BUSINESS AND PLEASURE: MIXED TRAFFIC ISSUES DRIVE NETWORK EVOLUTION

Robert L. Howald, Ph.D.
Motorola Broadband

Abstract

The term "triple play" originally was meant simply to convey the convergence of video, voice, and data in the network. It once encompassed everything that was important to know about network functionality. However, over the years the lines have blurred and the elements of the triple play have been further fine-tuned and splintered into a variety of different items. All three parts of the triple play are multi-faceted, and all contain important variables necessary for proper bandwidth and traffic management. Video has multiple, possibly interrelated, faces: analog, digital, narrowcast, on-demand, IP, and HDTV. Data, which once essentially meant DOCSIS 1.0, now includes the ability to support tiered data services that include best effort Internet traffic, like that offered via DOCSIS 1.0, as well as guaranteed and mission critical business services. Even simple residential Internet access is becoming a more complicated offering. With VoIP still in the wings, ideas such as online gaming have leap-frogged into play as part of the residential data mix. With data just as with video, possibilities abound that involve bandwidth and traffic management in both the access, backbone, and interconnecting points in the network. Successful deployment of tiered, prioritized, and guaranteed services include understanding aspects of the access network, including higher versions of DOCSIS, as well as non-DOCSIS solutions and technologies behind the HFC access network. Proper treatment of data services to and from the access network is a critical component of bandwidth management when considering architecture design options. Finally, voice circuits and IP voice also have a role in the redefining of the meaning of triple play.

This paper will analyze and characterize the traffic dynamics of the various service components above. Aggregation of these services in cases consistent with likely architectural scenarios will be discussed. Architecture and bandwidth conclusions will be drawn that align with the service and traffic mixes currently being offered. Finally, offerings such as gaming, security, and medical applications are some of the ideas among many potential services that have been mentioned recently. Their significance is magnified by the amount of highly interactive real-time voice and video needed to support them. The implications of such new service offerings will be discussed.

INTRODUCTION

At last year's NCTA show, a paper planted a stake in the quicksand [8], describing, as is self-described by the title, "A 10-yr Residential Bandwidth Demand Forecast and Implications for Delivery Networks." The presentation generated quite a few comments and questions, as any discussion related to predictions of bandwidth consumption might. In particular, the paper did not focus on perceived needs or the conventional "Field of Dreams" theory of bandwidth growth – that is, "if you build it, they will come." Instead, it was grounded in demand expectations based on market research, trend studies, and conversations with various individuals across the industry whom the information was shared with. The approach was to identify current and coming services that could be reasonably anticipated, evaluate and predict behavior, and aggregate the results. Obviously, the analysis was not done as an academic exercise, but as a tool that can be used to plan a business towards expected growth areas, and to engage operators in discussion that help them plan their networking needs.

As it stands today (now a third year into assessing the predictions), predicted behavior has deviated in a couple of areas, but none in earth-shattering ways. Only one service at this point is ramping more slowly than anticipated (VoIP). As a quick data point to encourage active minds, the study currently expects that, in 2010, there will be about 10x the bit rate demand on the forward path as there is in 2003, and about 25x in the return over the same period.

As valuable as this paper is to the company as a benchmark that is updated regularly, its original intent is only a piece of the puzzle needed to get a complete snapshot of the bandwidth and architectural evolution landscape. In particular, the results presented describe forecasts for the North American, residential market, and only the implications on the HFC access portion of the network. While there is certainly not a lot of extensive network infrastructure activity going on in the current slowdown, there is significant attention being given to services and equipment that have equivalent, if not more dynamic, impact on metro interconnect or backbone portions of the network. Examples include the growth in video-on-demand(VOD), the emphasis on supporting commercial services, and enhanced data aggregation and backhaul platforms. And, clearly, the paper described above, while concerning itself with evolution of the HFC plant aspects of service growth, implies impacts beyond HFC access. To complete the picture that enables the steadily predicted bandwidth growth requires a peek into what is going on outside the residential portion of the network, as well as at and behind the hubs that feed the distribution network. To do this, first we need to understand the relevant traffic engineering problems as well as access bandwidth problems for effective end-to-end system design.

This paper will introduce some of the concepts associated with the traffic engineering side of the problem. The problem at hand can actually be summarized quite easily. Never before has one network been asked to support so much content variety with such a wide range of quality of service (QoS) objectives. With this being the assumed case for growing cable operators, the following questions are being explored today:

1) What are the traffic implications of this multi-service, multiple goal situation?
2) What are the resulting architectural implications?

The fact that cable systems are able to encounter this type of problem at all and learn the important issues makes a strong statement for its competitive readiness in the larger picture of broadband providers. This paper will discuss scenarios that can be used to evaluate question one. There are as many answer to question two as there are opinions on optimal architectures. We will discuss common ones and general themes to be understood.

SOME MATHMATICAL TRAFFIC CONCEPTS

While widespread DOCSIS deployment brought data traffic management to the attention of the industry, the idea of understanding traffic characteristics and the effect on performance is not a new to cable. Archives at Motorola contain traffic studies aimed at understanding the response time of early settop IPPV request traffic, which used a basic ALOHA protocol. ALOHA is essentially a free-for-all that allows a user to send a message whenever ready and, basically, take their chances that no one else is doing so at the same time. If an acknowledgment is received prior to a time-out waiting for it, the user knows the message got through. The study goal was to understand how the settop loading and acknowledgement scheme effected the response time to a request from the user, determine the re-send likelihood, and understand the system breaking point. Implementation details were derived from this study. Through traffic modeling, the analysis was able to show that about 40% more settop returns could be accommodated at the HE equipment if the acknowledgments sent to the settop were modified in the way they were originally designed to be delivered. Clearly, equipment cost savings were directly obtained in this simple case.

As a second example, prior to full two-way activation of cable plants, Motorola had deployed an early SurfBoard® cable modem with a telephony return path. Traffic studies were commissioned to understand this drastic network asymmetry, the impact of PC hardware, and the effect of TCP/IP implementation on the PC. The analysis characterized how the telephony modems of the time (14.4 kbps and 28.8 kbps) compromised downloads that had much higher raw throughput capability. Performance of FTP transfer of large files was compared against symmetrical 10 Mbps and 100 Mbps point-to-point Ethernet to understand the user experience relative to, for example, the office environment. This data was use to support configuration guidelines for the product.

In the 1980's, Ethernet itself, standardized around a carrier-sense, collision detect multiple access (CSMA/CD MAC) protocol came under traffic analysis scrutiny. A widely referenced throughput analysis and testing study was performed as LAN technology began to explode during that period of time [3].

Telephone Network Simplicity

The purpose of traffic modeling is simple. By developing proper statistical models for data in the network, it is possible to predict pipe size requirements, bottlenecks, performance, and equipment requirements. Historically, there are two paradigms – telephone networks and data networks. The phone network is traffic engineered to minute decimal place (the "five nines") precision. The unit of Erlangs is used to describe voice traffic volume. The voice traffic arrivals are characterized statistically as a process with call arrivals that exhibit a Poisson characteristic. This information is used with the well-known Erlang formula to determine trunking capacity necessary to ensure that circuit availability can be guaranteed to the high level described above. Traffic engineering is possible with precision because of the well-understood nature of voice traffic with many years of historical precedent, and the single-service nature of that system at inception.

Data traffic, on the other hand, has not historically been heavily traffic engineered. Providing plenty of excess bandwidth has been the protection against performance degradation due to congestion, and there are still advocates of cheap bandwidth and less complexity as the way to continue. Others argue that, besides the inherent cost of higher performance equipment associated with under utilizing the network, flows are likely to encounter some bottleneck in an end-to-end system, particularly as the routes grow longer and more complex. Delivering repeatable QoS for high-performance services is not practical through pure bandwidth means in such cases.

We have mentioned the idea of Erlangs. Telephone system trunking curves – how many circuits must be deployed as a function of subscribers to assure a given blocking criteria – are available in many classic textbooks and papers. The results provide remarkably straightforward formulas for telephone network design – a formula that depends only on voice traffic offered (arrival of calls and duration) and the statistical assumption of a Poisson process for call arrivals. What statistical characterization can be used for other services, such as data? Are the answers as conveniently simple? Unfortunately, this answer is no.

Long-Range Dependence (LRD)

The finding that data traffic has a *self-similar* characteristic was one of the major traffic modeling discoveries to date for this relatively young discipline. Self-similarity – also called fractal or long-range dependent behavior – implies that, *regardless of time scale*, the traffic pattern has the same basic structure. When we say "traffic", we are talking about the baseband data volume and trends observed at the output of a CMTS or switch serving MPEG VOD streams, for example. This was an unusual finding, in that it indicates that there is correlation across much wider time scales than previously thought, and the assumption that smoothing occurs when observed over long periods was proven inaccurate. Another surprising way to envision this characteristic is to think in terms of our basic understanding that data traffic is bursty, which we usually associate with short time dependence in our minds. However, self-similar traffic indicates that long bursts separated by long time intervals are characteristic of the traffic as well. Intuitively, we would have expected the wider time scale to smooth out the peaks and valley around a mean.

The seminal paper showing self-similarity at work was based on an Ethernet analysis, but because of the astounding nature of the discovery, others were inspired to look closely at their own assumptions. Subsequent findings included self-similar properties of ATM traffic, metro area traffic (MAN), wide area network traffic (WAN), and also for multimedia traffic, such as compressed digital video streams and Web traffic.

The unearthing of self-similarity created a camp of network theorists that felt that the book on traffic theory now had to be re-written. Traditional models generally focused on Markovian behavior, which relies on limited memory of prior traffic – in other words, correlation is lost over time, and smoothing out occurs as the time scale is broadened. The stock market is a good example of this expected smoothing, although studying these curves are probably best avoided at this juncture. The impact of this correlation lasting over broad time periods has implications for policing, scheduling, congestion control, and statistical multiplexing gain.

The surprising finding naturally led researchers to search for the reasons for it. The cause of self-similar behavior was found to be associated with the fact that the distribution of transmission content is *heavily-tailed*. That is, the tails of the probability density function do not decay rapidly. What this means is that, rather than seeing the likelihood of the size of a transmission flow occurrence decreasing exponentially as the flow size increases, this drop-off in likelihood is not so drastic. There is a very wide variation in the size of packet flows that throws off traditional statistical models – large files, MP3's, JPEGs, database activity. In fact, it has been shown that such heavily-tailed characteristics are a *sufficient* condition for self-similarity. As important is to recognize what self-similarity is *not* caused by. The breadth of examples indicates that self-similarity is not associated with the delivery format generated to carry the information – i.e., it is not a protocol artifact.

Now, obviously, cable systems deploy equipment for carrying multimedia traffic. In particular, compressed video streams are on today's cable transport networks, with today's most relevant example in terms of equipment growth and network design being video-on-demand (VOD). Based on the above, the traffic characteristics will be the same whether the video delivery is MPEG over IP – such as GbE-based transport – or the other way around. And, of course, cable companies are interested in moving data around in the form of Internet traffic from CMTS's to ISP points of presence, and data from business services, both Internet directed or otherwise.

Summarizing, then, understanding the role of self-similar traffic patterns is valuable for the applications above in designing the HE to hub or hub-to-hub interconnects. As networks become more integrated, the value of understanding traffic increases as the aggregation pushes bit rates higher, making efficient use of resources yet more important. As movement of different types of traffic becomes integrated, there is the further need to ensure the QoS support for each. Providers must therefore understand the implications of traffic characteristics and the distribution of QoS needs of each.

<u>M/Pareto Model</u>

While we have explained and described a fundamental and surprising trait of many traffic types relevant to cable, making use of this model for statistical calculation requires fitting this knowledge into a distribution. The characteristic described has been shown to be a result of an aggregate of bursts of widely varying sizes. A model based on randomly arriving bursts with a heavily-tailed distribution is therefore called for. A Pareto distribution, commonly described in statistics texts, is combined with a Poisson arrival rate of overlapping bursts to create a mathematical realization of the situation. More specifically, data traffic is assumed to be bursts with a Poisson distribution and associated arrival rate, where each burst is of duration described by a Pareto distribution.

The M/Pareto model has several variables associated with it, including the Poisson arrival rate information. This portion of the model has been shown to be important to accurately curve fitting real traffic to it [2]. The essential "real" traffic property captured by varying the Poisson parameters is the amount of traffic being multiplexed onto a pipe for characterization. This approach is a valuable step to a traffic model comprised of an aggregate of multiple sources of independent information. The aggregation of traffic is not significant enough to use mathematical assumptions of Gaussian behavior driven by the central limit theorem. Models based on long-range dependence provide network designers with a tool for developing architectures and equipment requirements that support the aggregated traffic. This is important to capture, as issues associated with self-similarity drive network changes in queuing and congestion control mechanisms then Markov-based assumptions would imply.

<u>Gaussian Behavior</u>

What is occurring on the Internet and to a similar extent in breadth on HFC networks is the aggregation of more traffic and more traffic types from independent sources. It is not difficult to envision the challenge this growth entails; yet, this very growth and the evolution of integrated networks is potentially a blessing to the traffic modeler. The central limit theorem provides a fundamental statistical underpinning for what the nature of the traffic over time could evolve to – Gaussian behavior.

The central limit theorem is the basis for many natural phenomenons that exhibit Gaussian behavior, and, in the case of aggregated traffic, it becomes asymptotically so when many independent contributors - under some minor, but important, caveats - are aggregated. The convenience of this is that Gaussian statistics are very well-studied and understood, and if the traffic statistics can be assumed Gaussian, then many simplifications can occur and probabilities of occurrence characterized. Multiplexing gain can be predicted under Gaussian assumptions, and pipes designed efficiently for some pre-selected level of congestion avoidance. This can be used to support a desired set of policing, shaping, scheduling, and queuing mechanisms. While rapid traffic growth makes network evolution difficult, the bandwidth explosion, in general, is good for business, and the handling of traffic from a bandwidth boom potentially makes it more readily predictable.

SERVICE SET

DOCSIS

Since the wide acceptance and deployment of DOCSIS, all operators and vendors have interest in traffic characteristics of essentially this same basic system. As a result, there have been many articles on configuration of the CMTS and guidelines for a DOCSIS-based system setup. The paper described previously [8] suggests a doubling of DOCSIS return path traffic each year, a result that is a combination of take rates, modulation profiles, and usage of the medium by subscribers. This variable in the paper is dedicated to residential cable modems – i.e. Internet users at home. This doubling effect is corroborated by other, more general Internet traffic studies that suggest a "Moore's Law" for data traffic [7]. This analysis notes that this trend has been pretty reliable except for a period in 1995-96 where there was a burst of greater growth attributed to simplified web browser breakthroughs that led to mass acceptance, and the subsequent changes made by online providers to graphically rich interfaces. Further traffic related information from trend studies indicate that access to broadband via DSL or cable modem results in a user increasing their time online by 50-100%, and that the bytes consumed per month increase 5x to 10x as well.

A very informative paper based on a project at CableLabs and also presented at last year's conference [12] offered a first real comprehensive glimpse into DOCSIS traffic. Summarizing some of the key findings:

- Daily activity is a slow build throughout the day, with "busy hours" between 8 pm – 12 am (peak), and a subsequent rapid drop-off until beginning again at 5 am

- Traffic is seasonal – following school holidays and vacations

- Traffic asymmetry decreases from 3:1 to about 1.5:1 as familiarity and capabilities set in

- DOCSIS 1.1 enhancements to support voice traffic result in a 15% efficiency improvement over DOCSIS 1.0

The seasonal phenomenon represents the dominance of traffic by a younger generation of user. This particular phenomenon should become less pronounced over time as these kids become tomorrow's adults, although a generally heavier level of usage by the academic community may linger.

DOCSIS 1.1 provides the ability to support VoIP traffic. It does so by supporting multiple classes of service (CoS), whereas DOCSIS 1.0 supports only one – best effort. DOCSIS 1.1 also allows packet fragmentation to ensure that latency-sensitive voice traffic is not bogged down behind large "best effort" data packets. From a traffic generating standpoint, DOCSIS 1.1 also implements pre-equalization at the CM side, a physical layer technique similar to pre-distortion, but for bits. This feature permits more practical use of the 16-QAM mode, which doubles the bit rate compared to a QPSK channel of the same symbol rate. In other words, the 160 ksps mode, which results in 320 kbps for QPSK, provides 640 kbps in 16-QAM mode. The result at the hub or HE is more bits-per-second pouring out of a CMTS spigot.

DOCSIS 2.0 also speaks foremost to raw throughput enhancements, It provides for a dual, selectable, medium access control, or MAC (S-CDMA or A-TDMA), and an enhanced modulation profile capable of 64-QAM at twice the previous maximum symbol rate. Raw capability is now about 30 Mbps. Built-in as well are enhanced interference mitigation techniques for narrowband and burst interference make use of the newest modes realistic, and use of lower return channels possible, creating more effective bandwidth. The CMTS spigot therefore just got wider, or the pipe became more fully utilized.

In summary, with DOCSIS we can expect rapid, raw, bits-per-second growth, more efficient bandwidth consumption in access, self-similarity in backhaul, and both best-effort and class-of-service (CoS) mapping.

Peer-to-Peer (P2P)

Although the proliferation of peer-to-peer communication still generally falls under a residential data (DOCSIS) discussion, this phenomenon is significant enough to warrant special mention. It is, of course, certainly the case that broadband access has, in fact, *enabled* P2P traffic to become as significant as it has, freeing downloaders from the limitations of dial-up speeds to deliver multi-megabit files. Peer-to-peer traffic – pioneered by music file sharing through Napster, but subsequently followed up by similar services (Kazaa, Morpheus) – raises a significant flag to operators of broadband networks who observe and/or police their traffic patterns.

The impact of heavy P2P traffic on cable modem systems offering no byte count or rate limits is twofold. The raw bits-per-second load sees a "bias" around which the web browsing peaks and valleys vary. The effect is to have a constant offset or mean value associated with the streaming file content much like any constant bit rate (CBR) application. The difference in the modern case is the high bit rates associated with the CBR-like traffic are rates that are considered high speed. Of course, with no rate capping or policies in place to limit this type of traffic, a very small number of users can essentially dominate the throughput of the link. Enough heavy P2P users or enough sharing of single return channels among users can therefore create a congestion scenario, as the demand for transmission time slots upstream outstrips supply. Simulations [12] support this effect. A single users acting as a source of MP3's, when placed among a dozen or so other users, managed to consume up to half of the return capacity over significant period of observed time. This creates a clear forward-looking argument for tiered service offering based upon either rate or volume limitations with the necessary prioritization and policing schemes to enforce the tiered structure.

To summarize, we can therefore add to our prior discussion of DOCSIS the need or objective of supporting tiered services.

Enterprise Traffic

The play for commercial services can be attacked primarily in two ways – DOCSIS-based and fiber-based. Which solution is determined by service needs at the business, which basically boils down to the size of the enterprise. The key features that DOCSIS 1.1 provides that make it a reasonable solution for a subset of the business market, in addition to cost, are the support of voice, enhancements that offer multiple service classes, enhanced security, and, finally, a more realistic opportunity to achieve 10 Mbps type of performance due to physical layer improvements that add pre-equalization. Of course, 10 Mbps has a nice ring to it in light of comparison to 10Base-T LAN environments. Finally, DOCSIS 2.0 encompasses the key features of DOCSIS 1.1, but also offers the 3x increased capacity return due to symbol rate and modulation improvements. In addition, the protocol advancements inherent in A-TDMA signaling and S-CDMA are designed to expose more return bandwidth to the operator for high-speed services that previously had been unusable for this type of traffic.

Based on the above, we can summarize by saying that business services over DOCSIS 1.1 represents for the operator a need for QoS tiers, and service level guarantees, similar to previously mentioned DOCSIS needs for peer-to-peer traffic.

Fiber-based solutions can deliver higher levels of service, corresponding to larger businesses or business campuses. Of course, the target here is to provide cost-competitive voice and data services with the same level of service experience to the end customer – data rates, security, reliability. While residential data traffic has grown rapidly as previously described, business data traffic has taken a more modest trajectory, and business voice traffic is essentially flat. Thus, while the per-subscriber business needs are higher, the growth in this data sector is more gradual, meaning the pipes assigned to support a fiber-to-the-business application may have longer legs than expected. This is somewhat intuitive, in that, while residential users find new and bandwidth consuming ways to exchange and download content, the shift in business transaction content has not been this dramatic. In terms of security, virtual private networks (VPNs) translate to both security and bandwidth guarantee issues. In terms of reliability, resiliency and 50 msec recovery characterize common "carrier class" attributes expected by business customers.

Fiber-based systems that segment spatially or via wavelength – the only non-DOCSIS, HFC-centric approaches today, have no aggregation issues in the access network. As the data hits the first aggregation point in a hub, it is at this point that it may have to be managed, prioritized, and scheduled alongside other traffic requiring bandwidth. If both business voice and data exist, an architecture that guarantees bandwidth for voice and queues data packets may be necessary. Architecture decisions are made with these types of network crossroads in mind – the service mix at these points could consist of VOD content, CMTS activity that could include voice and tiered data, business voice and data, even broadcast digital video content.

In summary, then, overall higher symmetrical bandwidth and quality of service guarantees are important to this market, and traffic growth to plan for may be less dynamic.

Networked Gaming

The average age of a "gamer" is on the rise. The Gen-X and Gen-Y demographic that grew up with the phenomenon of Playstation, Nintendo, and now X-Box are now growing up themselves – if perhaps in age only. They are bringing their bad habits with them, including their passion for gaming. At last year's Western Show, not much was well-attended. But, one of the best attended sessions dealt with the gaming phenomenon and the impact on broadband.

Intuitively, what would we expect gaming needs to be? Certainly it is real-time interactivity, the magnitude of which needs quantification relative to the well understood baseline needs of voice traffic. Latency is of key importance, as thumb actions must be translated into game states and communicated rapidly to other players. As important yet is probably jitter variation among peers, so that the game can be fair to everyone, and no one has a built-in edge. Studies have shown that delays that are not noticeable to voice traffic users are quite noticeable to gamers, and, indeed, can affect game outcome.

Again, intuitively, the nature of gaming activity would not be expected to have the statistical look of web browsing or streaming media. Recent work has analyzed traffic distribution of the popular "Quake" application [4], observing both packet arrivals and sizes for clients and of the game server. The results of this study suggested that an Extreme distribution is a good fit for packet arrival of both servers and clients in most cases, as well as packet sizes of server traffic, where the server is the point from which game states are updated and broadcast to the clients.

The parameters of the distribution that define the exact shape of the broader family are functions of processing speed distributed among the servers and clients – another effect that would seem intuitive if there is no network bottleneck. Of course, making sure this is *not* the case is what our job is all about. Some observations suggest that a split statistical model – part deterministic and part exponential – is a better fit for client packet arrival in some cases. Client packet sizes were found to be deterministic. Future research is ongoing.

Today, gaming traffic is a minor contributor to overall volume. The number of gamers is relatively small, although the certainty of this growing is about a sure a bet is there is in predicting network usage. Secondly, however, gaming transmission volume is small, as the games merely transmit state information that is used by the software at the ends to actually render the complex video images. Migrating game activity to virtual reality based image transport could change the transmission volumes dramatically.

Video-on-Demand

Data growth has some historical precedent, including recent trends with cable modem users. These trends offer anecdotal evidence supporting the bandwidth growth assumptions predicted in [8]. Video-on-demand has not been characterized as carefully. However, that VOD has accelerated rapidly in the last couple of years is not news, with total VOD revenues increasing nearly 10x between 2000 and 2002 [14], and the expectation it will double again by the end of next year. It is also not news that the bandwidth needs of video content as a service to one subscriber greatly exceed that of a data service to that subscriber. VOD services are well down the path of some of deploying the widest-pipe and most cost effective technologies, including Gigabit Ethernet (GbE) and WDM.

The choice of technologies above are the same as those discussed to support the bandwidth growth in residential data. Because of this, carrying them intermingled with one another seems like a logical step for added efficiencies in the network. Important questions to consider are those associated with the QoS needs of each, and, if necessary, how Ethernet-based transport is augmented with added robustness to ensure such needs are met. For example, over 200 MPEG movies are carried on a single GbE pipe. The case for redundancy and fail-over becomes more compelling as pipes widen and carry more traffic. Whereas data implementations today may have Layer 4 mechanisms providing loss and flow control, loss of unidirectional streams represent a non-trivial protection situation, not to mention a bad day to be near the customer service center.

In terms of traffic, VOD has some obvious diurnal dynamics. Of course, most network decisions rely on demand requests during peak usage hours on peak days. The dynamic range is quite wide, and makes for some tempting unused bandwidth in off-hours. Fortunately, these behaviors are predictable, and fortunately, as an example, peak VOD hours will not coincide with, for example, peak busy hours of enterprise data traffic. By the same token, however, peak VOD busy hours may roughly coincide with peak residential Internet usage hours.

In summary then, aggregated VOD streams present to us a rapidly growing, high QoS, wide bandwidth application, the needs of which today are essentially met via silo networks. The technology trends, however, point towards the same technology choices expected for data growth. VOD traffic varies in daily and weekly trends in predictable ways over time.

<u>IP Video/Audio</u>

Streaming content has received quite a bit of air time in the past few years, while during that time P2P traffic was really what caught a buzz about it and took off. What constitutes streaming traffic and P2P begins to blur, but, in general, the concept of streaming media conventionally applied to the idea of a content providing service streaming IP to a computer terminal or settop box connected to a computer (or even a TV). The reference bandwidth projection predicts that this type of traffic will grow to be about 18 times as large between now and 2010. It current contribution in term of bandwidth consumption and traffic engineering is negligible, although it has potentially large architectural impacts if the enabling technology to light this fuse is all-IP, all the time. This is a visionary decision or timetable – depending on your perspective – every operator must make for the future growth and service providing capabilities of their system.

From a traffic standpoint, we have discussed what P2P traffic does to a dynamic set of flows of residential Internet. The effect is to create a steady “bias.” In other words, the mean bits-per-second increases, and the traffic dynamics exist on top of this mean. For streaming media, this same effect would be the expectation when it becomes significant enough to matter, except that this would be a downstream phenomenon. As such it could be more buried in the noise depending on the asymmetry experienced in the network. Obviously, if the mean is very large in comparison to the peaks and valleys, the impact of peaks and valleys on efficient pipe usage is very minor. In other words, if the streaming media content (or VOD content for that matter) dwarfs data transport content along the same conduit, the traffic engineering pipe size problem is simplified, since the ups and downs will be relatively small.

Now, a significant difference between streaming media and most P2P traffic today is that the former is real-time content, deserving of QoS capabilities that do not require the same level of sophistication as moving MP3’s and JPEG’s around. Since the content still exhibits LRD regardless of whether data or streaming media, mechanisms at the ends of the pipe that enforce policies and switch packets need awareness into the proper traffic models of this LRD, so that queues can be build and implemented to avoid dropped packets and blocking and supply the QoS expected for real time content.

Summarizing, streaming multimedia represents content characterized in prior work as having self-similar behavior. It is a relatively high bandwidth consumer on a single user session basis with the QoS needs of other real-time media, but the total bandwidth usage is low and growth path very dependent on architectural and service choices going forward.

Digital TV or HD

Broadcast digital TV has much the same characteristics as VOD. There are two main differences. First, in many cases, linear supertrunking is used rather than digital transport as a low-cost alternative when high bandwidth digital backbone is not in place. Second, the service group size is very large, making redundancy of path and equipment quite important. High Definition streams, for the most part, represent to the network engineer, digital TV traffic on steroids. The bandwidth hungry nature of HDTV is a promising possibility for inspiring networking bandwidth upgrades.

Similar to streaming media – and in fact analogous except for the more standardized format of delivery – digital TV provides a steady flow of packets and a nice averaging of bandwidth behavior to a bit rate number that is a function of compression and statistical multiplexing of a large number of streams. If anything rides along the same channel, such as VOD, then VOD dynamics would be superimposed.

ARCHITECTING FOR MULTIPLE SERVICES

QoS Parameters

There is no universal definition of Quality of Service QoS), just as there is no universal definition of “carrier class.” Nonetheless, QoS encompasses basically five parameters:

- Latency – End-to-end absolute delay
- Jitter – End-to-end variation in delay
- Loss – Dropped transmissions
- Throughput – Bits-per-second or Bandwidth
- Availability – Likelihood of the network being “up”

QoS has achieved buzzword status quite recently, and its footprint is all over standardization committees. What is essentially going on are efforts to bring to the data world something it has always lacked – guarantee-able QoS – but doing so on a “connectionless” network while keeping as much legacy frame and protocol structure intact as possible. The result is primarily profitable to the acronym maker (or marketing team). The effort is a logical outgrowth of the indisputable fact that Ethernet dominates the LAN. As the LAN aggregates to the MAN and WAN, the goal of leveraging the broad familiarity with Ethernet, its cost points, and the flexibility of features within Ethernet and IP as protocols have driven traditional Ethernet and IP network designers to innovative approaches to solving this classic QoS shortcoming in the hopes of scaling the local LAN to a broader market.

Not surprisingly, some techniques to enhance Ethernet resemble old ideas. For example, the use of the DiffServ protocol (differentiated services), when implemented over MPLS (Multi-protocol label switching) has the look and feel of ATM, but with a lot of different acronyms describing the details. DiffServ provides the ability to classify a packet with a forwarding class describing its priority on a per-hop basis. The latter fact actually limits the overall QoS strength of DiffServ on its own, but increases its practicality. The role of MPLS is to expedite the forwarding of packets through the network by creating label-switched paths via tags on the Ethernet frames, directing packets at Layer 2, rather than making route decisions through the network that create processing bottlenecks and subsequent latency and jitter problems. Thus, these schemes together offer prioritization of payload types, and create predefined and expedited paths through the network. This scenario has indisputable similarities with ATM.

So, why re-invent the wheel? The answer is simply because the dominance of Ethernet and flexibility of IP have made riding this wave a necessity in network design, to the extent that incrementally upgrading the technology to carry more than best-effort is more palatable than addressing major equipment and protocol overhauls and learning curves.

QoS – Who Needs It?

Let's list some of the services encountered or viewed as on the horizon. How do these compare as far as who needs what for QoS? Let's use a simple scale: High (H), Medium (M), or Low (L) need for the particular QoS parameter below. A qualitative summary of QoS needs is shown in Table 1 below. Certainly, there is plenty of room for debate (I adjusted this chart more than a dozen times), and some still require more learning and evaluation.

Table 1 – Services and QoS Need

	Latency	Jitter	Loss
Residential Data (DOCSIS Internet)	L	L	L
Residential Voice (VoIP)	H	H	L
Business Data (DOCSIS 1.1 or higher)	L	L	H
Business Data (fiber-based)	M	L	H
Business Voice (VoIP or T1/T3)	H	H	M
Business Video (videoconference)	H	M	L
MPEG or IP Video or VOD	M	H	M
HDTV Broadcast	M	H	M
IP Audio (Radio AOL)	M	M	M
Interactive Gaming	H	H	H

Clearly, we can recognize that some applications are real-time, while others are not. This fact primarily drives the latency and jitter QoS needs. Also, based on the nature of the service, it may be loss tolerant or not. In general, if the content itself is to be transduced for human senses, it is likely to be loss tolerant to some extent. Human senses are quite effective as filters. If the content is information for a computer to interpret and process, it is likely to be less loss tolerant.

What tools exist to assure the level of QoS desired is achieved? In the DOCSIS world, the use of DOCSIS 1.1 or higher, and a next generation CMTS [13] provide this capability. The CMTS is a key element between the access and transport network, acting as both a media converter at layer 1 and protocol delineation point for layers 2 and 3. DOCSIS 1.1 provides the class of service capabilities on the HFC side, while advanced layer three implementation such as per-flow queuing provide traffic management functionality on the network side. Thus, for the first three items in Table 1, DOCSIS 1.1 and next generation CMTS enable the providing of QoS mapping from access network to interconnect ports. For business voice and data based on fiber connectivity rather than DOCSIS 1.1, and segmented spatially or via wavelengths in the access network, QoS schemes must reside in the aggregation equipment and supported elsewhere in the architecture.

On the video transport side, such as VOD and HD transport, QoS is assured by the fact that these systems currently are essentially silo systems. Statistical multiplexing occurs as server content traverses switches, but the rules of engagement are simplified by the singular content and rules easily developed from this simplified, application-specific architecture. Should these services become part of an integrated triple-play transport network, the dynamics of the traffic situation could change significantly. For example, for a single IP pipe supporting video and data, there would be a heavy reliance on IP QoS schemes through some of the existing standardization efforts to assure the video QoS needs are met. The end-to-end capabilities are not of the "guaranteed" variety, and would certainly require some traffic engineering and modeling.

Streaming media has the same type of architectural implications in the network, with the difference being that DOCSIS supports the access portion of the network. Thus, mapping of QoS mechanisms from one side of the CMTS to the other – again using the CMTS both the due media and QoS transducing – is needed.

Architecture Technologies

Clearly, many services with many different needs are set to co-exist. Sound business practice means finding an efficient means to handle them by judicious choice of technologies, levels of integration, and a healthy concern for operational costs and scalability. The term "triple play" alone sounds like an abbreviated set of services, but the flavors within the triple play, as shown above, clearly make the problem more complex. The access network itself is constantly being re-thought for fresh ideas, such as data overlays, wireless interfaces, and intelligent processing. At the aggregation points in the hub and Headends, various technology choices exist, some of which are deployed already in silo networks as previously indicated in VOD cases. VOD represents a good reference example because of where it is in the cycle – basically just at or past the "knee" of one of those classic marketing hockey stick charts, depending on which analyst you ask. It is a service experiencing significant growth, and it is using technology on the move in both the server and multiplexing arena, as well as in the transport pipe. VOD transport has been migrating from DVB-ASI transport to lower cost, more-flexible, GbE links. And, again, use of GbE technology makes its way into the data world as well. VOD also has been a driving application for another key technological option - wave division multiplexing (WDM).

Gigabit Ethernet has some notable shortcomings as an all-inclusive answer for network design. Ethernet as well as IP over Ethernet – both designed for data – do not inherently offer the resiliency, availability and network management attributes that become important elements of a "carrier class" solution when so much content is riding on the success of a single link. IP over Ethernet was developed as a "best effort" technology, and most of the ongoing efforts today revolve around finding way to be better than best effort. The previously mentioned DiffServ and MPLS developments fall into this category, and there are others. Thinking in terms of end-to-end IP as an attractive end game, MPLS, in fact, can be viewed as a way to skirt the routing limitations of an all-IP network by avoiding the per-hop calculation of routes through the network. Not all developments aimed at traffic engineering and hardening data systems are completely new, however. TCP/IP itself is a kind of QoS feature, and type-of-service (ToS) header bits have been around since 1981. Limited capabilities of these features – invented still in a "data world" context – limit their power to meet the kind of diverse needs expected.

Another Ethernet issue is that it does not inherently support circuit-based voice. Technologies exist to create virtual circuits over Ethernet. Similarly, at layer 3, voice over IP (VoIP) has been developed technologically. Each today has cost penalties. In short, however, Ethernet, even GbE or 10 GbE, and even if we include wave division multiplexing (WDM), cannot go it alone. All-IP looks attractive from an interoperability and flexibility standpoint, but the jury is out on guaranteeing that all of the QoS needs can be met even with the traffic engineering tools brought to the table with MPLS and other tools designed to optimize packet transport. Actually, right now, use of enough wavelengths and 10 GbE would be essentially the oft-practiced lazy man's QoS – gobs of bandwidth assuring that nothing gets held up. For the price of this overcapacity, if well managed, there would be no congested routes, and no pipe traffic peaks requiring buffering and queuing delay. There are smarter ways to solve the problem rather than relying on this brute force approach.

WDM itself is, in general, a brute force capacity enhancing tool. It allows a single fiber to carry multiple data-bearing streams, such as GbE or 10 GbE, by using a different wavelength for transmission for each. However, WDM does not easily offer QoS consciousness. This is not a show-stopping issue - intelligent wavelength management exists as a relatively mature technology. Integrating wavelength management as part of system resource management has been an anticipated direction of the telecom sector for some time, and could see relevance in cable networks as well since the same premise drove that thinking – bandwidth growth and support for advanced services.

The classic shortcoming of QoS is precisely the reason for equipment based on Next Generation Sonet. There is no debating the QoS features of Sonet transport. The knock against Sonet had to do with its rigid structure that made it inefficient as a packet data transport system. Coarse bandwidth increments left excess unused bandwidth, driving up the effective cost of doing business by decreasing fiber usage efficiency.

However, precisely these data issues are addressed with Next Generation Sonet, all the while holding firm on the guaranteed, proven, mature resiliency and reliability that has no comparison today among alternative technologies. Furthermore, there is lots of existing Sonet-based infrastructure. What modern equipment does is build the data flexibility into formerly coarsely-grained all TDM-only platforms through virtual concatenation (VC) and link capacity adjustment (LCAS) which allows dynamic – i.e. supporting data – provisioning of VC carriage. Such platforms contain port interfaces natural to data handling, such as 100BaseF and GbE, as well as the traditional voice-related interfaces of a traditional Sonet platform. Other data-oriented features are included as the platforms evolve to meet the shift in traffic demand. Support of standard framing protocols for data is gathering momentum (Generic Framing Protocol or GFP), as well as standards-based packet classification and guarantee-able class-of-service mapping.

Next generation Sonet platform, then, offer the guaranteed resilience and inherent QoS parameters for both TDM and Ethernet services – the resiliency still unique to Sonet – and now offer the added granularity and flexibility to efficiently support data needs.

The networking world never has quite enough protocols, and one of the latest is aimed at providing efficiencies of packet-based transport natively, but having the resiliency characteristics of Sonet. Avoiding the *traditional* Sonet limitations was a key target of this development effort. A standards body has been formed, IEEE 802.17, that is developing the layer 2 protocol known as resilient packet ring (RPR). Not surprisingly, the basic frame structure of RPR was based on Ethernet, and adds to it MPLS and Class-of-Service header content, as well as other fields.

RPR is a ring protocol, with the objectives of optimally supporting all previously described traffic types, maximizing efficient use of counter-rotating ring bandwidth, and simplifying network provisioning. The critical importance of packet resiliency is recognized as a key focus, as the traffic mix and amount no longer fit into a "best effort" paradigm as was once the case. As a layer 2 technology, RPR can run on entrenched physical layers (i.e Sonet and Ethernet PHYs), which is important considering the amount of deployed infrastructure. The standard is still an emerging one, and the discussions break down into two camps – Cisco and others.

WHAT NEXT?

What to make of all of these services, technologies, and, in general, the many choices that face network designers today? Needless to say, there is no single magic answer, and recommending what works best for 2010 will simply create competition for someone's infamous "who needs more than 64k of memory" comment. The good news is that cable operators are in the drivers seat at the moment. The competition for service provider of choice is arranged in their favor. But it is unclear how long this will last given alternative solutions and the allure of residential broadband as something people will pay for, a few-and-far between uptick in a beaten down economy. To capitalize on a game that is cable's to lose, and since there is no one-size-fits-all solution at this juncture, what we can recommend is a five steps "keys to success" approach:

1) Have a planned roll out of services to provide or support. A comprehensive sample set is provided in this paper. *This first step is, as always, a business exercise.*

2) Develop a bandwidth forecast of your own, and comprehend the traffic and QoS aspects of each service. The given information and references provide guidance towards both. *This is a business* ***and*** *technical exercise.*

3) Understand the current and growth capabilities and limitations of any existing infrastructure in the context of 2). This especially includes the emerging standards and techniques suited to the evolving service mix, and the flexibility available where there is infrastructure to build. The previous section points the way towards much of the

relevant activity in this area. *This is a technical reading & research exercise.*

4) Develop a time-phased service integration and network evolution plan for the traffic mix (not necessarily the same as an integrated network) aligned with 1), 2), and 3); the hard work of the first three steps make this more straightforward than it sounds. *This* is *a business* ***and*** *technical exercise, and the one that decides whether you grow, maintain, or flounder and become exposed to competition.*

5) Know what end-to-end means to your network responsibilities, and modify any "silo" role & responsibility organizational definition to smoothly evolve across the network interfaces, features, standards, and mapping schemes. *This is a technical management exercise. Technical managers often used to be technical people and carry with them technical biases, so this one may be more difficult than it sounds.*

CONCLUSION

On the QoS front, one of the more compelling human stories in recent years occurred when a researcher in Antarctica found herself stranded with her team during an inaccessible time of year for rescue missions and in need of critical medical attention. Much like the space shuttle engineering techniques being developed and discussed herein.

And what of 2010, based on the forecast previously referenced? Architecturally, will the triple play be implemented in a unified network, with today's darling being an all-IP format all the way to the home? Will some services continue to ride over separate parallel networks optimized to the bandwidth and QoS required? According to the forecast referenced, 59% of the forward path digital traffic demand will be Internet access, 26% will be some form of VOD, 13% will be streaming audio or video content of the PC variety, and the remaining IP telephony. Do

Columbia tragedy this past winter, the video phone images and stories of the researcher and her team offered a glimpse into the risks of pushing the envelope. Fast-forwarding to the present, recent literature describes the wireless infrastructure in an Alabama hospital [9]. The article also describes the network services the hospital uses to transport mission-critical data. The concept of remote medicine involves supporting transmission of high-resolution images and data in real-time to doctors to serve live patients anywhere, such as remote locations, a third-world country, or on a battlefield. Obviously, administrative and legal obstacles abound anywhere medicine is involved and lawyers are prowling. But the advantages possible are access to expert advice in a timely manner, access to trusted medical advice internationally, access to observation of expert and advanced procedures by other doctors and students, and less dependence on getting out and about to receive care when seriously ill. Can we invent a higher QoS service need? The military designs and implements its own private networks for mission critical data, because there is absolutely no margin for error. But, it can afford to as well, while applications such as the above would rely on the kinds of QoS and traffic

you buy this perception of today? Would this aggregation mix lend itself to a particularly convenient model? Some of the constraints of the previously introduced central limit theorem, aside from a variety of independent sources, is that the independent distributions have finite variance, and that there not be a singularly dominant distribution. Thus, while the broad traffic mix implies central limit simplification, a single dominant service can disturb this convenience. Furthermore, a cornerstone characteristics of self-similarity – heavy tails – can also imply infinite variance, another central limit theorem killer. The jury is once again out, as research continues to classify traffic trends and distributions for network modeling and optimal architectural design. Similarly, networking technologies will simultaneously evolve, making putting a stake in the quicksand that much more perilous.

Contact Info:

Rob Howald
Director
Systems Engineering
Network Infrastructure Solutions
Motorola Broadband
rob.howald@motorola.com

References

1) Adas, Abdelnaser, *Traffic Models in Broadband Networks*, IEEE Communications Magazine, July, 1997.

2) Addie, Ronald, Moshe Zukerman, and Timothy Neame, *Broadband Traffic Modeling: Simple Solutions to Hard Problems*, IEEE Communications Magazine, August, 1998.

3) Boggs, David, Jeffrey Mogul, and Christopher Kent, *Measured Capacity of an Ethernet: Myths and Reality*, Digital Equipment Corp, Western Research Laboratory, Reseacrh Report 88/4, September 1988.

4) Borella, Michael, *Source Models of Network Game Traffic*, 3Com Corporation.

5) Brownlee, Nevil and KC Claffy, *Understanding Internet Traffic Streams: Dragonflies and Tortoises*, IEEE Communications Magazine, October, 2002.

6) Chu, Tam-Anh, *WAN Multiprotocol Traffic Management: Theory and Practice*, Communications Design Conference, September, 2002.

7) Coffman, K.G., and A.M. Odlyzko, Internet Growth: Is there a "Moore's Law for Data Traffic?", AT&T Labs Research, July, 2000.

8) Nash, Randy, *A 10-yr Residential Bandwidth Demand Forecast and Implications for Delivery Networks*, 2002 NCTA Technical Papers, May, 2002.

9) *Look Doc, No Wires*, Communications News, March 2003.

10) Robert, Jim, *Traffic Theory and the Internet*, IEEE Communications Magazine, January, 2001.

11) Sahinoglu, Zafer and Sirin Tekinay, *On Multimedia Networks: Self-Similar Traffic and Network Performance*, IEEE Communications Magazine, January, 1999.

12) Shaw, Terry, *Bandwidth Utilization on Cable-Based High Speed Data Networks*, 2002 NCTA Technical Papers, May, 2002.

13) White, Gerry, *Tiered Data Services to Enhance Customer Value and Revenue: Why are DOCSIS 1.1 and Advanced Queuing and Scheduling Required*, 2003 NCTA Technical Papers, June, 2003.

14) MultiChannel News, Broadband Data Book