High-Speed Data Multiplexing over HFC Networks

Mark Laubach
Com21, Inc.

Abstract

The deployment of high-speed data services over all-coax and Hybrid Fiber-Coax (HFC) networks is forcing cable operators to quickly develop an understanding of high-speed data multiplexing. Traditional Internet engineering focuses on local area network design, backend network technology and topology, and connection to wide-area or regional networks. In addition, the emergence of different varieties of Cable Modem Termination Systems at the head-end has met with some interesting challenges related to incremental growth, scalability, and matching data channels to RF trunks in a scalable and cost effective manner. This paper overview engineering design issues associated with aspects of high-speed data multiplexing and the incremental re-ordering of the networks to effectively meet growth starting from sparse subscriber deployment and continuing to large take rates.

Getting Started and Growing Larger

A cable operator or Internet Service Provider (ISP) working with a cable operator has a variety of ways to get started with deploying Internet services. The basic starter kit has simply been a router with one Ethernet port and a T1 connection (1.544 Mbps) to a larger Internet Access Provider (IAP) or ISP, a cable modem head-end system, an Ethernet LAN. This back-end LAN connects the cable modem equipment with the router, and a truck full of cable modems. This starter set has elements common to all Internet access scenarios since the early 80's: a connection to the bigger world, a back-end network, and a baseband or broadband local area network connecting the users to the backend network.

What is not the same as in the 80's (even with broadband modems) is that there has been a technology revolution (several times) and an explosion in the number of users who can potentially take advantage of new access to incredible amounts of bandwidth. Since these users are located in homes, the incremental growth from starter residential access network system to largely deployed systems is both staggering and highly variable on a market or fiber node basis. Fortunately, cable modem equipment has followed an anticipated advancement path however, their flexibility with respect to data multiplexing has in some cases outstripped the CATV plant.

Mentioned briefly, but not the subject of this paper are the other new services, abilities, and challenges that a cable operator and ISP have jointly discovered:

- Management and deployment of Dynamic Host Configuration Protocol (DHCP) servers
- Management of large blocks of IP addresses and changing these over time, web proxy and caching servers and where to place them
- When and how provide news, mail, and FTP servers
- How to manage the entire data network effectively
- What type and size routers needed
- Where and when to use ATM
- What to do about simultaneous data and voice transport?
- When to where to use multicast services
- What standards do I need to be concerned with?.

Generally, cable operators and their ISPs, seem to fall into one of three business model mentalities:

- “flat rate pricing is just fine”,
- “I want multi-tier pricing”, or
- “I want multi-tier pricing and the ability to sell bandwidth pipes over the RF to other ISPs or corporate concerns”

The requirements for high-speed data multiplexing are different for each model.

FLAT RATE PRICING IS JUST FINE

This business model begins with the Internet “starter kit” model. I also call this model Home Box Internet. The topology of the backend and regional networks form a tree. A cable head-end is connected to other head-ends via a large backend networks or each head-end is connected to a regional network center, if available. There may or may not be a nationwide backbone.. Large networks such as @Home do have their own backbones.

A data multiplexing point (or concentration point) exists where the “child” connects to the “parent” in the tree; i.e. link from head-end to regional hub, link from regional to nationwide access, etc. When the link costs money to provide, then the goal is: spend as little money as possible on the cost of the link but maximize the number of users accessing the link through the child. The link is sized to meet the needs of the statistical peak load of the users and not on the sum of bandwidth available to each individual users. This is called statistical multiplexing. Practically, this means putting caching servers and other servers on the child side of the link. This reduces the overall number of individual “connections” that need to flow through the link. So long as the price of the caching servers is less than the cost of the next size link, this model holds. Sizing the link is not rocket science. But we do have rocket science modeling tools to predict the size of links. If the link is too small, users will complain. If the link is too large, then the Chief Financial Operator (CFO) will complain. The trick is sizing within the subscriber/CFO tolerance envelope.

Within a head-end, the choice of local area network equipment is going to be chiefly decided by the Internet model being deployed. If the downstream world looks like one big Ethernet, than just a single high-speed Ethernet port is needed between the Cable Modem Terminal System (CMTS) and the router. However, this model varies greatly based upon the desires and design of each cable operator and ISP. For the basic system, one Ethernet is needed and one IP subnet is needed. As the system grows larger, than the number of Ethernet ports and IP subnets will vary. This effects the type and configuration of the cable modem equipment and will require changes while the size of the subscriber base grows. For North America, plan on a CMTS with a single downstream RF channel moving about 25 Mbps full-duplex when fully loaded. If using Ethernet, than getting started with 10BaseT works, but needs to go to 100 BaseT very soon. If using ATM as the backend connection from the CMTS, then the size is either OC3 (155 Mbps) or DS3 (45 Mbps) meaning the connection doesn’t need to change over time between the CMTS and the router - there is just unused bandwidth. Since that link is free from monthly cost, this is ok. However, if the ATM link (from a router or from the CMTS) is being back hauled over a SONET network, than there many be monthly costs and the data multiplexing ability of the router/CMTS must be capable of filling the OC3 link. This means supporting multiple downstream and upstream data channels over the RF via the single link.

I WANT MULTI-TIER PRICING

In the context of this paper, multi-tier means that the cable operator has the ability to give different subscribers different allocations of bandwidth based upon how much monthly service fees they are going to pay per month. This capability provides the operator with a much better revenue generation model than single flat tier pricing, as there are always customers willing to pay more for more bandwidth or reduced delay.

From a data multiplexing point of view, the model for supporting this type of billing and tiered ability is very similar to the flat rate model. It is just that the size of the link from the child to the parent has to be increased to accommodate that bandwidth that has been sold to the multiple tiers. The assumption is that the users at the same pricing level (billing class) can be statistically multiplexed together, so it is not a sum of all bandwidth per user game. The issue at hand is that users are paying more money for better perceived performance (lower delay) and there needs to be capacity allocated over the links between the head-end and the Internet in excess above a flat price scenario. Also, the same rule applies to stay within the subscriber/CFO tolerance envelope. The added dimension is that there are now several classes of subscribers (related to service tiers) and expectations must be set per tier.

I WANT THE ABILITY TO SELL BANDWIDTH PIPES OVER THE RF

This model is a bit different than the previous two types. While the single tier or multi-tier scenarios are part of the data multiplexing puzzle, there are added requirements:

- cable modem equipment must be able to “carve out” portions of data and/or RF bandwidth that can be allocated to a specific user, group of users, or virtual LAN
- there must be added routing and data segregation in the back-end networks
- companies or other ISPs may be connecting at different places in the network hierarchy

This type of business model is very close if not the same used by Competitive Access Providers (CAP). The cable operator may actually not be in the direct business of offering ISP services for subscribers, rather they are providing access to one or more ISPs, which in turn provide Internet, access to their subscribers. There may be several ISPs in operation over the same RF channels in this model. In addition, companies may buy bulk bandwidth for their own backend network needs or for supporting a closed group of employees for telecommuting. Every one is sharing the same RF channels, the cable modem equipment is providing segregated multiplexing to keep the traffic separate from one another: in a routing sense, a security/privacy sense, and from a bandwidth management sense.

The motivation for this type of business model is driven by the recognition that cable operators can compete with local telephone companies for providing bulk bandwidth connections. It is expected that the cost of service over cable is and will remain a better value deal than a corresponding telephone service.

For this business model cable modems must be able to be placed into closed user groups. Closed user groups need to be given different amounts of bandwidth over the RF, the data traffic from closed user groups must be able to be kept separately partitioned from the traffic from other closed user groups for privacy, management, and billing/accountability reasons. Cable modem equipment that supports this type of operation can be configured to support the single tier and multi-tier types. Cable modem equipment designed to support either of the other types cannot support this CAP model.

In addition to the above, connections to ISP or companies can occur as various places with the hierarchy of the cable operator’s network. There are many varieties of scenarios and too numerous to detail however, each one impacts the amount of segregated bandwidth that needs to traverse a link.

Also, the same rule applies to stay within the subscriber/CFO tolerance envelope. The added dimension is that there are now several classes of subscribers (related to service tiers) including ISPs or companies who have purchased bulk access.

Oop's, I've Added Another Different Type of Digital Service

This small section just mentions that if the cable operator wants to install an additional type of service (e.g. a digital toll quality voice system or digital video transport) down the road, then they might want to give consideration when they are building their networks for high-speed data for Internet. For example, if the cable operator used a backend ATM network, they would find that it costs a little more, but multiplexes a wide variety of services and scales to meet a variety of configurations and capacities for high-speed deployments. If they deploy an additional service type that could not be run over Internet service (e.g. true toll quality voice or MPEG2 video streams), the existing switching equipment should be able to handle the capacity, but they may need to buy additional port cards or re-size a link. The same network can be used for multiple services. If the cable operator has deployed Ethernet based services and IP links between routers for Internet use only, they are likely faced with having to replace the IP routing equipment with high capacity routing equipment. They likely must also beef up the interconnection links or they are going to have to deploy a second network just for the new service. Both options are costly when compared to the former ATM model.

The Impending Data Ports Versus Cable Plant Ports Mismatch Problem

Initial deployment of high speed data services on all coax plants can typically be accomplished using one CMTS for the entire plant. Existing CMTS equipment today come in one of two scalability architectures: “fixed” scale configuration with one downstream port with only one upstream port, or “flexible” scale configuration with one or more downstream ports with one or more upstream ports. Incremental growth to meet new subscriber demand or capacity is different for fixed versus flexible architecture.

For incremental growth with a fixed scale CMTS, when more high-speed data capacity is needed in either the downstream or upstream, a new CMTS is required; i.e. the cable operator needs to purchase another CMTS box.

For incremental growth with a flexible scale CMTS, when more upstream capacity is needed, the cable operator can add an additional upstream channel demodulator (demod) card to the CMTS in addition to current channels. When more downstream capacity is needed, the cable operator can add an additional downstream channel or purchase a new CMTS box.

There is a large difference between fixed and flexible scale CMTS systems. With a fixed configuration, the operator must recombine upstream trunks into as few as many ports as possible to avoid having to purchase and abundance of fixed scale CMTS boxes. With a flexible configuration CMTS, each upstream channel may be connected to a different upstream trunk, eliminating any need for the recombination of trunks. This has two benefits: firstly, the cost of an additional upstream channel is generally less than the cost of a fixed configuration CMTS box, and secondly, the noise floor is reduced at the upstream port. The operator is free to distribute to trunks and combine upstream trunks with a flexible scale system.

At some point in the growth of service deployment, more downstream capacity will be required to meet subscriber demand. In a fixed scale CMTS, a new downstream channel is required for every upstream channel added and vice versa regardless of whether the downstream or upstream channel capacity has been filled by demand. In a flexible scale CMTS, the relationship of the downstream channels to the upstream channels within a single CMTS box are separately scalable, allowing the addition of downstream or upstream channels to follow subscriber demand. In addition, this flexible scale-ability allows for capital expenditures to more closely match revenue growth, and also allows for noise impairment to be better controlled by use of more upstream ports per downstream channel. This latter point is very important in that the cable operator has much more flexibility in managing the recombination of upstream trunks and subsequent noise funneling issues.

When the downstream channel capacity has been exceeded and not enough RF spectrum is available in the cable plant, the operator has the option of upgrading the plant to HFC. The upgrade to HFC will produce more downstream trunks and more upstream trunks. If previous CMTS installments matched capacity and revenue growth, there is likelihood that the existing installed CMTS equipment will match the newly available trunks and subsequent ports. Note that in this incremental HFC upgrade scenario, the cable operator has the option to do upgrades only where high-speed data capacity is needed, i.e., where the active subscribers and revenue is coming from. Upgrading the entire plant to HFC is not required.

Recombining return trunks at greater than four to one (4:1) causes noise funneling contribution and reduces the Carrier-to-Noise Ratio (CNR) below a 25 dB margin at the upstream return port. Several cable operators use this ratio. Converting the upstream lasers from FP to Direct Feedback (DFB) lasers allows the upstream return trunks to be recombined at a ratio of up to ten to one (10:1) which is attractive. If the plant currently has FP lasers, the cost differential to go to DFB is substantial and in most cases prohibitive.

High noise floor interrupts all upstream modulation schemes in an HFC plant. The ability to recombine upstream return trunks is limited by the lowest capable interactive service; for example, impulse pay per view, interactive two-way node management protocols, etc. The recombination problem affects more than just high-speed data services for Internet.

Solutions for the initial sparse deployment scenario are few. Either buy sufficient CMTS equipment to cover the upstream return ports or look into solutions that recombine data but do not recombine noise. Look towards CMTS solutions that support a large number of upstream return ports per downstream port.

The use of a Reverse Path Multiplexer has been shown to be effective in allowing the data from a larger number of return plant ports to be recombined without recombining the noise as compared to existing passive combiner techniques. As an active device, the RPM is placed downstream from the upstream channel port card in a CMTS. The RPM receives communications from the port card in response to when packets are expected to be received on a particular trunk. This high-speed switching system is effective at allowing a 200 microsecond packet burst with a small lead time and a rapid turn off.

Summary

The multiplexing of high-speed data over CATV networks and the business models of the cable operator are closely tied together. Different business models have different requirements for growth and capacity planning. The single flat rate per month is very straightforward to implement and grow. Multi-tier models and bulk bandwidth models each require more backend multiplexing than the simple mode and have benefit of increased revenues. The bulk bandwidth model is the most comprehensive of the business models, but allows a great deal of flexibility in how data is handled in networkp. The future may bring the need for additional digital data services to be multiplexed through a head-end. It is desirable to share the same backend switching network. This is advantageous as capital is only needed for incremental capacity increase and not for building a second parallel network for the new service. Data multiplexing over the RF has challenges for initial sparse deployments that can be overcome with CMTS equipment that supports many upstream channels and reverse path multiplexing.

References

[1] S. Beauchamp, "Throughput Analysis of Com21 Service Levels and ISDN," Application Note, Number 280-0024-D1, August 4, 1997.

[2] Kathleen M. Nichols and Mark Laubach, "On Quality of Service in an ATM-based HFC Architecture," IEEE ATM Workshop 96, August 27, 1996

[3] Kathleen M. Nichols, "Improving Network Simulation with Feedback," Internal Paper, Bay Networks.

[4] Mark Laubach, "Deploying ATM Residential Broadband Networks", NCTA Cable 96 Conference, Los Angeles, CA, 30 April, 1996.

[5] Kathleen M. Nichols and Mark Laubach, "Tiers of Service for Data Access in a HFC Architecture", SCTE Emerging Technologies Conference, January 1997.

[6] Mark Laubach, "ATM over Hybrid Fiber-Coax Networks", 2nd International Workshop on Community Networking – Integrated Media to the Home, IEEE Communications Society, 20-22 June, 1995, pp 27-33.

[7] Mark Laubach, "To Foster Residential Area Broadband Internet Technology: IP Datagrams Keep Going, and Going, and Going….", ConneXions, Softbank Expos, 2 February, 1996.

[8] Mark Laubach, "High Speed Data Systems Architecture and Tools for Managing the Upstream Noise Floor in HFC Systems", IEEE/SCTE, High Integrity Hybrid Fiber-Coax Networks Workshop, 7-9 September 1997.

[9] Mark Laubach, "Serving Up Quality of Service", CED Magazine, December, 1996.

[10] Mark Laubach, "Avoiding Gridlock on the Data Infobahn: Port Mismatches Pose Challenges", CED Magazine, April 1997.

Author Information

Mark Laubach
VP of Technology, Chief Technical Officer
Com21, Inc.
750 Tasman Drive
Milpitas, CA 95035
Laubach@com21.com
(408) 953-9175 (direct)
(408) 953-9299 (FAX)