Worst Case Arrivals of Leaky Bucket Constrained Sources: The Myth of the On-Off source

Philippe Oechslin

Department of Computer Science,

University College of London

Gower Street, London WC1E 6BT, UK

+44 171 380 72 96,

+44 171 387 13 97 fax

p.oechslin@cs.ucl.ac.uk

Abstract

We have simulated a set of independent connections limited by leaky bucket shapers and fed into a buffered multiplexer. This scenario is typical of an ATM switch or in a looser sense of an RSVP capable router. We found periodic traffic patterns which result in much worse loss rates than the on-off or tri-state patterns found in literature to date. We give an intuitive justification for what we believe is the worst case and back this with an extensive set of simulations. Our results are important for Connection Acceptance Control when connections are known to be statistically independent. They clearly invalidate the widespread belief that on-off patterns are the worst case traffic of independent leaky bucket constrained sources.

Keywords

leaky bucket, worst case, traffic control, traffic modelling

1 INTRODUCTION

The leaky bucket traffic descriptor has been chosen as the traffic descriptor for ATM networks and in a looser sense in the integrated services Internet. The advantage of the leaky bucket is that it makes it easy to verify whether a source conforms to its traffic descriptor. However, it is very difficult, given the leaky bucket parameters, to estimate the exact amount of resources that a set of connections will require. This information is needed at connection setup time to know whether a new connection can be accepted or not. This is the problem of Connection Acceptance Control (CAC).

A recent overview of existing CAC schemes is given in (Perros and Elsayed 1996). The goal of a good CAC scheme is to accept as many connections

as are possible without disrupting the contracted Quality of Service (QoS) of accepted connections. Most CAC schemes make some assumptions about traffic, in order to be able to estimate the resources required by a connection. The most common assumption is that traffic in each connection conforms to a given stochastic process, usually some kind of on-off process.

In this paper we make no assumption about the traffic. Each connection may thus have an arbitrary distribution. The only assumption we make is that in each connection the traffic conforms to the declared leaky bucket parameters. Furthermore we assume that the connections are independent of each other, in other words that there is no correlation between them. Typically, connections can be assumed independent when they are unrelated (eg. originating and terminating at different hosts) and when the network does not introduce artificial correlation.

Under these assumptions we try to find the maximum loss rate which can occur in a multiplexer given the leaky bucket parameters of all connections. In particular we try to find the worst kind of input traffic (referred to as the worst case) which leads to the maximum loss rate.

1.1 Existing work

It is a common belief that on-off sources are the worst case of independent, leaky bucket constrained sources, as in (Kositpaiboon and Phung 1990), (Rathgeb 1991), (Kvols and Blaabjerg 1992), (Worster 1994), (Johri 1995), (Elwalid, Mitra and Wentworth 1995) and (Presti, Zhang, Kurose and Towsley 1997). This is due to the fact that on-off patterns have the highest variance of all possible patterns and because of the assumption that a higher variance always leads to a higher loss rate. However, (Doshi 1994) and (Yamanaka, Sato and Sato 1992) have shown counter examples where a pattern called tri-state pattern results in more loss. In (Worster 1994) the author points to some possible flaws in (Yamanaka et al. 1992) and tries to re-establish that on-off patterns are the worst case.

For the case of correlated, leaky bucket constrained sources exact solutions for delay and queue-length bounds have been found. First results can be found in (Cruz 1991), with an extended and more elegant form in (Le Boudec 1996). An example of worst case patterns for correlated traffic based on these results can be found in (Lee 1994).

2 THE EXPERIMENT

We consider a multiplexer with N incoming connections, all described by their leaky bucket parameters m for mean rate, p for peak rate and b for burst size.

Figure 1 the multiplexer

The multiplexer has a FIFO buffer of size X and outputs its traffic on a link with capacity $\frac{1}{\rho} \sum_{i=1}^N m_i$. For reasons of stability ρ must be smaller than 1.

We consider the homogeneous case where all connections have the same leaky bucket parameter. We apply a deterministic periodic pattern to all inputs of the multiplexer and measure its average loss rate. The loss rate is defined as the amount of data lost due to buffer overflow, divided by the amount of data offered to the multiplexer. Since the same pattern is applied to all inputs, the only difference between the traffic in the connections is the phase of the patterns. The connections are assumed to be independent and thus the phases have a uniform random distribution.

For this experiment we have chosen following conditions: $N = 20$, $p = 5$, $m = 1$, $b = 20$, and $\rho = 0.99$. Traffic is assumed to be fluid and the discrete size of packets is not taken into account. The effect of packetization can be made arbitrarily small by using adequate units for the different parameters.

2.1 Traffic Patterns

The traffic patterns used in our simulations are derived from the on-off pattern. An on-off pattern consists of a burst at peak rate with size b , followed by a silence of duration b/m after which a new burst can be sent. We derive new patterns from the on-off pattern and express their difference by a form factor, f or g . We report the results for the following patterns:

on_off

The maximum allowed burst followed by the shortest period of silence allowing a new burst.

tri_state(f)

An on-off pattern with an inserted plateau at mean rate between the burst and the silence period. f denotes the total duration of the pattern in units of the on-off pattern. Hence $\text{tri_state}(1) = \text{on_off}$.

shift(f, g)

The shift pattern takes a tri_state pattern and shifts the burst to a point between the beginning and the end of the plateau. g denotes the time between the beginning of the pattern and the beginning of the burst in proportion to the length of the plateau. We thus have $0 \leq g \leq 1$ and $\text{shift}(f, 0) = \text{tri_state}(f)$.

sym(f)

Sym is a symmetrical pattern which corresponds to an on-off pattern with two identical plateaus inserted before and after the burst. f denotes the length of the pattern in units of the length of the on-off pattern. We have $\text{sym}(1) = \text{on_off}$ and $\text{sym}(f) = \text{shift}(f, \frac{1}{2})$.

2.2 Simulation Results

We now describe the results obtained from the simulations. All plotted results have a 95% confidence interval smaller than $\pm 5\%$ of their value. Figure 2 shows the loss rate in the multiplexer as a function of its buffer size. We fix a reference point at buffer size 180. At this point we see that the average loss with on_off patterns is about 2.2×10^{-5} .

Figure 3 shows the loss rate at the reference point for tri_state patterns as a function of the form factor f . As hinted in the literature we find that the loss rate can be higher than for on_off patterns. Note that $\text{tri_state}(1)$ is equal to on_off, thus for $f = 1$ the loss rate of tri_state is the same as for on_off. We see that for small f the loss rate initially increases. It reaches a maximum of 1.05×10^{-4} at $f = 1.4$. tri_state patterns can thus be worse than on_off ones, in the particular case by almost a factor of ten.

Based on the worst tri_state pattern mentioned above we next investigate the effect of shifting the burst. Figure 4 shows the loss rate as a function of the form factor of $\text{shift}(1.4, g)$. We see that the curve is symmetrical and that the loss rate is maximal when the burst is in the centre of a plateau at mean rate. This finding motivates the next simulations using the sym pattern.

Figure 3 shows the loss rate for sym(f) as a function of the form factor f . Again $\text{sym}(1) = \text{on_off}$ and the loss rate at $f = 1$ is the same as for on_off. For $f = 3.5$ the loss rate reaches 1.05×10^{-3} , yet another factor of ten above

Figure 2 loss rate as a function of buffer size for selected patterns

Figure 3 loss rate for tri_state pattern as a function of f

the maximum loss rate of the tri_state pattern. Also, we see that after the maximum the loss rate does not decrease as rapidly as it does for the tri_state pattern. We will explain this effect in the next section.

Figure 2 also shows the loss rate as a function of the buffer size for tri_state(1.4) and sym(3.5). We can see that for buffer sizes greater than 100 tri_state(1.4) and sym(3.5) are worse than the on_off pattern. Note that the f which maximises the loss rate depends on the buffer size, thus sym(3.5) is only the worst case for a buffer of size 180.

Figure 4 loss rate of $\text{shift}(1.4, g)$ as a function of g

3 DISCUSSION

Let us first introduce the notion of a *full rate multiplexer*. We speak of a full rate multiplexer when the sum of the mean rates of its inputs is equal to its output rate ($\rho = 1$). Furthermore we define the set of *full burst sources* as the set of periodic sources which have plateaus at the mean rate m , one burst of size b at a rate between m and p and one period of 'rest' where they send at a rate smaller than m until the burst is compensated for. Note that all the patterns investigated belong to full burst sources and that full burst sources have a mean rate of m .

Consider a full rate multiplexer fed by N independent periodical full burst sources. The set of sources can be seen as N sources sending continuously at their mean rate plus N sources sending positive and negative bursts of size b . A single positive burst will occupy b amount of buffer space in the multiplexer. Since in absence of bursts the output rate of the multiplexer is equal its input rate, the buffer occupied by that single burst will only be released when a negative burst occurs. Full burst patterns all have alternating positive and negative bursts of the same size. Now define the centre of the bursts as the point in time where half of the burst has been transmitted. Call the interval between the centre of a positive burst and the centre of the following negative burst Φ_1 . Call the interval between the centre of the negative burst and the centre of the next positive burst Φ_0 . Consider the case where n positive bursts need to add up to overflow the buffer. Their centres need to be within a period smaller than Φ_1 to avoid the last positive burst being cancelled by the negative burst following the first positive burst. The centre of the bursts also need to occur in an interval smaller than Φ_0 . If not, the preceding negative burst of

the last positive burst falls into the interval. Thus for the buffer to be occupied by at least n bursts, n or more bursts must occur within $\min(\Phi_1, \Phi_0)$.

The above does not consider the cases where positive and negative bursts partially overlap. To study this effect let us introduce one last pattern, $\text{delta}(f)$.

Instead of having positive and negative bursts at peak rate and zero, $\text{delta}(f)$ has bursts of $b\delta(t)$ and $-b\delta(t)$ where b is the burst size and $\delta(t)$ is Dirac's impulse function. f is the length of the pattern measured in the length of the original on-off pattern and both burst are distant of $\frac{f}{2}$ from each other.

For the delta pattern, the loss *per period* only depends on the buffer size X and not on the period f . Indeed changing f has no effect on the shape of the pattern itself and thus on the probability of buffer overflow. We have

$$E[\text{loss per period}] = l(X)$$

On the other hand, the loss rate is inversely proportional to the period since the same loss occurs for every period. We thus have:

$$E[\text{loss rate}] = \frac{1}{f}l(X)$$

Figure 5 loss rate of the full rate multiplexer

The only difference between the delta pattern and the sym pattern is that

the bursts of the delta pattern cannot overlap because they have zero duration. Thus we can attribute differences in loss rate between the delta and sym patterns to the overlap of positive and negative burst in the connections with sym patterns. Indeed, consider two sets of delta and sym patterns with identical distribution of phases. If in the set of sym patterns there is no overlap of any positive and negative burst, then the sym patterns will produce exactly the same amount of loss. However, if a positive burst overlaps with a negative one, then the bursts will reduce each others effect. If that positive burst participates in the loss produced by the set of sym patterns, then the loss will be reduced.

To verify the above statements we have simulated the case of the full rate multiplexer with delta, sym and tri.state sources. The results are in figure 5. Note the linearity of the loss rate of $\text{delta}(f)$ which confirms its inverse proportionality to f . As expected the loss rate of $\text{sym}(f)$ converges asymptotically towards $\text{delta}(f)$ as f increases and the relative duration of the bursts decreases.

We now come back to our general non-full-rate multiplexer to explain a final effect of f on the loss rate. In a non-full-rate multiplexer the output rate is larger than the mean input rate. If all inputs send at mean rate and one burst is received, this burst will occupy buffer space only for a limited time. The buffer will be cleared at a rate equal to the difference between the output rate and the mean input rate. This adds another constraint on the series of bursts adding up to overflow the buffer. The longer the period during which the bursts accumulate, the more will their effect be attenuated by the extra output rate of the multiplexer. Thus increasing the period of a pattern – even when its shape is not modified, as for delta – reduces its loss per period in a non-full-rate multiplexer.

3.1 Summary

We have explained the following effects:

1. For full burst patterns, bursts can only accumulate if they happen in intervals shorter than $\min(\Phi_1, \Phi_0)$.
2. Patterns with shorter burst lengths compared to their period have a smaller probability of overlapping burst and can thus generate more loss per period.
3. For the same loss per period, the loss rate is inversely proportional to the period of the pattern.
4. An output rate larger than the sum of the mean input rates reduces the loss rate of patterns with long periods more than the ones with short periods.

Effect 1. is the reason why in all our simulations sym patterns generate more loss than tri.state patterns with the same f . Indeed $\min(\Phi_1, \Phi_0) = \Phi_0 = 1/2$

for the tri_state patterns, whereas it is equal to $f/2$ for sym. In other words, tri_state patterns produce less loss because there is always a negative burst immediately preceding positive bursts, thus preventing longer series of positive bursts to add up. Effect 4. explains the difference in simulations of the original multiplexer and the full rate multiplexer. Both loss rates of tri_state and sym decrease faster with f in the original multiplexer as in the full rate multiplexer. Finally, the opposite effects of 2. versus 3. and 4. are the reason why there is a maximum loss rate for some $f = f_{max}$.

sym(f_{max}) is the worst case pattern of the class of full burst patterns since it maximises $\min(\Phi_1, \Phi_0)$ by having $\Phi_1 = \Phi_0 = \frac{f}{2}$, minimises the probability of overlapping bursts by having the shortest allowed bursts (at peak rate and 0) and balances effects 2. against 3. and 4. by having $f = f_{max}$.

Note that on_off is a special case of sym and that it can be the worst case when $f_{max} = 1$. This could for example be the case in a very under-loaded multiplexer (exacerbating effect 4.) or when the burst size is small compared to peak and mean rate (reducing effect 2.)

4 CONCLUSION

We have shown in simulations, that sym patterns can generate significantly more loss than on_off or tri_state patterns. We have shown four different effects which explain the results of the simulations. Due to these effects the sym pattern is shown to be the worst case within the class of full burst patterns, which includes on_off patterns.

Our simulations definitively invalidate the belief that on-off sources are the worst case for independent leaky-bucket constrained sources. This has an important consequence on existing CAC schemes. Indeed, any CAC scheme which is based on the assumption that sources behave like a given on-off process may grossly underestimate losses if the sources chose to behave like the worst case we have identified.

The sym pattern we propose is not as easy to use as the on-off pattern since it has an additional parameter, the form factor f and we do not yet have an analytical way of finding f_{max} which will maximise the loss.

Further interesting questions are whether there exist worst case patterns outside the set full burst patterns and what the worst case pattern is in the heterogeneous case, where connections have different leaky bucket parameters.

REFERENCES

- Cruz, R. L.: 1991, A calculus for network delay, part I: Network elements in isolation, *IEEE Transactions on Information Theory* 37(1), 114–131.
- Doshi, B. T.: 1994, Deterministic rule based traffic descriptors for broadband ISDN: Worst case behavior and connection acceptance control,

- in J. Labetoulle and J. W. Roberts (eds), *Proceedings of ITC 14*, Elsevier Science B. V., pp. 591–600.
- Elwalid, A., Mitra, D. and Wentworth, R.: 1995, A new approach for allocating buffers and bandwidth to heterogeneous, regulated traffic in an ATM node, *IEEE Journal on Selected Areas in Communications* **13**(6), 1115–1127.
- Johri, P. K.: 1995, Estimating cell loss rates in high-speed networks with leaky bucket controlled sources, *International journal of communication systems* **8**, 303–311.
- Kositpaiboon, R. and Phung, V.: 1990, Usage parameter control and bandwidth allocation for B-ISDN/ATM variable bit rate services, *Proceedings of IEEE Multimedia '90*.
- Kvols, K. and Blaabjerg, S.: 1992, Bounds and approximations for the periodic on/off queue with applications to ATM traffic control, *Proceedings of Infocom '92*, pp. 487–494.
- Le Boudec, J.-Y.: 1996, Network calculus made easy, *Technical Report 96/218*, EPFL/DI, Ecole Polytechnique Federale de Lausanne.
- Lee, D. C.: 1994, Effects of leaky bucket parameters on the average queueing delay: Worst case analysis, *Proceedings of IEEE Infocom '94*, pp. 482–489.
- Perros, H. G. and Elsayed, K. M.: 1996, Call admission control schemes: A review, *IEEE Communications Magazine* **34**(11), 82–91.
- Presti, F., Zhang, Z., Kurose, J. and Towsley, D.: 1997, Source time scale and optimal buffer/bandwidth trade-off for regulated traffic in an ATM node, *Proceedings of IEEE Infocom '97*.
- Rathgeb, E. P.: 1991, Modeling and performance comparison of policing mechanisms for ATM networks, *IEEE Journal on Selected Areas in Communications* **9**(3), 325–334.
- Worster, T.: 1994, Modelling deterministic queues: The leaky bucket as an arrival process, in J. Labetoulle and J. W. Roberts (eds), *Proceedings of ITC 14*, Vol. 1, Elsevier Science B.V., pp. 581–90.
- Yamanaka, N., Sato, Y. and Sato, K.-I.: 1992, Performance limitation of leaky bucket algorithm for usage parameter control and bandwidth allocation methods, *IEICE Transactions on Communication* **E75-B**(2), 82–86.