

Aggressive Telecommunications Overbooking Ratios

Robert Ball, Mark Clement, Feng Huang, Quinn Snell, Casey Deccio
Computer Science Department

Brigham Young University
3361 TMCB

Provo, Utah 84602

Corresponding Author: Robert Ball (brg4q@yahoo.com)

Abstract

The Internet is comprised of vast networks of wires and fiber. A common misconception is that there is an unlimited amount of bandwidth; in reality there exists only a finite amount. Each length of wire and fiber is owned by a company, and every company wants to maximize its profit. One means of improving profit is to overbook existing transmission lines in order to increase income without increasing expenses. If too much overbooking is performed, the Quality of Service (QoS) seen by customers will decline. This paper explains a process to achieve an optimal Overbooking Ratio (OR) for admission control in network routers. By optimizing the overbooking ratio, profits can be increased while minimizing QoS problems for users.

1. Introduction

Overbooking is a method that telecommunications companies use to maximize the usage of their physical lines. Similar to airlines that book more seats on planes than are available, telecommunications companies reserve more bandwidth than is actually available. For example, given a T1 physical connection (1.544 Mbps), a company could sell usage of that line in pieces (e.g., 10 connections at 154.4 Kbps each) or as a whole only up to 1.544 Mbps. Intuitively, the company could not sell more bandwidth than what is physically available without overbooking the line.

Overbooking has become a popular way of making more bandwidth available for customers without actually installing more physical lines. Historically, finding the correct overbooking factor has been a hit-and-miss process. This paper proposes a process that systematically finds an appropriate overbooking factor through a process

of gathering data, analyzing that data, and verifying the result with simulations.

In this paper we refer to the overbooking ratio (OR) as the amount of overbooking that takes place. If a T1 line has an OR of 0.50 then the line's owner could sell 3.088 Mbps of bandwidth instead of only 1.544 Mbps, effectively selling the same bandwidth twice. The Overbooking Ratio is multiplied by the bandwidth request from the customer in order to determine actual bandwidth allocated by admission control. Or in terms of the router, when a customer requests usage on the T1 line, the router only logs 0.50 of what the customer requests, only logging a request for 772 Kbps (1.544 Kbps * 0.50). The following equation shows how much effective bandwidth is produced depending on the overbooking ratio:

$$(1/OR) * \text{Physical Link} = \text{Effective Bandwidth}$$

The nature of network traffic is rarely constant, but it is more commonly "bursty". For this reason, an entity seldom uses all of its purchased bandwidth all of the time. There are many physical links whose bandwidths are under-utilized at different times. Overbooking takes advantage of this under-utilization to sell more bandwidth than is available to customers.

Although the situation arises in which a customer utilizes all or more of the requested bandwidth, bandwidth overbooking takes this into account. If the personality of a line is such that the customers use as much or more than they requested on a continual basis, then overbooking can work in the opposite way; causing the effective bandwidth to decrease instead of increase. This research has found that customers rarely use all of the bandwidth that they request and that when multiple customers share the same line the "bursty high" traffic of different customers do not often align. The different personalities that a line might have underscore the need to have a process to develop a specific overbooking ratio for each line or group of lines.

Most modern routers and switches have the ability to overbook lines. The default OR for most routers and switches is 100%—effectively not overbooking at all [1,2]. Overbooking assumes that when overbooking a line the customers purchasing its bandwidth will not collectively use more of the line than is physically available at any particular time. When this assumption is violated, packets are lost. Dropping packets on a regular basis introduces a problem that many consider worse than idle links. This paper proposes a process for finding an optimal overbooking ratio that both reduces packet dropping and line idleness.

1.1. Previous research

Pazos and Gerla have performed research in overbooking in the use of ABR services on Internet backbones [6]. Their paper asserts that the use of ABR services permits better utilization of ATM resources than with CBR services, but that full utilization rarely occurs. They propose an approach to improve resource utilization that relies on bandwidth overbooking.

Urgaonkar, et al. demonstrate the feasibility and benefits of overbooking resources in shared platforms [10]. Other research has also turned to overbooking as a reliable alternative to increasing link utilization [5,7,8].

Overbooking can be an effective tool for increasing link utilization for many different types of applications. This research differs from previous research in that instead of proving that overbooking is helpful, it proposes a process that router administrators can use to determine an effective overbooking ratio for each particular link that they administer.

2. Process

The process described in this research consists of three steps: gathering information on the line to be overbooked, running statistical analysis on the gathered data, and running simulations to verify the results of the statistical analysis.

2.1. Gathering Data

The most important step in finding an optimal overbooking ratio is the gathering of data on the line that will be overbooked. Each data line, LAN, or WAN, has unique characteristics. Data gathered for a specific data line or group of lines cannot arbitrarily apply to another data line or group of lines. The distribution of packets of each line or group over a specific time is unique to it alone and cannot be applied to other lines that may have a different distribution of packets over time.

For instance, one line might normally have a very low load with high burstiness while another line might

normally have a high load with low burstiness. The overbooking ratio that fits the first line most likely will not fit the second line. Other characteristics of lines that might distinguish them one from another would be the actual traffic content. Consideration should be taken into account for lines that have mostly UDP traffic, TCP traffic, or high volumes of media traffic. In order to get the most usage out of any particular line, all of these factors should be considered.

In order to get optimal utilization out of a line, the data to be gathered should record at least the bandwidth that is requested from the customer and how much bandwidth is actually used for the line. The requested bandwidth is the amount of bandwidth that a customer is paying for and the bandwidth used is how much actual bandwidth was used during the data gathering period. Table 1 is an excerpt of data that shows the data of 5 virtual circuits (VCs) out of a total listing of 475,227 rows of data from a commercial vendor used for this research. For the particular data used in this research, each row represents a customer to a telecommunications company. In this particular case, a telecommunications company is looking for a general overbooking ratio that will work for a large collection of customers and network connections.

Customer	Requested Bandwidth	Physical Capability	Bandwidth Used
A	1088	2212	0
B	34	2212	18
C	1088	2212	295
D	272	2212	257
E	816	2212	239

Table 1: This table shows an excerpt of data used in this research. The most important columns for gathering data are the Requested Bandwidth or Committed Information Rate (CIR) which indicates how much bandwidth was requested and the Bandwidth Used.

Most networking hardware does not gather detailed statistics on bandwidth used. If this value were totally accurate, then it would be easier to generate an overbooking ratio as the average Bandwidth Used divided by Requested Bandwidth. Since many customers do not use their bandwidth very extensively during the night or during non-business hours, the average in this case would be highly skewed towards a much smaller bandwidth usage. If the overbooking ratio overbooked on what the line does on average then if customers used more than average bandwidth usage during the day, packets would be lost. Even if the customers were to use their bandwidth at a constant ratio during business hours, due to the overbooking ratio being based on the average time, any

bandwidth above the average bandwidth might result in loss of packets.

However, if the data were gathered based on maximum bandwidth utilization, or the peak bandwidth usage used during a particular time period, then instead of losing packets, the opposite effect would be true – idle lines. For instance, if customer X were a business that has a vast majority of its employees checking their email and visiting certain websites all at the same start of business hours, then a huge spike would be seen at that time. If customer X used only a small portion of that bandwidth during the rest of the day, then the line that customer X uses would be idle most of the time.

In order to overcome loss of packets and idle lines, a compromise is suggested. This compromise suggests that an average peak usage be recorded. For example, if customer X peaks at 100 Kbps, then instead of recording customer X's usage at 100 Kbps the average amount for a short period of time be recorded. So, if customer X peaked at 100 Kbps, then an average usage over a five-minute time period might be more appropriate. This five-minute average might be closer to 15 Kbps instead of 100 Kbps resolving problems with spikes.

Figure 1 shows the reason why not to simply take the peak bandwidth as the best indicator of bandwidth usage. As can be seen in Figure 1, the peak of 526 Kbps at time 318 minutes does not truly express how much bandwidth is normally used on this particular line. If only the spike is recorded, then an increased amount of bandwidth is reserved and the line sits idle. Also, if the average is taken, which is only 10.725 Kbps for the eight hour block, then an excessive amount of packets might be dropped. If the data gathered had recorded averages of five minute intervals instead, then the peak average would be 117.2 Kbps, a much better estimate on the data.

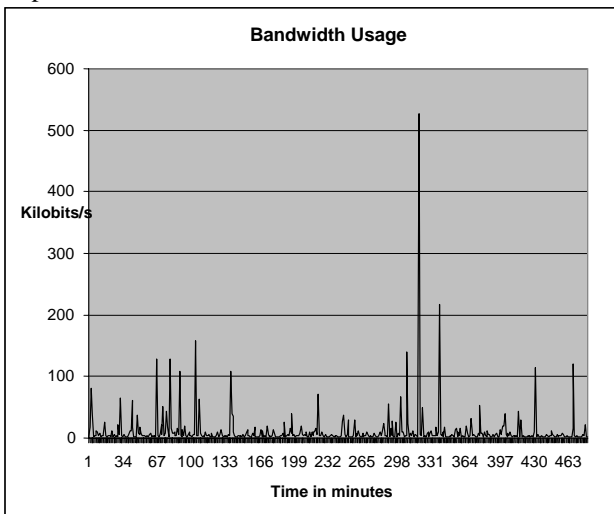


Figure 1: This figure shows a customer's usage in an 8 hour (480 minute) time period. This figure shows how the spike at minute 318 (usage of 526 Kbps) could skew the remainder of the data if used incorrectly.

The choice on how to record the data depends on what is desired. If just the peak amount of bandwidth is recorded, then the OR that is produced from the analysis will ensure that the minimal amount of overbooking is used while not losing any packets. This may be desirable for a company that is concerned more with quality of service than under-utilized lines. Even if the telecommunications company chooses this route, then a large amount of overbooking may still be obtained. For example, in the data that was used in this paper, only the peaks of a twenty-four hour period were recorded. Even so, the results from the analysis showed that for the particular data used, an OR of 25.78% could be safely used. That means that by taking the safe route, the lines that were analyzed could still be overbooked four times as much with minimal packet loss.

If under-utilized lines are more of a concern than packet loss, then the compromise of taking the peak over an average of time would be preferred. As can be seen by figure 1, if the average time over five minutes were used, and the peak were recorded at 117.2 Kbps then a much higher OR could be obtained resulting in overbooking the line at a much higher ratio with little packet loss. In this scenario, the period of time would not have to be five minutes but could be varied depending on what is desired. The period of time could range from one second to one hour. It should be remembered that the smaller the period of time, the closer the data will resemble the peak bandwidth used and the longer the period of time, the more the peak bandwidth will resemble the average bandwidth used.

The question of how often or long data should be gathered should also be taken into account. If data is collected only for one day, then it is presumed that all days are like that one day. On the other hand, recording data on a typical day may also be beneficial in that if every business day is much like any other then only one day is recorded. It should be clear to the reader that recording non-business days, or days that the customer regularly does not use significant bandwidth would be unwise as the OR would be highly skewed towards zero. If on the other hand a customer varies greatly from one day to another then several days of diversified use is recommended to get a more accurate feel on the personality of the data line in question.

The data to be gathered for analysis should be as consistent with the actual line as possible. Intuitively, the closer the data resembles the actual personality of the line, the better the results.

3. Statistical Analysis

Once the data has been gathered it should then be put into a logical form. Many times when routers log information about a particular data line, the log is all but incomprehensible. The first step to putting the data into a logical form is to understand what each value in the log actually means. As far as the statistical analysis is concerned, the algorithms used should have access to how much bandwidth was used and how much was requested. It may take considerable time to analyze and understand the logs. Once the meaning of the values is understood, then the extraction of the information should take place preferably into a format of columns and rows.

After the data has been put into a logical form, it is then necessary to analyze the data. Once again it is important to point out that if the gathered data does not closely match the personality of a particular line or lines that is being analyzed, then the analysis will not be very helpful.

When looking for an optimal OR for a data line, one is looking to see what portion of the requested bandwidth a router should log. As can be seen from table 2, customer A requested 1088 Kbps but used on average 0 Kbps. On the other hand customer B requested 34 Kbps and only 18 Kbps were on average used. In order to overcome the differences between the rows where customer A used 0% of the requested bandwidth and customer D used 94% of the requested bandwidth statistical analysis is used to reach a consensus.

The goal of the analysis is to find a common OR that would best match the bandwidth used for the entire line or lines. For example, in table 2 there are five different customers that request four different bandwidths that use bandwidth ranging from 0% to 94%. The goal for this group would be to find the best match between the requested bandwidth and the actual usage. Column 4 of table 2, Percentage Used, shows what the exact OR should be for each customer. That is to say that if the future were known to the router at the time the customer requested bandwidth then the router would be able to perfectly give the customer the correct percentage of that requested bandwidth; the OR that matched that customer perfectly. Unfortunately, routers do not currently have the ability to predict the future perfectly, so statistical analysis based on past history is currently the best predictor.

Customer	Requested Bandwidth	Bandwidth Used	Percentage Used
A	1088	0	0%
B	34	18	53%
C	1088	295	27%
D	272	257	94%
E	816	239	29%

Table 2: This table shows the diversity in amount of bandwidth used. The fourth column shows that the customers used on average anywhere from 0 – 94% of their requested bandwidth.

Also, it would be best if the router overbooked each customer based on their own personal history of bandwidth used instead of the entire line. If routers were able to keep a history of how much bandwidth each customer requested and used, then for each customer the router could keep a much more accurate OR. In this scenario, the router could allocate all the bandwidth requested for a new customer (OR of 100%) and slowly change the OR for that customer as time progressed and more data became available for analysis. As routers currently only have the capability of having one OR per data line, instead of one OR per customer, then the best possibility is not currently available.

As a result, the solution is to find an OR that matches a group's personality instead. This is similar to how professors teach in a classroom. Every student in the classroom is different and learns and reacts differently from each other. In response, a professor must use teaching methods that cater to the class as a whole which may not be suitable for any particular student, but is the best way to teach the students collectively. Similarly, as customers are like students and routers are like professors, the router cannot cater to any particular customer but to all the customers at once. This may not precisely benefit any particular customer, but will benefit all the customers as a group.

In order to find an OR that best matches a group's personality, linear regression algorithms were applied to the data. As this paper's data consisted of 475,227 rows of data, spreadsheets such as Microsoft Excel and Open Office were not able to handle the load. Also, such spreadsheets may lack sophistication with complex statistical analysis algorithms. As a result a more sophisticated statistical analysis application was needed. This paper used a statistical analysis application named SAS. Other applications such as R or Minitab could also be used for the same purpose.

Linear regression is used to "predict, or estimate, the value of one variable from known or assumed values of other variables related to it." [9] That is to say that in this case the linear regression algorithms try to predict how

much bandwidth will be used based solely on how much bandwidth is requested. Other more complex models might also be used but the required data may not be available. Also, log files may be large and often range into the gigabytes in size which makes them much harder to manage.

Many different theoretical models may be created based on any number of variables, but due to limited resources simply knowing the requested amount of bandwidth and the bandwidth used may be enough. To illustrate the point, routers are relied on heavily to send packets to the correct locations. Router administrators want routers to be as fast as possible and do not want a performance hit due to constant router logging and analysis. A balance of how much information can be stored for analysis must be met which tends to favor performance and not an excessive amount of evaluation.

A linear regression algorithm generally gives a solution with an intercept as default. For the purpose of obtaining an OR, an intercept is not wanted as only one value, or one percentage, can currently be used in routers. Although complex formulas would be more beneficial, routers at this time are only able to use one value, or one percentage.

As a basic guideline, the steps used to obtain the OR from SAS are outlined. After the data from the routers has been put into logical columns and rows, with each column labeled, the data is ready to be imported into SAS. The data may be delimited in any way, but a simple space between each column is sufficient. After the data has been imported into SAS it should then be analyzed. To analyze the data the following sample SAS code to find the desired OR without an intercept follows:

```
1 proc glm data=OverbookingData;
2 model BandwidthUsage=RequestedBandwidth
/solution ss3 noint;
3 output out=output1 r=Residual p=PredictedValues;
4 run;
```

The first word, “proc” invokes a method, or algorithm, in this case the “glm” or general linear regression algorithm. The algorithm uses data imported into SAS, in particular data labeled “OverbookingData.”

The next line, line 2, explains what model is wanted. In this particular example, it is proposed, as explained in this paper, that the bandwidth usage is somehow correlated to the amount of bandwidth requested. In this particular example, there are two columns of data named RequestedBandwidth and BandwidthUsage. The next part of line two “/solution ss3 noint” gives a resulting model or solution without an intercept. The result of line 3 is summarized as follows:

Parameter	Estimate
BandwidthUsage	0.2578129169

The meaning of the above two lines is that the parameter “BandwidthUsage” = “RequestedBandwidth” * 0.2578129169. That is to say that expected bandwidth usage for a customer on the line or lines that this data corresponds to is equal to the requested bandwidth multiplied by 0.2578129169. This number (0.2578129169) is the OR for the line, which means that for every Kbps (or whatever unit is being used) that is requested that the router would allot approximately 25.8% of what is requested.

This output means that with the given data the optimal OR is 25.8% of the requested bandwidth. Intuitively, then means that for every request given to the routers that corresponds to the data, only 25.8% of the requested bandwidth will be logged and that four times the amount of customers can use these lines.

However, this does not mean that customers will not be able use all of the bandwidth that they requested, but that they most likely will not use it all the time. In fact, what normally happens is that customers peak at different times. For example, consider figure 2 where each peak, or spike, occurs at different times for different customers (virtual circuits). The first column shows all four customers using the same bandwidth. The second column shows that customer VC1 spikes from 5 Mbps to 20 Mbps. VC1 later goes down to 19 Mbps then back to 5 Mbps. As is illustrated by the red, spikes often occur at different times allowing customers to still get the bandwidth they need, but at different times as what naturally occurs. If customers constantly spiked at regular intervals then that would be taken into account in the linear regression algorithm.

VC0	5Mbps	5Mbps	10Mbps	5Mbps	5Mbps	15Mbps
VC1	5Mbps	20Mbps	19Mbps	5Mbps	5Mbps	5Mbps
VC2	5Mbps	5Mbps	10Mbps	25Mbps	5Mbps	5Mbps
VC3	5Mbps	5Mbps	9Mbps	5Mbps	10Mbps	5Mbps

Figure 2: This figure shows four different virtual circuits at differing times. As time progresses from left to right, each virtual circuit spikes at different times. This shows the fluctuations of patterns in each virtual circuit and how they are not constant.

When the OR is evaluated to be a good fit or not, the OR should be compared to the actual data. The reason for this evaluation is that linear regression algorithms do not always merge due to uncorrelated data. Line 3 of the SAS code produces another table of data that has in addition to

the original data the residual and predicted values of the data. Figure 3 is a representation of that data.




Figure 3: This chart is an example that shows the results of the predicted OR compared to the actual usage minus the predicted usage (residual). If the predicted OR were completely correct every time then the result would follow a straight line along the x-axis at the origin of the y-axis.

The above figure is an example that shows a data plot of residual values versus predicted values. The residual value is a value of the actual usage that was logged in the data minus the predicted value. The predicted value is the requested bandwidth multiplied by the OR. If the predicted value were exactly correct then it would lie at the value of zero, the dotted line, along the y-axis (residual). That is to say that if in the above example all the predicted points were exactly correct then the predicted values would form a solid line at point zero on the y-axis. However, as can be seen, this is rarely (if ever) the case. Figure 3 shows a large cluster around the dotted line that represents the ideal solution, but there are several outliers.

It should also be pointed out that in this particular example there are more outliers that have a residual value less than zero. This means that more bandwidth would be given to those customers than they need in this area. This also means that there will be slightly more idle lines and less packet loss as a result. If the opposite result is wanted, that of less idle lines and more packet loss, then the OR value can be adjusted manually until the results are achieved and more of the plots are above the dotted line.

Requested Bandwidth	OR	Actual Usage	Predicted Usage (Requested Bandwidth * OR)	Residual (Actual - Predicted)
500	25%	244	125	119
500	25%	38	125	-87
500	25%	125	125	0
500	25%	85	125	-40

Table 3: This example shows the process of obtaining the predicted and residual values. The predicted usage is obtained by multiplying the requested bandwidth times the OR and the residual value is the actual usage subtracted by the predicted usage. The above values are in Kbps.

To also illustrate this, an example result of four points is shown in table 3. Table 3 shows that although all four customers request the same amount of bandwidth only one of them actually meets the exact prediction of 125 Kbps.

In fine, very rarely will the OR predict how much a customer will actually use, but will be accurate for the larger picture. Referring back to figure 3, notice that approximately half the plots are above the dotted line and approximately half the plots are below the dotted line. Recall that routers currently only have the capability to predict on an entire data line, there is not a better solution.

It is important to note that although the predicted usage is not always completely correct, it is better than what was given before. For example, in table 3 the four links have a combined requested bandwidth of 2000 kilobits/s. After analysis of the data, an OR of 25% was obtained a total of 500 kilobits/s was allotted, freeing room for 1500 kilobits/s more space for other customers. In this instance although the predicted usage is not always right and allocates more space than is needed, in total it only allocates 8 kilobits/s more than what is used between the four of them. So, though not correct for any one customer, the resulting OR is correct for the group as a whole.

4. Simulation

Simulations have two roles in this paper: first, the question of verification of the OR and second, verification of all other factors. Simulations to verify the OR should be done after obtaining the OR that resulted from the linear regression algorithm. By running simulations, the effect of the OR can be determined before actually applying the OR on the physical line. By running simulations in a controlled environment, the outcome can be seen without any damage done to the real system.

These simulations verify that the OR is in fact a good fit to the particular line being tested.

Simulators, such as *ns-2* [reference] can be used to predict real-time results without actually using a real router. As an example, figure 4 shows the packet loss results of running the simulator *ns* with varying OR's. Such simulations should include as much information about the real data line as possible. As *ns-2* tries to simulate as best as possible real conditions, the exact nature of the line should be taken into account. Figure 4 demonstrates this by showing two data lines. The top line, a T3 line, shows a packet drop of 11% at an OR of 0.3 (30%) while the T1 connection shows a packet drop of 2% at the same OR point. With all other factors being the same, same packet types, same buffer size, etc., the two links show a very different story depending on the bandwidth speed. This example demonstrates the need to be as precise as possible in the simulations. By only having one value off, that of the bandwidth capacity, the effects are very large. The same effects can be seen by things such as difference in buffer size, using red or not, the type of traffic such as all TCP, all UDP, 80% TCP with 20% UDP, and so on. The closer the simulation is to the actual personality of the data line the better.

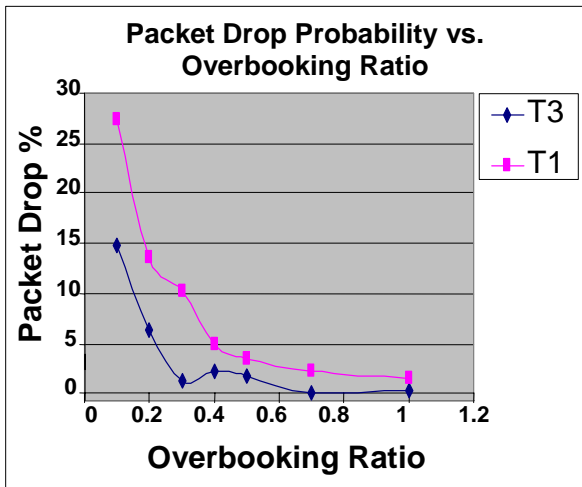


Figure 4: This figure is an example of how only one variable being different from another can drastically affect the results of a simulation. In this example, the only difference in parameters is that of a T3 line instead of a T1 line.

In an effort to validate the OR of 25.8% that was produced from the linear analysis algorithm, simulations on the data used in this paper were done. Figures 5 and 6 show a summary of the simulations in a chart form. As can be seen from figure 5, a much lower percentage for the OR could have been used for overbooking to increase the utilization of the routers, but in order to keep the loss

of packets from getting out of hand, figure 6 clearly indicates that less overbooking is more beneficial. These two figures show that the OR of 25% both increases utilization and keeps the loss rate of packets at a minimum.

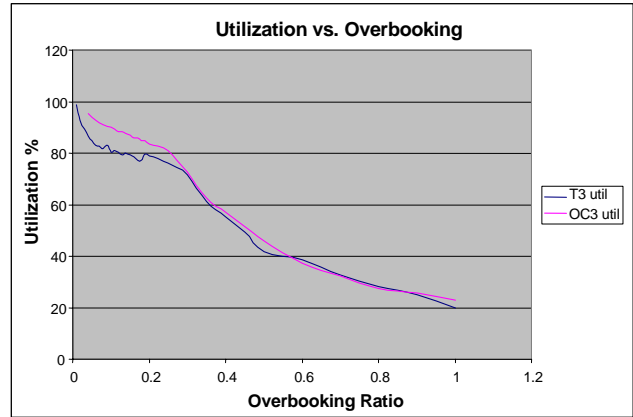


Figure 5: This figure shows that as overbooking is increased so too does the utilization of the router.

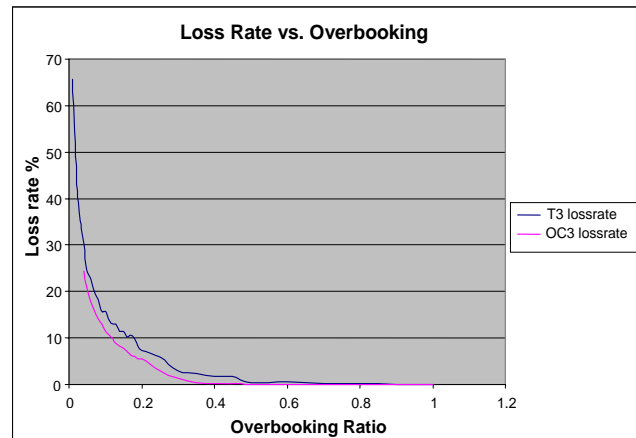


Figure 6: This figure shows that as overbooking is increased the loss rate of packets is exponential and should be avoided.

By combining the results of both simulations (Utilization vs. Overbooking and Loss Rate vs. Overbooking), the value of 25% for the OR can be seen to be optimal by visual inspection. By running simulations and producing visual feedback such as figures 5 and 6 to verify the optimality of the OR, the router administrators can feel more assured that the process proposed in this paper does in fact work.

With an actual OR that was produced through a process instead of guesswork to work with, router administrators could then verify that the rest of the

configurations of the router are correct or find that they are incorrect. By having an OR that matches the data line being used, a much better configuration can be achieved. With routers every small configuration can change the optimality of the router. With one less variable to worry about, router administrators can then focus their attention on other parameters.

5. Conclusion

Overbooking has been found to be a way to increase revenue for telecommunications companies and to overcome idleness in lines. By overbooking lines, the line is used by more people and telecommunications companies are able to gain more money without any additional cost on their part. Unfortunately, overbooking has been a hit and miss game in that network administrators have to initially guess what would be a good overbooking ratio (OR) and then continually adjust the ratio until they find an OR that seems optimal. This guesswork is very tedious and time consuming and fraught with error.

This paper has explained how to find an optimal OR without any guesswork. The process proposed in this paper consists of gathering data about the line or lines to be overbooked, analyzing the data, and then creating simulations to verify the result. By following the process proposed in this paper one can more quickly and reliably obtain an optimal OR based on well defined statistical algorithms instead of guessing.

6. References

- [1] Cisco Systems, Incorporated. Series Universal Broadband Routers Chapter 3: Configuring Cable Modem Interface Features. July 14, 2003.
http://www.cisco.com/en/US/products/hw/cable/ps2217/products_configuration_guide_chapter09186a00801170dd.html.
- [2] Cisco Systems, Incorporated. Series Switches Traffic Management. July 14, 2003.
http://www.cisco.com/en/US/products/hw/switches/ps525/products_configuration_guide_chapter09186a0080086fee.html
- [3] Kweon, S. and Shin, K. Real-time transport of MPEG video with a statistically guaranteed loss ratio of ATM networks. *IEEE Transactions on Parallel and Distributed System*, 12(4):387-403, April 2001.
- [4] Parulkar G., Schmidt, D., Kraemer, E., Turner J., and Kantawala A., "An Architecture for Monitoring, Visualization, and Control and Gigabit Networks". *IEEE Network*, vol. 11, September/October 1997.
- [5] Pazos, C. and Gerla. M. "A Rate Based Back-pressure Flow Control for the Internet", In *Proceedings of HPN'98*. Vienna, Austria - 1998.
- [6] Pazos, C. and Gerla. M. "Improving Internet Traffic Transport over ABR Backbones Through Bandwidth Overbooking", In *Proceedings of Globecom '98*. Sidney, Australia ,1998.
- [7] Pazos, C. and Gerla. M. "Improving IP over ATM Efficiency Using Bandwidth Overbooking". In *Proceedings of LANMAN'98*. Alberta, Canada - 1998.
- [8] Turner, Jonathan S. "Maintaining High Throughput During Overload in ATM Switches". *Proceeding of IEEE INFOCOM 1996*, San Francisco, California, April 1996, pp. 287 – 295.
- [9] Younger, Mary Sue. "A Handbook for Linear Regression". Duxbury Press, Belmont California. 1979, p. 1.
- [10] Urgaonkar, B., Shenoy P., and Roscoe T.. "Resource Overbooking and Application Profiling in Shared Hosting Platforms". *USENIX Symposium on Operating Systems Design and Implementation (OSDI 2002)*, Dec. 2002.