On-Off Sources and Worst Case Arrival Patterns of the Leaky Bucket

UCL/CS Research Note 97/125

Philippe Oechslin
Department of Computer Science, University College of London
p.oechslin@cs.ucl.ac.uk

October 31, 1997

Abstract

It has long been believed that on_off patterns where the worst case with respect to the loss rate for independent connections limited by leaky bucket constraints. We report simulation results of a buffered multiplexer fed by a set of independent connections limited by leaky bucket shapers. This scenario is typical of an ATM switch or in a looser sense of an RSVP capable router. Both for homogeneous and heterogeneous traffic we found periodic patterns that resulted in much more loss 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. We also give a formal proof that the investigated pattern generates more loss than on_off patterns. Finally we explain the importance of our results for Connection Acceptance Control when connections are known to be statistically independent.

Keywords: leaky bucket, worst case, connection acceptance control, traffic modelling

1 Introduction

The leaky bucket traffic descriptor has been chosen as the traffic descriptor for ATM networks [5] and in a looser sense in the integrated services Internet[1]. 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 [14]. 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 examine 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.

The rest of this paper is organised as follows. Section 2 describes the experiment we simulated and for which we want to find the worst arrival pattern as well as the set of pattern used. Section 3 reports and discusses the simulation results for the homogeneous case. It ends with an argumentation for what we believe to be the worst case. Section 4 presents the results of simulation with heterogeneous traffic. Section 5 is a formal proof showing that on_off is not the worst arrival pattern. Finally Section 6 discusses the implications of these findings and concludes the paper.

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 [7], [16], [8], [18], [6], [4] and [15]. In [11] we even find a attempted proof of this belief. This is due to the fact that on-off patterns have the maximum possible variance of all patterns, the largest possible burstsize, and because they maximise the chernoff bound, which is a bound on the probability that a sum of random variables exceeds a given limit. However, [19] [3] and have shown counter examples where a pattern called tri-state pattern results in more loss. The results of [19] have been widely ignored and in [18] the author even points to some possible flaws in [19] 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 [2], with an extended and more elegant form in [9]. An example of worst case patterns for correlated traffic based on these results can be found in [10].

Part of the work presented here has appeared in [12, 13].

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. The multiplexer has a FIFO buffer of size X and outputs its traffic on a link

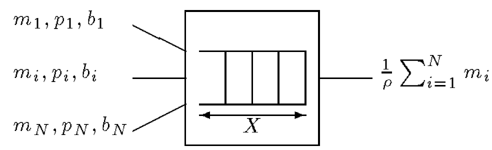

Figure 1: the multiplexer

with capacity $\frac{1}{\rho}\sum_{i=1}^{N} m_i$. For reasons of stability ρ (the output link utilisation) must be smaller than 1.

We first 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: $N = 28$, $p = 2.5$, $m = 1$, $b = 15$, and $\rho = 0.947$. 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 Chosing a Loss Measure

There are two main measures for loss in a multiplexer, the saturation probability and the cell loss rate. The first one, often found in queueing theory, is the probability that the length of an unbounded queue is longer than the size of the actual queue. Its advantage is that it can be calculated directly from the distribution of the queue occupation. The cell loss rate is the ratio between the amount of data lost due to buffer overflow and the amount of data offered to the queue. Although mathematicaly less tractable the cell loss rate gives a more concrete information to the user of a network. Both loss measures are believed to be asymptoticaly equivalent for large networks and they are believed to lead to the same CAC decisions. In this paper we have chosen cell loss ration as loss measure as because it is also the measure used to define general QoS in ATM networks [5].

2.2 Arrival Patterns

The arrival 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

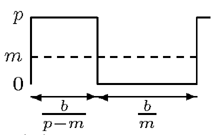

The maximum allowed burst followed by the shortest period of silence allowing a new burst.

tri_state(f)

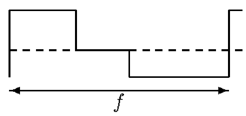

An on-off pattern with an inserted plateau at mean rate between the burst and the silence period. f denotes the total duration of the pattern in units of the on-off pattern. Hence tri_state(1) = on_off.

shift(f, g)

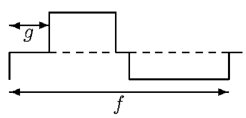

The shift pattern takes a tri_state pattern and shifts the burst to a point between the beginning and the end of the plateau. g denotes the time between the beginning of the pattern and the beginning of the burst in proportion to the length of the plateau. We thus have $0 \leq g \leq 1$ and shift(f,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

sym(f)

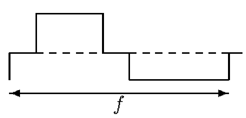

Sym is a symmetrical pattern which corresponds to an on-off pattern with two identical plateaus inserted before and after the burst. f denotes the length of the pattern in units of the the length of the on-off pattern. We have sym(1) = on_off and sym(f) = shift($f, \frac{1}{2}$)

3 The Homogeneous Case

3.1 Simulation Results

We now describe the results obtained from the simulations. Where no error marks appear on the plotts, the results have a 95% confidence interval smaller than ±10% 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 220. At this point we see that the average loss with on_off patterns is about $2.8\text{x}10^{-9}$.

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 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 $4.8\text{x}10^{-8}$ at $f = 1.3$. tri_state patterns can thus be worse that on_off ones, in the particular case by more than 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 shift($1.3, g$). We see that the curve is symmetrical and that the

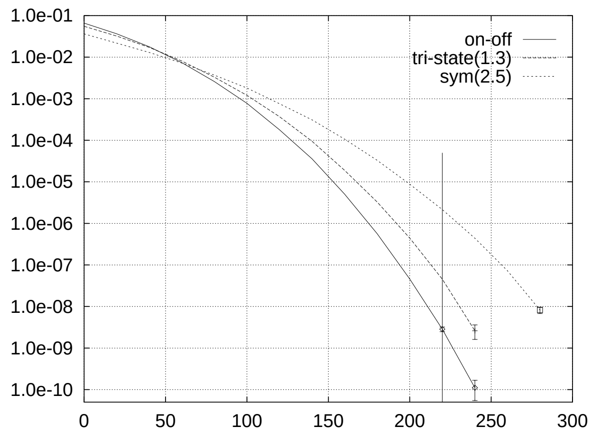

Figure 2: loss rate as a function of buffer size for selected patterns

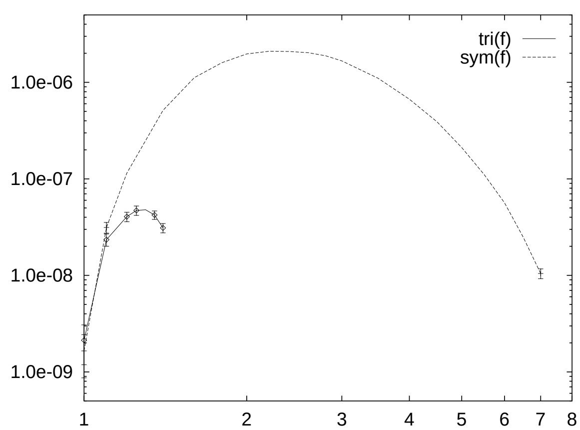

Figure 3: loss rate for tri_state pattern as a function of f

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.

The loss ratee for sym(f) is also plotted in Figure 3. Again sym(1) = on_off and the loss rate at $f = 1$ is the same as for on_off. For $f = 2.5$ the loss rate reaches 2.1 10^{-6}, about three orders of magnitude above the loss rate of the on_off 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.

The loss rate of tri_state(1.3) and sym(2.5) as a function of the buffer size are also plotted in Figure 2. We can see that for buffer sizes greater that 70 tri_state(1.3) and sym(2.5) are worse than the on_off pattern. Note that the f which maximises the loss rate depends on the buffer size, thus sym(2.5) is only the worst case for a buffer of size 220 and actually produces less loss than on_off patterns for buffer smaller than 50.

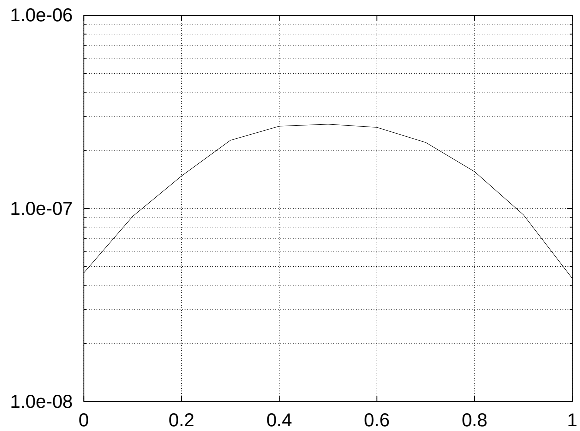

Figure 4: loss rate of shift$(1.4, g)$ as a function of g

3.2 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 between 0 and 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, delta(f).

delta(f) 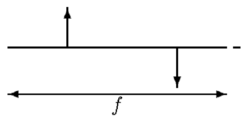

Instead of having positive and negative bursts at peak rate and zero, delta(f) has bursts of $b\delta(t)$ and $-b\delta(t)$ where b is the burst size and $\delta(t)$ is Dirac's impulse function. f is the length of the pattern measured in the length of the original on-off pattern and both burst are distant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

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

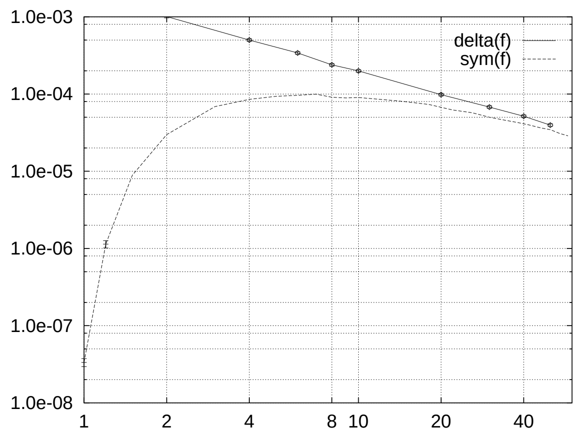

Figure 5: loss rate of the full rate multiplexer

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 bursts 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 delta(f) which confirms its inverse proportionality to f. As expected the loss rate of sym(f) converges asymptotically towards 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.

3.3 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 $f = f_{max}$.

We thus conjecture that 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 The Heterogenous Case

In this section we discuss the results of simulations with two different arrival patterns feeding the multiplexer simultaneously.

4.1 Heterogeneous f

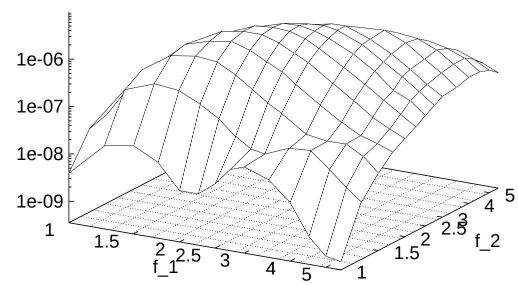

Figure 6: cell loss rate for varying f_1 and f_2, all connections having the same traffic descriptor

Previouly we only looked at the case where all connections have the same traffic parameters and the same form factor f. We assumed that the worst case would happen when all connections had the same f_{max}. We this simulation we try to give some support to this assumption. For the same experiment as above, we divide the 28 connections into two groups of 14, with form factors f_1 and f_2 respectively.

There are a couple of interestig facts to be seen in the resulting 3D plot (Figure 6). First, as expected, the diagonal section of the plot, defined as $f_1 = f_2$ is identical to the plot in Figure 3. Second, the highest point on that diagonal is also the highest point over the whole hyperplane, meaning that the maximum loss rate is reached when $f_1 = f_2 = f_{max}$. Third, we see a local maximum around $f_1 = 1$ and $f_2 = 3$ (and vice-versa). This is due to the fact that when one pattern has a length which is an odd multiple of the other one, both the positive and negative bursts of the longer one can coincide with positive and negative bursts of the shorter one. Finally, and most interstingly, the plot shows that f_{max} is not the worst case for one group of connections when the other group has an f different from f_{max}. This has a most important consequence for connection acceptance control. It means that knowing the traffic descriptor of all existing connections doesn't give us enough information to find out the worst behaviour of new connections. Indeed the f_{max} of that new connection will depend on the arrival patterns of the existing connections.

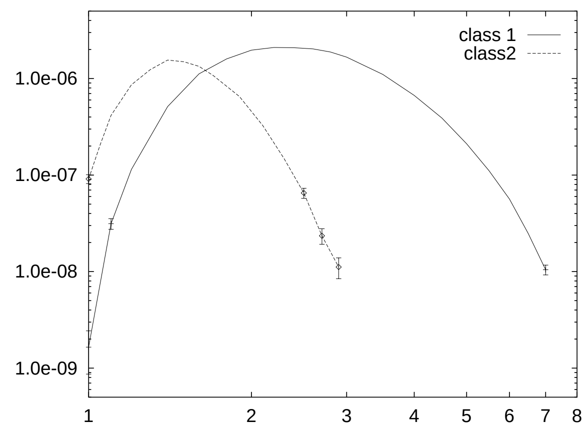

Figure 7: cell loss rate for 28 connections of class 1 or 20 conncetions of class 2 with varying f

4.2 Introducing a Second Traffic Class

We now step to a proper heterogeneous case where we have two traffic classes with different traffic descriptors. We chose a second traffic class which is more bursty. It has the following parameters: $p = 25$ $m = 1$ $b = 24$. From the plot in Figure 7 we see that the multiplexer can accpet about 20 connections of class 2 with about the same loss rate as for 28 connections of the previous class, class 1. In terms of equivalent capacity, we can say that for a loss rate of 2e-6, class 1 traffic has an equivalent capacity of one 28th of the output link (1.06) whereas class 2 has an equivalent capacity of one 20th (1.48). For our heterogeneous simulation we chose 14 connections of class 1 and 10 of class 2. This traffic mix should have the same total equivalent capacity than 28 connections of class 1 or 20 of class 2.

In Figure 8 we see the cell loss rate in function of the form factors of both traffic classes (f_1 and f_2 respectively). Note that for the same f both pattern happen to have the same length. The maximum loss rate of 5.5 10^-7 is reached around $f_1 = 1.4$ and $f_2 = 1.8$. We notice that the maximum loss rate is slightly smaller than for the homogeneous case, by a factor of two. More surprising is that at the maximum f_1 and f_2 are quite different from f_{1max} and f_{2max} and that f_2 is actually larger than f_1. Looking back at Figure 7, this comes quite unexpected. It shows again that with heterogeneous traffic, the worst case pattern of new connections depends on the behaviour of the existing connections and not only on their traffic descriptors.

5 Sym vs On-off: the Proof

In previous sections we have argued why we believe that the sym pattern is the worst case arrival pattern for our experiment. However, we are not able to give a formal proof of that. As a first step in this direction we present a proof which demonstartes that sym patterns are worse than on_off patterns. This contradicts the wide belief that on_off sources are the general worst case as well

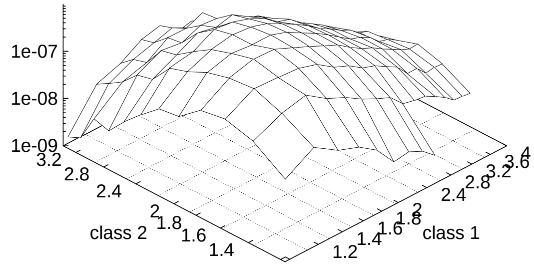

Figure 8: cell loss rate for 14 connections of class1 and 10 connections of class2 with varying f_1 and f_2

as a published proof of that belief. Now, since the on_off pattern is a special case of the sym pattern, there can still be circumstances in which on_off is the worst case. We will discuss this issue as well as the contradicting proof at the end of this section.

We want to prove that on_off pattern do not generally generate more loss than sym patterns. We measure the loss in cell loss rate, defined as the amount of data lost due to buffer overflow divided by the amount of data offered. A very similar proof can be done using saturation probability rather than cell loss rate and can be found in [13].

Define P^+ as an upper bound on the probability that on_off patterns with a set of random phases generate loss, given buffer size X, number of connections N and the parameters of the connections m, p, b. Define l^+ as the maximum loss rate when loss does occur. An upper bound for the cell loss rate $P_{\text{on\_off}}$ is thus :

$$P_{\text{on\_off}} < P^+ l^+$$

Similarly, define P^- as the a probability that sym patterns with a set of random phases generate the maximum possible loss, and l^- the loss rate when maximum loss occurs. $P^- l^-$ is thus a lower bound of P_{sym}:

$$P_{\text{sym}} > P^- l^-$$

Our goal is to show that there are cases where

$$P^+ l^+ < P^- l^- \tag{1}$$

thus proving that the on_off pattern is not always the worst case in buffered systems. To do this we consider the case of a *full rate multiplexer*, meaning a multiplexer with an output link capacity equal to the sum of input mean rates. Furthermore we analyse the cell loss rate for buffer sizes close to the total burst size of all sources. We thus have $\rho = 1$ and $X = Nb - \epsilon$ with $\epsilon < b$

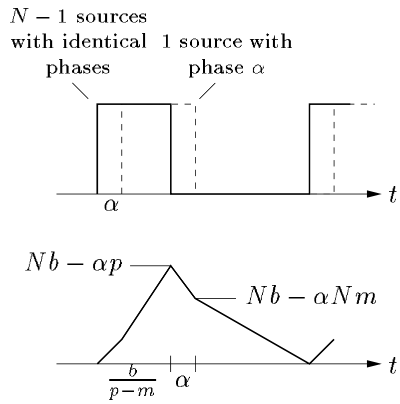

Figure 9: Buffer occupancy in function of time for N sources with phases within α. At the end of the busy period $\frac{b}{p-m}+\alpha$ the buffer occupancy is $Nb-\alpha Nm$. During the busy period a maximum occupancy of $Nb-\alpha p$ can be reached if $N-1$ sources are in phase. Note that $Nb-\alpha p$ is only larger than $Nb-\alpha Nm$ if $p<Nm$

5.1 An Upper Bound for on_off Sources

For $\epsilon=0$ on_off patterns can only fill a buffer of size Nb if all sources have exactly the same phase. For a small ϵ loss only occurs if all phases are within small interval α. P^+ is thus $N(\frac{\alpha}{\omega})^{(N-1)}$ where ω is the length of the on_off pattern, equal to $\frac{bp}{m(p-m)}$.

We now try to find the maximum α for which a buffer of $Nb-\epsilon$ can still be filled. As shown on Figure 9, when phases are within an interval α the buffer will fill up to $Nb-Nm\alpha$ during the active period. We thus have $\epsilon=Nm\alpha$ or $\alpha=\frac{\epsilon}{Nm}$. However, buffer occupancy can go above that level during the busy phase. The maximum occupancy happens if all but one sources are in phase. In that case buffer occupancy reaches $Nb-\alpha p$ which is greater than $Nb-Nm\alpha$ when $p<Nm$. ϵ is then equal to αp or $\alpha=\frac{\epsilon}{p}$. Since P^+ is an upper bound we take the maximum of both expressions which is $\alpha=\frac{\epsilon}{\min(p,Nm)}$ and $P^+=N(\frac{\epsilon(p-m)m}{bp\min(p,Nm)})^{(N-1)}$.

An upper bound of the loss rate that can occur is $\frac{\epsilon}{B}$, where $B=\frac{Nbp}{(p-m)}$ is the total amount of data transmitted during one period. We thus have:

$$P^+l^+=\left(\frac{\epsilon(p-m)m}{bp\min(p,Nm)}\right)^{(N-1)}\frac{\epsilon(p-m)}{bp} \qquad (2)$$

5.2 A Lower Bound for sym Sources

As can be seen from Figure 10, sym patterns can fill a buffer of size Nb only if all phases are within an interval of β which is equal to $\omega\frac{f-1}{2}$, f being the form factor of the pattern, or its length in units of the on_off pattern. For any ϵ, $P^-=N(\frac{\beta}{\omega f})^{(N-1)}$ is thus a lower bound for the probability of loss occurring. Indeed, if all phases are in an interval of β, the buffer would eventually reach Nb thus generating the maximum possible loss.

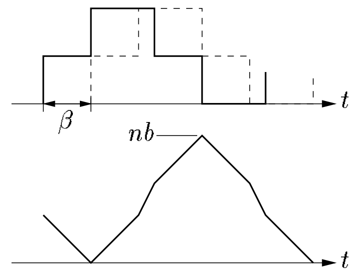

Figure 10: Buffer occupancy resulting of two sym with different phases. As long as the phase difference is smaller than $\beta = \omega \frac{f-1}{2}$ the buffer will fill up to Nb

$l^- = \frac{\epsilon}{fB}$ is the cell loss ratio when maximum loss occurs, the amount of data offered per period being fB for sym patterns. We have:

$$P^- l^- = \left(\frac{f-1}{2f}\right)^{(N-1)} \frac{\epsilon(p-m)}{fbp} \tag{3}$$

5.3 Comparing the Bounds

For $P^- l^-$ to be greater than $P^+ l^+$ the following condition must hold:

$$\left(\frac{f-1}{2f}\right)^{(N-1)} \frac{\epsilon(p-m)}{fbp} > \left(\frac{\epsilon(p-m)m}{bp\min(p, Nm)}\right)^{(N-1)} \frac{\epsilon(p-m)}{bp} \tag{4}$$

Decomposing and simplifying for both outcomes of $\min(p, Nm)$ we find:

$$p > Nm: \left(\frac{f-1}{2f}\right)^{(N-1)} \frac{1}{f} \quad > \quad \left(\frac{\epsilon(p-m)}{bpN}\right)^{(N-1)} \tag{5}$$

$$p \leq Nm: \left(\frac{f-1}{2f}\right)^{(N-1)} \frac{1}{f} \quad > \quad \left(\frac{\epsilon(p-m)m}{bp^2}\right)^{(N-1)} \tag{6}$$

We can chose f to be 2, such that the left hand side becomes $(\frac{1}{4})^{(N-1)}\frac{1}{2}$, a small positive number. The right hand side can be made arbitrarily small by chosing ϵ small enough, thus proving that, for any value of the other parameters, there are cases where on_off patterns can not generate a higher cell loss rate than sym pattern, *QED*.

5.4 Numerical Results

We first look at the example simulated in Section 3 for a full rate multiplexer with parameters $N = 28$, $p = 2.5$, $m = 1$, $b = 15$, $\rho = 1$, $f = 2$ and $\epsilon = 10$ (thus $X = 410$). Note that $p \leq Nm$. We refer to equation 4 to find the loss bounds.

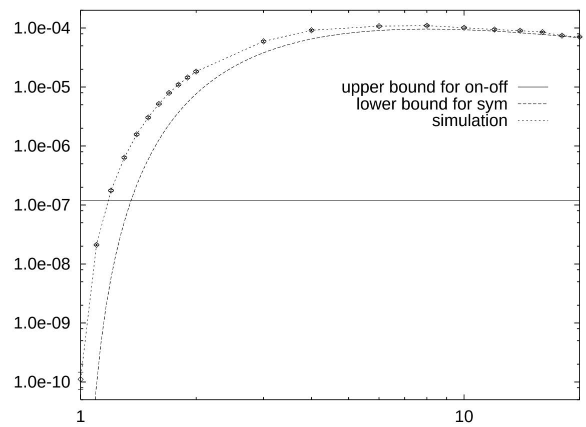

Figure 11: bounds and simulated cell loss rate as a function of f

$$2.8\ 10^{-17} > 3.2\ 10^{-22}$$

which confirms that at least for small values of ϵ the sym pattern produce a higher cell loss rate than on_off patterns. The values are too small to be verified by simulation but we can make a simpler example which we are able to simulate: $N = 8$, $p = 2$, $m = 1$, $b = 20$, and $\rho = 1$. We plot the cell loss rate in function of f for $\epsilon = 10$

Taking equation 4 and replacing the given parameters we find

$$(\frac{f-1}{2f})^7 \frac{1}{4f} > 1.19\ 10^{-7}$$

The right hand side is the upper bound for the cell loss rate of on_off patterns while the left hand side is a lower bound of the same for sym patterns in function of the form factor f. Figure 11 shows both bounds as a function of f as well as simulation results of the cell loss rate for this example. We see than for f between 1.5 and 20, the lower bound for sym patterns is higher than the upper bound of on_off patterns, proving that sym pattern can be worse than on_off patterns. The result is confirmed by the simulation which shows the cell loss rate to be $1\ 10^{-10}$ for the on_off pattern and $1.0\ 10^{-4}$ for sym pattern with form factor of 10.

5.5 Discussion

We have given a formal proof that the on_off pattern is not the worst case pattern for independent leaky bucket constrained sources feeding a buffered multiplexer. We have proven that the sym pattern can generate more loss measured as cell loss rate. The same proof can be made using the probability of saturation as shown in [13]. This result stands in opposition with the general belief and a formal proof [11] that on_off is the worst case.

5.5.1 What about the contradicting proof?

The proof in [11] proves that on_off is the worst case for a bufferless multiplexer by means of large deviation theory, namely by maximising the Chernoff bound. The proof is extended to the bufferful multiplexer by using a transformation of a bufferful multiplexer into an equivalent bufferless one described in [4]. This transformation is very appealing since it reduces a two resource problem (buffer and bandwidth) into a single resource one. Without going into the details of that transformation we note that it hinges on a decomposition of the multiplexer into virtual trunk/buffer systems which are all allocated slices of the total bandwidth and buffer space. The buffer occupancy in each trunk/buffer system is then upper bounded by its maximal value during the periods where the buffer is not empty. This procedure is pessimistic in that the set of virtual trunk/buffer systems with bounded buffer usage can only generate more loss that the actual multiplexer. We argue that the on_off source indeed maximises the loss in this pessimistic virtual systems but since the loss in the actual multiplexer can be smaller, another arrival pattern may still generate more loss than on_off in the real system. Because of this pessimistic upper bound it is not correct use the transformation of the buferfull multiplexer into a bufferless one to prove that bufferfull and bufferless multiplexers have the same worst case arrival pattern.

Being pessimistic, the decomposition into virtual trunk/buffer system is still a valid method for connection acceptance control, even if it is not based on the actual worst case pattern and might reject more calls than necessary.

5.5.2 So when is sym really the worst case?

Our proof is done by showing the existence of counter examples in a particular range of parameters (capacity of the multiplexer equal to the sum of mean input rates, large buffer). However, this does not imply that sym patterns are only worst than on_off patterns in these situations. It is rather a consequence of the chosen upper and lower bounds having been derived and being valid in this range. Simulations show that sym patterns generate more loss in a much larger range of the parameters of the experiment.

At this point we have to remember that the on_off pattern is a special case of the sym pattern and that there are cases where it is the worst case. In Section 3.3 we have described which effects could make the on_off pattern to be the worst case. We can now risk a more mathematical explanation. From large deviation theory we know that on_off maximises loss from bufferless multiplexers by maximising the Chernoff bound [17, 4]. As shown for example in [17] this result can also be applied to a queue in discrete time. However, it is only valid if there is no time dependence in the traffic, eg. if the rate of a flow in the next time slot does not depend on the values it took in previous time slots. Leaky bucket constrained sources can display some important temporal dependences. Indeed, if a source has just expired its token bucket, the probability that it will send at more than its mean rate is zero. On the other hand, as long as there are tokens left in the bucket, the past values have no influence on the future ones. The buffer of the multiplexer acts like a memory of the past values of the rates and the key parameter is the relative size of that buffer and the leaky buckets. If the buffer is small, such that it can easily overflow without any source exhausting its token bucket, then the sources can been considered as independent in time, the chernoff bound can be applied and on_off argued to be the worst case. When the buffer is large, the amount of bursts needed to fill the buffer significantly influences the probability of more bursts occuring. In that case large deviation theory does not apply and on_off patterns are not the worst case.

6 Summary and Conclusion

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

We have simulated the case of two groups of sources with same leaky bucket parameters but different arrival patterns. The results show that the worst arrival pattern of one group depends on the traffic pattern of the other group. The same effect was also observed from a simulation of two traffic classes with different leaky bucket parameters.

Finally we have given a mathematical proof that on_off patterns are not always the worst case. We have explained why a contradictory proof reached the opposite conclusion by ignoring the effect of an upper bound applied during a transformation of the problem.

Our simulations and proof definitively invalidate the belief that on-off sources are the worst case for independent leaky-bucket constrained sources. This has an important consequence on CAC schemes. Indeed, any CAC scheme which naively assumes that sources behave like an on-off process may widely underestimate losses if the sources turn out to behave like the worst case we have identified. This does not apply to the particular CAC scheme developed in [4] and used in [15] for reasons explained in Section 5.5.1.

However, we have also explained that on_off, being a special case of the sym pattern, can still be the worst case when the buffer is small in relation to the bucket sizes. This may particularly be the case in large networks where multiplexing gain achieved by sharing a link among a large number of links reduces the amount of buffering needed to absorb bursts.

Finally we would like to conclude by pointing out that the leaky bucket is not a very efficient traffic descriptor for large networks. Its only advantage is that it makes policing very easy as well as the calculation of bounds for the deterministic case. However, the advantage of building large networks is that a multiplexing gain can be exploited when multiple flows share common links. In this situation, the leaky bucket does not describe a flow with enough detail such that one could predict the amount of resources that will be consumed by the flow. This makes measurment based connection accpetance control all the more appealing.

References

[1] R. Braden, D. Clark, and S. Shenker. Integrated services in the internet architecture: an overview. Internet RFC, June 1994.

[2] Rene L. Cruz. A calculus for network delay, part I: Network elements in isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, January 1991.

[3] B. T. Doshi. Deterministic rule based traffic descriptors for broadband ISDN: Worst case behavior and connection acceptance control. In J. Labetoulle and J. W. Roberts, editors, *Proceedings of ITC 14*, pages 591–600. Elsevier Science B. V., 1994.

[4] A. Elwalid, D. Mitra, and R. Wentworth. 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, August 1995.

[5] ATM Forum. UNI signaling 4.0. ATM Forum Technical Specification af-sig-0060.00, July 1996.

[6] P. K. Johri. Estimating cell loss rates in high-speed networks with leaky bucket controlled sources. *International journal of communication systems*, 8:303–311, 1995.

[7] R. Kositpaiboon and V. Phung. Usage parameter control and bandwidth allocation for B-ISDN/ATM variable bit rate services. In *Proceedings of IEEE Multimedia '90*, 1990.

[8] K. Kvols and S. Blaabjerg. Bounds and approximations for the periodic on/off queue with applications to ATM traffic control. In *Proceedings of Infocom '92*, pages 487–494, 1992.

[9] Jean-Yves Le Boudec. Network calculus made easy. Technical Report 96/218, EPFL/DI, Ecole Polythechnique Federale de Lausanne, December 1996.

[10] D. C. Lee. Effects of leaky bucket parameters on the average queueing delay: Worst case analysis. In *Proceedings of IEEE Infocom '94*, pages 482–489, 1994.

[11] D. Mitra and J. Morrison. Multiple time scale regulation and worst case processes for ATM network control. In *Proceedings of the 34th Conference on Decision and Control CDC*, pages 353–358, 1995.

[12] Ph. Oechslin. Worst case arrival of leaky bucket constrained sources: The myth of the on-off source. In *Proceedings of the fifth IFIP International Workshop on Quality of Service*, pages 67–77, Columbia University, New York, USA, May 1997.

[13] Philippe Oechslin. In search of the worst case arrivals of independent leaky bucket constrained sources. Research Note RN/97/69, Computer Science Department, University College of London, August 1997.

[14] H. G. Perros and K. M. Elsayed. Call admission control schemes: A review. *IEEE Communications Magazine*, 34(11):82–91, November 1996.

[15] F. Presti, Z. Zhang, J. Kurose, and D. Towsley. Source time scale and optimal buffer/bandwidth trade-off for regulated traffic in an ATM node. In *Proceedings of IEEE Infocom '97*, 1997.

[16] E. P. Rathgeb. Modeling and performance comparison of policing mechanisms for ATM networks. *IEEE Journal on Selected Areas in Communications*, 9(3):325–334, 1991.

[17] Jean Walrand and Pravin Varaiya. *High-Performance Communication Networks*. Morgan Kaufmann, 1996.

[18] Tom Worster. Modelling deterministic queues: The leaky bucket as an arrival process. In J. Labetoulle and J. W. Roberts, editors, *Proceedings of ITC 14*, volume 1, pages 581–90. Elsevier Science B.V., 1994.

[19] Naoaki Yamanaka, Youichi Sato, and Ken-Ichi Sato. Performance limitation of leaky bucket algorithm for usaga parameter control and bandwidth allocation methods. *IEICE Transactions on Communication*, E75-B(2):82–86, February 1992.