AN ANALYTICAL MODEL OF A BANDWIDTH ALLOCATION SCHEME FOR A NETWORK SERVICE PROVIDER (NSP)

Daniel A. Menascé
Dept. of CS, MS 4A5
George Mason University
Fairfax, VA 22030-4444
Menasce@cs.gmu.edu

Aiqun E. Wang
Dept. of CS, MS 4A5
George Mason University
Fairfax, VA 22030-4444
Awang@gmu.edu

Robert Steele
Loral Orion, Inc
205 Perry Parkway, unit 8
Gaithersburg, MD 20877
Rsteele@LoralOrion.com

Gregory J. Weidman
SM&A Corporation
2600 Park Tower Dr suite 1000
Vienna, VA 22180
Greg.weidman@smawins.com

Network Service Providers (NSPs) provide Internet backbone access to ISPs and need to guarantee a committed quality of service to their customers. These customers share NSP resources such as routers and links. Admission and congestion control policies can be used to allocate resources to different users and guarantee quality of service. This paper presents an analytic model, based on Markov Chains, used to determine proper bandwidth allocation thresholds so that packet drop probabilities can be minimized and throughput maximized.

1. Introduction

Loral Orion, a subsidiary of Loral Space and Communications, belongs to a small cadre of "intermediary" Internet Service Providers (ISPs), who are referred to as Network Service Providers (NSPs). Rather than focusing on the consumer dial-up Internet access market, NSPs provide Internet backbone access to tens of thousands of primarily internationally based smaller ISPs (a.k.a. Tier 3 ISPs) via traffic aggregation and multiplexing techniques. In essence, NSPs provide Internet infrastructure services that are transparent to end-users. The international focus of Lorals' marketplace reflects the significantly higher telecommunications and Internet infrastructure costs outside of the United States.

In Fall 1998, Loral introduced a new Internet access service called eBurst. A dedicated 4 Mbps satellite channel (carrier) was allocated for the launch of the service. See Figure 1.1 for a pictorial view of the service.

eBurst is sold on a subscription basis with each customer receiving a guaranteed minimum bandwidth or Committed Information Rate (CIR) and the expected occasional use of available capacity beyond their CIR. This excess capacity is referred to as Excess Information Rate (EIR). The service was founded on the assumption that the naturally "bursty" nature of the predominantly Web-based ISP traffic would enable bandwidth sharing. Ideally, the entire channel bandwidth would be consumed on a real-time basis with "bursts" filling in the gaps from customers with instantaneous offered loads below their CIR.

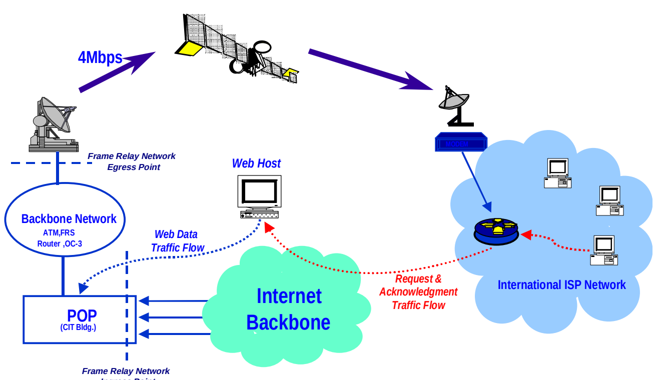

Figure 1.1 Architecture of eBurst Service

Several congestion problems were detected and policies for fair bandwidth allocation between different customers were deemed necessary. An immediate, interim fix was applied by simply scaling back the customer Excess Information Rates (referred to as "Be"). Research was then undertaken to identify the causes of the congestion and assess probable fixes. We describe in this paper an analytic model that was developed to evaluate the impact of different bandwidth allocation schemes on throughput and packet drop rate. Based on the results of the analytical process described in this paper, a more appropriate permanent fix was facilitated via the introduction of Class-Based Queuing (CBQ). The analytical and test results of the CBQ solution will be addressed in a follow-up paper.

The rest of the paper is organized as follows. Section 2 describes the problem faced by the NSP. Section 3 presents a Markov Chain-based analytic model used to determine packet drop probabilities, throughputs, and packet delays. The next section presents several numerical results obtained with the model. Finally section 5 presents some concluding remarks.

2. Problem Definition

Packets arrive from the Internet at the Ingress router (see Fig 2.1) and, according to an admission control policy at the router, they are either discarded or continue to the outbound router, called Egress on Fig. 2.1.

The admission control policy at Ingress works as follows. If the incoming traffic from a user exceeds its allocated Excess Information Rate (EIR), its packets are discarded. They are called red packets. If the incoming data rate exceeds the Committed Information Rate (CIR) but does not exceed the EIR for a user, packets are tagged as yellow packets and are sent to the Egress router. Finally, if the data rate does not exceed the CIR, packets are tagged as green packets and are sent to the Egress router (see Fig. 2.2)

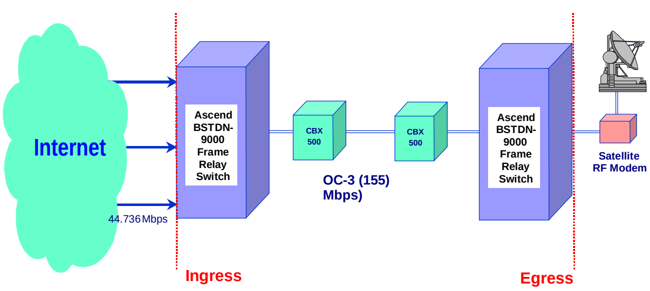

Figure 2.1 - Networking Diagram for eBurst

Egress implements a congestion control policy. Arriving packets are tagged as either green or yellow. If a yellow packet arrives and finds T or more packets (green or yellow) in the system, the arriving packet is dropped. If an arriving green packet finds M packets (green or yellow) in the system, it is dropped. Note that T < M.

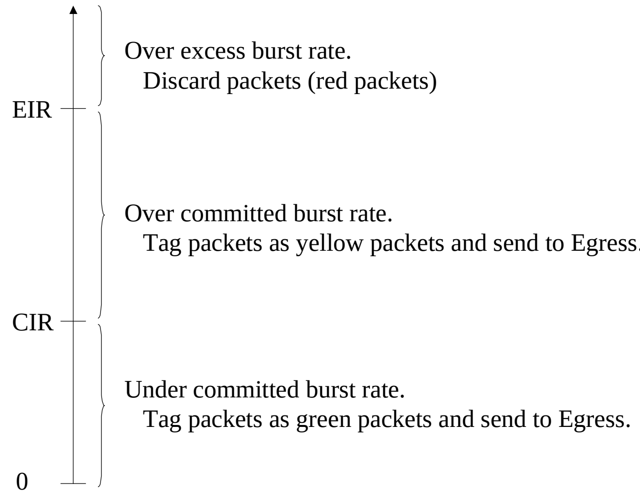

Figure 2.2 - Definition of Red, Yellow, and Green Packets.

Several anomalies were detected shortly after the introduction of the eBurst service. Of foremost concern is the inability of the current Frame Relay configuration to protect "guaranteed" CIR service availability. Customer PVCs (Permanent Virtual Circuits) with excess information rate traffic volumes are not supposed to encroach on offered traffic loads ≤ CIR of co-resident PVCs. Excessive congestion is the apparent cause of the condition. The following symptoms were detected: a) During busy hours, green packets are dropped, indicating an absolute congestion condition, and b) the packet drop rate varies from 6.8% to 15.3%, with an average of 9.1%, during busy hours.

By properly adjusting the CIR and EIR for each customer, one can affect the packet drop rate, the throughput, and packet delay values. Section 3 provides an analytic model of the Egress router. This model can be used to determine the proper bandwidth allocation settings for each user so that adequate quality of service is delivered.

3 The Model

To analyze the various policies for bandwidth allocation between various customers, we built an analytic model described in what follows. Figure 3.1 shows the Egress Router Queue.

We are interested in obtaining the following performance metrics:

- P_g^d : probability that a green packet is dropped,
- P_y^d : probability that a yellow packet is dropped,
- X_g : throughput of green packets, i.e., number of green packets transmitted per second (in pps).

- X_y: throughput of yellow packets, i.e., number of yellow packets transmitted per second (in pps).
- R_g: average delay of green packets at the Egress router (in sec).
- R_y: average delay of yellow packets at the Egress router (in sec).

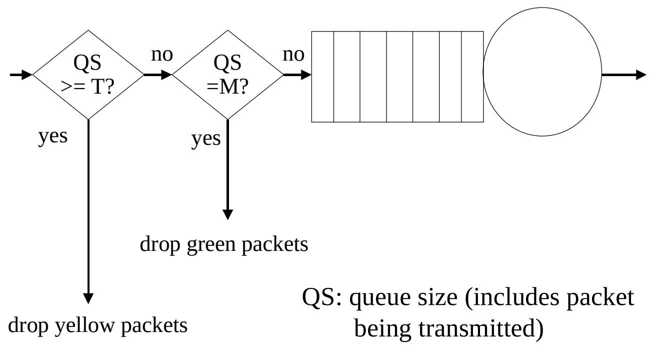

Figure 3.1 Egress Router Queue

The input parameters are:

- λ_g: arrival rate of green packets, in pps,
- λ_y: arrival rate of yellow packets, in pps,
- μ_g: service rate of green packets, i.e., the inverse of the transmission time of green packets. This is computed as the average size of green packets divided by the bandwidth of the outgoing link (in sec/packet), and
- μ_y: service rate of yellow packets, i.e., the inverse of the transmission time of yellow packets. This is computed as the average size of yellow packets divided by the bandwidth of the outgoing link (in sec/packet).

To model this queue we used a Markov Chain in which the state *k* is the number of packets in the system, including the packet being transmitted (see Fig. 3.2).

Since the packet being transmitted can be green or yellow, we need to find an average packet service rate, μ. We do that by applying Little's Law [Little61] to the server, considering first the green packets, then the yellow packets, and then considering that the average number of packets at the server (also equal to the server utilization) is the sum of the average number of green plus yellow packets.

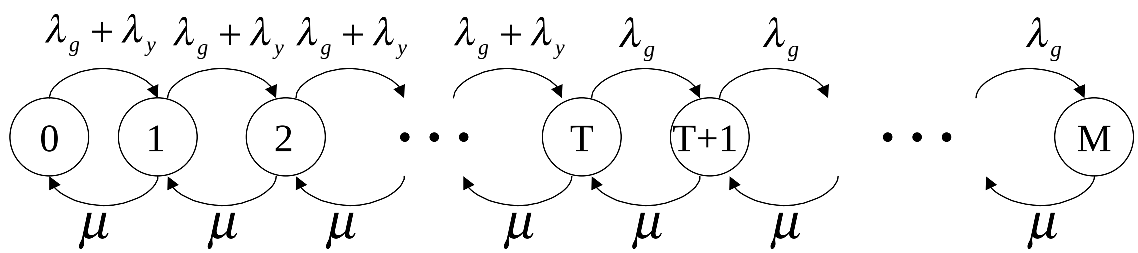

Figure 3.2 - Markov Chain for the Egress Router

So, if we let N_g and N_y be the average number of green and yellow packets at the server, respectively, then we have that

$$N_g = \lambda_g / \mu_g \text{, and} \quad (1)$$

$$N_y = \lambda_y / \mu_y . \quad (2)$$

Since $N_g + N_y$ is the average number of packets being transmitted and ($\lambda_g + \lambda_y$) is the total arrival rate of packets, we can use Little's Law again to obtain

$$N_g + N_y = (\lambda_g + \lambda_y) / \mu \quad (3)$$

So, combining (1)-(3) we get that the average packet transmission rate is given by

$$\mu = \frac{(\lambda_g + \lambda_y)\mu_g \mu_y}{\mu_g \lambda_y + \mu_y \lambda_g} \quad (4)$$

So, our Markov Chain has M+1 states, the death rates are all equal to μ and the birth-rates are given by ($\lambda_g + \lambda_y$) for *k* < T and λ_g for *k* = T, …, M-1.

The general formula for the probability P_k of finding the system at state *k* is given below [Klei75]

$$P_k = P_0 \prod_{i=0}^{k-1} \frac{\lambda_i}{\mu_{i+1}} \quad (5)$$

where

$$P_0 = \left[1 + \sum_{k=1}^{M} \prod_{i=0}^{k-1} \frac{\lambda_i}{\mu_{i+1}} \right]^{-1} . \quad (6)$$

If we plug in the expression for the birth and death rates into equations (5) and (6), we get after some algebraic manipulation the following expressions for P_k, *k* = 0, …, *M*.

$$P_k = P_0 \left(\frac{\lambda_g + \lambda_y}{\mu} \right)^k \quad \text{for k= 0, …, T,} \quad (7)$$

$$P_k = P_0 \left(\frac{\lambda_g + \lambda_y}{\mu} \right)^T \left(\frac{\lambda_g}{\mu} \right)^{k-T} \quad \text{for k=T+1, …,M}$$

(8)

$$P_0 = \left[\frac{1 - \left[\left(\lambda_g + \lambda_y \right) / \mu \right]^{T+1}}{1 - \left(\lambda_g + \lambda_y \right) / \mu} + \left(\frac{\lambda_g + \lambda_y}{\mu} \right)^T \frac{\lambda_g}{\mu} \left[\frac{1 - \left(\lambda_g / \mu \right)^{M-T}}{1 - \left(\lambda_g / \mu \right)} \right] \right]^{-1}$$

(9)

The probability P_y^d that a yellow packet is dropped is given by

$$P_y^d = \sum_{j=T}^{M} P_j \quad (10)$$

The probability P_g^d that a green packet is dropped is given by

$$P_g^d = P_M \quad (11)$$

The throughput X_y of yellow packets is given by

$$X_y = \lambda_y \left(1 - P_y^d \right) = \lambda_y (1 - \sum_{j=T}^{M} P_j) \quad (12)$$

The throughput X_g of green packets is given by

$$X_g = \lambda_g (1 - P_g^d) = \lambda_g (1 - P_M)$$

(13)

The average packet delay R_y for yellow packets is given by

$$R_y = \frac{1}{\mu_y} + \sum_{j=1}^{T-1} P_j \times \frac{1}{\mu} \times j . \quad (14)$$

Equation (14) can be explained as follows. The average delay of a yellow packet is equal to its service time (1/μ_y) plus the time needed to wait for all *j* packets found in the system by the arriving packet. Each packet in the queue has an average service time of 1/μ. Since arriving yellow packets can find at most *T*-1 packets in the queue because they are discarded after T, the summation goes from 1 to *T*-1.

The average packet delay R_g for green packets can be derived in a similar way:

$$R_g = \frac{1}{\mu_g} + \sum_{j=1}^{M-1} P_j \times \frac{1}{\mu} \times j . \quad (15)$$

4. Numerical Results

Using the model presented in section 3, one can obtain several interesting results as shown in this section.

Figure 4.1 shows the drop probability of green packets versus M for a fixed arrival rate of yellow packets equal to 200 pps. Various combinations of the arrival rate of green packets and of the value *T* are shown in the figure. In this, and in all other figures, we adopt the convention that a solid line is always paired to a dashed line with one of the parameters of the combination having the same value. Figure 4.1 has three pairs of curves. For each pair, λ_g is the same and *T* varies. In all cases, P_g^d decreases with an increase in *M*, as expected. However, for a fixed arrival rate, the drop probability reaches a horizontal asymptote after a certain value of *M*. The curves also show that for higher values of T one gets a higher value of P_g^d since less yellow packets are being dropped. This increases the queue length and therefore increases the probability of finding the system in state M.

Figure 4.2 displays the throughput X_g of green packets as a function of *M*. For each pair of curves, the throughput increases and then saturates at a maximum value which is the same for each pair of curves. The asymptotic value of the throughput corresponds to the value of *M* where the drop probability P_g^d reaches its minimum (see Fig. 4.1).

Figure 4.3 depicts the drop probability P_y^d as a function of *M*. It is interesting, and surprising at first sight, that as the value of *M* increases, the drop probability of yellow packets increases and then reaches a maximum value. This can be explained as follows. As *M* increses, the probability that green packets are dropped decreases (see Fig. 4.1). Therefore, arriving yellow packets will tend to see a bigger queue upon arrival, thereby increasing the probability that they are dropped. As the drop probability P_g^d reaches its minimum, P_y^d reaches its maximum.

Figure 4.4 shows the throughput X_y of yellow packets. Because of the behavior shown in Fig. 4.3, the throughput X_y decreases with *M* up to a minimum point, which happens when P_y^d reaches its maximum.

Figure 4.5 shows the drop probability P_g^d of green packets as a function of T. As it can be seen, as T increases, less yellow packets are dropped and therefore the queue seen by arriving green packets increases. This increases the probability of a green packet being dropped.

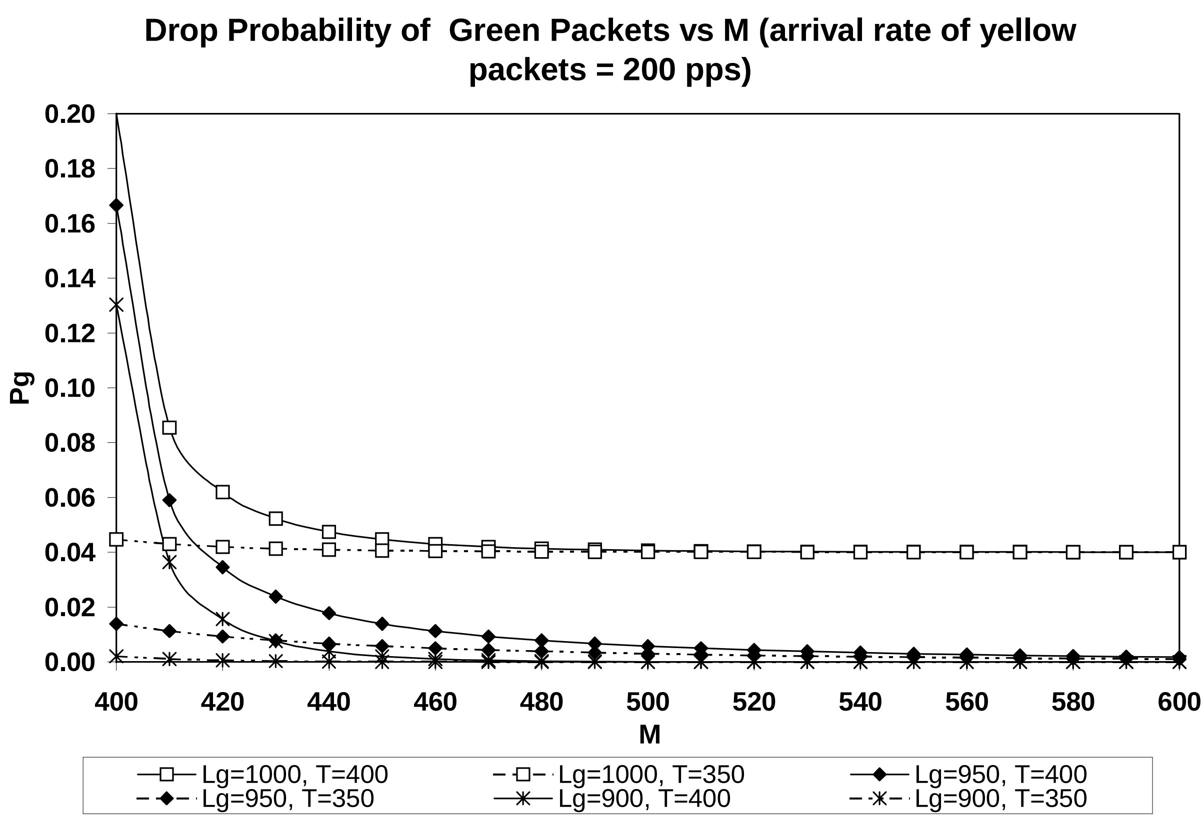

Figure 4.1 P_g^d **vs. M**

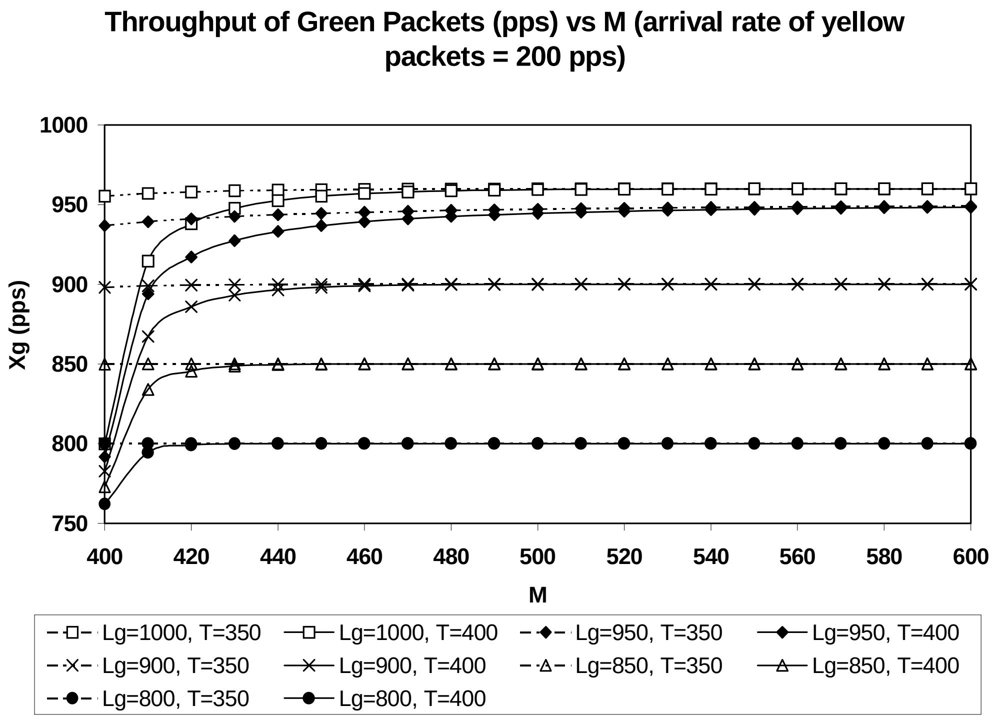

Figure 4.2 - X_g **vs. M.**

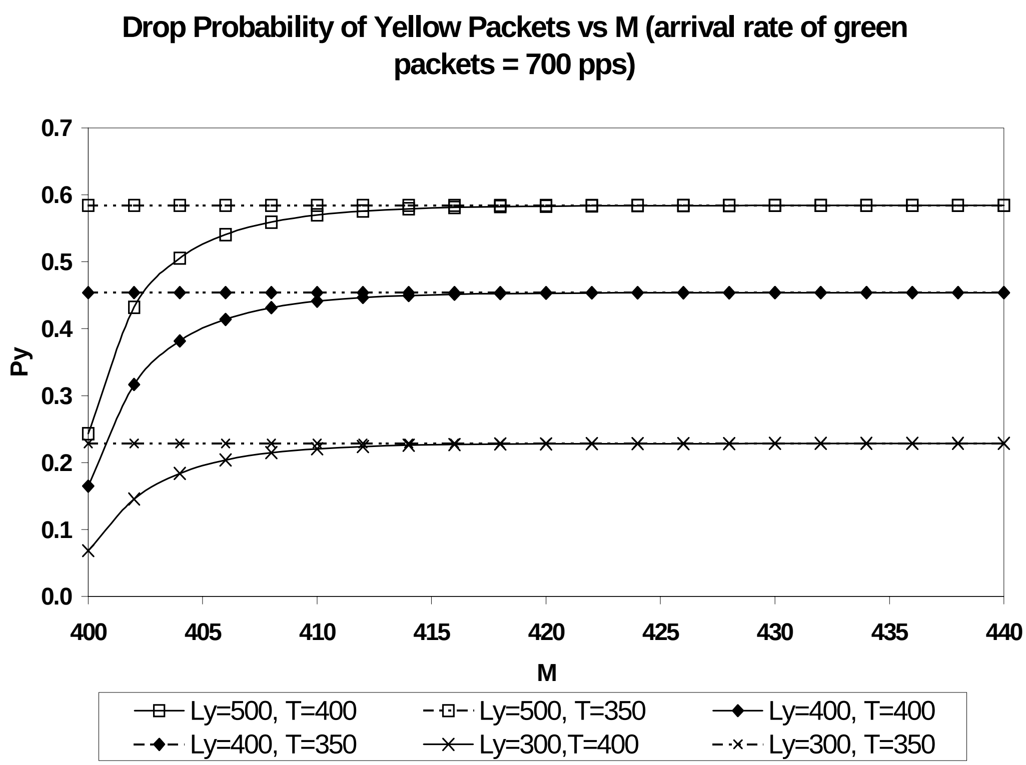

Figure 4.3 - P_y^d **vs. M**

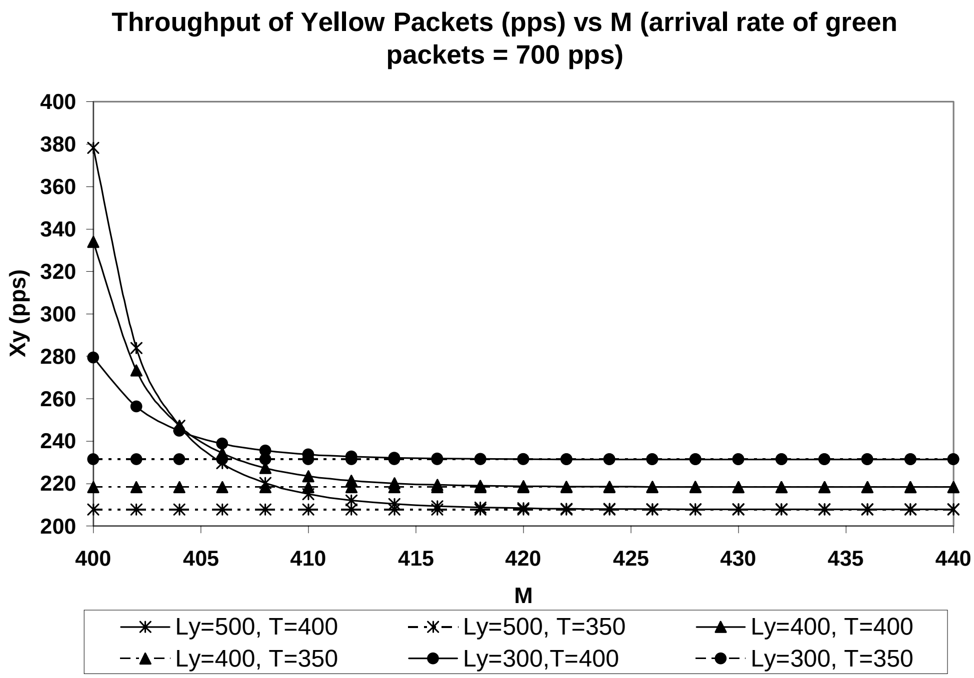

Figure 4.4 - X_y **vs M**

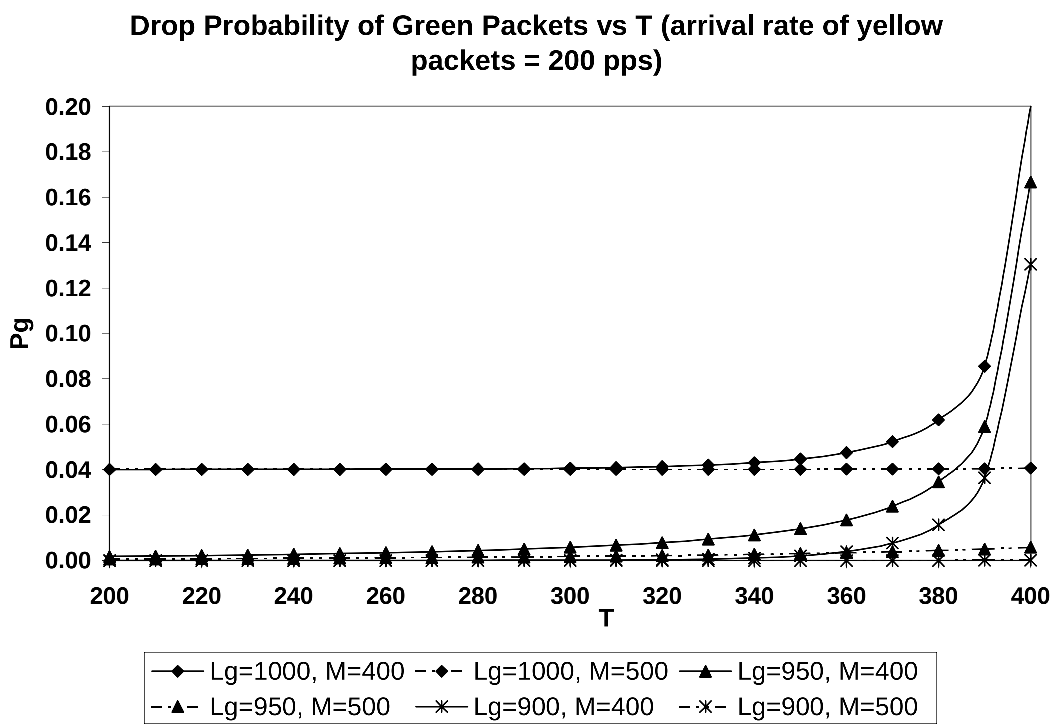

Figure 4.5 - P_g^d **vs. T**

The throughput X_g of green packets as a function of *T* is shown in Fig. 4.6. The figure reflects the behavior of P_g^d shown in Fig. 4.5. When P_g^d increases, X_g decreases.

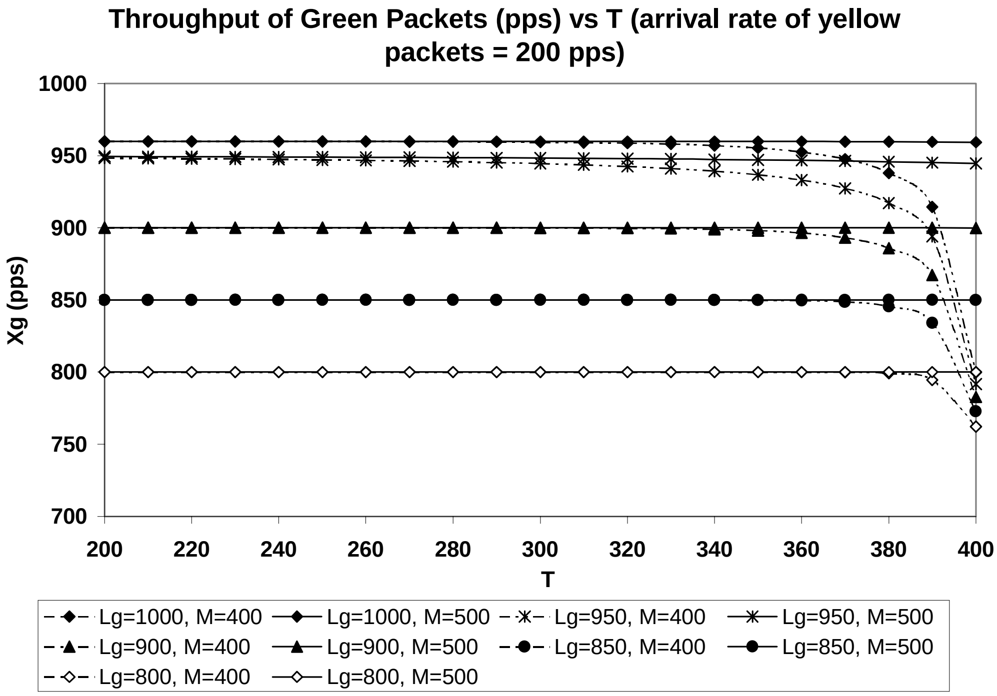

Figure 4.6 - X_g vs T.

Figure 4.7 shows the probability P_g^d that green packets are dropped as a function of the arrival rate λ_g of green packets. As it can be seen, for T = 350, green packets start to be dropped after the arrival rate λ_g exceeds 750 pps. This means that the NSP is not being able to guarantee the Committed Information Rate. The figure also shows that for values of M greater than or equal to 400, the drop probability P_g^d has its minimum value.

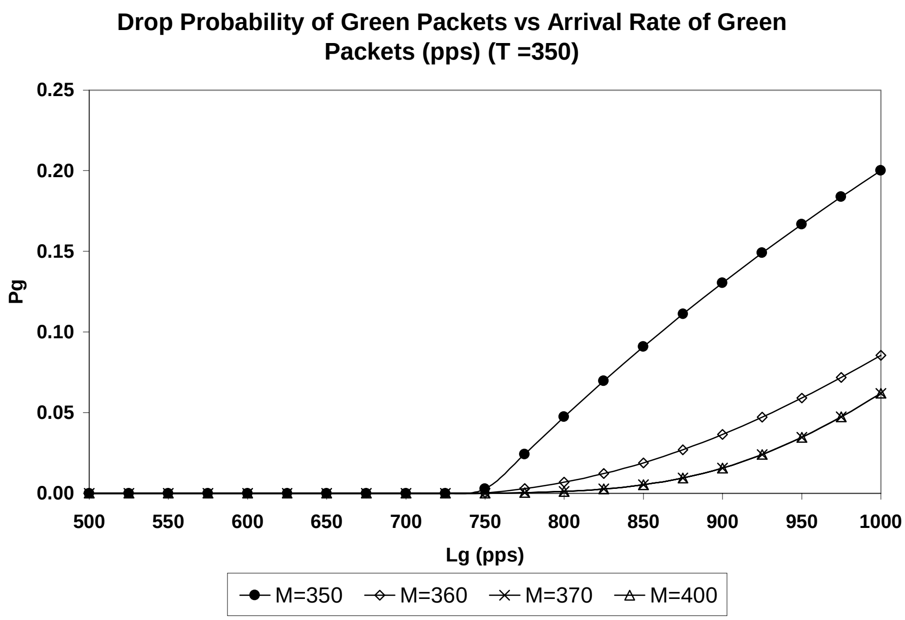

Figure 4.7 - P_g^d vs. λ_g

As the various figures show, there is a very interesting and sensitive interplay between the parameters *T* and *M*.

5. Concluding Remarks

Network Service Providers (NSPs) need to guarantee a committed quality of service to their customers. These customers share NSP resources such as routers and links. Admission and congestion control policies can be used to allocate resources to different users and guarantee quality of service. This paper presented a Markov Chain based analytic model used to determine the proper bandwidth allocation thresholds so that packet drop probabilities can be minimized and throughput maximized.

During the course of the analysis, it was determined that the FIFO-based, single queue architecture employed in traditional IP routers and Frame Relay Switches was not an appropriate choice for the WorldCast eBurst service. The requirement to share unused channel capacity, while ensuring the availability of customer CIRs, necessitated the use of a multiple-queue router architecture. Subsequent analysis and testing led to the adoption of a Class-Based Queuing bandwidth management scheme. The CBQ analytical and test results will be described in a forthcoming paper.

References

[Klei75] Kleinrock, L., Queuing Systems, Vol. I: Theory, John Wiley & Sons, 1975.

[Little61] Little, J. D. C., "A Simple Proof of $L=\lambda W$," Operations Research, 9, 1961, pp.383-387.