

Dimensioning high speed IP access networks

T. Bonald, P. Olivier, J. Roberts

France Telecom R&D

38-40, rue du Général Leclerc, 92794 Issy-les-Moulineaux Cedex 9, France

This paper discusses the definition of simple dimensioning rules for high speed IP access networks carrying data traffic. We notably provide formulas relating capacity, demand and performance allowing dimensioning for a target quality of service expressed in terms of useful per-flow throughput. These formulas derive from a data traffic model equivalent of the Engset model for telephone access networks. Performance is shown to be largely independent of precise traffic characteristics. The key dimensioning parameter is offered traffic defined as the average data rate a user would generate in the absence of congestion.

1. Introduction

To correctly dimension network links it is essential to understand the three-way relation between link capacity, expressed demand and realised quality of service. The typical example of such a relation for access networks is the Engset loss formula giving the call blocking probability in terms of the number of circuits and the number of sources and their offered traffic [6]. In this paper we show that similar relations exist for data traffic in access networks. In this case it is appropriate to model traffic at flow level and to express dimensioning objectives in terms of flow throughput.

We consider a link of capacity C shared by N users each having an access rate c . The link in question might typically connect a set of xDSL lines to the first router of an IP network. The access rate in this case corresponds to the modem rate (e.g., around 1 Mbit/s for an ADSL user). In the present paper we only consider *elastic* data traffic (traffic using TCP) and we assume the flows in progress share link capacity perfectly fairly. Performance is measured in terms of the useful rate d , a measure of expected throughput, while demand is expressed through the individual offered traffic a , equal to the overall rate a user would contribute in the absence of congestion.

The following constitutes the basic dimensioning formula. The capacity C necessary to handle per user offered traffic a with a useful rate d ($d < c$) is simply:

$$C = \frac{N}{1/a + 1/d - 1/c}. \quad (1)$$

If the target useful rate is equal to the access rate ($d = c$), it is sufficient to provide capacity greater than the overall offered traffic: $C > Na$. The above formulas in fact constitute an approximation derived from an accurate model of the considered stochastic system. This model is the data traffic equivalent of the Engset model for dimensioning telephone access networks.

We use a fluid traffic model and assume the data transport protocol (generally TCP) realizes perfectly fair bandwidth sharing between competing flows. This model was first proposed for a finite set of homogenous sources by Heyman et al. [7] and further developed by Berger and Kogan [2]. The insensitivity of the performance of the homogeneous model with respect to the distribution of flow size and think time duration was recognized in [7] where the authors drew on known results for the generalized Engset model [6]. The applicability of the model to a hierarchical model of data traffic where flows occur in sessions was identified in [1] and [5], extending the insensitivity properties to quite general and realistic flow arrival processes for both infinite and finite number of sources. The underlying stochastic model is a network of processor sharing queues. The general properties of this model and its application to the performance evaluation of data networks have been explored in recent papers by Bonald and Proutière [3,4].

Our objective in the present paper is to draw the practical consequences of the above models for dimensioning a data access network. We seek simple dimensioning rules and guidelines enabling the network operator to economically meet performance targets. In the following sections, we derive formula (1), discuss the performance implications and show how it can be applied. This formula is strictly applicable only in the case where users have homogeneous traffic characteristics. We consider extensions to a heterogeneous user population and derive the impact on overall performance of unbalanced load and unfair bandwidth sharing.

2. Dimensioning parameters

We discuss the relationship between the traffic and performance parameters introduced above based on an assumed model of user traffic.

2.1. Structure of a user session

Traffic generated by a user is typically composed of a succession of flows interspersed by periods of inactivity, as illustrated in Figure 1. Each flow corresponds to the transfer of a digital document (Web page, e-mail, stored video sequence,...) or several documents transferred in series or in parallel (elements of a Web page, successive e-mails,...). The inactivity period typically corresponds to the time the user consults the transferred document and is referred to as the “think time”.




Figure 1. Example of an ADSL user session

Clearly, in general, the size of successive flows, their duration and the duration of ensuing think times are highly variable and correlated random quantities. We use the following notation:

- V is the mean flow volume (in bits);
- T is the mean flow duration (in seconds);
- S is the mean think time duration (in seconds).

2.2. QoS parameters

User perceived quality of service depends essentially on the mean document transfer time or, equivalently, on the *useful rate* defined as the ratio of mean flow volume to mean duration: $d = V/T$. We choose d as a performance parameter rather than the expected throughput of an arbitrary document for reasons of tractability. It turns out that in the models presented below, d has a more general interpretation. The expected time to transfer a document of size s is s/d . Thus d is a throughput measure independent of the document size. In particular, d can be readily estimated in practice by observing the transfer of a long document and dividing its size by the realized response time.

2.3. Traffic parameters

The offered traffic corresponds to the traffic a user would generate if the access rate c were always available (assuming document size and think time duration are independent of realized performance). This is an intrinsic user characteristic which is independent of other users and of the QoS provided by the network. It is given by the following expression:

$$a = \frac{V}{V/c + S}.$$

The measured user traffic, the “carried traffic” denoted b , is always less than the offered traffic. We have:

$$b = \frac{V}{T + S}.$$

Carried traffic depends on the QoS provided by the network. From the above relations we deduce the following simple relation between carried traffic and useful rate:

$$\frac{1}{a} + \frac{1}{d} = \frac{1}{b} + \frac{1}{c}. \tag{2}$$

2.4. Summary of rate parameters

The following table summarizes the different traffic and QoS parameters introduced above distinguishing those that depend on link capacity (via dynamic sharing between active users) and those that do not. The arrows \rightarrow mean “is greater than”.

	Exogenous parameters (capacity independent)		Endogenous parameters (capacity dependent)
Traffic parameters	offered traffic, a	\rightarrow	carried traffic, b
	\uparrow		\uparrow
QoS parameters	access rate, c	\rightarrow	useful rate, d

2.5. Estimating offered traffic

The offered traffic a is the key parameter for network dimensioning. To attempt to measure a directly by identifying successive periods of activity and think times is particularly prone to error in view notably of the imprecise definition of these phases (depending on the assumed threshold distinguishing an inter-flow gap and an inter-packet gap within the same flow). Relation (2) linking a to quantities b and d provides a more convenient alternative, as carried traffic b can be estimated by dividing the total volume of data emitted in a busy period by the length of that period while the useful rate d can be estimated as discussed in §2.2.

3. Performance model

We now derive the relation between link capacity, offered traffic and realized performance (in terms of useful rate). We first describe the underlying assumptions.

3.1. Assumptions

We consider a link of capacity C shared by N users having the same access rate c . We assume bandwidth is shared perfectly fairly between flows in progress, i.e., when x flows are in progress, each flow realizes the instantaneous rate $\min\{C/x, c\}$. This instantaneous rate thus jumps to a new value on each flow arrival or departure. The fair sharing assumption is reasonable when each flow corresponds to a TCP connection that is not rate limited elsewhere in the network [1].

The impact of unfair bandwidth sharing, due for example to multiple parallel TCP flows or flows with different round trip times, is considered in §4.4. It turns out to be of secondary importance, especially when the access rate c is considerably smaller than link capacity C . We also ignore the exact operation of TCP by which the rate of a flow adjusts gradually and imperfectly to the current fair share. One impact of ignoring this behaviour is clearly to significantly underestimate the duration of short flows. It has a much smaller effect on the estimated performance of larger flows, however, which depends more on the overall volume of traffic than the precise rate at which particular flows are transmitted. The considered model derives from an idealization of system behaviour that allows us to conveniently appraise the impact on performance of different characteristics and parameters. We believe the general conclusion we draw from it would remain valid if it were possible to account exactly for the way TCP realizes bandwidth sharing.

We model each user as a *stationary* random succession of flows and think times. This corresponds to the traditional assumption of stationary traffic in the busy hour and is justified by the results of traffic measurements in operational networks. Successive flow volumes and think time durations can have arbitrary distributions and be mutually dependent (imagine, for instance, a succession of short flows in a Web session culminating in the download of a large document and long ensuing think time).

3.2. Relation between capacity, offered traffic and useful rate

Let a_i be the traffic offered by user i , for $i = 1, \dots, N$. This parameter is between 0 (if user i is inactive in the busy hour) and c (if user i continuously transfers data throughout the busy hour). Given the above assumptions, it is possible to express the useful rate of user i as a function of link capacity C and user offered traffics a_1, \dots, a_N . A highly

significant observation is that the useful rate does not depend on the distribution of flow volume or that of the think time duration or any correlation between the successive values of these random variables. This follows from results on the insensitivity of bandwidth sharing in a stochastic system introduced in [1] and generalized in [4].

Homogeneous demand.

First assume all users have the same offered traffic a . Let r denote the ratio V/S so that $a = rc/(r + c)$. Let $x(t)$ denote the number of flows in progress at time t . Under the assumptions of Section 3.1, $x(t)$ converges to a stationary regime x having the distribution:

$$\pi(x) = \pi(0) \binom{N}{x} \left(\frac{r}{c}\right)^x \prod_{i=1}^x \frac{i}{\min(i, C/c)}, \quad (3)$$

where $\pi(0)$ is determined by the usual normalizing condition.

The distribution π is independent of the distributions of flow volumes and think time durations and their correlation [1]. Moreover, it depends only on the ratio of their means. The useful rate can be deduced on applying Little's formula: the flow arrival rate is equal to the mean number of inactive users $N - E[x]$ divided by the mean think time. It is also equal to the mean number of active flows divided by the mean flow transfer time. We have therefore: $(N - E[x])/S = E[x]/T$, and deduce:

$$d = r \frac{N - E[x]}{E[x]}. \quad (4)$$

Instead of applying the above analysis we can derive approximations from relation (2) on making some further simplifying assumptions. We suppose there exist two operating regimes: a saturated regime where the link is always fully used and a transparent regime where all flows are transmitted at rate c . In the saturated regime we have $b = C/N$ and deduce from (2):

$$d \approx \frac{1}{N/C + 1/c - 1/a}. \quad (5)$$

This is just another way of writing (1). We further assume the transparent regime $d \approx c$ occurs when $a \leq C/N$.

Figure 2 left shows the normalized useful rate (expressed as a fraction of the access rate) as a function of capacity for an offered traffic $a = c/3$ (implying $r = c/2$). Note that approximation (5) provides an accurate estimate for any reasonably large value of N . Figure 2 right illustrates the quality of the approximation for different offered traffics with the number of users N set to 100. The approximation correctly predicts the two operating regimes and is more accurate as N increases.

Heterogeneous demand.

Consider now a system with K user classes where class k has N_k users all with offered traffic a_k , for $k = 1, \dots, K$. Let r_k denote the ratio V_k/S_k , where V_k is the mean flow volume and S_k is the mean silence duration for users of class k so that $a_k = r_k c / (r_k + c)$. Let




Figure 2. Relation between capacity and useful rate for various numbers of users N (left) and various mean offered traffics a (right).

$x_k(t)$ denote the number of active users of class k at time t . With the assumptions of §3.1, the vector $x(t) = (x_1(t), \dots, x_K(t))$ converges to a stationary state x with distribution:

$$\pi(x) = \pi(0) \prod_{k=1}^K \binom{N_k}{x_k} \left(\frac{r_k}{c}\right)^{x_k} \prod_{i=1}^{\sum x_k} \frac{i}{\min(i, C/c)}. \quad (6)$$

The normalizing condition again determines $\pi(0)$.

As previously, the stationary distribution does not depend on the distributions of flow volume and think times but only on the ratio of their means. Applying Little's formula, as above for homogeneous demand, we have $(N_k - E[x_k])/S_k = E[x_k]/T_k$, so that:

$$d_k = r_k \frac{N_k - E[x_k]}{E[x_k]}. \quad (7)$$

The useful rate d_k is again the same for flows of any size, i.e., the expected time for a user of class k to transfer a document of size s is s/d_k .

Noting that V, T, S are respective weighted means of V_k, T_k, S_k , with weights proportional to $N_k/(T_k + S_k)$, $k = 1, \dots, K$, we deduce that the *mean* useful rate over all flows of all users is:

$$d = V/T = \frac{\sum_k r_k (N_k - E[x_k])}{\sum_k E[x_k]}. \quad (8)$$

If we again define

$$r = V/S = \frac{\sum_k r_k (N_k - E[x_k])}{\sum_k (N_k - E[x_k])},$$

this formula reduces to (4) with $E[x] = \sum_k E[x_k]$.

In practice, whenever the traffic of any user is small compared to the overall offered traffic, the useful rate of that user is approximately equal to the mean useful rate d .

This observation can be explained by the MUSTA (“Moving Units See Time Averages”) property [8]. This states that the distribution of x at class k user flow arrival instants is the stationary distribution π conditioned on the fact that that user was inactive. If each user only counts for a small fraction of the overall traffic, the conditional distribution for class k users is nearly the same as the stationary distribution π . We deduce the following approximations:

- The useful rate is approximately the same for all users and is close to the *mean* useful rate d (equal to the ratio of mean volume to mean transfer duration for all users).
- When offered traffic $\sum_{k=1}^K N_k a_k$ is less than link capacity C , the mean useful rate d is close to the access rate c .
- When offered traffic $\sum_{k=1}^K N_k a_k$ is more than link capacity C , the carried traffic $\sum_{k=1}^K N_k b_k$ is close to the capacity C . Using (2), we deduce the following relation satisfied by the mean useful rate:

$$C \approx \sum_{k=1}^K \frac{N_k}{1/a_k + 1/d - 1/c}. \quad (9)$$

4. Applications

In this section we present a number of applications of dimensioning formula (9). We fix a target useful rate d and determine the required capacity per user C/N .

4.1. Dimensioning based on mean offered traffic

First suppose all users have the same offered traffic a . Using (1), the required capacity per user depends on a through the function:

$$f(a) = \frac{1}{1/a + 1/d - 1/c}. \quad (10)$$

For example, an access rate $c = 500$ Kbit/s, an offered traffic $a = 50$ Kbit/s and a target useful rate of 100 Kbit/s, lead to required per-user capacity of 36 Kbit/s. Note that these results are independent of the number of users N and there is no scale economy¹.

When users have distinct offered traffics a_1, \dots, a_K , we again use (9) to deduce the required capacity per user:

$$\frac{C}{N} \approx \frac{1}{N} \sum_{k=1}^K N_k f(a_k). \quad (11)$$

Noting that function f is concave, we deduce that required capacity per user is greatest for a given overall demand when all users have the same offered traffic a satisfying:

$$a = \frac{1}{N} \sum_{k=1}^K N_k a_k. \quad (12)$$

¹Scale economy would occur, as for the telephone network dimensioned using Engset’s formula, if we were to apply flow based admission control to guarantee a minimum per flow throughput and used a blocking probability criterion for dimensioning.

A conservative dimensioning rule is therefore to ensure capacity is greater than $Nf(a)$. The useful rate is guaranteed to be greater than the target d .

4.2. Impact of user activity

Formula (10) provides a simple dimensioning rule based only on the mean offered traffic (12). A more accurate rule would take account of the *user activity ratio*. Assume only a fraction τ of users is actually active during the considered busy hour. As the other users have zero offered traffic, the active users offer a/τ where a denotes the overall mean offered traffic. The activity ratio τ is necessarily greater than a/c . From (11), the required capacity per user is:

$$\frac{C}{N} \approx \tau f(a/\tau).$$

Clearly, $\tau f(a/\tau)$ is less than $f(a)$. The assumption of homogeneous user activity thus constitutes a worst case for dimensioning purposes. Figure 3 left shows the impact of the activity ratio on the required capacity for an access rate $c = 500$ Kbit/s and a mean offered traffic $a = 50$ Kbit/s.




Figure 3. Required capacity depending on the user activity ratio (left) and the proportion of insatiable users (right)

4.3. Unequal traffic distribution

Consider now the impact of an unequal distribution of offered traffic among users. Assume for instance that a certain proportion of users σ is insatiable in the sense that their offered traffic is equal to the access rate c . Mean offered traffic for the other users is then: $a' = (a - \sigma c)/(1 - \sigma)$. The fraction of insatiable users is necessarily less than a/c . Assuming only a fraction τ of the other users are active in the busy hour, we deduce from (11) that the required per user capacity is:

$$\frac{C}{N} \approx \sigma d + (1 - \sigma)\tau f(a'/\tau).$$

Note that the approximation $d_k \approx d$ remains valid as long as each user contributes only a small fraction of the overall traffic. This is true in particular when the number of insatiable users is not very small (more than 10, say). Figure 3 right shows how required capacity depends on the proportion of insatiable users for an access rate $c = 500$ Kbit/s, a mean offered traffic $a = 50$ Kbit/s and an activity ratio $\tau = 50\%$.

4.4. Unfair bandwidth sharing

Several factors can invalidate the assumption that bandwidth is shared perfectly fairly between active flows, including the fact that a flow may consist of several TCP connections in parallel and that realized throughput depends on the round trip time which is different for different flows. To evaluate the impact of unequal sharing we consider an extreme case where a fraction of users (class 1 users) have priority over the others (class 2 users). We retain the assumption that rate changes occur instantaneously, ignoring the impact of TCP slow start. By priority we mean that class 2 users have zero instantaneous throughput whenever class 1 users alone can saturate the link, i.e., when the number of simultaneous class 1 flows is greater than C/c . When the number of class 1 flows is less than C/c , each attains rate c while the class 2 flows share remaining bandwidth equally.

Let N_k denote the number of class k users, a_k their offered traffic and d_k their useful rate, for $k = 1, 2$. Applying the approximations introduced in §3.2, we deduce that when overall offered traffic $N_1a_1 + N_2a_2$ is less than C , the useful rates d_1 and d_2 are close to the access rate. When overall offered traffic $N_1a_1 + N_2a_2$ is greater than C , on the other hand, carried traffic is nearly equal to C . If the traffic offered by class 1, N_1a_1 , is greater than C , the useful rate of class 2 users is nearly zero and the useful rate of class 1 users satisfies:

$$C \approx \frac{N_1}{1/a_1 + 1/d_1 - 1/c}.$$

If $N_1a_1 < C < N_1a_1 + N_2a_2$, d_1 is close to c and d_2 satisfies:

$$C \approx N_1a_1 + \frac{N_2}{1/a_2 + 1/d_2 - 1/c}.$$

When 10% of users have priority over the others and all users have the same offered traffic 50 Kbit/s for instance, we observe that the useful rate of class 2 users suffers marginally (generally less than 5 Kbit/s) with respect to a situation where bandwidth is shared in a perfectly fair way, whatever the link capacity.

5. Conclusion

In this paper we have investigated the fundamental teletraffic relation linking demand, capacity and performance for a link in a high speed IP access network concentrating the traffic of a certain number of users. Demand is expressed through the speed of the user access line (the access rate), c , and the offered traffic, a . Performance for data traffic is measured through the useful rate, d , equal in practice to the average transmission rate realized for the transfer of a long document.

An exact relation, equivalent to the generalized Engset formula for telephone traffic, is available under an assumption of fair sharing. The useful rate is then insensitive to traffic

characteristics such as the size distribution of transferred documents and depends only on the offered traffic of each user.

System behaviour exhibits two main performance regimes: a transparent regime where the useful rate coincides with the access rate, and a saturated regime where the useful rate is derived deterministically from the fact that the link is always busy and is shared fairly between active users. Simple approximations derived for these regimes are shown to be accurate as long as the number of users is not too small.

The two performance regimes correspond to two possible dimensioning strategies. The network can be made virtually transparent by providing a link bandwidth somewhat greater than the overall offered traffic. If the target useful rate is less than the access rate, the network should be operated in the saturated regime. The required bandwidth can then be derived simply if the offered traffic is known (cf formula (9)). Note, however, that performance in the latter regime is extremely sensitive to the accuracy of the offered traffic estimate. The useful rate decreases very rapidly as overall offered traffic increases beyond the link bandwidth.

The models have been used to evaluate the impact on performance of a heterogeneous user population. Broadly speaking, overall performance tends to improve as the distribution of user traffic becomes less uniform. To assume a homogeneous population thus appears as a conservative dimensioning strategy. The useful rate experienced by all users tends to be roughly the same, independently of their particular traffic intensity. Unfairness has an impact mainly in the saturated regime. However, the reduction in useful rate experienced by the least favoured users tends to be slight under reasonable traffic assumptions.

REFERENCES

1. S. Ben Fredj, T. Bonald, A. Proutière, G. Régnié and J.W. Roberts. Statistical bandwidth sharing: a study of congestion at flow level, in Proceedings of ACM SIGCOMM, 2001.
2. A. Berger and Y. Kogan. Dimensioning bandwidth for elastic traffic in high-speed data networks. *IEEE/ACM Transactions on Networking*, 8-5 (2000) 643–654.
3. T. Bonald and A. Proutière. Insensitivity in processor sharing networks, *Performance Evaluation*, 49 (2002) 193–209.
4. T. Bonald and A. Proutière. Insensitive bandwidth sharing in data networks, *Queueing Systems*, 44 (2003) 69–100.
5. T. Bonald, A. Proutière, G. Régnié, and J.W. Roberts. Insensitivity results in statistical bandwidth sharing, in Proceedings of ITC17, 2001.
6. J.W. Cohen. The generalized Engset formula, *Philips Telecommunication Review* 18 (1957) 158–170.
7. D.P. Heyman, T.V. Lakshman, and A.L. Neidhardt. New method for analyzing feedback protocols with applications to engineering Web traffic over the Internet, in Proceedings of ACM SIGMETRICS, 1997.
8. R. Serfozo. *Introduction to Stochastic Networks*, Springer Verlag, 1999.