

Fun Factor Dimensioning for Elastic Traffic

Joachim Charzinski

Siemens AG, Information and Communication Networks,
Hofmannstr. 51, D-81359 Munich, Germany,
Tel. +49 89 722 46803, Fax +49 89 722 26877,
j.charzinski@ieee.org

Abstract

In this paper we argue that dimensioning for elastic traffic can be achieved in some network scenarios. The “fun factor” φ is introduced as a basic measure of perceived quality of service (QoS) for elastic traffic. As the impact of φ on the perceived QoS is easy to see, target values for φ for residential or business network access can be quantified. Several existing models for dimensioning data networks are compared, the M/G/r-PS model is generalized to arbitrary link rates and a rate convolution approach is presented which allows to dimension links when the distribution of bit rates available from the Internet is known. All approaches are compared using a consistent set of parameters derived from recent high-speed access measurement results, giving an indication of the general relation between the fun factor and the packet loss probability as a more classical QoS measure.

Index terms: Dimensioning; perceived QoS; Fun Factor; TCP; Elastic Traffic; Comparison; M/G/r-PS; Processor Sharing.

1 Introduction

It has been argued that due to its complexity and constant growth, every attempt at dimensioning parts of the Internet must fail. However, there are certain parts of the Internet and some other IP networks where dimensioning can successfully be achieved. One example is the tuning of the capacity of an access trunk line ② to the total capacity of all subscriber access lines ① in the scenario depicted in Fig. 1.

The trunk line must have enough capacity in the upstream direction (from subscribers to core network) to accommodate the sum of all upstream traffic as well as enough downstream capacity to satisfy the download needs of all subscribers.

If all subscribers fully utilized their access lines with full-rate stream traffic, the trunk line capacity would have to be the sum of all access line rates. However, judging from current Internet usage, a multiplexing gain can be achieved




Figure 1. Access architecture

by exploiting the statistical features of the download traffic, which is mostly elastic TCP traffic. The following investigations will help to find and understand the dimensioning target to aim for.

Other scenarios are non-Internet IP networks like company Intranets or the fixed access network part of future mobile radio systems, which will also carry a significant amount of data traffic. Here it is especially important to offer a certain level of quality also to the elastic data traffic as the mobile radio link is a scarce resource which should have a high utilization and therefore the trunk line should be dimensioned such that no radio link capacity is wasted.

In differentiated services networks, the following dimensioning criteria can be used for the higher quality classes in order to estimate the achievable QoS in a class or to find out which amount of capacity to allocate to a class.

Two extreme positions have been found in the literature as to which parameter values to choose when dimensioning for elastic traffic.

Some authors, e.g. [1], see elastic traffic equal to best-effort traffic which does not generate much revenue when operating a network and therefore argue that a minimum rate guarantee plus admission control is the right way to achieve a high network utilization while still maintaining a minimum necessary service level. The main argument in favour of this approach is that customers tend to pay only for the service they need and provided that elastic traffic can be transported with a minimum tolerable rate, no additional revenue can be generated from offering more than this necessary rate.

Other authors, e.g. [2], see the elastic traffic as the major

reason why subscribers are currently investing into high-speed access links and therefore the access system must be able to offer adequate capacity to this traffic. However, when traditional QoS measures and parameter values like a loss probability of 10^{-6} are employed, the resulting link capacities will be overdimensioned [3].

In Sec. 2 we will further discuss the dimensioning issues and introduce the *fun factor* as a parameter for perceived QoS with elastic traffic. In Sec. 3 different models for stream and elastic data traffic are summarized using a consistent notation, which helps comparing the corresponding results using measured traffic parameters in Sec. 4.

2 Dimensioning for elastic traffic

Traditional QoS measures like mean delay, delay quantiles, delay variation or loss probability miss the point when they are employed with elastic traffic. On the one hand, TCP is well able to adapt its transmission rate to bandwidth bottlenecks, so that the packet loss probability is kept at a certain low level, which is also evaluated for rate feedback information. On the other hand, a low bandwidth can have a significant impact on the perceived QoS of a connection as the main use of a TCP connection is to transport a given traffic volume from one host to another. A good QoS measure should therefore consider the achievable bandwidth as the most relevant for elastic traffic.

2.1 Dimensioning targets

As introduced in Sec. 1, the scope of this paper is dimensioning one access trunk line shared by a number of access lines. The main dimensioning target is to make sure (according to a metric to be selected) that this access trunk is not the bottleneck for the upstream or downstream traffic on the individual access lines. Extensions of this approach can be imagined in network planning where the same metrics can be used to describe the quality of service achieved on the different links in a network.

In the following sections, we concentrate on a single traffic class and stationary offered traffic. When dealing with non-stationary traffic, it is usually advisable to restrict the analysis to peak traffic periods as otherwise a long time of low traffic demand would smooth out the critical congestion problems occurring during a peak traffic period. For multiple traffic classes to be considered, e.g. methods from [4, 1] can be employed.

The models introduced later take a traffic characteristic and provide QoS or dimensioning results with respect to this traffic characterization. Care must be taken in order to focus on the right level of interest and on a situation applicable to the considered network scenario when specifying the traffic model. An xDSL (digital subscriber line) or cable mo-

dem network usually provides line terminations inside the network for each subscriber, regardless of their activity, so that also “always-on” connections can be provided. In this case, the maximum number of active subscribers is equal to the total number of subscribers. In contrast, dial-up access usually limits the number of subscribers sharing the same access trunk line to the number of access ports connected to this line, so that the amount of network resources (i.e. bandwidth) to be provided per access port will differ significantly in both cases. In the rest of this paper, we will use the traffic observed in *active* client sessions. In a dial-up network model, this is also the number of access ports whereas in an xDSL or cable modem network model the total number of subscribers will be higher than the number of active access lines during the peak hour.

2.2 Buffers

A main issue of distinction between models for stream traffic and models for elastic traffic is the way buffers are considered. For non-elastic data traffic, burst scale buffering is often used to reduce packet loss. On the contrary, elastic traffic is able to adapt to the available bandwidths and therefore large buffers are neither needed nor as effective as with non-elastic traffic. Long-range dependence (LRD) in data traffic [5], which is a consequence of heavy-tailed burst sizes and durations [6], further reduces the effectiveness of buffers because the packet loss probability cannot be significantly reduced even by huge buffers in the presence of LRD traffic [7, 8]. In those cases, large buffers drastically increase the packet delay without really reducing packet loss.

Consequently, it is a sensible approach to assume that all network nodes provide enough buffer space to cope with packet scale queueing¹ but buffers are small enough not to increase packet delays too much. In such an environment, bufferless fluid loss models or processor sharing models with infinite buffer are applicable – provided that in the real system there is a small buffer present for packet scale multiplexing. Using bufferless models has two major advantages over buffer models: Heavy tailed ON phase distributions do not influence the results and the actual buffer sizes of network elements need not be known.

In the following, we will be estimating message delays from bufferless models, in this way neglecting the packet queueing delays occurring in real buffers. This estimation is justified if the total message or burst transfer time is much larger than the delay experienced by the individual packets.

¹Packet scale queueing describes the effect that due to asynchronous multiplexing of packet traffic, multiple packets from different input links must be serialized to share the same output link.

2.3 Access control

There are several approaches towards quality problems in the Internet today. One extreme is to “throw bandwidth at them”, the other extreme is to solve quality problems “the old way”, i.e. by introducing intelligence at the network layer, using access control (AC) to ensure a connection can only be accepted when there is enough bandwidth available along the whole path it will take through the network.

[4, 1] argue in favour of AC in order to allow a high network utilization while still guaranteeing a minimum rate for each TCP connection. But if no access control is employed, TCP connections still receive a QoS that can be statistically quantified and dimensioned to have a certain mean or quantile value as we will show below. The differences between this dimensioning and employing access control are:

- Access control requires additional network layer functions.
- Access control allows to reach higher network utilization.
- Under unplanned overload, access control reduces the QoS by increasing the blocking probability for new connections whereas in networks without access control, the perceived QoS of all connections – including the already existing ones – is reduced. In the extreme, this can lead to a starvation of all connections.

Deriving dimensioning results without AC can therefore be seen as a step in quantifying how much bandwidth to “throw at the problem”, i.e. how much a link needs to be overdimensioned in comparison to AC in order to deliver a comparable total QoS comprising the blocking probability as well as the achievable rate within an existing connection. As mentioned above, this cost of overdimensioning has to be put in contrast to the savings made by not investing into network layer functions to support AC functions in every router.

2.4 Definition of the fun factor

In an effort to capture the quality of service for elastic traffic, the following measure is defined: A parameter φ , which should be as simple as possible, gives the QoS achieved in a TCP connection relative to a theoretically possible value. In addition, φ should have a maximum of one under ideal circumstances and a minimum of zero if there is no service at all.

Making this quantity a relative measure of QoS allows us to quantify the relative difference of systems. Consider again the access scenario depicted in Fig. 1. Here the “ideal” situation is as follows: Every subscriber can utilize the full bandwidth of their access line whenever they want. Relative

to this reference situation, the service degradation can be observed as an increase in transfer time or as a decrease in the available rate. Using Φ to denote a random variable and φ for its expectation $\varphi = E[\Phi]$, this leads to the following two equivalent formulations for the access fun factor Φ_A with the bit rate B and the capacity C_{AL} of an access line where t_{ideal} is the time a data transfer would have taken if C_{AL} had been utilized fully.

$$\Phi_A = \frac{t_{ideal}}{t_{observed}} \quad (1)$$

$$\Phi_A = \frac{E[B|B > 0]}{C_{AL}} \quad (2)$$

The quantity Φ is called “fun factor” to stress the fact that the corresponding elastic traffic service can usually also be used at much lower rates – with a corresponding loss in the enjoyment of usage. As such, it is also applicable for different service providers to use the fun factor as a means of differentiation.

In the following, only this *access* fun factor will be considered and consequently the index A will be dropped. Similar investigations can and should however be conducted for other fun factors as well.

For business scenarios, of course the same term can be changed into “access line usage efficiency”. Both time based and money based markets are in principle described adequately by this measure as it answers both questions “How long did I have to wait using my flat rate service?” as well as “How much could I have gained from a high-price service if the network had been perfect?”. The reciprocal questions addressing the cost, of course, cannot be answered in general as the answers depend on the actual tariffing structures.

The same approach as for the access fun factor can be taken towards other fun factors with respect to other limiting resources like e.g. the (variable) rate available in the real Internet, the rate available from a server or the rate that a client device gets with respect to the maximum rate it can consume.

Note that the idea of considering the rate decrease or delay increase as a QoS parameter for elastic traffic is not new. A mean delay factor $f_R \approx 1/\Phi$ to describe the service degradation due to less-than-ideal transmission rates for elastic traffic has been defined before [9, 1]. However, the attempt to quantify $\varphi = E[\Phi]$ in order to dimension for a certain target value and to find corresponding target values for loss models is novel.

Continuing the above arguments about time and money based markets, we can estimate target values for φ : Residential subscribers have bought new Internet access equipment increasing the bandwidth by factors of 2–3. Assuming that the old equipment was operated at the limits of its throughput capacity ($\varphi \approx 1$), the same rate would correspond to a

fun factor of $\varphi = 0.3-0.5$ with the new equipment. A noticeable improvement can therefore be felt if $\varphi > 0.7-0.8$, demanding that the access as well as the wide area service providers can offer the corresponding throughput. For business customers, on the other hand, the target values should be much higher, e.g. 95 or even 99%, depending on the degree of service requested. For high-quality business access, one could also imagine determining the distribution of Φ and demanding that the ratio of connections with $\Phi < 90\%$ be less than a certain percentage during busy hours.

3 Traffic models for dimensioning

In the following, a few models for dimensioning data networks are summarized. A more extensive summary and comparison of classical models is attempted in [10]. There are some fundamental differences between the traffic models such as a finite or infinite number of traffic sources, the assumption of greedy vs. rate limited sources and between the system models where loss models with or without buffer are opposed to rate sharing models.

3.1 Engset Fluid Flow

A simple approach for a burst scale loss model with homogeneous ON/OFF sources has been given by [11, 9, 1, 2]. It is based on fluid flow modelling and computes the fraction of fluid traffic that is lost when more traffic is offered than a bufferless system can handle. Being a fluid model, it does not consider packet scale queueing and loss and is therefore realistic for a real system with small buffers. The complete absence of a buffer in the fluid model allows concentrating on rate multiplexing effects without having to include any traffic correlation into the model. In essence, the model is robust with respect to correlation parameters like mean or distribution of burst (ON phase) duration or the long-term correlations described by the Hurst parameter.

A single source is described by the parameters r_{ON} for the data rate in the ON phase and β for the probability of the source to be in the ON state. If the traffic from n such sources is multiplexed on a link with a total capacity c_L , the fraction of fluid traffic lost, p_{loss} , is given by (3):

$$p_{loss} = \sum_{i=[c_L/r_{ON}] }^n \binom{n}{i} \beta^i (1-\beta)^{n-i} \frac{i \cdot r_{ON} - c_L}{\rho \cdot c_L} \quad (3)$$

Eq. (3) computes p_{loss} from a sum over all overload states, using the binomial distribution p_i for the probability of i sources being in the ON state, assuming independent state transitions of all sources.

$$p_i = \binom{n}{i} \beta^i (1-\beta)^{n-i} \quad (4)$$

The relative total offered load in this case is

$$\rho = \frac{n\beta r_{ON}}{c_L} \quad (5)$$

3.2 Fractional Brownian Motion

Norros [12] introduced a fluid model of a buffer fed by Fractional Brownian Motion traffic in an attempt to incorporate the long-range dependence properties observed in many data traffic measurements. The basic traffic model is a fluid source that generates traffic according to the arrival process

$$A_t = mt + \sqrt{ma} Z_t \quad (6)$$

with mean rate m , variance coefficient a and Hurst parameter H (property of the FBM Z_t).

On the time scale of $t_0 > 0$, the fractional Brownian Motion Z_{t_0} has an expectation of zero and a variance equal to one. According to the scaling law for self-similar traffic, the variance of average values taken over a longer time decreases dependent on the Hurst parameter H :

$$\text{VAR} \left[\frac{1}{t/t_0} Z_t \right] = \left(\frac{t}{t_0} \right)^{-2(1-H)} \quad (7)$$

Accordingly, the variance of A_t depends on the counting time t :

$$\text{VAR} \left[\frac{t_0}{t} A_t \right] = ma \cdot \left(\frac{t}{t_0} \right)^{-2(1-H)} \quad (8)$$

This equation also explains the role of the variance coefficient a and shows how it can be determined from a variance-time plot.

The link capacity needed to accommodate one source with parameters m , a and H is given as

$$C = m + \left(H^H (1-H)^{1-H} \sqrt{-2 \ln \varepsilon} \right)^{\frac{1}{H}} a^{\frac{1}{2H}} \varepsilon^{\frac{H-1}{H}} m^{\frac{1}{2H}} \quad (9)$$

if a loss rate ε is tolerable and a burst scale buffer of size x is present. For multiplexing n homogeneous sources, Norros [12] gives the scaling laws $m^{(n)} = n \cdot m^{(1)}$, $a^{(n)} = n \cdot a^{(1)}$ and $H^{(n)} = H^{(1)}$.

3.3 Processor Sharing

Processor sharing (PS) models provide a simple and effective way of dealing with the elastic properties of traffic carried by transport protocols which adapt the transmission rate to the available capacity in the network.

Lindberger [1] describes a PS queueing model M/G/r-PS for a link which can accommodate r times the peak rate r_{ON} of an individual source, i.e.

$$r = \frac{c_L}{r_{ON}} \quad (10)$$

for integer r and negative exponential interarrival time distribution (i.e. an infinite number of sources) offering a total offered load ρ to the link. Lindberger derives a mean *delay factor* f_R which is roughly $1/\varphi$ (except for the averaging process, as $f_R = E[1/\Phi] \neq 1/E[\varphi]$). In addition, generally there is a difference between time averages and per-connection averages.

For sources which have a maximum rate greater than the link rate, the corresponding M/G/1-PS model gives a delay factor of $f_R = 1/(1 - \rho)$ leading to $\varphi \approx 1 - \rho$.

For the M/G/ r -PS model, the delay factor is based on Erlang's formula for waiting probability in the corresponding M/G/ r -FIFO waiting system:

$$f_R = 1 + \frac{1}{r(1 - \rho)} \cdot \frac{\frac{(r\rho)^r}{r!}}{(1 - \rho) \sum_{i=0}^{r-1} \frac{(r\rho)^i}{i!} + \frac{(r\rho)^r}{r!}} \quad (11)$$

In order to use this model for dimensioning links, it is favourable to extend it to arbitrary link rates, i.e. non-integer values of r . Define r_g as the number of sources that can be served by the link without extra delay:

$$r_g = \left\lfloor \frac{c_L}{r_{ON}} \right\rfloor \quad (12)$$

Using X as a random variable for the number of sources in the ON state, the expectation of the random delay factor F_R is

$$f_R = E[F_R | X \leq r_g] \frac{P\{0 < X \leq r_g\}}{P\{X > 0\}} + \sum_{i=r_g+1}^{\infty} E[F_R | X = i] \frac{P\{X = i\}}{P\{X > 0\}} \quad (13)$$

While F_R is equal to one in the states $X \leq r_g$, it is i/r when more sources are active. The state probabilities are given by

$$p_i = P\{X = i\} = \begin{cases} \frac{(r\rho)^i}{i!} p_0 & , i = 0, 1, 2, \dots, r_g \\ \rho^{i-r_g} p_{r_g} & , i > r_g \end{cases} \quad (14)$$

and

$$p_0 = \left[\sum_{i=0}^{r_g-1} \frac{(r\rho)^i}{i!} + \frac{1}{1 - \rho} \frac{(r\rho)^{r_g}}{r_g!} \right]^{-1} \quad (15)$$

The resulting mean delay factor is

$$f_R = 1 + \frac{\rho}{r(1 - p_0)} \left(\frac{1}{1 - \rho} + r_g - r \right) \cdot \frac{\frac{(r\rho)^{r_g}}{r_g!}}{(1 - \rho) \sum_{i=0}^{r_g-1} \frac{(r\rho)^i}{i!} + \frac{(r\rho)^{r_g}}{r_g!}} \quad (16)$$

Approximating $1 - p_0$ by ρ and the waiting probability in (16) by the loss probability as computed from other models in this section, we obtain the following very coarse approximation:

$$f_R \approx 1 + \frac{p_{loss}}{r(1 - \rho)} \quad (17)$$

Note that (17) is too optimistic as the waiting probability can be significantly higher than the loss probability in a corresponding bufferless system.

3.4 Rate convolution

For a bufferless fluid loss computation, the loss probability is equal to the probability of the sum of all active sources' rates being greater than the link rate. If the rate distribution for each single access line is given by the rate distribution probability density $g_R(x) = \frac{d}{dx} P\{R_{ON} \leq x\}$ with Laplace transform (LT) $L_R(s)$, this probability is given by the n -fold convolution of $g_R(x)$ with itself, g_{R_n} , which has the LT

$$L_{R_n}(s) = (L_R(s))^n \quad (18)$$

Using the corresponding complementary distribution function $C_{R_n}(x)$, the system overload probability is

$$p_{ovl} = C_{R_n}(c_L) \quad (19)$$

the ratio of lost traffic is given by

$$p_{loss} = \frac{1}{nE[R_{ON}]} \int_{x=c_L}^{\infty} (x - c_L) g_{R_n}(x) dx = \frac{1}{nE[R_{ON}]} \int_{x=c_L}^{\infty} C_{R_n}(x) dx \quad (20)$$

This rate envelope multiplexing (REM) approach can be regarded as similar to a continuous state Engset model where the sources are specified with a distribution of instantaneous rates instead of a probability of activity and a fixed ON state rate. With this feature, (20) can be used for sources with independently varying rates. In the case of elastic sources sharing common links, the rates are obviously not varying independently but vary such that the aggregated rate peaks are smaller than what they could be under the assumption of independence, so that (20) gives an upper bound for p_{loss} : If connections do not share any common link in the wide area networks, they will be independent. If on the other hand they do share a common link, they will slow down each other leading to a reduction of the loss probability p_{loss} occurring on the access link.

The rate convolution approach differs from the other models in that it can take the distribution of achievable rates

from the Internet into account. This feature makes the formula both more realistic and at the same time more dependent on up-to-date measurement results than the other approaches. In addition, there is in practice a problem of finding the right averaging time resolution to determine the rate distributions.

3.5 Simulation models used

Two burst level (fluid flow) simulation programs have been used to allow more detailed statistical analysis, validation and extension of the analytical results given above. In both simulation models, the traffic is modeled by a fluid flow with an intensity equal to the sum of the rates of the currently active sources. Program A allows loss simulations with various buffer sizes.

In contrast, Program B simulates a processor sharing system where at each source ON and OFF instant the available capacity is re-distributed fairly among the currently active sources taking into account the source's current target r_{ON} , re-computing the residual traffic to be transmitted by each source during its ON phase. In simulation model B no explicit buffer is simulated as all sources are assumed to adapt to the new available rates without delay. The ON rate of each source is drawn from a rate distribution $C_R(x) = P\{R_{ON} > x\}$ at the beginning of a burst and is assumed to be a constant target value during the duration of the whole burst. This simulation model allows to investigate realistic distributions of rates available from the wide-area Internet but does not include TCP behaviour details like non-zero reaction times or phase synchronization effects.

3.6 Summary of models

Tab. 1 summarizes the properties of the different models with respect to the number of sources, the burst level buffer sizes, source rate limit, source type and rate sharing options.

Table 1. Model summary

model	sources	buffer size	r_{ON}	source type	rate sharing
Engset Fluid	n	0	r_{ON}	on/off	no
FBM		x	∞	FBM	no
M/G/r-PS	∞	∞	r_{ON}	on/off	yes
RC	n	0	$C_R(x)$	open	no
Simulation A	n	x	r_{ON}	on/off	no
Simulation B	n	0	$C_R(x)$	on/off	yes

4 Comparison of results

The results obtained from the different models are compared in sections 4.2 and 4.3 using common and comparable sets of parameters derived from measured data described briefly in Sec. 4.1.

4.1 Measured traffic

Comparable sets of parameters are needed to objectively compare the results obtained with different models. As the models require different types of parameters, measurement results have been used to derive the individual parameter values. In accordance with the idea of dimensioning an access trunk line, traffic traces have been collected close to clients during an ADSL (asymmetric digital subscriber line) field trial in Münster, Germany, where 100 subscribers have been observed from May to December 1998 [13, 14]. Measurements were taken cyclically in groups of around 7 users for a week each. Each subscriber accessed the Internet through a line with a configured rate limit of 2.5Mbit/s downstream and 384kbit/s upstream. The access trunk line (⊙ in Fig. 1) was a 100Mbit/s line offering bandwidth in abundance. Servers outside University of Münster's campus network could be reached via its 34Mbit/s connection to the German Research Network. The dependence of parameters on access link speeds is discussed in [15].

Table 2. Parameters

Parameter	Symbol	Value
Mean Rate	m	10.5kbit/s
Access Line Rate	r_{ON}	2.5Mbit/s
Activity Factor	β	0.0042
Variance Coefficient	a	$4.4 \cdot 10^5$
Hurst Parameter	H	0.915

The parameters observed for HTTP/TCP/IP/Ethernet traffic in *active* client access sessions are summarized in Tab. 2. A client access session has been defined as the interval in time in which a client computer transmits or receives packets with pauses of no more than 10 minutes. Mean rate m , variance coefficient a and Hurst parameter H have been determined from a variance-time plot covering the time scales from 10ms up to one day. The region for consistent a and H estimation was between around 50ms and 3min.

The measured one second average bit rates during active HTTP client sessions as defined above have been used

for the rate convolution model from Sec. 3.4. The n -fold convolution and the evaluation of (20) have been performed numerically on the measured distribution using linear interpolation between the distribution values.

Simulation B requires a rate and duration distribution for ON phases in order to allow the (possibly more realistic) investigation of fun factors under the assumption of a distribution of rates available from the Internet similar to when the measurement data were recorded. During a client session, it is realistic to assume that the duration of a phase of constant rate is distributed approximately like the duration of a TCP connection. In the measured data, this duration had a mean of 57s and a coefficient of variation of 16. For the simulation, a Pareto distribution truncated [16, 17] at 70000s has been used, which corresponds very well to the measured data. The rate distribution in the ON phases has been extracted from the active client rate distribution as the conditional distribution for rates that are greater than zero. A best fit was obtained using a truncated Pareto distribution with a minimum value of 19kbit/s, $\alpha = 1.5$ and a censoring bound of 2.5Mbit/s, i.e. a concentration of the remaining tail probability mass on the discrete value of 2.5Mbit/s, which is in good correspondence with the measured distribution. The probability for the rate to be zero for one second during an active session was 0.8. This has been used to determine the mean of a negative exponential OFF phase duration (228sec).

4.2 QoS indicator results

In Fig. 2, analysis and simulation results are depicted for 32 and 800 sources with parameters as described above. In order to plot it together with loss ratio results, the fun factor has been replaced by $1 - \varphi$ in the plots. The results in the graph are:

- p_{loss} from simulation A, given with 95% confidence intervals.
- p_{loss} according to the Engset analysis (3).
- $1 - \varphi$ from simulation B with ON/OFF sources and $r_{ON}=2.5\text{Mbit/s}$.
- $1 - \varphi$ from the M/G/r-PS analysis (16).
- $1 - \varphi$ from approximation (17) using p_{loss} from (3).

Additional simulations agree with the analysis result that the ON phase distribution has no influence on the loss ratio or mean fun factor results for the bufferless or processor sharing models even in the limited sources case.

In the parameter range chosen here, the results for p_{loss} and $1 - \varphi$ are almost indistinguishable and vary by about 10–20% in terms of the rate necessary to give the same level

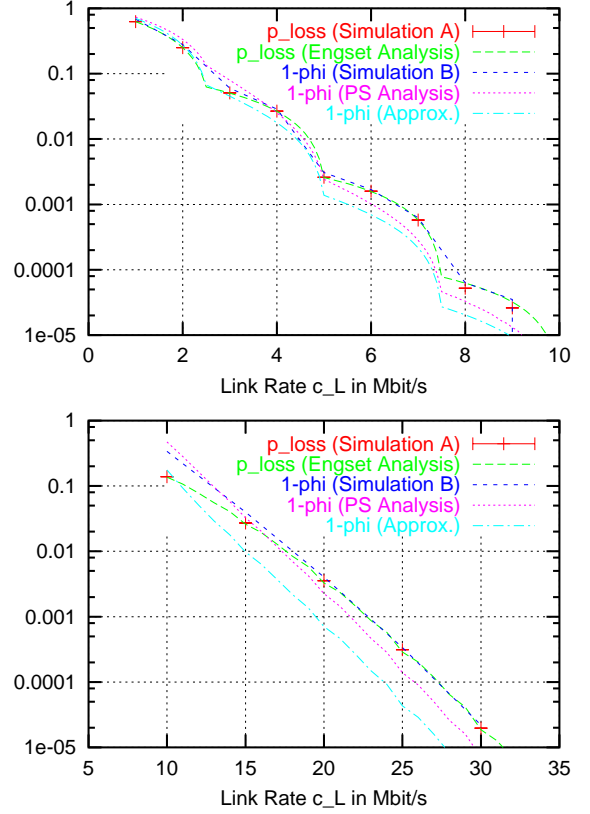


Figure 2. Transformed Fun Factor $1 - \varphi$ and loss probability for (a) 32 sources and (b) 800 sources. For abbreviations and parameters see text in Sec. 4.3.

of QoS. This property is found as long as the difference between time and per-connection averages is negligible. In simple ON/OFF cases where all sources have the same peak rate, the equivalence of p_{loss} and $1 - \varphi$ can also be shown analytically.

The picture changes when the real rates available from the Internet are included in the evaluation. As mentioned before, in order to work with such values, these rates must be measured frequently in order to keep track with the evolution of the Internet. For the rate distribution documented in Sec. 4.1, Fig. 3 gives the dependence of p_{loss} according to (20) and $1 - \varphi$ as obtained from simulation B with ON rates randomly chosen from the rate distribution.

With this real rate distribution, the achievable fun factors are much higher than would be expected from loss rate dimensioning, even if a target $1 - \varphi$ had been used for p_{loss} . This difference is due to the fact that sources with an achievable rate r_{ON} lower than the current fair share produce a contribution $\Phi = 1$ and only those with higher achievable rates are slowed down by the access trunk. However, it cannot be stressed too often that in order to profit from this additional gain, an operator must keep track of the evolution of

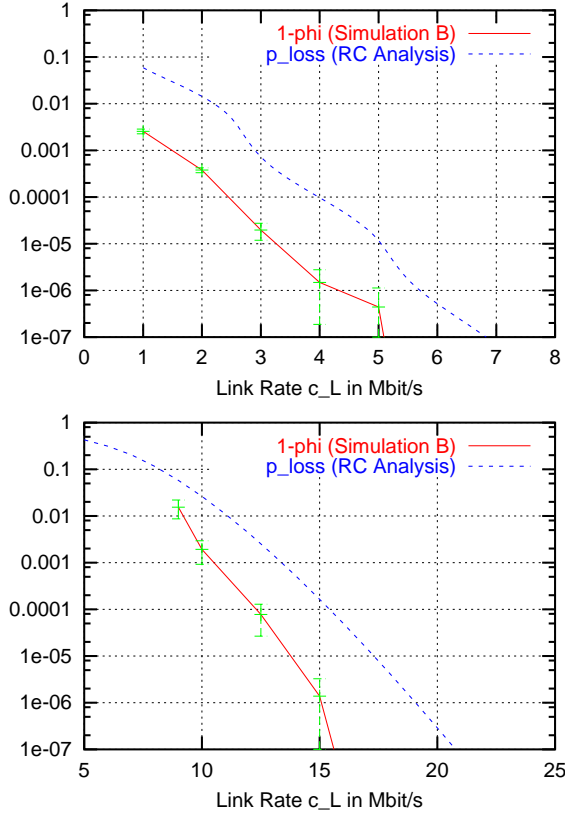


Figure 3. Transformed Fun Factor $1 - \varphi$ and loss probability for (a) 32 sources and (b) 800 sources with a rate distribution as given in Sec. 4.1.

bit rates available from the Internet and re-dimension their access trunk lines accordingly.

4.3 Dimensioning results

Fig. 4 depicts a summary of the results from the analytical models for a wider range of n . The comparison is done with respect to the predicted required link capacity $c_L(n)$ to serve n sources with a given QoS measure. The following models and parameters have been used:

- The required accumulate mean rate of all sources, $n \cdot m$.
- $p_{loss} = 10^{-6}$ in the Engset model (3).
- $p_{loss} = 10^{-6}$ and $x = 1$ in the Fractional Brownian Motion model (9).
- $\varphi = E[\Phi] = 0.99$ in the M/G/r-PS analysis (16).
- $p_{loss} = 10^{-6}$ for rate convolution (20).

The dimensioning results are in fairly good agreement with each other. The FBM result over-emphasizes the long

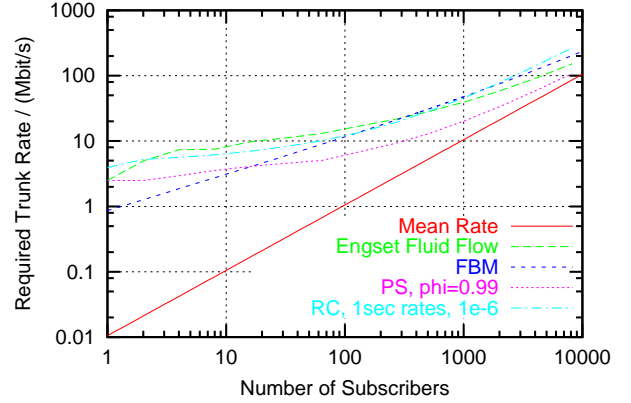


Figure 4. Comparison of dimensioning results. For abbreviations and parameters see text in Sec. 4.3.

range dependence and under-estimates the bandwidth required to satisfy a small number of access links. The results from (16) are nearly indistinguishable from the Engset results if $\varphi = 1 - 10^{-6}$ is used.

5 Conclusions

In this paper the “fun factor” Φ was introduced as a basic measure of perceived quality of service for elastic traffic. As the impact of $\varphi = E[\Phi]$ on the perceived QoS is easy to see, it is a simple task to define target values for φ for residential or business network access. Different existing models for dimensioning data networks were compared. The M/G/r-PS model from [1] was generalized to arbitrary link rates and a new approach for dimensioning links with a known distribution of the available rate from the Internet (rate convolution) was added. A comparison with a consistent set of parameters derived from recent measurement results helped to see the general relation between the fun factor and the packet loss probability as a more classical QoS measure.

If the distribution of rates available from the Internet is known, bandwidth can be saved due to better multiplexing of sources with lower ON rates. The time scale used for the rate distribution should be smaller than the duration of the connections of interest.

In order to validate especially the fluid simulation results, simulations including full TCP models should be performed. A continuous state model might help deriving the distribution or mean of Φ in the rate convolution context analytically.

It should be an easy task to define a measurement method for Φ based on fluid rate sharing assumptions. Such a method could help Internet service providers (ISPs) monitoring their access trunks towards customers as well as their interconnection points towards the rest of the Internet. A

future application could also be to include Φ limits in service level agreements or offer ISPs to advertise their good network performance and overdimensioning towards their subscribers as a differentiation feature.

Acknowledgements

The author would like to thank Phuoc Tran-Gia for bringing up the term “fun factor” and discussing the associated issues. Stefan Bodamer and the anonymous reviewers have helped to improve the presentation of this paper.

References

- [1] K. Lindberger, “Balancing Quality of Service, Pricing and Utilization in Multiservice Networks with Stream and Elastic Traffic” Proc. ITC 16, Edinburgh, UK, June 1999, pp. 1127–1136.
- [2] F. Hartleb, G. Haßlinger, “Ein Modell zur Analyse der QoS in DiffServ Netzen” (“A model to analyse the QoS in DiffServ networks”), presentation at ITG Fachgruppe 5.2.1, Nuremberg, Germany, Jan. 2000 (in German).
- [3] Å. Arvidsson, P. Karlsson, “On Traffic Models for TCP/IP”, Proc. ITC 16, Edinburgh, UK, June 1999, pp. 457–466
- [4] L. Massoulié, J.W. Roberts, “Arguments in favour of admission control for TCP flows” Proc. ITC 16, Edinburgh, UK, June 1999, pp. 33–44
- [5] W.E. Leland, M.S. Taqqu, W. Willinger, D.V. Wilson, “On the Self-Similar Nature of Ethernet Traffic”, Proc. ACM Sigcomm’93, San Francisco, CA, USA, 1993, pp. 183–193
- [6] M.E. Crovella, A. Bestavros, “Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes”, IEEE/ACM Trans. Networking, Vol. 5, No. 6, Dec. 1997, pp. 835–846
- [7] M. Grossglauser, J.-C. Bolot, “On the Relevance of Long-Range Dependence in Network Traffic”, IEEE/ACM Trans. Networking, Vol. 7, No. 5, Oct. 1999, pp. 629–640
- [8] G. Haßlinger, “Die begrenzte Wirkung von Puffern in diensteintegrierenden Kommunikationsnetzen” (“The limited effect of buffers in integrated services communication networks”), Proc. MMB’99, Trier, Germany, Sep. 1999, pp. 221–235 (in German).
- [9] D.P. Heyman, T.V. Lakshman, A.L. Neidhardt, “A New Method for Analysing Feedback-Based Protocols with Applications to Engineering Web Traffic over the Internet” Proc. ACM SIGMETRICS’97, Seattle, WA, USA; ACM Perf. Eval. Review Vol. 25, No. 1, pp. 24–38.
- [10] S. Bodamer, J. Charzinski, “Evaluation of Effective Bandwidth Schemes for Self-Similar Traffic”, ITC Specialist Seminar on Internet Traffic Modelling, Monterey, CA, USA, Sep. 2000.
- [11] J.W. Roberts (Ed.), “Performance evaluation and design of multiservice networks”, Final report of the COST 224 project, Office for Official Publications of the European Communities, Luxembourg, 1991.
- [12] I. Norros, “On the Use of Fractional Brownian Motion in the Theory of Connectionless Networks”, *IEEE JSAC* Vol. 13, No. 6, 1995, pp. 953–962.
- [13] J. Charzinski, “Internet Client Traffic Measurement and Characterisation Results”, Proc. ISSLS 2000, Stockholm, Sweden, June 2000.
- [14] J. Charzinski, “HTTP/TCP Connection and Flow Characteristics”, *Performance Evaluation*, Sep. 2000.
- [15] N. Vicari, S. Köhler, J. Charzinski, “The Dependence of Internet User Characteristics on Access Speed” University of Würzburg Technical Report No. 246, Jan. 2000, available at <http://www-info3.uni-wuerzburg.de/TR/tr246.pdf>
- [16] N. Vicari, “Measurement and Modeling of WWW-sessions” Technical Document COST 257(97)52, Sep. 1997, available at <http://www-info3.informatik.uni-wuerzburg.de/cost/cost257/september97/257td9752.pdf>
- [17] K. Lindberger, “Delays in ABR-like traffic and relations to processor sharing” Technical Document COST 257(98)10, Jan. 1998, available at <http://www-info3.informatik.uni-wuerzburg.de/cost/cost257/jan98/257td9810.pdf>