Unleashing the Speed with DOCSIS 2.0

Supporting the Burgeoning Commercial Market

John J. Downey
Broadband Network Engineer
Cisco Systems
jdowney@cisco.com

Originally, most Internet content was downstream (DS) related and required an asymmetrical design of DS speed versus upstream (US) speed. Most traffic was web-based downloads and e-mails, and very bursty in nature. Streaming content was added that caused average usage to increase. Peer-to-peer (P2P) traffic created more US and DS requirements that made traffic patterns appear somewhat symmetrical. This, along with voice services, creates a higher load and a lot of symmetrical traffic. Now we have services like Bit-Torrent where movies may be downloaded and uploaded or companies like Movielink providing downloadable movies for Ipods or any other type of multimedia device. Downstream throughput is once again becoming the "bottleneck". If cable systems want to eventually provide much more content over the DS DOCSIS channel, they either need to decrease the number of users sharing that channel with physical node splits or add more channels into each node. This is at the expense of frequency bandwidth and more CMTS ports.

The issue with this becomes, what to do with all the USs that are associated with those DSs. A MAC domain is a DS carrier and its associated US carriers. A 1x4 MAC domain may not be needed as US ports are unused or extremely under utilized. Also, modems that move to another DS, must move to a new US associated with that DS. This is where DOCSIS 3.0 comes to the rescue by breaking the 1xN MAC domain constraint and allowing MAC domains of MxN.

This paper explores the evolution of DOCSIS cable modem speeds and outlines some options available to provide greater speeds to combat competitive threats and to service the commercial market. One system in particular (Cablevision of Long Island) will be evaluated to explain how they approached this issue.

Competitive Outlook and Objectives

Cable systems need to provide faster network service to commercial and residential customers to compete against new fiber-to-the-home (FTTH) offerings from Verizon of 15 Mbps DS by 2 Mbps US. This service is being labeled as FiOS, which stands for fiber optic systems.

Consumers are concerned about speed, cost, ease of use, and customer service, but not technology. Whoever can provide the faster service for a cheaper cost, will win over the customer. Running fiber to a customer's house can be very expensive. Equipment is expensive and fiber placement is cumbersome in regards to "right-of-ways" and trenching. Using the existing hybrid/fiber coax (HFC) cable network is much more economical than placing a node in each house.

This material was originally presented at SCTE Cable-Tec Expo 2006. Reprinted with permission.

Cable systems need to provide separate tiers of service with DOCSIS 1.1 while using the existing HFC network. Some examples of tiers of service include:

1. Dial-up replacement 128 kbps x 128 kbps
2. Low speed 1 Mbps x 256 kbps
3. Medium speed 3 Mbps x 384 kbps
4. High speed 5-7 Mbps x 512-768 kbps

Upgrading to DOCSIS 2.0 increases US throughput to satisfy these new services, but 10% of your users could be "abusing" the service and causing 80% of the load. It may be prudent to utilize equipment to shape the traffic and/or redirect it to keep things manageable.

These tiers of service are just one variable when the age-old question is asked, how many subs can I put on the US or DS. The answer is still the same, it depends. Some systems may get away with 1000 devices on an US port with only 50% utilization during peak periods, while another system is relegated to 75 subs per US and their utilization is "pegged" at 80% for most of the day. Another system could be constrained to 100 subs per US because the MAC domain of 1x4 would put the DS at 400 subs and that was their "bottleneck".

Current and New Speed Offerings

With a typical rate of 3 Mbps DS per user, only 12 subs would be able to share a 256-QAM signal simultaneously. Some systems may allow an oversubscription of 100:1 on the DS leading to 1200 subs per DS. Oversubscription is very dependent on demographics; college town or suburb?

Increasing a customer's "cap" to 5 Mbps and using the same model would produce only 800 subs, but changing the "cap" may not have a linear affect as predicted or assumed. On one hand, average usage may be less than extrapolated, but on the other hand, customers that use a lot of P2P services may look more appealing to others outside the network. For example, offering a tier of service with 15 Mbps per DS would only allow 2*100 = 200 subs per DS and hard to justify a business case. Allowing much more oversubscription may be fine, but needs to be observed over time.

Customers that actually pay for a higher tier of service could feel compelled to get their moneys worth and use much more than previous. Also, as customers become more computer savvy or other applications become prolific, usage could increase exponentially. This equates to an oversubscription calculation that must be re-evaluated and probably decreased. Now the oversubscription of 100:1 may need to be 75:1.

Some cable systems have different offerings in different parts of North America. Cablevision has 2 tiers with one providing all basic subs with 10 Mbps DS by 1 Mbps US and Tier 2 providing high-end customers with 15 Mbps DS and 2 Mbps US. Cablevision also has a 64-QAM US trial going on. Advanced Time Division Multiple Access (ATDMA) as provided by DOCSIS 2.0 allows higher modulation schemes and channel widths and has been deployed in Japan and Korea for over a year, but just being deployed in the US now.

Comcast, in some parts of the US, now offer a 9 Mbps DS by 1 Mbps US. In other parts of the US they offer a standard 6 Mbps DS by 768K US. Cox have many systems now getting 8 Mbps DS by

1 Mbps US and in parts of the US are now rolling out 6.4 MHz US channel width with 64-QAM on the US. They have this in production and plan to start rolling this out more aggressively as they learn to use this more widely. Some current offerings include; 5x512, 9x1, 10x1or2, 12x2, 15x2, 5x5, and 10x10. Different markets have different requirements. Other customers are doing something like a 5 Mbps DS and 512K-1Mbps US and there are some European customers doing 8 Mbps DS and 1 Mbps US.

In all the cases above, customers are working toward 600-1000 subscribers on a single DS. Cablevision would like to get to 500 subscribers per DS and TWC in Houston is looking at 75 subs per US port. In all cases, customers are running 16-QAM on all USs and 256-QAM on the DS.

Option 1 - Same CMTS with Frequency Separation

Cable companies have some choices that consist of:

- Do nothing and watch the competition erode your subscriber base.
- Segment the fiber nodes, which could cost ~ $10,000 per node.
- Utilize FTTC (fiber to the curb), FTTH(P)(fiber to the home or premise), FTTWAP (fiber to the wireless access point). This is an expensive proposition.
- “Bonding” of US and/or DS DOCSIS channels. This is the option being promoted.

Frequency division multiplexing (FDM) allows multiple carriers to be on the same plant. Using 3.2 MHz channel width at 16-QAM on both USs can give ~9 Mbps per US frequency of usable rate, but this depends on the frame size. Using 256-QAM on both DSs can provide ~36 Mbps per DS frequency of usable rate, but also depends on the frame size.

Mapping two DS frequencies and two US ports on a linecard into the same node allows frequency A to serve residential subs and frequency B to serve new commercial subs. This is displayed in figure 1 below. The use of client-class processing "steers" residential subs to A and new commercial subs to B. This entails setting the DS frequency and/or US channel ID in the CM’s configuration file. The use of US and/or DS load balance allows the subs to share the combined USs and/or DSs without requiring modem configuration file changes.

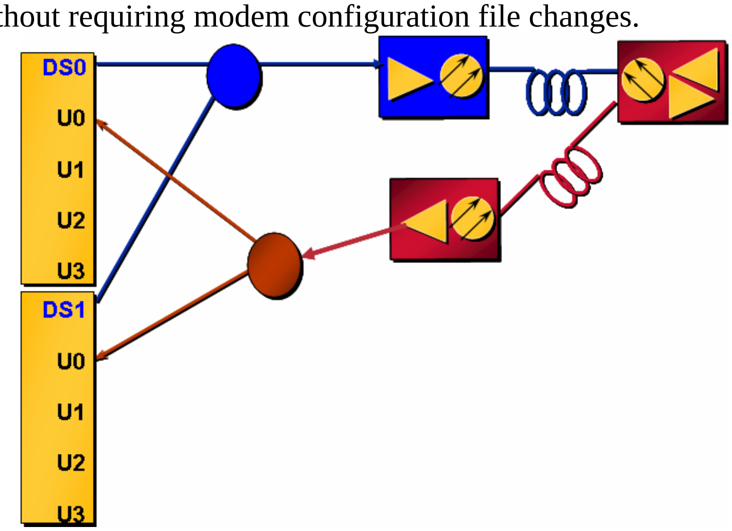

Figure 1 Separating CMs via FDM and Provisioning

In this example, DS0 = 453 MHz and DS1 = 459 MHz with both at 256-QAM. Upstream 0 for DS0 = 24 MHz and U0 for DS1 = 30 MHz with both at 3.2 MHz.

Two DSs at 256-QAM give ~ 72 Mbps of usable rate and two USs at 16-QAM /3.2 MHz give ~ 18 Mbps of total usable rate. This option provides a simple "get up and run" approach with the only real modifications of combining to map multiple DSs and USs to nodes. More involved modifications require the provisioning system to steer specific modems to the proper US and/or DS channel or load balancing can be used for a simpler approach.

The use of US and DS load balancing could be considered a “poor man’s” redundancy scheme. If one DS upconverter dies, then all the modems would register on the other DS upconverter. This assumes the modem doesn’t use provisioning to redirect it to a specific US channel ID or a specific DS frequency.

Option 1 - Cons

Some cons to this option include:

- Requires combining work such as cabling and/or splitters. Most systems already have rack-mount modular splitters with empty ports.
- Requires DS and US spectrum availability.
- Moving to a new DS requires a new associated US.
- A total node outage could make CMs register on an incorrect DS and affect each other in terms of how fast they come online.
- CMs may need to be client-class processed with info in their DOCSIS config files.

Some hurdles to overcome and questions to answer are:

- Does the system broadcast multiple DSs because of optic amplifiers? Broadcasting the CMTS DS signals makes it difficult to achieve 1:1 DS-to-node narrowcast combining later down the road.
- Do the nodes have one DS receiver & two US transmitters? Two DS carriers could be sent to the node where one US laser feeds an US of one MAC domain and the 2nd laser feeds an US from the other MAC domain. Using this architecture allows US frequency re-use, but it requires open upstream ports and combiner work. Adding insertion loss forces different levels from the modems and you must utilize two separate US frequencies.

Option 2 - Separate CMTSs

Option 2 just utilizes two separate CMTSs where one is for existing residential subs and the new CMTS would be for commercial subs.

The pros of this approach are:

- Hardware isolation.
- More processing power.
- Future expansion and “poor-man’s” high availability.
- Separates “high speed” customers for Network Operation Center (NOC) “clarity”.
- Some systems may use this today for open access.

The cons for this architecture include:

- Must purchase 2nd CMTS.
- Will the CPU max out if one chassis dies and the other takes over the full load?
- More power draw and rack space in HE/hub.
- Must integrate the 2nd chassis to the network and manage this device. This requires address space for the additional CMTS. Also, CMs could lock on the wrong CMTS causing IP address depletion.
- Modems could “bounce” between both CMTSs unless tight provisioning is configured.

Option 3 - Same CMTS Utilizing ATDMA

Option three is the same as option one, but the second US frequency in the node is ATDMA-only. Refer to figure 1 for a graphical display. Utilizing ATDMA-only on US ports for commercial CMs allows only 2.0 CMs to "see" that US port and register on it. Existing residential subs are blind to an ATDMA port because they don't understand MAC message 29, which is included in the upstream channel descriptor (UCD). UCDs are sent for each US that belongs to the specific DS every 2 seconds in a DS map.

One issue to address when using ATDMA is the need to configure provisioning to block commercial CMs from registering on the 1.x (residential) US port. Provisioning is used to force a specific DS frequency or US channel ID for a specific modem and is inserted in the modem’s config file.

In this scenario DS0/U0 is configured for TDMA (1.x mode) and DS1/U0 = ATDMA (2.0-only). This only allow 2.0 CMs to "see" DS1/U0 and register on it while existing residential customers are blind to DS1/U0. This allows two stacked US channels, but only appears as one for residential customers. If the US spectrum is available and “clean”, then 2.0 CMs could use a 6.4 MHz channel width at 64-QAM to give ~ 27 Mbps of usable speed. This is an easy way to add bandwidth to a service area.

An added feature with DOCSIS 2.0 is advanced phy features, which include ingress cancellation, 24-tap equalizer (EQ), analog-to-digital (A/D) conversion, and modulation profile advances. Although not part of the DOCSIS 2.0 spec, ingress cancellation allows higher orders of modulation. Ingress cancellation makes the US port robust against certain plant impairments, opens unused portions of spectrum, and is insurance for life-line services.

Another tool in the battle against linear distortions is adaptive equalization. DOCSIS 1.1 supports 8-tap adaptive equalization, and DOCSIS 2.0 supports 24-tap adaptive equalization. Unfortunately, the large installed base of DOCSIS 1.0 modems won't benefit from either, because DOCSIS 1.1 and 2.0-specified adaptive equalization is done using pre-equalization in the modem. DOCSIS 1.0 modems generally don't support adaptive equalization.

US equalization is supported for 1.0 and 1.1 with an 8-tap blind equalizer. Some people call this “post” equalization. DOCSIS 1.1 allows 'pre-equalization' where EQ coefficients are sent to the CM allowing it to pre-distort its signal. This is a configurable option using “cab up x equalization-coefficient”. The modem itself does not require a 1.1 config file, but it does require

1.1/2.0 capable CMs. DOCSIS 2.0 increases the equalizer tap length from 8 to 24 and supported in ATDMA mode, but may be off by default.

DOCSIS 2.0 Benefits

DOCSIS 2.0 increases US capacity to 30.72 Mbps. The modulation profiles add interval usage codes (IUCs) for atdma mode. This allows 2.0 modems to burst with different modulation schemes vs 1.x CMs in a mixed environment. DOCSIS 1.x modems use IUCs 5 & 6 for short and long grants while 2.0 CMs use IUC 9 for a-short, 10 for a-long, and IUC 11 for a-ugs.

DOCSIS 2.0 introduces a “docsis-mode” concept, “cable upstream x docsis-mode {}” to configure the US channel to a desired mode. These modes are; TDMA (traditional) mode, ATDMA-TDMA (mixed 1.x and 2.0) mode, and ATDMA (2.0-only) mode.

The advantage of DOCSIS 2.0 and ATDMA usage is wider channels, which results in better statistical multiplexing (a 6.4 MHz channel is better than 2, 3.2 channels), greater spectral efficiency, better use of existing channels and basically more capacity. This provides higher throughput in the US direction and greater per-modem speed with better packets-per-second (PPS) rates. Refer to the references section for more details about DOCSIS Throughput considerations.

Since DOCSIS 2.0 increase the US rate, a 1x1 MAC domain may make more sense. One feature that may be available is configurable MAC-domains also called virtual interfaces which enhances flexibility when used in combination with ATDMA. The references section has more information.

Option 3 - Cons

The cons to using ATDMA include:

- Requires new “high speed” users to have 2.0 CMs.
- Requires provisioning work to "block" 2.0 CMs from registering on the residential frequency.
- If residential subs buy their own 2.0 CMs, they could lock to the commercial US without provisioning interdiction. One way to avoid this is the use of TLV 39 = 0 for residential CMs, which forces 1.x mode even if they are 2.0 capable.
- Can’t utilize load balancing unless DOCSIS mixed-mode is configure with utilization-based load balancing. If using mixed-mode, then 2.0 CMs could burst at 64-QAM for ~ 13 Mbps while 1.x modems could burst with 16-QAM for ~ 9 Mbps. Another con to this is the channel width is limited to 3.2 MHz at the most.
- May require dynamic frequency hopping or modulation changes to provide insurance against plant impairments that cause high packet loss for the higher channel width and/or modulation schemes. For more information about advanced spectrum management and load balancing, visit Cisco.com, call TAC, or drop an e-mail to the author of this paper.

New Architecture Idea Using Option 3

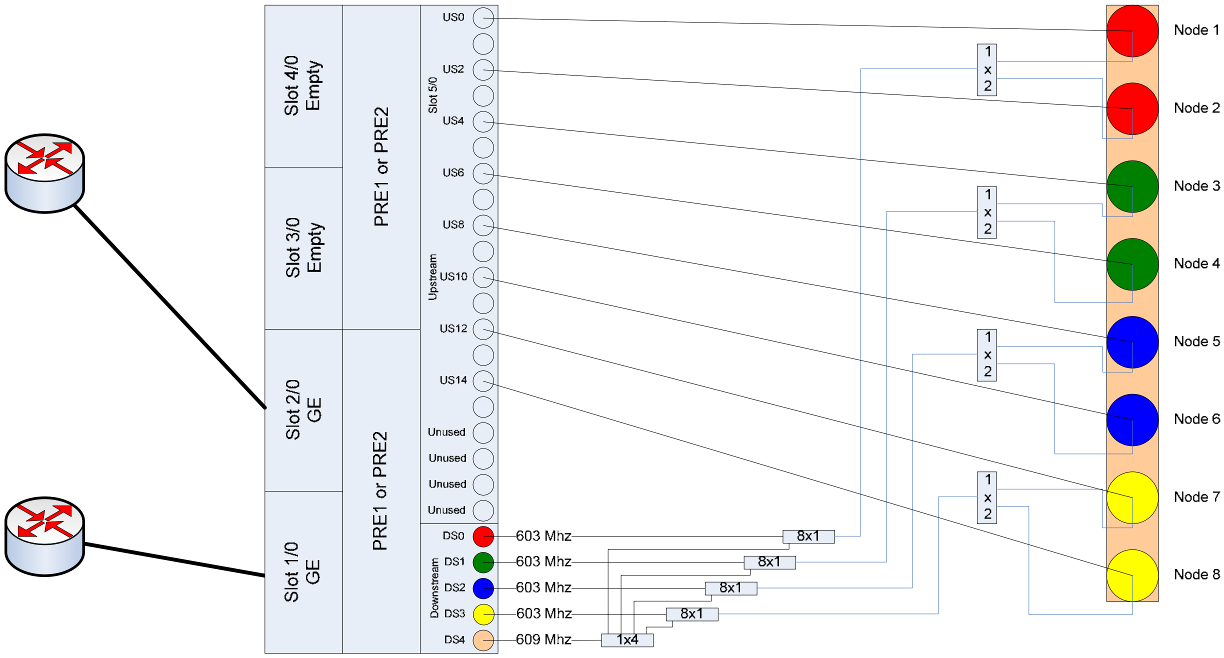

Figure 2 Current Architecture used by Cablevision

This architecture provides simple cabling and combining by configuring four, 1x2 MAC domains and one, 1x8 MAC domain for the commercial subs. This allows the commercial DS to be split four ways to be combined into the four residential DSs. Each of the residential DSs are split to feed two nodes. This “Boost” configuration allows an US rate up to 27 Mbps per node and a DS rate of ~35 Mbps per 8 nodes at 256-QAM. This could be further segmented down to four nodes if commercial take rates exceed expectations.

The residential tiers use DS0-DS3 and are 1x2 MAC domains. The power tier for commercial subs use DS4 with a 1x8 MAC domain. Upstreams are frequency stacked to allow one cable from the fiber optic receiver to feed one US port on the CMTS with two frequencies. One US frequency is for DS 0 and the other US frequency is for DS4. The only negative aspect to this design is four US ports are not used, but they could be combined in with the residential to make four, 1x3 MAC domains if needed. In this example, the residential USs are all configured for 24 MHz, while all the commercial US ports are configured for 30 MHz.

Questions to Answer

If the residential CM locks to the commercial DS4, it will receive all UCDs unless all UCDs contain MAC message 29 for atdma-only. Once the residential CM locks to DS4, it tries to transmit on all UCDs at 30 MHz. This doesn’t affect existing residential data at 24 MHz because it’s at a different frequency. It can affect commercial registration times slightly since the residential modem is using contention time and attempting to range on the commercial US frequency. If DS4 is from the same linecard and uses tdma-atdma mode, then DS and US load balancing is possible.

If the commercial CM locks onto the residential DS0 & transmits at 24 MHz, this doesn't affect existing commercial data at 30 MHz because it's at a different frequency.

If a power outage happens on an entire node or a fiber cut, then modems will start scanning DS frequencies and could lock to the wrong one. Modems cache the DS frequency, US frequency, and US power level. Some modems will "look back" for the cached DS frequency every so often, which will help in these situations.

If a residential sub buys their own 2.0 CM, they could lock to the commercial US without provisioning interdiction. One way to do avoid this is the use of TLV 39 = 0 for residential CMs, which forces 1.x mode even if they are 2.0 capable.

Some physical implementation issues concern DS spectrum availability for 256-QAM and can the US laser handle multiple carriers and higher modulation schemes. There could be laser clipping from adding additional frequencies and higher modulation schemes. Distributive Feedback (DFB) lasers are a must and/or possibly the usage of baseband digital reverse (BDR) technology.

Concerns

Some concerns include: How to scale 1x8 and 1x2 MAC domains as utilization increases and take into account RF connections, provisioning, etc.

- Are there bandpass filters in plant that would negate an US frequency needed and are US diplex filters in the range of 5-30, 5-40/42, 5-55, 5-65, … If I decide to use amplifiers with 5-65 MHz diplex filters, can I use the entire bandpass? As long as the DS channels, 2 - 4, are moved and CMs are used that can be programmed for the frequency range, then it is possible. Vendors make DOCSIS CMs that can do 5-55 MHz, but not 5-65. Only Euro-DOCSIS CMs use 5-65 MHz and they won't work in North America. The issue at hand is existing customer premise equipment in the field that could overload from this new high power signal above 50 MHz.
- How to prevent IP address exhaustion during ranging onto an incorrect DS frequency.
- Are there any known CM vendor incompatibilities or issues when using multiple DS & US frequencies? One issue with this is that it could take a long time to register since CMs cache the DS frequency and if the frequency is new, where does the modem start to scan?
- What happens from a power level perspective on an US when the second US frequency is added? The CMTS performs power on per-US frequency & channel width, but an analog front-end of a CMTS could overload and cause harmonics internally.

Another thing to keep in mind is per-modem speeds. DS speed is affected by frame size and the transport layer 4, which is TCP or UDP. TCP is affected by US speeds because TCP requires US acknowledgements and this affects TCP windowing. Upstream speed is affected by the DOCSIS protocol, map advance, DS interleaving, concatenation, max concatenation and traffic burst, CM and CMTS fragmentation, and modulation profiles. It may be prudent to upgrade the routing engine in the CMTS to meet performance and new feature requirements.

New Technology Cornerstones

New technologies are being pursued to address the DS bottleneck conundrum. DOCSIS 3.0 uses a channel bonding technique to achieve higher capacity links, enable faster high speed data (HSD) service, and provide M x N MAC domains to enable Video over IP solutions.

The modular CMTS (M-CMTS) architecture is promoted to achieve better DOCSIS economics, lower cost DS PHY, and de-couple DS and US ports. One day we may see fiber optic nodes with DOCSIS physical layer chips embedded so we can use ingress cancellation at the node, digital links from the node back to the headend without the need to amplify, and no more laser clipping. Of course, this means all traffic needs to be DOCSIS-based on the US!

Other Commercial Offerings and Summary

Other commercial offerings include a T1-type of service called Circuit Emulation over IP (CEoIP) to provide latency and jitter sensitive services to business customers. This may be offered with existing PacketCable equipment or with new PacketCable Multimedia equipment.

For small office/home office (SOHO) customers, a Wan interface card (WIC), may be the answer to provide a quality of service (QoS) required by this type of customer.

Another fruitful offering is overselling older frame-relay services with either an MPLS-VPN service or a layer 2 VPN service.

Doing the math for data throughput and implementing new architectures are only beginning steps. One must monitor actual traffic loading to determine how to manage fair use of the network, then determine when additional capacity is necessary to satisfy customer needs.

References:

Advanced PHY
http://www.cisco.com/en/US/products/hw/cable/ps2217/products_white_paper09186a008017914d.shtml

Understanding DOCSIS Throughput Issues
http://www.cisco.com/en/US/tech/tk86/tk168/technologies_tech_note09186a0080094545.shtml

DOCSIS 2.0 ATDMA Configuration on MC5x20S and MC28U Linecards
http://www.cisco.com/en/US/tech/tk86/tk168/technologies_white_paper09186a0080231fc3.shtml

Virtual Interfaces and Frequency Stacking Configuration on MC5x20S and MC28U Linecards
http://www.cisco.com/en/US/tech/tk86/tk804/technologies_white_paper09186a0080232b49.shtml