Appropriate Teletraffic Framework in IP over WDM

Helio Waldman[1], Fernando L. Trazzi Jr.[2], Moisés R. N. Ribeiro[2]

[1]DECOM/FEEC/UNICAMP Caixa Postal 6101 - 13083-970 - Campinas, SP - Brasil

[2]LABTEL/DEL/CT/UFES Caixa Postal 01-9011- 29060-970-Vitória, ES – Brasil

waldman@decom.fee.unicamp.br, {fernando,moises}@labtel.ele.ufes.br

Abstract. *Current WDM networks see the traffic burstiness only after it has been shaped by the aggregation process in upper layers (ATM or IP) and by the mapping process into a SDH transport layer. However, present trends point towards direct mapping of the IP layer into WDM. The mean residual waiting time of service by single servers of self-similar Pareto traffic is shown to dominate over the mean of any finite queue service time. It is derived in this paper that a sufficiently high number of wavelengths (multiple servers) used in pool will bring this mean to finite and even small values. However, higher moments of the residual waiting time distribution will always run away to infinity unless limits are placed on the wavelength holding time.*

1. Introduction

The last decade has seen the emergence of the first bandwidth mining techniques, which have converged to wavelength division multiplexing (WDM). They have widened the transmission pipes in a much larger scale than the routing nodes. More and more pressure is then concentrated on the electronic processing capacity of such nodes, thus shifting the electronic bottleneck from transmission to routing functionalities. Easing this new electronic bottleneck has then become crucial to provide, or at least maintain, the expected quality-of-service (QoS) of applications carried by an ever increasing volume of aggregated traffic.

Mitigating the electronic routing bottleneck is the aim of the emerging art of optical networking. Its purpose is to devise network architectures that combine, in the most effective way provided by current and new technologies, the functions that may be realized by electronics and by photonics. This is to be done for a traffic environment characterized: a) by an explosive growth in volume; and b) by profound changes in the nature of such traffic, due to bursty behavior of traffic sources and aggregation of traffic generated by applications with different QoS requirements.

Characterization of burstiness has always been an elusive proposition. Classical models for data packet transmission have used markovian, memoryless traffic models, which have been considered adequate for many traditional applications. More recently, however, it has become clear that such models are not suitable to describe the behavior of traffic generated by emerging Web-based applications. Instead, it has been found that emerging data communication environments generate traffic with long-range dependence.

The purpose of this paper is to put forward teletraffic approach for the emerging networking environments. The next Section discusses some of the main features (e.g. self-similarity, heavy-tailedness) vis-à-vis their impact on network resources. Section 3 discusses the effect of such features on the WDM environment, taking the Pareto distribution as a paradigm for the heavy-tailedness of the Web file (and hence burst duration in a queueless environment such as WDM) size distribution. Finally, Section 4 discusses the networking implications of the results along with some of the evolving issues relating to IP over WDM.

2. Self-similarity and Service Time Distribution in WDM Networks

Let us consider N independent, identically distributed (i.i.d.) random variables with mean μ and variance σ^2. Their sum will have mean $N\mu$ and variance $N\sigma^2$. Their sample mean will have mean $N\mu/N=\mu$ and variance $N\sigma^2/N^2=\sigma^2/N$. The standard deviation of the mean will therefore decrease with N as $N^{-1/2}$. This simple result is imbedded in a well-known theorem from the theory of probability, called the law of large numbers, which also states that the distribution of the sample mean is asymptotically Gaussian. Its validity, however, is constrained to the case where the i.i.d. variables have a variance. Not all random variables have variance, though. An interesting class of distributions that do not have variance is the class of heavy-tailed distributions. Consider a random variable K for which the probability of K exceeding x falls off with x as $cx^{-\alpha}$, where $0<\alpha<2$ is a shape parameter. If $0<\alpha<1$, T will have neither variance nor mean, but this case is not relevant to the framework of our discussion. If $1<\alpha<2$, however, K will have a finite mean but no second moment, hence no variance. In both cases, K is considered a heavy-tailed distribution.

When N i.i.d. heavy-tailed variables with shape parameter $\alpha>1$ are averaged, the deviation from the statistical mean does not fall off like $N^{-1/2}$ as prognosticated by the law of large numbers for lighttailed variables. Instead, it falls off like N^{H-1}, where $H=(3-\alpha)/2$ is called the Hurst parameter of the self-similar cumulative process. Such averages (or sums) arise, of course, when we take the accumulated traffic generated by on-off traffic sources. It has been shown that an aggregation of on-off traffic sources will exhibit self-similarity if the duration distribution of either the on or off periods is heavy-tailed [2]. The self-similarity property consists of the preservation of the distribution of a random process at different time scales, except for a normalization factor that is the H-power of the time scaling factor. For $H = ½$, self-similarity reduces to the classic behavior prognosticated by the law of large numbers.

Heavy-tailedness in the distribution of the time between new connection arrivals may also result in self-similarity of the traffic. However, if connection requests are generated by independent entities, the arrival process is expected to be Poissonian, so interarrival times are exponentially distributed. Since this seems to be the case with traditional Internet applications (e.g. ftp and telnet), traffic self-similarity has been more often associated with heavy-tailed service time distributions. However, the rapid emergence and predominance of Web-based applications has changed this picture. Typical Web usage is motivated by search in a "surfing" mode, in which one download operation (i.e. a "connection") may lead to another in a concatenated fashion. This enhances both very short and very long (with respect to the mean) interarrival times, generating heavy-tailed inter-arrival time distributions [5].

The distribution of file sizes transferred through a network through Web-based applications has been surveyed [3]. The file sizes were found to follow a heavy-tailed distribution with α approximately 1.2 through at least three orders of magnitude. Future Internet-oriented transfer protocols in the WDM network will likely support such file downloads with some kind of flow switching or burst switching. In such protocols, the transfer is free from store-and-forward and buffering operations as much as possible, so that the file size is essentially proportional to the duration of an on period at some WDM server. The discussion of the nexus between traffic characteristics and networking performance has been mostly focused on the correct dimensioning of buffering resources in the context of ATM networks, and is still evolving. Two distinct methodologies are available for this purpose: simulation; and emerging mathematical theories, such as large deviation theory [7]. Classic networking frameworks are basically classified into two categories: queuing networks (i.e. data networking), and blocking networks (e.g. the telephone network). Currently emerging networking paradigms, however, do not fit well into any of these models. Newly defined frameworks should then be used in traffic performance studies. With the current shift in networking philosophy, new approaches are clearly needed to understand the implications of traffic shapes – and shaping - on performance.

Fig. 1(a) brings an illustrative case for a WDM network that is being used to serve a connection (or bursts) from node 1 to node 3, which follows the route highlighted. In the meantime, a connection request from node 5 to node 4 arrives. This connection will demand one out of n-1 servers (wavelengths) available between node 5 and 4 to use the route shown with the dashed line. Suppose n=1, and this connection request may wait in a buffer until the ongoing connection (or burst) ends. Fig. 1(b) shows that the ongoing connection has been active for t_0 seconds when the connection request arrives. This paper is particularly concerned with the ongoing connection residual lifetime seen by the connection request. This will allow, for instance, QoS polices to decide whether a connection request should be simply blocked or await resources to be freed based on inference of residual lifetimes (Z) of ongoing connections according to its elapsed time (t_0). However, we show that there might be surprisingly large discrepancies between process residual lifetime itself (S), presented in Appendix 1, and the residual lifetime seen by the connection request (client view) addressed in Section 3.2 and 3.3. Moreover, we investigate the tradeoff between the number of servers (n) required for a QoS target and the minimum (t_{min}) and maximum (t_{max}) connection holding times.

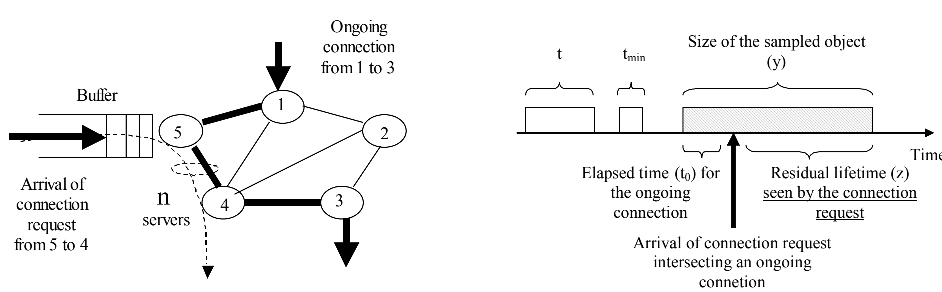

(a) (b)

Figure 1. Problem under analysis: (a) WDM network. (b) Intersected connection.

In order to gain an insight into these issues, consider a collection of objects with i.i.d. sizes following a random variable T for which the probability of T exceeding x falls off with x as $cx^{-\alpha}$. If the objects are disposed on a linear (e.g. time) axis, the probability of a point in this axis belonging to an object with size larger than x is proportional:

a) to the probability $cx^{-\alpha}$ of T exceeding x; and

b) to the size of the sampled object, that exceeds x.

Hence it must fall off at least as slow as $c'x^{-\alpha+1}$. However, any p.d.f. must fall off faster than x^{-1}, otherwise it would not be summable. Therefore, α must be at least 2 in order that the objects may be feasibly accommodated along a one-dimensional axis. This means that the probability of sizes being larger than x must fall off by at least a factor of 100 when x increases by a decade (factor of 10). Therefore, the rate of occurrence of sizes must decrease by at least 1000 times per decade. Because current pricing methods in the Internet do not punish users for requesting very large objects, such strongly diminishing rates for the occurrence of large sizes cannot be expected to occur, and indeed they do not occur. For this reason, this kind of traffic cannot be accommodated by a single server (n=1). Section 3 will show that multiple servers (n>1) can do it, but some higher moments of the queue length will still run away to infinity, causing high traffic variability. Alternatively, traffic shaping (or even policing) may impose a limit for the object sizes on the WDM network. While t_{min} can be used to make minimum holding time worth compared with the time spent setting up and tearing down connections, t_{max} might prove useful in reducing the number of servers required in order to bringing down waiting times.

3. Pareto Servicing, Residual Lifetimes, and Waiting Time in Queues

3.1. Truncated Pareto Distribution for Connection Servicing Time

The Pareto distribution is usually specified by its cumulative distribution function (c.d.f.), given by:

$$F_T(t) = prob(T \geq t) = \begin{cases} 0, & t < t_{\min} \\ 1 - \left(\dfrac{t_{\min}}{t}\right)^{\alpha}, & t \geq t_{\min} > 0 \end{cases} \tag{1}$$

The positive parameters t_{min} and α specify the distribution. The probability density function (p.d.f.) may be obtained by differentiation of (1) and it can be truncated at t_{max}:

$$p_T(t) = \begin{cases} 0, & t < t_{\min} \\ \dfrac{\alpha\, t_{\min}{}^{\alpha}}{t^{\alpha+1}} \left(\dfrac{1}{1-r^{\alpha}}\right), & t_{\min} \leq t < t_{\max} \end{cases} \quad \text{, where } r = \frac{t_{\min}}{t_{\max}} \tag{2}$$

If and when it exists, the mean τ of the Pareto truncated distribution is given by:

$$\bar{t} = \tau = \frac{\alpha}{\alpha - 1} t_{\min} \left(\frac{1 - r^{\alpha-1}}{1 - r^{\alpha}}\right), \qquad \alpha > 1 \tag{3}$$

Clearly, the mean does not exist if $\alpha<1$. For this reason, we will assume that $\alpha>1$ in the ensuing discussion. This assumption implies some loss of generality, but no loss of practicality in the framework of our discussion.

Let us now consider the second moment of the distribution:

$$E(t^2) = \frac{\alpha}{\alpha-2} t_{\min}{}^2 \left[1 - r^{\alpha-2}\right] \quad , \quad \alpha > 2$$

By the same token as for the mean, the second moment, and hence the variance, will also be infinite for $\alpha \le 2$ for $r \to 0$ ($t_{max} \to \infty$). However, we are not willing to disregard this case, since the empirical evidence indicates that it is exactly the case of the Internet traffic. If $\alpha \le 2$, we will say that the Pareto distribution when $r \to 0$ is heavy-tailed and has no variance (or that it has infinite variance).

Let us consider a single server that is servicing a connection (burst or packet) at some time and has another finite queue of connections (burst or packets) to be served at the moment when a new connection request arrives. Let q be the number of requests in the queue, i.e. the queue length. Let us normalize all waiting times with respect to the mean service lifetime τ by making $\tau = 1$. According to residual lifetimes presented in Appendix 1, the mean waiting time the new connection will have to wait to be serviced is effectively proportional to: q+1/2, if the service time is deterministic; q+2/3, if the service time is uniformly distributed; q+1, if the service time is exponential (markovian case); $q+\infty = \infty$, if the service time is heavy tailed Pareto ($\alpha < 2$); and q+1, if the service time is light tailed Pareto with $\alpha \cong 2.4142$.

Comparing the five cases considered above, one can see that, except for the heavy-tailed Pareto case, in all other cases one can assume that the total waiting time is dominated by the queue length q whenever q >> 1. This is of course the reason why classic queueing system analysis is so much focused on characterization of queue lengths. When the packet duration distribution is heavy-tailed Pareto, however, waiting time in a M/P/1 queue is not dominated by the queue length for any finite queue, but rather by the single connection (or packet) that is currently being served. Therefore, the classic focus on queue length is not warranted in this situation.

When a newcomer connection request (client) arrives at a queue shown in Fig. 1(a) that is waiting for one or more such entities to be served, we will say that the newcomer intersects each ongoing service interval at some age t_0 (see Fig. 1(b)). The client view for the case of a single server (n=1) is presented in Section 3.2 while generalization for the multiple servers (n>1) to represent WDM networks will be addressed in Section 3.3.

3.2 – A Client View of Single Server

Let Y be the total duration of the intersected interval, and let $p_Y(y)$ be its p.d.f.. The probability $p_Y(y)dy$ of Y belonging to the {y,y+dy] interval must be proportional:

a) to the probability $p_T(y)dy$ that $y<T<y+dy$, since Y is drawn from a population of intervals with distribution $p_T(.)$; and

b) to the duration y of the intersected interval, since a randomly chosen time will be proportionately more likely to intersect longer intervals.

One may then write:

$$p_Y(y) = Ayp_T(y) \ , \tag{4}$$

where A is a normalization constant, which may be determined by simple integration:

$$1 = \int_{-\infty}^{\infty} p_Y(y)dy = A\int_{-\infty}^{\infty} yp_T(y)dy = A\tau$$

$$\therefore \ p_Y(y) = \frac{1}{\tau} yp_T(y) \tag{5}$$

Without knowing about the residual lifetime paradox [4], one might be led to think that a randomly arriving newcomer would see a mean residual lifetime equal to one half of the mean lifetime at birth. Actually this is true only for periodic traffic (fixed duration). However, in all other cases, there is a discrepancy between these two values. Suppose $p_T(.)$ is exponential with mean τ. As a result, (5) yields E(y) = 2τ. This result is a paradox since the client is intersecting connections with residual lifetime τ, as seen in Appendix 1. The discrepancy is only a factor of 2 in the Markovian case (exponential distribution), and a factor of 4/3 for uniformly distributed durations. Therefore, the mean waiting time for the residual lifetime may in general be considered of lesser importance when compared with the waiting time for large queues. In the case of heavy-tailed durations, on the other hand, the paradox has much stronger implications: the discrepancy runs away to infinity, making the residual lifetime contribution to the waiting time more important than the queue itself.

Now let us consider what the intersected interval p.d.f. will be when the i.i.d. queued intervals are α-Pareto (i.e. Pareto distributed with shape parameter α). From (2), (3) and (5):

$$p_Y(y) = \begin{cases} 0, & y < t_{\min} \\ \dfrac{\alpha}{\tau(1-r^{\alpha})}\left(\dfrac{t_{\min}}{y}\right)^{\alpha}, & t_{\min} \le y < t_{\max} \end{cases} \tag{6}$$

Eq. (6) states that <u>if the i.i.d. intervals are α-Pareto with $t_{max}\to\infty$, then the intersected interval will be (α-1)-Pareto with the same minimal duration t_{min}</u>. If 1<α<2, the i.i.d. intervals will have finite mean τ and infinite variance when r→0, but the intersected interval will have infinite variance <u>and mean</u>. Therefore, the mean residual lifetime after any finite age must be infinite, For α>2, we have from (3) and (6):

$$\bar{y} = \frac{\alpha-1}{\alpha-2} t_{\min}\left(\frac{1-r^{\alpha-2}}{1-r^{\alpha-1}}\right) = \frac{\alpha}{\alpha-2}\frac{t_{\min}^{\ 2}}{\tau}\left(\frac{1-r^{\alpha-2}}{1-r^{\alpha}}\right) \tag{7}$$

For the light-tailed Pareto (α>2) case, (6) will yield a finite mean for Y. In addition, if α>3, Y will have a finite variance. In general, α>j+1 assures the existence of a finite j-th moment of residual lifetime even with $t_{max}\to\infty$.

The intersected interval distribution (5) is the effective duration of the ongoing process as seen from a newcomer client. Therefore, it may provide a basis for analyzing residual lifetime by averaging over the time axis, and not over the more cumbersome sample space as done in the Appendix 1.

If Z is the residual lifetime (see Fig. 1(b)) of the process when it was intersected at age t_0, we may write [4, Chap. 5]:

$$t_0 + Z = Y \tag{8}$$

Since the intersection time is uniformly chosen in time, Z is uniformly distributed between 0 and Y for any known value of Y, so:

$$prob[Z \le z \mid Y = y] = \frac{z}{y} \qquad z < y \tag{9}$$

$$\therefore\ prob[z \le Z < z + dz, y \le Y < y + dy] = \frac{dz}{y}\frac{\alpha}{\tau(1-r^{\alpha})}\left(\frac{t_{\min}}{y}\right)^{\alpha} dy,$$

$t_{\min} < y < t_{\max}, y > z$

$$\therefore\ p_{YZ}(y,z) = \frac{\alpha}{\tau(1-r^{\alpha})}\frac{t_{\min}^{\alpha}}{y^{\alpha+1}}$$

$$\therefore\ p_Z(z) = \int_0^{\infty} p_{YZ}(y,z)dy = \frac{\alpha}{\tau(1-r^{\alpha})} t_{\min}^{\alpha} \int_a^{t_{\max}} \frac{dy}{y^{\alpha+1}}, \qquad \text{where } a = \max(z, t_{\min})$$

Then,

$$p_Z(z) = \begin{cases} \dfrac{1}{\tau}, & z < t_{\min} \\ \dfrac{1}{\tau(1-r^{\alpha})}\left[\left(\dfrac{t_{\min}}{z}\right)^{\alpha} - r^{\alpha}\right], & t_{\min} \le z < t_{\max} \end{cases} \tag{10}$$

The mean residual lifetime will then be:

$$\bar{z} = \int_0^{\infty} z p_Z(z)dz = \int_0^{t_{\min}} \frac{z}{\tau}dz + \int_{t_{\min}}^{t_{\max}} \frac{z}{\tau(1-r^{\alpha})}\left[\left(\frac{t_{\min}}{z}\right)^{\alpha} - r^{\alpha}\right]dz \tag{11}$$

The second integral on the right-hand side is not summable for $\alpha<2$ and $t_{max}\rightarrow\infty$, indicating that the mean residual lifetime will be infinite for heavy-tailed Pareto traffic. On the other hand, finite mean residual lifetime is obtained when connections are truncated at finite t_{max}. For $\alpha>2$, (11) may be further manipulated to yield:

$$\bar{z} = \frac{\alpha}{2(\alpha-2)}\frac{t_{\min}{}^{2}}{\tau}\left(\frac{1-r^{\alpha-2}}{1-r^{\alpha}}\right) = \frac{1}{2}\bar{y} \tag{12}$$

This reconciles the intuitive notion that the mean residual lifetime should be one half of the mean lifetime at birth. But it shows that this idea is correct only when referred to the intersected interval (y), and not to a randomly picked interval (t) seen in Fig. 1(b).

3.3. WDM Networks

Let us now consider a queue generated by Poissonian arrivals and Pareto service time duration, with n servers as a simplified, i.e. one hop connection request, model for the WDM network illustrated in Fig.1. The waiting time for the first server to be released is the minimal residual lifetime of the n busy servers intersected by the newcomer client.

Let Z_1, Z_2, ..., Z_n be the residual lifetimes of the n processes intersected by the new arrival. The effective waiting time for a server to be released will then be:

$$Z = \min\left[Z_1, Z_2, ..., Z_n\right] \tag{13}$$

The random variables Z_i, i=1, 2, ..., n, are i.i.d. with p.d.f. given by (10):

$$p_{Z_i}(z) = \begin{cases} \dfrac{1}{\tau}, & z < t_{\min} \\ \dfrac{1}{\tau(1-r^{\alpha})}\left[\left(\dfrac{t_{\min}}{z}\right)^{\alpha} - r^{\alpha}\right], & t_{\min} \le z < t_{\max} \end{cases} \tag{14}$$

If $F_z^{\{n\}}(z)$ is the cumulative distribution function (c.d.f.) of Z, we have:

$$1 - F_z^{\{n\}}(z) = prob\left[Z > z\right] = prob\left\{\min\left[Z_1, Z_2, ..., Z_n\right] > z\right\} =$$

$$= prob\left[Z_1 > z, Z_2 > z, ..., Z_n > z\right] = \left[1 - F_{Z_i}(z)\right]^n$$

$$\therefore \mathrm{F}_Z^{\{n\}}(z) = 1 - \left[1 - \mathrm{F}_{Z_i}(z)\right]^{n}$$

Differentiating with respect to z:

$$p_z^{\{n\}}(z) = n\left[1 - F_{Z_i}(z)\right]^{n-1} p_{Z_i}(z) = np_{Z_i}(z)\left[\int_z^{t_{\max}} p_{Z_i}(z_i)dz_i\right]^{n-1} \tag{15}$$

From (14), we then have:

$$p_z^{\{n\}}(z) = \begin{cases} \dfrac{n}{\tau}\left(1 - \dfrac{z}{\tau}\right)^{n-1}, t < t_{\min} \\ n\left(\dfrac{1}{\tau(1-r^{\alpha})}\right)^{n}\left[\left(\dfrac{t_{\min}}{z}\right)^{\alpha} - r^{\alpha}\right]\left[z\left(\dfrac{1}{\alpha-1}\left(\dfrac{t_{\min}}{z}\right)^{\alpha} - r^{\alpha}\right) - \dfrac{\alpha}{\alpha-1}t_{\max}r^{\alpha}\right]^{n-1}, t_{\min} \le t < t_{\max} \end{cases} \tag{16}$$

4. Implications on Networking

The implications of these traffic analyses on networking principles and guidelines are a much tougher issue. Early appraisals of the evolution of the WDM network seemed to imply that it might lead to the supply of dedicated wavelengths to bandwidth-hungry customers. However, it is shown in this Section that our results on the M/P/n queue suggest that a dedicated wavelength is not a good deal for a client with Pareto traffic, nor very likely to any client with long-range dependent traffic. Instead, sharing a pool of wavelengths with other clients with the same kind of traffic is probably more effective to fight the self-similarity syndrome that would otherwise plague each of these clients. We also suggest the use of traffic shaping (or policing) limiting holding time to t_{max} in order to reduce the number of wavelength needed to mitigate the instability effects in networks operating with $\alpha \le 2$.

4.1. Unbounded Holding Time

From (16) one can see that in order for Z to have a finite mean for $r \to 0$, we must then have:

$$n(\alpha - 1) + 1 > 2, \quad \therefore n > \frac{1}{\alpha - 1}, \tag{17}$$

and in order that Z has a second moment and a finite variance:

$$n(\alpha - 1) + 1 > 3, \quad \therefore n > \frac{2}{\alpha - 1} \tag{18}$$

In general, we need $k/(\alpha-1)$ servers for the waiting time to have a finite k-th moment. For example, if $\alpha = 1.2$, we need more than 5 servers for the waiting time to have a finite mean, more than 10 servers for a finite variance, and over 15 servers for a finite third moment, and so on. Increasing the size of the wavelength pool would first produce a finite mean waiting time, then make it reasonably small, and then do the same to higher moments of the waiting time distribution. However, although larger and larger wavelength pools would bring higher and higher moments of the waiting time distribution to finite values, there would always be some sufficiently high moments running away to infinity. The effect of the infiniteness of such high moments of the residual waiting time on the queuing behavior is not clear yet, but it seems reasonable to expect that both waiting times and queue lengths would become more predictable as more and more moments are brought down to finite values.

4.2. Bounded Holding Time

Traffic shaping (or policing) may impose t_{max} as the maximum holding time allowed in the network as a means of bringing moments to finite values with fewer servers. Fig. 2-4 show mean and variance for the residual lifetime (assuming t_{min}=1s) for t_{max}= 100, 10^3,10^4, and 10^6 s. For a network under traffic with α=1.2, suppose there is limit, set at 10s, connection requests, on average, can wait to be served due to QoS constraints. Fig.2 shows that a single server could cope with this requirement provided the maximum holding time of connection is kept below 100s while 6 servers are required in the unbounded case.

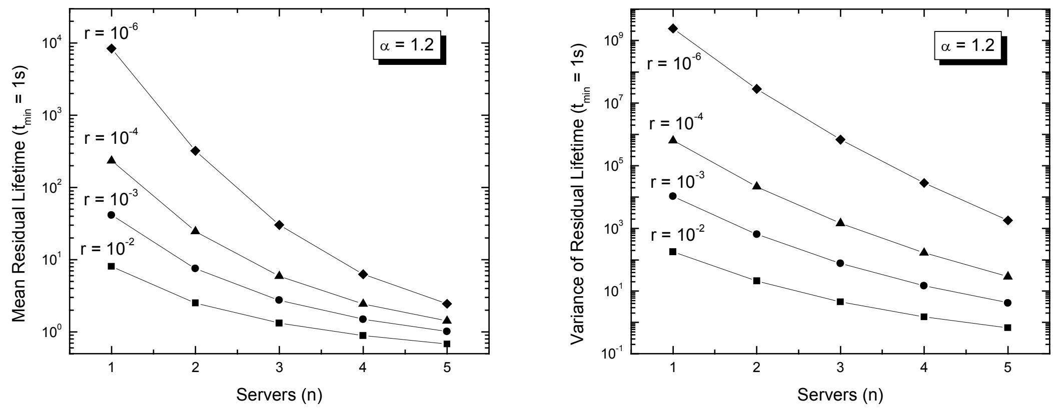

(a) (b)

Figure 2. Residual Lifetime for α = 1.2. (a) Mean and (b) Variance

Even in cases connections are allowed to be held for very long periods compared with t_{min}, e.g. 10^6 which corresponds to eleven days and a half, the QoS requirement can be matched with just 4 wavelengths. As expected, variance reduction in Fig. 2(b) is less sensitive to the increase in the wavelength pool. However, limiting t_{max} is proving an effective way to bring down variance in case this is a QoS issue.

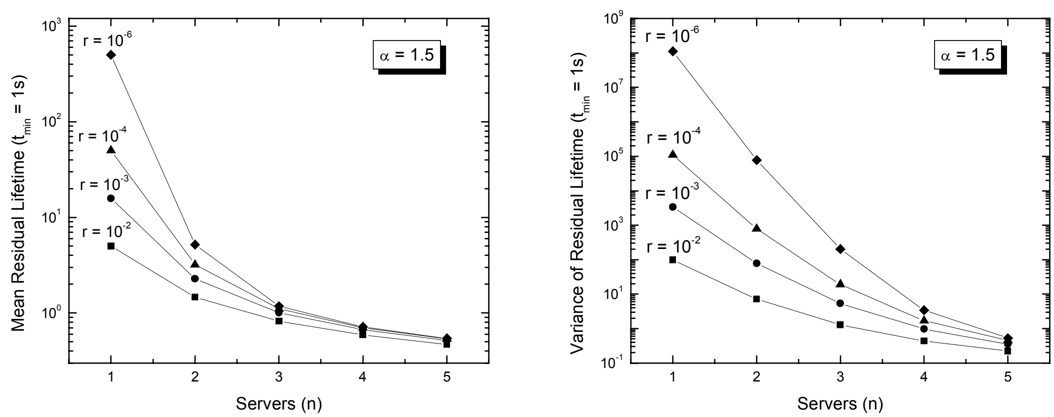

(a) (b)

Figure 3. Residual Lifetime for α = 1.5. (a) Mean and (b) Variance

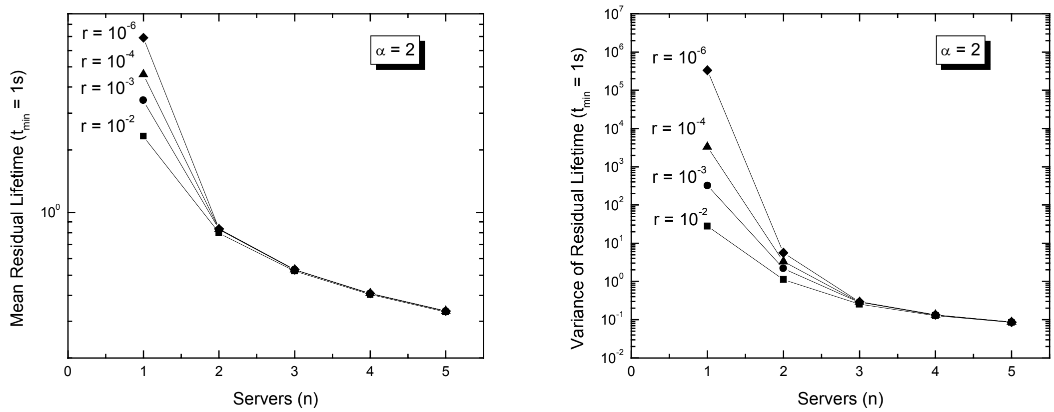

(a) **(b)**

Figure 4. Residual Lifetime for α = 2.0. (a) Mean and (b) Variance

As it could be expected, when the number of servers reaches values found in (17) and (18), t_{max} becomes irrelevant for mean and variance, respectively. This can be clearly seen in networks under less demanding traffic features, e.g. $n \geq 3$ in Fig. 3(a), $n \geq 2$ in Fig. 4(a), and $n \geq 3$ in Fig. 4(b). Nevertheless, limiting the holding time remains a relevant issue as far as reducing higher moments is concerned.

5. Conclusion

We have derived some properties regarding moments of the residual lifetime distribution when the process duration is Pareto-distributed, and have suggested that such residual lifetime may have an important role in the behavior of waiting time in queueing systems, in accordance with reported simulations of unstable behavior of such queues [3]. Current WDM networks operate only on the wavelength granularity, so they can see the traffic burstiness only after it has been shaped by the aggregation process in upper layers (ATM or IP) and by the mapping process into a SDH transport layer.

The current trend towards direct mapping of the IP layer into WDM, however, will change this picture. A likely evolution will replace the current overlay model with a peer-to-peer networking philosophy in which IP routers, enabled by MPLS, will do both traffic routing and wavelength assignment and routing. This will create an opportunity to combine traffic routing and network configuration into the same functionality, for a more effective optimization of traffic in the network. New algorithms must then be devised in order to explore this extended capability. Good traffic modeling may be critical to reliably evaluate these algorithms. One important feature of emerging networks is the aggregation of traffic with different QoS requirements. Appropriate service disciplines must be designed to conciliate these requirements between themselves and with the efficient utilization of network resources. Several studies have shown that some service disciplines are able to drastically reduce waiting time in queues of objects with heavy-tailed size distribution [6].

If WDM networks are considered bufferless and transmitted payload is known in advance (e.g. as in file downloads), session duration is also known at admission time. Under this framework, traffic scheduling is possible, either to accomodate competing QoS requirements or to improve network utilization, or both. This is one emerging networking framework to be considered for future algorithmic studies.

Acknowledgements

This work has been supported by Ericsson Telecomunicações S.A. and by CNPq, Brazil.

References

[1] W. Leland, M. Taqqu, W. Willinger, and D.V. Wilson, "On the self-similar nature of Ethernet traffic", IEEE/ACM Transactions on Networking, vol. 2, pp. 1-15, 1994.

[2] K. Park and W. Willinger, "Self-Similar Network Traffic: an Overview", in Self-Similar Network Traffic and Performance Evaluation, Ch. 1, edited by K. Park and W. Willinger, Wiley, 2000.

[3] M.E. Crovella and A. Bestavros, "Self-Similarity in World Wide Web traffic: evidence and possible causes", IEEE/ACM Transactions on Networking, vol. 5, n. 6, pp. 835-846, 1997.

[4] L. Kleinrock, "Queueing Systems, vol.1: Theory", Wiley, 1975.

[5] A. Feldmann, "Characteristics of TCP Connection Arrivals", in Self-Similar Network Traffic and Performance Evaluation, Ch. 1, edited by K. Park and W. Willinger, Wiley, 2000.

[6] J.W. Roberts, "Engineering for Quality of Service", in Self-Similar Network Traffic and Performance Evaluation, Ch. 1, edited by K. Park and W. Willinger, Wiley, 2000.

[7] A. Schwartz and A. Weiss, "Large Deviations for Performance Analysis", Chapmand and Hall, 1995.

Appendix 1. Residual Lifetimes

Pareto residual lifetime distribution is here considered as a function of the elapsed time t_0. For the sake of comparison, we will discuss the residual lifetimes for some other duration distributions more commonly found in classical frameworks. The residual lifetime S of a process with duration T after an elapsed time t_0 is the remaining duration of the process after it has already lasted a time t_0. Given the p.d.f. $p_T(t)$ of the process duration T, the p.d.f. $p_S(s|\, t_0)$ of the residual lifetime S will be:

$$p_S(s \mid t_0) = \frac{p_T(t_0 + s)}{\int_{t_0}^{\infty} p_T(t)dt} \qquad \text{(A1)}$$

We will often be interested in the mean residual lifetime, given by:

$$\bar{s}(t_0) = \int_0^{\infty} s p_S(s \mid t_0) ds = \frac{\int_{t_0}^{\infty} t p_T(t)dt}{\int_{t_0}^{\infty} p_T(t)dt} - t_0 \qquad \text{(A2)}$$

Let us derive the mean of $\bar{s}(t_0)$ over the distribution $p_0(t_0)$ of the "age" t_0 of the processes. This distribution depends also on the inter-arrival time distribution, i.e. on the statistical regime that rules the birth of new processes. Let us assume, for the sake of simplicity, that arrivals are poissonian, with a rate of λ arrivals per second. Then, the average number $n(t_0)dt_0$ of surviving processes (in sample space) with age between t_0 and (t_0+dt_0) at any given time t is the number of processes born between $(t- t_0-dt_0)$ and $(t- t_0)$ that are still active at t, i.e. that have lifetimes longer than t_0:

$$n(t_0)dt_0 = \lambda dt_0 \int_{t_0}^{\infty} p_T(t)dt \qquad \text{(A3)}$$

Normalization of (A3) yields the age p.d.f.:

$$p_0(t_0) = \frac{\int_{t_0}^{\infty} p_T(t)dt}{\int_0^{\infty} dt_0 \int_{t_0}^{\infty} p_T(t)dt} \quad , \qquad t_0 > 0 \qquad \text{(A4)}$$

The global mean residual lifetime will then be:

$$\bar{\bar{s}} = \int_0^{\infty} \bar{s}(t_0) p_0(t_0) dt_0 \qquad \text{(A5)}$$

Let us now apply Eqs. (A1)-(A5) to a few examples in Table 1.

Table 1. Mean residual lifetime ($\bar{s}(t_0)$) and global mean residual lifetime ($\bar{\bar{s}}$) for some distributions.

Distribution	$\bar{s}(t_0)$	$\bar{\bar{s}}$
Deterministic	$\tau - t_0$	$\frac{\tau}{2}$
Uniform	$\tau - \frac{t_0}{2}$	$\frac{2\tau}{3}$
Exponential	τ	τ
Pareto	$\bar{s}(t_0) = \tau - t_0, \quad t_0 < t_{\min}$ $\bar{s}(t_0) = \frac{t_0}{\alpha - 1}, \quad t_0 \geq t_{\min}$	$\bar{\bar{s}} = \frac{(\alpha-1)^2}{2\alpha(\alpha-2)}\tau$
Truncated-Pareto	$\bar{s}(t_0) = \tau - t_0, \quad t_0 < t_{\min}$ $\bar{s}(t_0) = \frac{\alpha}{\alpha - 1}\left(\frac{t_0^{-\alpha+1} - t_{\max}^{-\alpha+1}}{t_0^{-\alpha} - t_{\max}^{-\alpha}}\right) - t_0$ $, t_{\min} \leq t_0 < t_{\max}$	$\bar{\bar{s}} = \frac{t_{\min}}{\tau + t_{\min} r^{\alpha}}\left[\tau(2 - r^{\alpha}) - t_{\min}\left(\frac{\alpha(\alpha-3)}{2(\alpha-1)(\alpha-2)} + \frac{\alpha r^{\alpha-2}}{2(\alpha-2)}\right)\right]$

Numerical examples are also given in Fig. (A1) in order to illustrate the behavior of the expected residual lifetimes against elapsed time. Pareto distribution bears $\alpha = 1.2$, and t_{min} = 5s, while truncated Pareto bounds the same distribution at 100s (t_{max}). The other distributions use τ found for the truncated Pareto using the parameters above, i.e. 13.9s.

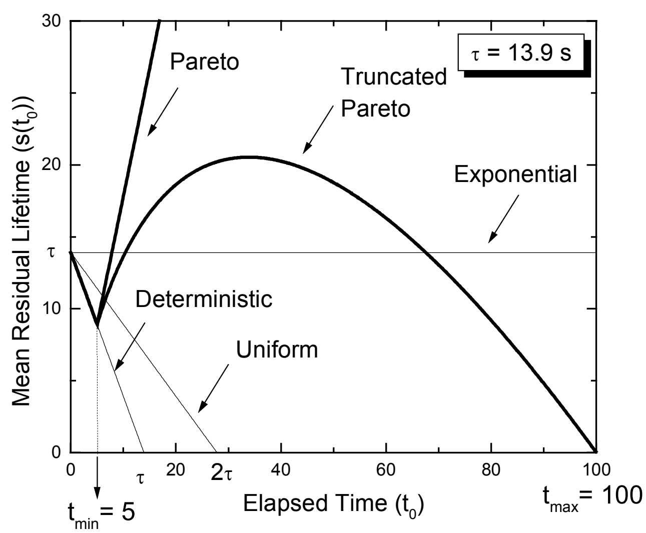

Figure A1. Expected residual lifetime vs. elapsed time. τ is the mean lifetime at birth