

# ON WORST CASE TRAFFIC IN ATM NETWORKS

Bahadir Erimli, John Murphy and Jennifer Murphy

## Abstract

Asynchronous Transfer Mode ( ATM ) networks, which are the emerging standards for the future broadband-ISDN networks, allows for the input traffic from users to vary both from one call to another and within the call. ATM specifies a method for controlling the traffic flow across the user network interface ( UNI ). This involves each user negotiating call parameters with the network. Once these parameters have been decided then a contract is made between the user and network. The network must then enforce the contract in order to guarantee performance and quality of service to other users. What is of interest to the network is given a particular set of users and contracts what is the worst traffic that the users could input to the network while still abiding by their contracts. This type of input traffic would be called the worst case traffic as it would produce the lowest performance in the network. The reason why this is of importance is that for simulating network performance we would like to have the worst case traffic inputs. Furthermore, it's important for the traffic controller, network, to know the possible worst case traffic so that it can assign parameters accordingly.

The contract parameters have been decided upon by the standards organisations and what is needed now is to decide what type of traffic can pass these tests and produce the lowest network performance. This problem has been studied in the literature. We give some theoretical background to explain some of the results in the literature, and we further look at some examples of types of worst case traffic sources. We show that for the two most common types, the greedy on-off and the three state source that either can be worse, depending on the situation. This would imply that there is no general worst case traffic as has been considered for the last number of years.

## 1 Introduction

The primary role of traffic control and congestion control is to protect the network and the user in order to achieve network performance objectives. An additional role is to optimise the use of network resources. The traffic control and congestion control mechanisms should not rely on other higher layer protocols which are either application or service specific. However protocols may make use of the information in the ATM layer to increase their efficiency.

There are two levels of congestion and control involved with ATM, the call level and the cell level. With ATM connections there is a unidirectional specifying of the Quality of Service

---

<sup>1</sup>Bahadir Erimli, Electrical Engineering, 116-81, California Institute of Technology, Pasadena, CA 91125, USA. email: erimli@systems.caltech.edu

<sup>2</sup>John Murphy and Jennifer Murphy, Electronic Engineering, Dublin City University, Dublin, Ireland. email: murphyj@eeeng.dcu.ie

( QOS ) parameters. These parameters are specified at connection setup and are guaranteed by the network. To guarantee the QOS the network must be able to obtain enough information from the user about the connection and be able to ensure that no other connections that share the resources degrade the QOS. A user must enter into a contract with the network about the parameters of the call. Then the network will implement traffic control to avoid problems with degraded QOS before they occur. This includes Network Resource Management ( NRM ), Call Admission Control ( CAC ), Usage Parameter Control ( UPC ), and selective cell discard. The network will also control the case where congestion does occur by implementing Explicit Forward Cell Indication ( EFCI ), selective cell discard and reaction to UPC failure.

Within the ATM cell there are a number of bits available for congestion and priority setting. These include the Payload Type Indicator ( PTI ) and the Cell Loss Priority bit ( CLP ) both of which are contained in the header of the ATM cell. The first bit of the PTI tells that the cell is a user cell and the next bit tells if congestion has been experienced by the cell in the network. The last bit differentiates between two different types of ATM-SDU's. The CLP bit is for high and low priority setting of the cells. This can be done by the user and/or by the network. A cell entering the network with low priority is subject to being discarded by the network in times of congestion.

Traffic control is necessary to protect the network so that it can achieve the required performance objectives. UPC enforces a contract between the user and the network about the nature of the call. This prevents any one user from causing excessive traffic and hence degrading the quality of service provided to the other users. It is necessary to determine what is the worst traffic a user can inflict on the network while still abiding by UPC. The Leaky Bucket Algorithm is commonly used to implement UPC.

## 1.1 Asynchronous Transfer Mode Model

A simplified ATM switch model consists of  $N$  users feeding a finite FIFO buffer (  $B$  places ) and is shown in Figure 1.




Figure 1: Simplified Model

The arrival process from the users is random, but the UPC algorithm for worst case analysis makes the arrival process to the FIFO buffer deterministic. There are two types of arrival patterns that we consider after the UPC and both of them are periodic. The service process is also deterministic ( at rate  $r$  ). The worst case traffic is that which creates the highest cell loss for a certain type of UPC, like the leaky bucket. A cell is lost when a cell arrives and the buffer is full.

## 1.2 Contract and Usage Parameter Control

The traffic contract specifies the negotiated characteristics of the connection. A connection traffic descriptor is the set of traffic parameters in the source traffic descriptor, the cell delay variation tolerance and the conformance definition. The conformance definition is used to decide which

cells are conforming in the connection. A typical conformance definition is the leaky bucket [1] or Generic Cell Rate Algorithm ( GCRA ) although many such algorithms may be used in tandem. The CAC will use the connection traffic descriptor to allocate resources and to derive parameters for the UPC. Any connection traffic descriptor must be enforceable by the UPC. Even though a cell is found to be nonconforming that does not mean that the connection is not conforming. The precise definition of a compliant connection is left to the network operator. However a connection where all the cells are conforming is compliant. The traffic contract consists of the connection traffic descriptor and a requested QOS for each direction of the connection. This includes the definition of a compliant connection. The private UNI may support a different traffic to the public UNI.

The contract must contain the Peak Cell Rate ( PCR ) of the source traffic, the cell delay variation and the cell delay variation tolerance. Sustainable cell rate and burst tolerance are optional parameters. For best effort traffic the only parameter specified is the PCR and the network may not reject the call because that bandwidth is not available but it may impose a different PCR. CAC is used to decide if a connection should be accepted or continue to be accepted in the network. It is required that the traffic contract be accessible to the CAC. The prime concern is to achieve the required QOS for the new or renegotiated connection as well as to ensure that the connection will maintain the QOS of all the other connections in the network. As well as deciding to accept the connection the CAC must determine the parameters needed by the UPC and route and allocate the resources to the connection. Even if high and low priority are not set the network may set them for nonconforming cells.




Figure 2: Leaky Bucket or UPC Algorithm

UPC is the set of actions the network take to monitor and control traffic. This includes the validity of the connection. The operation of the UPC shall not violate the QOS objectives of a compliant connection. However the excessive policing actions on a compliant connection are part of the overall network performance degradation and so safety margins should be engineered to limit the effect of the UPC. The UPC can also fail to take action on a non compliant connection. Policing actions on the non conforming cells are not to be allocated to the network performance degradation of the UPC. At the cell level the UPC may pass a cell, change the priority of the cell or discard the cell. A low priority cell is discarded by the UPC if it is non conforming. Following the UPC shaping may be implemented on the conforming cells to reduce cell clumping. It is optional for the network operator to allow the UPC to initiate the release of a non compliant connection. When two levels of priority are used the UPC may discard high priority cells even though if the UPC were performed on the high priority alone the cells would be conforming.

The UPC and CAC are operator specific and should take into account the traffic contract to operate efficiently. It is specified that the signaling should take into account experimental traffic parameters that could be proprietary to either the manufacturer or network operator. It is optional to allow the the operation of these parameters across the UNI by mutual agreement. It is optional for the user to be allowed to mark cells as low and high priority. It is also optional

for the network to mark cells as low priority if they are not adhering to the traffic contract. The cell loss ratio for low priority cells must be higher than for high priority.

The leaky bucket algorithm is a UPC standardised by the ATM Forum [2] and is shown in Figure 2. The operation of the leaky bucket is that a splash is added to the bucket (counter increment) for each incoming cell when the bucket is not full. When the bucket is full cells cannot pass through to the network unmarked but the bucket leaks away at a constant rate. The important parameters to be defined in this system are the leak rate of the bucket ( $R$ ), the bucket capacity ( $M$ ) and the peak cell emission ( $p$ ). We assume that if the cells cannot pass the leaky bucket unmarked then they are lost and are not taken into account in the cell loss rate. This is because the network will only give guarantees to the marked cells and the unmarked cells will not interfere with the marked ones.

The UPC standard of the ATM Forum [2] is actually a double leaky-bucket that controls the two different cell rates  $R$  and  $p$ , each with a separate bucket. The peak rate,  $p$ , is controlled by a bucket of size 1 and leak rate,  $p$ , and the second bucket operates as explained above. For a cell to be a conforming cell, it needs to be conforming with both buckets at the time it's transmitted.

### 1.3 Traffic Types

When considering which type of sources will produce the worst performance in the network while still maintaining the contract the first point to note is that the type of source will not allow the leaky bucket to ever overflow. In other words the total available number of cells that are allowed to enter the network will enter to produce the lowest performance. Two types of sources that have been proposed that could be the worst types are a two state source and a three state source as shown below in Figure 3.




Figure 3: Source Types

The two state source (greedy on-off) emits a burst of cells at the peak cell emission until the bucket is about to overflow and then falls silent waiting for the bucket to empty. This occurs periodically depending on the parameters of the system. The three state source is similar to the greedy on-off source except that it keeps emitting cells at the leak rate after a burst. Therefore its operation is to emit a burst of cells at the peak cell emission rate until the bucket is about to overflow and then emit cells at the leak rate of the leaky bucket for some time and then fall silent to allow the bucket to empty. This would then be repeated again. What can be seen is that the three state source would have a longer period than the greedy on-off for the same system parameters.

The general belief is that the greedy on-off source gives rise to the worst case traffic as it would have the largest variance possible. However the three state source has been proposed as producing longer queues and therefore larger loss.

### 1.4 Finite and Infinite Buffers

A study of the three state source [5] compared its performance to that of the greedy on-off source, for a number of sources into an infinite buffer. What was found in those simulations is that the three state sources produced higher buffer occupancy than the greedy on-off sources. The survivor function,  $P[Q > q]$ , was found where  $Q$  was the buffer occupancy. This function was then assumed to approximate to cell loss in a finite buffer. However this we believe would only be true for input traffic that would be statistically independent of the queue lengths, which is not the case here as the input traffic is periodic and deterministic. We can demonstrate this fairly simply with a small example as is shown in Figure 4. Here we have plotted the buffer occupancy for the infinite case of both the greedy on-off and the three state on the same axis. Note that we are considering only a single source here on its own.




Figure 4: Buffer Occupancy in an Infinite Buffer

We have chosen the period  $T$  of the greedy on-off to be half of the period of the three state source so that we can just examine a single period of the three state for cell loss. There are 2 cells arriving per cell time while the source is emitting at the peak rate  $p$  and this continues until the bucket is full to capacity  $M$ , which in our example is 4. Then the bucket leaks away at a rate of 1 cell per cell time  $R$  for the greedy on-off but for the three state source it emits at this rate  $R$  for 6 time units, which is denoted by  $T^*$  and here is equal to the period of the greedy on-off source. Therefore we look over 12 time units as both repeat exactly the same after this. Over this period there are 12 cells inputted to the buffer which is the same as the leak rate  $R$  multiplied by the time period. The three state source gives rise to higher buffer occupancy than the two state source as can be seen in the infinite buffer plot. This would lead us to believe because of the higher queue occupancy for the infinite case that we would have higher cell loss in the finite case. However if we now consider a finite buffer of capacity 2, we find that that the three state source losses 2 cells, one at time period 1 and one at time period 2. However the greedy on-off source losses one at time period 1 and one at time period 2 and again one at time period 7 and one at time period 8. So the cell loss rate for the three state source is 0.166 while the cell loss rate for the greedy on-off is 0.333.

The fact of higher buffer occupancy in the infinite case does not necessarily mean higher cell loss rates in the finite buffer when the arrival process is deterministic. For deterministic arrivals the loss in a finite queue is not necessarily related to the buffer occupancy in the infinite case but more on the method of arrivals, or the process of arrivals, past the finite places in the queue. This problem was simulated [4] and these conclusions help to explain the results obtained therein. Furthermore, the survivor function  $P[Q > q]$  can provide more conflicting results if viewed more

carefully. As an example, let's look at a situation with  $M = 2$ ,  $p = 1$ ,  $R = 0.5$  and  $r = 1$ . For the three-state source, let's leave  $T^*$  as a variable and the overall period of the source be  $T$  ( $T = 6$  for the two-state case). Then we can calculate the exact frequency that each possible queue occupancy is going to occur at.

- $P[Q = 0] = (4T - 7)/T^2$
- $P[Q = 1] = (.75T^2 - 4.5T + 12)/T^2$
- $P[Q = 2] = (.25T^2 - 3)/T^2$
- $P[Q = 3] = (.5T - 2)/T^2$

Using these formulas, we can easily calculate the survival function and compare it for different values of  $T^*$ . For the probability  $P[Q > 2]$ , starting from  $T^* = 6$ ,  $P[Q > 2]$  for the three-state sources is less than that of the two-state source as seen in Figure 5. On the otherhand, the average queue occupancy of all the three-state sources in this case are greater than that of the two-state source, but the cell loss rate of the two-state source is greater than any of the same three-state sources for a buffer length of 2 cells ( $B = 2$ ).

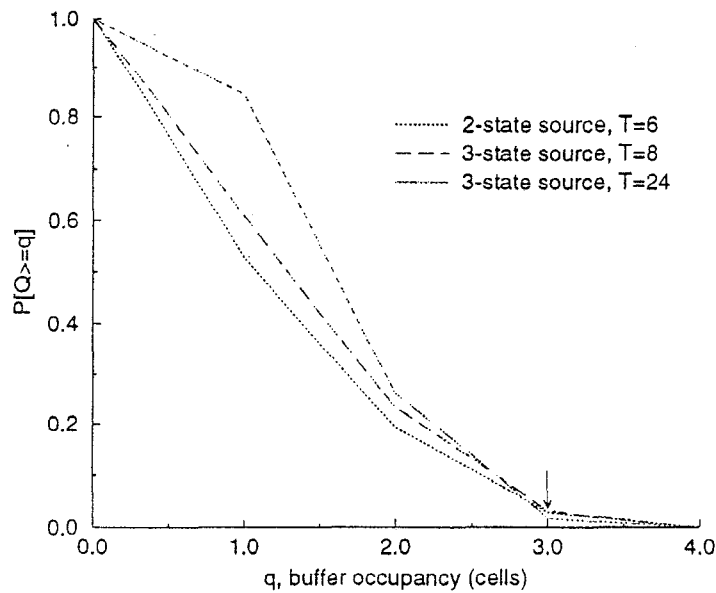


Figure 5: Comparison of Two State and Three State Sources

This shows that the survival function for an infinite buffer cannot be used to make conclusions on the finite buffer situation for this deterministic traffic pattern. In fact, if the tail-end of the simulation in [5] is calculated, it can be seen that  $P[Q > 898]$  for the two-state is higher than that of the three-state source, indicating that there is a crossover point in the graph that was unlikely enough not to appear during the simulation process.

### 1.5 Continuous and Discrete Variables

The problem contains both continuous and discrete variables and therefore it is important to distinguish between the two types. We consider a simplified model similar to [3] where there are two identical, independent users feeding a finite buffer. If the users were not independent

then the worst case would be the greedy on-off source with all sources in phase emitting together. However when we consider that the users or sources are independent then the phase between the sources is random. This then gives rise to the probability of cell loss because there is a probability of phase difference between the sources. We assume that the two sources are randomly phased, which means we can consider one source as a reference source and the other is than an amount out of phase. The amount out of phase will determine the number of cells lost. The more in phase the more loss we expect. The probability of each of the possible combinations of out of phase is just the reciprocal of the number of the possible combinations.

We consider two possibilities, either two greedy on-off sources or two three state type sources. The cell losses for all distinct combinations of the two traffic patterns are calculated and averaged. The cell loss rate for each scenario is compared. There are a number of constraints on the problem due to the discrete nature of the variables.

1. The sources may be out of phase only in one cell time units which means that we cannot use integration as was used in [3].
2. The on period of the two state source is the time taken for the buffer to just about to overflow whilst emitting at the peak rate and also to emit any cells which arrive during this time. In the discrete case this is equal to  $\lceil M/p - R \rceil - 1$ .
3. Cells are only lost as units - no fractional cell loss. Therefore it is necessary to ensure that in using general formulae the cell loss is truncated to an integer. We investigate the case of  $M \geq 1$ .
4. The discrete nature of the cell also implies that the buffer size must be an integer.
5. The service process can be assumed to be continuous. This means a cell can be served as it arrives. Alternatively, can think of the server waiting until the cell has fully arrived before starting to serve it.

## 2 New Approach to Cell Loss

The comparison of cell loss from both source types in the two identical user system is revisited assuming the cell arrival and service processes to be discrete. The cell loss rate is computed as the total number of cells lost in one period divided by the total number of cells emitted in one period. The total number of cells lost in any one period is the average of the cells lost by each combination of the two traffic patterns. It is assumed that the minimal phase difference between two sources is one cell time and there are  $T$  ( $T$  is the period) combinations of the two traffic patterns. The total number of cells emitted in one period  $T$  equals  $2RT$ . Furthermore the following stipulations are placed on the parameters :

- The service rate is at least equal to the peak cell emission rate,  $p \leq r$
- The leaky bucket rate is less than or equal to half the peak cell rate ( and hence the service rate ),  $R \leq p/2$
- Together the peak cell emission and the leaky bucket rate exceed the service rate,  $p + R \geq r$
- The buffer size must be small enough so that cell loss is guaranteed when both sources are in phase,  $(2p - r)T_{on} \geq B$

These stipulations ensure that cell loss only occurs when at least one of the sources is emitting at the peak rate. If the service process is assumed to begin as the cell is arriving we call this

a fast server. Then the cell loss for two greedy on-off sources,  $x$  places out of phase with each other, is denoted by  $CL(x)$  and is calculated in Equation 1

$$CL(x) = [(2p - r)(T_{on} - x) - B] \quad (1)$$

where the symbols have the usual meanings. Similarly the cell loss for the three state source can be represented by Equation 2.

$$CL(x) = [(2p - r)(T_{on} - x) + (p + R - r)(x) - B] \quad (2)$$

Alternatively if it is assumed that the server waits until the cell has arrived in the buffer before starting to serve the cell, which we call the slow server, then the expressions for cell loss are modified for the greedy on-off source as follows in Equation 3

$$CL(x) = [(2p)(T_{on} - x) - (T_{on} - x - 1/p)r - B] \quad (3)$$

and for the three state slow server the cell loss will be given in Equation 4.

$$CL(x) = [(2p)(T_{on} - x) - (T_{on} - x - 1/p)r + (p + R - r)(x) - B] \quad (4)$$

To calculate the cell loss ratio we allow  $x$  to vary over all phase possibilities and then divide by the number of combinations and also divide by the number of cells transmitted by both sources. The number of cells transmitted by both sources will be  $2RT$  and the number of phase combinations will be  $T$ , so the cell loss ratio,  $CLR$ , is given by Equation 5.

$$CLR = \frac{\sum_{x=0}^{T-1} CL(x)}{T \cdot 2RT} \quad (5)$$

### 3 Counter Examples

To show that there is no single worst type traffic for two identical sources we present a counter example to the traditional theory of the greedy on-off being the worst case. We show that the three state source can produce higher loss for integer values of variables chosen. As mentioned previously we consider two different types of servers, the fast serving server and the slow serving server.

#### 3.1 Fast Serving Server

Assuming that the cell is served as it arrives we have an example showing the three state source to give rise to greater cell loss in a finite buffer than the two state source. The following system parameters are used :

- $p$ , peak cell emission = 1
- $r$ , service time = 1
- $R$ , leaky bucket rate = .5
- $M$ , leaky bucket capacity = 8
- $B$ , buffer size = 12
- $T^*$  (for three state source)  $\geq ((2p - r)T_{on} - B)/(p - R) = 6$




Figure 6: Fast Server, (A) Greedy On Off Source, (B) Three State Source

It takes 8 cell times to fill the bucket up but in that time 4 more cells are allowed through because of the constant leak rate of the bucket and while they try to fill the bucket 2 more arrive and then finally one arrives and the bucket is full. Therefore the amount of time that the source is on and emitting cells at the peak rate is given by  $T_{on} = \lceil M/p - R \rceil - 1 = 15$ . When the bucket is full it takes  $M/R$  seconds to empty normally, however here one time period has already elapsed so  $T_{off} = 15$ .

There are 30 phase combinations for the greedy on-off source patterns, however sources that are too far out of phase do not give rise to cell loss. For the greedy on-off sources by examining Equation 1 we conclude that sources that are 3 or more time units out of phase do not produce any cell loss. Remember that this loss can occur when the second source is a little advanced from the reference and also when it's so advanced that it is almost back in phase with the reference. This can be seen in Figure 6. The cell loss for the two state sources is in total 9 cells, which is calculated from 3 cells lost when in phase, 2 cells lost when either one out of phase and also 29 out of phase, and 1 cell lost when either two out of phase or 28 out of phase. For

all other phasings there is no cell loss. Therefore using Equation 5 the cell loss ratio can be calculated to be 0.01.

For the three state sources cell there are 36 possible phase combinations and by examining Equation 2 loss can occur up to 4 units out of phase. The total number of cell lost over all possible phasings can be seen in Figure 6 and is 15 cells. By using Equation 5 we can calculate the loss to be 0.01157. Therefore the three state source produces higher loss than the greedy on-off source for the fast server.

### 3.2 Slow Serving Server




Figure 7: Slow Server, (A) Greedy On Off Source, (B) Three State Source

If the service process is assumed to begin serving a cell only after it has fully arrived into the buffer we have a similar counter example. The following system parameters are used :

- $p$ , peak cell emission = 1
- $r$ , service time = 1
- $R$ , leaky bucket rate = .5
- $M$ , leaky bucket capacity = 8

- $B$ , buffer size =14
- $T^*$  (for three state source)  $\geq ((2p - r)T_{on} - B + r/p)/(p - R) = 4$

Similar to the fast serving server  $T_{on} = T_{off} = 15$ . For the greedy on-off sources by examining Equation 3 we conclude that sources that are 2 or more time units out of phase do not produce any cell loss. This can be seen in Figure 7. The cell loss for the two state sources is over all possible phase combinations equal to 4 cells and by using Equation 5 we can calculate the cell loss ratio to be 0.00444.

For the three state sources cell loss can occur up to 2 time units out of phase. This is concluded by examining Equation 4 for the cell loss for a three state source and this is also seen in Figure 7. Here in total 6 cells are lost over all possible phase combinations and so again by using Equation 5 the cell loss ratio can be calculated to be 0.00519. Therefore the three state source produces higher loss than the greedy on-off source for the slow serving server.

Therefore regardless of how the service is achieved in the buffer there is an example of the three state source producing more loss than the greedy on-off source. It is therefore shown that there is no single type of worst case traffic source for the two identical source problem, considering these types of sources.

## 4 Conclusions

From our analysis it may be concluded that the greedy on-off source does not always give rise to the worst case traffic. The first point that we make is regarding the infinite versus finite buffer and we show that the occupancy in an infinite buffer will not relate to the loss in a finite buffer when the traffic is deterministic. The second point we make is that some of the variables are discrete and therefore attention must be paid when selecting examples and formula. We then concluded by showing two counter examples against the greedy on-off source and were able to say that there is no single worst case traffic type for the two identical source problem.

The problem is far from totally solved, though. The result for this two source case is not binding on an  $n$  source situation, indicating one area to work towards. On the other hand, the traffic controller still needs to know in some manner what the worst case, or even a supposed or an asymptotic worst case, should be to be able to make decisions for QOS. This brings up the questions of under what conditions is one case worse than others or if there is a single case that we don't know about that performs universally badly for a given traffic controller.

## Acknowledgments

We would like to thank Dr. B. T. Doshi for assistance with this problem. We are also grateful to Dr. Mike Mandell for helpful discussions. The first author acknowledges Prof. R.J. McEliece, California Institute of Technology, and the second and third authors acknowledge Prof. Charles McCorkell, Dublin City University, for their continuing encouragement and support of this work.

## References

- [1] S. Akhtar, "'Congestion Control In A Fast Packet Switching Network', Master's thesis, Washington University, 1987.

- [2] *ATM Forum*, 'ATM User--Network Interface Specification', Prentice Hall, 1993.
- [3] B. T. Doshi, 'Deterministic Rule Based Traffic Descriptors For Broadband ISDN: Worst Case Behaviour And Connection Acceptance Control', International Teletraffic Congress, ITC-14, Vol 1a, pp 591-600, Antibes, France, June 1994.
- [4] T. Worster, 'Modelling Deterministic Queues: The Leaky Bucket as an Arrival Process', International Teletraffic Congress, ITC-14, Vol 1a, pp 581-590, Antibes, France, June 1994.
- [5] N. Yamanaka, Y. Sato and K. Sato, 'Performance Limitation of Leaky Bucket Algorithm for Usage Parameter Control and Bandwidth Allocation Methods', IEICE Trans. Commun., Vol E75-B, No. 2, February 1992.