HSD Traffic Behavior in HFC Networks and RDC Interconnects

David Brown, Wayne Ebel, Esteban Sandino and Oleh Sniezko
AT&T Broadband

Abstract

The high-speed data access provided by HFC networks has been a great success story. In many systems and municipalities, the penetration within the first year exceeded all expectations. Moreover, the early adopters were Internet and computer networking savvy and generated high traffic in both downstream and upstream directions of the HFC access network. The same high capacity utilization was experienced in the interconnects between CMTSs, proxy server locations and regional data centers (RDCs).

This paper analyzes capacity utilization of several major components of the high-speed data access in HFC networks:

- *access side of CMTSs (downstream and upstream channels),*
- *network side of CMTSs,*
- *interconnects between CMTS locations and proxy server locations,*
- *RDC LANs, and*
- *interconnects to global Internet.*

This analysis accounts for the number of customers served and for the user behavior. The purpose of this exercise is to develop simple tools for initial network design and capacity engineering for different levels of penetration and for different behavior of the users. The data for this analysis has been collected over a period of several months. This data by itself is interesting and representative of diurnal and weekly traffic patterns, and can be used for capacity engineering in networks shared by different user categories: residential and business.

INTRODUCTION

The explosive growth of the demand for high speed data (HSD) access services, although welcomed and partially anticipated by the HSD access network operators and HSD access service providers, introduced an element of surprise. Traffic generated by the early adopters did not follow the expected patterns. At low penetration rates on the access side of the HSD plant and unpredictable traffic patterns of the early adopters, the efficiency of caching and proxy servers in traffic containment and traffic load reduction was low. This could lead to unexpectedly high capacity utilization if this inefficiency were disregarded. This in turn could result in unexpected capacity exhaustion in an under-engineered HSD network segment.

As the penetration increases, the more predictable user behavior and traffic patterns prevail, and engineering of the interconnect capacity based on the average user behavior becomes well grounded. Moreover, it is also expected that the server and caching efficiency in traffic containment will increase as the number of customers served by the server location increases.

Other factors such as BER, rate shaping and service tiering, customer traffic patterns based on the customer type, and aggressiveness of the IP protocols must be considered and monitored beside using historical data for traffic and utilization prediction and capacity engineering. Therefore, traffic and capacity utilization monitoring as well as development of the demand-extrapolation tools must be continuous in nature.

The authors present several metrics and statistics that may be useful as simple predictive tools. However, these proposals are very preliminary. Although the tools must be simple and intuitively interpretable, they may become quite more accurate and adaptive with the use of today's data analysis and processing engines.

HSD NETWORK ARCHITECTURE EXAMPLES

Three systems were selected for the analysis presented in this paper. The systems/markets were selected to represent three different sizes.

Single-CMTS Architecture

The first system (see Figure 1) is based on a single CMTS with collocated proxy servers. The interconnect with the RDC was initially engineered for a capacity of four T-1s and has been upgraded to 22 Mbps capacity during the period of collecting data for this paper.

An HSD NOC reports *In Data Rate* and *Out Data Rate* for the CMTS on a weekly basis. The *In Data Rate* statistics represent the traffic collected from either the larger Internet or from any local or proxy servers within the logical data network segment, and forwarded to customer cable modems (CMs). This is equivalent to "downstream traffic". The *Out Data Rate* statistics represent the traffic collected from customer cable modems within the HFC service areas and destined to either the global Internet or to any local or proxy servers within the logical data network segment. This is equivalent to "upstream traffic". The NOC also provides statistics for each downstream data transmitter and upstream data receiver on the access side of the HSD HFC plant.

Figure 1: HSD Architecture for Single-CMTS Network

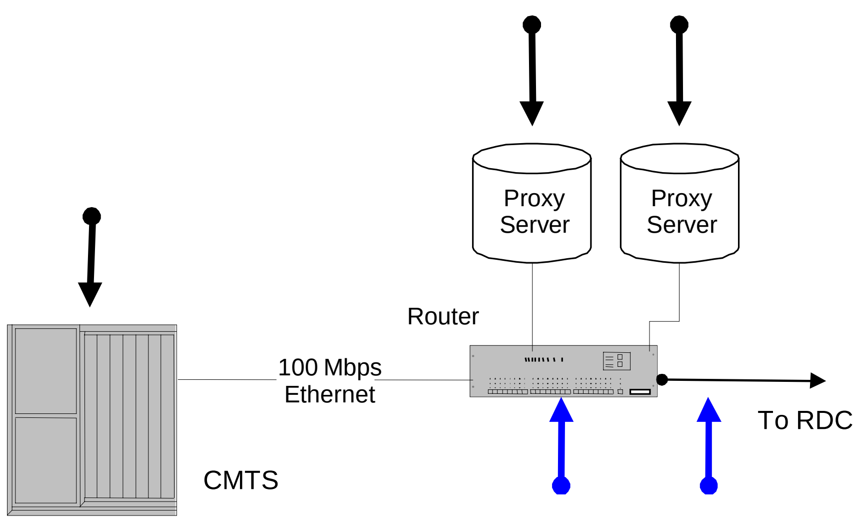

Medium Size CMTS Configuration

The second network (see Figure 2) consists of three CMTSs connected via a Fast Ethernet link with a proxy server location. This location is in turn interconnected with an RDC via a 100 Mbps EtherRing.

As described above, the HSD NOC reports *In Data Rate* and *Out Data Rate* for each CMTS 100BaseT input/output interface on a weekly basis.

Figure 2: Medium-Size CMTS Configuration

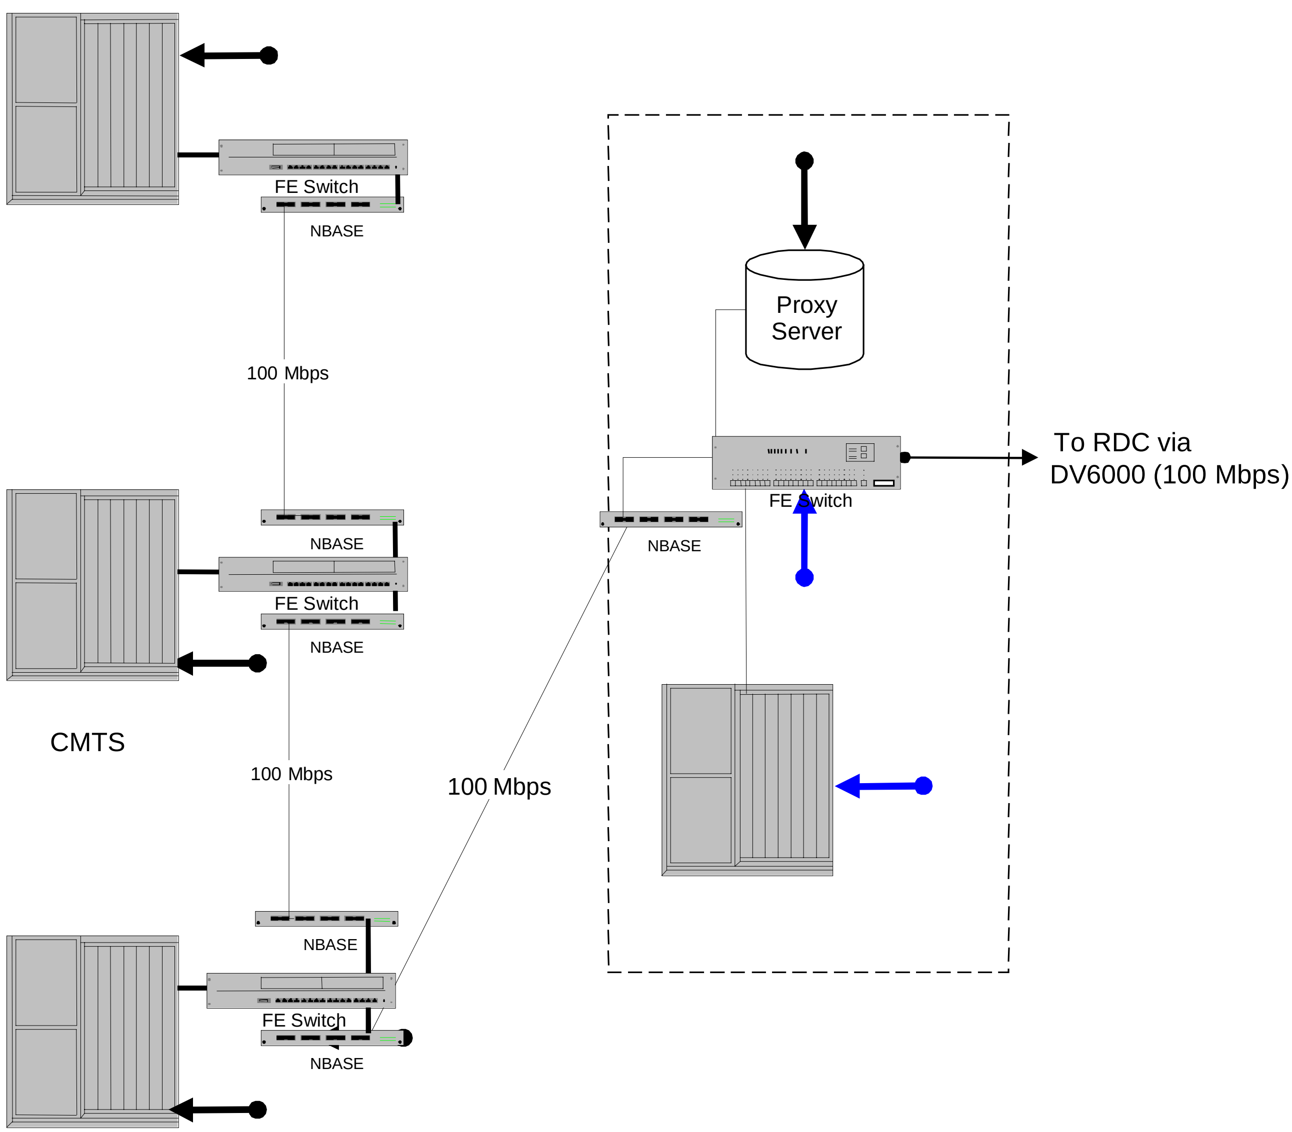

Large-Size HSD Network

An example of the large-size network consists of 12 headends and 32 CMTS units. However, traffic data for this network was available only for 20 of the CMTS units. Traffic statistics were also available for the Gigabit Ethernet internal RDC LAN at the master headend. The internal RDC LAN aggregates traffic for the entire market data network. This Gigabit RDC LAN also provides access to the global Internet for all customers in this market.

The distribution data network consists of a single 100BaseT Fast Ethernet backbone traversing from the northwest headend through all west headends and hubs (CMTS locations) to the master headend, and finally heading through east headends and hubs (CMTS locations) into the hub in the northeast service area. Additional 100BaseT spurs along this main route feed into the main backbone. There is a single proxy server at one headend in the West and two proxy servers at another headend in the East.

Logically, the main Fast Ethernet backbone is split into two main segments: West and East. The West Fast Ethernet segment services CMTSs in several headends and hubs in the West. All Internet content requests generated from customers within the West Ethernet backbone service area are first directed to the proxy server in one of these headends. At this location, a determination is made on whether to forward those requests to the RDC and the larger Internet cloud.

The East Fast Ethernet segment services CMTSs in several headends and hubs in the East. All Internet content requests generated from customers within the East Ethernet backbone service area are first directed to the two proxy servers in one of these headends. As is the case with the West portion of the Fast Ethernet backbone, the two proxy servers determine whether to forward content requests to the RDC and the larger Internet cloud.

As already described, for each CMTS 100BaseT input/output interface, the HSD NOC reports *In Data Rate* and *Out Data Rate* on a weekly basis.

The NOC also reports *In Data Rate* and *Out Data Rate* for the Gigabit Ethernet internal RDC LAN in the master headend. In this case, the *Out Data Rate* represents incoming traffic from the global Internet that is forwarded to both the West and East Fast Ethernet backbone segments for distribution to all customers in the market data network, i.e., “downstream traffic”. The *In Data Rate* represents incoming traffic originating from all customer data terminals arriving at the RDC from both the West and East Fast Ethernet backbone segments and destined to the global Internet, i.e., “upstream traffic”.

CMTS TRAFFIC ANALYSIS

Due to the concerns of the industry outsiders, the access network capacity engineering has been the focus of the HFC engineering effort, despite repeated reports on the results of simulation and traffic modeling that showed significant capacity margins in the access plant. These results can be found in several publications (two are referenced in this paper).

The traffic statistics presented in the next several subheadings will support the earlier results and will also allow to draw some preliminary conclusions that may be useful during capacity engineering effort and in capacity exhaustion predictions.

Single-CMTS System

Channel Utilization — Downstream

The CMTS in this market served a maximum of 1313 active modems during the week for which the traffic statistics are presented in Figure 3. The transmitter s6.p6 served a maximum of 820 active modems. As can be seen from Figure 3, downstream traffic did not reach at any point 50% of the capacity utilization (downstream capacity equal to approximately 27 Mbps).

After traffic aggregation on the network side of the CMTS, the maximum traffic levels were much below the capacity of the CMTS.

An interesting conclusion can be derived from the analysis of the downstream data rate per CM. It seems that the log-linear plot of this statistic (see Figure 5) fits very well a power trendline. However, the data represent too few observations to draw binding conclusions.

Channel Utilization — Upstream

The transmitter s6.p6 served a maximum of 820 active modems. As can be seen from Figure 4, upstream traffic did not significantly exceed 50% of the capacity utilization for any receiver at any point in time (upstream capacity equal to approximately 2.5 Mbps for each of the six upstream receivers).

The upstream data rate per CM does not follow the same pattern as data rate in the downstream direction. However, the upstream packet rate per CM (see Figure 4 for raw data plots) for this particular example follows similar pattern (see Figure 6) to the pattern for the data rate per CM in the downstream direction.

Diurnal and Weekly Traffic Patterns

The downstream traffic reaches the highest data rates at midnight, tapers down fast to reach a minimum at 6 a.m. and then starts increasing at a slower or faster rate to close the daily cycle at midnight. This pattern is discernable in both downstream traffic data on the access side and downstream aggregated traffic data on the network side of the CMTS as long as the number of users is sufficient to generate sizeable traffic and perform statistical traffic aggregation. This pattern is not clearly visible on the access side for the upstream traffic. However, the aggregated data rates for the upstream traffic on the network side as well as the packet rates for the upstream traffic on both sides of the CMTS show similar patterns. The difference for the upstream data rate diurnal behavior on the access side can be explained by statistically smaller group of users and by higher variability of the packet sizes in the upstream direction. In most cases, the downstream and upstream data and packet rate traffic reaches its peak during weekends. Moreover during weekends, the traffic reaches high levels already at noon.

To engineer the access side of the HSD plant, one must account for the peak values of the data rates unless traffic shaping and throttling schemes are implemented either for downstream or upstream traffic or in both directions.

Figure 3: Single-CMTS System Data Rate Statistics

a) Access Side of CMTS — Downstream

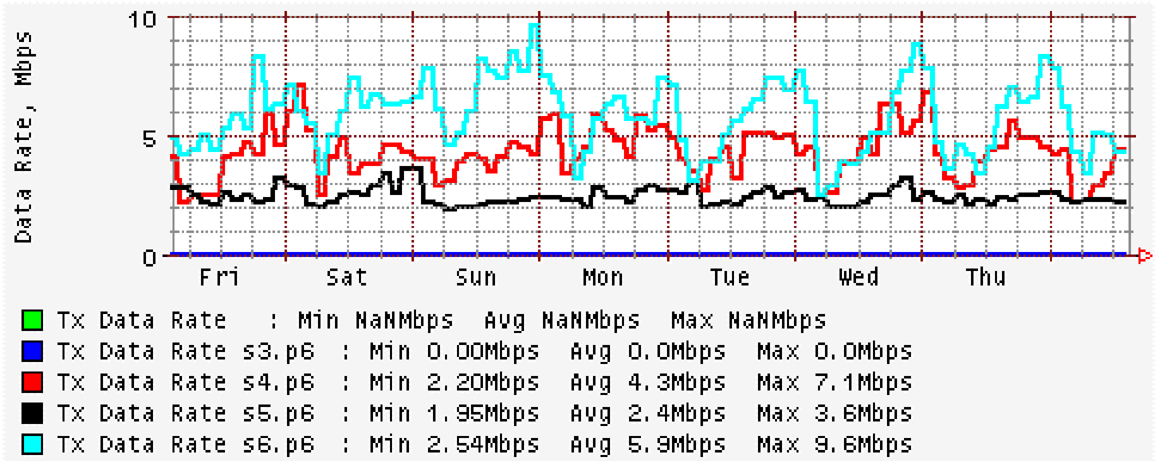

b) Access Side of CMTS — Upstream Receivers Corresponding to TXs6.p6

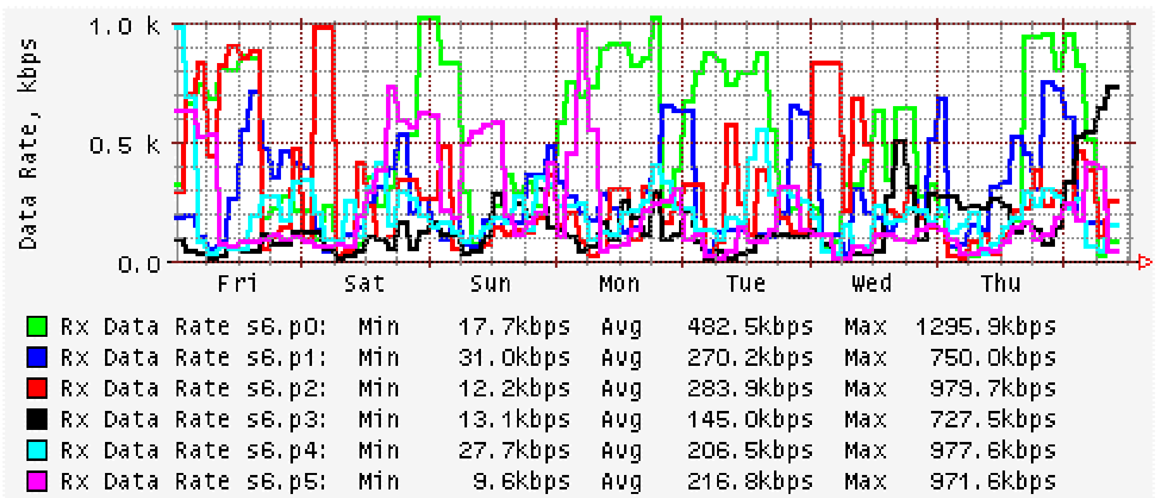

c) Network Side of CMTS

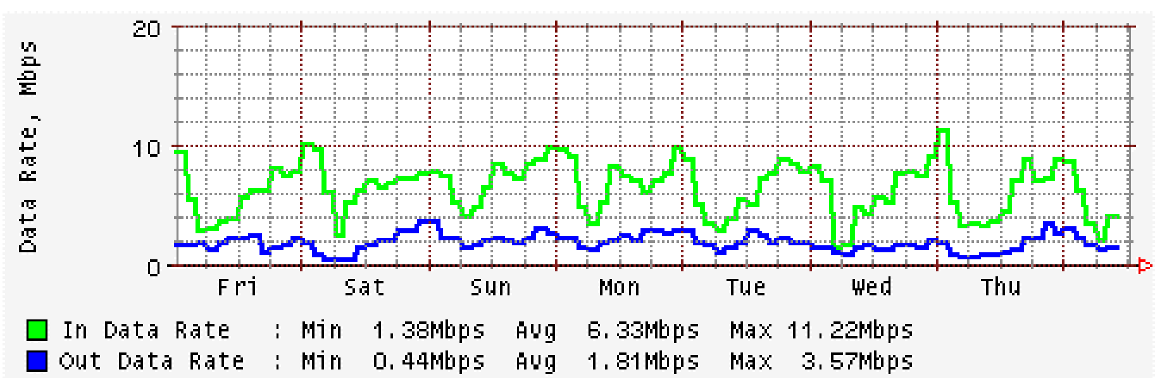

Figure 4: **Single-CMTS System Packet Rate Statistics**

a) **Access Side of CMTS — Downstream**

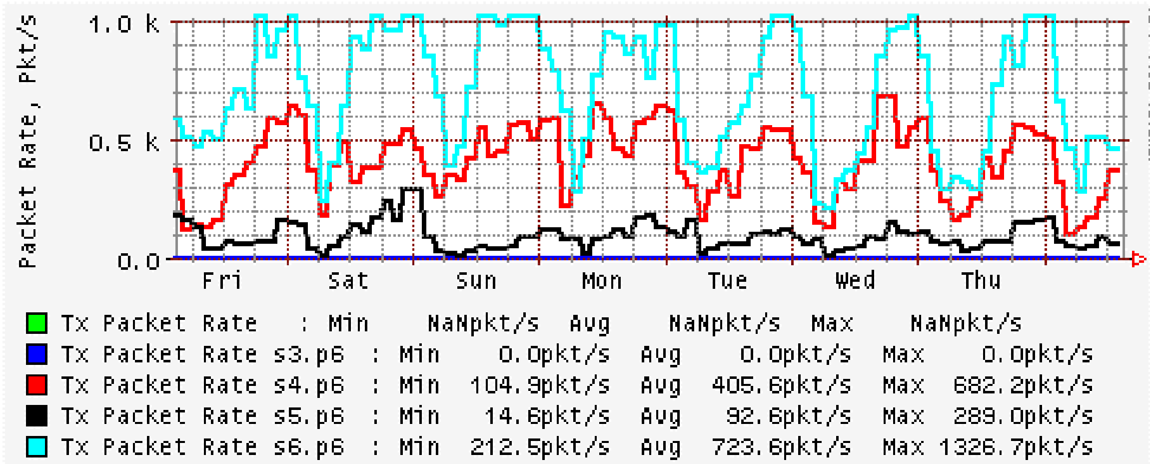

b) **Access Side of CMTS — Upstream Receivers Corresponding to TXs6.p6**

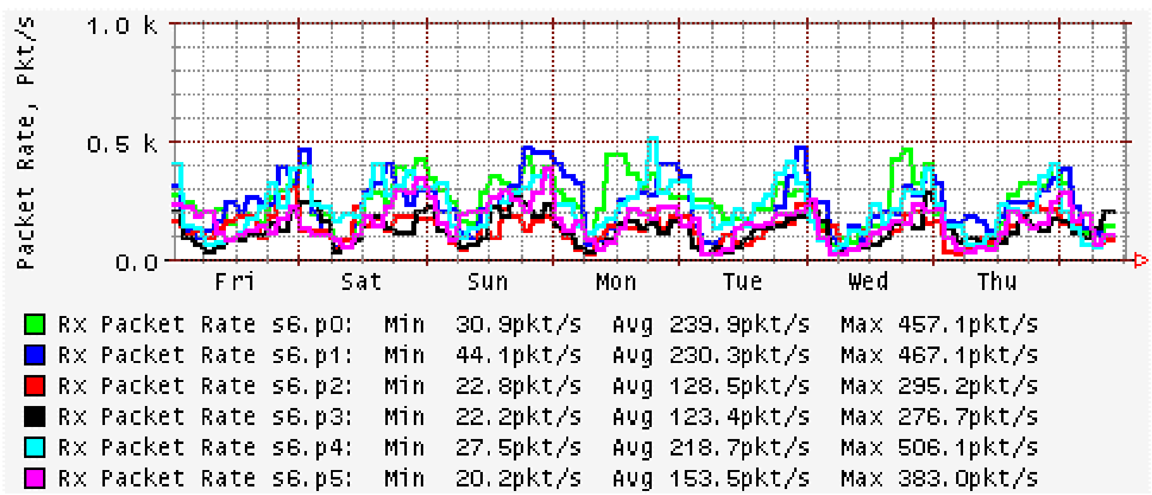

c) **Network Side of CMTS**

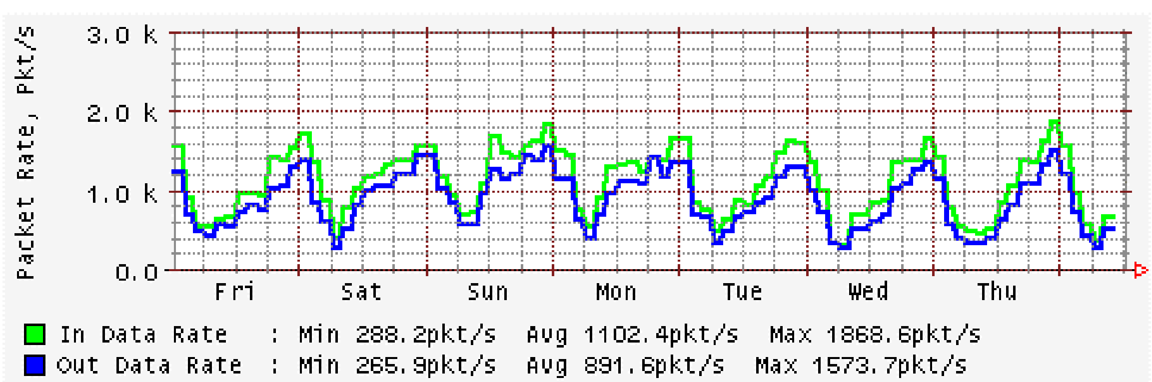

Figure 5: Downstream Data Rate/CM Behavior

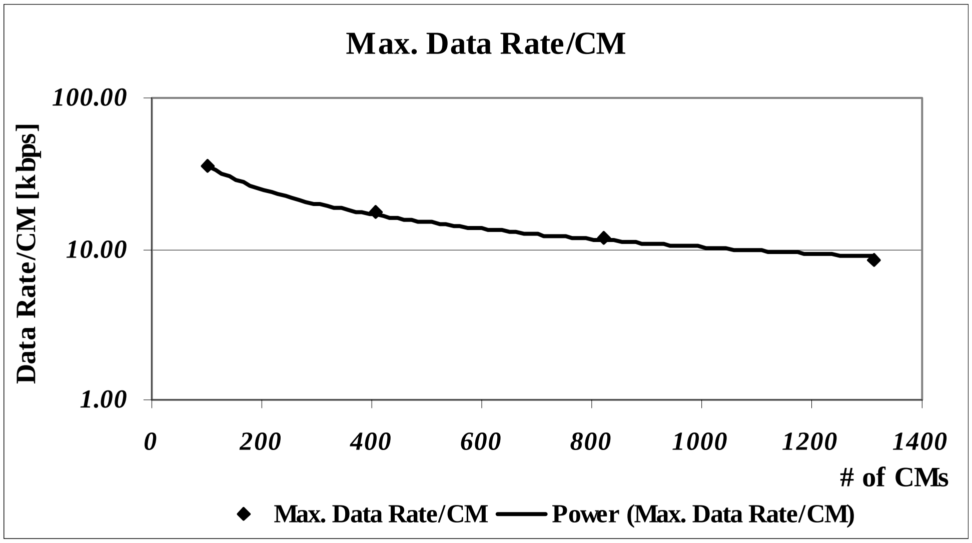

Figure 6: Upstream Packet/CM Behavior

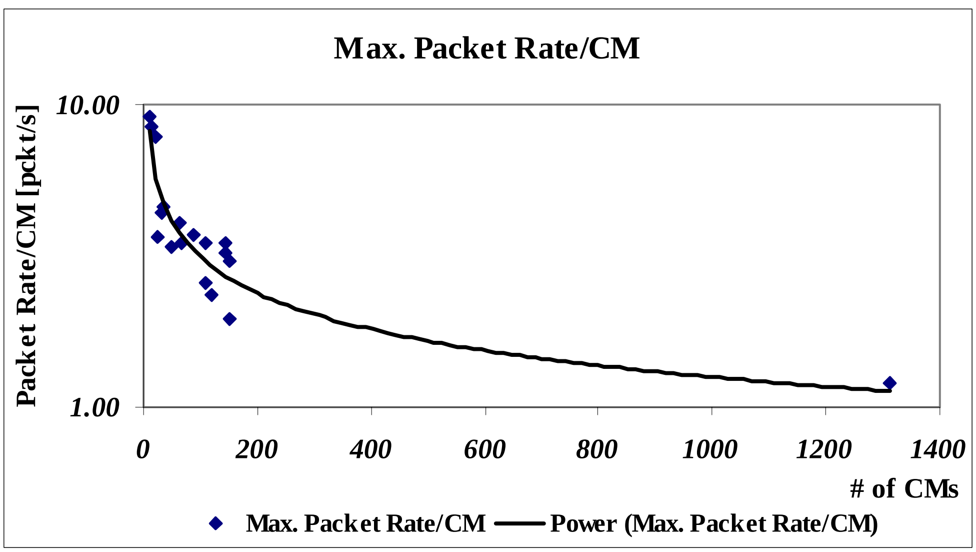

Medium-Size HSD Configuration

The traffic patterns from each CMTS have similar characteristics as the traffic characteristics for the CMTS in the single-CMTS system. There are some small differences in the upstream traffic patterns that will be analyzed later in this paper.

Of interest is the data presenting statistics on the traffic in the link between the RDC and the proxy server location for this system. These traffic records show that the diurnal and weekly patterns described above for a single CMTS are even more pronounced at higher levels of traffic aggregation. Moreover, the aggregate traffic reached above 20 Mbps (peak value of 5-minute averages) over the period of six months since the service launch. The number of active modems reached 4,000 over that period of time.

Downstream Data Rate/CM Behavior

The results of the regression analysis show that there is a reasonably good fit between the collected data points and power trendline described by the following equation:

$$y = cx^b$$

where c and b are constants.

In the case depicted in Figure 7, the data rate per CM drops approximately by a factor of two for each quadrupling of the number of the active CMs (total data rates increase twice for each quadrupling of the CM number). For 1,000 of CMs, total data rate reaches 10 Mbps (10 kbps/CM); for 4,000 CMs, total data rate reaches approximately 20 Mbps (5 kbps/CM). These statistics should be continuously verified as the customer behavior may change as well as new, more bandwidth demanding applications may start dominating the downstream traffic. However, they can be used today for capacity engineering of the access and network side of the CMTSs.

Figure 7: Downstream Data Rate/CM Behavior for Large Sample of Observations

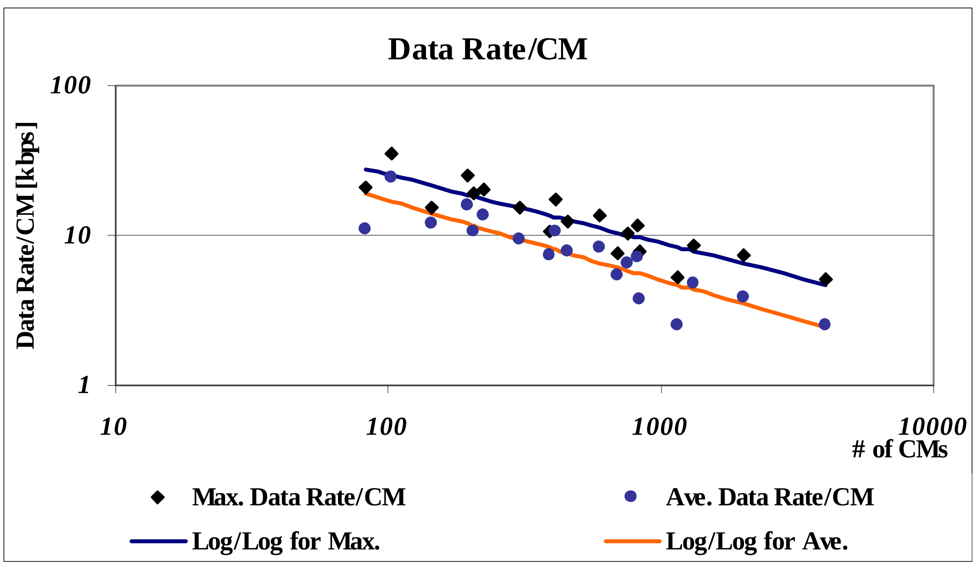

Table 1: Linear Regression Analysis Results for Max. Downstream Data Rates (Log-Log Scales)

Regression Statistics	
Multiple R	0.8915
R Square	0.7947
Adjusted R Square	0.7827
Standard Error	0.1064
Observations	19

	df	*SS*	*MS*	*F*	*Significance F*
Regression	1	0.7457	0.7457	65.81875	3.01446E-07
Residual	17	0.1926	0.0113		
Total	18	0.9383			

	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*
Intercept	2.3264	0.1537	15.1401	2.67E-11	2.0022	2.6506
X Variable 1	-0.4602	0.0567	-8.1129	3.01E-07	-0.5799	-0.3405

Upstream Packet/CM Behavior

The results of the regression analysis for the upstream traffic are less conclusive. There is no good fit for data rates as a function of the active CM number. There is only acceptable fit between the packet rate and the number of active CMs (see Figure8). The only convincing conclusion based on the data collected is that, for CM numbers exceeding 500, the packet rate is approximately 1 pkt/s/CM and the data rate does not usually exceed 3 kbps/CM. These numbers can be used for capacity engineering, especially when supported by the previous conclusion that for the number of CMs approaching 1,000 units, upstream channel capacity utilization rarely exceeds 50%. With four to six channels per single downstream channel of 27 Mbps (this channel can serve approximately 6,000 active modems), there is sufficient capacity of 10 to 15 Mbps in the upstream direction. Based on the numbers listed above, this capacity could serve 8,000 to 12,000 active modems.

Figure 8: Upstream Packet Rate/CM Behavior for Large Sample of Observations

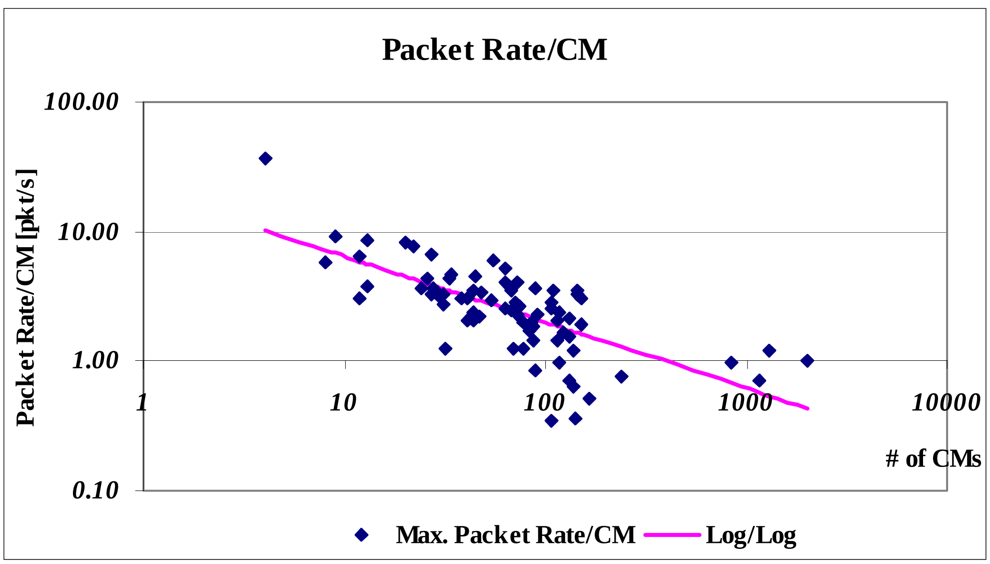

Table 2: Linear Regression Analysis Results for Max. Upstream Packet Rates (Log/Log Scales)

Regression Statistics	
Multiple R	0.7052
R Square	0.4973
Adjusted R Square	0.4906
Standard Error	0.2375
Observations	78

	df	*SS*	*MS*	*F*	*Significance F*
Regression	1	4.2412	4.2412	75.1696	5.72298E-13
Residual	76	4.2881	0.0564		
Total	77	8.5293			

	Coefficients	*Standard Error*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*
Intercept	1.3189	0.1101	11.9792	3.51E-19	1.0996	1.5382
X Variable 1	-0.5096	0.0588	-8.6700	5.72E-13	-0.6266	-0.3925

Assumptions and Definitions

The downstream channel capacity is estimated for 6 MHz 64 QAM DOCSIS channel, the upstream capacity is estimated for 1.6 MHz QPSK DOCSIS channel.

All the conclusions for the downstream and active traffic and customer behavior are based on historical data and are may be dependent on other factors not included in the analysis.

TRAFFIC AGGREGATION BY PROXY SERVERS IN LARGE-SIZE MARKET

Available Data

Data for the large-size market has been collected from individual 100BaseT interfaces at each of the CMTSs polled. As such, this data reflects the raw traffic before the local headend data routers direct it to either the proxy servers (if available) or into the appropriate 100BaseT Fast Ethernet backbone segment. The 100BaseT Fast Ethernet backbone segments carry all traffic into the master headend and RDC location. The RDC location is the main gateway to access the global Internet. Access to the Internet cloud is via two OC-3 packet-over-Sonet (POS) interfaces.

The traffic statistics for the Gigabit Ethernet internal RDC LAN that aggregates all incoming data from the 100BaseT Fast Ethernet backbone are also available. Unfortunately there was no meaningful data available from the proxy servers in this market at the time the data was collected.

Data Analysis Results

From among 32 CMTSs in this market, data was available for 20 units only. The resulting estimated traffic statistics database allows for aggregation of the specific traffic sources within the data network.

Traffic aggregation has been done separately for the sources feeding the West Fast Ethernet and the East Fast Ethernet logical backbone segments. This allowed for an estimate of the current utilization for each of these two segments. Traffic aggregation for all combined sources has also been done. This latter aggregation allowed for direct comparison against the reported traffic statistics available for the internal RDC Gigabit LAN. It also allowed for an assessment of how effectively the proxy servers at both West and East proxy server locations perform traffic containment and traffic load reduction on each of the two Fast Ethernet logical backbone segments. The following sections summarize the findings to date.

Utilization of West and East Fast Ethernet Backbone Segments

Figures 9 and 12 for the West and East Fast Ethernet backbone segments respectively illustrate the Internet traffic transported from the RDC Gigabit LAN to each CMTS. The West Fast Ethernet backbone peaks at 90% utilization (90 Mbps). This peak utilization starts at around 6:00 p.m. everyday and reaches its maximum level of 90 Mbps at just before midnight. Utilization tapers off rapidly after that to its minimum level at around 6:00 a.m. everyday. The East Fast Ethernet backbone has a worst-case peak of 100% utilization (100 Mbps). A similar pattern of traffic utilization is present here.

The charts in Figures 9 and 12 indicate that both Fast Ethernet backbone segments are heavily utilized. Figures 10 and 13 for the two Fast Ethernet backbone segments illustrate the traffic transported from each CMTS location to the RDC Gigabit LAN and destined to the global Internet (upstream traffic). This upstream traffic utilization is very uniform. Maximum utilization for this case is less than 20 Mbps for the West Fast Ethernet backbone and less than 30 Mbps for the East Fast Ethernet backbone. Figures 11 and 14 illustrate the ratio of downstream to upstream traffic for the West and East Fast Ethernet backbone segments respectively and have been included here for information purposes only. The observed ratios fall between 3:1 and 6:1. The higher ratios appear to track the volume of traffic and occur at approximately the same time as the peak utilization for each Fast Ethernet backbone segment.

Figure 9: Aggregate Downstream Traffic for West Backbone

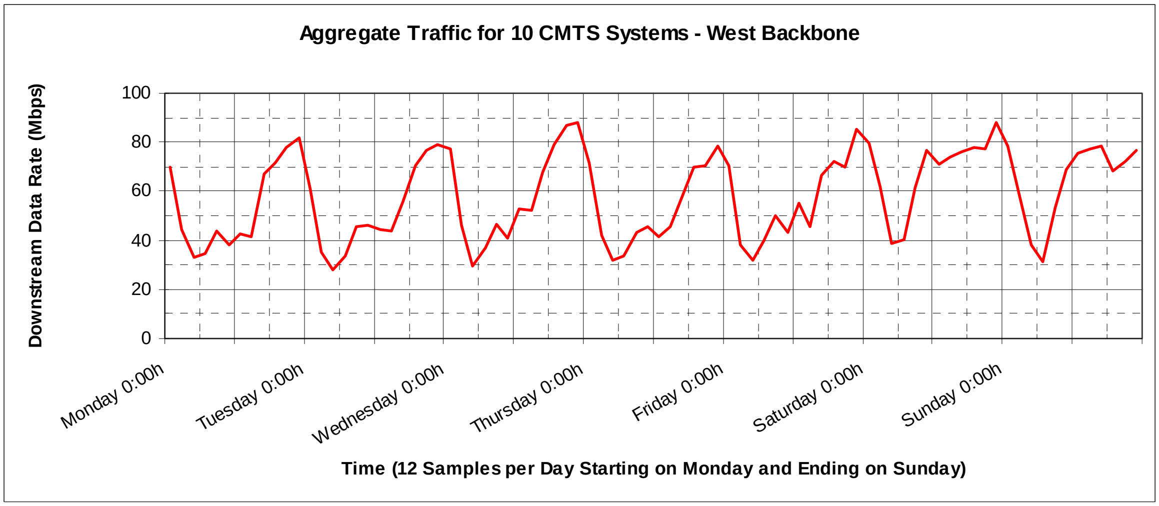

Figure 10: Aggregate Upstream Traffic for West Backbone

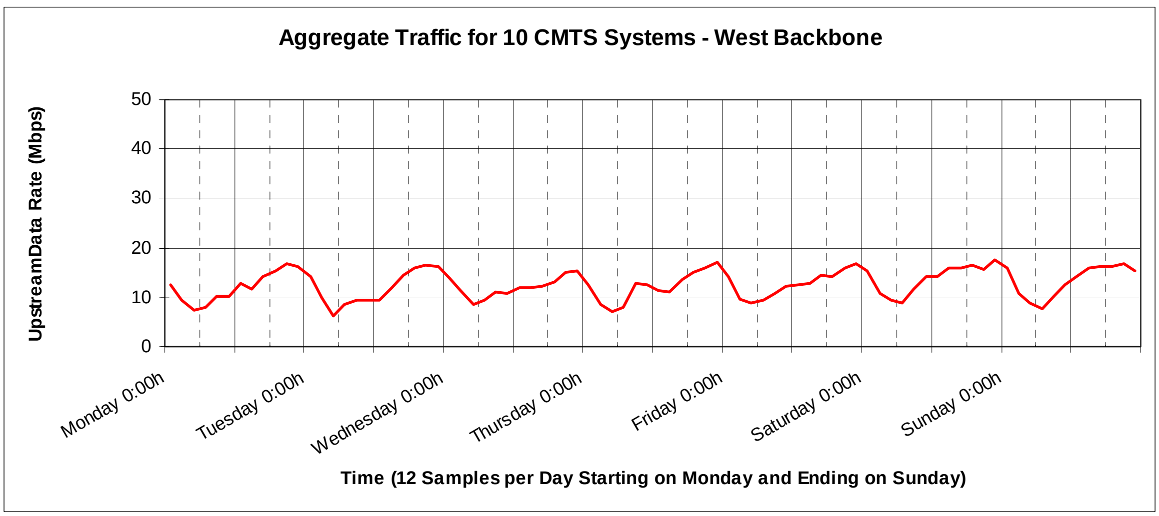

Figure 11: Downstream-to-Upstream Data Rate Ratios – West Backbone

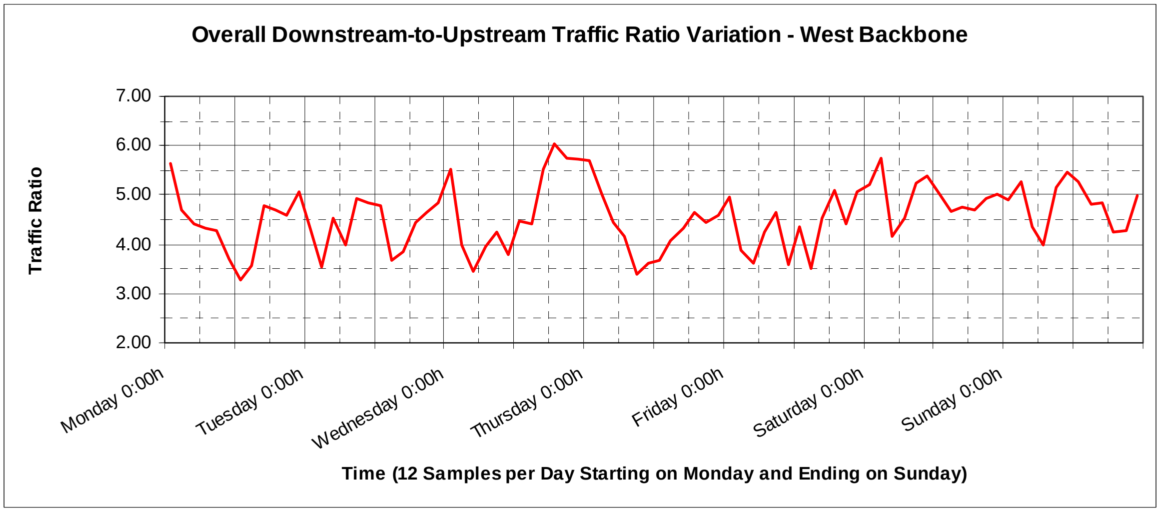

Figure 12: Aggregate Downstream Traffic for Eat Backbone

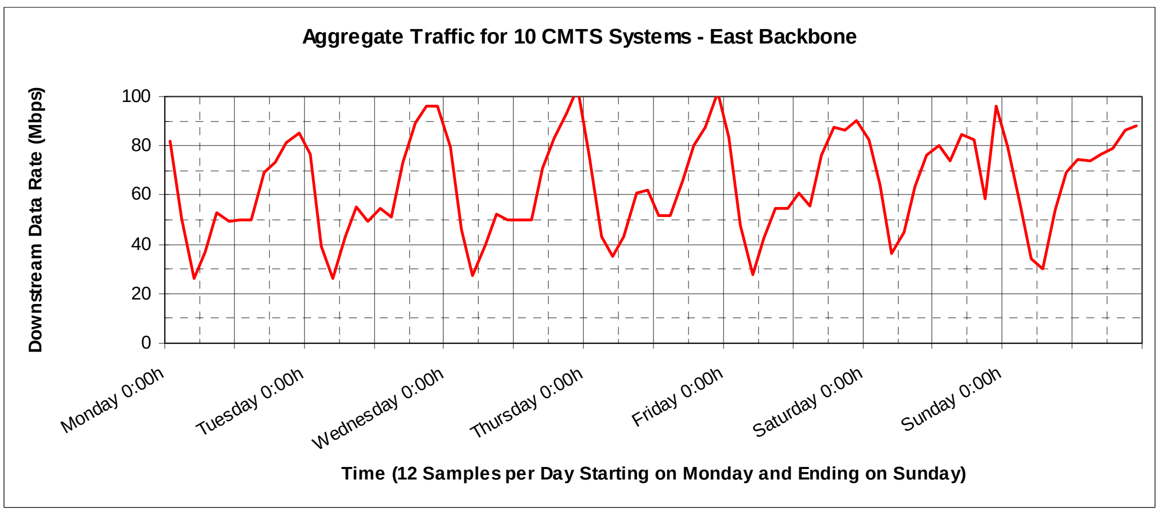

Figure 13: Aggregate Upstream Traffic for East Backbone

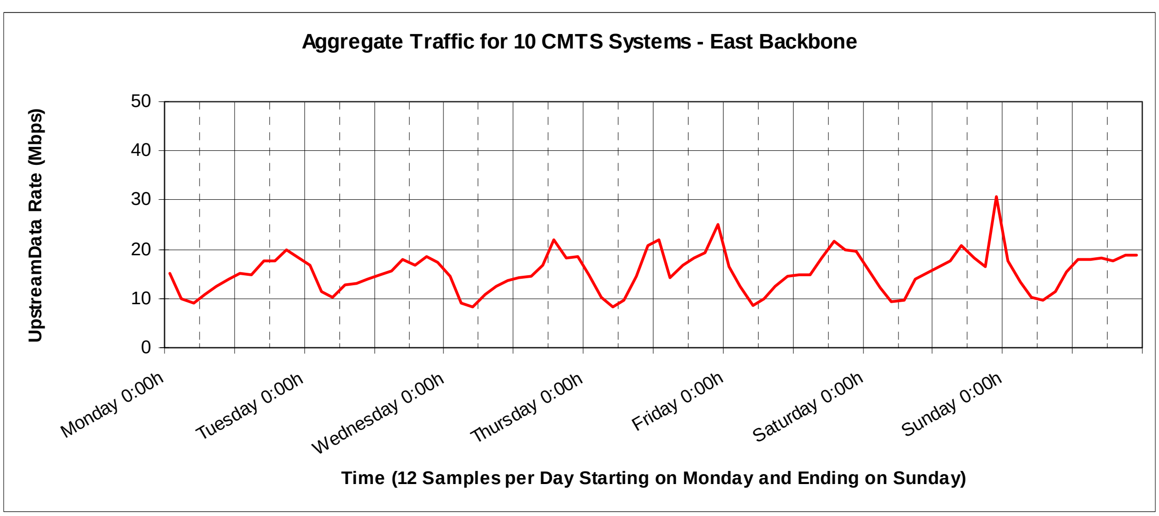

Figure 14: Downstream-to-Upstream Data Rate Ratios – East Backbone

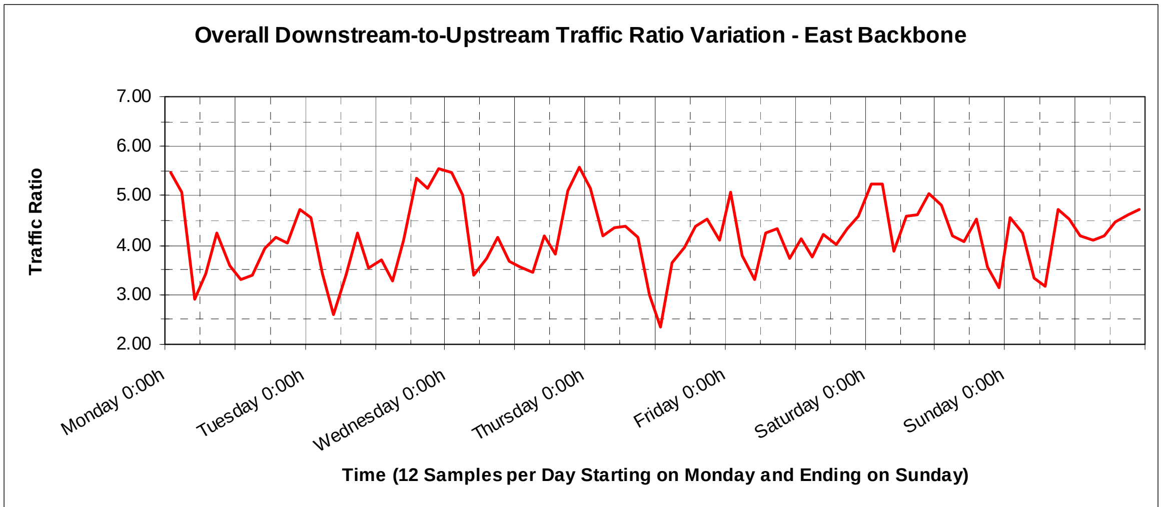

T

Utilization of Internal RDC Gigabit LAN

Figure 15 shows the aggregation of all the traffic entering each of the monitored 100BaseT interfaces for all 20 CMTS units under study, i.e., downstream traffic into the subscriber modems. This is the aggregation of all the reported *In Data Rate* traffic in the charts for the CMTS units.

Figure 16 shows the aggregation of all the traffic leaving each of the monitored 100BaseT interfaces for all 20 CMTS units under study, i.e., upstream traffic generated from all the subscriber modems. This is the aggregation of all the reported *Out Data Rate* traffic in the charts for the CMTS units.

Figure 17 shows a representative weekly traffic chart tracking utilization for the Gigabit Ethernet internal RDC LAN. The Out traffic trace should represent all incoming Internet data traffic leaving the RDC LAN for distribution into each of the CMTS units via the 100BaseT Fast Ethernet backbone segments. The In traffic trace should represent aggregation of all data traffic entering the RDC LAN from the 100BaseT Fast Ethernet backbone segments and destined for the global Internet.

The proxy servers in the network are intended to contain Internet-related traffic (web page traffic) within specific sub-network domains and away from the 100BaseT Fast Ethernet backbone segments. The objective is to minimize the amount of duplicate packets on each of the backbone segments and eliminate duplicate user requests to access the same information over a period of time. The proxy servers in this scenario would store the most requested Web pages and Internet content locally. Local storage closer to the end user should also result in faster access to the requested information, i.e., reduced delays.

Under the current scenario, we would expect to find the utilization of the Gigabit Ethernet internal RDC LAN for both the *Out* and the *In* traces to be somewhat less than what we would find if straight traffic aggregation from all CMTS sources took place. The proxy servers would be expected to limit traffic on each of the Fast Ethernet backbone segments feeding into the RDC LAN. The data collected so far does not support this expectation. A quick visual check of Figures 15 through 17 shows that almost all In and Out traffic for the Gigabit RDC LAN is aggregated in a straightforward way. However, the above analysis is by no means complete. We are missing some of the traffic contributions for the additional CMTS sources. The current analysis represents only 20 out of the 32 CMTS unit universe in the distribution network.

Fast Ethernet Backbone Utilization

Another interesting finding is that, although a Gigabit Ethernet LAN is implemented at the RDC, the two Fast Ethernet backbone segments connecting all CMTS traffic aggregation points are still point-to-point Fast Ethernet links. From the RDC perspective, there are two separate 100BaseT links feeding traffic into the Gigabit Ethernet LAN. However, these two Fast Ethernet links are approaching the limit of their capacity.

Figure 15: Aggregate Downstream Traffic for 20 CMTSs

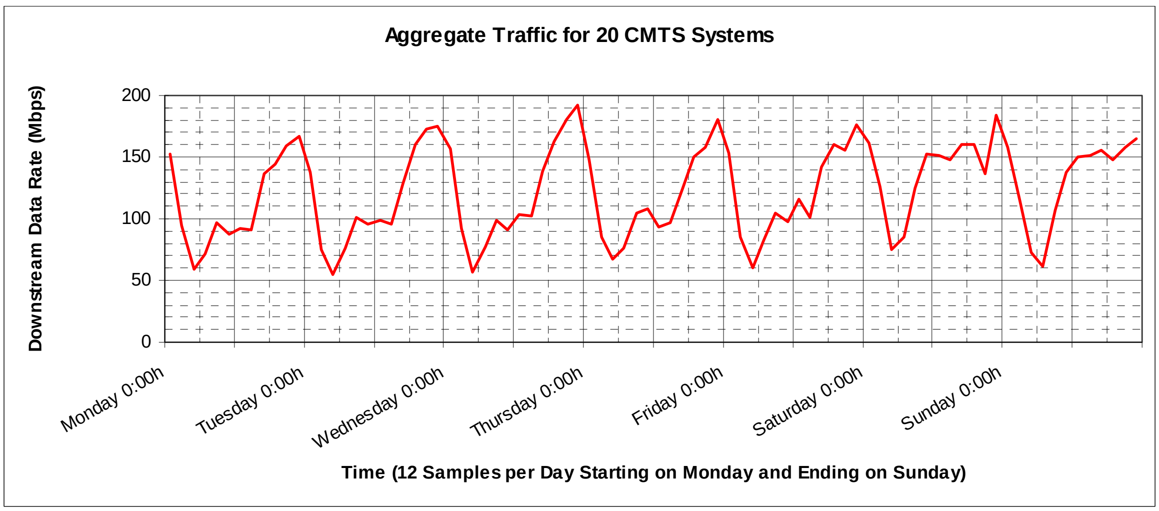

Figure 16: Aggregate Upstream Traffic for 20 CMTSs

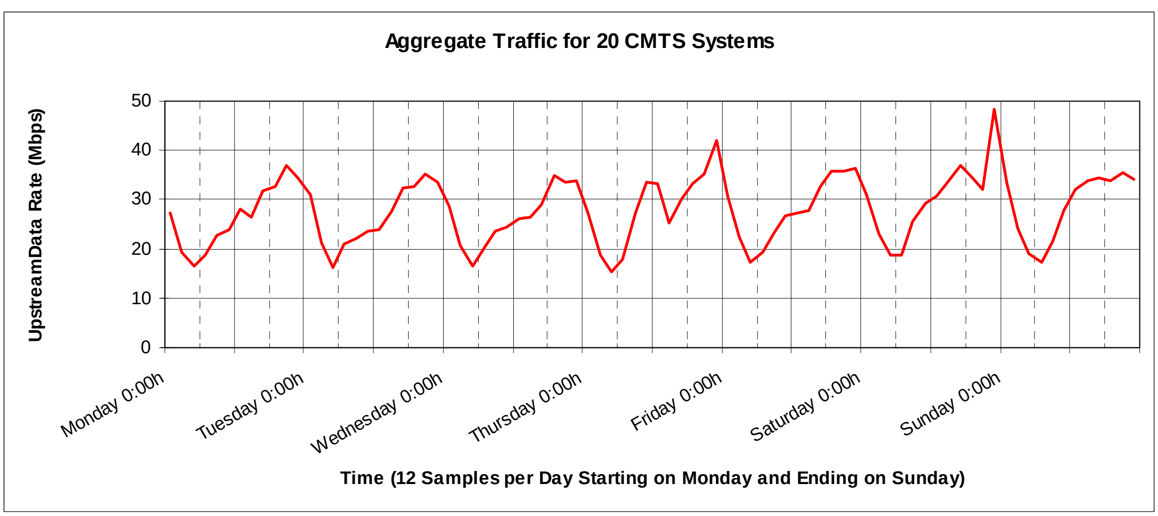

Figure 17: Traffic Activity in Gigabit Ethernet RDC LAN

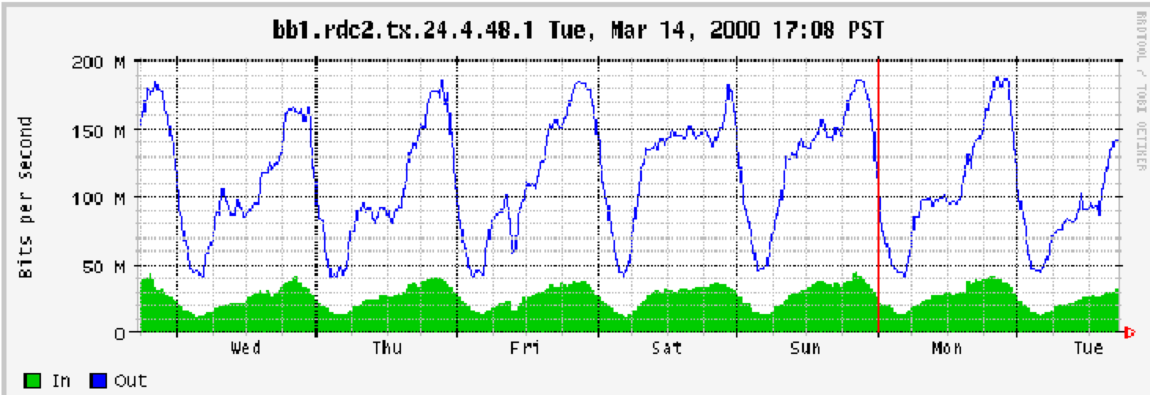

TRAFFIC STATISTICS ANALYSIS

Table 3 below contains some data rate and packet rate statistics that can be useful in capacity engineering.

Table 3: Simple Traffic Statistics

Statistic	Average Value	Standard Deviation
Downstream/Upstream Packet Rate Ratio	1.20	0.0571
Downstream/Upstream Data Rate Ratio	3.20	0.70
Downstream Packet Size (Network Side)	746 bytes	42 bytes
Downstream Packet Size (Access Side)	1295 bytes	277 bytes
Traffic Aggregation Gain for Downstream Packet Rates	121%	8%
Traffic Aggregation Gain for Downstream Data Rates	224%	55%
Upstream Packet Size (Network Side)	305 bytes	60 bytes
Traffic Aggregation Gain for Upstream Packet Rates	224%	91%
Traffic Aggregation Gain for Upstream Data Rates	360%	57%

The data shows that the upstream traffic aggregation gains (also called the multiplexing gains) are much higher than those for downstream traffic are. This fact indicates that upstream packet rate and data rate peaks from different customer groups do not coincide in time as well as the downstream rates do. Moreover, these gains in both direction are much lower than previously expected (usually for traditional LAN applications, the access to trunk capacity ratios are much higher). This may be caused by the fact that traffic peaks from different customers coincide in time.

Also, traffic rate asymmetry is much lower than expected. For peak and average traffic, the asymmetry for packet rates is 1.2:1 and for data rates is 3.2:1 downstream to upstream.

OTHER FACTORS TO CONSIDER

BER Impact

Figure 18 shows an example of the access network with very low BER and FEC activity. Figure 9, on the other hand, shows an example of the access network with very high BER and FEC activity. The impact on the upstream traffic pattern and behavior is visible. Some impact can be also observed in the downstream traffic (not shown in Figures 18 and 19).

Figure 18: Low BER and FEC Activity Access Plant

a) Data Rate

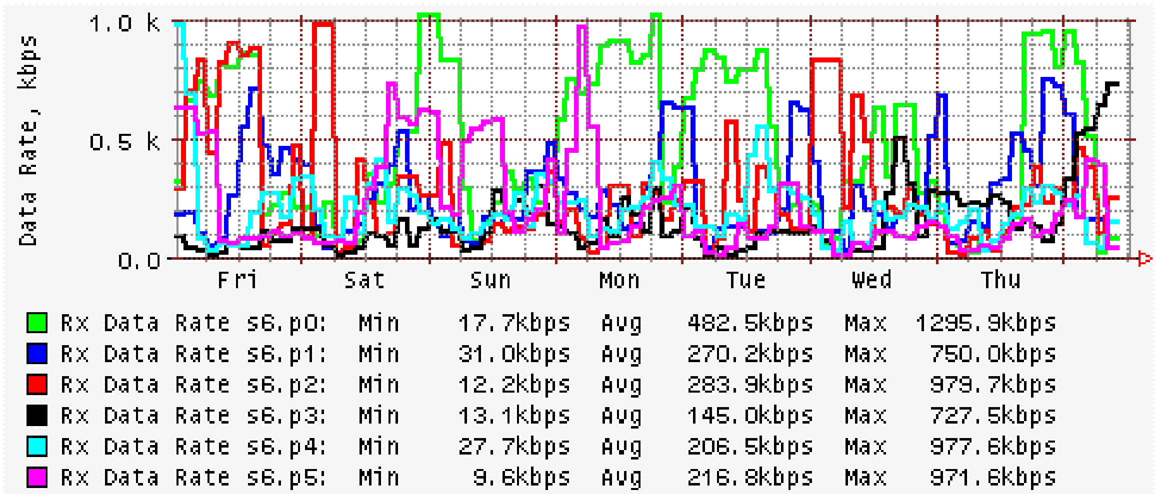

b) Packet Rate

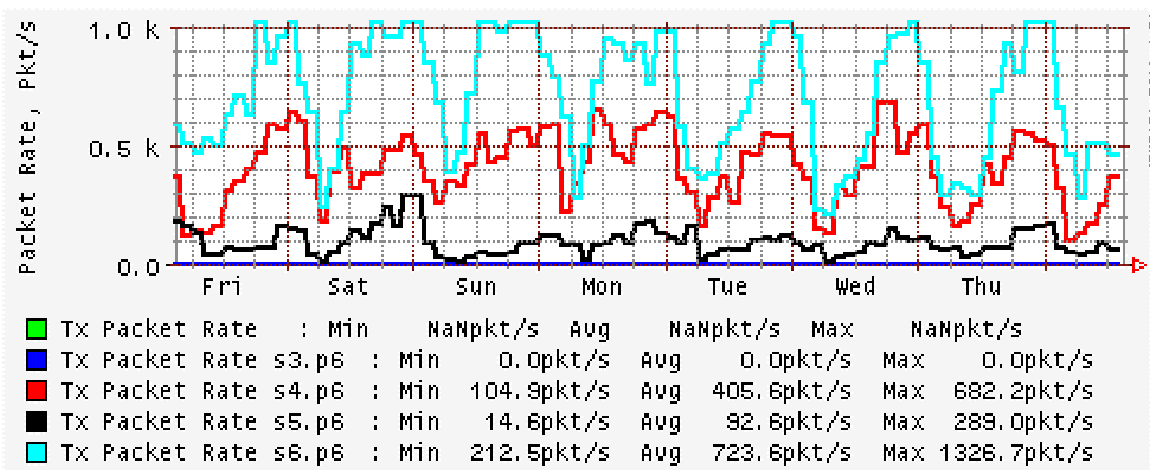

c) BER and FEC Activity

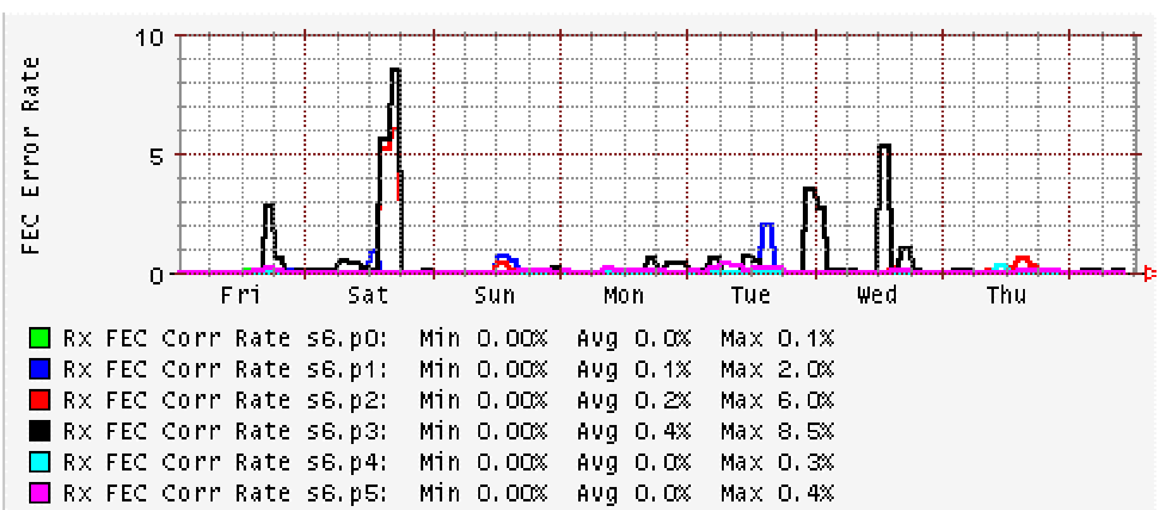

Figure 19: High BER and FEC Activity Access Plant

a) Data Rate

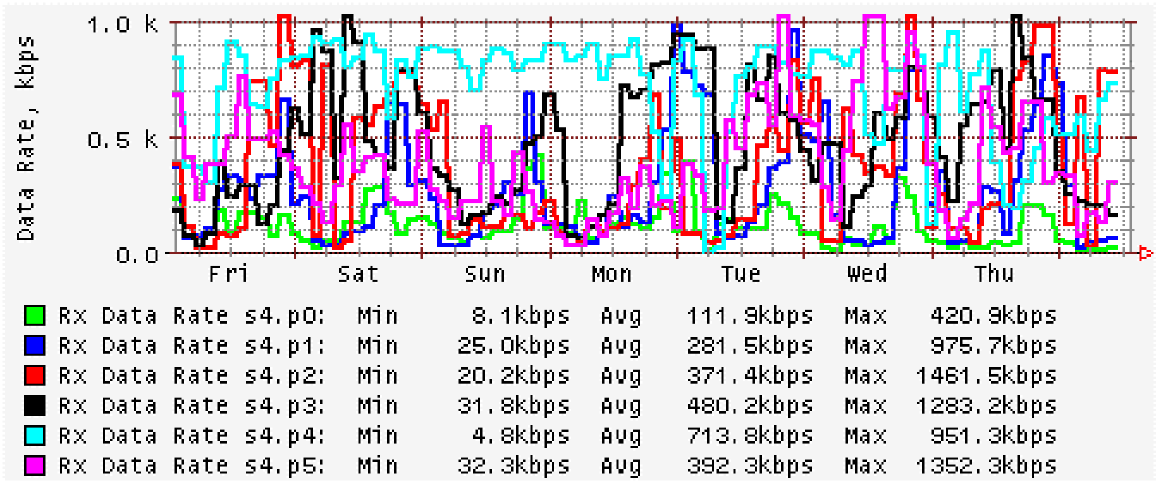

b) Packet Rate

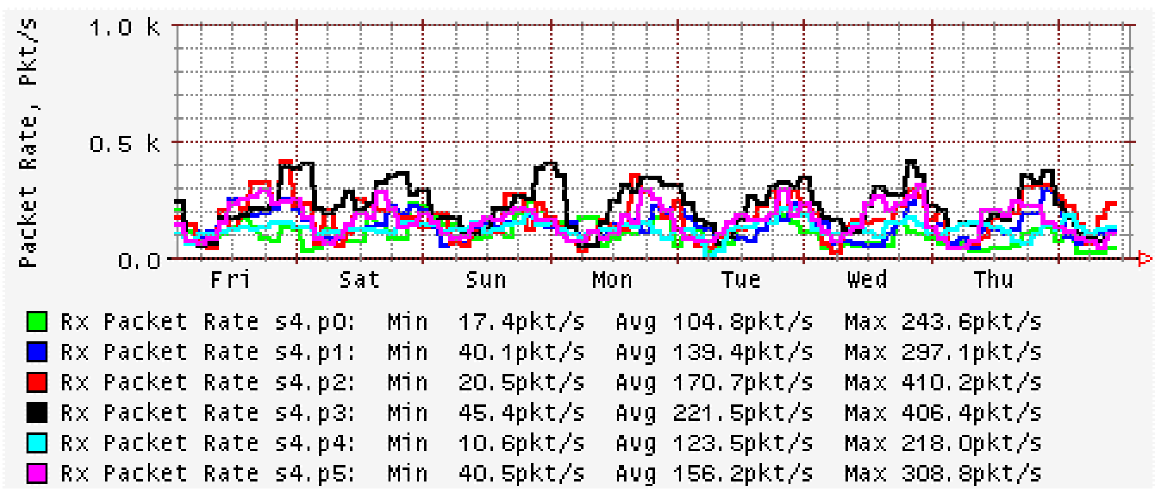

c) BER and FEC Activity

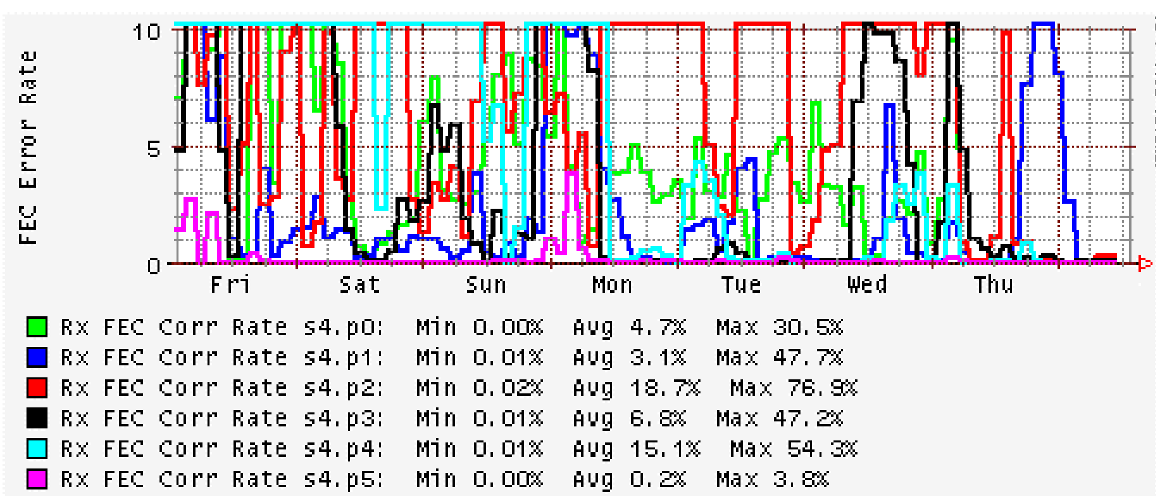

Capacity Exhaustion

Figure 20 shows data rate on the network interface of the CMTS connected to the RDC via a link with 100% utilization for most of the time. Its impact on traffic is visible in flattening the diurnal patterns. After the link upgrade, the typical diurnal patterns were restored.

Applications

Downstream traffic data rate peaks and packet rate peaks coincide with each other almost perfectly. The downstream packet sizes are more uniform (see Table 3). This is not the case in the upstream traffic where data rate peaks do not coincide with packet rate peaks (compare Figures 3 and 4). This may be caused by different packet composition in the upstream direction as well as by high BER and packet retransmission rates.

Miscellaneous

Other factors such as traffic shaping and service tiering as well as data rate throttling may soon start playing a dominant role in influencing traffic rate patterns and behavior. However, they were not included in the analysis presented in this paper.

New applications and more aggressive protocols (for media streaming) may also affect the traffic patterns and behavior. For example, if the media streaming will dominate in the downstream (as opposed to upstream), the asymmetry of data rates and especially packet rates may increase significantly. Other traffic statistics may also be affected (for example, packet sizes and data/packet rates per CM).

Figure 20: Capacity Exhaustion Influence on Diurnal Traffic Patterns

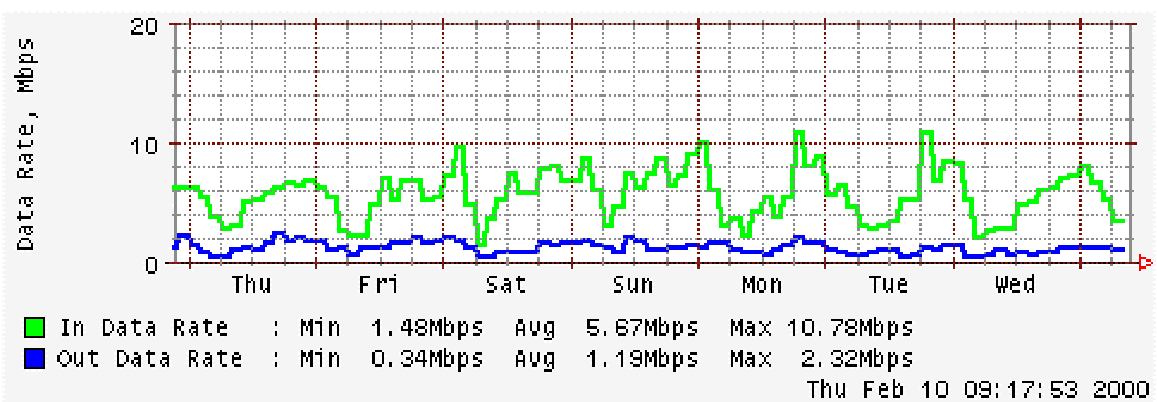

CONCLUSION

The first step to understanding all the elements of the HSD networks in the metro areas is to monitor the traffic behavior in its:

1. CMTS access and network interfaces,
2. Interconnects between CMTS locations and proxy server locations,
3. RDC LANs, and
4. Interconnects to global Internet.

Based on traffic behavior and pattern databases, a simple set of data rate, packet rate and traffic statistics and metrics can be developed for capacity engineering of all the HSD network elements listed above. Moreover, trend monitoring will allow for early adjustments of the statistics and discovery of unusual traffic patterns and behaviors. As an example, this monitoring can and should be used in elimination of high BER.

This paper presented preliminary results of the traffic database analysis from several HSD systems of different sizes. Although further analysis is still required, the paper presented some analysis tools as well as the analysis results. The most useful are:

1. Downstream data rate/CM of 5 kbps for the number of CMs higher than 1,000;
2. Ratio of downstream-to-upstream data rates of 3.2:1 for weakly peak traffic values (not necessarily coincidental in time);
3. Significant multiplexing gains (albeit lower than expected) for upstream traffic after the traffic integration from the access side to the network side of the CMTSs (in excess of 300%) and some gains for the downstream traffic (in excess of 200%).

These results are based on historical data and their application must always be accompanied by verification whether the traffic patterns and behavior assumptions remain the same. The data collected proved that the capacity of the downstream and upstream DOCSIS HFC channels can support up to 5,000 active modems. On the other hand, the results of the analysis on effectiveness of the proxy servers and information caching in traffic containment and traffic load reduction are inconclusive and further analysis on richer traffic parameter databases is warranted.

REFERENCES

1. Dan Pike and Tony E. Werner, "On Plant Renewal Strategies," 1998 NCTA Technical Papers.
2. Harold S. Fluss, "Effective Performance of Shared Bandwidth Data Channels in Hybrid Fiber/Coax Networks," 1999 Emerging Technologies Conference.

THE AUTHORS

David Brown

David Brown, M.TEL., University of Denver, B.S.E.E., University of Colorado at Denver, is a Director for Broadband Engineering at AT&T Broadband (formally TCI Communications). He has over 15 years of experience in distributed communications architecture and engineering. He has held engineering and telecommunications positions at Martin Marietta Corporation, now Lockheed Martin Corporation. He is also a member of IEEE and ACTE.

Wayne Ebel

Wayne T. Ebel is the Director of Network Implementation of High Speed Data at AT&T Broadband (formerly TCI Communications). He has 20 years of telecommunications experience in HFC operations, engineering, and advanced services. Prior to joining TCI, he worked at Pacific Bell where he was involved in the deployment of advanced services for HFC and MMDS. He is actively involved in SCTE as former Chapter Secretary and DOCSIS speaker.

Esteban Sandino

Esteban Sandino, Bachelor in Engineering Physics, McMaster University in Hamilton, Ontario, is a Director of HFC Engineering at AT&T Broadband (formally TCI Communications). He joined Rogers Cablesystems in Toronto, Canada in 1989 and was later involved with CableLabs in the initial effort to characterize two-way cable plant for high-speed data applications. Esteban is also a chairman of the Hybrid Management Sublayer (HMS) subcommittee under the SCTE Engineering committee for the integration and management of HFC outside plant network elements.

Oleh J. Sniezko

Oleh Sniezko, MBA, York University in North York, Canada, B.S.E.E, Wroclaw Polytechnic University, Poland, is the Vice President of Engineering for AT&T Broadband (formerly TCI Communications). Prior to joining AT&T, Oleh worked 8 years at Rogers Engineering in Toronto. For his achievements in implementing fiber optic technology, Oleh received the 1997 Polaris Award sponsored by SCTE, Corning, and CED. Oleh is a licensed Professional Engineer in Ontario and a long-term member of IEEE and SCTE. He has been active in a series of CableLabs and NCTA projects, and represents the cable TV industry in NEC Panel 16.