MODELING BROADBAND TRAFFIC STREAMS

Timothy D. Neame, Moshe Zukerman
The University of Melbourne
Parkville, Vic. 3052, Australia.
Email: {t.neame, m.zukerman}@ee.mu.oz.au

Ronald G. Addie
University of Southern Queensland
Toowoomba, Qld. 4350, Australia.
Email: addie@usq.edu.au

Abstract

This paper presents the M/Pareto model as a potentially useful tool in the modeling of broadband traffic streams. We show that this model can be used to accurately estimate the queueing performance of a variety of realistic multimedia traffic streams. We also show a practical way to fit the parameters of the model with those of the real traffic stream. We also point out that it it is not practical to seek a perfect model and that a consensus around a model such as the M/Pareto is important.

1 Introduction

Over the past two decades there were many proposals for broadband and/or multimedia traffic models (see, for example [2, 5, 6, 7, 8, 9, 11, 14, 15] and references therein). Despite the extensive effort to find a perfect model, there is still no consensus on a traffic model which is used by all practitioners for analysis and performance evaluation of new products and protocols, and for network dimensioning. There is no model which fills the role that the Erlang formula has in traditional telephony. Lack of such consensus makes it difficult to compare and choose between competing products, as different benchmark measures make it impossible to compare "apples with apples." In recent work [3, 4, 12, 13] the M/Pareto process has been examined as a model for broadband traffic. Because of its simple and realistic features, we propose the M/Pareto as the standard model for broadband traffic, even though it is not "perfect,"

In this paper we summarise the results of the previous work, and present a critical evaluation of the capabilities of the M/Pareto process. We show that the M/Pareto process can accurately predict the queueing performance of a wide range of broadband traffic sources. We also take the opportunity to correct formulae relating to the M/Pareto process which were quoted incorrectly previously, and show derivations for the corrected expressions.

We use a form of the M/Pareto model given in [11] as a model for broadband traffic streams. Using a discrete time modeling framework, we show that when the M/Pareto process is correctly matched to a realistic broadband traffic stream, it produces queueing performance results which match those of the original traffic.

In Section 2 we explain our requirements for a useful traffic model, and define the queueing framework used in evaluating our models. In Section 3 we describe the M/Pareto model and explain how the properties of the M/Pareto ¡ process make it possible for us to fit multiple M/Pareto processes ¡ with differing levels of aggregation to a given traffic stream. Section 4 presents a summary of results showing that the M/Pareto model can accurately model realistic broadband traffic sources. In Section 5 we discuss some of the difficulties in using the M/Pareto model which have yet to be overcome.

2 Modeling a Traffic Stream

We model a single link in a packet switched network as a FIFO single server queue (SSQ) with an infinite buffer. We divide our time scale into fixed length sampling intervals, and let A_n be a continuous random variable representing the amount of work entering the system during the nth sampling interval. Let the server have a constant service rate of τ per interval. The process $\{A_n\}$ is assumed to be stationary and ergodic. We define $\mathrm{E}(A_n)$ and σ^2 to be the mean of $\{A_n\}$ and variance of $\{A_n\}$ respectively. We further assume that $\{A_n\}$ is a long range dependent (LRD) process with Hurst parameter H.

We base our decisions on the suitability of a model upon its ability to characterise the buffer overflow probability of measured traffic in the SSQ system just described. For most practical purposes, an accurate prediction of queueing performance in a given scenario is sufficient. Of course, an ideal model will also correctly model the marginal distributions and correlations of the modeled traffic, but these niceties tend to come with a price of increased model complexity.

3 The M/Pareto Model

We have described the M/Pareto model in [3, 4, 12]. The M/Pareto process is a particular type of the general $M/G/\infty$ process considered in [9] and can also be considered a type of Poisson burst process, as described in [14]

As for any Poisson burst process, the Poisson rate, λ, of the M/Pareto process controls the frequency with which new bursts commence. The superposition of two independent Poisson burst processes with identical burst length distributions will itself be an Poisson burst process with Poisson arrival rate equal to the sum of the arrival rates of the two constituent processes. Thus, increasing λ can be considered to represent an increase in the number of sources which make up an M/Pareto stream.

The cell arrival process for each burst is constant for the duration of that burst, and has rate r. All bursts generate cells at the same rate r. The burst duration is taken from a Pareto distribution. The complementary distribution function for a Pareto-distributed random variable is given by

$$\Pr(X > x) = \begin{cases} \left(\frac{x}{\delta}\right)^{-\gamma}, & x \geq \delta, \\ 1, & \text{otherwise}, \end{cases} \tag{1}$$

$1 < \gamma < 2$, $\delta > 0$. The mean of X is $\frac{\delta\gamma}{(\gamma-1)}$ and the variance of X is infinite.

Thus the mean number of cells within one burst is: $\frac{r\delta\gamma}{\gamma-1}$. The mean amount of work arriving within an interval of length t in the M/Pareto traffic model is $\frac{\lambda t r \delta\gamma}{(\gamma-1)}$.

Although the Pareto process has infinite variance, the variance of the M/Pareto process is finite. The variance function of the M/Pareto process is

$$\sigma^2(t) = \begin{cases} 2r^2\lambda t^2\left(\frac{\delta}{2}\left(1 - \frac{1}{1-\gamma}\right) - \frac{t}{6}\right), & 0 \leq t \leq \delta \\ 2r^2\lambda\left\{\delta^3\left(\frac{1}{3} - \frac{1}{2-2\gamma}\right.\right. & \\ \quad \left. + \frac{1}{(1-\gamma)(2-\gamma)(3-\gamma)}\right) & \\ \quad +\delta^2\left(\frac{1}{2} - \frac{1}{1-\gamma}\right. & \\ \quad \left. + \frac{1}{(1-\gamma)(2-\gamma)}\right)(t-\delta) & \\ \quad \left. - \frac{t^{3-\gamma}}{\delta^{-\gamma}(1-\gamma)(2-\gamma)(3-\gamma)}\right\}, & t > \delta \end{cases} \tag{2}$$

This is a more general expression of the variance function given for processes of this type in [2] and represents a correction to the variance function quoted in [3, 4, 12]. A full derivation of the variance function for an M/Pareto process is given in Appendix A.

Throughout our modeling we make use of the fact that, as a type of Poisson burst process, the M/Pareto model has the useful property that the superposition of multiple independent M/Pareto processes is itself an M/Pareto process, with increased Poisson arrival rate, λ. We use of this property to increase the level of aggregation in the M/Pareto process. We follow the fitting method described in [12] to produce a family of M/Pareto processes with given mean net arrival rate, variance and Hurst parameter, but with varying levels of aggregation. When we do so, we observe that increasing the level of aggregation in the M/Pareto process, i.e. increasing λ, alters the behaviour of the M/Pareto process. In fact, as λ increases the M/Pareto process behaves more and more like a Gaussian process.

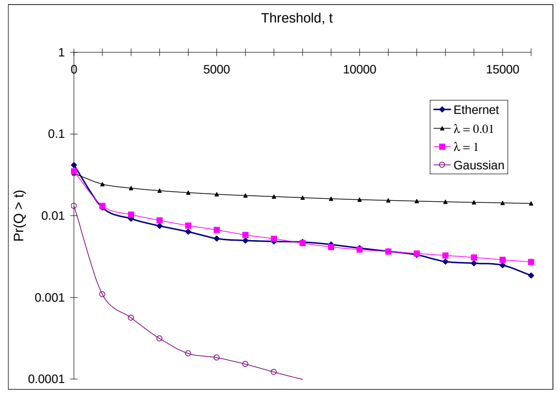

Figure 1: Modeling an Ethernet trace

4 Modeling Results

In [4] we showed some success in using the technique described above to fit an arbitrary Ethernet trace. As Figure 1 shows, for a correct choice of the value of λ, the M/Pareto model accurately fit with the queueing performance of a real Ethernet trace.

Figure 1 illustrates this result. In the figure, the $\lambda = 0.01$ curve represents an M/Pareto process fitted to the three parameters (m, variance and Hurst parameter) of the real traffic, with an arbitrary choice of λ. The $\lambda = 1$ curve represents a fitting of the same values of m, σ^2 and H, but with a more careful choice of λ. The *Gaussian* curve in the figure shows that the Ethernet traffic stream cannot be modeled by a Gaussian process.

As was suggested in [1], where a sufficiently large number of independent sources contribute to an aggregate traffic stream, i.e. when λ becomes sufficiently large in our M/Pareto model, we expect that stream to assume the properties of an Gaussian process. It seems that in the Ethernet traffic, the number of contributing sources is insufficient for the Gaussian behaviour to manifest itself. Using the M/Pareto model with a correct choice of λ, it is possible to accurately estimate the queueing performance for this non-Gaussian stream.

In [13] we showed similar results for an IP packet trace. As with the Ethernet trace, a Gaussian process was not able to accurately predict the queueing performance of the IP stream. The M/Pareto process did provide accurate prediction of the queueing performance, when the value of λ was chosen correctly. It is worth noting that the value of λ required to characterise the IP stream was larger than that required for the Ethernet trace. Our interpretation of λ suggests that this would imply a greater number of sources contributed to the IP trace than to the Ethernet trace.

Given the interpretation of λ as the level of aggregation in the stream, it is perhaps not surprising that it needs to be correctly fitted in order to create a model which can be used to match an aggregated traffic stream such as Ethernet or IP traffic. However aggregated data streams are not the only source of broadband traffic streams. In [12] we examined another major source of broadband traffic: VBR video traces. As for the data traffic, our ability to predict the queueing performance of the VBR video traffic was found to depend upon the choice of λ. Thus it seems that some representation of the level of aggregation is important in predicting the queueing performance even of traffic streams in which the level of aggregation has no physical meaning.

5 Limitations of Our Approach

As we stated in Section 2, we have evaluated the accuracy of the M/Pareto model only in terms of its ability to predict queueing performance. It is not necessarily useful in any other characterisation of traffic or of network performance.

As Arvidsson and Karlsson have observed [6], the impact of retransmissions in the real data traffic streams will affect the accuracy of our model. We assume that the measured traces contain only user data, but in reality a proportion of the packets will be the result of retransmissions of lost packets. If we dimension based on the assumption that all packets are user packets we will over-dimension the network. Nevertheless, if we then measure traffic from the over-dimensioned network, it should be free of retransmissions, allowing us to make a better revised decision.

Finding the correct parameters for our M/Pareto process is not a simple task. Evaluating H from a real (and therefore finite) traffic stream is never easy. As discussed in [8], it is not always possible to distinguish between an LRD process and a non-stationary one. Our assumption throughout has been that the traffic under examination is both stationary and LRD, but testing these assumptions is no easy task.

We have seen that choosing the right value for λ is vital in creating an M/Pareto process capable of matching the queueing curves of real traffic streams. However the choice of λ is complicated by the fact that the correct value of λ differs depending on the service rate τ (or equivalently the value of m).

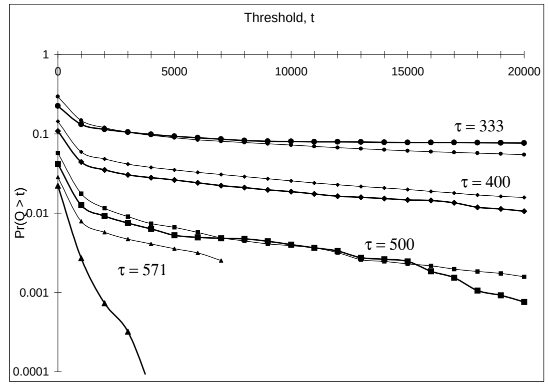

Figure 2: Fitting an Ethernet trace for different service rates

In Figure 2, we show the queueing curves produced when a pair of traffic streams are fed into SSQs with a variety of services rates. The figure shows four pairs of curves. In each pair the heavier line represents the queueing performance of the Ethernet trace when fed into an SSQ with service rate τ. The lighter line represents the performance of an M/Pareto process matched to the properties of the Ethernet trace in an identical SSQ. The M/Pareto process used has λ chosen so as to provide a good fit with the Ethernet traffic when $\tau = 500$. As Figure 2 shows, while a given value of λ may give an acceptable fitting for a range of service rates, in general a different value of λ must be determined for each service rate considered.

As yet we have no systematic method for determining the value of λ to be used in modeling a given traffic stream. Trial and error must be used for each different traffic source, and for each different service rate. In every case we have considered so far it has been possible to find a value of λ which is appropriate, but a systematic method would greatly enhance this process. Until a heuristic for determining λ is developed, this will limit the practical usefulness of the M/Pareto process.

6 Conclusions

In this paper we have presented the M/Pareto process as a suitable practical model for broadband traffic. We have shown that the M/Pareto process is capable of accurately matching the queueing performance of a range of different broadband traffic types. The simplicity and accuracy of the M/Pareto model make it a good candidate for the role of the standard model for broadband traffic.

We have also discussed some of the limitations of our use of the M/Pareto process. Some of these limitations, such as the difficulty in estimating the Hurst parameter from a finite data set, are common to all models, but others are specific to the M/Pareto model. One of the key limitations of the M/Pareto process is the lack of a simple formula or heuristic to determine the correct value for λ. If such a formula can be developed, the M/Pareto process will present a useful practical approach for the modeling of multimedia and data traffic.

Appendix A: Derivation of M/Pareto Variance

In [14] the variance function for a Poisson burst process with bursts arriving with Poisson rate λ, and burst rate r is given by

$$\sigma^2(t) = 2\lambda r^2 \int_0^t du \int_0^u dv \int_v^\infty \Pr(X > x)dx \qquad (3)$$

where $\Pr(X > x)$ is the complementary distribution function describing the burst durations.

We use an M/Pareto process with complementary distribution function given by

$$\Pr(X > x) = \begin{cases} \left(\frac{x}{\delta}\right)^{-\gamma}, & x \geq \delta, \\ 1, & \text{otherwise}, \end{cases} \qquad (4)$$

where $1 < \gamma < 2$, $\delta > 0$.

Define $f(v)$ as the first stage of the integration given in Equation (3). When $v \geq \delta$

$$\begin{aligned} f(v) &= \int_v^\infty \Pr(X > x)dx \\ &= \int_v^\infty \left(\frac{x}{\delta}\right)^{-\gamma} dx \\ &= \frac{-v^{-\gamma+1}}{\delta^{-\gamma}(-\gamma+1)} \end{aligned}$$

When $v < \delta$

$$\begin{aligned} f(v) &= \int_v^\infty \Pr(X > x)dx \\ &= \int_v^\delta 1dx + \int_\delta^\infty \left(\frac{x}{\delta}\right)^{-\gamma} dx \\ &= \delta - \frac{\delta}{(-\gamma+1)} - v \end{aligned}$$

Now define $g(u)$ as the integral of $f(v)$ over the bounds required by Equation (3). When $u \geq \delta$

$$\begin{aligned} g(u) &= \int_0^u f(v)dv \\ &= \int_0^\delta \left(\delta - \frac{\delta}{(-\gamma+1)} - v\right) dv \\ &\quad + \int_\delta^u \frac{-v^{-\gamma+1}}{\delta^{-\gamma}(-\gamma+1)} dv \\ &= \left[\left(\delta - \frac{\delta}{(-\gamma+1)}\right) v - \frac{v^2}{2}\right]_0^\delta \\ &\quad + \left[\frac{-v^{-\gamma+2}}{\delta^{-\gamma}(-\gamma+1)(-\gamma+2)}\right]_\delta^u \\ &= \delta^2\left(\frac{1}{2} - \frac{1}{(-\gamma+1)}\right) \\ &\quad - \frac{u^{-\gamma+2}}{\delta^{-\gamma}(-\gamma+1)(-\gamma+2)} + \frac{\delta^2}{(-\gamma+1)(-\gamma+2)} \end{aligned}$$

When $u < \delta$

$$\begin{aligned} g(u) &= \int_0^u f(v)dv \\ &= \int_0^u \left(\delta - \frac{\delta}{(-\gamma+1)} - v\right) dv \\ &= \left(\delta - \frac{\delta}{(-\gamma+1)}\right) u - \frac{u^2}{2} \end{aligned}$$

The final integral is

$$\sigma^2(t) = 2\lambda r^2 \int_0^t g(u)du. \qquad (5)$$

For $t \geq \delta$,

$$\begin{aligned} \sigma^2(t) &= 2\lambda r^2 \int_0^t g(u)du \\ &= 2\lambda r^2 \left\{ \int_0^\delta \left(\left(\delta - \frac{\delta}{(-\gamma+1)}\right) u - \frac{u^2}{2}\right) du \right. \\ &\quad + \int_\delta^t \left(\delta^2\left(\frac{1}{2} - \frac{1}{(-\gamma+1)} + \frac{1}{(-\gamma+1)(-\gamma+2)}\right)\right. \\ &\quad \left.\left. - \frac{u^{-\gamma+2}}{\delta^{-\gamma}(-\gamma+1)(-\gamma+2)}\right) du \right\} \\ &= 2\lambda r^2 \left\{ \left[\left(\delta - \frac{\delta}{(-\gamma+1)}\right)\frac{u^2}{2} - \frac{u^3}{6}\right]_0^\delta \right. \\ &\quad + \left[\delta^2\left(\frac{1}{2} - \frac{1}{(-\gamma+1)} + \frac{1}{(-\gamma+1)(-\gamma+2)}\right) u \right. \\ &\quad \left.\left. - \frac{u^{-\gamma+3}}{\delta^{-\gamma}(-\gamma+1)(-\gamma+2)(-\gamma+3)}\right]_\delta^t \right\} \\ &= 2\lambda r^2 \left\{ \frac{\delta^3}{2}\left(1 - \frac{1}{(-\gamma+1)}\right) - \frac{\delta^3}{6} \right. \end{aligned}$$

$$
+\delta^2\left(\frac{1}{2}-\frac{1}{(-\gamma+1)}+\frac{1}{(-\gamma+1)(-\gamma+2)}\right)(t-\delta)
-\frac{t^{-\gamma+3}-\delta^{-\gamma+3}}{\delta^{-\gamma}(-\gamma+1)(-\gamma+2)(-\gamma+3)}\Bigg\}
$$

So when $t \geq \delta$

$$
\begin{aligned}
\sigma^2(t) = \quad & 2r^2\lambda\Big\{\delta^3\left(\tfrac{1}{3}-\tfrac{1}{2-2\gamma}+\tfrac{1}{(1-\gamma)(2-\gamma)(3-\gamma)}\right) \\
& +\delta^2\left(\tfrac{1}{2}-\tfrac{1}{1-\gamma}+\tfrac{1}{(1-\gamma)(2-\gamma)}\right)(t-\delta) \\
& -\tfrac{t^{3-\gamma}}{\delta^{-\gamma}(1-\gamma)(2-\gamma)(3-\gamma)}\Big\}
\end{aligned} \tag{6}
$$

When $t < \delta$

$$
\begin{aligned}
\sigma^2(t) &= 2\lambda r^2 \int_0^t g(u)du \\
&= 2\lambda r^2 \int_0^t \left(\left(\delta-\frac{\delta}{(-\gamma+1)}\right)u-\frac{u^2}{2}\right)du \\
&= 2\lambda r^2 \left(\left(\delta-\frac{\delta}{(-\gamma+1)}\right)\frac{t^2}{2}-\frac{t^3}{6}\right)
\end{aligned}
$$

$$
\sigma^2(t) = 2\lambda r^2 t^2\left(\frac{\delta}{2}\left(1-\frac{1}{(-\gamma+1)}\right)-\frac{t}{6}\right) \tag{7}
$$

If we choose the case where $\delta = 1$ and set $\beta = -\gamma$ in Equations (6) and (7) we see that these are equivalent to the expressions given in [2].

References

[1] R. G. Addie. On the Weak Convergence of Long Range Dependent Traffic Processes. In *Proceedings of the International Workshop on Long Range Dependence*, January 1997. Accepted for publication in *Journal of Statistical Planning and Inference*.

[2] R. Addie, P. Mannersalo and I. Norros. Performance Formulae for Queues with Gaussian Input. In *Proceedings of ITC 16*, June 1999, pp 1169–1178.

[3] R. G. Addie, M. Zukerman and T. D. Neame. Broadband Traffic Modeling: Simple Solutions to Hard Problems. *IEEE Communications Magazine*, Vol. 36, No. 8, August 1998.

[4] R. G. Addie, T. D. Neame and M. Zukerman. Modeling Superposition of Many Sources Generating Self Similar Traffic. In *Proceedings of ICC '99*, June 1999.

[5] D. Anick, D. Mitra and M. M. Sondhi. Stochastic Theory of a Data Handling System with Multiple Sources. *Bell System Tech. Journal*, Vol. 61, No. 8, 1982, pp 1871 – 1894.

[6] A. Arvidsson and P. Karlsson. On Traffic Models for TCP/IP. In *Proceedings of ITC 16*, June 1999, pp 457–466.

[7] H. Heffes and D. M. Lucantoni. A Markov Modulated Characterization of Packetized Voice and Data Traffic and Related Statistical Multiplexer Performance. *IEEE Journal on Selected Areas in Communications*, Vol. SAC-4, No. 6, September 1986, pp 856–867.

[8] K. Kobayashi and Y. Takahashi. Tail Probability of a Gaussian Fluid Queue under Finite Measurement of Input Processes. *Performance and Management of Complex Communication Networks*, Eds. T. Hasegawa, H. Takagi and Y. Takahashi, Chapman & Hall, 1998, pp 43 – 58.

[9] M. M. Krunz and A. M. Makowski. Modeling Video Traffic Using $M/G/\infty$ Input Processes: A Compromise Between Markovian and LRD Models. *IEEE Journal on Selected Areas in Communications*, Vol. 16, No. 5, June 1998.

[10] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson. On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Transactions on Networking*, Vol. 2, No. 1, 1994, pp 1–15.

[11] N. Likhanov, B. Tsybakov and N. D. Georganas. Analysis of an ATM Buffer with Self-Similar ("Fractal") Input Traffic. In *Proceedings of Infocom '95*, April 1995.

[12] T. D. Neame, M. Zukerman and R. G. Addie. Applying Multiplexing Characterization to VBR Video Traffic. In *Proceedings of ITC 16*, June 1999.

[13] T. D. Neame, M. Zukerman and R. G. Addie. A Practical Approach for Multimedia Traffic Modeling. In *Proceedings of Broadband Communications '99*, November 1999.

[14] J. Roberts, U. Mocci, and J. Virtamo. *Broadband Network Teletraffic, Final Report of Action COST 242*. Springer, 1996.

[15] K. Sriram and W. Whitt. Characterizing Superposition Arrival Processes in Packet Multiplexers for Voice and Data. *IEEE Journal on Selected Areas in Communications*, Vol. SAC-4, No. 6, September 1986, pp 833–846.