ON PLANT RENEWAL STRATEGIES

Dan Pike
Senior Vice President
Science and Technology
Prime Cable
One American Center
600 Congress Avenue
Austin, TX 78701

Tony E. Werner
Executive Vice President
Engineering and Technical Services
TCI Communications, Inc.
Terrace Tower II
5619 DTC Parkway
Englewood, CO 80111-3000

Abstract

Just as many plants were built originally in a short time in the early 1980's, many now face renewal issues. There are many new variables beyond the traditional technological ones to consider in plant renewal strategy, including competitive risk factors and financial risk factors. This work integrates the concerns known to be influential in the process, and suggests ways to balance between them in the overall adjudication of available technology interventions for plant renewals.

This paper focuses on macro trends as well as specific considerations relating to economics, performance and network design. The Appendix includes extensive work relating specifically to node size determinants, based on most recent theory and experience in traffic engineering, noise accumulation, service levels and network availability.

Introduction

The duty of the cable industry to its constituency has always been to provide entertainment to the residential consumer. The form of technology used for this purpose has changed regularly, from the original tall towers and tube type amplifiers, to technologies such as microwave relays, satellites, fiber and computer controlled terminals, but the purpose has always remained the same--to provide more and better.

The satellite platform of the later 70's was a watershed technological platform. It enabled many new services and engendered urban viability not possible with only the previous terrestrial microwave platforms. Throughout the 80's the industry made optimum use of the technological opportunities, and today all the nation's homes are substantially passed by cable service. Now, in the 90's we contemplate the new platform with which to renew the cable facilities to continue to satisfy the fundamental duty of providing more and better information and entertainment.

This paper discusses factors important to the task, and strategies useful for the purpose.

Today's Environment

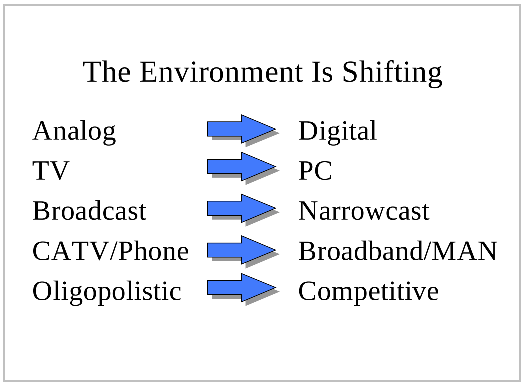

Figure 1

As today's plant renewals are being planned there are some shifts happening in the telecom environment that must be taken into consideration. (*Figure 1*) In the environment of today, traffic moves more to digital from analog form, and information and entertainment previously viewed on a television set is being viewed on a home computer. Broadcasted information is moving more to narrowcasted information. Traditional industries such as the cable industry or the telephone industry are migrating to more of a wide area data structure of a metropolitan area network and are becoming more broadband in nature, and the public policy of the nation today encourages industries that were once isolated by rules or economics to become competitive.

The Impact

These shifts will have an impact as television moves more to a digital format and information becomes indistinguishable between text, imagery, graphics and moving video. (*Figure 2*) The social trend seen with audio will move quickly to video in the form of individual consumption. Market expansion in the form of product and packaging will follow the differentiation and stratification.

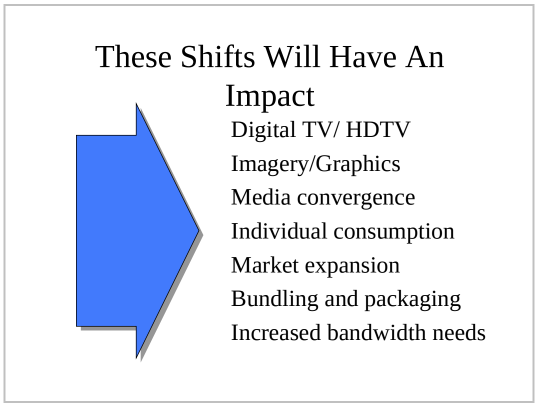

Figure 2

The common denominator of this generally manifests itself in the need for more bandwidth, but provided in ways that do not necessarily have a concomitant cost structure. Today's consumer wants more, but doesn't necessarily want to pay more.

The consumer of today is perhaps best characterized by the 3C's. They desire to be in control of their decision spending, they desire to be able to select among choices for their best value fit, and they desire to bring convenience to a busy life. Selection of plant renewal strategy must be done in consideration of these consumer attributes.

Though it has been discussed for a decade, nearly all would agree that the industry is at the verge of some substantial digital television usage. The implications of that along with the broadcast issues of must-carry and multicasting will have an impact on plant renewal strategies. Up till now, channel bandwidth used for television distribution had serious linearity constraints necessary for transmission of VSB-AM television signals. As more digital service is contemplated, the concern shifts to one of a unity gain bandwidth solution rather than one constrained by linearity. Better silicon transistors with higher F_T products, and better device technologies, such as GaAs, make unity gain easier and cheaper to attain for digital service. [1, 2] Similarly, there is little debate about the rate of increased usage of the internet and as people use the internet with its imagery and graphics in a more interactive form, the demands of the transport facility to achieve that will be more digital in nature.

Mentioned earlier, but expanded here is the concept of individual consumption. In the last twenty years audio devices have gone from one or less per person per household to many, and as video trends follow audio trends, the individual consumption of video services is an attribute that must be considered with new plant initiatives. Rather than one or two it is likely that several devices will be in use simultaneously per household for the asset life contemplated with today's renewal strategies.

Some will require the lowest cost, most ubiquitous distribution through the household, for example the breakfast area television that only requires some news and information sources fed by an additional analog outlet. Others will be more demanding of choices, as for example, the entertainment devices in the media room where pay per view, near video on demand services, and internet response features are enjoyed. All of this has a common denominator requiring more bandwidth, yet as the digital trend moves forward, bandwidth can be thought of as more than one kind, and with more than one cost structure.

Legacy Approach

In the cable industry the previous practice to consider plant upgrades was to decide at what point along one axis one wanted to go. (*Figure 3*) One simply increased the bandwidth of the cable plant as governed by economics or technology. Generally, the performance of the fundamental amplifier component was the limiting consideration, particularly the third-order distortion characteristics.

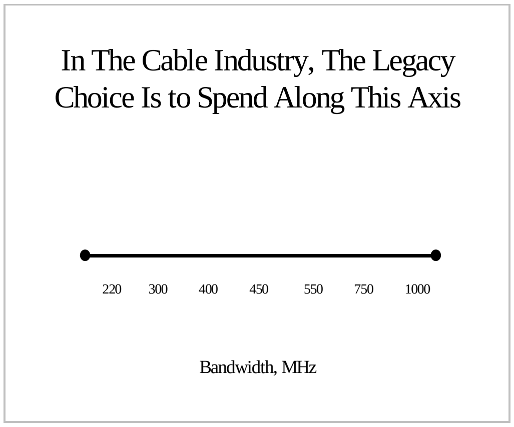

Figure 3

Now There's Another Dimension

Now there is another, digital, dimension to consider. Digital bandwidth may be considered and expressed in megabits per second, and contrasted to analog bandwidth which can be expressed in megahertz. One can equate the two, for purposes of this discussion, though not rigorously, with the aid of a modulation conversion terminal or expression of efficiency(bits/Hz).

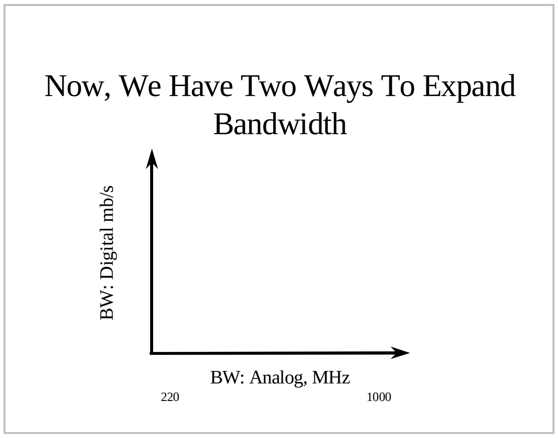

Figure 4

This means there are two axes now to consider in capital spending decisions. (*Figure 4*) We must arbitrate whether to allocate capital for digital bandwidth accomplishment or plant analog bandwidth accomplishment. *Figure 5* is an illustration of the many varieties of ways that capital can be deployed both today and in the future to achieve the situational best answer for any given set of circumstances, such as plant renewal timeframes driven by asset useful life concerns, demographics, financial considerations, and product availability.

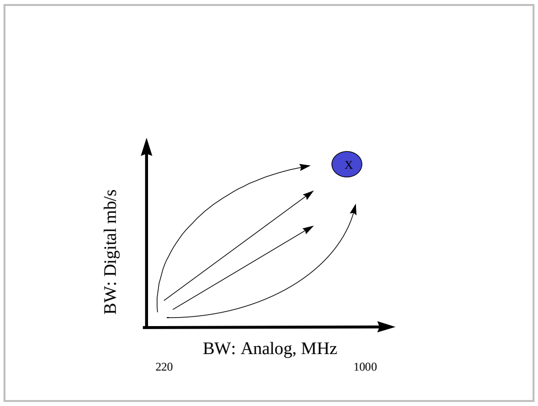

Figure 5

One could take a more digital oriented approach with the purchase of more terminals that achieve digital bandwidth, or one might take a more bandwidth oriented approach to deliver more analog oriented services, or any mixture in between. In any case it all adds up to total capital spent on total bandwidth capacity and each market consideration will determine the best trajectory for that particular location and market.

Risk Elements

Between the historical comfort of known services and plant platforms and the trepidation of new technology and services and the uncertainty it brings lies the concept of decision making under risk. This is today's case far more than in the original decisions for the cable plant platform. Compensation for imperfect knowledge may take the form of safety factors, overcapacity, or capacity built too early.

Figures 6 and 7 illustrate risk factors necessary to consider in upgrade strategy. Risk of too little and too much capital is illustrated as well as the combined risk levels of service acceptance and technological life cycles
.

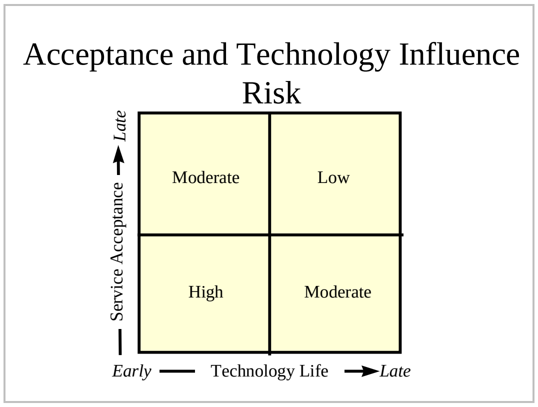

Figure 6

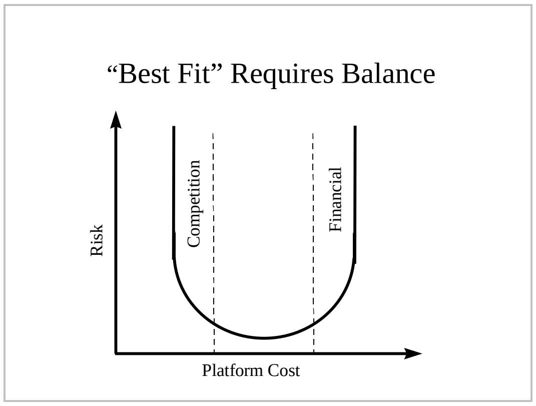

Figure 7

Conceptually illustrated here is the risk tolerance to be observed and realized in plant renewal strategy. Too much capital spending produces risk of financial nonviability; too little implies risk from a noncompetitive service offering. A broad "sweet spot" of lowest risk is the desired operating region but many different approaches are possible within this region. A possible strategy of this risk region is a digital intervention, making early use of digital transport structure enabled by digital terminals, with options of digital-grade additional bandwidth.

<u>Digital Encoding Advantages</u>

Figure 8 illustrates the resultant efficiency from digital modulation, or encoding, and transport, which as stated earlier can be considered casually as the numerical bridge between analog bandwidth in megahertz and digital bandwidth and megabits per second.

Judicious use of digital technology can provide several methods to mine the embedded value from existing coaxial networks. Thorough testing across several systems reveals incremental bandwidth that can be obtained for the fraction of the cost required for a full 750 MHz. This combined with the capacity offered from digital technology provides for robust service offering while reducing cost burdens.

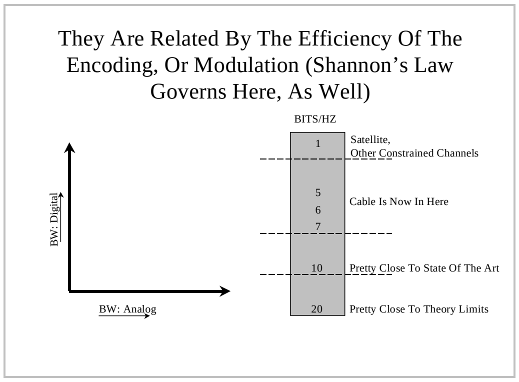

Figure 8

The efficiency of the encoding is treated in the body of science called information theory[3] which began with the publication of Claude Shannon's work about the time the transistor was invented. This very lengthy body of work describes how information can be encoded and transported in any communications channel and discusses the relative theoretical efficiencies. It is not easily summarized[4,5] into a single simple equation, but a useful concept for this discussion can be drawn to express the relationship of digital capacity (C) to bandwidth and channel quality as

$$\boxed{C = Bw\ \mathrm{Log}_2\,(1+{}^{S}/_{N})}$$

The satisfactory transport of the traditional VSB television signal in recent years has required a pristine and very linear channel for consumer satisfaction as described in previous investigations.[6, 7] Typical urban-grade carrier to noise ratio (typically) in the high 40's and distortion requirements in the low to mid 50's leave a very linear, very noise free channel for conversion to digital bandwidth. The total conversion range that is possible, considering the Shannon capacity limits plus the efficacy of concatenated coding[8], can be expressed as a linear range of bits per hertz multipliers beginning at about one and ending at about twenty. Our industry's linearity advantage over propagated or more channel constrained narrowband services means that the cable industry can operate in the five to ten b/Hz range while satellite and other generally more noisy channels must operate below this area. Above this area is another doubling of payload capacity before the theory limit is matched in practice as digital technology and software algorithms improve. This would be analogous to the Moore's Law progression of computer chips and the increasing speed of computer modems of recent years. Thus one can expect cost effective advantages along the digital axis brought on by fundamental Silicon Valley improvements while on the analog axis the linearity of amplifiers remains the governing criteria.

Cable vs Competitive Capacity

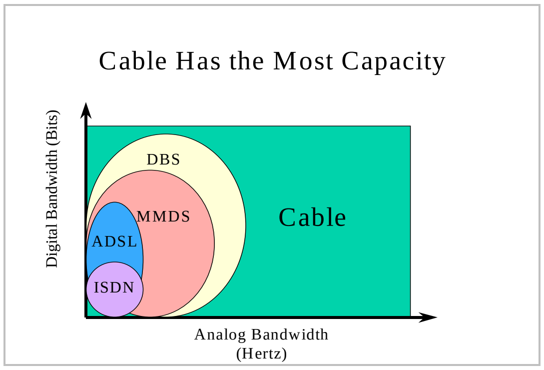

Figure 9

Many of the industry's broadband competitors have made early use of the digital axis available to them. The purpose of the conceptual illustration in *Figure 9* is to show that the cable industry is more advantaged by its intrinsic linearity for both the digital axis bandwidth conversion and the traditional analog bandwidth usage. The cable industry can go further and do better.

Best Fit Engineering

The "Best Fit" engineering approach is used to optimize the balance between network capabilities/performance, market requirements and cost. The process involves a thorough market evaluation, detailed system diagnostics and network modeling. Once completed the options are catalogued so that business and strategic decisions can be made.

This process provides several benefits. The first as alluded to above is an optimized cost structure. The second is a well defined scope of work that provides for an accurate budgeting process. The third is assurance that bad components get identified and replaced during the plant renewal and that good components do not.

Drivers on the market requirement side include current and future must-carry services, competitive offerings, personal computer penetrations and market expansion potential for new services.

There is embedded value in most of the existing networks. The objective of this engineering approach is to maximize this embedded value. Most operators for example re-use the majority of the existing coaxial cable during network renewals. The cable that is added (typically 5-20%) is to make the new design more efficient, not because the existing cable is defective or incapable of the new bandwidth.

There are several additional leverage points that can provide significant value above and beyond the coaxial cable. These largely result from the differential specification performance of passive devices and the initial design. When these are combined with new technologies including fiber optics, digital video and advanced hybrid designs, the result is often dramatic as the following cases illustrate.

Case Study 1

This case studies an existing 300 MHz system without fiber technology in place.

After significant testing of the embedded cable, passive devices and active devices it was determined that the cable provided no practical limitation to bandwidth, that the taps limited bandwidth to just past 450 MHz, the splitters limited the bandwidth to 400 MHz and the amplifiers limited the bandwidth to 300 MHz.

Distortions were analyzed based upon several different scenarios for loading and cascade length. Based upon these findings a network design was developed that installed fiber to limit the cascade lengths and a digital loading of a minimum of 50 MHz was assumed. All of the splitters and couplers will be replaced as will the amplifiers with 750 MHz equipment. The fiber architecture was designed so that if required to later move to 750 MHz, this portion of the network would not have to be redesigned. The network included two way service activation. The costs for this case, in $/home passed, are outlined below and contrasted to a 750 MHz industry standard renewal.

	450 MHz	**750 MHz**	**Diff**
Fiber Cost	$31.00	$31.00	$0.00
Coax	$33.00	$212.00	$179.00
Upstream	$18.00	$28.00	$10.00
Total	$82.00	$271.00	$189.00

Case Study 2

This case studies an existing 450 MHz system without fiber technology in place.

After significant testing of the embedded cable, passive devices and active devices it was determined that the cable provided no practical limitation to bandwidth, that the amplifiers, taps and splitters limited bandwidth to just past 450 MHz. The spacing and design of the system, however, allowed for 550 MHz operation with 50 MHz of digital loading with all of the devices remaining in their exact same location. This meant that amplifier modules and tap plates could be simply changed out which gave least time to completion and least customer disruption. Again fiber was installed to the same architecture as required for a 750 MHz system and two way was activated from day one.

The cost comparison is similarly outlined as follows:

	550 MHz	750 MHz	Diff
Fiber Cost	$39.00	$39.00	$0.00
Coax	$42.00	$124.00	$82.00
Upstream	$22.00	$26.00	$4.00
Total	$103.00	$189.00	$86.00

As stated earlier, the public policy of the nation is to encourage competition among broadband providers, thus any capital spending decision must be considered with all the salient inputs of that environment. There are several ways to do this, but one of the more common is the Internal Rate of Return, which is sensitive to the expected sales success enabled by the capital platform. Using like assumptions, and a Hurwicz criterion[9] of 0.5, three cable systems were considered for three renewal strategies, with cost and IRR information detailed in Table 1. Careful study will show that each case is optimized differently, where IRR is considered as a significant indicator of competitive comparison. Moreover, the standard deviation of renewal efforts is large. Significant bandwidth improvements can be found for as low as $7/home passed to values between one and two orders of magnitude beyond. Illustrated in Table 2 are recent results from 46 projects studied. Clearly no single simple solution platform comes of this disparity, which is why the Best Fit solution is advocated.

Table 1

		Option 1 450 MHz	Option 2 550 MHz	Option 3 750 MHz
System A	Cost Per Passing	$131	$183	$233
	IRR	27%	26.6%	20.1%
System B	Cost Per Passing	$132	$152	$171
	IRR	28%	36%	32%
System C	Cost Per Passing	$151	$167	$184
	IRR	19%	25.7%	23.7%

Table 2

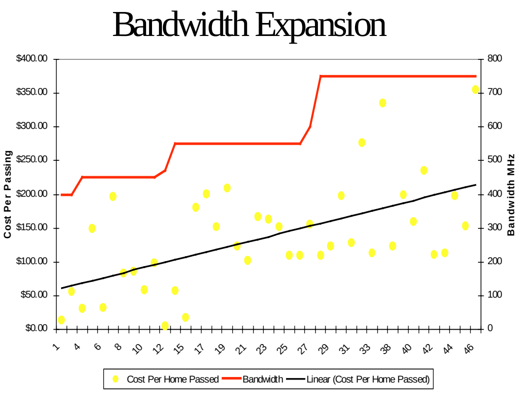

These studies identify situations where a reduced bandwidth approach met the capacity requirements of the local markets and provided substantial savings in implementation time and in cost.

This is not always the case though. Sometimes there is only a modest savings that can be obtained because of the existing network, density or other factors. Or in other cases there are revenue potentials that can only be realized through the additional analog bandwidth.

Figure 10 provides a histogram of selected architectures by bandwidth. These were selected based upon internal rate of returns. It should be noted that the 750 MHz architectures rarely had a more favorable rate of return, but they were close enough to the network with the best IRR, that the marginal difference was overlooked for the extra capacity.

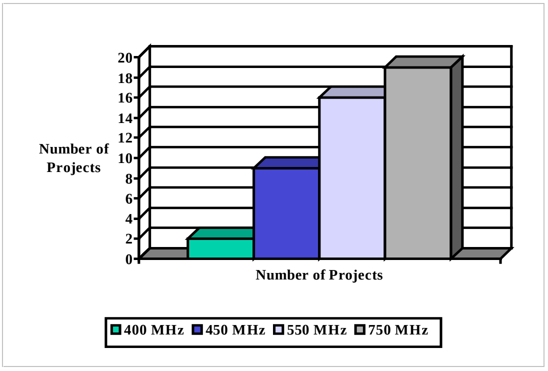

Figure 10

OPTIMAL NODE SIZE STRATEGY

The node size employed in these projects ranged from 600 homes passed to 1500 homes passed per node. Again it was determined by economics and the specific layout of the particular network including housing density.

Discussion of node size requires in-depth examination of trafficking, noise accumulation, and other assumptions. Due to its considerable length this discussion is presented in Appendix 1.0. Latest theory and experiential data is used.

The Appendix provides a substantial review of node size calculations. *Figure 11* provides an overview of the bandwidth required for various node sizes, even when modem penetration reaches 100%.

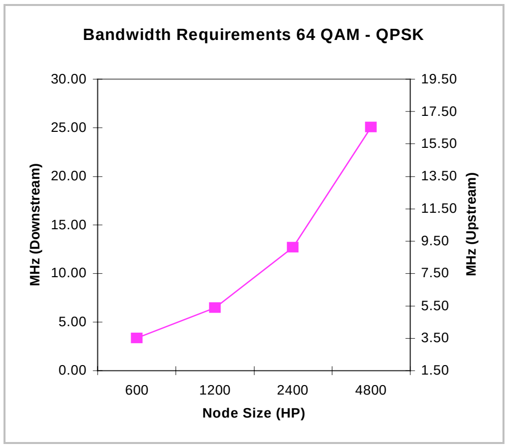

Figure 11

Other graphs are included in the Appendix report that identify the requirements when 256 QAM/16 QAM are employed in the downstream and upstream respectively.

The Appendix report concludes that HFC networks have significant capacity for data communications, so much capacity that even large nodes can support a model based upon 100% modem penetration. This, combined with technology advances in Dense Wave Division Multiplexing (DWDM) and Dense Frequency Division Multiplexing (DFDM) technologies provide for almost unlimited bandwidth capabilities.

Neither noise contributions nor service level performance limit node size designs. High performance can be achieved even with large nodes. Cascade length should be a consideration and must be defined based upon the long loop AGC capabilities of the modem technology combined with the thermal characteristics of the network.

There is no imperative from traffic, performance, maintenance or economics to design nodes below 600 homes passed. In fact, in many cases there are several arguments to design nodes to larger sizes especially in high density areas where reasonable cascade lengths are still being maintained.

Allocating two fibers for every 600 homes passed is more than ample to support the residential business applications. Because of advancements in DWDM technology, these counts may even be reduced in the future.

Conclusion

This paper integrates a range of strategies into a renewal calculus useful to achieve both digital and analog bandwidth additions to existing plants considering the latest theoretical traffic models and empirical evidence, and today's competitive environment. General conclusions supported in the discussions are that Best Fit Analysis produces the best IRR performance, node size is not a significant design variable in plant renewal, and traffic modeling for broadband facilities is highly nonlinear and not yet well understood.

APPENDIX 1.0

Considerations For Determining Optimal Node Sizes in HFC Networks

Node Size Considerations

Communication companies have been deploying Hybrid Fiber Coax (HFC) networks for the last ten years. Construction of these networks started to gain momentum in the early 90's and today, they are the network of choice for video and multi-media communications in the local loop. These networks have been designed and constructed with node sizes varying from 200 to 5,000 homes passed. Until recently, the industry has not had solid theoretical models or empirical data required to determine optimum node sizes. This paper integrates the current theoretical information that is available with as empirical data that has been derived from CableLabs experiments and actual field experience.

Factors that must be considered to determine optimum node sizes include:

- ☐ Services
- ☐ Traffic Engineering
  - ✓ Today's Requirements
  - ✓ Future Requirements
- ☐ Noise Accumulation
- ☐ Service Levels
- ☐ Technological Trends
- ☐ Maintenance Considerations

1.0 Services

Network designs in the late 1980's and early 1990's anticipated that a wide array of technologies and payloads would be utilized to provide advanced communications services over an HFC network. This included different protocols as well as different modulation and multiplex schemes for applications such as telephony, high-speed data, teleconferencing and games. It was envisioned that these different services would be frequency division multiplexed onto the HFC network.

Due of the broad acceptance of the IP protocol, several of these services can now be provided over a single, common, time division multiplexed network employing the IP protocol. This results in several advantages, including;

reduced CPE and office terminal equipment, ubiquity through national IP networks, reduced requirements for additional gateways, significantly more efficient use of bandwidth and a consolidated traffic model.

While most of the envisioned residential services have stayed the same, the method of providing these services has converged to either the cable modem or the OpenCable[10] set top device.

2.0 Traffic Engineering

This section focuses on IP traffic modeling for today's services and extrapolates the bandwidth requirements for emerging and future services. While certain discussions will relate to the global Internet backbone, the predominant focus of this paper is on the HFC portion of this architecture.

Simple mathematics do not apply to shared access mediums. If they did, 10 Mbps Ethernet LANs with 50 to 100 users would perform only marginally faster than dial-up modems. Yet Web pages representing several million bits of information are retrieved and displayed in a few seconds. The reason for this lies in the bursty nature of the data and in the manner that traffic aggregates.

IP Traffic engineering is influenced by several factors, including user behavior patterns, applications, home platforms and quality of service requirements. Residential traffic today, consists mainly of news groups and web surfing. Electronic mail, file transfers and Telnet are popular applications, but they do not account for much traffic. Traffic types will change though, as high-speed networks become more ubiquitous, promoting new services.

Several challenges exist in predicting IP based traffic over the HFC network. These challenges result from the HFC network providing an extension to the global Internet. The Internet itself is largely heterogeneous in both applications and behaviors. While the IP architecture unifies many different network technologies, it does not unify behavior in network administration or by the end user.

Adding to the challenge is the fact that the Internet is big. It was most recently estimated at more than 16 million computers (January 1997, Lottor-97). This large scale means that rare events now occur routinely in certain parts of the network.

Change is another challenge in determining Internet traffic. Change in size, with the Internet nearly doubling on an annual basis. Change in traffic characteristics including median packet length and protocol mix. But, the largest change that pales all others is the change in applications and use.

This is a moving target that is fueled by several factors. The following lists parameters that can significantly affect traffic characteristics and are causing continual change:

1. Continuous increases in processor and communication speeds;
2. Pricing structure;
3. Equitable distribution of resources instead of FCFS (first-come-first-served) scheduling algorithms (fair queuing);

4. Native (source) multicasting over the network (instead of predominantly unicasting;
5. QoS (quality of service) protocols for different classes of traffic (a must for IP telephony) and other services requiring better than "best-effort" performance;
6. Local Traffic caching, and
7. New killer applications

While change has to be accounted for in any forward looking traffic model, we have to start with a basis for how we aggregate traffic from various sources. Today there are two common methods for modeling traffic arrivals. The first is using traditional Poisson processes. The second is based upon statistically self similar modeling.

Both techniques have their appropriate application, largely based upon traffic type.

2.1 Network Characteristics

You must first analyze the characteristics of the network itself. This includes characteristics of the users, the protocols and the network medium.

2.1.1 Users

Traffic may be generated by the aggregate of several small connections (Web "mice"). Or it may be dominated by a few extremely large, one way, rate adapting bulk transfers ("elephants") or long-running, high volume data streams containing audio and video that are "multicasted" from one sender to multiple destinations. It can also consist of bidirectional multimedia traffic generated by interactive gaming.

2.1.2 Protocols

The IP (fundamental Internet protocol) is responsible for routing packets through the network. Other protocols such as TCP - transmission control protocol are built on top of the IP layer and based upon implementation differ significantly. These protocols define how to divide streams of data into packets such that the original data can be delivered to the receiver. Certain protocols will deliver this data even if some of the packets are dropped or damaged.

On top of these two layers there are application specific protocols providing such network services as email (SMTP), WWW (HTTP), file transfer (FTP), and others. Traffic modeling and simulation must take into account interactions between all the protocols in a stack.

2.1.3 Network Medium

Both the network (comprised of several, differing, networks) and the services continue to evolve, adding to the complexity and dynamics. The IP protocol allows for uniform connectivity but not uniform behavior. Moreover, it interconnects millions of users (computers) and even if some of them behave atypically, the Internet can still include thousands of these atypical users. The phenomenon, of course, changes over time from the growth in the number of users and from the growth in service types.

Additionally, new technologies and topologies add to the growing heterogeneity of the network. An example is satellite links (geostationary satellites with constant, large latency or LEO satellites with continually varying latency). New protocols will emerge for these technologies to compensate for the issues resulting from their latency and other disadvantages, this again will change the network characteristics. Other examples include fast modems such as ISDN, XDSL and ultra fast modems such as DOCSIS modems connected over HFC networks.

2.2 Traffic Characteristics

Traditionally, network arrivals have been modeled as Poisson processes. If this model is applied consistently to all network processes, it would lead to the conclusion that packet interarrivals are exponentially distributed. However, traffic studies showed this not to be true. Especially, in LANs and WANs, the distribution of packet interarrivals clearly differs from the exponential (no natural length for a burst, traffic bursts appear on a wide range of time scales).

The network can be a perfect aggregate of traffic for different applications or can be dominated by one or a couple of applications. The following are the most popular protocols and applications:

- TELNET: interactive client server communications traffic,
- FTP: file transfer traffic,
- SMTP: email is a small machine generated and sometimes timer driven, bulk transfers of data,
- NNTP: network news transport is a small machine generated and sometimes timer driven bulk transfers of data,
- HTTP: World Wide Web traffic,
- IP Telephony

There are many other processes (RLOGIN and X11, for example) that also contribute to the network traffic.

2.2.1 Aggregation Levels

The traffic will show different characteristics at different levels of aggregations. The empirical data indicate that some models used for traffic on LANs (for example self-similar process models) are not applicable to WANs. It is not known precisely what level of aggregation can be applied to the traffic on HFC networks with different levels of segmentation.

2.2.2 Distribution of Traffic Processes

2.2.2.1 TELNET

TELNET connection arrivals within one hour intervals are well described by Poisson model with fixed hourly rates. This is valid after neutralizing the dominating 24-hour pattern. Each arrival represents an individual user starting a new session.

Packet arrivals are well described by an empirical, heavy-tailed Tcplib process. This process preserves the burstiness of the packet arrivals over many time scales. This distribution is primarily network invariant and is determined by human typing characteristics. Packet inter-arrivals are far from an exponential distribution and are more accurately described by a Pareto distribution with infinite mean and variance.

Connection sizes in bytes have log-extreme distribution.

Connection sizes in packets have log-normal distribution.

2.2.2.2 FTP

User generated FTP session arrivals within one hour intervals are well described by a homogeneous Poisson model with fixed hourly rates. This is valid after neutralizing the dominating 24-hour pattern. Each arrival represents an individual user starting a new session.

Data connection arrivals, on some aggregation levels, behave as self-similar processes. These are initiated within a session, whenever the user lists a directory or transfers a file. This traffic has a long-range dependency and comes in bursts.

Byte-number distribution in every burst has a heavy upper tail. A small fraction of the largest bursts carry almost all FTP bytes during data transfer. Therefore, all that matters in analyzing the network behavior is the behavior of a few large bursts. A 5% tail for these bursts fits well into a Pareto distribution. Arrivals of the upper tail bursts are not well described by a Poisson model. Generally speaking, the FTP connection arrival processes show large-scale correlation indicating long-range dependency for data connections.

2.2.2.3 SMTP and NNTP

Connection arrivals for NNTP are not well modeled using Poisson processes. This is explained by the periodicity of machine generated IP traffic that can result in traffic synchronization which is not characterized by the Poisson model.

The SMTP connection arrivals within 10 minute periods, are well described by the Poisson model.

2.2.2.4 HTTP

Connection arrivals are decidedly not well modeled with Poisson processes.

2.2.2.5 IP TELEPHONY

This traffic will most likely follow Poisson models for telecommunication traffic.

2.3 Traffic Simulation

Traffic simulation is important, as testing with collected traffic traces does not take into account the adaptive congestion control used by the network and the sources attached. Traffic simulations and protocol simulations also allow us to test the network under different scenarios. Moreover, simulations protect the network from an unintentional overload of the network by a very successful application (so-called success disaster).

Simulation processes must be based on network invariance and judicious exploration of the parameter space. Invariance include long-term correlation's in the packet arrivals in aggregated (LAN) traffic. Internet traffic is well described in terms of "self-similar" (fractal) processes instead of traditional approach (Poisson or Markovian modeling). Long term here is defined as hundreds of milliseconds to tens of minutes. Short term behavior is affected by network transport protocols which affect short-term correlations, and longer term is affected by non-stationary effects such as diurnal patterns.

Network users session arrivals are well described using Poisson processes. Examples are remote logins and initiations of a file transfer (FTP) dialog. Unlike the packet arrivals, session arrivals are at a much higher level (exchange of hundreds of packets). This is true after accounting for diurnal patterns in session arrivals. Individual network connections within a session are not well described by Poisson processes.

Connection sizes or duration can be described by log-normal distribution with mean and variance based on empirical measurements. These, however, are not generally useful due to variations in connection characteristics from site to site.

Network activity can be characterized by a distribution with heavy tails (for example, Pareto distribution with shape parameter $\alpha<2$ with infinite variance). This conclusion is supported by such examples as Unix processes (CPU time consumed), sizes of Unix files, compressed video frames, WWW items, bursts of Ethernet and FTP activity.

The pattern of network packets generated by a user typing (TELNET) has an invariant distribution with a Pareto upper tail and a Pareto body.

Parameter space must identify all parameters that have to be exercised during simulation. This includes selecting one and varying it over several orders of magnitude and in very small steps, to account for non-linear feedback, during the simulation process. Additionally, one must exercise all parameters to determine network sensitivity to each of them independently.

2.4 CableLabs Experiment

In 1996 CableLabs conducted research into high-speed data transmission over HFC networks and concluded the following:

- Large node sizes (~2,000 homes passed) work well for data services;
- Nodes can be subdivided in the frequency domain long before a physical reconfiguration is necessary, and
- That the quality of service does not degrade noticeably as up to several hundred users are added

CableLabs designed a high-speed surfing experiment to provide an understanding of WWW impacts on HFC based, shared mediums. In a controlled laboratory setting containing both the PC performing the surf and the server holding the requested pages, the engineers performed a graphically intense web surf. During this surf, the engineers analyzed the data transmissions over the broadband pipe that connected the machines. *Figure 1* provides a representative, single user web surf derived from this work.

Figure 1

Binned Throughput (Mbps)

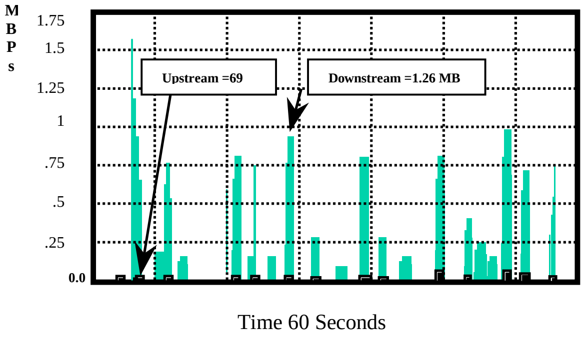

The solid bars indicate downstream data bursts and the dashed lines indicate upstream data. In a surf, this is usually a request or a search that creates the upstream data. The downstream burst of data is a web page being downloaded to the client. In this case, the web pages were graphicaly rich and included significant spacial and temporal information. As can be seen in *Figure 2*, bursts hit in the range of 1.6 Mb/s in the downstream and less than 50 Kb/s in the upstream. The total data downloaded over a 60 second period was 1.26 Mbytes. Concurrently the upstream data over this 60 second period was 69 Kbytes providing for an 18:1 assymetry.

Based upon this data and the MAC layer behavior of the MCNS protocol, CableLabs created a predicted target architecture performance based upon the previous WWW surf. This result was based upon a sophisticated computer model and then validated with real modems. It is illustrated in *Figure 2*.

Figure 2

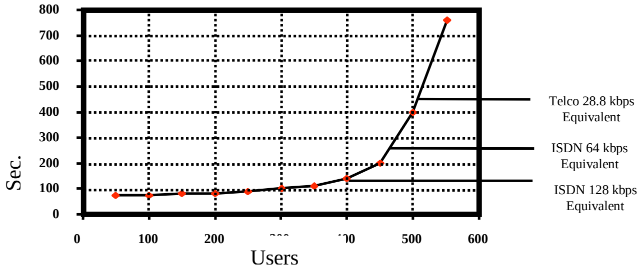

Based upon *Figure 2*, the knee in the performance curve starts at approximately 400 simultaneous users on a single MCNS modem channel.

It should be noted that this experiment identifies the possible upper limit of Web surfing, not typical behavior. The pages were downloaded faster than what could be absorbed by the user.

2.5 Empirical Data

While a number of comprehensive studies have been published over the years, there are relatively few recent studies to on internet traffic and characteristics. Some of the early traffic studies include CBP93 which investigated detailed NNstat and SNMP NSFNET data in the summer of 1992. At that time, application usage was reported as being dominated by FTP, SNMP, NNTP, DNS, Telnet and a growing amount of “other” traffic.

Merit’s final NSFNET report summarized statistics on the backbone from 1988 to 1995 where it was noted that information retrieval services such as Gopher and Web were beginning to overtake mail and file transfers in their share of network traffic. Traffic make up in packet counts in the spring of 1995 are outlined in Table 1.

Table 1

Other	27%
WWW	21%
FTP-data	14%
NNTP	8%
Telnet	8%
SMTP	6%
IP	6%
Domain Name Server	6%
IRC	2%
Gopher	2%

This distribution is based upon packet count and would change significantly based upon byte count because of the distribution of packet sizes.

Paxson published a comprehensive report in 1997 that analyzed 20,000 TCP transfers among 35 sites and over 1,000 internet paths. His findings examine such areas as route pathologies, loss characteristics, specific TCP implementation behaviors and packet queuing and delay differences. His research is the basis of several of the preceding sections that describe the nature of different traffic sources.

2.5.1 MCI (vBNS)

In the fall of 1997, Thompson – Miller - Wilder published results from monitoring the MCI very-high speed Backbone Network Service (vBNS) for the National Science Foundation (NSF).

The MCI paper provides daily and weekly traffic patterns and composition for two traffic points on the commercial internet backbone. In this section we will highlight some of the key findings from point one, which was a domestic traffic point versus point two which was an international traffic point. Point one was within a node that serves as a junction point for several backbone trunks as well as an access point for local customer traffic near a major U.S. East Coast city.

Figure 3 below is an approximate recreation of the byte volume for a one day record on the domestic link. Note that there are some differences in communication symmetry, which also changes over time.

Figure 3

Byte Count for 1 day on domestic link

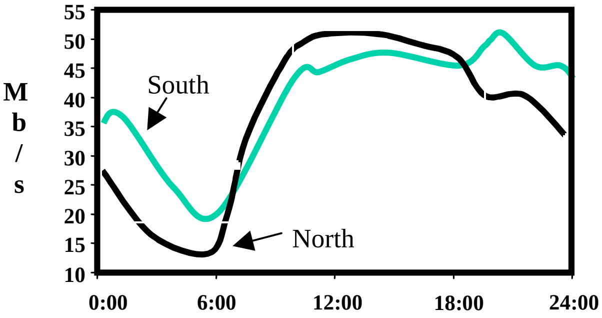

These curves show consistency with other patterns whereby the traffic volume nearly triples from 6:00 am to noon and then holds a moderately steady pattern until near midnight when it starts to trail off.

Weekly traffic patterns indicated that weekday traffic was much heavier than weekend traffic, but the daily patterns were similar, especially when time zone impacts are considered. An HFC network will differ as it will be much more asymmetrical as it relates to byte counts and it will likely not diminish as much during the weekends.

Again the MCI study showed significant modal distribution of packet sizes based upon protocols. Over 50% of the packets were smaller than 45 bytes and nearly all were smaller than 1500 bytes. This modality again brings into question the value of average packet lengths when discussing traffic characteristics. This conclusion may be different for HFC networks, where by definition they will have packet sizes that will range from a minimum of 64 bytes to a maximum of 1500 bytes.

Traffic composition for the MCI study revealed that for IP protocols, TCP was the predominant protocol accounting for over 75% of the flows, 90% of the packets and over 95% of the bytes. UDP accounted for the majority of the remainder accounting for 20% of the flows, 10% of the packets and 5% of the bytes. The remainder was accounted for by emerging protocols, such as IPv6, encapsulated IP, ICMP and others.

Within TCP, the most predominant application is WWW, accounting for 75% of the bytes. FTP accounted for 5% of the bytes, NNTP accounted for 2% of the bytes and Telnet accounted for 1% of the bytes. The remaining traffic is distributed over several other types of traffic. UDP traffic mainly consisted of DNS and RealAudio, RealVideo apps.

2.5.2 Rogers Engineering

Rogers Engineering monitored HFC and ISP gateway utilization during the Spring and Summer of 1996. During this time, they had 471 early adopter data customers connected to Zenith high speed modems over an HFC network. This network was connected to the internet via a T-1 connection and during this time traffic caching was not performed. The T-1 utilization is graphed in *Figure 4* for a typical day out of the analysis. You will notice that even with 471 customers that the peak aggregate bit rate never exceeded 900 kb/s. The pattern is, however, quite similar to that identified in the backbone study by MCI where by dramatic traffic increases occurred between 6:00 AM and Noon. The bottom trace indicates the outgoing data. Asymmetry in byte count, increases as traffic increases, largely based upon the traffic types. It starts out as low as two to one and grows to as high as ten to one during the peak usage. The CableLabs WWW experiment showed that this asymmetry could grow as high as eighteen to one during exaggerated peak usage, especially if it was based upon pure WWW traffic.

Figure 4

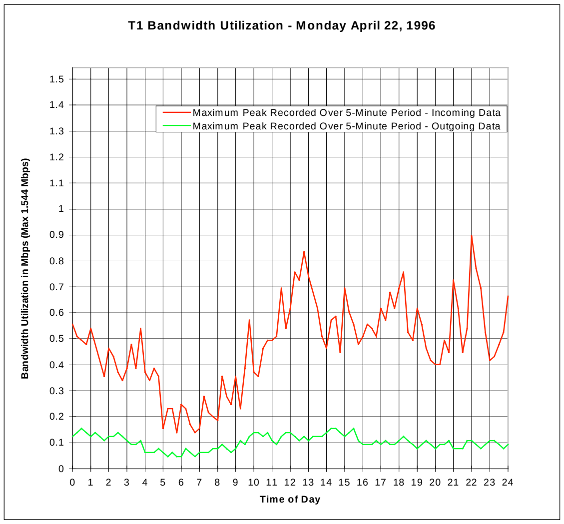

Active users are plotted for a typical day in *Figure 5*. In it, you will notice that the curve differs from byte count graphs shown previously. This is largely accounted for by users that leave their computers on all evening, even though they are not being actively used.

This graph depicts the highest simultaneously active hosts from the study, whereby 194 of the 471 (41%) hosts were active. This peak occurs at 9:30 PM and represents a growth from 5:00 AM when 19% were active. It must be noted that active terminals do not represent simultaneous use as it is common practice to have active, but idle computer terminals. This pattern and its growth supports the typical industry standard for a 25% peak simultaneous usage.

Figure 5

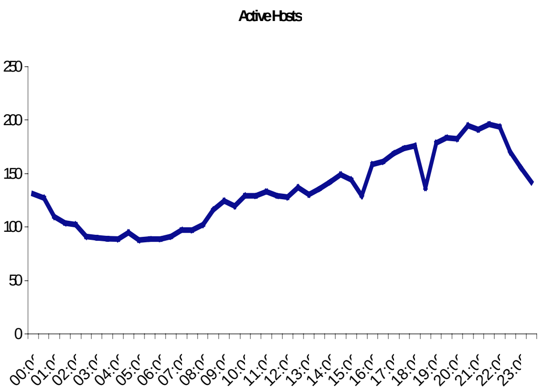

Traffic composition based upon the Rogers experience indicated that WWW (HTTP) and News (NNTP) accounted for 98% of the traffic volume in bytes. News (NNTP) traffic accounted for a higher percentage of the HFC traffic than it did on the backbone. This became especially true as the operation implemented local caching of the News group data. This combined with the early adopter profiles is the predominant reason for this characteristic.

User profiles indicate that the majority of the users are similar and in the case of Rogers, 62% of the customers consumed less than 25 MB of data over a one month period. The remaining heavy users form a fairly long tail. This characteristic becomes even more predictable as larger numbers of users are present and the early adopter impact has less of an effect on the total usage.

Figure 6

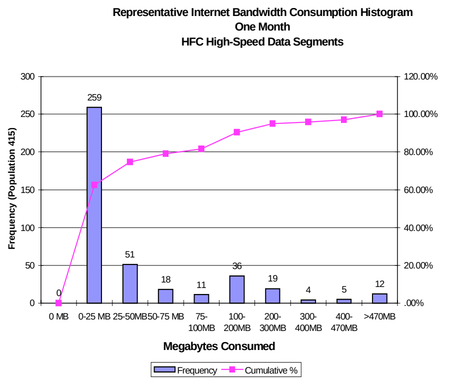

2.6 HFC Capacity

The preceding sections have illustrated that it is nearly impossible to develop a traffic model that can account for all of the dynamics that will affect traffic and user behavior as we move into the future. One can, however, model an upper bound that will demonstrate the ample capacity of HFC networks even with relatively large node sizes. Operators have typically taken a significantly conservative approach up until this time in traffic provisioning. This has resulted in many under utilized ISP connections and the practice of node combining in the HFC networks themselves.

As can be seen from the previous studies, even though peak burst rates are often high, the aggregation of this traffic is quite modest. This is what allows the backbone to operate with relatively small circuits and the 471 customers in the Rogers early experience to only require 900 kb/s.

The remainder of this section demonstrates the significant capacity of HFC networks for two-way data-communications. This section demonstrates that even large nodes can support 100% modem penetration and incredibly high consumption and peak characteristics while allocating only a modest portion of the up and downstream bandwidth.

The following provides assumptions based upon a 100% modem penetration as a result of the OpenCable initiative. It should be noted that while there may be 100% modem penetration, only a percentage will be attached to stand alone computers. The remainder of the modems will provide communications to the OpenCable terminal, which will operate as a network computer. This device will have a powerful processor, but it will not contain substantial local storage or rapid input mechanisms. Because of this architecture, the upstream communications will mainly consist of short packets used to request information. It will also serve a much different need and application than that of the personal computer.

Table 2 provides assumptions baselines for the calculations. We start with the assumption that 100% of the cable households have a high-speed modem installed and operational. We use today's penetration of 62% for cable penetration. We base PC penetration at 50% which is much higher than average penetrations today.

We have divided the users into power users and casual users. Power users would be those who spend considerable time on the PC and who have the latest generation processor, etc. The baseline uses a 40%-60% split respectively. In the Rogers experience, 62% of the customers fell into the bin of 0 to 25 MB per month. From there, the usage distribution had a long tail, with 80% of the users consuming less than 75 MB per month.

For simultaneous peak usage, we have used 25%. This number is based upon today's low penetration levels and early adopter make up. This number should actually go down as penetration increases. With the move to OpenCable types of services, TV viewing habits will drive a substantial amount of the usage. Typically the maximum TV on rate is in the high 60% range. This usage is split based upon market share, which today, usually breaks this down into small segments. Major interest items will typically use a multi-cast approach and also reduce peak data requirements.

Telephone penetration is calculated at a 30% penetration and a 3 CCS traffic load. This is the average residential usage today. This is likely conservative as these homes will typically have two phone lines and high-speed data access through the DOCSIS modem.

IP video telephony is estimated at a 20% penetration and a 1 CCS utilization. In other words, consumers will use video connections one minute for every three minutes that they are using the standard POTS line. Again this is likely extremely conservative as these services will have a rate differential and they will require that you connect to another customer who has this video capability. Just like wireless communications, customers will likely use wireline POTS if it will satisfy the business or personal need, before they will use the more expensive service.

Table 2

Penetration Baseline

Customer Penetration	62%
PC Penetration of Customers	50%
Power User	40%
Casual User	60%
Simultaneous Peak	25%
OpenCable Penetration	100%
Simultaneous Peak	25%
IP Telephony Penetration	30%
Usage CCS	3
IP Video Telephony Penetration	20%
Usage CCS	1

Table 3 provides estimates for peak average data rates. For WWW, FTP and NNTP applications, we have used the 168 kbps downstream and 8 kbps upstream results that were derived from the CableLabs high-speed surfing experiment. Again this is not burst speed but aggregate speeds that are then averaged. We have used 80 kbps and 1,544 kbps for real-player audio and video respectively. Both of these are typical of today, but likely conservative for the future. IP telephony is based upon 32 kbps for voice and 128 kbps for video. This should likely be conservative as well. Today the majority of the IP voice is compressed to levels 50% to 60% lower than this 32 kbps rate.

Table 3

Peak Rate Baseline Kbps

	Peak Rate	
	DN	UP
Protocol/Application	Stream	Stream
WWW	168	8
FTP	168	8
NNTP	168	8
RealPlayer Audio	80	2
RealPlayer Video	1,544	8
IP Phone	32	32
IP Video Phone	128	128

Tables 4 and 5 provide the capacity of HFC networks based upon 64 and 256 QAM downstream modulation schemes and QPSK and 16 QAM upstream modulation schemes.

Equipment is already being shipped with 64 QAM and QPSK schemes today with 256 QAM and 16 QAM equipment likely available in the summer of 1998.

Table 4

Downstream Capacity Baseline

	64 QAM	256 QAM
Downstream Bandwidth MHz	6	6
64 QAM Gross Data Rate Mbps	30	43
Minus Error Correction Mbps	3	5
Minus MAC and Messaging Mbps	1	1.5
64 QAM Net Data Rate Mbps	26	36.5

Table 5

Upstream Capacity Baseline

	QPSK	16 QAM
Upstream Bandwidth MHz	6.4	6.4
Gross Data Rate Mbps	10.2	20.4
Minus Error Correction Mbps	1.4	2.8
Minus MAC efficiency and OH Mbps	3.7	7.5
Net Data Rate Mbps	5.1	10.1

Table 6 derives the number of users for four different node sizes based upon the preceding baselines. This table initially indicates the total number of customers and then extrapolates this for simultaneous use.

Table 6

User Calculations

Node Size (HP)	600	1200	2400	4800
Customers 62%	372	744	1488	2976
OpenCable Customers	372	744	1488	2976
PC Customers	186	372	744	1488
Power Users	74	149	298	595
Casual Users	112	223	446	893
Simultaneous Active *(not including telephony customers)*				
OpenCable Customers	93	186	372	744
PC Customers	47	93	186	372
Power Users	19	37	74	149
Casual Users	28	56	112	223

Tables 7 and 8 calculate the number of lines required to provide a P.01 grade of service based upon the preceding baselines. This is based upon Poisson addition and then discounted the trunking efficiency for small statistical groups. The trunk efficiency typically would be higher as one moved from the 600 home example up to the 4800 home example. Under this IP example, however, blocking will technically be determined by the bit priority that is assigned.

Table 7-IP Phone

Node Size (HP)	600	1200	2400	4800
IP Phone Customers	112	223	446	893
Total CCS	335	670	1339	2678
Lines Required for P.01	26	46	82	152
Trunk Efficiency	0.9	0.9	0.9	0.9
Lines	29	51	91	169

Table 8-IP Video Phone

Node Size (HP)	600	1200	2400	4800
IP Video Phone Customers	74	149	298	595
Total CCS	37	74	149	298
Lines Required for P.01	8	11	16	26
Low Trunk Efficiency	0.9	0.9	0.9	0.9
Lines	9	12	18	29

Table 9 defines the characteristics of Power users, again based upon the preceding baselines. Notice how conservative the model is. We are now up to an average peak usage of 322 kbps in the downstream. If one reinserted this into the previous data on the Rogers customers, the peak usage of 900 kbps that they recorded would have jumped to 38 Mbps or an increase of 42 fold. A customer of this characteristic would consume as more data in ½ hour than the average Rogers customer did in an entire month. This is an extremely aggressive assumption for any user, and contributes to make this model highly conservative.

Table 9

Power User Characteristics (kbps)

Application	DN Stream	UP Stream
WWW/FTP/NNTP	168	8
Real Player Audio 20%	16	2
Real Player Video 1 - 6 Sec. Clip	154	1
Total	322	10

Tables 10 and 11 define user characteristics for casual PC users and couch surfers. Again these are average peak usage, not burst speeds. Again these extrapolate into extremely aggressive usage patterns.

Table 10

Casual User Characteristics (kbps)

Casual User	DN Stream	UP Stream
WWW/FTP/NNTP	84	4
Real Player Audio 20%	8	0.2
Total	92	4.2

Table 11

OpenCable Characteristics (kbps)

Open Cable	DN Stream	UP Stream
Web Retrieval	42	2

Table 12 now uses these preceding baselines and calculations to determine cumulative data requirements for four different node sizes.

Table 12

Cumulative Bandwidth Requirements (Mbps)

Bandwidth Requirements Mb/s								
Node Size (HP)	600		1200		2400		4800	
Stream Direction	DN	UP	DN	Up	DN	UP	DN	UP
Services								
IP Phone	0.9	0.9	1.6	1.6	2.9	2.9	5.4	5.4
IP Video	1.1	1.1	1.6	1.6	2.3	2.3	3.7	3.7
Personal Computer								
Power User	6.0	0.2	12.0	0.4	23.9	0.7	47.9	1.4
Casual User	2.6	0.1	5.1	0.2	10.3	0.5	20.5	0.9
TV Based Web Services								
Open Cable	3.9	0.2	7.8	0.4	15.6	0.7	31.2	1.5
Total	**14.5**	**2.5**	**28.1**	**4.2**	**55.0**	**7.1**	**108.7**	**13.0**

A 600 home node only requires 14.5 Mbps in the downstream and 2.5 Mbps in the upstream to satisfy 100% modem penetration and aggressive and even perhaps unrealistic data use assumptions.

Figure 7 and 8 graph the RF bandwidth requirements for networks based upon 64 QAM/QPSK and 256 QAM/16 QAM modulation schemes.

Figure 7

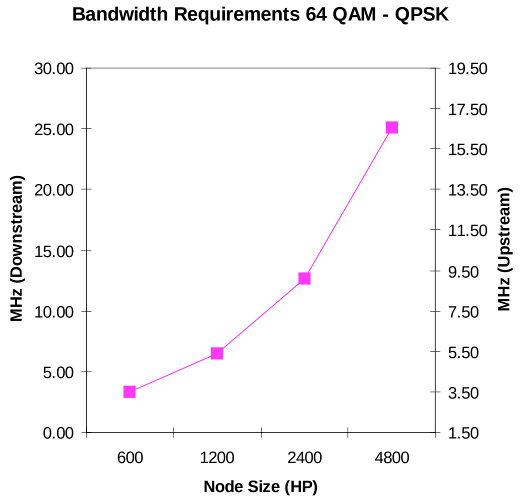

Figure 8

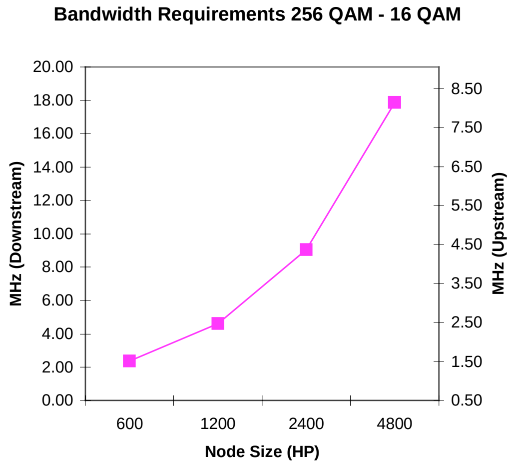

Note that even with QPSK modulation schemes and a 4800 home passed node you can support 100% modem penetration with less than 18 MHz of upstream spectrum.

Figure 9 charts the number of downstream channels that are required to support various sized nodes based upon 256 QAM and 64 QAM modulation schemes.

Figure 9

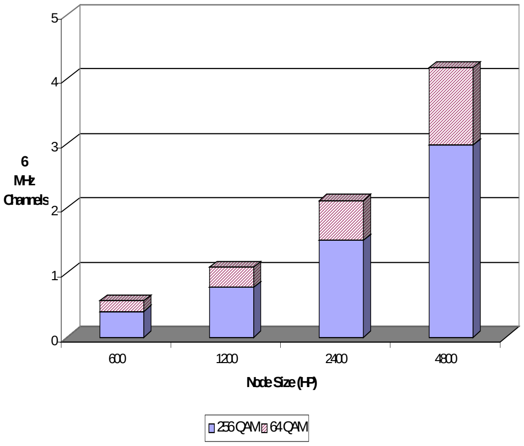

3.0 Noise Accumulation

Larger nodes do have a higher accumulation of noise from active device contribution. The increase, however, is not a material design factor. Let's first look at the active components, their operational levels and noise contribution as it applies to the upstream network.

The design assumptions for an upstream network are listed in Table 13.

Table 13

Laser Type	DFB
Laser Operating Levels	-56.5 dBmV/Hz
Return Amplifier	Hybrid
Amplifier Operating Levels	-54 dBmV/Hz
Laser Equivalent Noise Factor	27.6 dB
Hybrid Noise Figure	4.5

These design guidelines provide for a 15 dB operating to clipping ratio. Because the dynamic environment of the upstream network and the effect of laser clipping on digital signals, an adequate amount of headroom is critical to return path operation. This allows for level changes and impulse noise effects.

Based upon these design parameters, the carrier to noise resulting from the optical network would be:

Table 14

Noise Floor	-125.1 dBmV/Hz
Minus Input Level	-56.5 dBmV/Hz
Plus Equivalent Noise Figure	27.6 dB
Optical Network CNR	41 dB[11]

The carrier to noise ratio resulting from a single hybrid would be as follows:

Table 15

Noise Floor	-125.1 dBmV/Hz
Minus Input Level	-54 dBmV/Hz
Plus Noise Figure	4.5 dB
Single Station CNR	66.6 dB

Figure 10 depicts an HFC network indicating nodes that vary from 500 to 2,000 homes passed.

Figure 10

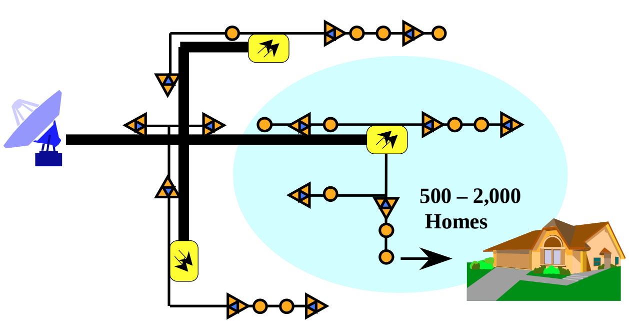

Based upon 5 amplifiers per mile and 100 homes per mile, a 500 home passed node would include 25 amplifiers. The 10xLog$_{10}$ of 25 is 13.97 providing a 52.6 dB CNR for the entire 25 amplifiers. When this is added to the 41 dB CNR from the optical network we arrive at a 40.7 dB CNR for the 500 home node.

A 2,000 home passed node would consist of four times as many amplifiers or 100. The 10x log$_{10}$ of 100 is 20, providing for a 46.6 dB CNR for the entire 100 amplifiers. When this is added to the 41 dB CNR from the optical link, we arrive at 39.9 dB CNR. Notice that quadrupling the node size provides only a 0.8 dB loss of CNR performance. This analysis is based upon the best laser that is practical to deploy in the upstream network today. Several operators are employing lower cost lasers such as cooled FP lasers. Under these situations there would be practically no improvement in CNR performance of a 500 home passed node over that of a 2,000 home passed node.

Because of the cost and space required for cable modem termination (CMTS) equipment, most operators are serving between 2,000 and 5,000 homes passed from a single channel termination. Some operators are doing this from a single optical node, others have actually combined smaller nodes to arrive at this level. Provisioning a single cable modem termination for a 500 home or even 1,000 home passed node offers poor economics at penetration levels below 15%.

Figure 11 depicts the node combining that is performed to group four 500 home nodes into a 2,000 home passed area to satisfy the CMTS economic and space criteria.

Figure 11

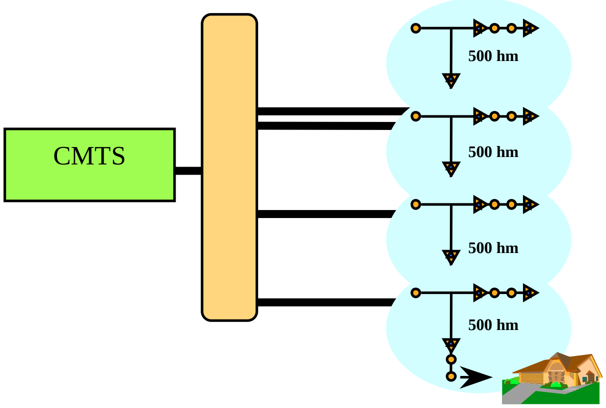

Under this example, the output of four optical receivers is combined at RF and fed into the CMTS. The CNR performance for this example now consists of 100 amplifiers plus four lasers. The four lasers now produce a combined CNR of 41 dB –10Log$_{10}$ of 4 or 35 dB. Note the combined effect of the optical transport and the RF Transport is 34.7 dB CNR. Note that this configuration offers the poorest performance of all of the examples.

Table 16 depicts noise based upon four examples, the three previously illustrated and a 10,000 home passed RF only network.

Table 16
CNR (dB)

	500 HP	2,000 HP	4x500 HP	10,000 HP
RF Network	52.5	46.5	46.5	40
Optical Network	41	41	35	N/A
System	40.7	39.9	34.7	40

Table 16 illustrates the dominant noise contribution from the optical network in upstream applications. This is largely caused by the low optical modulation index requirements for return applications. After laser combining, the four 500 home passed node example provides the poorest noise performance of any of the examples.

4.0 Service Levels

Early on, many network engineers and system designers hypothesized that smaller nodes would yield much higher levels of service availability. This was based upon the lower levels of network contributed noise, the lower level of impulse noise accumulation from home appliances and shorter cascades of RF amplifiers.

Actual implementations, however, have demonstrated that there is minimal if any correlation between node size and service performance of two-way services. *Figure 12* is provided by CableLabs. In it, CableLabs tested and charted the performance of several different networks across North America. After tabulating the results, they categorized the network performance of these systems by node size. As the chart illustrates there is no correlation between service performance and nodes size. In fact in this case the 250 home passed node had an equivalent number of bit errors as did the 2,000 home passed node.

Figure 12

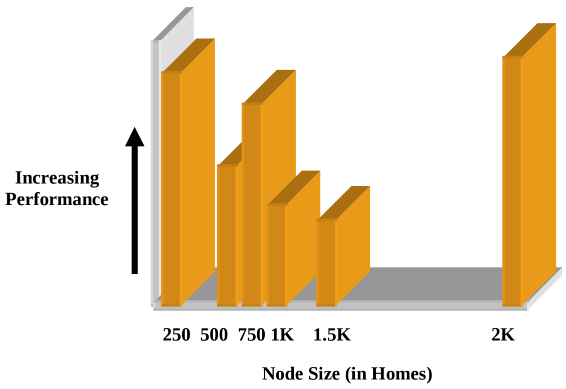

Table 17 is also provided by CableLabs and provides testing results of a 250 home passed node, a 600 home passed node and one 2,000 home passed node. Again you will notice that there is no correlation between node size and bit error rate performance. In these results the 2,000 home passed node exhibited superior performance to the 600 home passed node and nearly identical performance to the 250 home passed node.

Table 17

Node Size Homes Passed	% Performance Above Threshold	
	10^{-4}	10^{-3}
250	99.94	99.95
600	99.87	99.93
2,000	99.93	99.95

The explanation for these results is fairly simple. It isn’t that node size doesn’t have an impact on two way service performance, it is simply that the node size impact is extremely small when compared to other factors which include maintenance standards and component selection.

System amplifier noise accumulation improvements from a 500 home passed node versus a 2,000 home passed node is only 0.8 db and this is when a high grade, DFB laser is used in the return path. If a Fabry Perot laser is used, this improvement becomes even lower.

Because of the random arrival nature of impulse noise, it does not show statistical addition based upon node size. Either you have impulse noise at a level that is impacting service performance or you don't. Several non service impacting impulse noise sources do not aggregate to cause a service affecting impairment.

Unless RF amplifier cascade lengths exceed the thermal range that can be compensated for by the long loop AGC circuit, they will not materially reduce service level performance. RF amplifiers today typically experience MTBF of 60 years and MTTR of 60 to 90 minutes.[12]

5.0 Technological Trends

Experienced design engineers know that they must understand the current state of technology, but even more importantly the future state of technology. This is extremely applicable to node design as laser prices are falling at a rate of 3% per month and new technologies such as dense wave division multiplexing (DWDM) and small form factor dense frequency division multiplexing (SFF-DFDM) technology is rapidly progressing.

These two latter technologies allow system designers to employ scaleable designs that can exploit the nearly unlimited bandwidth of HFC networks.

Figure 13 illustrates equipment that is available today to provide DWDM for up to 8 separate wavelengths

Figure 13

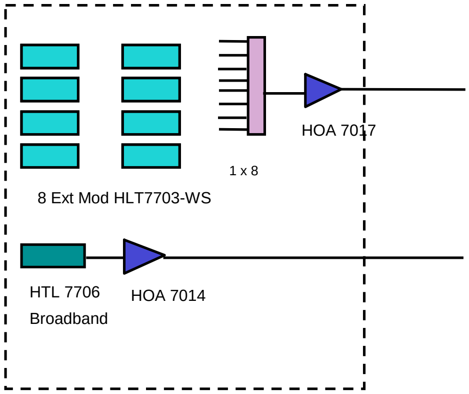

DFDM technology allows for 8, 5-40 MHz sub-low segments to be multiplexed into each wavelength. If one concatenated the two technologies it would allow for 64 nodes to be carried on a single fiber.

Initially, standard WDM technology can be deployed that will allow for two wavelengths per fiber (1310 nm and 1550 nm). These devices are under $500 today and on a rapidly declining cost curve. Use of these devices is always less expensive than installing additional fiber and usually less expensive than deploying higher fiber counts initially.

6.0 Maintenance Considerations

Maintenance and reliability considerations must be taken into account when determining network architecture and node sizes. Fiber installation reduces maintenance costs and improves reliability. HFC reliability has been demonstrated in theory and also in practice through many commercial launches. Fiber plays an important part in this, but node size correlation has not been identified. Nodes of 300, 600, 1200 and 2500 have been compared with no visible correlation to maintenance costs.

The majority of maintenance improvements result from cascade lengths, not node size. This is true also of reliability. Reliability has been mainly the result of power loss, cut cables and craft error. Each of these items are a function that multiplies with cascade length, not homes passed. It needs to be pointed out that increased cascade length for a given home density increases the total homes. The point is that there is no demonstrated reason to artificially limit cascade length to preserve a low number or homes passed per node.

7.0 Node Size Considerations Summary

HFC networks have significant capacity for data communications, so much capacity that even large nodes can support a model based upon 100% modem penetration. This, combined with technology advances in DWDM and SFF-DFDM technologies provide for almost unlimited bandwidth capabilities.

Neither noise contributions nor service level performance limit node size designs. High performance can be achieved even with large nodes. Cascade length should be a consideration and must be defined based upon the long loop AGC capabilities of the modem technology combined with the thermal characteristics of the network.

There is no justification from traffic, performance, maintenance or economics to design nodes below 600 homes passed. In fact, in many cases there are several arguments to design nodes to larger sizes especially in high density areas where reasonable cascade lengths are still being maintained.

Allocating two fibers for every 600 homes passed is more than ample to support the residential business applications. Because of advancements in DWDM technology, these counts may even be reduced in the future.

Acknowledgments

We would like to thank Rogers Communications for their contributions of empirical user behavior data and CableLabs for their contributions relating to WWW surfing and to service levels based upon node size. We acknowledge helpful discussions with Mark Laubach of Com21 and Tom Moore of CableLabs on the traffic models and Dr. Terry Shaw of CableLabs on the Shannon limits. We appreciate the review and counsel given by colleagues.

We would like to specially thank Oleh Sniezko and Esteban Sandino of TCI for their contributions to the sections on theoretical traffic engineering and empirical data. We note heavy use of the traffic references in that section.

REFERENCES

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

[2] CableLabs, *"Data Over Cable Interface Specification,"*

[3] CableLabs, *"Cable Data Modem Performance Evaluation,"* Summer Conference 1996

[4] Chariot Console 2.1 Specification Sheet

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

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

[7] Kevin Thompson, Gregory Miller and Rick Wilder "Wide-Area Internet Traffic Patterns and Characteristics" IEEE November/December 1997

[8] Vern Paxson and Sally Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling," IEEE/ACM Transactions on Networking**,** 3(3) (June 1995): 226-244.

[9] Vern Paxson and Sally Floyd, "Why We Don't Know How To Simulate The Internet," Proceedings of the 1997 Winter Simulating Conference, Lawrence Berkeley National Laboratory, University of California

[10] Vern Paxson,"Fast Approximation of Self-Similar Network Traffic," April 20, 1995, Lawrence Berkeley Laboratory, U.C. Berkeley.

[11] James L. Riggs, Economic Decision Models for Engineers and Managers, McGraw Hill, 1968.

[12] Balakrishnan, V. Padmanabhan, and R. Katz,"The Effects of Asymmetry on TCP Performance", University of California at Berkeley, Berkeley, CA.

[13] Frances Cairncross, The Death of Distance: How the Communications Revolution Will Change Our Lives, Harvard Business School Press, 1997.

ENDNOTES

---

[1] Phil Miguelez, Gary Picard and Fred Slowik, "The Advantages of Gallium Arsenide Hybrid Power Amplifier Technology in Cable Television Applications," 1998 Conference on Emerging Tech Proceedings Manual: Collected Technical Papers (January 1998) :221-244.

[2] C. Day, "Heterojunction Device Technologies: Characteristics and Advantages for Wideband Linear Hybrid Amplifier Design," 1998 Conference on Emerging Tech Proceedings Manual: Collected Technical Papers (January 1998) :245-254.

[3] Claude Elwood Shannon, Claude Elwood Shannon Collected Papers; ed.. N.J.A. Sloane and A.D. Wyner, IEEE Press, 1993.

[4] Robert S. Burroughs, "What Shannon Really Said About Communications And Its Implication to CATV," 1990 NCTA Technical Papers (May 1990) :23-49

[5] Jerry I. Gibson, ed., The Communications Handbook, Florida: CRC Press and IEEE Press, 1997.

[6] Bronwen L. Jones, "Subjective Assessment of Cable Impairments on Television Picture Quality -- A Preliminary Report," 1991 NCTA Technical Papers (March 1991) :92-94.

[7] Bronwen L. Jones, "Subjective Assessment of Cable Impairments on Television Picture Quality," 1992 NCTA Technical Papers (May 1992) :1-25.

[8] Sergio Benedetto, Dariush Divsalar and Joachim Hagenauer, ed., "Guest Editorial: Concatenated Coding Techniques and Iterative Decoding: Sailing Toward Channel Capacity", IEEE Journal on Selected Areas in Communications, (February 1998) :137-139.

[9] The Hurwicz criteria is a measure sometimes used in modeling decisions under risk. A 0.5 criteria is a neutral case, assuming neither more optimistic nor pessimistic conditions than previous.

[10] OpenCable is a trademark of CableLabs and refers to the generation of digital cable terminals that are based upon open standards.

[11] This based upon a high quality (DFB) return laser.

[12] Werner/Sniezko "Network Availability Requirements and Capabilities of HFC Networks" 1996 Conference on Emerging Tech Proceedings Manual: Collected Technical Papers (January 1996) :123-135.

[13] Werner/Gatseos "Network Availability, Consumer Expectations, Plant Requirements and Capabilities of HFC Networks," 1996 NCTA Technical Papers (April 1996) :313-324.