Performance Limitations of a Banyan-based ATM Switching System under Multiple, Shaped Traffic Flows

PAULO CARVALHO
University of Minho
Largo do Paco
Dept. of Informatics
4719 Braga Codex, Portugal
paulo@uminho.pt

PETER LININGTON
University of Kent at Canterbury
Computing Laboratory
Canterbury
Kent CT2 7NF, United Kingdom
pfl@ukc.ac.uk

ABSTRACT

This document concentrates on the behaviour of a blocking, input-buffered ATM switch when carrying multiple block-based traffic flows. The impact which the periodicity inherent to certain traffic patterns may have for the performance of ATM networks and systems is studied using both simulation and experiments. The results show that for concurrent traffic flows the switch, which implements an independent round-robin scheduling algorithm on a switching element basis, can be extremely unfair for some traffic patterns. In fact, a traffic flow can be significantly delayed comparing to other competing flows. In the presence of MPEG video, the scheduling process can be influenced by a memory effect which can persist over several frames, leading to unbalanced delays per VC.

KEYWORDS: ATM, switching, performance, scheduling, periodicity

INTRODUCTION

The asynchronous transfer mode (ATM) is gradually taking a prominent role in the provision of broadband communication services worldwide. Based on a cell technology, ATM is able to integrate efficiently both real-time and non-real-time traffic, while providing quality of service guarantees per connection. In this scenario, the performance evaluation of network elements assumes a vital role.

This paper concentrates on the behaviour of an ATM switching system when carrying block-based traffic, studying the effects which the traffic periodicity may have for the performance of ATM networks. Few studies investigated the impact of traffic flow's periodicity on the overall performance of communication networks [1, 2, 3]. These studies concentrate mainly on how cell loss can be minimised. Our study concentrates on the effects the interaction of regular flows can have in the performance of an ATM switch, giving special attention to the unfairness introduced by the scheduling scheme. The results show that traffic flows can be significantly delayed as result of unfairness introduced by the switch.

NETWORK MODEL

The network model encompasses generically a group of N independent sources, a single ATM switch and a group of N destinations. Each source is seen as a set of m independent primitive traffic sources which are multiplexed onto a single output link. Each traffic source generates time-stamped cells according to a pre-defined traffic model. Access to the output link is determined by the multiplexer in a round-robin (RR) fashion; in a single time unit only one cell can get through the multiplexer. Traffic shaping functions can be performed optionally at the multiplexer.

The ATM switch under study is based in the Fairisle switch design [4]. The configuration adopted corresponds to a 16x16 blocking switch using eight 4x4 switching elements (SEs) interconnected in two stages. This switch is a particular case of an $N * N$ Banyan switch using $n * n$ basic SEs organised in a network with $S = log_n(N)$ stages. The switch is input buffered having one port controller per input; no internal queuing nor output queuing capacity is provided. The input queues are served on a FIFO basis, and when internal contention occurs (at any switching stage), "loser" cells remain at the top of the corresponding input queue. Each SE runs a private RR arbitration algorithm with two priorities, and the outputs of each SE hold a reference to the last entry served in order to increase fairness in the arbitration. The role of a sink is to consume cells, making appropriate measurements (on a VC basis). Time is measured in cell units or time slots (TS).

TRAFFIC LOAD

In this study, the emphasis is put on traffic loads inherently periodic such as video. The use of deterministic traffic patterns (due to its simplicity) aims to allow a better understanding of the switch behaviour. Traffic presenting a regular, block-based structure may result from traffic shaping actions taken by an ATM source.

In order to simplify the problem of video modelling without losing too much generality, our study focuses on a particular type of video sources using the MPEG standard coding algorithm [5].

Model for MPEG video sources

Modelling a video source at frame level leads in general to a distribution of cells per frame which is much harder to understand and reproduce [6]. Both Pancha et al. and Izquierdo et al. [6, 7] recognise that the most appropriate approach for video modelling is at the slice level since MPEG is based essentially on slices. Taking into account these observations, the present study proposes and uses a model based on slices. Our MPEG source is modelled as a periodic process with two distinct periods (P_1 and P_2). P_1 is determined by the frame rate and P_2 is determined by the time required to code a slice (T_s) and has a number of cycles depending on the number of slices in a frame.

The distribution of the number of cells per slice is characterised by having heavier or shorter tails depending on the degree of intraframe/interframe coding [6]. In order to reflect such behaviour it is assumed that the number of cells per slice follows a Gamma distribution with parameters α and β. This distribution has a probability density function given by $f_\Gamma(x) = \beta^\alpha x^{\alpha-1} e^{x\beta}/\Gamma(\alpha)$ with mean α/β and variance α/β^2. The parameters α and β change the shape and scale of the distribution, respectively, e.g. make it more or less bell-shaped with heavier or lighter tails. Here, instead of defining α and β directly, they are devised from two more sensitive parameters: i) the mean number of cells per slice v and ii) its coefficient of variation C. The values of α and β in $f_\Gamma(x)$ are then given by $\alpha = C^{-2}$ and $\beta = \alpha/v$. The values assumed for v have the same order of magnitude as the ones found by Pancha and El Zarki [6]. The values assumed for C can equally be derived from the standard deviation (σ) quoted in [6]. An MPEG source is also called MPEG process with parameters v e C, or MPEG(v, C) in short.

Deterministic on-off traffic

A deterministic on-off source is parameterised by: i) the length of the active period ($0 < b_0 < \infty$) in TS; ii) the global mean slot occupation probability m ($0 \leq m \leq 1$). As a consequence, the length of the inactive period will be $b_1 = b_0 * (1/m - 1)$ for $m \neq 0$. Here, this source is called a deterministic process with parameters m and b_0, or DP(m, b_0) in short.

TRAFFIC INTERACTION

As reported in [8], traffic from identical and synchronous MPEG sources can experience significative variations on mean end-to-end delay (D) and jitter (J). The case study reported here corresponds to the interaction among traffic flows for a configuration comprising four independent MPEG sources, multiplexed on a single SE and competing for the same destination. The sources have identical parameters and generate an aggregate load of 95% without using traffic shaping. Fig. 1 shows the values obtained for D when a deterministic offset a is introduced between adjacent MPEG sources. The intercept in Fig. 1 shows clearly that for a null offset the traffic in VCI6 experiences a higher delay than in the remaining VCIs[1]. The variation exists either for a limited or a full service rate and may affect any of the VCIs. In order to further study this effect, a deterministic load was submitted to the switch.

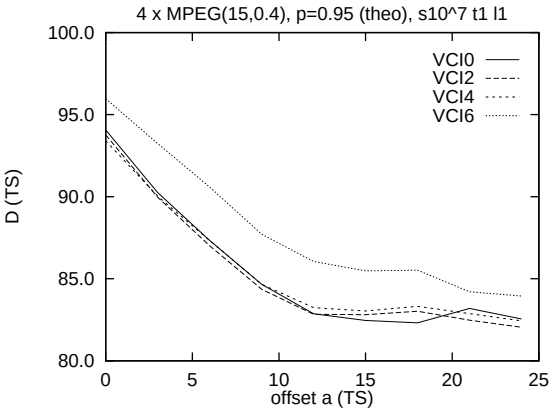

Figure 1: Mean delay for identical MPEG sources when their relative phase (offset) varies

SCHEDULING AT THE SWITCH

The results in Fig. 2 correspond to a configuration including four deterministic sources in-phase (balanced switch configuration). Fig. 2 illustrates clearly how the scheduling may become unbalanced when the burst size (length of the on period) varies. The curves show that traffic going through one of the SEs (src4 and src5) suffers a repetitive and increasingly unbalanced scheduling. For the common case where cells are served at

[1]Here, the same VCI is used from source to destination. Thus, it is used in an end-to-end basis in the same way as a VC.

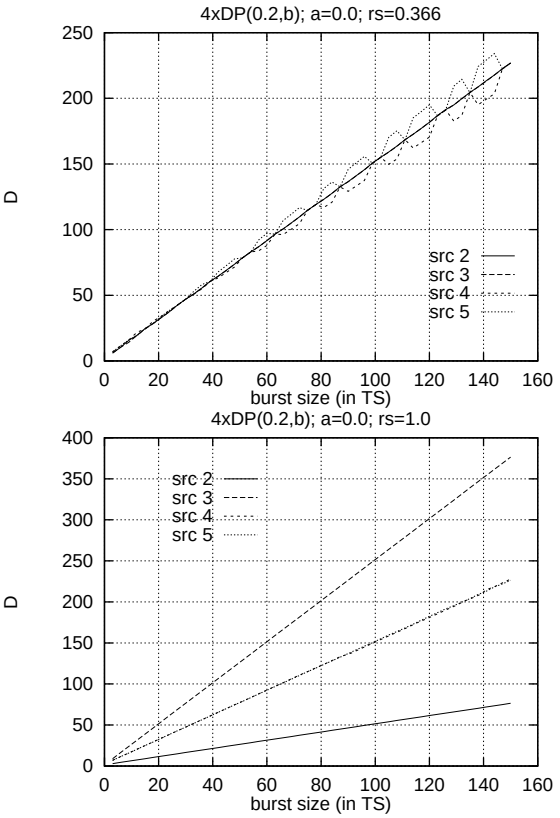

Figure 2: Scheduling sensitivity to the burst size: a) bounded and b) full service rate

full rate, the unfairness effect is even stronger, the traffic flows suffer substantially different delays. These effects were verified and confirmed experimentally both for balanced and unbalanced traffic source distribution over the SEs. The experiment reported in Fig. 3 illustrates how the switch ports are selected for an unbalanced switch load. When the traffic submitted is of the type 2), the sequence a2) shows that the scheduling algorithm is completely unfair as one of the losing contending flows is repeatedly delayed, resulting in a clear benefit for the single port. In the example, traffic on port 2 will be delayed until one of the competing entries is cleared up. For MPEG traffic, it was also found that the scheduling process is influenced by a memory effect which affects the input queues. This effect may persist from slice to slice and over several frames which explains the variation in D reported in Fig. 1.

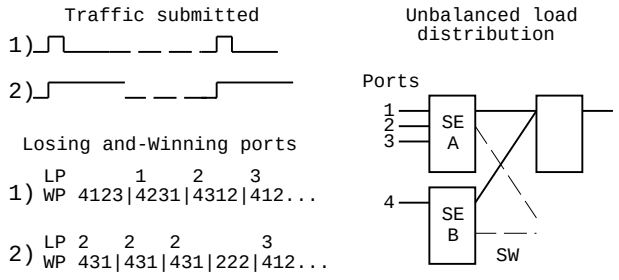

Figure 3: Unfairness in the scheduling process

CONCLUSIONS

As result of traffic interaction, it was found that the scheduling process can be affected by the regularity of traffic patterns such as MPEG video. A point of major concern is that, when multiple SE are involved in the scheduling process, running an independent RR algorithm for each SE output can lead to extremely unfair situations. Through simulation it was shown that, if more than one SE is active in the first switching stage, correlation of the RR systems cause some of the active entries systematically to fail to be selected, resulting in unbalanced delays per VC. Experiments with a Banyan-based switch have also confirmed that cells in a burst can be consecutively blocked until the remaining competing traffic is cleared. This behaviour is a consequence of the independence in the decision process between stages and between SEs.

REFERENCES

[1] V. Ramaswani and W. Willinger. Efficient Traffic Performance Strategies for Packet Multiplexers. *Computer Networks and ISDN Systems*, 20(1-5):401..407, 1990.

[2] Daniel P. Heyman, Ali Tabatabai, and T. V. Lakshman. Statistical Analysis and Simulation Study of Video Teleconference Traffic in ATM Networks . *IEEE Transactions on Circuits and Systems for Video Technology*, 2(1):49..59, March 1992.

[3] A. F. Dagiuklas and M. Ghanbari. Effect of Video Synchronisation in an ATM Multiplexer. *Electronic Letters*, 30(9):683..685, April 1994.

[4] Ian Leslie and Derek McAuley. Fairisle: An ATM Network for the Local Area. In *SIGCOMM'91 Conference - Communications, Architectures and Protocols*, September 1991.

[5] ISO (JTC1/SC2/WG11). ISO 11172 - MPEG (Moving Pictures Expert Group).

[6] Pramod Pancha and Magda El Zarki. MPEG Coding for Variable Bit Rate Video Transmission. *IEEE Communications Magazine*, 32(5):54..66, May 1994.

[7] Michael R. Izquierdo and Douglas R. Reeves. Statistical Characterization VBR MPEG at the Slice Layer. Technical report, North Caroline State University, July 1994.

[8] P. Carvalho. *Interaction of Periodic Bursty Traffic Flows in ATM Networks*. PhD thesis, University of Kent at Canterbury, 1996.