Experimental Investigation of CAC and Effective Bandwidth for Video and Data

EGIL AARSTAD
Telenor Research & Development, Kjeller, Norway

SØREN BLAABJERG
Technical University of Denmark, Lyngby, Denmark

FERNANDO CERDAN
Polytechnic University of Catalonia, Barcelona, Spain

SANDRA PEETERS, KATHLEEN SPAEY
University of Antwerp, Antwerp, Belgium

Abstract. With the availability of commercial ATM switches implemented with large buffers and strict priorities between different traffic classes, it becomes possible to study ATM traffic control methods as well as integration strategies in an experimental environment where both real time and non-real-time traffic are present. In this paper, the connection admission control (CAC) issue is investigated in the Expert ATM Testbed with real switches and as realistic traffic sources as possible. Two types of sources have been used, MPEG video sources which are artificially generated based on real MPEG traces, and data sources which are modelled as traditional on/off sources. The video traffic represents real-time traffic and is given priority over the data traffic. To complement and verify the hardware experiments, two simulation tools have been developed. The first tool evaluates some shortcommings in the hardware generation of the artificial MPEG sources while the second tool uses the real MPEG traces and thereby constitutes an evaluation of the way the artificial model is made. Results from both simulation tools are compared with the results of the hardware experiments. Since the last few years have shown solid progress in the development of a concept for effective bandwidth which reduces the complexity of the CAC problem to the same level as for circuit switched networks, this paper also compares effective bandwidth results obtained by three methods from the literature with the results obtained from the experiments.

1 Introduction

One of the most important benefits of ATM is the integration of different services into one multi-service network. In order to realise this objective, both ATM Forum and ITU-T have specified different service categories, tailored to the requirements of various present and future services and applications: Constant Bit Rate (CBR), real-time Variable Bit Rate (rt-VBR), non-real-time Variable Bit Rate (nrt-VBR), Available Bit Rate (ABR) and Unspecified Bit Rate (UBR). In ITU-T terminology a slight modification of these are denoted ATM Transfer Capabilities.

An ATM-based network must be able to support services with stringent Quality of Service (QoS) requirements (CBR and VBR), as well as best effort services like UBR and ABR. This requires an integrated control framework which on one hand can give performance guarantees to real-time services and on the other hand ensure high network utilisation.

An important part of ATM traffic control is handled by the Connection Admission Control (CAC) function. CAC determines whether a new connection can be accepted or not, depending on the availability of the necessary resources. Since the traffic types differ not only in their peak and mean bandwidth requirements but also in their burstiness, the CAC function has to be designed very carefully to achieve efficient utilisation of network resources. To realise this it is necessary to have a good understanding of statistical multiplexing of the traffic from different service categories. This understanding enables us to make correct bandwidth allocation decisions in order to avoid wasting bandwidth.

The concept of *effective bandwidth* has over the last few years had an increasing influence on how it is believed that CAC could be done. If it is feasible to assign to each connection on a link an effective bandwidth describing the load the connection puts on the link then the difficult determination of CAC can be moved from the time critical events of call arrivals to an off line investigation of assigning the correct effective bandwidth in the multiplexing scenarios at hand.

With the increasing interest in multimedia and the tremendous growth in the internet traffic with word wide web (WWW) it is expected that video as well as data traffic will constitute a significant part of the future broadband traffic. A better insight of statistical multiplexing of video and data therefore constitutes a step in a more solid understanding of statistical multiplexing in ATM in general. This in turn will form a basis for providing a powerful traffic integration framework.

Most literature on video and data so far considers statistical multiplexing (by simulation and analytically) of video and data separately, see e.g. [1], [2], [3], [4] and [5]. With the availability of ATM switches with delay priorities it is now possible to study this integration framework issue in an experimental setting.

In this paper we investigate the CAC issue with and without delay priorities for integrating real-time and non-real-time VBR services. We present results from experiments performed within the framework of the ACTS (Advanced Communications Technologies and Services) EXPERT project at the EXPERT Testbed in Basle, Switzerland. Because video services are expected to become an important broadband service we have chosen MPEG (Moving Pictures Expert Group) sources as a representative of the rt-VBR service category. In the experiments we have used artificial MPEG sources, derived from real MPEG traces, which enable us to investigate traffic scenarios with more sources than would have been practically possible using only real MPEG sources and which make our results repeatable. The non-real-time data sources are modelled as traditional on/off sources. The difficult issue of how to determine the policing parameters of MPEG or data traffic and the issue of whether a control loop between the source and the policing device should be implemented is not dealt with in this paper.

Section 2.1 presents the artificial MPEG traffic model that is implemented in the traffic generators as well as the on/off model. In section 2.2.1 a description of the configuration is given. The obtained CAC boundaries can be found in section 2.2.2. Additional to the hardware experiments a simulator has been developed. The results of these simulations and a comparison with the experimental results are given in section 2.2.3.

Section 3 presents a trace driven simulation tool by which multipexing of the real MPEG traces, which the artificial models are based on, with on/off traffic has been carried out and compared to the hardware results of section 2.

In section 4, the effective bandwidth obtained from the measurements is compared with the effective bandwidth computed from three different methods presented in the literature.

2 Experimental CAC boundaries

2.1 Modelling video and data

In order to reduce the required bandwidth to acceptable levels, coding algorithms for the compression of video data streams are needed and several coding schemes have been developed. The last few years have shown MPEG as the most promising scheme. It comes in two versions, MPEG-I (see [6], [7]) and MPEG-II (see [8]) where the MPEG-I functionalities are a subset of those in MPEG-II where e.g. layered coding is possible. In this paper it is MPEG-I which is dealt with.

MPEG encoding of a video sequence is realised using three different compression levels. Thereby three different types of frames namely I-, P- and B-frames are generated. A periodic generation of these frame types, usually of length 12 frames in the pattern IBBPBBPBBPBB, is the result. I-frames have the largest amount of data (lowest compression level) while B-frames have the smallest amount of data (largest compression level).

From a traffic modelling point of view MPEG has traditionally either been modelled as Markovian [3], autoregressive [9], or more recently fractal [10]. Since the traffic generators available at the EXPERT platform implement a Markovian model at the burst level, the starting point for our model is also Markovian. We have chosen to follow the approach of B. Helvik in [11] where the following properties are taken into account:

- periodic generation of the three frame types (in the sequence IBBPBBPBBPBB which constitutes one Group of Picture, GoP)
- correlation between Group of Pictures (GoP)
- frame size distribution for I-frames and distribution for the sum of B- and P-frames

No intention to capture long range correlations is made, mainly due to our multiplexing scenarios where only short or moderate sized buffers are used. The periodic generation is compromised in the sense that the duration in individual states is assumed exponential instead of constant (a requirement if more than one MPEG source on each module is to be generated). The sum for B- and P-frames is approximated by using four or five load levels and the I-frame size distribution is further refined by defining two sublevels within each load level. The duration of each frame is 45 msec. For further details see [11].

In the actual implementation four or five load levels and two I-frame sizes in each load level are used implying a model with either 52 or 65 states. Parameters in the model are the transition probabilities between the 5 levels as well as the I, B and P rate at each load level. A method developed by B. Helvik (and quite similar to the simple Markov method presented in the video modelling section in part I of [4]) to generate these parameters from an arbitrary MPEG video trace has been used.

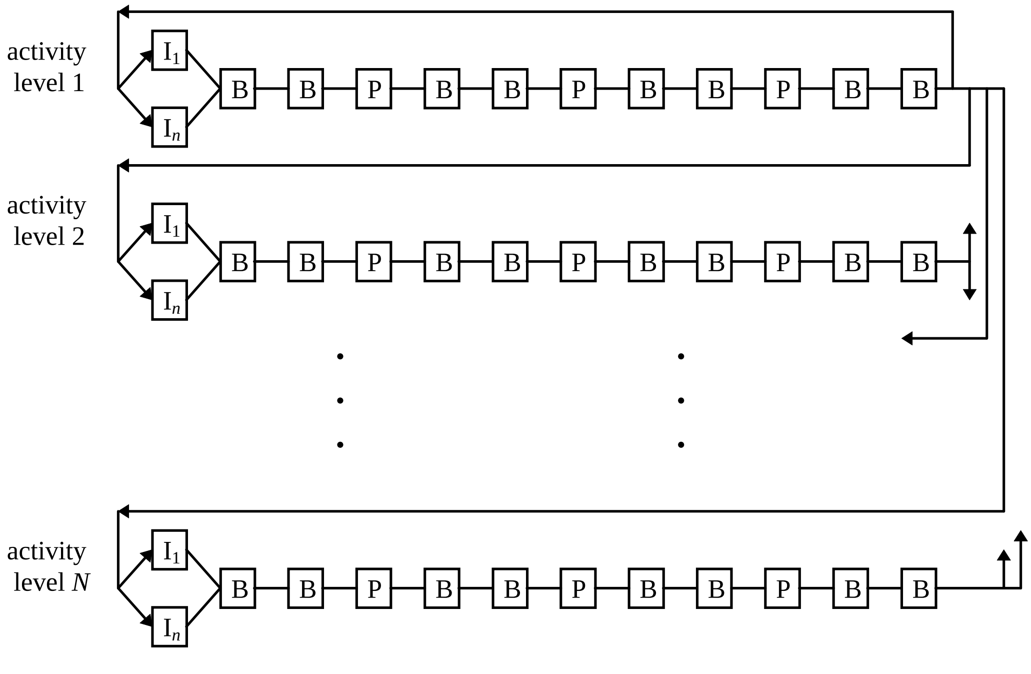

Figure 1: Structure of the artificial Markovian MPEG source. Each square represents a state in which the duration is exponentially distributed with mean 45 msec.

The two MPEG sources used are based on a 24 minutes trace from the Bond movie "Goldfinger" and a 24 minutes long trace from an Asterix cartoon which have been made available by O. Rose at the Institute of Computer Science at the University of Wuerzburg. The source based on the Bond movie which is implemented as a 65 state model has a mean rate of 0.59 Mbps while the Asterix cartoon has a mean rate of 0.51 Mbps and is implemented as a 52 state model. Thereby we are able to multiplex both MPEG sources in one generator without exceeding the upper limits for what the equipment can handle.

The on/off data source used in the experiments, has a peak rate 7.78 Mbps, a mean rate 0.389 Mbps, a mean on duration of 10 msec and a mean off duration of 190 ms. Thereby this source represents a bursty source with a peak rate that is big compared to the capacity of the multiplexing link. The duration in both on and off states is exponentially distributed and mutually independent (traditional 2-state Markov source).

2.2 Experimentally derived boundaries

2.2.1 Experimental configuration

The experimental results were obtained in the EXPERT Testbed which contains several ATM switches and a range of end-systems together with test and measurement equipment. Figure 2 shows the experimental configuration involving a Cisco LS1010 ATM switch and a test instrument called ATM-100 which gives the possibility to generate and analyse quite general random traffic sent through the ATM switch. The ATM-100 is equipped with two Synthesised Traffic Generators (STGs) [12] that are used for generating the artificial real-time MPEG traffic and the non-real-time on/off traffic. The multiplexing of the traffic sources takes place in two stages. First the MPEG sources and the on/off sources are multiplexed separately inside the two traffic generators. Then the aggregated real-time and the non-real-time traffic streams are multiplexed on an output port of the Cisco switch. Due to hardware constraints in the traffic generators a pacing rate function has been used to limit the output port capacity to 37.44 Mbps, thereby reducing the number of sources required to adequately load the system. The resulting aggregated traffic stream is analysed in the ATM-100, permitting cell loss and delay measurements.

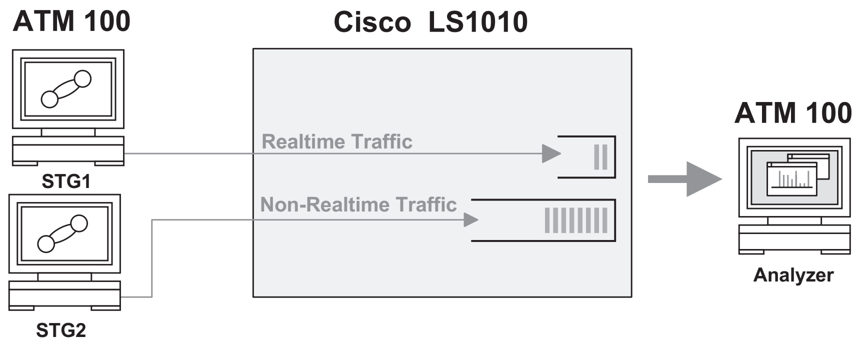

Figure 2: The experimental configuration

The buffer architecture of the Cisco switch implements delay priorities to protect the real-time traffic from influence of the non-real-time traffic. A short buffer has been used for the real-time traffic. For the non-real-time traffic, which can tolerate longer delays, a longer buffer has been used. The service discipline is such that the real-time traffic is served as long as there are cells in the short buffer. Only when the short buffer is empty, the non-real-time traffic will be served.

2.2.2 Measurement results

Three sets of experiments have been carried out. In the first two sets of experiments a number of MPEG sources of each type (Bond and Asterix) and on/off sources have been multiplexed in a common buffer. In the first set we use a large buffer of 1260 cells, in the second set a small buffer of 256 cells. In the third set of experiments the same experiments have been carried out but this time the MPEG traffic has been multiplexed in a small high priority buffer (256 cells) while the on/off traffic is multiplexed in a larger low priority buffer (1260 cells). The number of sources has been varied to obtain a cell loss ratio (CLR) below, but as close to 10^{-4} as possible. By changing the traffic mix it has been possible to obtain the CAC admission boundary valid for each of the three mixes MPEG-Bond + On/Off, MPEG-Asterix + On/Off and MPEG-Bond + MPEG-Asterix.

Figures 3 and 4 give the 2-dimensional CAC boundaries for MPEG-Bond + On/Off and MPEG-Asterix + On/Off and this for the three different cases: non-priority with small buffer and large buffer and the delay priority case.

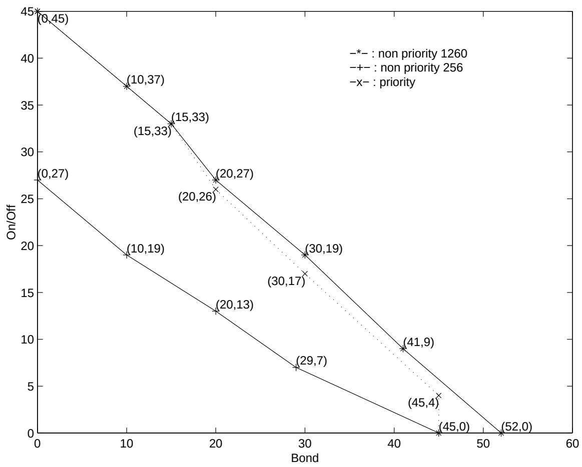

Figure 3: CAC boundaries for Bond + On/Off

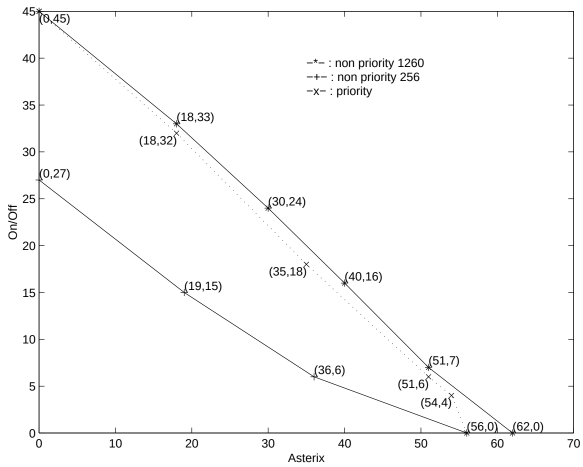

Figure 4: CAC boundaries for Asterix + On/Off

Experiments with a mix of the three different sources, have also been performed. From these measurements some points on the plane (the CAC boundary of MPEG-Bond + MPEG-Asterix + On/Off) have been derived. Figure 5 gives the three-dimensional admission boundary for the non-priority case, this time a small buffer of 256 cells was used. The three-dimensional CAC boundary when delay priorities are used, can be found in figure 6.

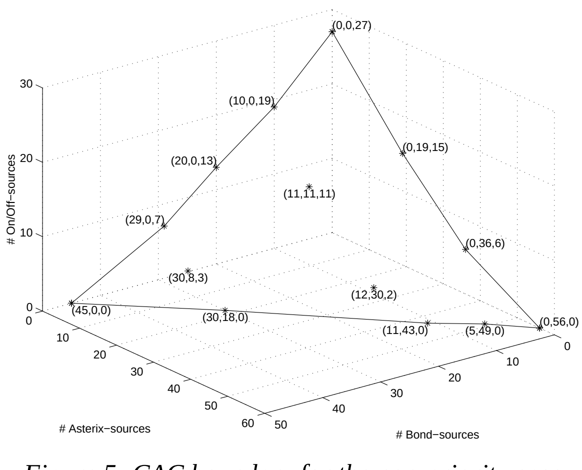

Figure 5: CAC boundary for the non-priority case with buffer = 256 cells

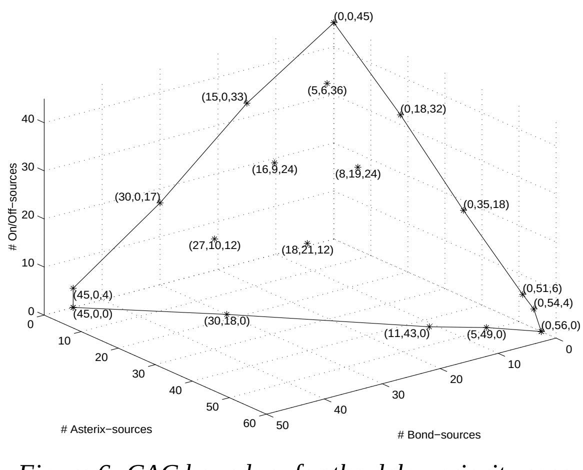

Figure 6: CAC boundary for the delay priority case

In the non-priority case the CAC boundaries (see figures 3 and 4) are close to linear. The small deviations from linear regions in the figures are due to general uncertainties which exist in any measurement. This linearity supports the feasibility of the concept of effective bandwidth [13].

As expected we can allow more MPEG sources in the delay priority case than in the non delay priority small buffer case. Keeping the number of e.g. on/off sources at six we can allow 51 Asterix sources, instead of 36.

In the priority case almost the same admission boundaries as in the non-priority case with large buffer come out, except that the number of MPEG sources that can be supported is less. 45 Bond source instead of 52 and 56 Asterix sources instead of 62 can be supported. The reason is that in the delay priority case the multiplexing of MPEG takes place in a buffer of size 256 while in the other (and non priority case) the size of the buffer is approximately 5 times bigger and when no or only a small amount of on/off traffic is present this large pool can be used by the MPEG traffic but at the expense of larger delay and delay variation.

2.2.3 Simulation results

Some aspects related with the STG traffic generator have to be taken into account before comparing experimental and simulation results. Both STGs have some hardware limitations which can make the implemented Markov models slightly different from the exact Markov models. Depending on the number of sources generated, these slight differences may become important.

The first limitation is that the STG can only provide transition probability values in integer multiples of 1/256. The second limitation is that the peak rate must divide the link rate such that the interarrival time between cells in a given state is always the same integer number of slots. The simulation tool presented in this section can provide transition probabilities with any accuracy as well as any peak bit rate value. The above two limitations give rise to the following differences with the MPEG sources (with a 95% confidence level measured on 30 independent sources) :

	STG		SIMULATOR	
	Bond	Asterix	Bond	Asterix
Mean	0.59221 Mbps	0.51318 Mbps	0.59764 ± 0.003 Mbps	0.54943 ± 0.002 Mbps

Table 1: Comparison of the mean bit rate of MPEG sources from the ATM100 and from the simulator

As can be seen in table 1, the exact model for Asterix sources generates 0.03625 Mbps more than the implemented model i.e. for a certain experimental point the CLR obtained by simulation can be worse depending on the number of Asterix sources used, since in a congestion situation small increases in load can lead to big differences in CLR. Nevertheless, in those points where Bond sources have more influence, negligible differences between experimental and simulated CLR should be seen.

2.2.3.1 Simulator description

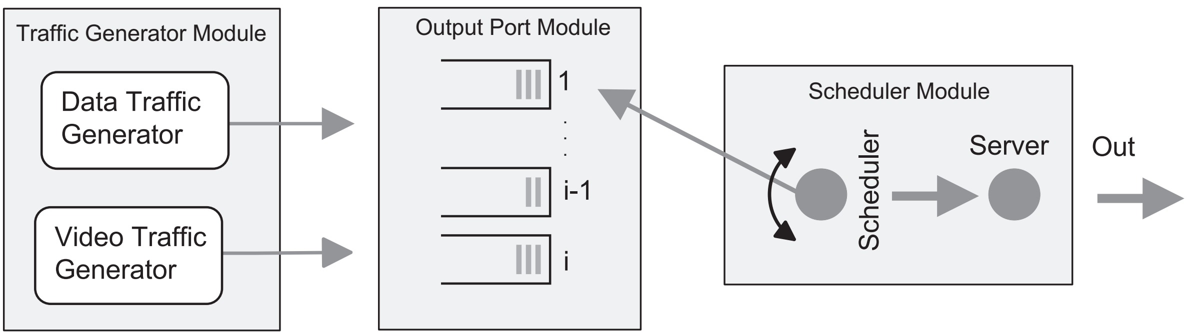

Figure 7: Simulation model

The simulation model of figure 7 has a high level of detail since it is intended to behave like the real equipment in the Testbed in order to give a fair comparison with the experimental results and predict the performance of new alternatives. It consists of three different modules: the first is the traffic generator module which implements the exact Markov models for the MPEG video sources as well as the data traffic as has been explained in section 2.1.

The second module corresponds to the output port of the switch where the traffic is multiplexed and consists of as many queues as there are service classes of traffic defined. We usually assign a different priority for each service class although this depends on the scheduling mechanism used in the third and last module. The scheduler module chooses one head of line cell from the queues in the second module according to the scheduling algorithm and delivers it to the link server.

2.2.3.2 Confidence interval for the mean

Due to the exponential distributions with long mean sojourn times in each state, we have chosen to calculate the sample mean from a sample of a few very long observations of the system in order to guarantee enough number of events such that the simulation trajectory is representative of the system behaviour in a steady state.

The small confidence intervals obtained for all measurements performed show that the correct length of the simulation has been chosen for most of the cases, as it can be seen in tables 2, 3 and 4 below. For all cases we have estimated the confidence interval with a 90 % confidence level assuming the observations come from a Student t-distribution.

Bond	Asterix	On/Off	Experimental CLR	Simulated CLR
20	0	13	$0.5329\ 10^{-4}$	$0.7714\ 10^{-4} \pm\ 0.0510\ 10^{-4}$
30	18	0	$0.6477\ 10^{-4}$	$0.7412\ 10^{-4} \pm\ 0.1608\ 10^{-4}$
45	0	0	$0.5001\ 10^{-4}$	$0.5747\ 10^{-4} \pm\ 0.1058\ 10^{-4}$
0	56	0	$0.4420\ 10^{-4}$	$2.9964\ 10^{-4} \pm\ 0.4807\ 10^{-4}$

Table 2: Non-priority case with 256 cells buffer size

Bond	Asterix	On/Off	Experimental CLR	Simulated CLR
10	0	36	$0.8707\ 10^{-4}$	$0.5987\ 10^{-4} \pm\ 0.2427\ 10^{-4}$
41	0	9	$0.4742\ 10^{-4}$	$0.5331\ 10^{-4} \pm\ 0.1373\ 10^{-4}$
52	0	0	$1.800\ 10^{-4}$	$0.8954\ 10^{-4} \pm\ 0.0913\ 10^{-4}$
0	62	0	$0.2915\ 10^{-4}$	$12.999\ 10^{-4} \pm\ 2.5547\ 10^{-4}$

Table 3: Non-priority case 1260 cells buffer size

Number of Sources			Experimental CLR		Simulated CLR	
Bond	Asterix	On/Off	Bond-Asterix	On/Off	Bond-Asterix	On/Off
18	21	12	0.0	$0.9249\ 10^{-4}$	0.0	$1.75\ 10^{-4} \pm 0.58\ 10^{-4}$
5	6	36	0.0	$0.8935\ 10^{-4}$	0.0	$0.89\ 10^{-4} \pm 0.22\ 10^{-4}$
0	54	4	$0.7869\ 10^{-4}$	$0.5773\ 10^{-4}$	$1.13\ 10^{-4} \pm 0.07\ 10^{-4}$	$13.7\ 10^{-4} \pm 1.16\ 10^{-4}$
45	0	4	$0.5189\ 10^{-4}$	$0.3612\ 10^{-4}$	$0.53\ 10^{-4} \pm 0.06\ 10^{-4}$	$1.20\ 10^{-4} \pm 0.67\ 10^{-4}$

Table 4: Priority case 256/1260 cells buffer size

The agreement between experiments and simulations are satisfactory except for the cases where traffic of the Asterix type is dominant which is due the discrepancy of the mean rate, see table 1.

3 Real MPEG traces as input traffic

In section 2.1 it was noted that a number of compromises and approximations had to be adopted in order to obtain an MPEG like traffic generator which can be implemented in the ATM-100 test tool. To evaluate the accuracy of these artificial MPEG sources, a multiplexing scenario where real MPEG traces are mixed with on/off traffic has been simulated and compared with the experimental results. It is of course the same Bond and Asterix traces which also the artificial models are based on. As in the experiments two scheduling disciplines, FIFO and strict delay priority have been used.

3.1 Description of the trace driven program

The input to the trace driven simulation program is a number of real MPEG traces each given as sequences of integers representing the number of bits in the video frames. Furthermore, it is possible to specify a number of on/off sources with either Erlangian or hyper exponential burst and silence durations.

In order to obtain a sufficient number of MPEG sources the program can chop a video sequence into a number equally sized subsequences implying that it is possible to multiplex a high number of MPEG sources on the expense of a shorter simulation time. In order to avoid visible dependence between consecutive pieces of the chopped video traces there is a limit to how many pieces a video sequence can be chopped into. In this study we have considered 1 minute as a desirable minimum and 20 seconds as an absolute minimum. This corresponds to 24 (resp. 72) identical MPEG sources. For the segmentation and transmission of the content of a video frame, first the size of the payload must be determined (47 bytes if AAL5 is used) and furthermore, it must be decided whether the transmission of cells is spaced to take the whole frame duration or the transmission rate is determined by the amount of data in the frame with largest content.

Due to the periodic nature of the MPEG coding, multiplexing experiments are highly dependent on how the actual phasing between the MPEG sources are. To avoid this property jeopardising the multiplexing experiments a phasing parameter can be specified. Used with the correct value random phasing between the MPEG sources is obtained, conditioned that a relatively large number of simulation runs are performed and relatively large confidence intervals are acceptable.

3.2 Results

First the 3-dim admission boundary with Bond MPEG on the x-axis, Asterix MPEG on the y-axis and on/off sources on the z-axis is determined where the two types of MPEG sources are directly from the real traces. It is the case where MPEG traffic has strict priority over on/off traffic that is displayed and as in previous figures a cell loss of 10^{-4} is the target cell loss. Figure 8 shows the result which does not differ in any significant way from the admission boundary given by the testbed experiments that were based on artificial MPEG sources. The admission boundary here indicates that we are able to carry two more Asterix source but one less Bond source when close to the homogeneous corners with minimal on/off traffic.

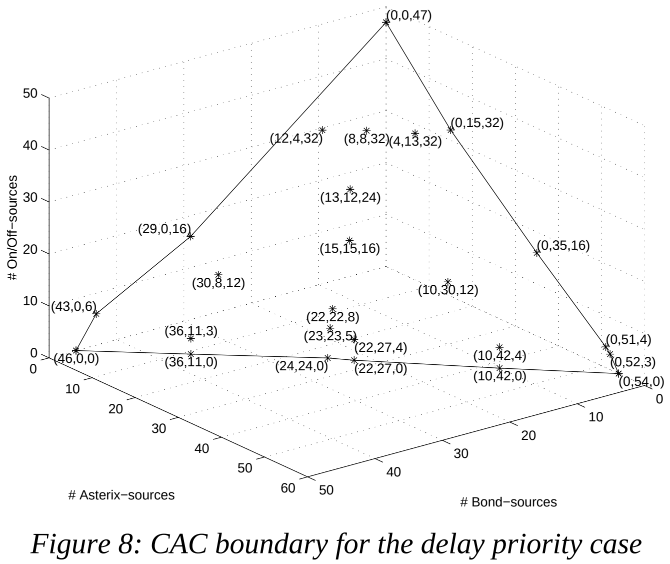

Figure 8: CAC boundary for the delay priority case

To take the comparison between the testbed results and the tracedriven simulations a bit further table 5 shows the actual cell loss figures for three of the points on the experimental admission boundary with a small buffer and no delay priority.

Buffer	Bond	Asterix	On/Off	Experimental CLR MPEG	Experimental CLR On/Off	Trace driven simulated CLR MPEG	Trace driven simulated CLR On/Off
small	11	11	11	0.064 10^{-4}	0.154 10^{-4}	0.12 $10^{-4} \pm$ 0.2 10^{-4}	0.53 $10^{-4} \pm$ 0.7 10^{-4}
small	12	12	12	0.446 10^{-4}	1.418 10^{-4}	0.43 $10^{-4} \pm$ 0.4 10^{-4}	1.70 $10^{-4} \pm$ 1.4 10^{-4}
large	16	16	16	0.072 10^{-4}	0.154 10^{-4}	0.35 $10^{-4} \pm$ 0.4 10^{-4}	0.78 $10^{-4} \pm$ 1.0 10^{-4}
large	17	17	17	0.380 10^{-4}	0.740 10^{-4}	1.41 $10^{-4} \pm$ 1.0 10^{-4}	3.03 $10^{-4} \pm$ 2.1 10^{-4}
priority	16	16	16	0	0.595 10^{-4}	0	1.56 $10^{-4} \pm$ 1.8 10^{-4}

Table 5: Comparison of cell loss figures between experiments and trace driven simulations

There is, taking into account the compromises in the artificial models, a reasonable agreement between the trace driven simulations and the experimental results with a tendency of the trace driven results to yield higher losses. The confidence intervals are quite big which is due to the tremendous influence the phasing of the individual MPEG sources has on the performance in the multiplexing buffer. Each of these confidence intervals is based on 200 simulation runs.

In the second plot (figure 9) a 2-dim admission boundary is displayed where the x-axis contains x Bond + x Asterix sources (i.e. we have constrained the number of Asterix and Bond sources to be equal) and the y-axis contains the number of on/off sources.

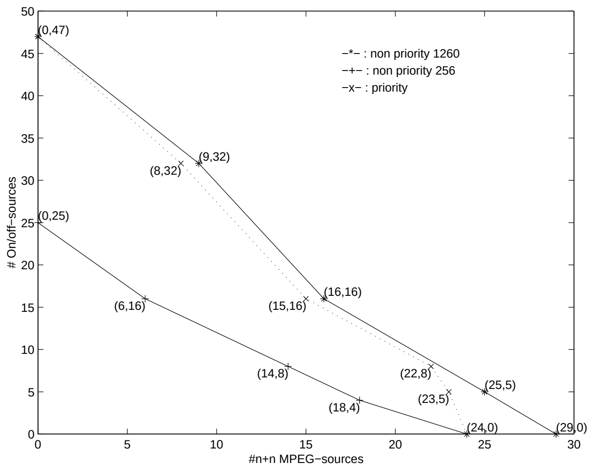

Figure 9: CAC boundaries for MPEG + On/Off

Here both non priority cases (small and large buffer) as well as the delay priority case are displayed. The boundaries have small discrepancies from linear which is due to the large confidence intervals in the simulations. However, the results do indicate that the admission set in the delay priority case, as a first approximation, can be determined as the intersection of the admission set in the non-delay priority case with large buffer and both types of traffic and the set determined by the high priority traffic multipexed in the small buffer alone.

4 Evaluation of three methods for effective bandwidth

The admission boundaries obtained in the former two sections are in the non priority case reasonable close to linear thus supporting the concept of effective bandwidth. It therefore appears natural to compare the effective bandwidth that comes out of our experiments with a few of the more significant methods for computing the effective bandwidth that the literature has presented.

The first method we have chosen to evaluate is a very simple emperical formula first suggested by K. Lindberger in a special case and at several occcations generalised within the Cost 224 and Cost 242 project. An underlying assumption in this first method is that the buffer is very small such that it can be neglected on the burst scale. In section 5.2.2 of the final report from the Cost 242 project [4], the latest version is presented and the formula is

$$e_{BWD} = am + b\frac{\sigma^2}{c}$$

where m is the mean rate, c is the bandwidth on the link, while the coefficient $a = 1 - \frac{\log_{10} P_{loss}}{50}$ and $b/a = -6\, log_{10} P_{loss.}$ The σ^2 quantity is the variance of the stationary rate distribution which in case of on/off traffic with peak rate p equals $m(p\text{-}m)$. The variance of the rate distribution of the two MPEG sources is computed numerically.

The second method assumes large buffers and a system that is not lighly loaded. Under these conditions section 17.1.1 in [4] derives the following formula for the effective bandwidth

$$e_{BWD} = m - \log P_{loss} \frac{v}{2B}$$

where v is the asymptotic variance of the arrival process and is defined as $\lim_{t \to \infty} \frac{Var\{W(0,t)\}}{t}$ where $W(0,t)$ is the amount of arrivals in the time interval $(0,t)$. For an on/off process with exponential on- and off duration it can be shown that $v = \frac{2{t_{on}}^2 {t_{off}}^2 p^2}{(t_{on} + t_{off})^3}$ where t_{on} and t_{off} are the mean on and off durations. For the two MPEG streams we have evaluated the asymptotic variance by first differentiating the log moment generating function of $W(0,t)$ twice, then evaluate the obtained expression at $s = 0$, subtract the square of the mean and taking the limit of the resulting function in the t variable when t tends to infinity. The last step was performed numerically.

The third method is the most complicated one. It is based on the theory of large deviation applied in a limiting regime where both buffer space and number of sources tend to infinity simultaneous. An effective bandwidth is first defined as a function of two variables, a space variable s and a time variable t. The definition is as follows

$$e(s,t) = \frac{\log E\{\exp(sW(0,t))\}}{st}$$

where E denotes expectation, and $W(0,t)$ is the amount of arrivals in the time interval $(0,t)$. In [14] the effective bandwidth function is introduced, a nice overview of the properties of $e(s,t)$ is given as well as a number of examples of how it can be computed. Furthermore, explanation why it can be interpreted as an effective bandwidth is also presented.

In order to apply the method in our context, we first have to determine the space and time pair (s_0,t_0) which characterises our multiplexing system. In [15] and [16] it is shown that the pair is the one that solves the following optimisation problem, i.e

$$\sup_t (\inf_s \{st \sum_{j=1}^{J} n_j e_j(s,t) - s(b+ct)\})$$

where J types of sources in which there are n_j sources of type j with effective bandwidth $e_j(s,t)$ that are multiplexed on a link with capacity c and associated output buffer of size b.

It turns out that a cell loss requirement of P_{loss} is fulfilled if

$$\sup_t(\inf_s\{st\sum_{j=1}^{J} n_j e_j(s,t) - s(b+ct)\}) \le \log(P_{loss})$$

Thus in order to be comparable with the other methods for effective bandwidth we need to adjust $ej(s_0,t_0)$ the following way:

$$e_{BWD} = \frac{e_j(s_0,t_0)}{s_0 + \frac{b/t_0 - \log P_{loss}/(s_0 t_0))}{c}}$$

Table 6 summarises our results for the non-priority case with a small buffer.

Source Type	Experimental Results	Trace Driven Results	Emperical zero buffer effective bwd	Large buffer effective bwd	Large deviation effective bwd
MPEG Asterix	0.669 Mbps	0.693 Mbps	0.658 Mbps	1.446 Mbps	0.705 Mbps
MPEG Bond	0.832 Mbps	0.814 Mbps	0.798 Mbps	2.018 Mbps	0.848 Mbps
On/Off	1.498 Mbps	-	2.411 Mbps	2.948 Mbps	2.122 Mbps

Table 6: Effective bandwidth comparison in the case with a small (256 cells) multiplexing buffer

Table 7 summarises our results for the non-priority case with a large buffer.

Source Type	Experimental Results	Trace Driven Results	Emperical zero buffer effective bwd	Large buffer effective bwd	Large deviation effective bwd
MPEG Asterix	0.604 Mbps	0.635 Mbps	0.658 Mbps	0.690 Mbps	0.620 Mbps
MPEG Bond	0.720 Mbps	0.669 Mbps	0.798 Mbps	0.840 Mbps	0.718 Mbps
On/Off	0.797 Mbps	-	2.411 Mbps	0.909 Mbps	1.068 Mbps

Table 7: Effective bandwidth comparison in the case with a large (1260 cells) multiplexing buffer

In table 6 and 7 the effective bandwidth for each traffic type is obtained under the assumption that this traffic type is alone on the link. In order to see how sensitive the effective bandwidth is to the traffic mix, experimental results and large deviation effective bandwidth results are recomputed under the assumption that the number of sources of all three traffic types are the same. The result is given in table 8.

Source Type	Buffer Size	Experimental Results	Trace Driven Results	Large Deviation Results
MPEG Asterix	256 cells	0.759 Mbps	0.784 Mbps	0.815 Mbps
MPEG Bond	256 cells	0.944 Mbps	0.922 Mbps	0.932 Mbps
On/Off	256 cells	1.700 Mbps	1.696 Mbps	2.265 Mbps
MPEG Asterix	1260 cells	0.627 Mbps	0.707 Mbps	0.642 Mbps
MPEG Bond	1260 cells	0.748 Mbps	0.745 Mbps	0.728 Mbps
On/Off	1260 cells	0.828 Mbps	0.888 Mbps	1.114 Mbps

Table 8: Effective bandwidth comparison in the case where all three kinds of traffic is present

As a comparison between table 6-7 and table 8 shows, the effective bandwidth computed in a heteronegeous mix is sligthly higher than when computed in idealised homogeneous environment. This indicates that the complement of the admission region is convex as both numerical studies as well as more theoretical studies have demonstrated.

Neither the emperical formula nor the large buffer formula can be applied in the case with strict delay priorities. The method based on large deviation can however be applied if the following approximate situation is assumed. High priority traffic sees the whole link associated with a small buffer and without interference from low priority traffic. Low priority traffic sees the whole link associated with a large buffer and it is sharing the link capacity with the high priority traffic.

The admission region that comes out of such a system is exactly the intersection of the admission region obtained from multiplexing all the traffic in the large buffer and the admission region of the high priority traffic multiplexed alone in a small buffer. For the large deviation method the effective bandwidth of the high priority traffic is to be based on the (s_0,t_0) pair determined by the multiplexing of high priority traffic alone in a small buffer while the effective bandwidth of the low priority traffic is to be obtained from the (s_0,t_0) pair determined by the multiplexing of the total traffic in the large buffer, and these two (s_0,t_0) pairs will in general differ significantly.

4 Conclusions

Extensive multiplexing experiments using model-based traffic generators have been carried out at the Expert ATM testbed in Basel to obtain CAC boundaries for traffic scenarios involving MPEG video and data traffic. Both FIFO as well as strict delay priority scheduling schemes have been studied, and the experimental results have been compared with simulation results and results based on effective bandwidth methods.

The approximate linearity of the experimental CAC boundaries shows that the effective bandwidth concept can be applied. However, for the delay priority case both the 2 and 3 dimensional CAC boundaries show that the linear allocation scheme should be modified by constructing the acceptance region from an intersection of two linearly bounded regions which is in accordance with what one of the methods for computing effective bandwidth says.

The flexibility of software simulations has been exploited to investigate any possible measurement error introduced due to differences between the derived MPEG models and MPEG models as implemented in the traffic generators. The simulation results support the validity of the traffic source implementation.

Trace driven simulations have been performed to compare the multiplexing behaviour of the artificial MPEG sources with the multiplexing of the video traces which the artificial sources are based on. Although this procedure is associated with some difficulties due to the dependency on phasing between individual traces, some evidence is given to the validity of the chosen MPEG modelling approach.

Furthermore, three of the methods to compute effective bandwidth known from the literature have been evaluated and the one based on large deviations has proven to be versatile enough to be applied both in the FIFO as well as in the strict delay priority case and to be sufficiently accurate in multiplexing scenarios with both small and large buffers.

By combining results from hardware experiments, software simulations and effective bandwidth methods we have demonstrated a coherent approach for investigating multiplexed video and data traffic. Application of these methods can be used for developing efficient network resource management schemes.

Acknowledgements: The achievements being made within the ACTS project EXPERT are only possible with the conscientious co-operation of all partners, together with the support of the Project Officer and funding from the Commission of the European Union. The contribution of all these players is therefore hereby gratefully acknowledged. We also acknowledge O. Rose of the University of Wuerzburg for providing us the video traces, as well as H. Christiansen of the Technical University of Denmark for his trace driven program with which results of section 3 have been derived.

References

[1] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, J.D. Robbins: *Performance Models of Statistical Multiplexing in Packet Video Multiplexers*, IEEE Trans. on Commm., vol. 36, pp. 834-844, July 1988.

[2] D. Reininger, B. Melamed, D. Raychaudhuri: *Variable Bit Rate MPEG video: Characteristics, Modelling and Multiplexing*, ITC-14, pp. 295-306, Antibes 1994.

[3] O. Rose: *Discrete Time Analysis of a Finite Buffer with VBR MPEG Video Traffic Input*, ITC-15, pp. 413-422, Washington, June 1997.

[4] J. Roberts, U. Mocci, J. Virtamo (ed.): *Broadband Network Traffic*, Final Report of Action Cost 242, Springer Verlag, 1996.

[5] W.E. Leland, M.S. Taqqu, W. Willinger, D.V. Wilson: *On the Self-Similar Nature of Ethernet Traffic*, IEEE/ACM Trans. on Networking, vol. 2, pp. 1-15, February 1994.

[6] D. Le Gall: *MPEG: A Video Compression Scheme Standard for Multimedia Applications*, Comm. ACM, 34(4), pp. 46-58, April 1991.

[7] Draft International Standard ISO/IEC DIS13818-2: *Generic Coding of Moving Pictures and Associated Audio Part 2 Video*, 1993.

[8] International Standard ISO/IEC DIS 13818-2: *Coding of Moving Pictures and Associated Audio for Digital Storage Media up to 1.5 Mbit/s Part 2 Video,* 1993.

[9] G. Ramamurthy, B. Sengupta: *Modelling and Analysis of a Variable Bit Rate Video Multiplexer*, IEEE Infocom 92, Session 6C,1992.

[10] J. Beran, R. Sherman, M.S. Taqqua, W. Willinger: *Variable-bit-rate Video Traffic and Long Range Dependence*, IEEE Trans. on Comm. , 43(2/3/4), pp. 1566-1579, 1995.

[11] B.E. Helvik: *MPEG Source Type Models for the STG (Synthesized Traffic Generator)*, SINTEF report STF40 A96016, February 1996.

[12] B. Helvik, O. Melteig, L. Morland: *The Synthesized Traffic Generator: Objectives, Design and Capabilities*, Proceedings of the IBCN&S, Copenhagen, April 20-23 1993.

[13] F.P. Kelly: *Effective Bandwidth at Multi-Class Queues*, Queueing systems 9, pp. 5-16, 1991

[14] F.P. Kelly, *Notes on effective bandwidth*, Stochastic Networks:Theory and Applications, Oxford University Press, 1996, pp. 141-168.

[15] C. Courcoubetis, R. Weber: *Effective bandwidth for stationary sources*, Prob. Eng. Inf. Sci, **9** pp 285-296.

[16] A. Simonian, J. Guilbert, *Large deviation approximation for fluid sources fed by a large number of on/off sources*, ITC-14, France, June 1994.