Multimedia Traffic Engineering

The Bursty Data Model

John T. Chapman
Cisco Distinguished Engineer
jchapman@cisco.com
john@johntchapman.com
408 526-7651

January 8, 2002
(February 14 - editorial updates)

Presented at SCTE Emerging Technologies 2002

Abstract

This paper describes a simple and useful model for modeling a data user on an access network such as an HFC network. The paper uses this model to project how many subscribers can be supported on a particular network such as a DOCSIS network. The paper also shows the same model can be used to describe networked applications and tiers of service, and how to relate those concepts to bandwidth used by a single user.

Introduction

How many cable modems can a CMTS support? This has to be one of the most often asked questions. A popular approach to answering this question is to put as many CMs on an upstream and downstream that will fit, and wait for customer complaints. This works for low data penetrations, but obviously does not work for a serious network.

Is the right answer a fixed number like 1000 CMs per downstream or 250 CMs per upstream [1]? It might be, but only for a given set of circumstances. If one MSO wants to offer twice the data rate to his customers as another MSO, then that 250 number might only be 125 for the other MSO.

The real answer is that it depends. If the bandwidth available at the CMTS and the bandwidth required by the user could both be accurately described, then there should be a simple answer to this question which looks like:

$$\frac{\text{CMTS Bandwidth Available}}{\text{CM Bandwidth Required}} = \text{Number of CMs Supported} \qquad (1)$$

Although this equation is simple, the derivation of the cable modem bandwidth is not. It depends on both the activity level of the user and the applications that are being run or will be run on the network in the future.

To perform this calculation, a model is required for the bandwidth usage of the user behind the cable modem. The applications that run or will be run by the user today generally fit into three distinct categories:

- **Data:** This would include applications such as web traffic, e-mail, file transfers, and low speed audio and video streaming.
- **Voice:** This would specifically cover constant bit rate type traffic which is carried over IP or circuit switching equipment.
- **Video:** This would refer to broadcast quality video, usually MPEG encoded, which is of sufficient bandwidth that usually multiple downstreams per fiber node are required. The transport could be either MPEG-TS or IP.

Models for all three of these traffic types were presented by the author in the white paper *Multimedia Traffic Engineering* (MMTE) [2]. This white paper builds upon this foundation, and focuses specifically on the data model.

Bandwidth Modeling

Modeling traffic is not easy. The process of creating a model for voice traffic has had several advantages. First, the bandwidth of a voice call is constant bandwidth and thus its bandwidth properties are well known. Second, voice has an 80+ year history to it with plenty of real world statistics to rely on for validity. And finally, voice models already existed in the circuit switched world which could be adapted for the IP world. In particular, a model was need to predict the arrival rate of new calls. This was addressed by using the Erlang formulae. When the voice model gets overlaid on an HFC network with the addition of VoIP, compression, and voice activity detection, the model gets more complex but is manageable.

Modeling data traffic is quite different. Models for specific data behaviors such as constant bit rate (CBR) and variable bit rate (VBR), but there is no one commonly used model for the aggregation of several of these data services, combined with some kind of usage pattern by the subscriber. Data characteristics of different applications are diverse and often unknown. As the recent experience with Napster [3] has shown, the dominant data application on a network might not even exist yet while the network is being planned.

Some of the data models that are being researched are very complex. There is work being done using self-similar traffic based upon fractal mathematics [4] as well as a wide array of other mathematical gymnastics. While quite interesting, these models are usually difficult to use and each have their own limitations.

All models are inherently flawed because they are models and are not reality. For example, the Erlang models for voice developed in the early 1920’s don’t account for recent inventions such as answering machines which create a large number of short calls, and data modems which create a large number of very lengthy calls. However, as long as the limitations of a model are understood, models can be an effective way of solving traffic engineering problems.

The Need for a Basic Model

What is really needed is a model that is both simple and useful. The measurement of a simple model would be that the equations:

- must fit on the back of an envelope,
- must be easy to put into a spreadsheet,
- must be easy to use by all.

Approximations are fine as long as they are understood and complications should be avoided.

The measure of usefulness would be that the model:

- must relate to measured parameters,
- must be usable for bandwidth calculations.

A model would not be of much use if it cannot be verified by field data, and if it cannot be actually used for traffic engineering to generate results.

The Role of a Model

Modeling what already exists is interesting, but modeling what does not yet exist is what is most important.

A good model allows current measured results to be combined with theory and growth projections to predict future results. This is illustrated in Figure 1.

The measurement data is good for analyzing current networks and verifying model calculations. However, the real value of modeling is the promise of predicting the future.

It is relatively straight forward to determine if an existing network is working. However, it is the role of the model to predict if the network will continue to work at some future point in time.

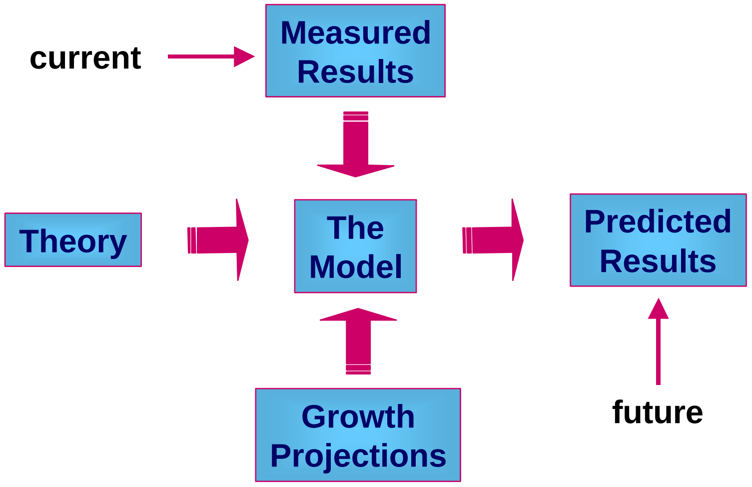

Figure 1: The Role of a Model

The Bursty Data Model

The Bursty Data Model is an attempt to satisfy the above criteria. This model is derived in large part from a behavioral description of what the subscriber and operator see when looking at the network.

The Bursty Data Model defines three scenarios. For each scenario, a measurement interval is assigned and the number of unique users and their bandwidth requirements during that measurement interval is determined. That's about it.

The Three Scenarios

The three scenarios are conveniently called average, peak, and max.

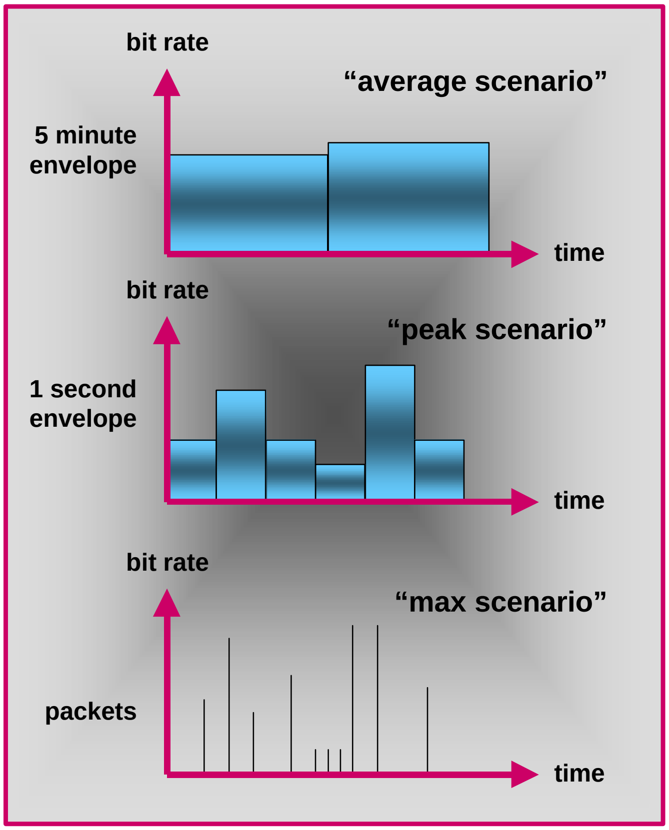

Figure 2: Graph of Avg and Peak Scenarios.

The average scenario is the performance seen by the subscriber over a longer interval of time on a loaded network. The exact time frame is a parameter of the model. It is usually picked to coincide with some counter on the CMTS or on a network analyzer. A typical value might be 5 minutes.

The average case is sometime referred to as the "short packet" scenario because its packet size is shorter than the peak scenario.

The peak scenario is the performance seen by the subscriber over a shorter time on a loaded network. The exact time frame is also a parameter of the model. It is usually picked to relate to some measurable user experience. A typical value may be 1 second.

The average measurement interval of 5 minutes is a ratio of 300x to the peak measurement interval of 1 second.

For example, let's say a downstream channel that has a payload bandwidth of 26.25 Mbps (refer to Appendix A) has a measurement interval of 1 second. How many users would share 26.25 million bits, and how much would each one get? Note that 1 second of 26.25 Mbps would be 2160 packets if each packet was 1518 bytes.

The peak case is sometimes referred to as the "long packet" scenario because the applications which are demanding peak bandwidth such as FTP or Video Streaming typically use large packet sizes. By using a different packet size for the peak calculation, the packets per second (PPS) requirement of the network may be partially relaxed.

The max scenario is the rate seen by the user when the network is not loaded. This is the value that the CMTS uses to rate shape the traffic to a CM. Rate shaping is a DOCSIS feature in which the CMTS regulates the bytes within a DOCSIS service flow with a token bucket algorithm which regulates a peak and average value. When the flow of bytes exceeds the bounds of this equation, the CMTS introduces latency to the packet flow instead of dropping the packet. This latency causes protocols such as TCP to slow down. Thus, the CMTS is able to “shape” the TCP flow without actually causing any data loss.

The average and peak scenarios are the two main scenarios that are used. The max scenario is more for completeness and to predict performance when rate shaping. The measurement interval for the max scenario is equal to the burst time of the token buck rate shaping algorithm, and is usually in the order of milliseconds. The bit rates used for the max scenario is enforced. The bit rates for the peak and average scenarios are observed, but not enforced.

The average, peak, and max scenarios are repeated separately for the upstream and the downstream. Then all six cases vote to see who is the worst case scenario. The worst case scenario then sets the operating limit for the CMTS.

Wall Street Analogy

To give a feel to the differences between the three scenarios, there is an interesting analogy that can be made with Wall Street. The analogy would be:

- The average scenario is the equivalent of quarterly sales
- The peak scenario is the equivalent of weekly sales
- The max scenario is the equivalent of daily sales

Quarterly sales can be predicted reasonable well. Weekly sales have large variations, whereas daily sales can be anything. Yet, a good factory must be able to respond well to daily and weekly fluctuations to be efficient.

And so it is with this model. The average scenario is somewhat predictable and measurable, the peak scenario is more difficult to predict, and the max scenario could be anything. Still, the network must have the headroom to be able to respond to a variety of packet arrival rates and peak rates, just as manufacturing must be able to respond to daily and weekly variations.

Over-Subscription

One of the ultimate goals of traffic engineering is to over subscribe the network aggressively enough to run at high bandwidth efficiency, but with a high enough degree of confidence that the network will still operate with a specified level of performance.

The session density parameters are the inverse of over-subscription. For example, say that during the average measurement interval, the session density is 25%. This means that 25% of the users are using the network. If the network is designed such that they use all of the available bandwidth (assuming the average scenario is the limiting case), then the network is over-subscribed by the inverse of the session density, which is 4x (1/25%). If the peak scenario had a session density of 20%, then the network would be oversubscribed by 20x (1/(25%*20%)) during a 1 second measurement interval.

Thus, the session densities may be chosen with the objective of setting the network over subscription.

The User Data Profile

The combination of the three scenarios and the corresponding packet size produce a user data profile as shown in Figure 3.

Down Stream	Data Profile		
	Bit Rate *kbps*	Pkt Size *bytes*	Pkt Rate *pps*
Max	2000	1518	165
Peak	256	1518	21
Average	80	400	25

Up Stream	Data Profile		
	Bit Rate *kbps*	Pkt Size *bytes*	Pkt Rate *pps*
Max	384	1518	32
Peak	100	1518	8
Average	24	64	47

Underlined values are inputs

Figure 3: The Bursty Data Model User Profile

The most important value is the bandwidth numbers. The packet size is carried along so that a packet per second calculation for the system can be made. Also in DOCSIS, the amount of overhead and thus the efficiency of the channel is related to the packet size, so this value will get used in the calculations. The packet rate is also important for sizing the processing power of the CMTS and the IP equipment upstream of the CMTS.

When these bandwidth numbers are chosen, it is critically important that the measurement interval is kept in mind. This paper suggests a measurement interval of 5 minutes for the average scenario, and 1 second for the peak scenario. Bear in mind that all packet transmissions in the downstream actually transfer at the instantaneous line rate which is 30.34 Mbps for 64QAM, ITU J.83 Annex B. Thus, all three of the bandwidth numbers above are an average over some interval.

For example, let’s say that for the average scenario, the downstream rate is 80 kbps. That means in this model that over a 5 minute interval, that user has received

$$\frac{80,000\ \text{bits}}{1\ \text{sec}} \times \frac{1\ \text{byte}}{8\ \text{bits}} \times \frac{60\ \text{secs}}{1\ \text{min}} \times \frac{5\ \text{min}}{1\ \text{interval}} = 3,000,000\ \text{bytes per 5 min interval}$$

Those bytes may have come in multiple bursts, but over a 5 minute interval, it was 3,000,000 bytes which equates to 80 kbps.

Carving up the Households Passed.

Before proceeding, let’s review how these three scenarios relate to the overall number of households passed (HHP). Figure 4 shows a graphic representation of how these variables relate to each other, along with an example for a 2,000 HHP fiber node. The variables included are the absolute market penetration of cable subscribers (MP_c) and data subscribers (MP_d), and the relative session density for the average (SD_a), peak (SD_p), and max scenarios (SD_m).

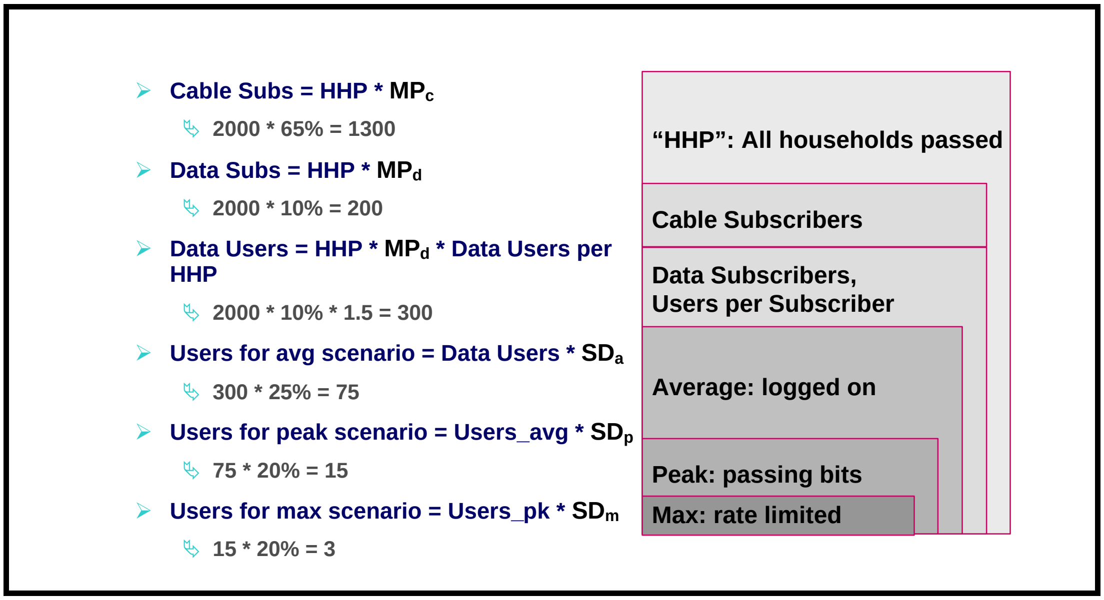

- Cable Subs = HHP * MP_c
  - 2000 * 65% = 1300
- Data Subs = HHP * MP_d
  - 2000 * 10% = 200
- Data Users = HHP * MP_d * Data Users per HHP
  - 2000 * 10% * 1.5 = 300
- Users for avg scenario = Data Users * SD_a
  - 300 * 25% = 75
- Users for peak scenario = Users_avg * SD_p
  - 75 * 20% = 15
- Users for max scenario = Users_pk * SD_m
  - 15 * 20% = 3

Figure 4: Carving up HHP

Of all the HHP, only a percentage of those are cable subscribers. A typical number is 65%. These subscribers take different services, such as data, voice, and/or video. The market penetration of these services is usually quoted with respect to HHP. For example, the market penetration from data may be 10% of HHP.

This model recognizes that there is usually more than one user located at a household. Since the word subscriber is usually associated with a paying subscriber, the word user will be used to denote the actual data user. If the model applies to the user, then the number of users per HHP must be specified. If, however, the data model is intended to be a service agreement for a single subscriber, regardless of the number of users at that location, then this variable should be set to 1.

The percent of users for the average, peak, and max scenarios in this model are all relative percentages to each other as this is a mentally easier to manage. The percentages for these three scenarios are referred to as session densities.

Session Density

Session Density	
Relative %	Direct %
20% of peak	1%
20% of avg	5%
25% of users	25%

Figure 5: Session Density

In the previous section, three scenarios were presented which described how much bandwidth a single user used. To finish the model inputs, we now need to know how many users are using the network for each of these three scenarios. This is referred to as the session density, and is easily the most subjective part of the model.

Recall that these are scenarios. Do not get hung up on the names that have been given to them (average and peak). What is more important is the length of the measurement intervals.

So, the question to answer is “How many users will be active during the measurement interval”. For example, if the measurement interval for the average interval is 5 minutes, than how many users will send bits according to their user data profile in that 5 minutes? For the peak scenario, how many users will be transferring data, according to their user data profile, during that 1 second interval?

If the measurement interval for the average scenario was 60 minutes instead of 5 minutes, then it is obvious that many more users would have sent data, and thus the percentage would be higher. If the measurement interval was 12 hours, then the percentage could easily be 80%. However, the measurement intervals are intentionally limited, and meant to line up with the user data profile.

Let’s look at an example. If there were 100 users, and 20 of them were active (20%), but they only transferred data every 10 minutes, then the session density for the average scenario (with a 5 minutes measurement interval) would be 10%, not 20%.

It is the session density numbers of the MMTE Bursty Data Model that allow the traffic engineer to compensate for the statistical (or just plain random) nature of the network. There are two approaches of choosing these number, behavioral and measured.

Using an Intuitive Approach

With this method, the number is chosen based upon the performance the operator wants his customer to see. This is best illustrated by example. Let’s take the downstream peak scenario.The downstream payload rate is 26.25 Mbps (refer to Appendix A) and the peak measurement interval is 1 second.

The first way is to just “wing it”. By choosing a number like 20% of average, that is stating that within the 5 minute measurement interval, the worst case scenario for a the 1 second interval is that 20% of those users will try and use the network. That does not mean that this has to be true for every 1 second internal, just for the scenario in which you want to guarantee the bandwidth specified in the user data profile.

Another way is to get more down and dirty with the numbers. That means during the 1 second measurement interval, there is 26.25 Mbits or 3.28 MBytes that can be appropriated between some number of users. If we further assume a packet size during this interval of 1518 bytes, allow for some DOCSIS framing overhead, use the bit rate from the user data profile (256 kbps), and divide this into the payload rate, we find that 93 users can access the network. This is then divided by the number of users on the network to get the percentage. This technique is really just reverse engineering the calculations that we are going to go through in the next section. In other words, you can start with the answer that you like, and work backwards to get the right inputs.

The important point is that these percentage numbers should have an intuitive feel to them.

Choosing Session Density Numbers from Field Measurements

The following information can be supplied when monitoring IP packets on the network during peak and average measurement intervals for both upstream and downstream:

- Mix of applications by examining the TCP/UDP port number
- Number of users by looking for unique IP DA addresses in the downstream and unique IP SA addresses in the upstream
- Size of packet per application

These measurements can be taken over the two measurement intervals of 1 second and 5 minutes. From these measurements, the following can be calculated for each scenario:

- Bandwidth per user.
- Nominal packet size per user
- Nominal PPS per user.

These measurements provide the bases for the Bursty Data Model.

Wrap Up

The Burst Data Model is intentionally simple. The input parameters for the model can be derived either intuitively, or through measurement. Thus, there is a natural method for correlating theoretical models with measured results. This meets the fundamental requirement for any practical model.

Later in this paper, we will look at how these bandwidth numbers can be related in theory to real world applications and to multiple service tiers. But first, let’s see how this profile can be used to calculate the number of cable modems and households passed a CMTS can support.

The Bursty Data Model Calculations

The following three sections explain the process of the calculations. The formulas which are used for these calculations are listed in the section after that. The actual numbers used are for the sake of example only. These numbers will not only vary with each MSO, but also vary drastically over time as new applications are introduced into the network. What is more important is the methodology used.

Downstream Calculations

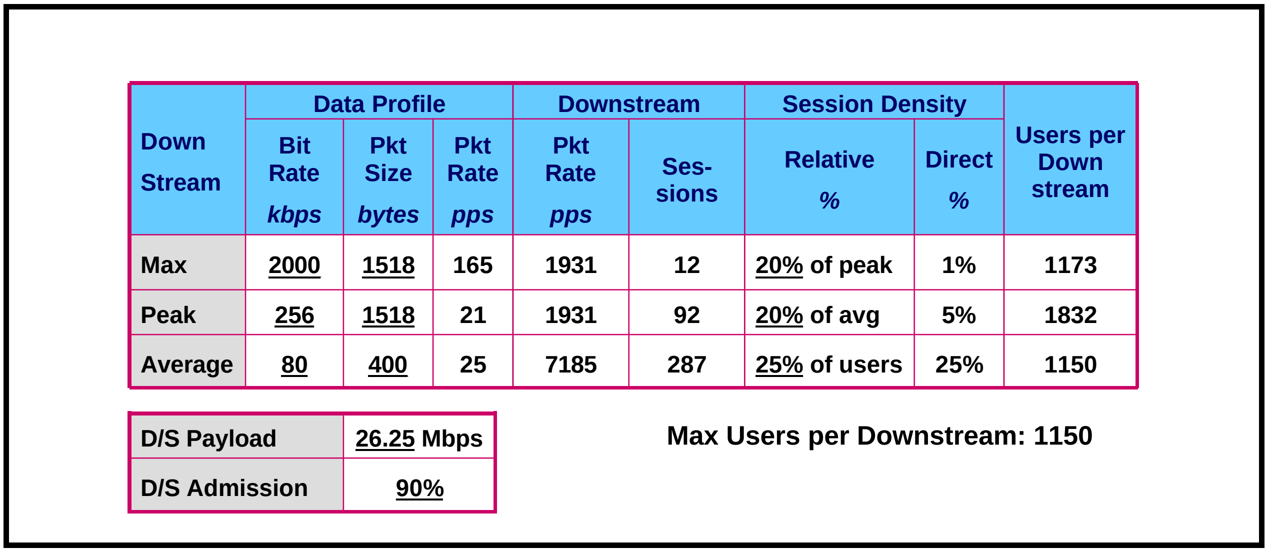

Down Stream	Data Profile			Downstream		Session Density		Users per Down stream
	Bit Rate *kbps*	Pkt Size *bytes*	Pkt Rate *pps*	Pkt Rate *pps*	Ses-sions	Relative *%*	Direct *%*	
Max	2000	1518	165	1931	12	20% of peak	1%	1173
Peak	256	1518	21	1931	92	20% of avg	5%	1832
Average	80	400	25	7185	287	25% of users	25%	1150

D/S Payload	26.25 Mbps
D/S Admission	90%

Max Users per Downstream: 1150

Figure 6: Downstream Calculations

The first inputs to the calculation come from our user data profile. The Max scenario is the rate limiting scenario, while the Peak and Average scenarios are intended to reflect a mixture of applications run by the user. For each scenario, the inferred packet rate is found by dividing the bit rate by the packet size.

Two new variables are introduced. The first is the downstream bit rate derived in Appendix A. The second is an admission control level. The downstream bandwidth must be shared between various models for data, voice, and video, as well as signaling. This variable permits the traffic engineer to carve out bandwidth for the data model. In this scenario, 90% of the bandwidth is being used for data. The remaining 10% is reserved for signaling, and there is no VoIP or high bandwidth video over IP on the system.

The packet per second rate can be calculated by dividing the packet size into the de-rated channel bandwidth, or by multiplying the packet rate per user by the number of sessions. In this calculation, the DOCSIS overhead per packet is added in.

The session column represents how many simultaneous sessions can be supported in the downstream. This value is obtained by dividing the channel packet capacity by the user's packet rate requirement.

The next part of the model takes the session density numbers which are relative to each other and multiplies them out to direct percentages. The number of users for each scenario is just the simultaneous sessions allowed divided by this direct percentage.

The maximum number of users is the lowest value from these three scenarios. Note that the number calculated is the number of users, not subscribers. From earlier, it was noted that although there may be one HHP per paying subscriber, there may be multiple users per HHP.

Upstream Calculations

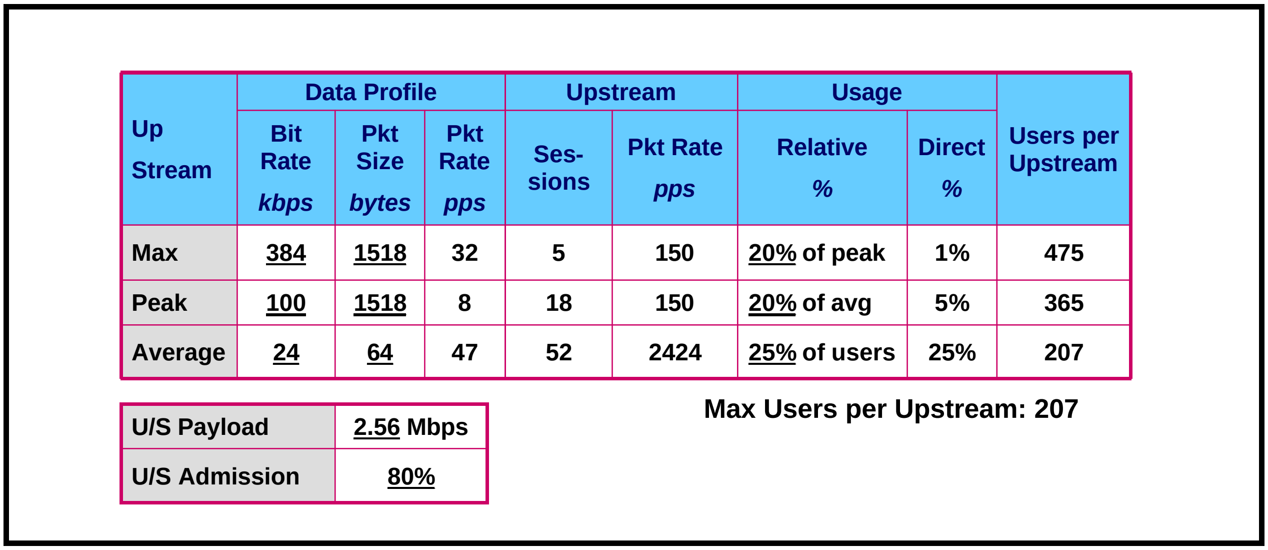

Up Stream	Data Profile			Upstream		Usage		Users per Upstream
	Bit Rate *kbps*	Pkt Size *bytes*	Pkt Rate *pps*	Sessions	Pkt Rate *pps*	Relative *%*	Direct *%*	
Max	384	1518	32	5	150	20% of peak	1%	475
Peak	100	1518	8	18	150	20% of avg	5%	365
Average	24	64	47	52	2424	25% of users	25%	207

U/S Payload	2.56 Mbps
U/S Admission	80%

Max Users per Upstream: 207

Figure 7: Upstream Calculations

The upstream calculations are identical in nature to the downstream calculations. All three scenarios are individually calculated, and the lowest result of the three is taken as the final answer.

Combining the Results

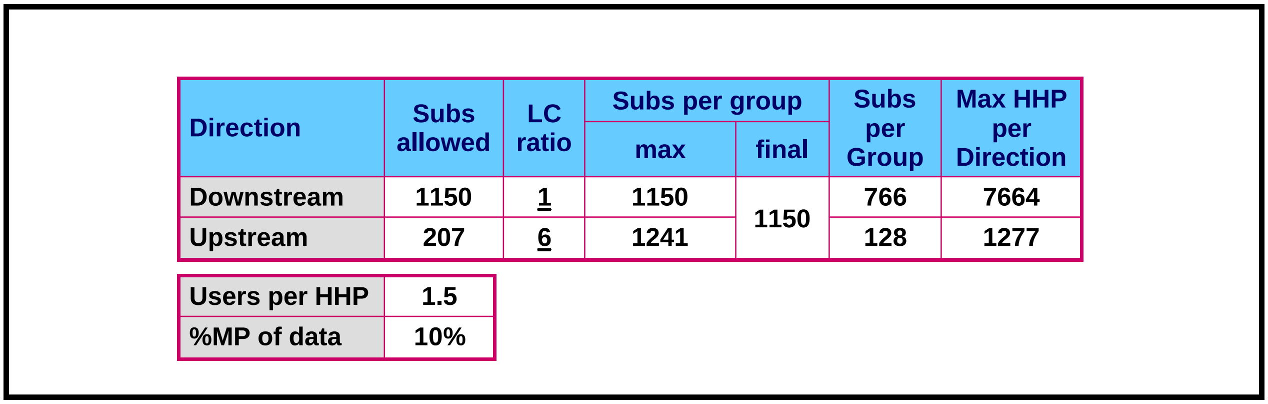

Direction	Subs allowed	LC ratio	Subs per group		Subs per Group	Max HHP per Direction
			max	final		
Downstream	1150	1	1150	1150	766	7664
Upstream	207	6	1241		128	1277

Users per HHP	1.5
%MP of data	10%

Figure 8: Combined Calculations

In DOCSIS, downstreams and upstreams come in domains. A typical domain might be 1 downstream and 6 upstreams. In this final calculation, the number of users that the downstream and upstream will support is multiplied by the domain ratio. The lowest of the two results sets the user limit. This user limit is then divided by users per HHP to get the number of subscribers that can be supported. That number is then divided by the market penetration of data to determine the number of HHP that the DOCSIS domain can support.

An alternative scenario is to determine the maximum data market penetration allowed by dividing the number of subscribers allowed by the number of HHP.

The Math

Here is the math that was used for the downstream, upstream, and combined scenarios. The downstream average scenario is used as the example calculation.

The following set of equations were used to perform the downstream and upstream calculations. The expression F(DC_user_pkt_bytes) adds in per packet overhead and is explained in Appendix B. This per packet overhead is different for the upstream and downstream.

For each direction (D = d/s, u/s) and for each case (C = Avg, Peak, Max):

$$\mathbf{DC\_user\_pps} = \frac{\mathbf{DC\_user\_kbps} \times \mathbf{1000}}{\mathbf{DC\_user\_pkt\_bytes} \times \mathbf{8}} \tag{2}$$

$$ds\_avg\_user\_pps = \frac{80 \times 1000}{400 \times 8} = 25\ pps$$

$$\mathbf{DC\_pps} = \frac{\mathbf{D\_mbps} \times \mathbf{D\_adm\_level} \times \mathbf{1,000,000}}{\mathbf{F(DC\_user\_pkt\_bytes)} \times \mathbf{8}} \tag{3}$$

$$ds\_avg\_pps = \frac{26.25 \times 90\% \times 1,000,000}{F(400) \times 8} = 7,185\ pps$$

$$\mathbf{DC\_sessions} = \frac{\mathbf{DC\_pps}}{\mathbf{DC\_user\_pps}} \tag{4}$$

$$ds\_avg\_sessions = \frac{7,185}{25} = 287$$

$$\mathbf{DC\_users} = \frac{\mathbf{DC\_sessions}}{\mathbf{DC\_session\_density}} \tag{5}$$

$$ds\_avg\_users = \frac{287}{25\%} = 1150\ users$$

$$\mathbf{D\_users} = \mathbf{MIN(\ D\_avg\_users, D\_pk\_users, D\_max\_users)} \tag{6}$$

$$ds\_users = MIN(\ 1150, 1832, 1173) = 1150\ users$$

The following set of equations were used to perform the combined downstream and upstream calculations.

For each direction (D = d/s, u/s)

$$\mathbf{D\_users\_total = D\_users \times domain\_ratio} \tag{7}$$

$$ds\_users\_total = 1150 \times 1 = 1150 \text{ users}$$

$$us\_users\_total = 207 \times 6 = 1242 \text{ users}$$

$$\mathbf{users\_per\_domain = MIN(ds\_users\_total, us\_users\_total)} \tag{8}$$

$$users\_per\_domain = MIN(1150, 1241) = 1150 \text{ users}$$

$$\mathbf{D\_subs = \frac{users\_per\_domain}{users\_per\_hhp \times D\_domain\_ratio}} \tag{9}$$

$$ds\_subs = \frac{1150}{1.5 \times 1} = 767 \text{ subs per d/s}$$

$$us\_subs = \frac{1150}{1.5 \times 6} = 127.7 \text{ subs per u/s}$$

$$\mathbf{D\_hhp = \frac{D\_subs}{mp\_data}} \tag{10}$$

$$us\_hhp = \frac{127.7}{10\%} = 1277 \text{ hhp per u/s}$$

In this example, it is the downstream that is limiting the system. Changing the user data profile or the admission levels on the CMTS can change it into an upstream limited system. Out of the six scenarios, one scenario will always set the operating limit.

If the user data profile in this example was intended to be the subscriber profile, so that the variable users_per_hhp was 1, then the resulting HHP per upstream would of resulted in more HHP per upstream.

An alternative and useful scenario is to calculate the maximum allowable market penetration if the size of the fiber nodes is known. Assuming 2000 HHP per upstream,

$$\mathbf{mp\_data = \frac{D\_subs}{D\_hhp}} \tag{11}$$

$$mp\_data = \frac{128}{2000} = 6.4\%$$

Generation of the Bursty Data Model User Profile

The Bursty Data Model user profile has allowed us to describe a user and calculate how many users can fit onto a CMTS. We have also discussed how to generate this profile both intuitively and from field measurements. To meet the requirement of being useful, the bursty data model must relate somehow to the applications, such as e-mail and web traffic, that the user will use. This would be for both existing applications, and for new applications that may arise.

Accommodating Applications

Here is the concept. If the user data profile can reasonably describe a user, then it should be able to describe a single application as well. If each applications uses the same model but with different parameters, the models can be easily combined to generate the user data profile.

When combining the separate application profiles, they must be weighted. For example, if e-mail is used 10% of the time and web traffic is used 90% of the time, then the e-mail profile is multiplied by 10%, the web traffic profile is multiplied by 90%, and the results are added. This concept is illustrated in Figure 9.

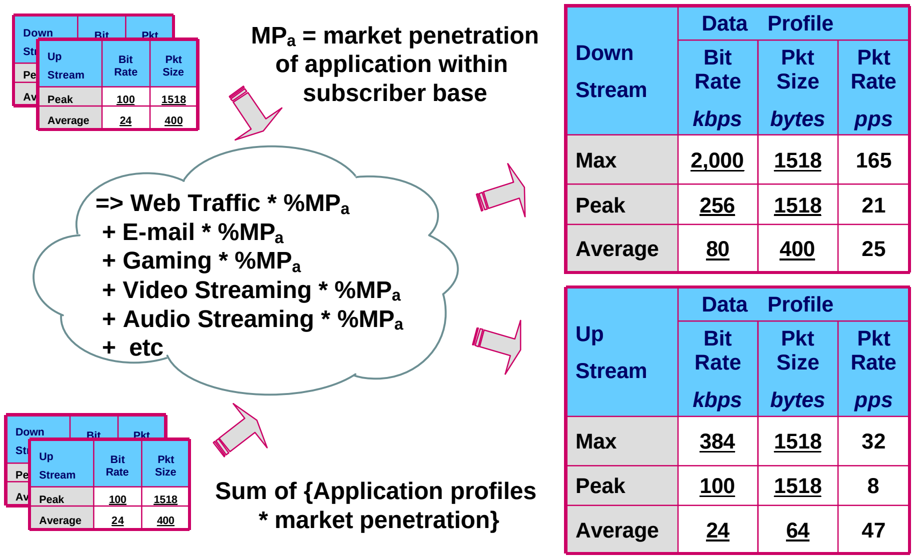

MP_a = market penetration of application within subscriber base

=> Web Traffic * $\%MP_a$
+ E-mail * $\%MP_a$
+ Gaming * $\%MP_a$
+ Video Streaming * $\%MP_a$
+ Audio Streaming * $\%MP_a$
+ etc

Sum of {Application profiles * market penetration}

Down Stream	Data Profile		
	Bit Rate *kbps*	Pkt Size *bytes*	Pkt Rate *pps*
Max	2,000	1518	165
Peak	256	1518	21
Average	80	400	25

Up Stream	Data Profile		
	Bit Rate *kbps*	Pkt Size *bytes*	Pkt Rate *pps*
Max	384	1518	32
Peak	100	1518	8
Average	24	64	47

Figure 9: Generation of the User Data Profile from Application Profiles

The application profiles provide have average and peak scenarios into the user data profile. The max scenario in the user data profile is taken from the rate limiting for the CM that is programmed into the CMTS.

The weighting factor [MP_a] is applied to each application. If more than one application is running at once in a user environment that consumes network bandwidth, then the sum of all the MP_a variables may exceed 100%. A good value is probably somewhere between 100% and 200%, indicating that across the user base, the nominal number of network bandwidth consuming applications that are active is between 1 and 2.

Accommodating Tiers of Service

There is a good potential market for providing different tiers of service to a customer. For example, three levels may be used which are called silver, gold, and platinum. Silver might provide low bandwidth services such as e-mail and basic web traffic, while platinum may have higher bandwidth for video streaming and gaming. How are these tiers of services accommodated by the model?

First, a tier profile is constructed using a weighted average of applications, just as was done before. Then, the three tier profiles are combined with a weighted average, MP_t, into one user data profile. This concept is illustrated in Figure 10.

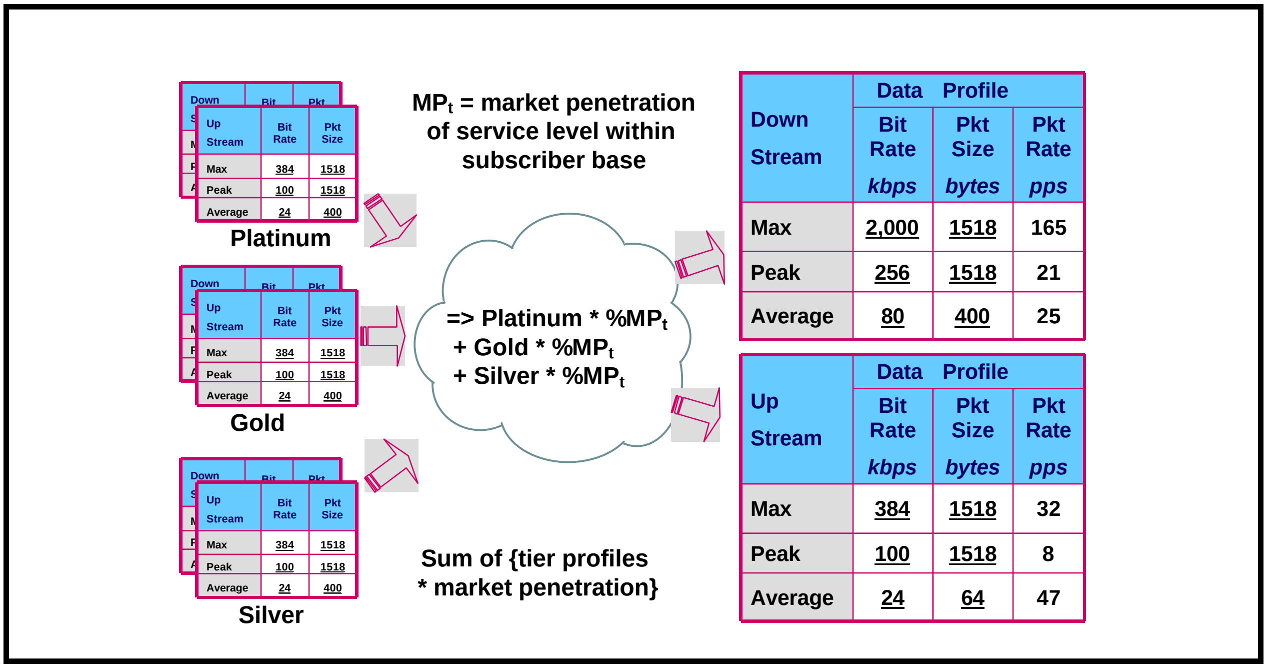

Up Stream	Bit Rate	Pkt Size
Max	384	1518
Peak	100	1518
Average	24	400

MP_t = market penetration of service level within subscriber base

=> Platinum * %MP_t
+ Gold * %MP_t
+ Silver * %MP_t

Sum of {tier profiles * market penetration}

Down Stream	Data Profile		
	Bit Rate *kbps*	Pkt Size *bytes*	Pkt Rate *pps*
Max	2,000	1518	165
Peak	256	1518	21
Average	80	400	25

Up Stream	Data Profile		
	Bit Rate *kbps*	Pkt Size *bytes*	Pkt Rate *pps*
Max	384	1518	32
Peak	100	1518	8
Average	24	64	47

Figure 10: Combining Multiple Tiers of Service

Note that in this case, the max level is defined in each tier and combined into the user data profile since rate shaping is a per cable modem parameter.

Putting it All Together

Figure 11 shows a sample calculation of combining multiple applications with multiple service tiers to achieve a single user data profile. The actual numbers used are for illustrative purposes only. It is the methodology that is significant.

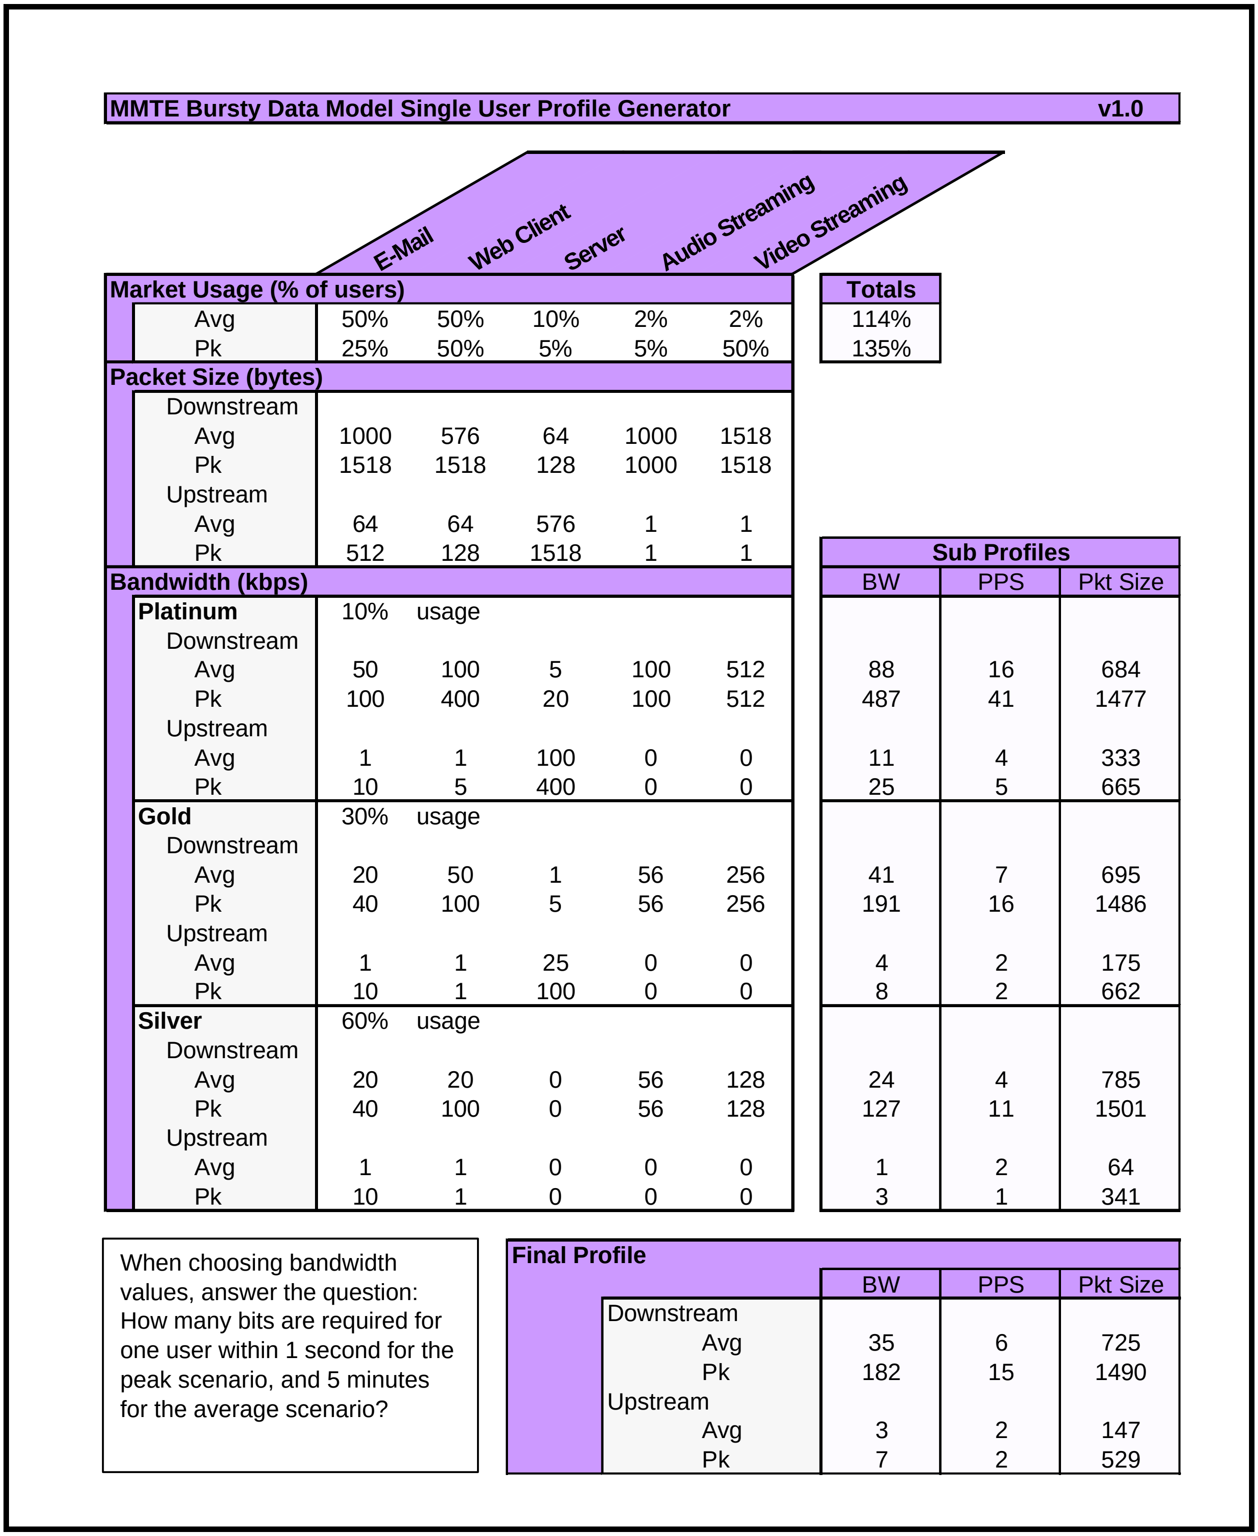

MMTE Bursty Data Model Single User Profile Generator **v1.0**

Market Usage (% of users)	E-Mail	Web Client	Server	Audio Streaming	Video Streaming	Totals
Avg	50%	50%	10%	2%	2%	114%
Pk	25%	50%	5%	5%	50%	135%

Packet Size (bytes)	E-Mail	Web Client	Server	Audio Streaming	Video Streaming
Downstream					
Avg	1000	576	64	1000	1518
Pk	1518	1518	128	1000	1518
Upstream					
Avg	64	64	576	1	1
Pk	512	128	1518	1	1

Bandwidth (kbps)	E-Mail	Web Client	Server	Audio Streaming	Video Streaming	Sub Profiles BW	Sub Profiles PPS	Sub Profiles Pkt Size
Platinum	10%	usage						
Downstream								
Avg	50	100	5	100	512	88	16	684
Pk	100	400	20	100	512	487	41	1477
Upstream								
Avg	1	1	100	0	0	11	4	333
Pk	10	5	400	0	0	25	5	665
Gold	30%	usage						
Downstream								
Avg	20	50	1	56	256	41	7	695
Pk	40	100	5	56	256	191	16	1486
Upstream								
Avg	1	1	25	0	0	4	2	175
Pk	10	1	100	0	0	8	2	662
Silver	60%	usage						
Downstream								
Avg	20	20	0	56	128	24	4	785
Pk	40	100	0	56	128	127	11	1501
Upstream								
Avg	1	1	0	0	0	1	2	64
Pk	10	1	0	0	0	3	1	341

When choosing bandwidth values, answer the question: How many bits are required for one user within 1 second for the peak scenario, and 5 minutes for the average scenario?

Final Profile	BW	PPS	Pkt Size
Downstream			
Avg	35	6	725
Pk	182	15	1490
Upstream			
Avg	3	2	147
Pk	7	2	529

Figure 11: User Data Profile Generator from Applications and Service Tiers

The Traffic Barometer

There is another, more basic traffic model that does not have the ability to relate to applications like the MMTE Bursty Data Model, but it does a reasonable job of indicating growth. This model is referred to as the MMTE Traffic Barometer.

In general, it is very easy for a CMTS to measure:

- the number of CMs per upstream & downstream
- the average bandwidth in terms of Mbps and PPS per u/s, d/s, and WAN port.

With these numbers, usually which are available through SNMP, the following statistics may be calculated for the upstream and the downstream:

- Average Mbps per CM: *for example, 20 kbps*
- Average PPS per CM: *for example, 5 PPS*
- Average packet size: *for example, 400 bytes*

These metrics may be tracked as the network grows, and used to measure trends and predict new growth. Other trends such as the number of CM per upstream over time are also useful trends to follow. MRTG, the Multi Router Traffic Grapher [5], is a popular graphing program available on the internet, and is often used for plotting these trends.

Note that the Traffic Barometer is useful for measuring trends such as growth, but due to its low sampling, it is not granular enough to pick up on traffic peaks. Although a good rule of thumb, it should be combined with theory and the MMTE Bursty Data Model to correctly predict new traffic patterns.

The Complete Profile

Bandwidth is one aspect of the user data profile. A complete profile must take into consideration:

- Bandwidth
- Latency
- Jitter
- Allowable Packet Loss

The MMTE Bursty Data Model addresses the first issue of bandwidth. The last three issues are generally handled with QOS policies.

Conclusions

Networks which carry data, voice, and video must be engineered if they are to operate properly. Multimedia Traffic Engineering is a series of techniques and models for engineering data, voice, and video over an IP network. The Bursty Data Model is a particular technique for engineering a network for data.

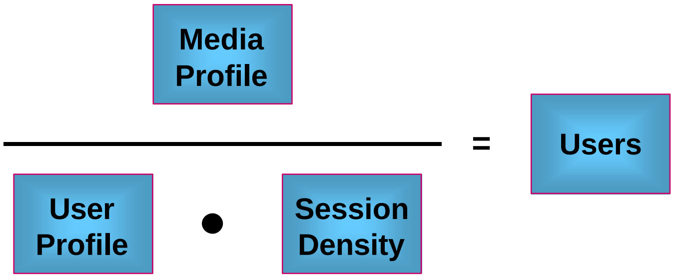

Figure 12: The Bursty Data Model

Figure 12 shows a high level view of the parameters of the Bursty Data Model. The media profile refers to the DOCSIS upstream and downstream and relatively constant. If the user data profile and session densities are known, then the number of users a network will support can be determined.

The model may also be used in reverse. If the number of users and the user data profile is known, the session densities may be calculated. Likewise, if the number of users is known and session density is known, the user data profile may be calculated.

There is certainly room for other models, either simpler or more complex. However, to be useful, they should incorporate the same objectives. The model will:

- have parameters that can be determined through intuition, calculation, or measurement
- relate to all traffic types and accommodates service tiers
- be used to calculate the network loading and the number of users that can be supported.

The MMTE Bursty Data Model provides a simple and useful method for establishing a profile for a data user. This profile is intuitive, it can be calculated, and it can be measured. Further, the profile can be used to describe both applications and tiers of service. These application profiles and tier profiles can be combined to generate a user data profile. This user data profile can then be divided into the CMTS bandwidth to provide the number of CMs per CMTS, thus solving equation (1).

Appendix A: DOCSIS Bandwidth

For the bandwidth model, we need to convert between the actual payload of the channel and the raw bit rate of the channel. The difference between these two bit rates is overhead.

Overhead Defined

There are two types of overhead, overhead per channel, and overhead per packet. It is important for accuracy that these two types of overhead get treated separately in the calculations. Figure 13 shows the overhead for the downstream and Figure 14 shows the overhead for the upstream.

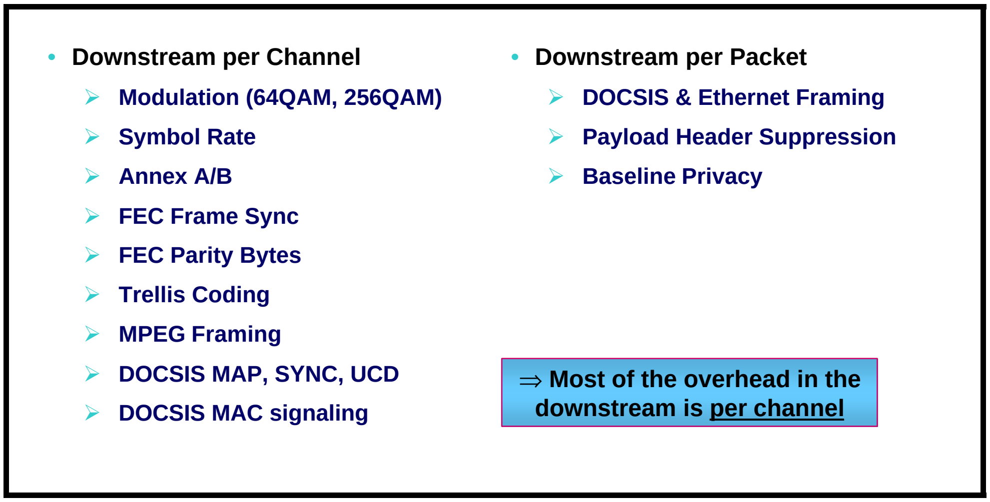

- Downstream per Channel
  - Modulation (64QAM, 256QAM)
  - Symbol Rate
  - Annex A/B
  - FEC Frame Sync
  - FEC Parity Bytes
  - Trellis Coding
  - MPEG Framing
  - DOCSIS MAP, SYNC, UCD
  - DOCSIS MAC signaling
- Downstream per Packet
  - DOCSIS & Ethernet Framing
  - Payload Header Suppression
  - Baseline Privacy

⇒ Most of the overhead in the downstream is per channel

Figure 13: DOCSIS Downstream Bandwidth Parameters

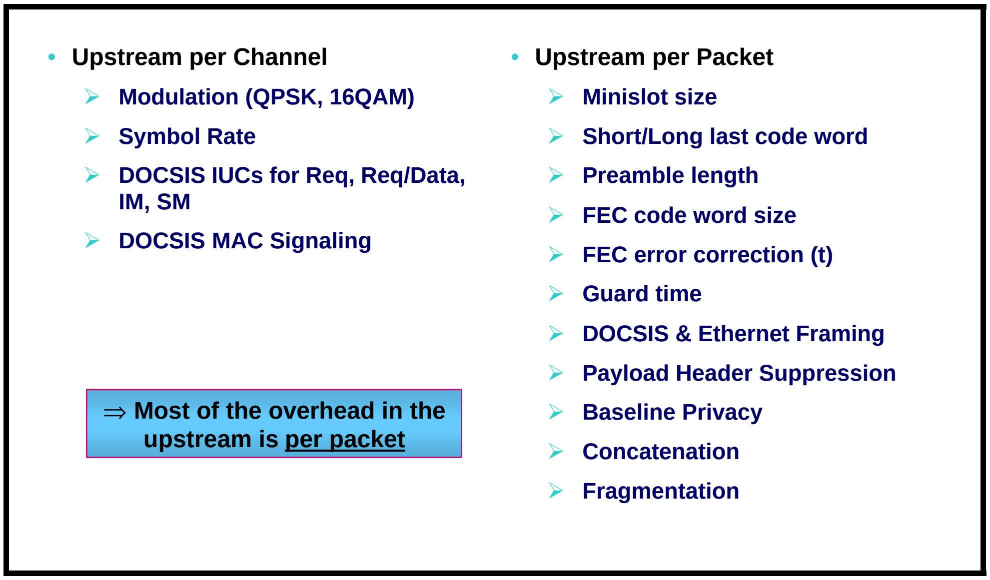

- Upstream per Channel
  - Modulation (QPSK, 16QAM)
  - Symbol Rate
  - DOCSIS IUCs for Req, Req/Data, IM, SM
  - DOCSIS MAC Signaling
- Upstream per Packet
  - Minislot size
  - Short/Long last code word
  - Preamble length
  - FEC code word size
  - FEC error correction (t)
  - Guard time
  - DOCSIS & Ethernet Framing
  - Payload Header Suppression
  - Baseline Privacy
  - Concatenation
  - Fragmentation

⇒ Most of the overhead in the upstream is per packet

Figure 14: DOCSIS Upstream Bandwidth Parameters

Note that in the downstream, most of the overhead is per channel, whereas in the upstream, most of the overhead is per packet. That has the interesting result that the numbers quoted for downstream bandwidth tend to be closer to the payload bandwidth, while the bandwidth numbers that get quoted for the upstream tend to be closer to the raw bandwidth. While this means that the upstream and downstream channels get quoted differently, to do anything else would mean introducing assumptions on packet rates and sizes, and would thus introduce inaccuracies.

Figure 15 lists the additional overhead which enters into VoIP calculations.

- **VoIP Parameters**
  - **CODEC type (G.711, G.729, …)**
  - **Packet length (10ms, 20 ms)**
  - **Voice Activity Detection**
  - **Grade of Service (1%)**
- **Customer Parameters**
  - **Lines per HHP**
  - **Avg calls per hour**
  - **Avg length of call**
- **CMTS Parameters**
  - **Static versus dynamic load balancing**
  - **Admission control limits**
  - **HHP per upstream**
  - **Ratio of upstreams to downstreams**

Figure 15: VoIP Bandwidth Parameters

Per Channel Calculations

Figure 16 shows the calculations for the per channel overhead for the downstream for 64 QAM and 256 QAM and for J.83 Annex A and Annex B.

Country	**USA**		**Europe**	
J.83 Annex	**Annex B**	**Annex B**	**Annex A**	**Annex A**
Bandwidth (MHz)	6	6	8	8
Constellation size	64	256	64	256
Symbol Rate (Msps)	5.056941	5.360537	6.952	6.952
Alpha	0.186	0.119	0.151	0.151
Bits per Symbol	6	8	6	8
FEC Frame Sync	0.08%	0.05%	0.00%	0.00%
FEC Parity Bytes	4.69%	4.69%	7.84%	7.84%
Trellis Coding Overhead	6.67%	5.00%	0.00%	0.00%
MPEG Header	2.13%	2.13%	2.13%	2.13%
MPEG Pointer Byte	0.54%	0.54%	0.54%	0.54%
PHY Layer Raw BW	30.34	42.88	41.71	55.62
PHY Overhead	13.5%	11.9%	10.3%	10.3%
PDU Layer BW (Mbps)	**26.25**	**37.78**	**37.42**	**49.89**

Figure 16: DOCSIS Downstream Channel Overhead

The overhead in Figure 16 was calculated assuming:

- FEC Frame sync is Annex B only. For 64 QAM, it is 42 bits out of 53,802 bits, and for 256 QAM, it is 40 bits out of 78,888 bits.
- FEC Parity Bytes are 6 bytes out of 128 bytes for Annex B, and 16 bytes out of 204 bytes for Annex A.
- Trellis Coding is Annex B only. For 64 QAM, it is 1 bit out of 15 bits, and for 256 QAM, it is 1 bit out of 20 bits.
- MPEG Header is 4 bytes out of 188 bytes
- MPEG Pointer is 1 byte out of 184 bytes. This is a worst case number and true for packets less than 184 bytes. For 1518 byte packets, 1 byte is added out of 184 * 9 = 1656 bytes, which is 0.06%. However, MAC messages such as MAPs bring the number back up.

Note that the individual PHY overheads are multiplied together to arrive the final channel overhead. The formula that was used was:

$$\begin{aligned}&= 1 - (\\&(1 - ds\_oh\_fec\_sync) *\\&(1 - ds\_oh\_fec\_parity) *\\&(1 - ds\_oh\_trellis) *\\&(1 - ds\_oh\_mpeg\_hdr) *\\&(1 - ds\_oh\_mpeg\_pointer)\\&)\end{aligned} \tag{12}$$

Figure 17 shows the calculations for the per channel overhead for the upstream for DOCSIS 1.0/1.1 and the advanced PHY of DOCSIS 2.0.

RF Bandwidth (MHz)	**1.6**	**3.2**	**3.2**	**6.4**
Alpha	**0.25**	**0.25**	**0.25**	**0.25**
Constellation Size	4	4	16	64
Symbol Rate (ksps)	1280	2560	2560	5120
Bits per Symbol	2	2	4	6
PHY layer BW (Mbps)	**2.56**	**5.12**	**10.24**	**30.72**

Figure 17: DOCSIS Upstream Channel Overhead

Note that all of the equivalent types of overhead that were listed as per channel for the downstream tend to be per packet for the upstream and do not show up here. The reason for this is that the transmission of the DOCSIS downstream is continuous, while the DOCSIS upstream is bursty, where each burst is a DOCSIS frame.

Per Packet Overhead

Figure 18 shows the overhead per packet in DOCSIS.

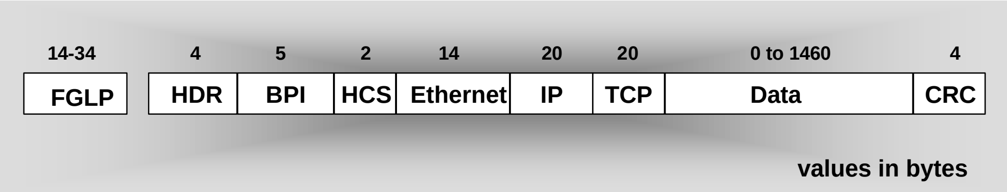

Figure 18: DOCSIS Per Packet Overhead

The expression FGLP refers to FEC, Guard Time, Last Codeword, and Preamble. These values are unique to the upstream. The other fields are the DOCSIS header (HDR) and header checksum (HCS), baseline privacy (BPI). The bytes values shown are typical values.

To calculate the payload capacity of the DOCSIS channel, a boundary must be choosen between those bytes in Figure 18 which are overhead bytes, and those bytes which are payload bytes. The choices for this boundary include starting at the DOCSIS header, Ethernet header, TCP/IP header, or the data payload. Anyone of these choices are valid. The choice that this paper is going to make is at the beginning of the Ethernet frame since the CM delivers Ethernet frames to the user. The second choice would of been the beginning of the TCP/IP frame.

To make reading of the formulas convenient, the expressions F(*payload_bytes*) will be used to denote the conversion of payload bytes to a full DOCSIS frame with all overhead.

For the downstream, most of the variables which influence bandwidth have been taken care of with the Per Channel calculations. For an approximation, a good assumption would be:

- a DOCSIS header of 6 bytes and a BPI extended header of 5 bytes.
- no other extended headers.

The formula is simple in this example. 11 bytes need to be added to cover the DOCSIS header and BPI.

$$\mathbf{F(\ ds\_payload\_bytes\ )\ =\ ds\_payload\_bytes + 11} \tag{13}$$

In the upstream, all of the overhead is variable in size and should be set according to the CMTS configuration. Also, the FEC overhead adds a percentage overhead instead of a fixed number of bytes. For an approximation, one case would be:

- a DOCSIS header of 6 bytes and a BPI extended header of 5 bytes.
- no other extended headers
- no concatenation,
- no fragmentation,
- no payload header suppression
- minislots are 8 bytes. This provides a round-off error between 0 and 7. Assume 3.5 bytes.
- use of shortened last code word. This provides a round off error between 0 and 15. Assume 7.5 bytes.
- a Reed-Solomon Forward Error Correction (FEC) of 10%
- a preamble of 72 bits (typical for QPSK) which is 9 bytes
- a guard time of 8 symbols which is 2 bytes (for QPSK)

In this example, these variables combine together to yield a formula that is:

$$\mathbf{F(\ us\_payload\_bytes\ ) = (us\_payload\_bytes + 22) \times 1.1 + 11} \qquad \mathbf{(14)}$$

Appendix B: Glossary

Expression	**Definition**
CM	Cable Modem
CMTS	Cable Modem Termination System
DOCSIS	Data Over Cable Service Interface Specification
D/S	Downstream
HFC	Hybrid Fiber/Coax
MMTE	Multimedia Traffic Engineering
MRTG	Multi Router Traffic Grapher
MP_a	Market penetration of an application within the user base. Relative to the number of data subscribers. The sum of all MP_a may exceed one.
MP_c	Market Penetration of all cable subscribers. Relative to HHP.
MP_d	Market Penetration of all data subscribers. Relative to HHP.
MP_t	Market Penetration of a tier of service. The sum of all MP_t equals one.
MP_v	Market Penetration of all voice subscribers. Relative to HHP.
PPS	Packet Per Second
SD_a	Session Density during the average interval. Relative to the number of data subscribers.
SD_m	Session Density during the maximum measurement interval. Relative to the number data subscribers during the peak measurement interval.
SD_p	Session Density during the peak measurement interval. Relative to the number of data subscribers during the average measurement interval.
U/S	Upstream
VoIP	Voice over Internet Protocol
WAN	Wide Area Network

Appendix C: References

[1] "*What is the Maximum Number of Users per CMTS?*", by Steve Lee, Cisco Tech Notes, November 20, 2000 *http://www.cisco.com/warp/public/109/max_number_cmts.html*

[2] *"Multimedia Traffic Engineering for HFC Networks, A White Paper on Data, Voice, and Video over IP"*, by John T. Chapman, presented at SCTE Emerging Technologies 2000 Conference, January 2000 *http://www.cisco.com/warp/public/cc/so/cuso/sp/hfcn_wp.pdf*

[3] Napster, *http://napster.com*

[4] "*Modeling Telecommunication Systems with Self-Similar Traffic*", by Pierre M Fiorini, University of Connecticut, September 1998 *http://citeseer.nj.nec.com/fiorini98modeling.html*

[5] "*MRTG: Multi Router Traffic Grapher*", by Tobias Oetiker and Dave Rand, *http://people.ee.ethz.ch/~oetiker/webtools/mrtg/*