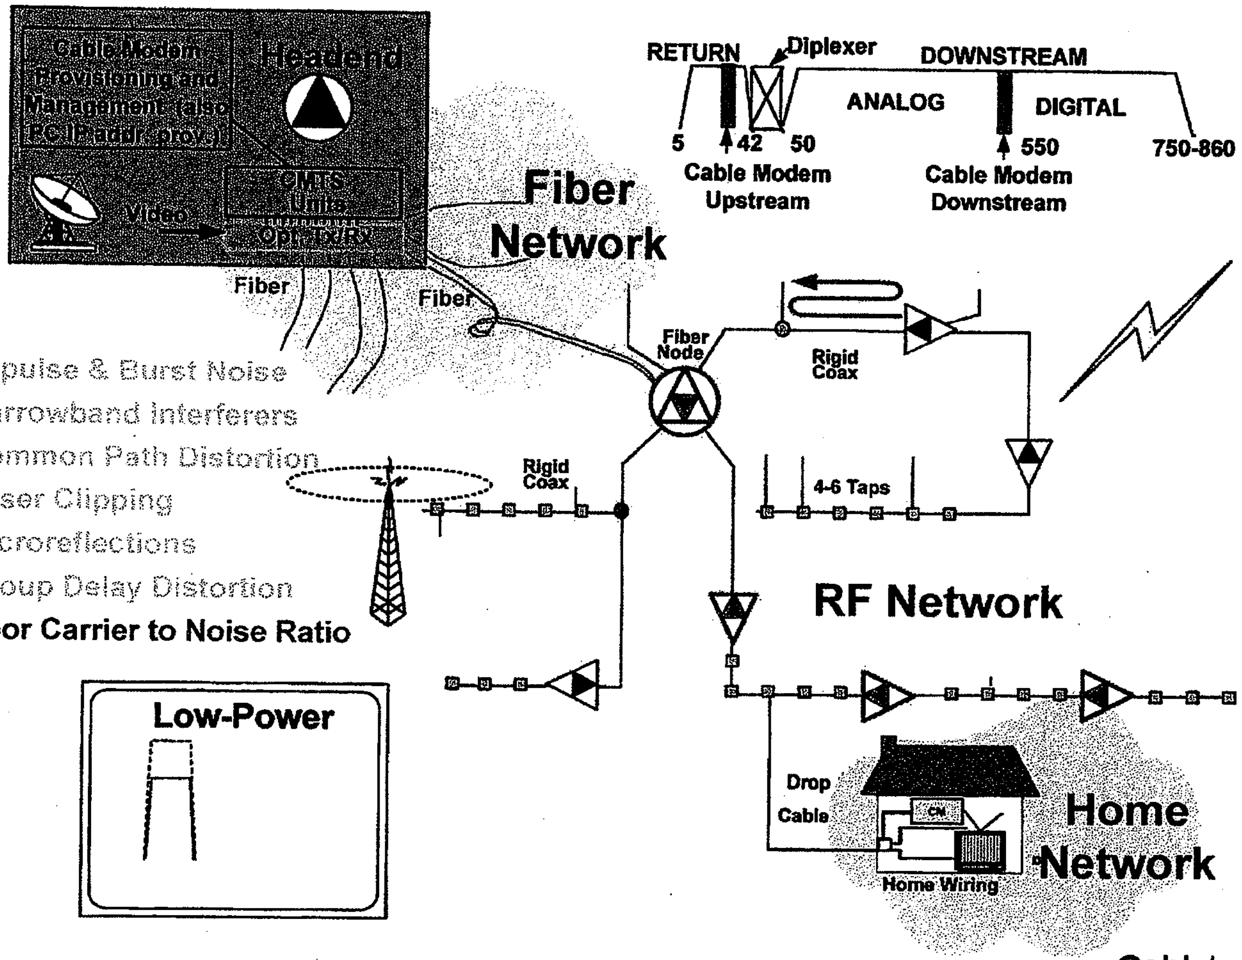

Impulse & Burst Noise

Narrowband Interferers

Common Path Distortion

Laser Clipping

Microreflections

Group Delay Distortion

Poor Carrier to Noise Ratio

[redacted]

HFC Network Impairment Characterization

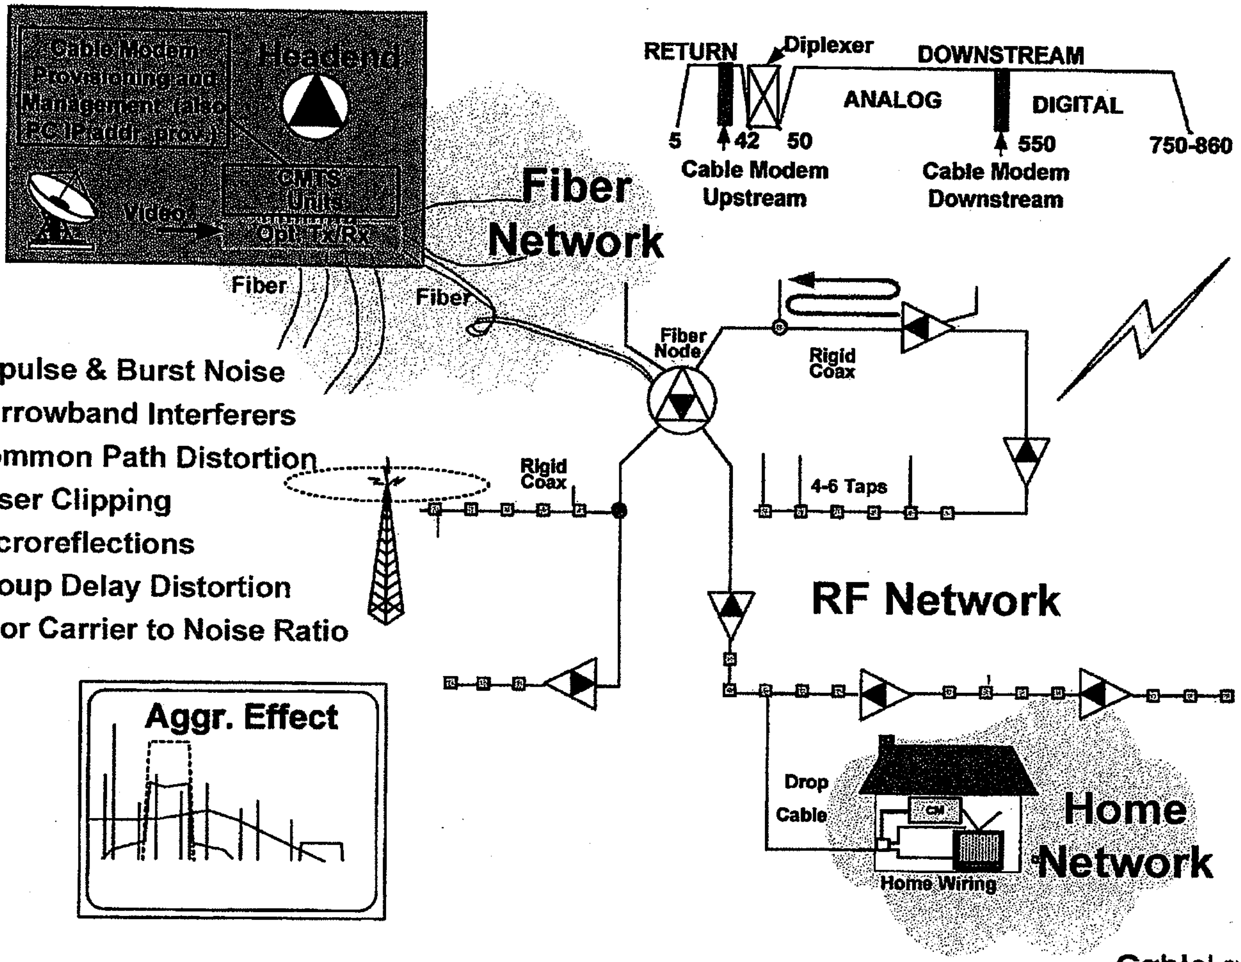

Impulse & Burst Noise
Narrowband Interferers
Common Path Distortion
Laser Clipping
Microreflections
Group Delay Distortion
Poor Carrier to Noise Ratio

[redacted]

CableLabs®
...Revolutionizing Cable Technology®

Upstream Path Loss
Assessed through CM Tx Power Stats

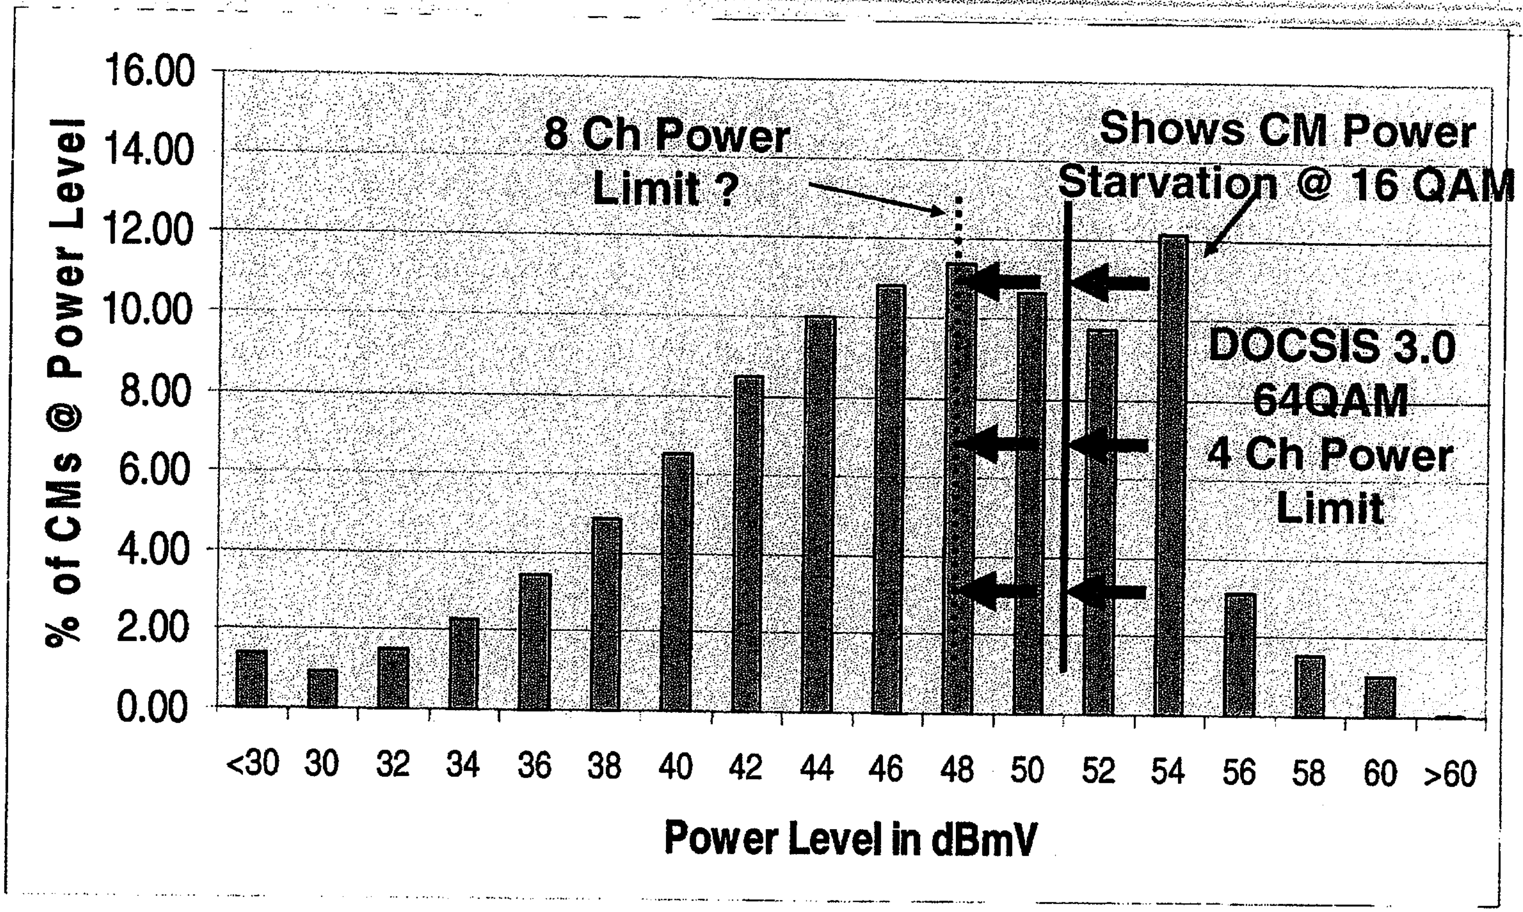

257624 CMs / Average = 46.3 dBmV
0.5 Million CMs Sampled

[redacted]

Requirements to Enable Multiple Ch.

- Increased CM Tx Power Levels
- Tighter CM Tx Power Level Control
- Lower Noise in CM US Transmitters
- Increased CATV Network US Dynamic Range

[redacted]

Plant & CM Transmission
Limited to current DOCSIS Technology

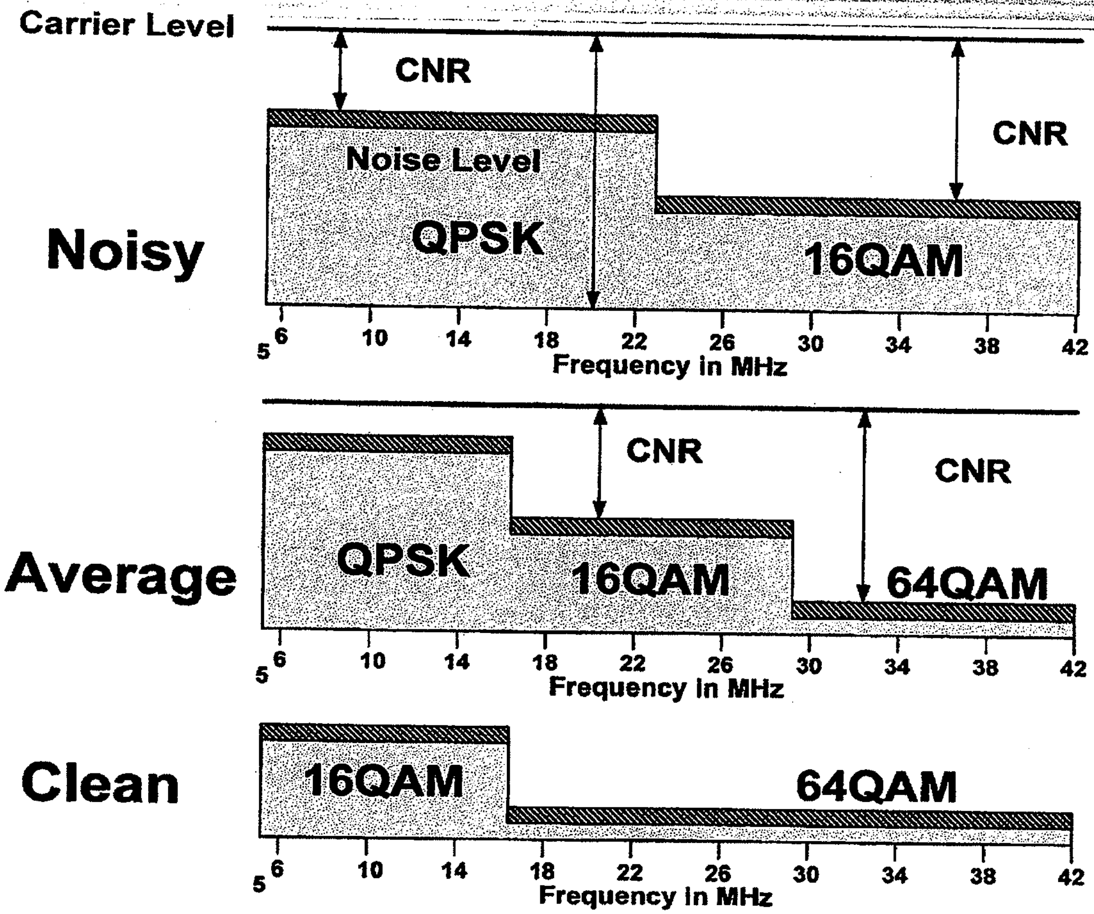

Clean plants are better suited to carry larger amounts of information

[redacted]

CableLabs®
...Revolutionizing Cable Technology®

Capacity Estimate Example
Limited to current DOCSIS Technology

Small % of plants meet noise mask resulting in this ch lineup

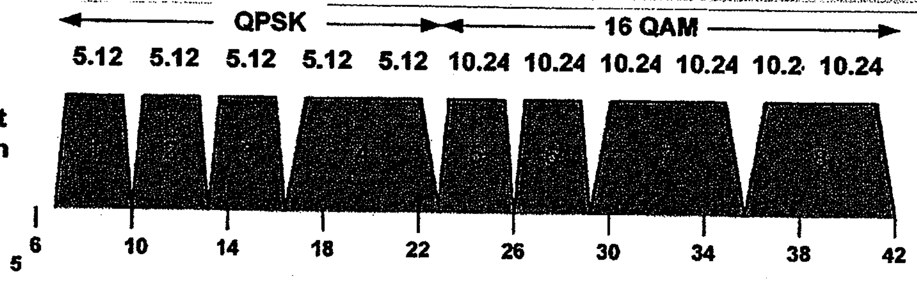

Raw Ch. Capacity

87.04 Mbps

Large % of plants meet noise mask resulting in this ch lineup

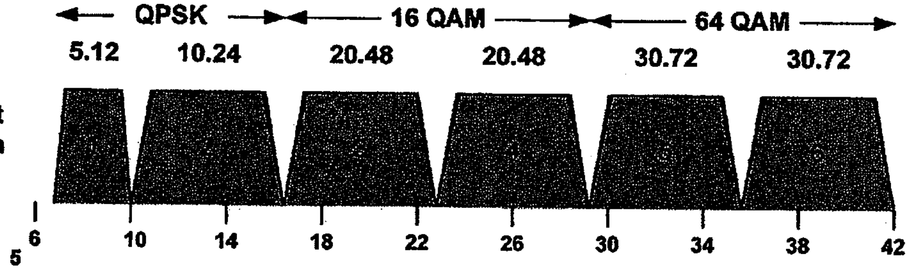

Raw Ch. Capacity

117.76 Mbps

Clean plants meet noise mask resulting in this ch lineup

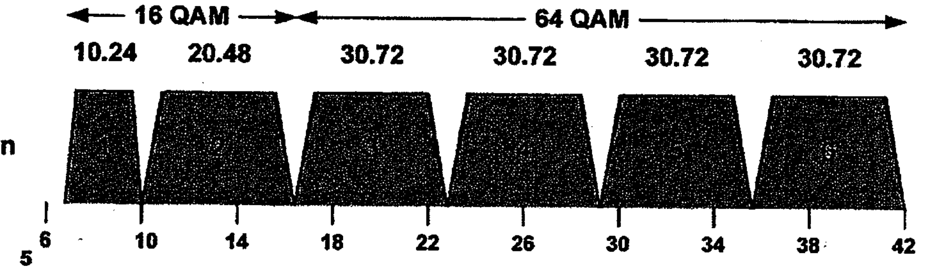

Raw Ch. Capacity

153.6 Mbps

[redacted]

[redacted]

CableLabs
...Revolutionizing Cable Technology

US Capacity Assessment

Theoretical Raw US Capacity Scenarios

- 5-42 MHz
  - » 150 Mbps* on clean plant & current DOCSIS technology
- 5-85 MHz
  - » 350 Mbps* on clean plant & optional DOCSIS 3.0 technology

* 85% long packet efficiency not included in estimate

New wideband downstream channel

- 1Gbps building block for downstream service
- 144 MHz (1 Gbps) or 192 MHz (1.3 Gbps)
- Common channel width between US & Euro
- Eliminates equivalent of 24 or 32 channels of bonding logic

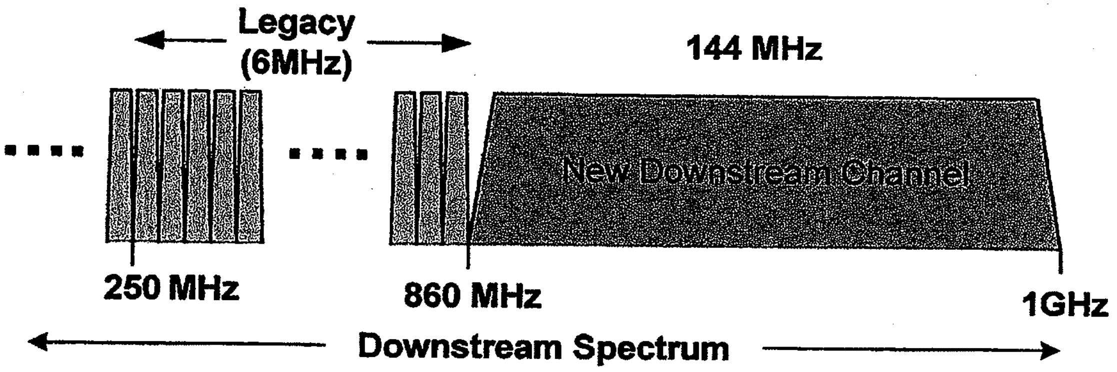

[redacted] [redacted]

Optical CMTS output

- Utilize DWDM to node, combine with video channels at node
- Eliminates expensive (and power hungry) linear power amplifiers on CMTS linecard
- Simplifies combining network in hub

- <insert graphic>

[redacted]

Upstream capacity issue

- 256QAM upstream
  - » Could extend upstream capacity by ~20%
- High Freq. Return
  - » High freq return (above 1 GHz) not attractive
    - – Caps downstream band (limits growth potential)
    - – Tri-plex filter cost and inefficiency
    - – Tx power budget for CM would be more costly

[redacted] [redacted]

New frequency split (much higher than 85 MHz)

- One possibility:
  - » 0-20 MHz: blocked to avoid wasting laser dynamic range
  - » 20-200 MHz: upstream band
    - 20-42 MHz: legacy upstream
    - 42-200 MHz: new upstream band
  - » 250+ MHz: downstream band
    - 250-860 MHz: legacy downstream
    - 860-1000 MHz: new wideband downstream channel
    - >1 GHz: potential future expansion
- Standardized Digital Reverse to bring FN cost down

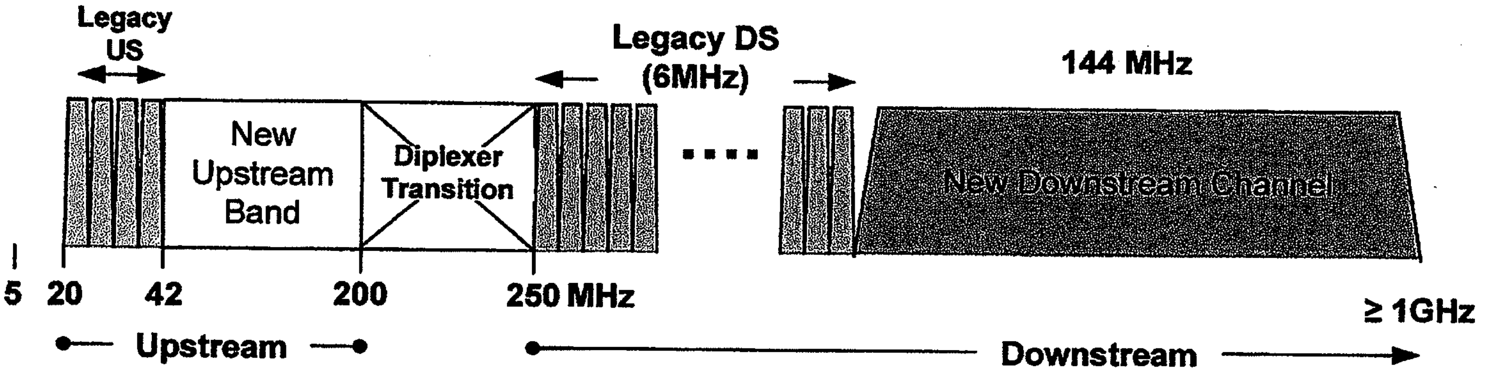

[redacted]

CableLabs® ...Revolutionizing Cable Technology®

Additional Upstream Ideas

- New Upstream PHY
  - In conjunction with a new upstream split (and cleaner upstream spectrum)
  - 256QAM, tighter roll-off, simpler scheduling
  - Alternative modulation formats?
  - Simpler channel aggregation: OFDM? Upstream convergence layer?
- Cable Modem at Coax Point-of-Entry
  - Reduces transmit power req. by avoiding in-home attenuation
  - Incorporates 42-200MHz filter to protect CE devices from upstream leakage

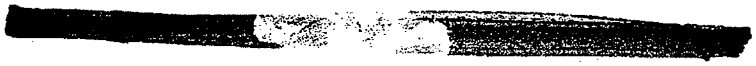

CABLELABS®

PHY Layer Design for Efficient Multi-Gigabit Transport over CATV Networks

prepared by

L. Alberto Campos, Ph.D.
a.campos@cablelabs.com

Jennifer Fang, Ph.D.
j.fang@cablelabs.com

[redacted]

HTTP://WWW.CABLELABS.COM

[redacted] DRAFT

Table of Contents

Figures

1 INTRODUCTION

The Cable TV industry has benefitted from the delivery of data services over their infrastructure for over 14 years. The deployment of CMs has been widespread and operators have been steadily improving the health of the network and therefore their transport characteristics. DOCSIS, the main mechanism of data transport has gone through 3 generations where transport efficiency, peak rates and total capacity has steadily improved over the different DOCSIS versions. Currently, in DOCSIS 3.0, operators have the capability to offer services with peak rates of at least 160 Mbps in the downstream direction and 120 Mbps in the upstream direction. Traffic in CATV networks has doubled every three years. This has prompted the use of additional channels as well as node splitting. Some projections of the type of peak service rates expected for the 2015-2020 timeframe reach the Gbps range in downstream direction.

It is questionable whether the current DOCSIS system will scale well in the gigabit per second range. The upstream limitation of having only a 5 to 42 MHz spectrum available seems to be the first limitation to be encountered as the demand for capacity increases.

This document describes a potential MAC layer evolution scenario for data services over the CATV infrastructure. The focus of this document is in the MAC layer. A separate document describing the PHY layer has been created and is used as basis for the MAC layer proposal described here. The assumptions for a data system to scale well and be cost effective are as follows:
The bandwidth resources provided from the PHY layer are constant and stable. There is significant decoupling between MAC and PHY. The MACs deals with a single large high capacity channel. There are modes of operation, one for a transition phase that relies on a 5-42 MHz upstream and a second one to be the ultimate end game that uses a modified upstream split.

The new data system MUST coexist with legacy data systems although it is not required to be backward compatible. This next generation MAC layer domain is intended to support a maximum of 255 end stations and a much shorter delay difference between end-stations.

The above assumptions can lead to a drastic simplification of the MAC layer and should reduce implementation and operations costs.

2 ORIGINAL DOCSIS ENVIRONMENT

There may be several PHY scenarios that could have met the criteria detailed in the previous section. To illustrate an example of an advanced MAC implementation we have selected an OFDM scenario. The next generation MAC archite

2.1 LARGE NUMBER OF END STATIONS IN MAC LAYER DOMAIN

A Metropolitan Area Network like DOCSIS based networks, initially intended to cover a very large area and a very large number of CMs. The potential of 8192 traffic flows and of CMs per US was built into the protocol. Earlier deployments easily included 2000 CMs per MAC layer domain. Since the early deployments, operators have been splitting nodes, mostly because of capacity demand. As a by-product the smaller plants have been managed more effectively and the capability of more efficient transport mechanisms have been possible. As traffic demand increased over the years the number of CMs per upstream has dropped to about 250.

2.2 WIDE TRANSMIT POWER RANGE & LONG ROUND TRIP DELAY

When DOC SIS was originally designed, the possibility of round trip delays of up to 200 miles were contemplated. In addition, it was assumed that in the same MAC Domain some customers could be located close to the Headend or Hub while others could be at a location of maximum reach. This results in a very different timing compensation ranging from about 0 to 1.6 ms within a MAC layer domain.
From a power perspective a DOCSIS CM was also required to adjust for a very large upstream power variation. Upstream transmit power level ranges from 8 dBmV to 58 dBmV were specified for the earlier cable modems. Deployed cable modem transmit typically at power levels covering a range of about 30 dB (Figure 1).

This variation in upstream power is due to the attenuation differences found in the home, drop as well as tap value differences and differences of signal levels as they hit the first amplifier. In particular it is worth mentioning the power level limitation exhibited by many DOCSIS enabled STBs that reside many times deep within the home network. These variations in delay as well as in transmit power level in addition to the very diverse set of upstream path distortion scenarios all of which vary in time, resulted in the need of an elaborate ranging process.

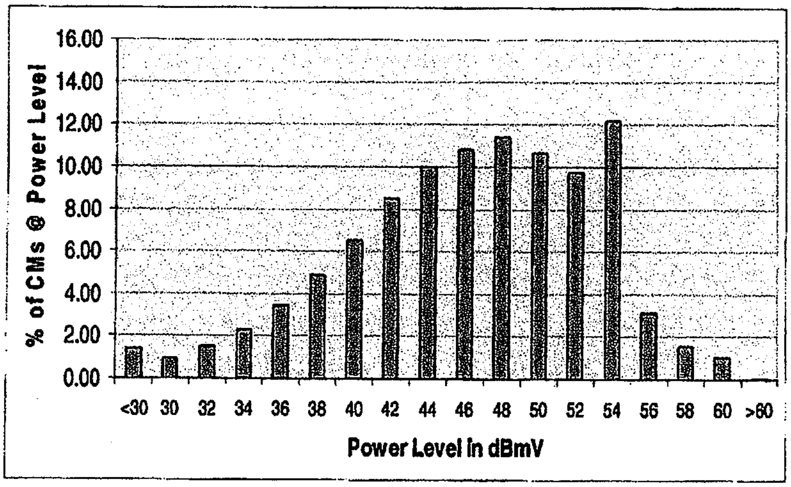

Figure 1 CM Tx Power Level Variation

3 DRIVERS FOR NEXT GENERATION SYSTEM

There may be several PHY scenarios that could have met the criteria detailed in the previous section. To illustrate an example of an advanced MAC implementation we have selected an OFDM scenario.

3.1 DOUBLING OF CAPACITY REQUIREMENTS EVERY 3.5 YEARS

Network traffic has doubled over every 3 years, there is no reason to believe that this trend won't continue in the foreseeable future. Figure 2 shows a historical view of the peak data service rates offered to residential data subscribers.

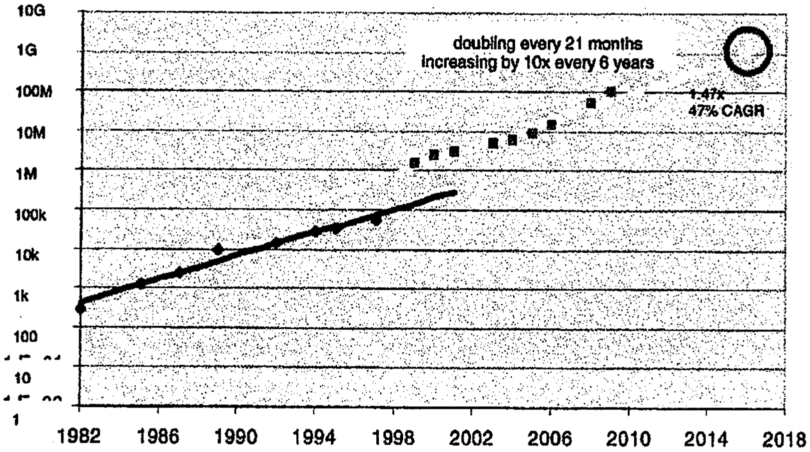

Figure 2 Service Capacity Timeline Estimate

The 47% CAGR in highest speed offered trend from this chart indicates that in the 2015-2020 timeframe, peak data service rates will approach or reach the gigabit per second range. A comparable 40% CAGR in peak bandwidth required per subscriber, further emphasizes this trend. Gigabit per second service will happen first in the DS, followed soon after in the US as a DS/US ratio of 3 would indicate. The DS has a lot of spectrum potentially available. The upstream is confined to the 5-42 MHz portion of the spectrum. Assuming that S-CDMA proves to be robust in the lower portion of the upstream spectrum and an efficient use of DOCSIS carriers is possible, between 120 to 150 Mbps could be achieved is the upstream with DOCSIS technology. Traffic demand beyond that requires a move of the US split.

3.2 DOCSIS COMPLEXITY AND CMTS COST TRENDS

An area of concern is that DOCSIS systems won't be able to scale well for the large number of upstream and downstream channels needed to provide services in the Gbps range. Historically the annual decline in cost per channel of CMTSs and Edge QAMs has been around 15%. This decline in cost cannot compensate the growth in highest speed offered of 47% as shown in Figure 2. That cost per QAM should decrease significantly or the growth demand for capacity show decrease in order for DOCSIS technology to meet future trends. It is also believed that DOCSIS management complexity will increase scenarios with large number of channels. The DOCSIS protocol has about 8500 technical requirements and the PHY layer allows in theory greater than 10^23 configurations combination options. A key problem with it is that the DOCSIS MAC layer is aware of and burden by all these configuration options.

In the following sections we examine potential approaches that could meet the requirements set forth by operators.

4 CURRENT & 5+ YEARS OUT ENVIRONMENT

The environment that was initially assumed for the design of DOCSIS in 1996 has been changed. In order to plan for an environment scenario expected in 5+ years we conducted an evaluation of field conditions and examine trends over the past few years. The implications of those trends are as follows.

4.1 LOWER NUMBER OF END STATIONS PER MAC LAYER DOMAIN

As a result of node splitting, fiber has been penetrating deeper in the network. This has reduced the number of devices a MLD has to serve. A serving area of 500 homes passed typically has about 3 to 4 actives in cascade. The average number of home passed for operators is about 600 homes passed. Assuming a 60% service penetration an average of 360 customers are served today within a node. This number is expected to keep reducing in the coming years as operators continue to split nodes as a means to add network capacity. A key parameter when designing a future network is the number of end stations expected to support.

The architecture envisioned for this future system is a gateway architecture this means that the end device will be the one responsible for forwarding data to other home devices through a separate protocol such as MoCA or a wireless protocol. This architecture results in one end station per home.

It is reasonable to assume that in 5+ year's time frame the number of end stations served per MAC layer domain would be less than 256 over nodes with 2 actives in cascade or less.

4.2 IMPULSES AND BURST NOISE

Cable operators have been very effective at improving their maintenance practices and the health of their plants with time. In the upstream in particular operators are beginning to deploy 64-QAM in key areas. Nevertheless, even though advancements in operational practices are making the HFC environment more desirable, impairments such as burst and impulse noise are still present and any system has to be robust against this type of impairment.

Impulse and burst noise can be generated through internal or external sources to the HFC network. They typically last less than 20 microseconds and they cover the entire upstream spectrum or a portion of it. Sometimes this time dependant phenomena exhibits a continuous increase in noise frequency wise and sometimes consists of a group of discrete interferers. One key characteristic is that in most cases the plant has very good transmission characteristics before and after the occurrence of the burst or impulse. Figure 3 shows a typical view of the upstream spectrum. It shows some narrowband interferers in the low portion of the spectrum as well as a DOCSIS carrier burst. This DOCSIS carrier is used to determine the carrier to noise ratio of the signal. Except for the presence of a few narrowband interferers, Figure 3 shows a carrier to noise ration of 45 dB. This is very suitable for 256-QAM operation.

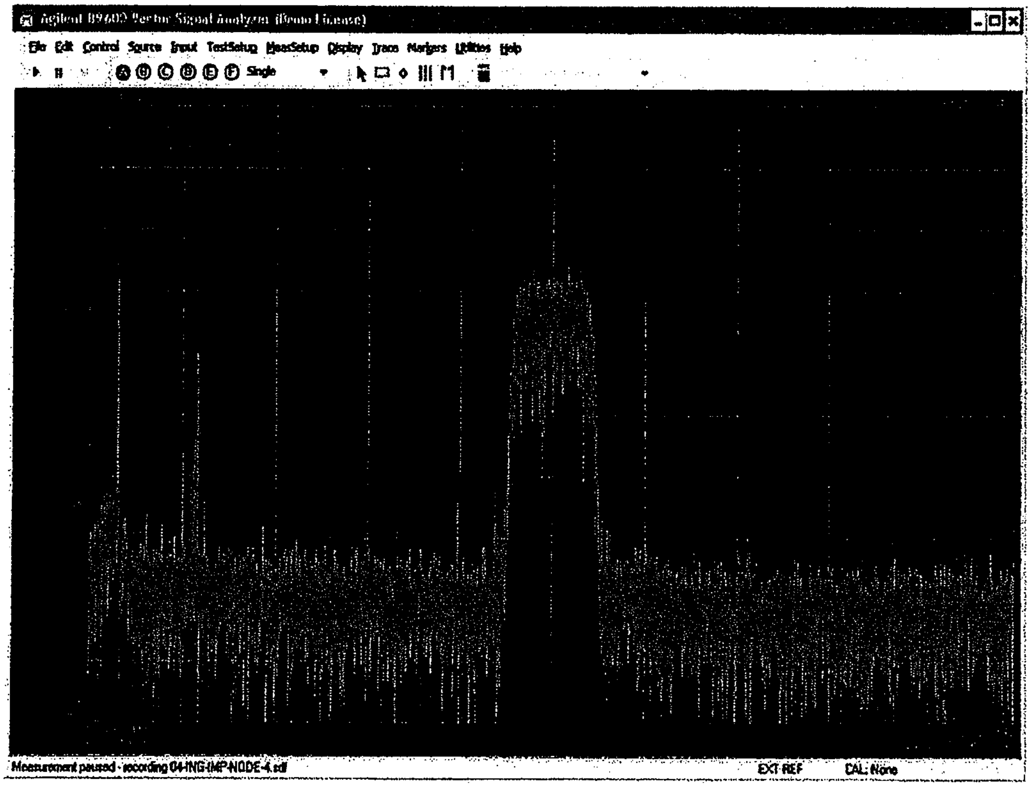

Measurement paused - recording 04-ING-IMP-NODE-4.sdf

Figure 3 Typical US Spectrum

Figure 4 shows the upstream spectrum of the same node as Figure 3 but at a different instant in time.

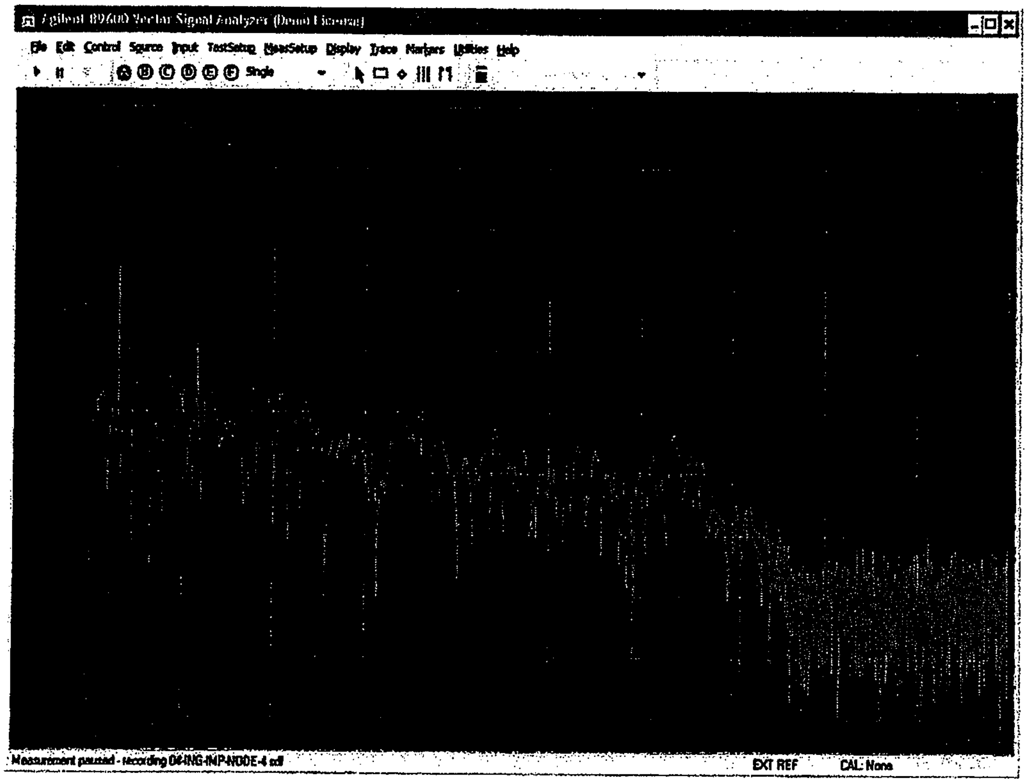

Measurement paused - recording 04-ING-IMP-NODE-4.sdf

Figure 4 US Spectrum with Impulse

This short impulse lasted less than 20 μs and covered the entire upstream spectrum. Immediately after the impulse the spectrum exhibited again the same level of noise shown in Figure 3. These burst and impulse events occur randomly at average intervals that are typically greater than 10 ms. These event intervals are greater than the typical access latencies that we expect from modern systems. If in a future design we used symbol periods that are much longer in duration than these 20 μs bursts we will have an automatic robustness built in.

FEC that corrects a single symbol will capable to correct for any burst when the duration of the burst is much shorter than the duration of the symbol. Also systems that have a long symbol duration can implement very simple interleaving mechanisms. In a worst case condition the interleaving mechanism may have to multiply the robustness against bursts by a factor of two.

4.3 NARROWBAND INTERFERERS

HFC networks also have in the upstream narrowband interferers which are typically present in the lower portion of the upstream spectrum (Figure 3). There are narrowband interferers that occur over a brief moment in time in which case it can be handled like a burst although in more cases the upstream interferers are fairly static in time.

4.3.1 INGRESS MITIGATION TECHNIQUES

There are many robustness mechanisms against interferers that can be implemented. There is no one ingress mitigation technique that is optimal for all transmissions systems. Some mechanisms may be more suitable than others for specific transmission systems.

For example, using digital filtering that is dynamically controlled based on where the interferer is located may be suitable for some wideband systems, although there is a level of complexity that relates to adjusting the notch at the right frequencies and generating the number notches to match the numbers of interferers.

An ingress mitigation example that is applicable to DOCSIS S-CDMA is the one that relies on selecting the better performing S-CDMA codes. This mechanism takes advantage of the fact that S-CDMA codes have distinct frequency sensitivities. For example, a narrow band interferer may disrupt the operation of one or a few codes while a narrowband interferer of the same level but slightly shifted in frequency may have absolutely no impact on the operation of those codes. Using the principle named selectable active codes (SAC), the CMTS has the capability of selecting the codes that it plans to use and the ones that will remain unused so that only high performance codes are used. Since there are 128 codes, the fraction of effective bandwidth lost for each code that is not used is one or multiples of 40 KHz (5120 KHz/128) bandwidth segments for every narrowband interferer. In a 3.2 MHz channel scenario the efficiency in granularity is improve at the expense of having to manage a larger number of bonded channels in order to achieve equivalent peak bandwidths as a wider channel.

A third approach which is applicable for multi-subcarrier or multi-tone systems is quite simple and efficient when the subcarrier or tone width is very narrow. This approach requires the silencing or suppression of the subcarrier(s) that coincides with the interferer. The subcarriers that are not impacted by interferers are used. This mechanism is similar to the DOCSIS SAC approach except that a narrowband interferer may impact a few S-CDMA codes while it typically impacts one or at most two subcarriers. The efficiency obtained by this mechanism improves with the use of narrower subcarriers as less bandwidth is lost.

4.4 TECHNOLOGY ADVANCEMENTS

In the last 15 years since DOCSIS was originally conceived, technology advances in processor speeds, memory capacity and signal processing capabilities have been increasing steadily following Moore's law. These advancements have resulted in very high ADC & DAC (Analog-to-Digital & Digital-to-Analog) speeds reaching now around over 2 Giga-samples per second.

At these sampling speeds the entire downstream (and upstream) spectrum can be generated and detected using single integrated circuits. These advancements have also enabled very large count multi-subcarrier or multi-tone systems with very narrow subcarriers. As transistor gate technology has gotten smaller, devices have gotten faster and power consumption has also decreased.

A new system should take advantage of these advances to optimize performance and flexibility for implementing systems.

4.5 COST CONSIDERATIONS

In the implementation of network systems a significant portion of the cost resides in the development of the system. If a new system that uses a well known and mature technology is selected, the development effort is reduced. In addition if there are existing technologies that are very similar or that can be adopted than this would also be beneficial to maintaining the costs low. In particular if the technology adopted enjoys large magnitudes of scale that already enjoys a low cost.

Keeping the number of variables and configuration options low reduces complexity and cost. Low complexity implementations mean smaller die area and more chips per wafer. In addition if the technological requirements are not challenging and its implementation is not pushing the envelope. The yield, which is the percentage of useful chips per wafer that can be obtained, is going to be much higher.

In relation to the CPE or end-station cost, another cost driver is the upstream transmit power. It is important to keep transmit power low in particular in an environment where the transmit bandwidth could be significant (>160 MHz). If transmit target levels approximately equal to DOCSIS 3.0 are maintained, low cost implementations could be achieved.

5 GENERAL ASSUMPTIONS FOR NEXT GENERATION SYSTEM

There may be several PHY alternatives that could have met the general criteria detailed in the previous section. To illustrate an example of an advanced MAC implementation we have selected an OFDM scenario

Physical Plant	Hybrid Fiber Coax Network
Coexistence	Must coexist with legacy systems (Not backward compatible)
Service Goals	Downstream: Multiples of 1 Gbps Upstream Short term: 200 Mbps (5-42 MHz option) Upstream Long term: ≥ 400 Mbps (>42 MHz option) Goal of 1 Gbps
Complexity	Need to keep everything as simple as possible Minimize cost of chips, products; and operations
PHY Layer Requirements	Robust against narrow band interferers & impulse noise Single, constant PHY configuration
MAC Layer Assumptions	Single PHY entity presented to MAC Max number of end stations per MAC Domain < 256 Granular classification of traffic flows

Table 5-1 General Assumptions for Next Generation System

5.1 RELYING ON HFC NETWORK ARCHITECTURE

HFC networks have been evolving. Cable operators have been penetrating deeper with fiber in order to split nodes. In spite of the ever increasing demand for bandwidth, there is still plenty of capacity left in HFC networks. In a 1 GHz HFC network which many operators have migrated to in significant number of systems, you can carry 6 Gbps with the same standard modulation techniques we have been using for years. In fact it is still possible to expand this capacity farther with minor upgrades to the actives and passives. The fact the the nodes are smaller makes its health easier to manage with less noise and impairments enabling even more transmission systems. The large expense into migrating to all fiber networks is not warranted unless it is a green field scenario. The proposed advanced PHY approach presented here relies on the continuing use of HFC networks. Nevertheless it also enables a smooth future migration to a fiber to the home (FTTH) architecture.

Figure 4 show a typical access network topology that includes the proposed system.

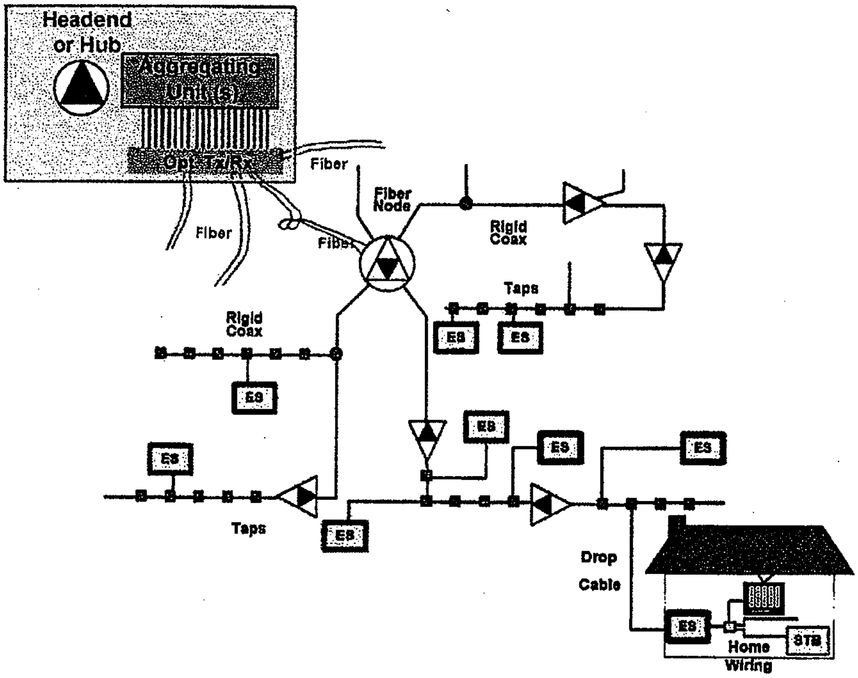

Figure 5 HFC Access Network with Proposed Next Generation System

In the headend or hub location there is an aggregating unit that includes the protocols necessary to manage the end-stations deployed in the field. The nodes are small enough to meet the proposed requirement ofr maintaining the number of end-station per MAC layer Domain to < 256. In the home network the end station is located at the drop-home boundary. Within the home legacy devices such as Set Top Boxes, TV sets and even legacy CMs can coexist.

5.1.1 LIMITED INTELLIGENCE AT FIBER NODE

One of the characteristics of this proposed system is that is does not add intelligence at the fiber node. Adding intelligence at the fiber node, result in the commitment to the technologies installed at the node. This to some extend binds the operator to certain technological choices for some time. If the components used to transport data services, rather than being located at the node, are located at Headends or hubs and at the customer premises, then it is easier to upgrade and evolve services.

5.2 COEXISTENCE WITH LEGACY SYSTEMS

One of the key requirements that a new PHY should have in order to guarantee a smooth transition is the coexistence with legacy systems. Coexistence with legacy systems does not imply backward compatibility. In fact, backward compatibility burdens any implementation with significant complexity and cost. The fact that a system doesn't have to be backward compatible provides freedom in the selection of key technologies, processes and parameters.

5.3 END STATION LOCATED AT DROP-HOME BOUNDARY

Another important characteristic of this proposed new system is the gateway architecture where the signal typically travels from Headend to the drop home boundary. Within the home environment, another transport system such as MOCA or WIFI would be used to carry the information from the end station (gateway) to other devices within the home. This approach has several advantages. It leads to a lower number of end stations as you would just have one end station per home. A second very advantageous consequence is the fact that by placing the end-station at the drop-home boundary the loss within the home network is avoided.

Following this strategy could lead to a power budget improvement of up to 12 dB, as the home network is bypassed. The fact that a much wider spectrum is needed to reach gigabit per second speeds, will also demand much higher power levels assuming the power spectral density in relation to today's systems will be maintained. Smaller HFC plants are easier to manage and maintain cleaner. In future networks where a greater majority of nodes is expected to be small the lower noise levels could lead to lower transmit power requirements. Nevertheless because in the proposed system it is assumed that a single and efficient modulation scheme of 256 QAM is used, keeping the power density at the same levels ensures that transmission robustness at those modulation efficiencies can be maintained.

In fact if we compare the current DOCSIS 3.0 transmit power requirements to a gigabit per second upstream system using a spectrum of 160 MHz and located at the drop-home boundary, we find that the power requirements are very similar.

Placing the end-station at the drop home boundary also enables the isolation between the home and HFC network. A lot of the sources of noise originate in the home network. A filter that isolates the home and HFC network can be used to keep unwanted noise from entering the HFC network. Such filters can be designed to support legacy systems while blocking portion of upstream spectrum that contribute more heavily to noise. Such filters can also be designed to enable the reclamation of significant amount of spectrum through the isolation of HFC and home networks (Figure 5 shows an example of the gateway configuration and some filtering approaches).

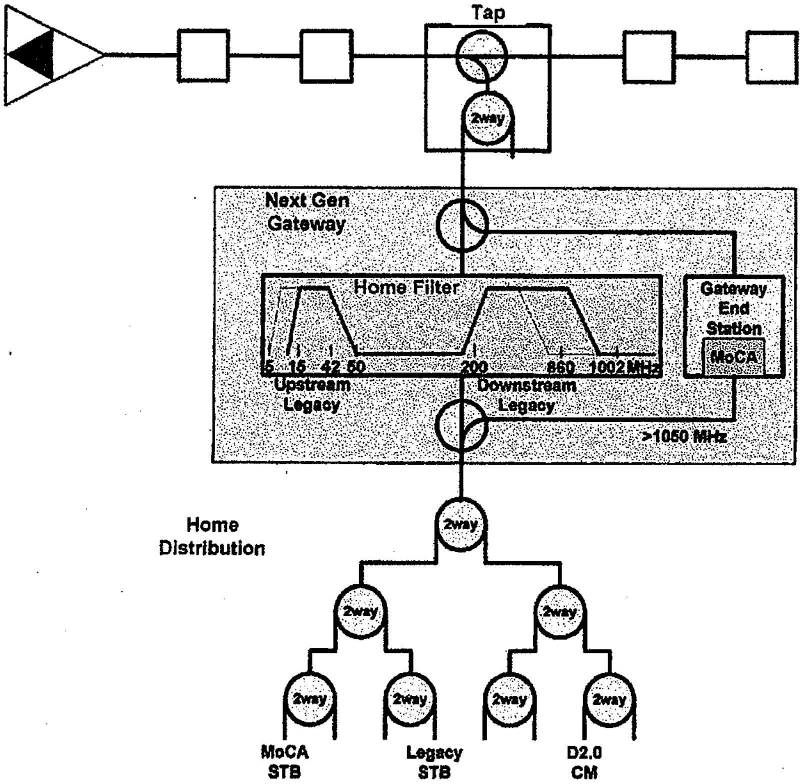

Figure 6 End Station Located at Home Edge Boundary

5.3.1 TRANSMIT POWER VARIABILITY ANALYSIS

CATV networks fiber nodes covering a large serving area exhibit significant transmit upstream power variability. Figure 1 highlights the power variability observed in current HFC networks. This variability is due differences in attenuation in the home coax, drop, tap coupling and insertion losses, and variability in the input power driving the upstream amplifier and rigid coax loss. In order to maintain upstream and for that matter downstream power variability low, cable operators have at their disposal the adjustment of tap values which impact both upstream and downstream transmissions as well as maintaining tighter amplifier input power values through proper power alignment.

Maintaining the power levels closer together is much easier when the home attenuation such as what you would have in a gateway architecture is not relevant. Figure 6 contains the elements that impact power variability in a gateway architecture.

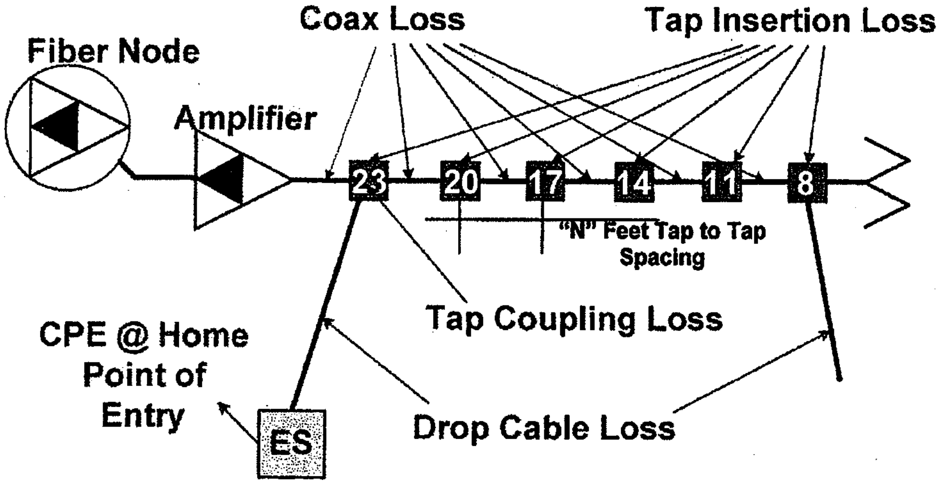

Figure 7 Loss Components In Coaxial Cable Segment

In this modeling, we also assume that the home density between homes that are sharing a tap is pretty similar, therefore exhibiting little difference in drop loss. In this scenario the adjustment to coupling loss that an operator can exert by changing a face plate can go a long way towards reducing the upstream an downstream power variability. With these assumptions that include tap coupling loss optimization our analysis indicates that the power variability can approach 6 dB. This could lead to simplification in ranging mechanisms.

In addition limiting the requirements for transmit power dynamic range can help in keeping implementation costs low.

5.3.2 TOTAL UPSTREAM TRANSMIT POWER LEVEL REQUIREMENT

Assuming the same power spectral density as in a DOCSIS 3.0 system with four 64 QAM channels. That is that for each 6.4 MHz channel 51 dBmV would be available if there would be a home network. If we assume that a home network contributes to an average loss of 7 dB (loss of a 4 way splitter), it means that a gateway architecture would need to deliver 44 dBmV for each 6.4 MHz of bandwidth. In case 160 MHz are needed to deliver 1 Gbps that is equivalent to 25 6.4 MHz channels. To maintain the power density of 44 dBmV per 6.4 MHz a total power of 44 dBmV + 14 dB (10 Log_{10} (25)) = 58 dBmV are needed. This upstream power level requirement is similar to the power capabilities of DOCSIS 3.0. Since the upstream transmit power level requirements is a key cost indicator it is encouraging that a gigabit per second system can maintain similar total power level as in the case of DOCSIS 3.0 CMs.

Assuming that 58 dBmV is delivered by the end station and assuming an upstream loss of 30 dB (27 dB TAP and drop loss and some rigid coax loss), the fiber node or amplifier would receive 28 dBmV of power. In the case of a fiber node, an internal 4 way combiner (typical of a four port fiber node) incurs a loss of 7 dB resulting in 21 dBmV input power available to drive the upstream laser (Figure ??). High power DFB lasers with a typical operating input range from 8 dBmV to 23 dBmV have a signal to noise and distortion ratio (SINAD) of 38 dB, which is very suitable to support 256 QAM operation (Figure ??)

5.4 REDUCED NUMBER OF END STATIONS

Cable operators have been gradually making their fiber nodes smaller by penetrating deeper with fiber and splitting nodes. In addition in this proposed approach the gateway architecture assumption leads to single end-station per home. These characteristics result in a reduced number of end-stations per node. A low number of end-station per node simplifies many processes in particular as it relates to the MAC Layer Domain (MLD).

6 PHYSICAL LAYER SYSTEM CHARACTERISTICS

There may be several PHY scenarios that could have met the criteria detailed in the previous section. To illustrate an example of an advanced PHY implementation we have selected a scenario using OFDM. This selection was made in part because OFDM is a well known technology and also because advances in signal processing have made possible to implement very large number of narrow subcarriers. Narrow subcarriers exhibit advantages compared to wideband systems in robustness against impairments.

6.1 OFDM – ORTHOGONAL FREQUENCY DIVISION MULTIPLEXING

OFDM consists of a number of subcarriers that are orthogonal to each other in the frequency domain. This means that one subcarrier's data does not depend and can be discriminated from that of other subcarriers. A spectral view of OFDM subcarriers highlights this characteristic (Figure 8). In this figure at one frequency point when one subcarrier achieves maximum transmission, the power levels of all other subcarriers are null. A key requirement for that is to maintain the same frequency spacing between subcarriers.

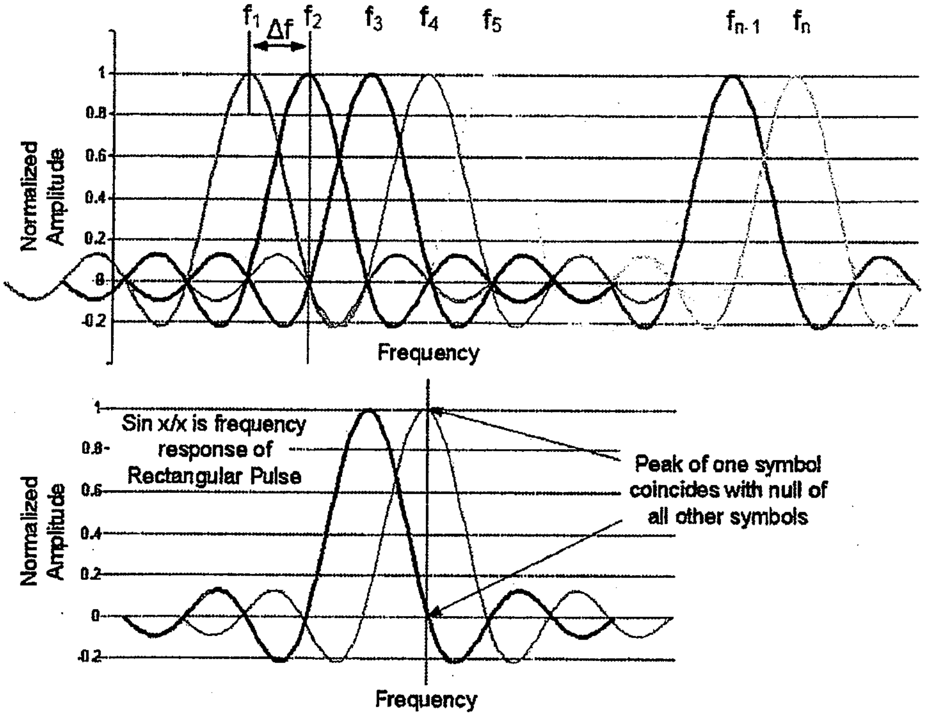

Figure 8 OFDM Subcarrier Frequency Orthogonality

6.2 UPSTREAM AND DOWNSTREAM PHY CHARACTERISTICS

The PHY layer approach proposed in this paper assumes the use OFDM technology in both the upstream and downstream directions.

In the downstream direction, the RF environment is more benevolent than in the upstream direction, resulting in more efficient modulation schemes. The modulation order proposed here for the downstream is 1024 QAM with the option of higher order modulations if plant conditions and suitable error correction encoding schemes allow it.

In the upstream, a harsher environment is anticipated. Nevertheless the design decision of very narrow subcarriers and long symbol periods go a long way in providing needed system robustness. In the upstream, it is assumed that the network characteristics would allow a modulation order of 256 QAM.

There are also two implementation phases that are assumed based on upstream spectrum availability; one before a split is modified and an ultimate after the split is changed. The details of these options are discussed next.

6.3 UPSTREAM FREQUENCY SUBCARRIER ALLOCATION 5-42 MHZ SCENARIO

In the initial implementation phase when upstream resources are bounded by the 42 MHz split, only 37 MHz (42-5 MHz) is available. Of the 37 MHz, the subcarriers that couldn't be used because of legacy transmissions being present, as well as the subcarriers not suitable for optimum transmission have to be subtracted. In order to simplify PHY configuration options the upstream portion of this system uses 256-QAM as a single modulation scheme for every data-carrying symbol.

In the system described here, we propose to use 4096-FFT blocks, with a subcarrier spacing of 10 KHz. In theory out of the 4096 tones in the FFT block we would be able to use at most 3700 subcarriers. From those 3700 subcarriers the following subcarriers are subtracted;

1) subcarrier locations where narrowband interferers are located
2) subcarrier locations where less than 256 QAM is achievable
3) subcarrier locations where DOCSIS legacy channels are being transmitted

The 10 KHz subcarrier spacing corresponds to a symbol period of 100 microseconds.

Sections 4.2 and 4.3 discussed HFC network impairments where plant measurement show that impulses or bursts of up to 20 microseconds could be expected. In addition microreflections in the cable environment are typically less than 2 microseconds in duration. Both impairments represent a small fraction of the symbol period and will have negligible impact in the operation of this system. Even in the case of implementing an interleaver to increase robustness against impulse, this interleaver would be very simple as you would only need to multiple by a small factor to have burst noise immunity. In the case of narrowband interferers, most narrowband interferers are less than 10 KHz in width. The fact that silencing a subcarrier is a very simple task compare to other ingress mitigation approaches (Section 4.3)

In the initial upstream phase scenario where the split is unchanged and a 5 to 42 MHz is available configuration decisions have been made regarding how these resources are used. This paper assumes that an operator will maintain in this initial phase one 6.4 MHz DOCSIS channel. Section ## goes demonstrates how 660 10 KHz subcarriers are sufficient to carve out enough upstream spectrum to support DOCSIS legacy operation. It is also assumes that the number of narrowband interferers plus the number of subcarriers that are in an environment where 256 QAM is not feasible is less than or equal 136. This represents an additional overhead of 1.36 MHz.

The total number of subcarriers available in this OFDM implementation is 2904 (3700 – 660 – 136). In addition in an FEC of 15/16 is assumed as well as a total number of 88 pilots (1 pilot every 32 subcarriers). Subtracting the pilots and multiplying by 15/16 we have a total of 2640 subcarriers available as a constant and stable set of resources for the MAC layer. Not taking into account guard time overhead (which will be discussed in detail later), the above assumptions result in a total capacity of 211.2 Mbps (2640*8bits/100*10^-6 sec) available to the MAC.

Figure 8 shows a representation of subcarrier allocation scenario in the 5-42 MHz (number of subcarriers represented in true scale).

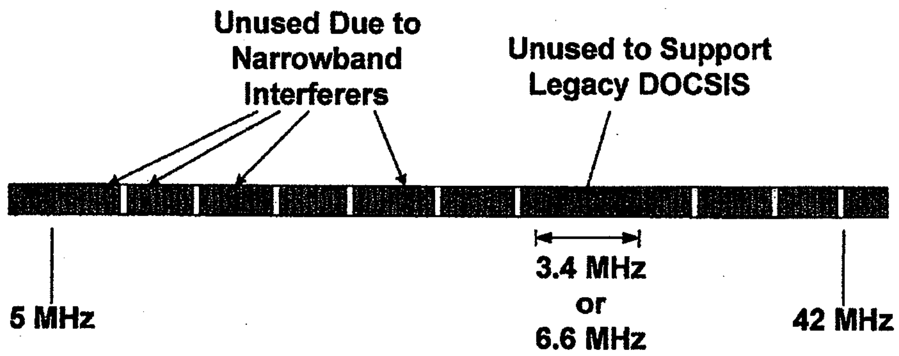

Figure 9 OFDM IFFT/FFT Block Subcarrier Allocation Covering 5-42 MHz

This upstream subcarrier allocation configuration is not the only one that would be supported but it is one of 16. Additional upstream allocation scenarios will have other options regarding how many legacy channels are to be supported as well as how many subcarriers are needed to reserve for suboptimal (non-256-QAM capable) performance.

6.4 IMPLEMENTATION PARAMETERS TO ENSURE COEXISTENCE WITH LEGACY DOCSIS

In the presence of narrowband interferers it is enough to have the subcarrier that coincides with the narrowband interferer silenced. However in order to allow intelligent carriers around OFDM tones, the OFDM sidelobes have to be a value low enough to allow for the intended modulation. For example in order to allow for QPSK transmission approximately a carrier to noise ratio of 17 de is needed while for 64 QAM transmission around 27 dB is required. Figure 9 shows the power drop for the different sidelobes adjacent to OFDM transmission. A margin better that 30 dB is obtained when a guard frequency of 10 subcarriers is used.

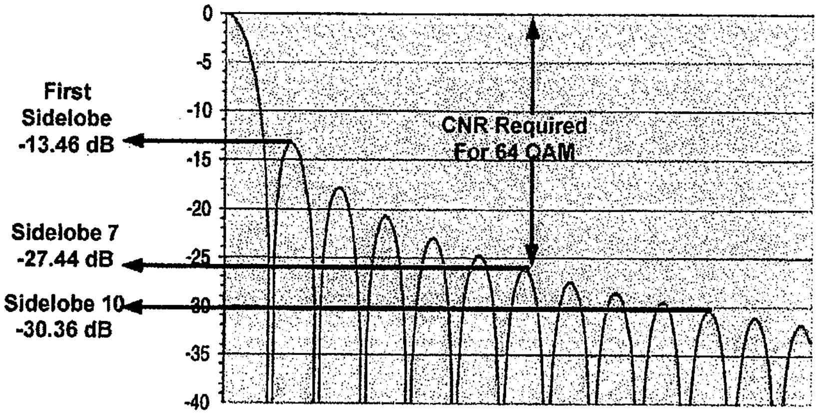

Figure 10 OFDM Spectrum on Adjacent Subcarrier Location

In particular for the DOCSIS case, the carrier does not occupy the entire channel but a portion. In the case of a 6.4 MHz channel the carrier uses 5.12 MHz. This means that there is about 640 KHz on each side of the carrier to the edge of the channel. Figure 10 highlights the difference between the channel and the carrier. It also shows how the sidelobe level of an OFDM carrier decreases as it is separated several subcarriers apart in frequency. The analysis shows that the a separation of 64 subcarriers brings the sidelobe level down to 46 dB.

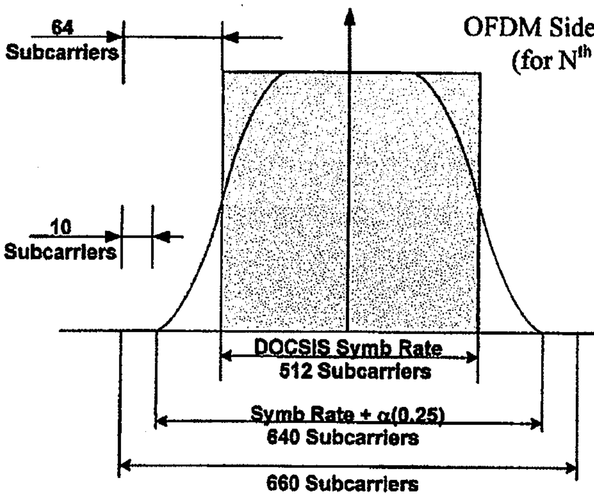

$$\text{OFDM Sidelobe Level (for } N^{th} \text{ Lobe)} = 20 \cdot Log_{10}\left[\frac{Sine[(2N+1)/2]}{(2N+1)/2}\right]$$

N	SideLobe (dBc)
1	-13.46
7	-27.44
10	-30.36
64	-46.13

Figure 11 Spacing between Band-Edge and DOCSIS Carrier in Number of OFDM Tones

6.5 UPSTREAM FREQUENCY SUBCARRIER ALLOCATION > 42 MHZ SCENARIO

It is assumed that at some point in time, cable operators will be able to move the split at a higher frequency. This will make more spectrum available for upstream transmission. It may be hard to decide what that future split frequency would be at this time, nevertheless the system described here is intended to be modular and flexible to adapt to whatever split option cable operators may decide to migrate to. In the frequencies above 42 MHz there are significantly less interferers present than below 42 MHz. It is expected that a larger number of subcarriers be available. Also in that frequency range, there is no spectrum to carve out for upstream legacy DOCSIS. Perhaps one reason to have a portion of the spectrum silent would be if there are sensitive frequencies such as aeronautical frequencies that operators may want to avoid if there is concern for interference.

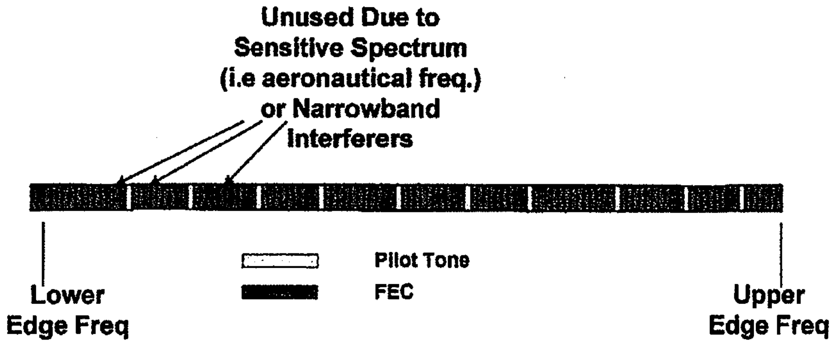

Figure 12 OFDM IFFT/FFT Block Subcarrier Allocation for Frequencies > 42 MHz

If out of the 4096 subcarriers that are theoretically available in the proposed FFT building block, 136 subcarriers are not used because they are allocated to either suboptimal transmission or to be silent because of protection of a sensitive portion of the spectrum. The total number of useful subcarriers becomes 3960 (4096-136). If we subtract pilot tone and FEC the resources that are available to the MAC layer would be 3600. Operating the useful subcarriers at 256 QAM the resources available for this block to the MAC without considering guard time overhead result in 288 Mbps. (3600*8bits/100x10^-6 sec).

In the upstream depending on the split selection either two or more FFT blocks can be used. In the case of a 85 MHz split with 2 FFT blocks a total of 480 Mbps can be provided by the system. In the case of a 200 MHz split 4 FFT blocks could be used and provide more that 1 Gbps service.

The upstream OFDM subcarrier usage is communicated by the MAC through an Upstream Channel Descriptor (UCD) message. This message is only understood by this particular PHY. If a future or alternative PHY would needs a different UCD, this proposed approach would accommodate by allowing multiple UCD types. For this UCD type, the message provides a data string with the sequence of the number of silent subcarriers followed by the number of active subcarriers and so on. This particular information would benefit from a run-length type of encoding, such as how facsimile data is encoded.

6.6 DOWNSTREAM FREQUENCY SUBCARRIER ALLOCATION > 42 MHz SCENARIO

In the downstream assuming the same 3600 subcarriers are available to the MAC and assuming an operation at 1024 QAM the resources available for this block results in 360 Mbps (3600*10bits/100x10^-6 sec). If a higher order modulation such as 4096 QAM is possible then a total of 432 Mbps (3600*12bits/100x10^-6 sec) would be available per FFT block. In the case of 1024 QAM modulation 3 FFT blocks could provide a service of 1 Gbps. Since in the downstream direction all the information comes from a single source, the aggregating unit and there is no handshake needed. The determination of what subcarriers are silent and which are not can be done using special tones/subcarriers that indicated the lower edge of a group of silent tones/subcarriers as well as a special tone that indicates the upper edge or the beginning of active subcarriers. DVB-C2 has defined this type of process.

6.7 MODULARITY AND FLEXIBILITY

Cable operators have not decided how to implement a upstream split change in the network or even if a split change is needed. This unknown requires flexibility in the potential plant migration strategies. For the purposes of supporting different migration alternatives, this approach describes a modular implementation that can adapt to several migration scenarios. Figure 12 highlights different example migration scenarios.

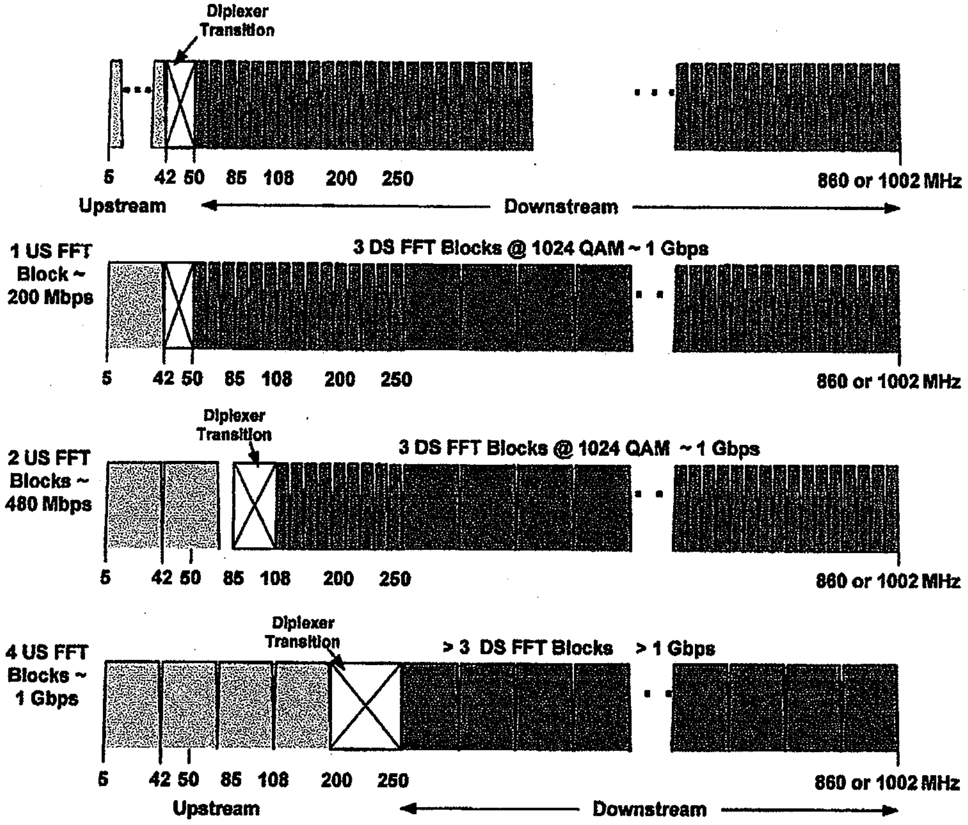

Figure 13 FFT Block Configuration Options depending on US Split Frequency

The first row in Figure 12 represents the current scenario in N.A. plants prior to a split change. The second, third and fourth rows in Figure 12 shows upstream splits at 42, 85 and 200 MHz carrying the proposed implementation. These implementations rely on the use of one or more modules or blocks to achieve the desired capacity and peak rates. Earlier sections describe in detail the upstream blocks for frequencies < 42 MHz and > 42 MHz as well as the downstream block. In this paper the upstream transmission was chosen to be at lower frequencies than the downstream. In principle, the upstream could have been placed above the downstream, at frequencies higher than 1 GHz, nevertheless to reduce power requirements on the end-station, about a 20 dB difference and to avoid a diplexer bandwidth overhead of about 150 MHz a low upstream was chosen.

6.8 UPSTREAM TRANSMISSION EFFICIENCY

In the early DOCSIS days, some of the thinking regarding how the narrower channels (200 KHz & 400 KHz) would be used, was in spectrum regions where narrowband interferers were not present. This approach to work around the interferers did not materialize in practice because cable operators always preferred to use the clean portion of the upper upstream spectrum or to split nodes. The use of narrower channels in the gaps of clean spectrum would have been operationally taxing on operators as each node would have unique spectrum regions and operatrors would have to customize each channels modulation profile and detail channel parameters. Because of this complexity operators have concentrated in operating in the cleaner upper portion of the spectrum at the modulation efficiency that the aggregate noise in that channel would allow. Figure 14 shows an example of how upstream modulation schemes are selected following this approach. That is, using higher order modulation schemes in the upper portion of the spectrum and more robust schemes at the lower portion of the spectrum where noise level are typically higher.

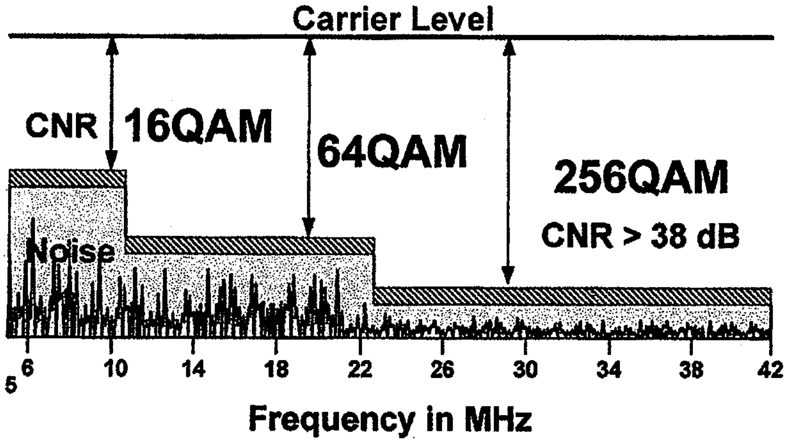

Figure 14 Sample US Modulation Order Selection in Legacy DOCSIS

In a scenario were the transmission system consists of many very narrow subcarriers that can in principle be silenced based on whether the CNR level is acceptable a more efficient use of resources is possible. A key characteristic of this proposed system is that the selection of what carriers to use or not is performed by the PHY layer. This way the MAC layer is not burdened by this process and the operator is not burden by the configuration and management of subcarriers either. Figure 15 highlights how with very narrow subcarriers a system can work around the narrowband interferers. This represents a minimum overhead from the subcarriers not used compared to the large number of subcarriers used at high modulation order.

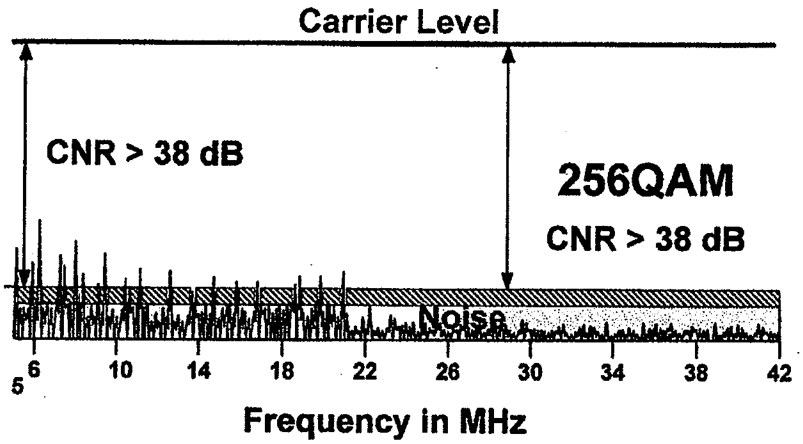

Figure 15 US Modulation Order in Narrow Subcarrier Implementation

6.9 OFDM SYMBOL

The OFDM symbol duration that this approach assumes for the 10 KHz subcarriers is in the order of 100 microseconds. This long symbol is much longer than the 2 microseconds micro-reflections that you see in the cable environment. Also, even if the guard time equal to 1/32 of the symbol duration, the guard time results longer than the micro-reflection. A cyclix prefix in the guard time can make the system more robust although it may not be needed. Figure 13 shows the time representation of a symbol including guard time. It also how a reflected symbol delayed by up to the guardtime does not impact the next symbol.

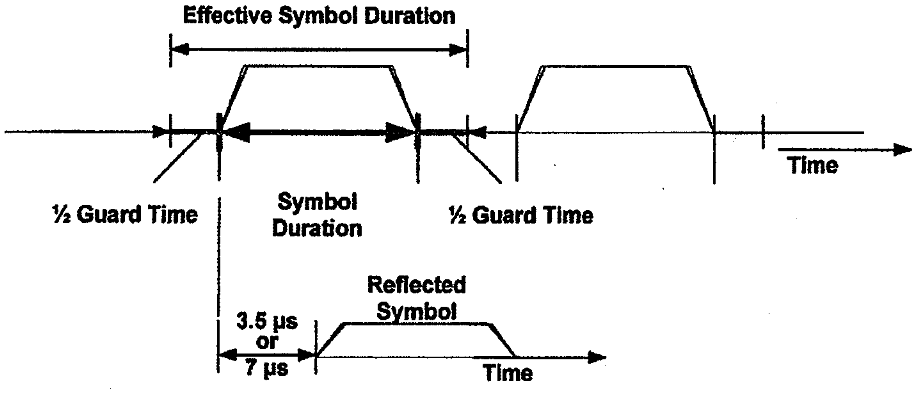

Figure 16~~161616~~ Effective OFDM Symbol Duration

Formatted: English (U.S.)
Field Code Changed
Formatted: English (U.S.)
Formatted: English (U.S.)

Even though the subcarrier spacing of 10 KHz implies a symbol period of 100 microseconds, the effective symbol period has to include the guardtime. In that case the effective symbol period and effective symbol rate is given respectively by:

$$T_{Eff} = T + T_G,$$

$$SR_{Eff} = 1/T_{Eff}.$$

For the scenario of an symbol period of 100 microseconds and a guard time of 1/32 the the effective symbol rate is 9.6969 KHz.

6.10 OFDM PARAMETERS

Additional OFDM parameters can be defined from the characteristics of the environment available, expected impairment characteristics and CNR available. Also the target performance levels and or service target as well as technology characteristics can be used to zoom into design parameters. The 10 KHz subcarriers spacing was chosen to be robust against narrowband interferers and to support a simple narrowband ingress mitigation technique. At the same time the resulting 100 microsecond symbol period provides robustness against micro-reflections and impulses and burst noise. Initial upstream spectrum analysis indicates that probably around 20 subcarriers will need to be silent due to narrowband interferers and a conservative assumption of 116 subcarriers that would not meet 256 QAM CNR transmission levels as well as 660 subcarriers intended for legacy DOCSIS transmission could also remain unused. This leaves 2904 subcarriers available. Using an FEC coding rate of 15/16 and one pilot tone for every 32 tones we have a total of 2640 subcarriers that can be provided to the MAC layer for use in the 5-

42 MHz scenario. At a modulation of 256 QAM this represent a capacity of 204.8 Mbps (Figure 17).

Modulation	256-QAM
Coding Rates	15/16
# of Subcarriers (inc. Pilot Tones)	2904 (for 15/16)
# of Pilot Tones	88
Subcarrier Spacing	**10 kHz**
Symbol duration (no guard time)	**100.0 µs**
Guard Interval (1/32)	**3.125 µs**
Eff. Symb. duration (w guard time)	**103.125 µs (Eff SR = 9.6969 KHz)**
Aggr. Data Bandwidth	**29.04 MHz**
# of Silent Subcarriers	136
Capacity Available	**204.8 Mbps = (15/16)*2816*(9696.96)*8)**

Figure 17 OFDM Configuration Parameters for Upstream FFT Blocks 5- 42 MHz

Modulation	256-QAM
Coding Rates	15/16
# of Subcarriers (inc. Pilot Tones)	3960 (for 15/16)
# of Pilot Tones	120
Subcarrier Spacing	**10 kHz**
Symbol duration (no guard time)	**100.0 µs**
Guard Interval (1/32)	**3.125 µs**
Eff. Symb. duration (w guard time)	**103.125 µs (Eff Symb Rate = 9.6969 KHz)**
Aggr. Data Bandwidth	**39.6 MHz**
# of Silent Subcarriers	136
Capacity Available	**279.2727Mbps = (15/16)*3840*9696.96*8)**

Figure 18 OFDM Configuration Parameters for Upstream FFT Blocks > 42 MHz

Figure 18 goes through the same exercise for the FFT blocks that occupy the spectrum above 42 MHz which could be available after a migration to a new split. In these upstream scenarios there is no need for operation around legacy device transmission and since the RF environment is expected to be much cleaner and free of narrowband interferers the 136 silent subcarriers could be reduced or used to eliminate transmission in sensitive spectral regions such as aeronautical frequencies. Using these modified assumptions the PHY could provide the MAC layer with 279.2 Mbps for each of the upper FFT blocks being used. As four upstream blocks are aggregated a total of 1042 Mbps is provided to the MAC.

In the downstream direction higher modulation order are expected. Assuming a fixed 1024-QAM scenario it is estimated a total capacity per FFT block of 349 Mbps (Figure 19). The aggregate of three of these blocks would result in a capacity of 1047 Mbps. If all the downstream spectrum up to 1 GHz is available a total of 19 FFT would be able to fit providing a capacity graeter than 6.5 Gbps.

Modulation	1024-QAM
Coding Rates	15/16
# of Subcarriers (inc. Pilot Tones)	3960 (for 15/16)
# of Pilot Tones	120
Subcarrier Spacing	**10 kHz**
Symbol duration (no guard time)	**100.0 µs**
Guard Interval (1/32)	**3.125 µs**
Eff. Symb. duration (w guard time)	**103.125 µs** (Eff Symb Rate = 9.6969 KHz)
Aggr. Data Bandwidth	**39.6 MHz**
# of Silent Subcarriers	136
Capacity Available	**349.09 Mbps = (15/16)*3840*9696.96*10)**

Figure 19 OFDM Configuration Parameters for Downstream FFT Blocks

7 CONVERGENCE SUBLAYER

7.1 MAC – PHY DECOUPLING

A convergence sublayer is proposed in this approach in order to decouple the PHY layer from the MAC layer. There is almost no PHY parameters that ist passed to the MAC layer. The intend is to provide the MAC with a constant amount of resources, regardless of the varying conditions the PHY layer may be experiencing.

The downstream MAC layer will pass to the PHY upstream channel descriptor (UCD) information that indicates which subcarriers are expected to be silent and which subcarriers will be carrying data. This design allows in principle different types of UCDs in case a future PHY layer is incorporated. This particular UCD is understood just by the proposed PHY layer. In this environment alternating groups of subcarriers that are silent with groups of subcarriers that carry data are expected. The number of silent subcarriers, are expected is expected to be significantly lower than the number of subcarriers that carry data. The UCD message is updated periodically as the subcarriers that are used may vary dynamically depending on plant conditions.

In the convergence sublayer, based on the UCD information the convergence layer knows which subcarriers it will have available for transmission. Moreover based on this information it will know at a particular time which subcarriers will be dedicated for forward error correction and which subcarriers will be dedicated to be used as pilot tones. Using a known pre-determined pattern the ES knows how the use of subcarriers dedicated for error correction and pilot tones change in time. Figure ## shows a simplified representation of that time dependant process.

The MAC is only receiving a constant number of minislots or data elements and is unaware of the PHY configuration parameters and adjustments in time.

There are some static configuration alternatives that are allowed by this approach which are based on operator's assessment on 3 criteria:

1. How clean the plant is
2. How much legacy is intended to support
3. What is the split configuration

The first criteria item is probably the one, operators are going to be quite certain about as the lower number of end stations per node imply a smaller serving area and reduced number of actives in cascade as well as plant that is easy to troubleshoot and managed. In the examples shown earlier 136 subcarriers are assumed to be dedicated to frequency regions where optimum, 256-QAM, transmission is not feasible due to narrowband interferers or higher noise levels. The second criteria item relates to the transition strategy. DOCSIS legacy may still be used to support certain service tiers. If for example the highest rate service tier to be supported has a peak upstream rate of 30 Mbps or less that a single 6.4 MHz DOCSIS channel would have to be supported. Depending on how much upstream legacy is needed the amount of reserved spectrum can be adjusted.

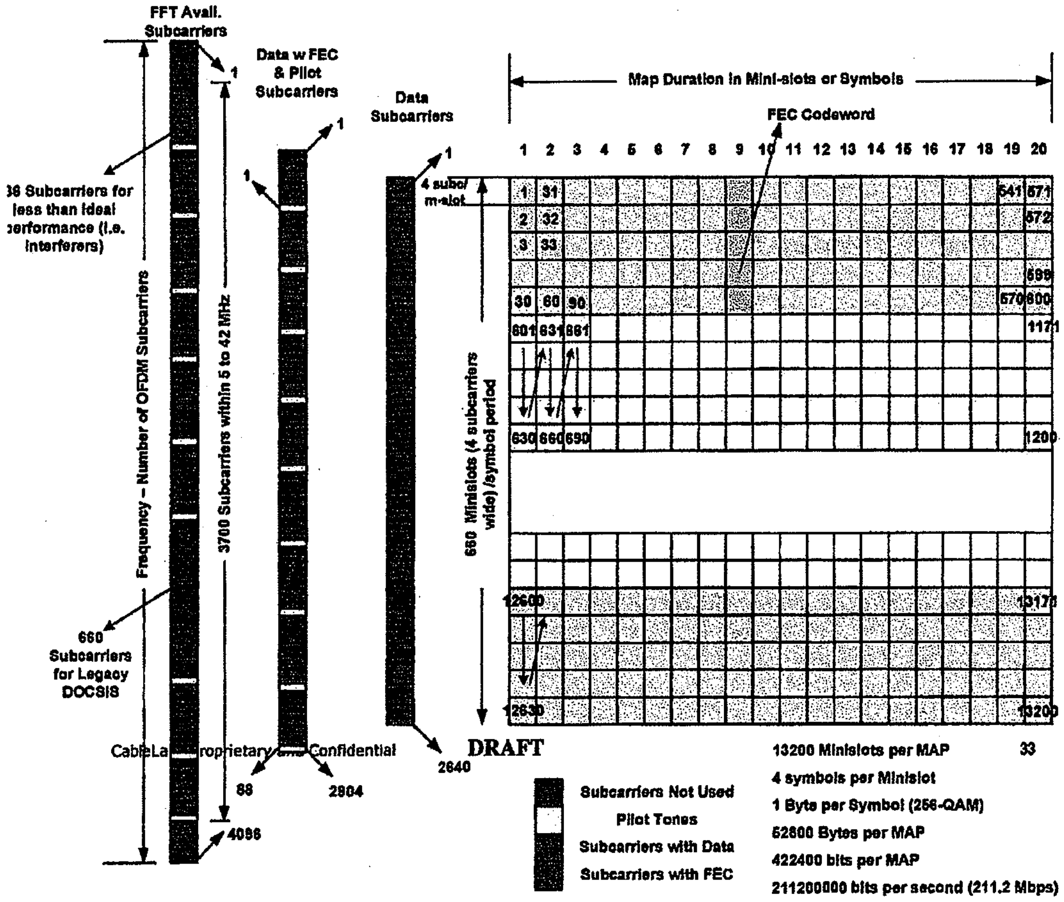

Figure 20 Convergence Layer Decoupling of MAC & PHY

a

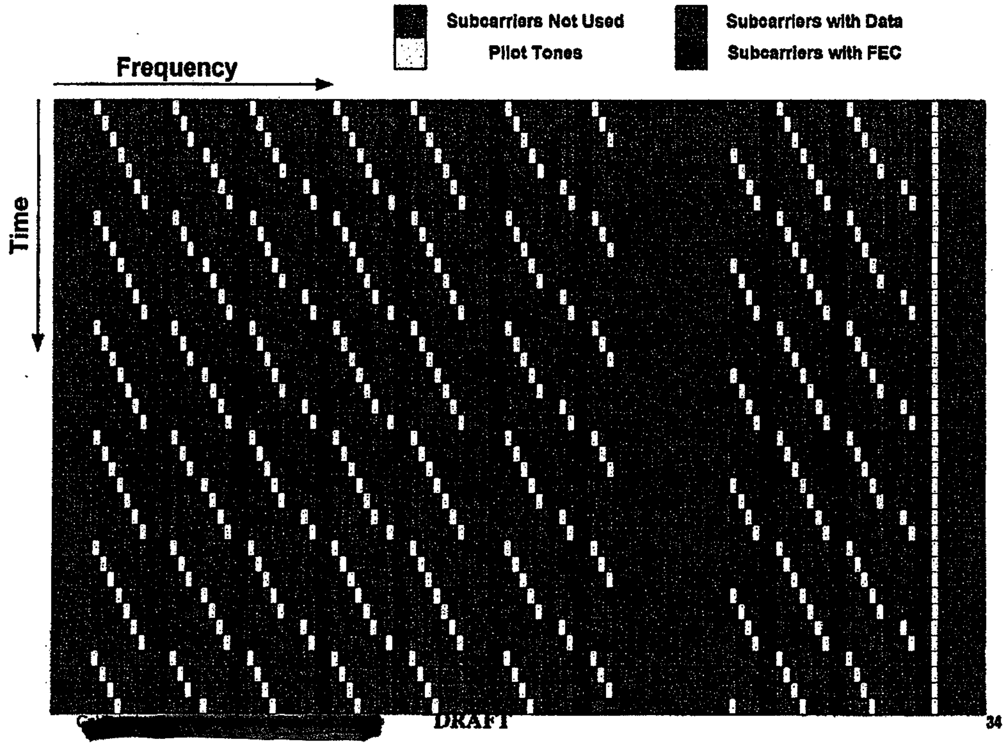

Figure 21 Pilot Tones shift in frequency with time to cover spectrum with high granularity

CABLELABS®

MAC Layer Design for Efficient Multi-Gigabit Transport over CATV Networks

prepared by

Jennifer Fang, Ph.D.
L. Alberto Campos, Ph.D.

[redacted] **DRAFT**

Table of Contents

1 INTRODUCTION

The cable television industry has benefitted from the delivery of data services over their infrastructure for over 14 years. The deployment of cable modems has been widespread and operators have been continuously improving the health of the hybrid fiber coaxial network and therefore their transport characteristics. DOCSIS®, the main mechanism of data transport, has evolved through 4 generations where transport efficiency, peak rates and total capacity have steadily increased. The most dramatic increase in data rates came about with the advent of DOCSIS 3.0, where, through channel bonding, operators have the capability to offer services with peak rates of at least 160 Mbps in the downstream direction and 120 Mbps in the upstream direction.

At the same time, the amount of traffic from the data services that utilize the CATV networks in both directions has been increasing at a rapid rate, doubling every three years [SHAW]. It is projected that peak service rates will reach the gigabit per second range in the 2015-2020 timeframe in the downstream direction. In the upstream direction, the demand has also been growing rapidly, albeit at a lower rate of doubling in the last five years. However, if the trend continues, and with the ever increasing popularities of peer-to-peer applications which accounts for three quarters of the traffic on the upstream, the bandwidth demand will reach the half gigabit per second mark around 2020 [SHAW].

While in DOCSIS 3.0 the use of additional bonded channels as well as node splitting may be able to ameliorate the need for capacity in the downstream direction for up to the gigabit per second range, the limited spectrum of 5 to 42 MHz on the upstream presents an insurmountable challenge to provide enough resources to meet the expected demand. It is also questionable whether a DOCSIS based system will scale well in the gigabit per second range in the upstream direction.

Realizing the expected explosive growth of consumer demand, and the limited ability of the existing DOCSIS 3.0 system to provide enough bandwidth resources in a cost efficient manner, operators have defined the following requirements a next generation PHY and MAC architecture would have to meet.

Categories	Requirements
Physical Plant	Hybrid fiber coax Network
Coexistence	Must coexist with legacy DOCSIS systems
Service Goals	Downstream: multiples of 1 Gbps
	Upstream short term: 200 Mbps (5-42 MHz option)

Categories	Requirements
	Upstream short term: ≥400 Mbps (>42 MHz option), with a goal of 1 Gbps
Complexity	Need to keep everything as simple as possible
	Minimize cost of chips, products, and operations
PHY Layer Requirements	Robust against narrow band interferers & impulse noise
	Single, constant PHY configuration presented to MAC
MAC Layer Requirements	Max number of end stations per MAC Domain ≤ 256[1]
	Granular classification of traffic flows

Table 1 – Next Generation Architecture Requirements

This report presents a potential design for next generation data services over the CATV infrastructure, with focus on the MAC layer. A separate report [AMP-PHY] provides details of the PHY layer design. They PHY assumptions serve as the basis for the MAC layer design proposed here. To design a system that achieves the cable operators' service goals, and at the same time that is low cost, robust, efficient, scalable and simple to implement and operate, assumptions are made to ensure the new architecture can coexist with the existing DOCSIS systems but does not need to be backward compatible. This new MAC architecture presented in this report represents a clean slate without the constraint of previous assumptions of protocols, technologies and tools that were incorporated in the design of the existing DOCSIS systems. While comparisons will be made with the legacy DOCSIS systems, the next generation architecture presented here is not a gradual evolution from them.

In addition to conforming to the list of design requirements and service goals defined by the cable operators, the MAC layer architecture presented in this report follows two additional design principles:

- The MAC and PHY are not tightly coupled.
- This system has been designed with modularity in mind to provide scalability and configuration flexibility based on operators' network migration paths and service goals. Multiple upstream and downstream configurations are possible but only a couple of representative ones are analyzed in detail: one for a transition phase that relies on a 5-42

---

[1] This does not necessarily mean a Node + 0 plant requirement. Most likely it will include N+1 and N+2 scenarios as well.

MHz upstream, and a second which is an example next generation configuration using a modified upstream split.

These assumptions and principles lead to a dramatically simplified MAC layer design intended to help reduce implementation and operations costs.

In this report, upstream and downstream assumptions and design proposals are treated separately. This provides the operators with the flexibility to implement either direction along separate timeframes.

The rest of the report is organized as follows: Section 2 describes the PHY layer assumptions that impact the MAC layer architecture; Section 3 presents several key MAC layer assumptions and their implications; Section 4 describes in detail the proposed architectural elements such as ranging, minislot structure, upstream requesting, downstream transmission, and quality of service, and makes a brief comparison to the EPON MAC architecture; and the conclusion Section 5.

2 PHYSICAL LAYER ASSUMPTIONS

Advances in signal processing have enabled highly granular segmentation of spectrum. This results in long symbol periods that provide significant robustness against impairments. Orthogonal Frequency Division Multiplexing (OFDM) is a well known technology that takes advantage of DSP technology advances and drastically reduces complexity that is present in wide channel implementations. While there may be multiple PHY transmission systems that meet the requirements detailed in the previous section, OFDM is used in this report due to its high performance at a low implementation complexity.

Details of the PHY layer characteristics are discussed in the PHY [redacted]port [AMP-PHY]. The following is a list of upstream PHY characteristics that impact the MAC architecture design:

- The end station (ES) and the aggregating unit (AU) utilize OFDM as the transmission scheme.
- The US frequency spectrum is occupied by OFDM subcarriers generated by one or more 4096 FFT blocks. Many upstream split configurations based on various deployment scenario can be achieved, thanks to the flexibility and modularity of the implementation:
    - For example, operators may want to utilize only the 5-42 MHz frequency range during transition period to help bridge the gap between existing DOCSIS network and higher upstream split. In this case, a single 4096 FFT block can be used.
    - In the modified US split mode where the 5-185 MHz frequency block is available, 4 FFT blocks can be used to carry OFDM subcarriers. Likewise, an 85 MHz split can be implemented using two 4096 FFT blocks.
- The separation between subcarriers is 10 KHz.
- The upstream PHY layer allows one legacy DOCSIS channel anywhere within the 5–42 MHz frequency range.
- The OFDM subcarriers and channel blocks are aggregated either at the PHY layer or at a convergence sublayer between the MAC and the PHY layers. This enables the PHY layer or the convergence sublayer to present the MAC layer with a single US channel, compared to the multiple US channels as in the DOCSIS 3.0 MAC layer bonding architecture. From a MAC perspective, its set of functions is only performed over a single, large capacity entity.
- The upstream modulation scheme is 256QAM.

In the downstream direction, the following PHY characteristics are assumed:

- The ES and the AU utilize OFDM as the transmission scheme.

- In the downstream direction, an available frequency range between 250 MHz – 1 GHz is assumed, where 4 to 16 4096 FFT blocks are used at the AU to generate 1-4 Gigabits per second transmissions. Since the implementation is modular, alternative frequency ranges can be accommodated. The separation between subcarriers is 10 KHz.

- The OFDM subcarriers and channel blocks are aggregated either at the PHY layer or at a convergence sublayer between the MAC and the PHY layers. This enables the PHY layer or the convergence sublayer to present the MAC layer with a single DS channel, compared to the multiple DS channels as in the DOCSIS 3.0 MAC layer bonding architecture. From a MAC perspective, its set of functions is only performed over a single, large capacity entity.

- The downstream uses more efficient modulation scheme of 1024QAM.

The following physical plant characteristics are assumed:

- CPEs are located at the boundary of subscriber home to eliminate in-house power variability of up to 12 dB.

- This next generation MAC layer domain is intended to support a maximum of 256 ESs and a much shorter delay difference between end-stations compared to DOCSIS.

3 MAC LAYER ASSUMPTIONS AND ASSOCIATED IMPLICATIONS

3.1 UPSTREAM

3.1.1 SUBCARRIER ALLOCATIONS

In the upstream, each subcarrier can be used to carry data or serve one of the following purposes controlled by the PHY or the convergence sublayer:

- Pilot Tones: to maintain frequency orthogonality between subcarriers it is necessary to add Pilot Tones. Any type of distortion that impacts an accurate determination of the subcarrier frequency can result in a loss of orthogonality and system degradation. Pilot tones serve as frequency reference for the subcarriers.

- Silent Tones: narrowband interferers may be present in the frequency spectrum that is occupied by the cable access network. To avoid using subcarriers with poor performance, those subcarrier(s) that coincide with the narrowband interferers may be excluded. In addition, subcarriers that coincide with sensitive transmissions such as aeronautical frequencies can also be silent.

- Reserved: some frequency spectrum used for US transmission may optionally be reserved for usage by non-OFDM/DOCSIS CMs. The amount of spectrum to be reserved and its location is configurable by the operator, and is communicated to the ES via the use of Upstream Channel Descriptor (see Section 3.1.2).

- FEC: to improve robustness to data traffic, FEC corrections need to be added.

The PHY layer is responsible to keep track of how data, FEC and pilot subcarriers are mapped to the upstream frequencies. Figure 1 shows an example of the frequency mapping for the 5-42 MHz frequency block where data subcarriers are shown in green.

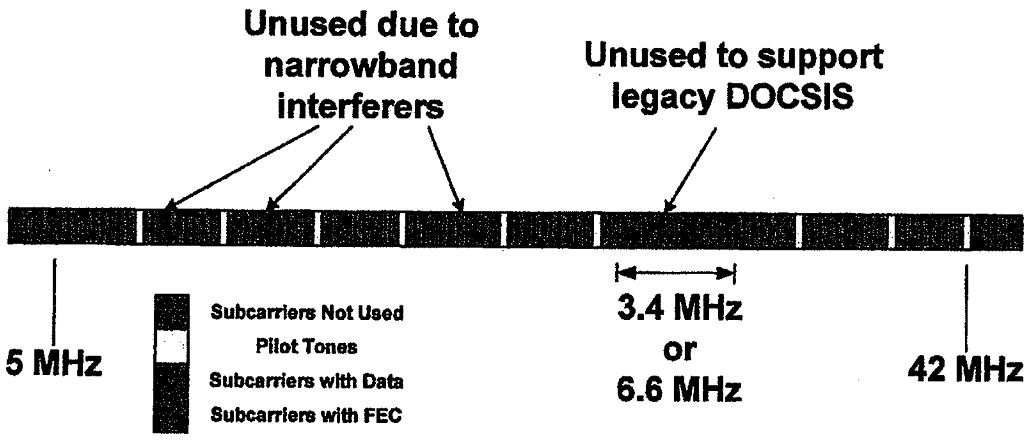

Figure 1 – An Example of Upstream Frequency Configuration for Transition Deployment (not to scale)

3.1.2 UPSTREAM CHANNEL DESCRIPTOR (UCD)

Upstream Channel Descriptors are necessary to indicate the position and the range of the subcarriers that are silent to support narrowband interferer mitigation and legacy operation. They are transmitted by the AU at fixed time intervals. To ensure that ES of any capabilities can receive the UCD, the AU should be transmitted in the downstream FFT block that is defined by the minimum downstream channel configuration option.

3.1.3 CHANNEL CONFIGURATION OPTIONS

Operators may configure variable sized spectrum to customize their next generation operation scenarios. However, there is a fixed set of number of options for the different FFT block capability scenarios that an ES and an AU have to support. This provides flexibility to operators in planning services while limiting design complexity. The total upstream throughput will vary depending on the upstream configuration scenario.

In the remainder of this section, we present an upstream configuration option to match a potential deployment scenario where the upstream split is extended to 185 MHz, with a block of 6.4 MHz spectrum in the 5-42 MHz frequency range reserved for DOCSIS operation. Since each subcarrier occupies 10 KHz of spectrum, there are 3700 subcarriers in total for the 5–42 MHz spectrum that may be used in the upstream. Among these subcarriers:

- 136 subcarriers are reserved to remain silent mostly due to narrowband interferers or spectrum locations with sub-optimal CNR. The number of subcarriers provides an upper bound on the number of subcarriers expected to be silent, and was chosen to match the assumed MAC & PHY configuration parameters such as number of FEC corrections and pilots stated below, and still serve to provide robustness for a diverse set of plant conditions.

- A 6.6 MHz block of spectrum is reserved for legacy DOCSIS CMs, including 100 KHz frequency guard band on each side. This amounts to 660 subcarriers reserved for a legacy system.

The remaining subcarriers are used for data, FEC, and pilots with the following configurations:

- 2 subcarriers for every 30 subcarriers transmitted are used for forward error correction (FEC). FEC setting is 15/16.

- 1 pilot tone is transmitted for every 32 subcarriers used for data transmission.[2]

The FEC and pilot tone setting is not unique to this example, and is applicable to all upstream channel configuration scenarios. Within the remaining spectrum in the 42-185 MHz range, 3 additional FFT blocks will fit. For each of the 3 upper FFT blocks, the same assumptions as the lower block for the transition period apply, except for the requirement on frequency reservation for legacy operations. This configuration option is summarized in Table 2. The number of silent tones for the upper blocks can be reduced for the upper blocks to match the operators' deployment plans.

	Lower Block of 5-42 MHz, with6.4 MHz is reserved for DOCSIS	Each of the 3 Upper Blocks
Total number of subcarriers	3700	4096
Reserved for legacy systems	660	-
Silent tones	136	136
FEC	176	240
Pilot tones	88	120
Data subcarriers remaining	2640	3600

Table 2 – A Configuration Option for Modified US Split with 6.4 MHz DOCSIS Reserve

3.2 DOWNSTREAM

In the downstream, the subcarrier allocation is similar to the upstream except the requirement o the frequency spectrum reserved for legacy deployment is removed. Silent tones are necessary to remove some OFDM subcarriers that coincide with aeronautical transmissions or other sensitive types of transmissions. However, broadcast messages similar to the UCD that are used to indicate the blocked spectrum for the upstream are not necessary for the downstream. Instead, the AU PHY layer can transmit standardized tones just before and after the silent subcarriers to indicate the range of the blocked spectrum. Since the ES may also have varying downstream capabilities, a fixed set of number of options for the different FFT block capability scenarios that an ES and an AU have to support can also be defined.

---

[2] This can also be modified to 64 subcarriers or others, depending on PHY layer capability.

4 KEY MAC LAYER FEATURES

4.1 RANGING

In DOCSIS systems, once the initial operation parameters have been determined after the process of broadcast initial ranging, the CMTS initiates unicast initial ranging with the CM to fine tune the timing and power offsets, among other parameters. Because timing reference and power level may drift as time goes on, the CM utilizes periodic station maintenance to correct the drift. In this section we consider in the context of timing offset whether periodic station maintenance is necessary in the next generation architecture. For discussions on other parameter adjustments such as power level and equalization, see the PHY report [AMP-PHY].

4.1.1 PERIODIC STATION MAINTENANCE

The timing offset is adjusted to compensate for variations as a result of expansion and contraction of the coaxial plant due to temperature changes. This timing variation will cause no inter-symbol interference through an appropriate selection of guard time.

Since each OFDM subcarrier size is fixed at 10 KHz per PHY layer assumption, the useful symbol period is 100 μsec. A typical guard time setting is 1/32 of the useful symbol period, which translates to 3.2 μsec. This is the amount of time it takes for the RF signal to travel 2740 feet[3] in the coaxial environment. The amount of plant contraction and expansion due to temperature change represents only a small fraction of the round trip distance traveled by the RF signal and is easily absorbed by the guard time. Therefore periodic timing adjustment that is usually performed during station maintenance in DOCSIS can be safely omitted.

We also note that since the periodic timing adjustment only corrects the variation due to temperature change regardless of plant configuration, removing this procedure does not place any restriction on the HFC plant node size.

4.1.1.1 Micro-Reflections Compensation

Impedance mismatches between devices in the coaxial cable network such as amplifiers, couplers, and the coax cable produce micro-reflections. As shown above, with OFDM subcarrier size of 10 KHz, and guard time setting of 1/32 of useful symbol time, the guard time is then 3.2 μsec. This is larger than the micro-reflections measured in a typical coaxial plant, which is less than 2 μsec [AMP-PHY].

---

[3] A propagation factor of V_prop = 0.87*C has been used. See [PHY-3.0] Annex III.2.2.

4.1.1.2 Bandwidth Saving

DOCSIS CMTS's reserve a portion of the upstream bandwidth to perform periodic maintenance with the CMs. Table 3 shows the average percentage of minislots reserved for periodic maintenance (IUC4) as a portion of the total number of minislots granted, under moderate amount of load on the upstream. CMTS's from 4 vendors were tested, with various PHY level setups, such as 3.2 MHz and 6.4 MHz channels, varying modulation schemes of QPSK or 64 QAM.

CMTS	IUC3 (broadcast ranging)	IUC4 (periodic maintenance)
vendor 1	2.2 %	< 0.15 %
vendor 2	< 1.5 %	< 0.1 %
vendor 3	TBD	TBD
vendor 4	TBD	TBD

Table 3 – IUC 4 Overhead

4.1.1.3 Recommendations

Based on the discussion presented in this section on the viability of removing periodic maintenance, as well as the bandwidth usage data collected from various CMTSs, we conclude with the following observations:

1. Since a typical guard time setting can compensate well for changes in the coaxial plant due to temperature variations, as well as for expected micro-reflections, the process of timing adjustments that is performed as part of periodic station maintenance should be removed.

2. Power adjustment that is performed during station maintenance in DOCSIS may still be needed in order to achieve optimal power level and received SNR required for an efficient upstream modulation scheme. This may be done at infrequent time intervals. See [AMP-PHY] for greater detail.

3. While station maintenance does not take up significant bandwidth on the upstream, removing or simplifying it will reduce MAC layer complexity.

4.1.2 INITIAL RANGING

Initial ranging acquires the correct timing offset between the ES and the AU at the time the ES registers, whereas station maintenance corrects the variations in delay after the initial registration. Since most of the delay variation after registration is caused by the temperature effect on the coaxial portion of the plant, it is expected that this delay variation will be reduced as fiber extends deeper in the network. However, the new plant environment will not have any effect on initial ranging. Therefore, it is still needed to obtain the initial timing correction.

4.2 REGISTRATION

The MAC layer proposed in this report does not impose any restrictions regarding the registration process that takes place to allow access to the network.

4.2.1 END STATION CAPABILITY REPORTING

During registration, each ES reports its upstream capabilities, including number of FFT blocks it supports on the upstream, as well as on the downstream. This will allow the AU to grant and transmit on the region of the MAP where the ES is capable of utilizing.

4.3 CHANNEL AGGREGATION

4.3.1 UPSTREAM

The existing DOCSIS systems utilize multiple upstream channels, either simultaneously as in DOCSIS 3.0, or one at a time with the ability to change channels as in pre-DOCSIS 3.0. The following MAC layer features are designed to manage the existence of multiple channels:

- Dynamic Channel Change
- Dynamic Bonding Change
- Load balancing
- Topology resolution
- CM-STATUS
- Logical upstream channels

As stated in Section 1, the next generation MAC layer architecture, as well as the transition architecture, assumes no backward compatibility with DOCSIS 1.x, 2.0, and 3.0 systems. In addition, the OFDM subcarriers and channel blocks are assumed to be aggregated either at the PHY layer or at a convergence sublayer between the MAC and the PHY layers, as stated in Section 2. These assumptions remove the need to keep track of multiple channels or perform channel aggregation at the MAC layer, as in existing DOCSIS systems. The proposed next generation MAC layer assumes that the task of channel aggregation is performed at the PHY layer or the convergence sub-layer, in order to present the MAC layer with a single US channel. Therefore, the DOCSIS features required to manage multiple channels at the MAC layer are not present in the new architecture. The result is a dramatically simplified MAC layer compared to existing DOCSIS systems.

4.3.2 DOWNSTREAM

Similar to the upstream, the existing DOCSIS systems utilize multiple channels in the downstream direction, and channel bonding and channel change are performed at the MAC layer level. In the next generation architecture, since either the PHY layer or the convergence sublayer is assumed to take on the task of aggregating OFDM subcarriers and channel blocks, and no backward compatibility is assumed, the bonding and channel change related functions at the MAC layer can be removed. These include:

- Dynamic Channel Change
- Dynamic Bonding Change
- Load balancing
- Topology resolution

4.4 MINISLOT AND ALLOCATION MAP STRUCTURE

Similar to the existing DOCSIS systems, the AU in the next generation MAC architecture utilizes allocation MAPs to grant each ES slots to transmit its bandwidth requests. However, unlike the existing systems where the allocation MAP is a time-based linear description, the new system employs a two-dimensional description that provides the flexibility to allow each ES to be allocated specific time slots (TDMA), as well as a subset of OFDM subcarriers (OFDMA[4]).

In this report, a MAP duration of 20 symbol periods, or approximately 2 ms is assumed. This setting can be optimized for each bandwidth request methods. We define a minislot as 1 symbol period in duration, and 4 data subcarriers wide, or a total of 4 symbols. The FEC setting is 15/16 with a FEC codeword of 32 symbols, or 8 minislots.

As in DOCSIS systems, the next generation MAC uses the concept of minislot to define the unit of granularity for upstream transmission opportunities. Instead of granting and transmitting in units of subcarriers, using minislot enables the MAC layer to be more decoupled from the PHY. Since the MAC layer is only aware of transmission opportunities in terms of minislots, instead of the underlying technology, it allows the potential adoption of other PHY layer technologies than the proposed OFDM based approach.

---

[4] Orthogonal Frequency Division Multiple Access (OFDMA) is a multi-user version of OFDM, where multiple access is achieved by assigning subsets of subcarriers (aka, subchannels) to users. This allows simultaneous transmissions from multiple users.

Figure 2[redacted]shows a visualization of PHY and MAC allocation maps for a potential transition deployment scenario where only the 5-42 MHz spectrum range is utilized. The 2 vertical maps on the left display the PHY layer mappings between various subcarriers and transmit frequencies. The MAC layer does not perform any frequency mapping. The third vertical map is shown to merely indicate that the MAC layer is only aware of the data transmit opportunities, in minislot form, which is shown through the 2-dimensional MAP on the right that represents the minislot structure for the entire MAP duration.

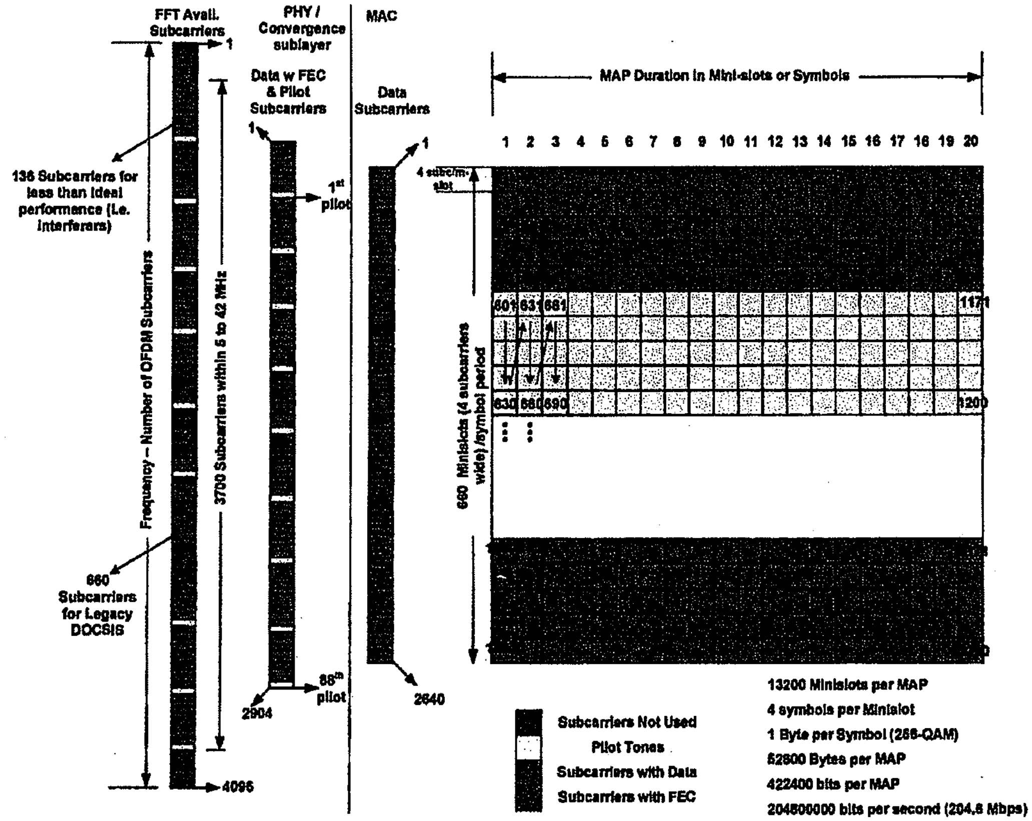

Figure 2 – Frequency Mapping and MAP for Transition Deployment

If operators choose to utilize only the 5-42 MHz frequency range during transition period, with 6.4 MHz reserved for DOCSIS operation, based on the data subcarriers computed for the in Table 2, there are 660 x 20 minislots per MAP for this scenario. We further assume that the modulation type is 256 QAM, or 8 bits / symbol. Then an ES can achieve a maximum upstream MAC throughput of 204.8 Mbps during the transition period scenario described here.

For a next generation deployment scenario where higher upstream split of 5-185 MHz is available, there are 900 x 20 minislots per MAP, for each of the 3 upper FFT block beyond 42 MHz. With same modulation assumptions as above, an ES can achieve a maximum upstream

MAC throughput of 1.042 Gbps, when using 4 FFT blocks with a 6.4 MHz block of spectrum in the lower block reserved for DOCSIS.

Since a number of the data subcarriers will be used to make initial bandwidth requests, the MAC throughput may be reduced, depending on the bandwidth request method.

4.4.1 MINISLOT NUMBERING

During registration, each ES reports its upstream capability. Although each ES may support different number of upstream FFT blocks, the AU still transmits the entire MAP that covers the spectrum supported by the most capable ES in a single MAP message.

Each MAP contains a counter that is incremented, and starts with minislot number one. The minislot numbering shown in Figure 2~~Figure 2Figure 2~~ provides an example of the numbering strategy[5]. During registration, the ES's MAC layer is responsible to compute the max minislot number it should read, based on its FFT block support capability. Alternatively the maximum number of minislots that the ES will be able to use is computed by the AU, and communicated to the ES during the registration process.

4.5 REQUEST FRAMES

Similar to DOCSIS systems, the next generation MAC adopts the concept of queue-depth based request. Each ES is required to request in multiples of 30 minislots, or 120 symbols.

Figure 3 – Next Gen Initial Request Frame w/ End Station Enforced QoS

Figure 4 – Next Gen Initial Request Frame with either Prioritized or Service Flow Based QoS

---

[5] When granting multiples of 30 minislots, the AU can either grant a single block of 30 minislots over multiple symbol periods, or granting multiple blocks of 30 minislots over a single symbol period. Using the former, the AU can take advantage of the pilot tones transmitted over the same frequency blocks. This may provide higher robustness compared to the latter method. The numbering system proposed here supports the former granting method more seamlessly. For more discussion, please see [AMP-PHY].

If the next generation MAC utilizes end station enforced QoS (for details see section below on Quality of Service), the SID field or a queue number indicator is not necessary. In this case, the request frame for a single queue reporting is shown in Figure 3~~Figure 3Figure 3~~, and is 4 bytes: 12 bits for requested amount, 1 byte for MAC layer FEC with the setting 2/3, and 1 sync byte using a known bit sequence. Then the 12-bit length for requested amount ensures the ES can request up to 1.9 Gbps.

When either service flow based QoS or prioritized QoS are provided, the 4-bit reserved field in the end station enforced QoS scenario can now be used to indicate queue number or SID field, and is included in each request frame, as shown in Figure 4~~Figure 4Figure 4~~. The requested amount is a 12 bit field and allows the ES to request up to 1.9 Gbps per queue or service flow. Multiple queue or service flow requesting is discussed in the QoS section below.

In both cases, the FEC and sync byte are both assembled at the MAC layer. Instead of transmitting pilot tones with each request frame in order to verify signal integrity, MAC layer sync byte is used. This enables the request frames to be allocated anywhere in the MAP, rather than requiring to be lined up with the pilot subcarriers, providing effective separation between the MAC and the PHY. With the use of 256QAM upstream modulation, the request frame fits exactly into 1 minislot.

As in DOCSIS systems, the next generation MAC allows piggyback requests to be used for bandwidth requesting for subsequent transmissions. The details of the piggyback request structure will be left for a future revision of this document.

4.6 UPSTREAM BANDWIDTH REQUESTING

Existing DOCSIS systems provide an array of network access or scheduling options for upstream bandwidth requests according to traffic types. The Unsolicited Grant Service (UGS) is designed to support VoIP traffic where the CM does not actually request, but is provided with fixed size grants on a periodic basis. The Real-Time Polling Service (rtPS) is designed to support video traffic where the CM is provided with opportunities to request variable size grants on a periodic basis. There are additional options to support VoIP traffic with silence suppression, as well as high bandwidth FTP traffic. Altogether, the existing DOCSIS systems must implement at least 5 scheduling options.

In order to reduce the complexity of the MAC layer implementation, the next generation system proposed in this report utilizes just one scheduling algorithm which when coupled with quality of service and bandwidth request mechanisms, will work well for all types of traffic under the PHY and MAC assumptions specified in earlier sections. This is possible due to the significant increase in the available network capacity in the next generation networks, compared to the existing DOCSIS systems. (See QoS Section below for details)

In this section, we explore four upstream bandwidth request mechanisms for scheduling the transmission of initial bandwidth requests in particular to optimize the network access latencies, and evaluate their comparative merits in the cable access network based on the access delay and bandwidth overhead. (In order to compare the bandwidth overheads for each of the scheme, we assume a single queue reporting) The four proposed here are:

- Dedicated subcarrier (Section 4.6.1)
- Dedicated minislots (Section 4.6.2)
- 2-Tier polling (Section 4.6.3)
- Contention (Section 4.6.4)

All four mechanisms are application agnostic. In the first 2 mechanisms, each customer is guaranteed to be able to transmit their initial request at least once per MAP, regardless of whether there is any traffic queued at the modem. The third mechanism utilizes a polling based method, similar to that is used in the EPON MAC standard. To provide some level of comparative analysis, some assumptions are made for the latter 2 mechanisms on the expected amount of upstream traffic during peak hours.

4.6.1 DEDICATED SUBCARRIERS

This request mechanism dedicates a fixed number of data subcarrier(s) for each customer, regardless of whether the customer has been active or inactive, to transmit the initial bandwidth requests at any time. The number of dedicated data subcarrier for each customer is set to 1 in order to minimize the bandwidth overhead. Each symbol in the request should be sent in serial on a single subcarrier, and therefore does not need to fit into the single minislot discussed in Section 4.4.

As discussed in Section 4.4.1, in addition to transmitting the 3-byte requests, a pilot signal must be transmitted for every packet transmitted upstream. Instead of allocating additional subcarrier as a pilot subcarrier for each customer, the pilot signal can simply be transmitted before the request itself.

After subtracting subcarriers allocated for FEC correction, pilots, narrowband interferers, and legacy DOCSIS 1.x/2.0, the number of data subcarriers during the transition period is 2640. Assuming 1 data subcarrier is dedicated to each of the 256 customers, the bandwidth overhead is 256 / 2640, or close to 10 % of the total upstream bandwidth. The bandwidth overhead is reduced to 2 % for the next generation deployment using a 4 FFT block configuration. The resulting access delay for the initial request is 4 symbol periods.

While this request mechanism provides the least amount of access delay for initial bandwidth requests compared to the other 3, it is a less attractive solution due to the amount of bandwidth overhead incurred during the transition period. Additionally, this approach demands stronger coupling between MAC and PHY.

Enabling multiple queue reporting and/or multiple service flows does not affect the bandwidth overhead.

4.6.2 DEDICATED MINISLOT

Instead of allocating dedicated subcarrier(s) to each customer at all times, this request mechanism guarantees to allocate 1 unicast request opportunity to each customer in every MAP.

The AU includes the bandwidth request allocation in the regular grant MAP. Since we require the ES to request and therefore AU to grant in multiples of 30 minislots as stated in Section 4.4.1, it is desirable that the AU groups the minislots reserved for bandwidth requests accordingly, although it is free to allocate it the 30-minislot groups anywhere within each MAP. The AU allocates 9 groups of 30 minislots for 256 customers for transmitting request frames.

Additionally, due to upstream transmission time and processing delay incurred at the ES and the AU, the AU should allocate the request minislots early on in every MAP, to ensure the access delay achieves a maximum of 1 MAP time, or about 2 msec.

During transition period, the bandwidth overhead incurred is 270 / (660 x 20), or 2 %, while that amount is reduced to 270 / (900 x 20 x 4), or 0.3 % for next generation deployment if a 4 FFT block configuration is used.

When multiple queue reporting and/or multiple service flow is enabled, bandwidth overhead will increase at most proportional to the number of queues and service flows that must be reported.

4.6.3 2-TIER POLLING

Since average bandwidth usage may be much lower than during the peak hours, polling all customers regardless of their activities incurs unnecessary waste of upstream bandwidth. Instea of polling every customer every MAP, the AU implementing the 2-tier polling dynamically classifies each ES as "active" or "inactive" based on vendor-defined criteria, and decreases the polling frequency for "inactive" ESs, thereby conserving upstream bandwidth. The polling frequencies are also vendor-defined.

While this mechanism is independent of the application types and traffic levels, in order to evaluate the bandwidth efficiency of this mechanism and provide some level of comparative analysis with other frameworks proposed in this report, assumptions need to be made on the expected amount of upstream traffic during peak hours. The assumptions stated in Section 4.6.3.1 are required only by 2-tier polling and contention mechanisms. In addition, while each vendor may determine its own optimal activity criteria and polling frequencies, we provide an example analysis in Section 4.6.3.2, in order to compare the bandwidth overhead with other request mechanisms discussed in this paper.

4.6.3.1 Traffic Assumptions

Among the services that the cable operators offer currently, or in the future, the following applications have stringent requirements based on bandwidth, delay or frequent upstream messages:

- Voice over IP
- Channel surfing
- Large file download (usage from ACKs)
- File upload / P2P (may or may not be offered by the operators)

From the above applications only file uploads can take advantage of piggyback request mechanisms, which alleviates the burden on the network. The temporal characteristics of the other applications do not benefit from piggybacking and is the focus of the analysis that follows.

Table 4 – Table 6 present the assumptions used to derive the average number of initial bandwidth requests during peak hours on a fiber node, for the 3 applications listed above, i.e., VoIP, Channel Surfing, and large file download.

	Assumptions and resulting numbers
Customers	256
# of customers subscribes to IPTV services	60% of Customers = 154
# of active subs at peak time	60% of subs = 92
Total # of IP video streams	Total streams = 120, assuming 30% of subs have 2 TVs in-house
Total # of linear video streams	50% of IP video streams = 60
# of clicks / sec / sub	3
Total # of clicks / sec	3 x 60 = 180
Average # of clicks or REQs / MAP	0.4 (MAP ≈ 2 msec)

Table 4 – IP Video Traffic

	Assumptions and resulting numbers

	Assumptions and resulting numbers
Customers	256
# of customers subscribes to VoIP services	50% of Customers = 128
# of simultaneous calls at peak hour	30% of subs = 39
Sampling interval for VoIP Codec in cable environment	20 msec (This is longer than the 2-msec MAP time. Therefore, no piggyback requests can be made, and each VoIP packet requires separate requests)
Average # of REQs / MAP	4 (MAP = 2 msec)

Table 5 – Voice over IP Traffic

	Assumptions and resulting numbers
Customers	256
# of customers downloading (files, movies etc)	20
Download rate / customer	20 Mbps
DS Packets / sec / customer	1647 (1518-byte pkt)
# packets / ACK	4
US ACKs / sec / customer	412
Total ACK rate for all downloaders	8235 ACKs/sec
Average # of ACKs or REQs / MAP	17.3 (MAP = 2 msec)

Table 6 – ACKs Resulting from File Download Traffic

Summing up the number of initial requests per MAP for each application, the expected total number of initial requests per MAP is 22, during peak hours.

4.6.3.2 Polling Criteria

VoIP based telephony is a key service that cable operators offer. The sampling interval of a common G.711 codec is 20 msec, or approximately 10 MAPs. The inter-arrival time of the VoIP packets is long enough to disallow any piggyback requests, and the arrival of each VoIP packet requires fresh bandwidth requests. In order to support interactive applications, network access delays of less than 10 ms are desirable. This means that the AU must poll the "inactive" ESs at a minimum of once every 5 MAPs.

To determine the activity classification criteria, we analyze the file downloading traffic. The ACK inter-arrival time per customer is 2.4 msec, longer than the 2 msec MAP duration. Therefore, piggyback requests are to a good extent also disallowed. However, because the delay in receiving the ACKs can greatly reduce the TCP throughput, polling intervals should be short enough to accommodate ACK traffic. Therefore, we define that an ES becomes "inactive" if it has not transmitted on the upstream in the last 3 MAPs. This should provide fast access time to high speed applications. Even those applications with periodicities longer than a MAP can achieve short network access latencies via request messages.

4.6.3.3 Bandwidth Efficiency

Even though the expected total number of initial requests per MAP was 22 during peak hours, as derived in Section 4.6.3.1, there may be a number of ESs that are active, yet may not have any requests to send. These ESs will need to be polled every MAP for as long as they remain classified as "active". We assume the number of such ESs is 22 in each of the last 3 MAPs. Using the same request frame structure, ie, 1 request packet fits into exactly 1 minislot, and the activity criteria and polling frequencies derived above, the average total number of subcarriers to be reserved to poll all active and non-active ESs per MAP is 487 (0.9 % of bandwidth overhead during transition, or 0.2 % for next generation deployment).

When multiple queue reporting and/or multiple service flow is enabled, bandwidth overhead associated with 2-tier polling will increase, but to a lesser degree than with dedicated minislot scheme. The polling mechanisms can be optimized according to the queue priorities.

4.6.4 CONTENTION

While 2-tier polling provides better bandwidth efficiency, the complexity of the AU may increase as a result of dynamically keeping track of each ES's activity level and polling frequency intervals. In a contention-based bandwidth request method, instead of polling each ES individually, the AU reserves a block of minislots such that any ESs with requests to send may transmit following a contention resolution process, similar to the one implemented in DOCSIS systems.

In order to reduce or eliminate collision of multiple request frames, for each expected request in a MAP, the AU reserves multiple minislots for each expected request. An appropriate collision avoidance factor of N is selected to ensure that the request frames have low probability to collide. Based on the assumptions and analysis from the previous section, the bandwidth overhead for contention-based method is N times as much as that in 2-tier polling.

An important note is that the request frame structure presented in Section 4.4.1 can no longer be used in contention-based mechanism, because the request frame does not allow the AU to identify the sender of the frame. The request size should be 4 bytes, instead of 3, to account for an ID assigned to an ES.

Two-tier polling keeps track of each ES's activity state in real time and polls accordingly. Contention on the other hand keeps of the aggregate traffic activity to determine the number of slots allocated for data and contention. Contention is simpler to implement although it may not result in optimal use of bandwidth.

4.6.5 OBSERVATIONS AND ASSESSMENTS

Table 7 summarizes the comparison of the 4 bandwidth request mechanisms in terms of bandwidth overhead and access delay, when a single queue reporting is enabled.

	Bandwidth overhead, transition deployment	Bandwidth overhead, next gen deployment	Access delay
Dedicated subcarrier	10 %	2 %	4 symbol periods (0.36 msec)
Dedicated minislot	2 % Will be at most proportional to the # of queues / flows	0.3 % Will be at most proportional to the # of queues / flows	1 MAP time (2 msec)
2-Tier polling	0.9 % Increases with # of queues / flows, but to a lesser degree than dedicated mslot	0.2 % Increases with # of queues / flows, but to a lesser degree than dedicated mslot	1-5 MAP time (2-10 msec)
Contention (N=8)	7.2 %	1.6 %	1-5 MAP time at best

Table 7 – Bandwidth Efficiency of Various Initial Bandwidth Request Mechanisms

Two methods presented here, dedicated minislot and 2-tier polling, incur similar amount of bandwidth overhead. When most customers on the node are active with some traffic to send, 2-tier polling needs to poll more ESs more frequently, thereby approaching the bandwidth overhead incurred by the dedicated minislot method. Two-tier polling is more efficient when most customers on the node frequently change their activity state, in other words, have bursty traffic to send. It also works well when most customers have long periodicity traffic (period > 3 MAP duration) to send. On the other hand, dedicated minislot is more efficient when most customers are active, transmitting some amount of periodic traffic with longer periodicity.

One of the major goals of the next generation design is to reduce the complexity of the MAC layer implementation. DOCSIS systems have many scheduling options with each option tailored to work optimally in very specific scenarios. This proposed approach utilizes one scheduling algorithm in conjunction with QoS and initial bandwidth requesting mechanisms that works well under most scenarios. This is intended to reduce implementation complexity. Based on the four methods presented here, 2-tier polling is the recommended scheme, due to its bandwidth efficiency, as well as the ability to adapt to the network traffic conditions in real time.

4.7 DOWNSTREAM TRANSMISSIONS

In order to accommodate IPv6 traffic, as well as IP video, the next generation provides mechanisms for transmitting both unicast and multicast downstream traffic. Similarly to DOCSIS 3.0, downstream multicast packet forwarding at the ES may be achieved by filtering and forwarding packets based on labels (termed DSID in DOCSIS). The following are 3 main functions that need to be performed by either the AU or the ES:

1. After the AU receives multicast join requests, it assigns a label (unique or non-unique) to the multicast stream, and communicates the label and the associated group forwarding attributes to the ES;

2. When AU forwards a multicast stream on the downstream, it tags the packets with appropriate labels determined earlier for that stream;

3. When the ES receives multicast packets on the DS, it performs filtering and forwarding utilizing the label and the associated group forwarding attributes.

In remainder of this section, several alternatives for labeling multicast streams are presented.

4.7.1 SINGLE LABEL PER ES

This mechanism requires a single, unique label to be assigned to an ES, regardless of the number of multicast flows it receives. Under this approach, when multiple IP video streams are being consumed by the CPEs behind the same ES, data from all video streams, as well as postregistration well-known IPv6 multicast traffic such as neighbor discovery and router solicitation, will be tagged with the same single, unique label assigned to the ES during registration.

The group forwarding attribute is a 4-tuple of (label, S, G, D), where S and G are source and group in Source Specific Multicast, and D is the destination such as CPE MAC address or the set of ES interfaces.

For all downstream multicast packets, the ES performs filtering based on the label to determine whether to forward or drop the packet. Since the ES is required to learn the CPE MAC address and the corresponding interface after registration, the ES then performs filtering based on S and G and forwards the packet to the appropriate interface or MAC address.

4.7.2 SINGLE LABEL PER MULTICAST FLOW(S)

This approach requires each multicast stream that is forwarded on the HFC network to be tagged with a unique label. For example, in case of IP video, each channel that is replicated can be tagged with a distinct label. IPv6 neighbor solicitation and router advertisement traffic is tagged with a separate label.

The AU is responsible to communicate the group forwarding attribute (label, D) to the ES after receiving a multicast join message. Upon receiving the multicast packets, the ES performs filtering based on the label to determine whether to forward or drop the packet. Unlike the previous scheme, the ES does not need to filter based on the source and group fields in the packet, instead forwards the packet directly to the appropriate interface or MAC address.

Under this approach, the AU can optionally tag multiple multicast flows with a common label. This may be useful in service scenarios such as bundled IPTV offerings. The IP video streams are only forwarded on the HFC network when a subscriber requests the channel. This approach may significantly reduce the number of downstream queues and labels that must be supported by the AU, since the number of IPTV bundles may be significantly fewer than the number of channels.

4.7.3 NO LABEL

This approach does not utilize the concept of labeling any downstream traffic. The AU communicates the group forwarding attributes (S, G, D) to the ES. Upon receiving downstream packets, the ES performs filtering on the destination MAC address to either forward or drop the packets. The ES also needs to perform filtering on S and G to forward packets to appropriate interface or MAC address.

Instead of requiring the AU to communicate the group forwarding attributes to the ES, the ES can alternatively perform IGMP snooping on upstream packets to obtain multicast group information. However, snooping must be performed on every upstream packet, which results in large processing overhead.

4.7.4 UNICAST DOWNSTREAM TRAFFIC

If a labeling approach is used for multicast traffic, it may be advantageous to use labeling for unicast traffic as well.

Table 8 compares the list of traffic labeling approaches that have been discussed in this section

	Pros	Cons
Single label per ES	Minimizes number of labels required (need only 1 byte to identify a label)	Does not take advantage of bandwidth savings offered by multicast when multiple ES requested same multicast stream (same stream is transmitted more than once, because each ES has its own label).
Single label per flow(s)	Takes advantage of bandwidth savings offered by multicast.	For the option where single label is applied to multiple streams, CPE will receive more streams than necessary.
No label	Small bandwidth savings.	More elaborate filtering process compared to label-based forwarding. IGMP snooping option can significantly increases upstream process time.

Table 8 – Analysis of DS Labeling Methods

4.8 QUALITY OF SERVICE

4.8.1 QOS ON THE UPSTREAM

In this section, several alternative ways of providing quality of service guarantees are presented.

4.8.1.1 End Station Enforced Prioritized Transmission

The existing DOCSIS network provides QoS on a per service flow basis by classifying traffic to a large extent by type into separate service flows, and guaranteeing QoS by carving out bandwidth for certain traffic types such as VoIP. This is due to the limitation that the capacity of the access network is considerably lower than the aggregate demand from customers. However, such limitation is no longer the case for next generation DOCSIS network, due to a fundamental increase in upstream capacity, as well as a dramatic decrease in number of customers per fiber node. As a result, differentiating traffic into service flows on the DOCSIS access network is no longer necessary. Instead, all traffic is sent in a best effort fashion in this proposed approach, and without distinguishing traffic types into service flows.

Unlike the environment assumed for the original DOCSIS system, the access network in the next generation environment is no longer the bottleneck. This enables alternative options to provide suitable levels of QoS. In this new environment, traffic prioritization mechanisms can be pushed to the ESs. One such mechanism is to use the IP packet tagging. At the ESs and in the home network, applications already have the ability to tag IP packets with Differentiated Service Code Points in order to ensure QoS. The ES may use DSCP to prioritize packets. Alternatively, a set of classifiers similar to DOCSIS can be used by the ES to classify upstream traffic into prioritized queues. The use of a best effort scheduling algorithm in the AU and the prioritization performed by the ES will provide enough resources for applications and services in most scenarios, without the need to introduce a flow based QoS strategy.

4.8.1.2 Centralized, Prioritized Network Access

Instead of assigning each ES the sole responsibility of prioritizing its own upstream transmission, an alternative approach of providing QoS is to provide the AU the ability to prioritize traffic on the access network, across the subscriber population on the same fiber node.

Similar to EPON, each ES implements multiple output queues on the upstream with priority levels. Ingress traffic from the CPE is classified into each queue utilizing classifiers such as IP address, port numbers, or DSCP. Each ES reports its queue capability during registration and is provided with reporting opportunities for each of its queue. The frequency of polling each queue can be optimized by each vendor. Based on the queue backlogs and their priorities, the AU can utilize a variety of vendor specific scheduling algorithm to allocate grants to each ES.

In the interest of reducing scheduling complexity, the maximum number of queues that each ES is allowed to report should be specified. The structure of the initial request frame proposed in this report allows for up to 16 queue reporting from a single ES. While this limit could serve as maximum each ES is allowed to report, an alternative is to place a recommendation on the minimum number of queues that an ES should support. Alternatively, a maximum number of queues shared by all ES on the fiber node can be specified. As an example, some EPON system allows the OLT to process up to 2048 queues for 256 ONUs.

Even though multiple queue reporting is enabled, each ES aggregates traffic from all egress queues together to form a single flow. During registration, the AU needs to assign a single Station Identifier. Even though the ES does not need to identify its upstream transmission with the assigned Station ID, because of the absence of contention in the next gen system, the grant MAP that is broadcasted needs to include the Station ID, as well as the queue ID so that each ES will be able to identify its unique transmission opportunity.

Compared to the end station enforced QoS, the centralized, priority-based media access protocol provides additional flexibilities for the operators in the context of specific service scenarios. For example, multiple operators are interested in deploying operator-managed Wi-Fi gateway that provides wireless high speed data service to their residential subscribers as well as roaming subscribers of partner networks. When both subscriber and roaming traffic are present, the priority-based QoS mechanism can provide AU a way to prioritize subscriber traffic against roaming traffic during the event when upstream is operating at capacity.

4.8.1.3 Service Based QoS

A similar, but not completely equivalent method to the prioritized network access QoS is service based approach. Instead of aggregating all upstream traffic (with different priorities) from an ES to form a single service flow, the ES can be allocated with multiple service flows, each with its own priority. This is equivalent to the existing DOCSIS approach, where service flows are usually service based.

Despite the similarities between the prioritized and service based approaches, an advantage of service based method is in the case of WiFi roaming service scenario. If the residential subscriber is usage capped and rate limited, the roamers' traffic should not be accounted towards the residential subscriber's quota. In this case, the roamers' traffic can be allocated a separate service flow with its own rate limiting parameters and usage accounting.

Under this approach, rate limiting parameters may be defined for each service flow, or for an aggregate of multiple service flows. The aggregate rate limiting may be useful when a subscriber is assigned multiple service flows. For example, a customer can be assigned a service flow for the peer-to-peer application, and a separate service flow for best effort VoIP. However, the implementation complexity increases as the rate limit violation needs to be policed, compared to the priority based QoS that allows only a single service flow be assigned to each ES. In DOCSIS, the CMTS polices the rate limiting by either ignoring the requests during the grant allocation process, or issuing the grants as usual and dropping the upstream packets later when violation is discovered. The next generation AU MAC can utilize the DOCSIS approach, in conjunction with the priority parameter assigned to each service flow.

Similar to the priority based QoS mechanism described in previous section, either the number of service flows that can be allocated to each ES, or the total number of service flows per fiber node should be capped by an upper limit like the one implemented in some EPON systems (eg, maximum of flows per MAC Domain, ie 2048) in order to control scheduling complexity.

4.8.1.4 Multi-Service, Prioritized Network Access

Another alternative to provide upstream QoS is to combine the service based and the priority based approaches proposed in the previous sections. Each ES can be assigned multiple service flows, and each service flow can enable multiple queue reporting.

This method may be useful when the total number of service flows that can be allocated per ES is limited, and consequently multiple streams of traffic with different QoS requirements (high bandwidth vs low latency, for example) are transmitted on the same service flow. Using this approach, each stream is assigned to separate priority queues that are part of the allocated service flow. By allowing the ES to enable multiple queue reporting, the AU is provided with the information required to satisfy the QoS requirement of each stream.

4.8.1.5 Observations

Table 9 summarizes the pros and cons of each QoS mechanism proposed in this section.

	Description	Pros	Cons
End station enforced prioritized QoS	Single SF per ES. Multiple queues per SF. QoS parameters defined per SF.	No centralized grant assignment based on QoS is required. This translates to simpler AU US scheduler requirements. Covers situations when US has abundant capacity, a likely scenario	US capacity may still be limited in specific deployment scenarios such as in transition phase. Does not allow traffic prioritization across the subscribers.
Centralized, prioritized network access	Single SF per ES. Multiple queues per SF. Per queue reporting enabled. QoS parameters defined per SF.	Prioritization is useful in scenarios when capacity is still limited. Queue backlog insight enables flexible configuration measures. For deployments such as the WiFi roaming gateway, if US capacity is limited, this ability may be essential. By assigning multiple prioritized queues for roaming traffic per ES, this method can provide another layer of prioritization for WiFi roaming. Some roaming traffic such as voice may be further prioritized against other roaming traffic. US grant allocation mechanism may be leveraged from all EPON vendors as multiple queue reporting is required by the EPON standard.	More complex to implement than ES enforced scheme, as the AU needs to apply prioritization mechanisms. Higher complexity can result from large # of queues reported by the ES. Cannot enable separate usage accounting, as well as separate rate limiting when including WiFi roaming. If residential subscriber is rate limited and usage capped, roaming users' traffic should not count towards residential sub's usage and rate limits. Lower priority queues can potentially starve. If AU implements mechanism to avoid starvation, this approach becomes similar to service based approach.

	Description	Pros	Cons
Service based	Multiple SFs per ES. Single queue per SF. QoS parameters defined per SF.	Provides true traffic separation between multiple customers utilizing the same ES. This might be of special interest in WiFi roaming. Can provide traffic prioritization in services such as WiFi roaming by using a priority parameter as part of the QoS config set for the SF. Enables separate usage accounting, as well as separate rate limiting on a per flow basis. US grant allocation mechanism may be leveraged from certain EPON vendors who already provide multiple LLID scheduling.	More complex to implement than ES enforced scheme, as the AU needs to apply prioritization mechanisms and police QoS parameters. Higher complexity can result from large # of service flows. Rate limit can be defined for an aggregate of SFs, although additional complexity will occur due to more complex policing mechanism.
Multi-service, prioritized network access	Multiple SFs per ES. Multiple queues per SF. QoS parameters defined per SF.	Useful when a large number of customers are served by an ES and their service is independently managed and billed and their individual traffic is prioritized. US grant allocation mechanism may be leveraged from certain EPON vendors who already provide multiple LLID and multiple queues per LLID scheduling.	Most granular QoS scheme among all that are proposed. More complex to implement than all other schemes, as the AU needs to maintain a set of queues for the received US traffic. The # of queues must be supported at the AU and/or the ES level should be limited to reduce complexity. The flexibility provided is easily accommodated using service based method. Unclear where the additional flexibility can be utilized.

Table 9 – Analysis of Upstream QoS Mechanisms

4.8.2 QoS ON THE DOWNSTREAM

Similar to the upstream, the existing DOCSIS network provides downstream QoS on a per service flow basis by classifying traffic to a large extent by type into separate service flows, and guaranteeing QoS by carving out bandwidth for certain traffic types such as VoIP. The significant increase in downstream capacity in the next generation systems to the order of multiples of gigabit per second, coupled with a dramatic decrease in number of customers per fiber node, may render differentiating traffic into separate service flows no longer necessary. Under this assumption, the ES is allocated one service flow, where all traffic destined for it, is multiplexed on to the same service flow. QoS parameters for each service flow are still defined in order to provide differentiated Service Level Agreement for each ES. The AU is responsible to implement and enforce QoS parameters.

A more flexible approach is to enable multiple downstream service flows per ES. With unicast video streams becoming more of the norm, multiple concurrent unicast HD video sessions may occasionally stretch the capacity of the network. Additionally, business services require stringent bandwidth guarantees. Allowing traffic to be classified into separate service flows with parameterized QoS is advantageous under these scenarios.

In order to keep complexity at a minimum, a reasonably small number of service flows that the AU is required to support should be recommended.

4.9 EPON MAC AND MIGRATION PATH TO FIBER TO THE HOME

The proposed MAC layer architecture bears resemblance to the EPON standard. The absence of channel bonding at the MAC layer and its associated functionalities signifies a key departure of the next generation design from the DOCSIS 3.0 systems. In addition, similar to the EPON MAC standard, the next generation system advocates the use of just one scheduling framework in conjunction with QoS and initial network access mechanisms that work well under a multitude of traffic conditions. The polling-based initial bandwidth request method that is proposed and recommended in this report is similar to the GATE and REPORT processes in EPON. Multiple queue reporting which is part of two of the proposed QoS mechanisms is also supported by EPON.

The similarities between the proposed next generation MAC design and EPON MAC standards potentially enables operators to take advantage of EPON's economies of scale.

Using an EPON friendly MAC layer as part of the next generation system as well as having an environment that shares many characteristics as EPON enables several migration paths to an all fiber network.

5 CONCLUSION

This report proposes a potential MAC layer architecture for data services over the CATV network. The architecture proposed here places heavy emphasis on the decoupling between MAC and PHY layers, which is a major design requirement, as well as modularity, which makes it scalable depending on the amount of spectrum available. By removing the backward compatible requirement, the next generation data system represents a clean slate regardless of previous assumptions, technologies and tools, and thus freeing itself from the burden of complex design required to support legacy systems. However, the new architecture can coexist with legacy systems to ensure smooth transitions.

Channel bonding and unicast ranging are absent from the proposed next generation architecture. This has resulted in dramatically simplified responsibilities performed by the MAC in addition to bearing closer resemblance to the EPON standard.

In addition, the environment which the next generation architecture is expected to operate in is also similar to the EPON systems. Through node splits, the size of a node in the HFC network has been steadily decreasing. Node sizes of 256 ESs are reasonable to assume in a 5+ year time frame. The fundamental increase in network capacity and reduction in number of users sharing the same bandwidth resources result in lower complexity alternatives than DOCSIS Quality of Service models. Given the dynamics of the network capacity and size, the detailed classification of packets into service flows designed to carry specific types of traffic may no longer be necessary. Instead of relying on the AU to prioritize upstream traffic, each ES can be delegated to prioritize its own upstream traffic. Consequently, it is possible for the next generation system to use just one QoS framework that work well under a multitude of traffic conditions. On the other hand, having the capability to manage traffic with a high degree of granularity may enable new ways of providing services. WiFi roaming may be one of the scenarios that could take advantage of such capabilities. Similarly on the downstream, increase in network capacity and reduction in number of users lead to simplified downstream parameter-based QoS, where a single service flow per ES is defined. However, a more flexible approach that enables multiple service flows on the downstream is also considered, and may be very relevant under some scenarios.

Multiple initial network access methods have also been proposed. A polling-based method has been recommended in this report, similar to the one used in EPON. To facilitate the bandwidth request and grant process, the minislot structure proposed here provides a bandwidth-efficient way for the ESs to transmit their existing queue length using request messages. Using MAP messages, the minislot structure also enables the aggregation unit to grant the ESs minislots for unicast transmissions.

With the features and characteristics presented in this report, the MAC layer design for the next generation systems will be able to achieve low implementation complexity and high modularity while maintaining high performance.

6 REFERENCES

[SHAW] Terry Shaw, Bandwidth Utilization Update, first quarter 2009. https://www.cablelabs.com/doczone/strategic_assessment/currentprojects/bandwidth_management/project_information_reports/2009/

[PHY-3.0] DOCSIS 3.0, Physical Layer Specification, CM-SP-PHYv3.0-I08-090121, January 21, 2009, Cable Television Laboratories, Inc.

[AMP-PHY] Alberto Campos, Jennifer Fang, PHY Layer Design for Efficient Multi-Gigabit Transport over CATV Networks, CableLabs internal report, 2H,2009.