

Dimensioning Bandwidth for Elastic Traffic

Zhong Fan

Marconi Labs Cambridge
William Gates Building, JJ Thomson Avenue
Cambridge CB3 0FD, UK
`zhong.fan@marconi.com`

Abstract. In this paper, we discuss the issue of dimensioning Internet access lines for elastic traffic. This is important for Internet service providers (ISPs) because over-dimensioning wastes precious bandwidth resources, while under-dimensioning generally leads to less satisfactory quality of service (QoS) perceived by subscribers. Our discussion is based on the M/G/R processor sharing model which characterizes TCP traffic at flow level. Our analysis demonstrates the impact of a number of key factors (and their relations) on the dimensioning procedure. We consider two dimensioning methods based on different QoS criteria. It is found that the method based on the delay factor is superior in that both the average delay (throughput) and blocking performance targets can be satisfied. Numerical and theoretical analyses also illustrate that significant multiplexing gain can be achieved for elastic flows and this gain increases with burstiness.

1 Introduction

It has been recognized that there are generally two classes of traffic in the current Internet, namely, stream traffic and elastic traffic [1]. Typical stream services are real-time video and voice services, while elastic services could be file transfers, emails, web pages and other data traffic based on TCP. At the moment a large amount of Internet traffic is elastic and therefore it is essential for ISPs to dimension Internet access lines properly to cater for the service needs of elastic traffic. Dimensioning should also allow for statistical multiplexing to achieve better utilization of network resources.

Over the past few years there have been extensive studies on IP traffic characterization, in particular, long-range dependence and self-similarity (see e.g., [2]). While many papers have focused on packet level behavior, recently a number of studies show that processor sharing (PS) models provide a simple and accurate characterization of elastic IP traffic at flow level [3] [4] [5] [6]. Nabe et al. [3] use an M/G/1 PS model to discuss a design methodology of the Internet access network as well as document caching at a proxy server. A drawback of the M/G/1 PS model is that it assumes that one TCP connection is able to utilize the total link capacity by its own, which is not true in reality. In [5] and [7], PS models have been used to demonstrate the need of admission control for TCP

flows. In general, PS models are able to successfully capture the elastic properties of traffic generated by closed loop control transport protocols (e.g., TCP) without going into complicated details of packet level traffic characteristics.

In this paper, we discuss the issue of dimensioning Internet access lines for elastic traffic. This is very important for Internet service providers because over-dimensioning wastes precious bandwidth resources, while under-dimensioning generally leads to less satisfactory quality of service (QoS) perceived by subscribers. Our discussion is based on the M/G/R processor sharing model which characterizes TCP traffic at flow level. The performance criteria of dimensioning could be average transfer delay and throughput, both of which are related to a so-called delay factor. We also consider parameters such as blocking probability and multiplexing gain. Among the two dimensioning methods in which we are interested here, we have found that the method based on the delay factor is superior in that both the average delay (throughput) and blocking performance targets can be satisfied. Both numerical and theoretical analyses illustrate that significant multiplexing gain can be achieved for elastic flows and this gain increases with burstiness.

2 The M/G/R PS Model

Here we consider a simple scenario (shown in Figure 1) where subscribers are connected to an access multiplexer via customer access lines (e.g., ADSL lines) and then to the core network (where servers reside) via an access trunk line. In this context, the trunk line must have enough capacity to accommodate both upstream and downstream traffic loads.




Fig. 1. Access network

We assume that elastic traffic is generated by file transfer applications. The flow (TCP connection) arrival process is Poisson¹ and the file size distribution

¹ The Poissonian assumption is appropriate when the considered link is shared by a very large number of users [5].

has heavy tails, e.g., Pareto distribution. Actually, an important advantage of PS models is that results derived from PS models are insensitive to file size distributions [8]. The file transmission rates are controlled by the TCP feedback algorithm as a function of network congestion. When TCP works ideally, the access trunk line can be modelled as a processor sharing queue. It is known that with the PS scheduling discipline large files do not delay small ones too much when compared with FIFO scheduling [9].

Let r_p and C denote the limited peak rate of an individual subscriber (peak access line rate, e.g., modem speed or the rate limited by the maximum TCP window size) and the trunk line capacity respectively ². Then the link appears like a PS system with R servers where R is an integer and $R = C/r_p$, hence the name M/G/R PS queue. If θ is the average file size and λ is the average flow arrival rate, then the traffic load (or utilization) ρ is $\theta\lambda/C$. It has been shown in [9] [10] that the mean conditional sojourn time $T(x)$ for a file of size x is given by

$$T(x) = \frac{x}{r_p} \left(1 + \frac{E_2(R, R\rho)}{R(1-\rho)} \right) \quad (1)$$

where E_2 represents Erlang's second formula:

$$E_2(R, R\rho) = \frac{A}{B(1-\rho) + A} \quad (2)$$

where $A = (R\rho)^R/R!$, $B = \sum_{i=0}^{R-1} (R\rho)^i/i!$.

As in [9], define a delay factor f_R as

$$f_R = 1 + \frac{E_2(R, R\rho)}{R(1-\rho)}. \quad (3)$$

Then $T(x)$ can be re-written as $T(x) = \frac{x}{r_p} f_R$. And the mean throughput γ is given by

$$\gamma = x/T(x) = r_p/f_R. \quad (4)$$

The delay factor f_R represents the increase of the average file transfer time (and decrease of the average throughput) due to link congestion. For the special case of $R = 1$ (M/G/1 PS), $f_R = \frac{1}{1-\rho}$, and $\gamma = C(1-\rho)$. Note that in [8], a similar demerit factor is introduced.

It has been advocated by Roberts et al. [7] that TCP admission control should be implemented so that flows sharing a bottleneck can achieve some minimum throughput. Let r_m denote the minimum fair share which can be used as an admission control threshold. Then the upper limit of the number of admitted flows, N , is C/r_m . In this case, the blocking probability, F , is a function of the parameters N , R and ρ [9]:

$$F(N, R, \rho) = \frac{AD(1-\rho)}{B(1-\rho) + A(1-D\rho)} = \frac{E_2D(1-\rho)}{1 - E_2D\rho} \quad (5)$$

² Here for simplicity, we assume that all subscribers have the same maximum access rate. For different access rates, we could use an average value of R in dimensioning as suggested by [9].

where A, E_2 and B are as defined previously, and D is given by $D = \rho^{N-R}$. Note that when $R = N$, (5) reduces to Erlang's first formula. When $R = 1$, (5) becomes

$$F(N, \rho) = \rho^N(1 - \rho)/(1 - \rho^{N+1}). \tag{6}$$

Figure 2 shows delay factor f_R as a function of ρ with different R . It can be seen that when load is low, $f_R \approx 1$, therefore $\gamma \approx r_p$. In this case, throughput is nearly full access line rate. However, as load increases, f_R increases dramatically, thus throughput drops sharply. The delay improves while R increases. This can also be seen in Figure 3, in which normalized throughput (γ/C) is shown as a function of ρ . It is easy to see that as ρ is very close to 1, $\gamma \approx C(1 - \rho)$.




Fig. 2. Delay factor vs. load

Figure 4 shows the relationship between blocking probability and f_R with different N . Here $R = 10$. Since both F and f_R are monotonically increasing functions of ρ , F increases with f_R . We can see for more elastic traffic ($N \gg R$), the blocking probability is smaller. In terms of dimensioning, we could provision the trunk line so as to have both a small blocking probability (say, 0.001) and a desirable delay factor (say, 1.01).

In [10] the accuracy of the M/G/R PS model is studied using simulations and the basic applicability of this model to access link dimensioning is confirmed. We further their work by considering two different dimensioning procedures and investigating in more depth the impact of the delay factor and other parameters (such as blocking probability) on network performance.

3 Access Trunk Line Dimensioning

One possible dimensioning method uses the blocking probability as a QoS criterion [1]. Similar to the situation in telephony networks, blocking probability




Fig. 3. Normalized throughput vs. load




Fig. 4. Blocking probability against delay factor

has a significant impact on user satisfaction. As an alternative, we can also base our dimensioning decision on the delay factor described in the previous section. For elastic Internet services such as file transfer and web traffic, large delay contributes greatly to the user-perceived quality degradation.

3.1 Dimensioning Method One

For link dimensioning purposes, the above model needs to be extended to arbitrary link rates, i.e., R does not have to be an integer [5]. In this case, let R denote the integer part of C/r_p . Then we have

$$T(x) = \frac{x}{r_p} f_R \tag{7}$$

and

$$\gamma = r_p / f_R, \tag{8}$$

where

$$f_R = 1 + \frac{(1 - (\frac{C}{r_p} - R)(1 - \rho))r_p E_2(R, C\rho/r_p)}{C(1 - \rho)}. \tag{9}$$

To dimension the trunk line for elastic traffic, we use the delay factor f_R as the QoS measure since f_R determines both average transfer time (Eqs. (1) and (7)) and throughput (Eqs. (4) and (8)) for a given flow. Obviously, f_R has to be chosen greater than 1 (but preferably close to 1). For a given f_R , we can numerically solve (9) to obtain the desired capacity value, C .

Figure 5 shows the dimensioning result of the trunk capacity (normalized capacity w.r.t. r_p) for medium to high loads with different target delay factors. As shown in the figure, f_R has significant impact on capacity, especially at high loads. For instance, the required capacity for f_R of 1.2 is roughly double that for f_R of 1.5 at the load of 0.95. Therefore the delay factor is indeed an appropriate QoS measure for elastic traffic. Assume admission control is implemented to ensure that the load is smaller than 1. In this case, the blocking probabilities as function of ρ are shown in Figure 6 for the dimensioning case of $f_R = 1.5$. All the blocking probabilities are very small (close to zero) when $\rho < 0.8$. However, at high loads (say, $\rho = 0.95$), blocking probabilities differ significantly for different N . It can be seen that for a dimensioned capacity satisfying the delay factor criterion, it is possible to achieve a target blocking probability (especially at high loads) by choosing the value of N .




Fig. 5. Trunk line capacity vs. load

As another more realistic example, we use the data traces from [11], which was obtained in an ADSL field trial in Germany in 1998. Some of the parameters




Fig. 6. Blocking probability vs. load when $f_R = 1.5$

observed for HTTP over TCP/IP traffic in active client access sessions are: downstream access line rate $r_p = 2.5$ Mbps, mean rate $m = 10.5$ kbps. Figure 7 shows the dimensioning result for the trunk capacity with $f_R = 1.01$. Also shown is the required capacity calculated based on the sum of the mean rates of all sources.

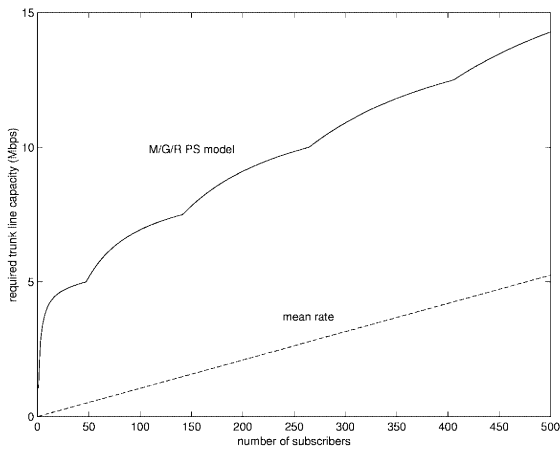


Fig. 7. Trunk line capacity vs. number of subscribers

3.2 Dimensioning Method Two

Figure 8 plots the dimensioning result when we use the blocking probability as the QoS criterion. The target blocking probability is 0.001 and the peak access rate r_p is 2.5 Mbps. It can be seen that as the minimum throughput r_m increases, the required capacity increases. Figure 9 shows the corresponding delay factor f_R under the dimensioned trunk capacities. As expected, f_R improves (decreases) as the minimum throughput r_m increases. For larger r_m , f_R remains more or less the same over a wide range of ρ . However, there are some oscillations for the case of $r_m = 0.5$ Mbps and f_R increases quite significantly as ρ approaches 1. This is because as ρ grows, the required trunk line capacity increases, therefore R increases. Thus, the fact that ρ and R have opposite impact on f_R (as shown in Figure 2) explains the oscillation phenomenon.

From the above discussion, we can see that for the dimensioning based on the blocking probability and minimum throughput, it is difficult to obtain a desirable delay factor. Hence, it can result in unsatisfactory average delay and throughput performance. For example, for $r_m = 0.5$ Mbps, f_R is always above 1.25. In this sense, the dimensioning scheme based on the delay factor is more suitable because the blocking performance can also be tuned therein.




Fig. 8. Trunk line capacity vs. load

4 Multiplexing Gains

Define the multiplexing gain G as

$$G = \frac{nr_p}{C}, \tag{10}$$

where n is the number of sources. The maximum possible value of gain is obtained when dimensioning is based on the mean rate, i.e., $C = nm$. Therefore, $G_{max} =$




Fig. 9. Delay factor vs. load

r_p/m . Intuitively, this means that for highly bursty traffic, with $m \ll r_p$, G can be very large. However, this maximum gain cannot be attained in reality, because the average bandwidth dimensioning is unacceptable in terms of QoS.

Figure 10 shows G versus n with an f_R of 1.01, using the data in [11]. It can be seen that G increases with n , which demonstrates the benefit of exploiting the statistical features of elastic traffic. In fact, when ρ is close to 1, $r_p/f_R \approx C(1-\rho)$, hence we have

$$C \approx r_p/f_R + nm. \tag{11}$$

So,

$$G \approx \frac{nr_p}{r_p/f_R + nm} = \frac{r_p}{n/f_R + m}. \tag{12}$$

It is clear that as $n \rightarrow \infty, G \rightarrow G_{max}$. (12) also shows that G increases with f_R .

Next we investigate the impact of the source activity factor $p = m/r_p$ on the multiplexing gain G . Figure 11 shows G as a function of p when n is 100 and f_R is 1.01. G decreases with p . In other words, multiplexing gains increase with burstiness. As a special case, it can be seen in (12) that G is indeed decreasing with p . This is also consistent with the results in [12] where multiplexing gains for stream traffic in an ATM QoS context is considered.

5 Discussion

For small documents, TCP slow start dominates the file transfer phase. In [6], the effect of the slow start phase and round trip times (RTT) are taken into consideration. The transfer time is thus a number of RTTs more than the result predicted by the PS model. However, when the document size is sufficiently




Fig. 10. Multiplexing gain vs. number of sources




Fig. 11. Multiplexing gain vs. source activity factor

large, the M/G/R PS model is accurate enough for dimensioning purposes [10]. Moreover, since the mean transaction time for small files is small anyway, the inaccuracy of the model does not affect users' perception of quality very much. On the other hand, further refinement of the model to deal with short file transfers is a direction of future research.

In some cases, the TCP window control mechanism is not “ideal”. For example, during congestion when packets are lost and retransmissions become necessary, the successful transfers of files proceed at a total rate that can be well below C . In this case, we need to introduce a link efficiency factor α ($\alpha \leq 1$) such

that the average file transfer rate is αC . A discussion on how to estimate α can be found in [8]. From a dimensioning point of view, we can have a safety margin for the link capacity, i.e., set the link capacity C as C'/α , where C' is obtained using the PS model described before.

6 Conclusion

In this paper, we discuss the issue of dimensioning Internet access lines for elastic traffic. Our discussion is based on the M/G/R processor sharing model which characterizes TCP traffic at flow level. Our analysis demonstrates the impact of a number of key factors (and their relations) on the dimensioning procedure. We consider two dimensioning methods based on different QoS criteria. It is found that the method based on the delay factor is superior in that both the average delay (throughput) and blocking performance targets can be satisfied. Both numerical and theoretical analyses illustrate that significant multiplexing gain can be achieved for elastic flows and this gain increases with the number of sources and traffic burstiness.

Acknowledgements. I would like to thank A. Berger of Akamai and D. Heyman of AT&T for their useful communications.

References

1. J. Roberts. Realizing quality of service guarantees in multi-service networks. In *IFIP PMCCN*, 1997.
2. A. Feldmann et al. Data networks as cascades: explaining the multifractal nature of Internet WAN traffic. In *ACM SIGCOMM*, 1998.
3. M. Nabe et al. Analysis and modelling of world wide web traffic for capacity dimensioning of Internet access lines. *Performance Evaluation*, 34:249–271, 1998.
4. A. Berger and Y. Kogan. Dimensioning bandwidth for elastic traffic in high-speed data networks. *IEEE/ACM Trans. Networking*, 8(5):643–654, 2000.
5. S. Ben Fredj et al. Statistical bandwidth sharing: a study of congestion at flow level. In *ACM SIGCOMM*, 2001.
6. J. Beckers et al. Generalized processor sharing performance models for Internet access lines. In *9th IFIP Conference on Performance Modelling and Evaluation of ATM and IP Networks*, 2001.
7. L. Massoulié and J. Roberts. Arguments in favour of admission control for TCP flows. In *ITC 16*, 1999.
8. D. P. Heyman et al. A new method for analysing feedback-based protocols with applications to engineering web traffic over the Internet. In *ACM SIGMETRICS*, 1997.
9. K. Lindberger. Balancing quality of service, pricing and utilization in multiservice networks with stream and elastic traffic. In *ITC 16*, 1999.
10. A. Riedl et al. Investigation of the M/G/R processor sharing model for dimensioning of IP access networks with elastic traffic. In *First Polish-German Teletraffic Symposium*, 2000.

11. J. Charzinski. Fun factor dimensioning for elastic traffic. In *ITC Specialist Seminar on IP Measurement, Modeling and Management*, 2000.
12. Z. Fan and P. Mars. Multiplexing gains in ATM networks. In D. Kouvatsos, editor, *Performance Analysis of ATM Networks*, pages 377–395. Kluwer Academic Publishers, 2000.