Delay Constraints and Admission Control in ATM networks

ABSTRACT

This article first reviews the architecture of a typical ATM switch and considers the problem of handling both continuous bit oriented and bursty traffic with low loss and delay requirements. We examine a buffer architecture that is capable of handling both delay sensitive and loss sensitive traffic, and present a scheme to allocate network resources, i.e. bandwidth and buffer space, to individual circuits to ensure acceptable delay and loss performance.

1. Introduction

The main resources in the ATM network are the link bandwidth and the buffer storage for the queues in the ATM switch. Bandwidth allocation is a very critical issue since it greatly influences the network performance (delay/cell rejection etc.). To avoid cell losses due to statistical variations in cell arrivals that exceed the transmission capacity of the output ports in the switch, it is necessary to provide buffers in an ATM switch.

2. Buffer Memory in an ATM switch

The buffer memory organization is determined by the interconnection network architecture and the switching principle employed. The ATM switch can be modeled as a set of input queues (incoming links) which through an interconnection network are switched to a set of output queues(outgoing links).

Usually the only purpose of the input queues is packet/bit synchronization and clock recovery, so the input queues size is minimal. Nevertheless, it should be noted that statistical switching techniques will necessitate in some cases large input queue sizes. The output queues are needed since cells from different incoming links that are to be routeed towards the same outgoing link may arrive concurrently. A variety of architectures have been proposed for the buffering of cells in an ATM switch.

The output buffer architectures proposed in [1] and [2] are able to select a cell from any port for transmission, using only one buffer per port.

The original descriptions of input buffered architecture assumed only one sequential buffer, so temporarily untransmitted cells impede transmission of later cells and reduce the effective throughput [3]. This difficulty is resolved by using random-access to N buffers at each input port, so that every input-output port pair has an individual buffer [4]. This architecture has the constraint that cells cannot be sent to all output ports from one input port.

3. Service requirements in Broadband-ISDN

As a target solution for B-ISDN, ATM supports a wide range of voice, data and video services. Narrowband services require low bandwidth or bit rates up to 2 Mbits/second, e.g. audio applications (telephony, hi-fi sound, audio library etc.) and short duration data (teletex, facsimile, telemetry, alarms, electronic mail). Broadband services require high bandwidth or bit-rates greater than 2 Mbits/second, e.g. video (video-telephony, broadcast TV, teleconference, pay TV, HDTV etc.) and long duration data (word-processing, home-computing, database access, computer communication etc.).

According to CCITT Rec.I.121, these services can be continuous bit stream oriented (CBO) or packet oriented (burst) depending on the way of generation of data on the source side.

A CBO source generates cells periodically, one by one. On the other hand, a burst source intermittently generates a series of cells. Within the B-ISDN, CBO services require much higher quality of service than do voice communications. The quality of CBO services has to be at least the same as that for existing circuit switching networks, i.e. no delay fluctuations, cell loss rate : 10^{-8}, switching node delay 0.5 ms [CCITT G.142].

Most burst-oriented services are loss critical providing for high integrity data transmission.

Having in mind that ATM network must satisfy individual service requirements, the strategy used for resource allocation minimizing the delay and loss probability is very important concerning quality of services offered within the B-ISDN.

Based on the service requirements in a B-ISDN, two bearer services can be defined:

1) *quasi-static (synchronous)* for continuous bit stream oriented services, and

2) *burst-mode* bearer service for packet oriented services

The quasi-static connection, through network resource allocation guarantees maximum/average bandwidth with minimal delay. This service will provide for time-critical services.

Two types of burst-mode connections can be defined.

low delay burst mode connection providing bounded delay services, but some packets can be lost.

low loss burst-mode connection providing for high integrity of data (no packet loss) at the expense of an increased delay.

4. Delay Characterization in an ATM network

The delay through an ATM network is determined by the following parameters:

Transmission delay: This delay depends on the distance between end-users. It is independent of the type of switching node used and by no means controllable.

Cell assembling delay: This is the delay introduced every time a source bit stream, generated at a sending terminal is converted into cells.

Switching delay The switching delay is composed of a fixed part, implementation dependent, but very low[2] and a queueing delay. The queueing delay is determined by the queues in the ATM node and can be controlled by restricting the number of connections through the switch.

Cell disassembly delay The delay introduced because the arriving cells are stored in buffers at the receiving terminals in order to regenerate the original bit stream.

Overall delay, called the end-to-end delay is equal to the sum of all these delays.

Since the transmission delay is uncontrollable and the fixed switching delay is negligible, only the cell assembly delay, the queueing and the cell disassembly delay are introduced in the model proposed below. Of these queueing delays are the most significant and admission control strategies should ensure that excess queueing delay is avoided.

5. Model

The physical carrier for the information transfer in an ATM network is the pipe. It transmits in time division, the information blocks belonging to different connections. Comparing with synchronous time division, a label is allocated to each block (cell) transmitted in a time slot of duration t identifying the virtual circuit used.

In our model, we assume that there are several input pipes and one output pipe. This is a reasonable assumption because after call-setup the path between two customers is fixed. Thus the destination pipe for each block of information is known immediately at its arrival in the switch and no further routing is to be done after the call is established. Thus we concentrate on one output pipe and compare different output buffering schemes for them.

Quasi-static bearer service packets should be given priority over burst-mode bearer service packets taking into account the fact that delay constraints are more stringent for quasi-static services. On the other hand, the influence of burst traffic on quasi-static traffic is eliminated [5]. For burst mode connections, priority is given to low delay connection over low loss connection.

Thus, a buffer architecture with 3 priority classes appears to be most suitable for handling traffic sensitive to delay and losses [6]. CBO traffic gets the highest priority, low delay burst mode traffic the next highest and low loss burst mode traffic cells the lowest priority. Cells at a given priority level are shipped only when there are no cells present at any higher priority levels (See Figure 1).

5.1. Service Specification

In order to reserve resources for a connection, we require that the their traffic characteristics be declared at the connection establishment time. For a quasi-static service, we need to know the average interval x_{ave} between two consecutive cells on that connection. In order to make the average meaningful, we also require that an averaging interval I be declared by the service client. In the simplest case, I can be a network wide constant. with a value of a few tens of milliseconds. With a payload of 48 bytes, a quasi-static service requiring a bandwidth of 2 Mbps (5208 ATM cells per second) will have an x_{ave} of 191 μs.

Burst mode bearer services will need two parameters, the average inter-cell arrival time x_{ave}, with the same meaning as for the quasi-static services, and a minimum inter-cell arrival time, $x_{\min}$. The ratio between these two intervals is a simple measure of burstiness. Suppose that TV-quality video is compressed to provide a bursty traffic at about 2 Mbps with a burstiness of 3, then, the x_{ave} value is 191 μs while the value of $x_{\min}$ will be about 64 μs.

Notice that the two interval values correspond to the average and peak bandwidths required for the connection. For a quasi-static connection, we can define the average inter-arrival time to be the same as the minimum inter-arrival time.

6. Admission Control

With the specification and buffer model mentioned in the previous section, we now present a set of tests which can be used to verify if a particular set of channels can be accepted in an ATM pipe. We assume that quasi-static connections have a maximum upper delay bound requirement of d_{cbo} and burst mode connections have a delay bound of d_{brst}, which is somewhat larger than the requirements of quasi-static traffic. If t is the time required to ship a cell on the output pipe, the first test which needs to be satisfied is

$$\sum_j (t / x_{min,j}) < 1, \tag{1}$$

where the sum extends to all the quasi-static channels in the output pipe. This condition simply ensures that the total bandwidth required by all the connections does not exceed the bandwidth of the output pipe.

If N_{cbo} is the total number of quasi-static channels in the output pipe, then condition (1) ensures that the maximum delay of a cell arriving on a quasi-static connection can not exceed $(N_{cbo}+1)t$. A new channel can be admitted only if this value is lower than the maximum acceptable delay for a cell on a quasi-static traffic, d_{cbo}.

Burst mode connections can be handled in the same manner as quasi-static services, but it may lead to a very poor utilization of resources. In order to achieve higher utilization of the output pipe, it is necessary to overbook the bandwidth required for bursty channels. The overbooking may cause some of the cells in these connections to be delayed or lost beyond their required bounds. However, one may restrict the probability of this undesirable event, to a certain low bound p_{ovf}.

Overbooking essentially means a violation of condition (1). Due to the bursty nature of traffic on the burst mode connections, different combinations of connections will be active (sending cells) at different times, some combinations satisfying condition (1) and others not. If we make the pessimistic assumption that cells are going to be delayed and lost whenever condition (1) is not met, we must restrict the probability that of any of these dangerous connections are active below p_{ovf}. A simple way to evaluate this probability, with the prescribed traffic specifications, is described in [7].

Figure 2 shows the maximum possible utilization of the output pipe in the presence of only low delay burst mode connections, with a delay bound d_{brst} of $100t$ and p_{ovf} of 10^{-9}. The horizontal axis shows the maximum number of connections that can satisfy the quasi-static condition (1). Utilization is poor for bursty channels with high peak bandwidth, but fair for channels with lower bandwidth, less than $1/30$ of the total output pipe bandwidth. From figure 2, we see that utilization when the peak bandwidth is $1/40$ of the output pipe bandwidth, is about 40 % with a burstiness of 4. Thus a 45 Mbps output link can support traffic with burstiness of 4 at a utilization of 40% if the maximum peak bandwidth is about 1 Mbps. In practice the actual utilization would be higher in the presence of quasi-static and low loss burst mode connections.

Let us consider the largest combination of channels that satisfies condition (1). This set will consist of all the N_{cbo} quasi-static connections in the output pipe and some N_{brst} low delay burst mode connections, selected in the decreasing order of $x_{\min}$ values. We now need to verify that

$$d_{brst} \le \frac{N_{cbo} + N_{brst}}{(1 - \rho_{cbo})}, \tag{2}$$

where ρ_{cbo} is the utilization of the output pipe by all the existing quasi-statisc connections.

A shared buffer architecture requires much less buffer space than one in which all channels have separate buffers [6]. If a rate control policy to ensure that channel clients are not misbehaving, similar to one in [7] can be implemented, it would suffice to allocate N_{cbo} buffers for the quasi-static traffic. For low delay burst mode connections, let λ_{brst} be the combined peak bandwidth of the low delay burst mode connections for the N_{brst} connections selected for condition (2). We require either N_{brst} buffers or $d_{brst}\lambda_{brst}$, buffers, whichever is greater. The low loss burst mode connection may have a delay as long as I, although it is likely to be much less depending on the utilization of the output switch, and thus $I\lambda_{rest}$ buffers suffice to eliminate buffer losses; λ_{rest} being the total bandwidth of all low loss burst mode connections.

We verified the correctness and performance of the admission control policy by means of simulations written in CSIM [8]. We measured the actual delays and buffer requirements observed in the simulations in the presence of all three kinds of traffic and compared them against the predicted upper bounds given by our admission control rules. Channels were created with a peak bandwidth $1/20$ of the output pipe bandwidth, and bursty channels had a burstiness of 4. A channel was equally likely to be of any type, and were accepted for a p_{ovf} value of 10^{-9}. An on/off model was used for generating bursty traffic. In order to keep the simulation runs to reasonable lengths, we observed the amount of delay observed in our simulations with a probability of 10^{-4} and extrapolated it to 10^{-9}. Figure 3 compares the simulated delays with the predicted maximum delays, and Figure 4 compares the buffer allocations with the actual utilized buffers. Our predicted bounds are always met, confirming the correctness of the admission control schemes. While the actual buffer utilizations for low loss bursty traffic is poor, a result of our simplistic buffer allocation policy, the other bounds appear to be quite close to the observed values.

7. Conclusions

Comparing with the previous works we introduce a common buffer storage allocation architecture, differentiating CBO and burst mode services. We introduce priority schemes corresponding to the real traffic characteristics taking into account all the delay component between end users. We find that a buffer architecture which differentiates between quasi-static traffic and two types of bursty traffic has the best possible performance among all the examined architectures. We develop a set of admission control rules to accept connections with this buffer architecture. Our control rules are easy to implement and evaluate, and lead to fair utilization of network resources in the presence of bursty traffic.

8. References

[1] J.P.Condruse and M.Servel, "Prelude: an Asynchronous Time-Division Switching Network", Proc. of ICC '87, June 1987, Seattle, pp 769-773.

[2] Y.S.Yeh, M.G.Hlyuchi and A.S.Acompora, "The knockout switch: a simple modular architecture for high-performance packet switching", ISS '87, March 1987.

[3] M.J.Karol, M.G.Hlyuchi and S.P. Morgan, "Input versus output queueing on a space-division packet switch", IEEE Trans. on Commun., COM-35, No 12, Dec. 1987, pp 1347-1356

[4] G.J.Fitzpatrick and E.A.Munter, "Input Buffered ATM switch traffic Performance", Multimedia '89 Ottawa, April 20-23, 1989 paper 4.2

[5] T.Murase, H.Suzuki, T.Takenchi, "Continuous Bit Stream Oriented Services in ATM network", C&C Systems Research Laboratoories, NEC Corporation.

[6] P. Todorova and D. C. Verma, "Resource Allocation and Delay Constraints in ATM Networks", to appear in Second IEEE Workshop on Future Trends of Distributed Computing Systems, Cairo, 1990.

[7] D. Ferrari and D. C. Verma, "A Scheme for Real-Time Channel Establishment in Wide-Area Networks", IEEE J. Selected Areas Communic., vol. SAC-8, n. 3, pp. 368-379, April 1990.

[8] H. Schwetman, "CSIM Reference Manual (Revision 12)", MCC Tech. Rept. No. ACA-ST-252-87, November 1987.

Figure 1

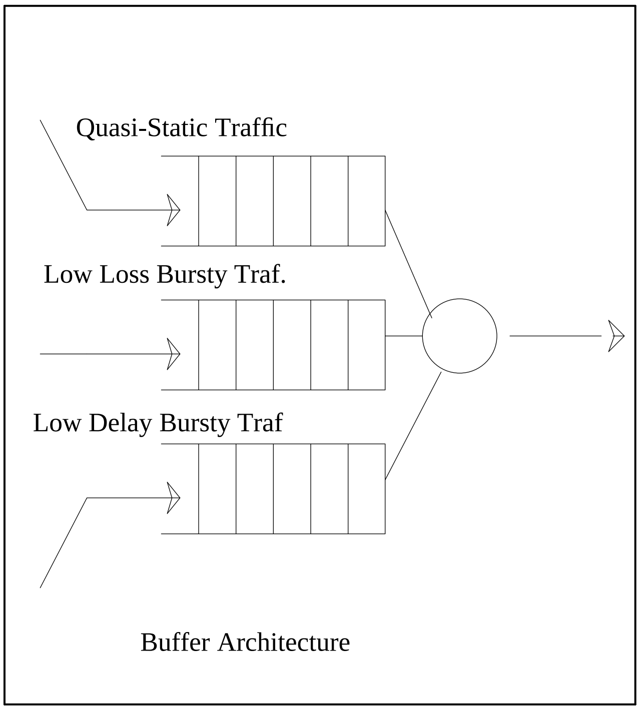

Figure 3

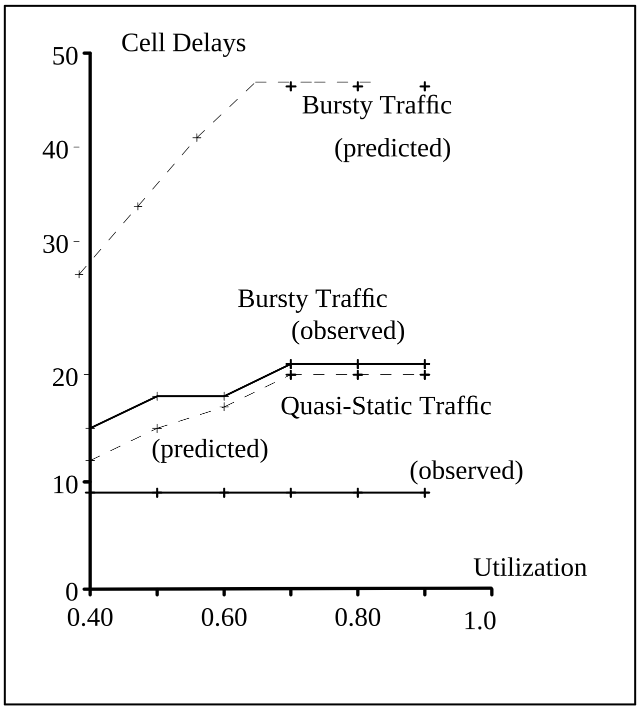

Figure 2

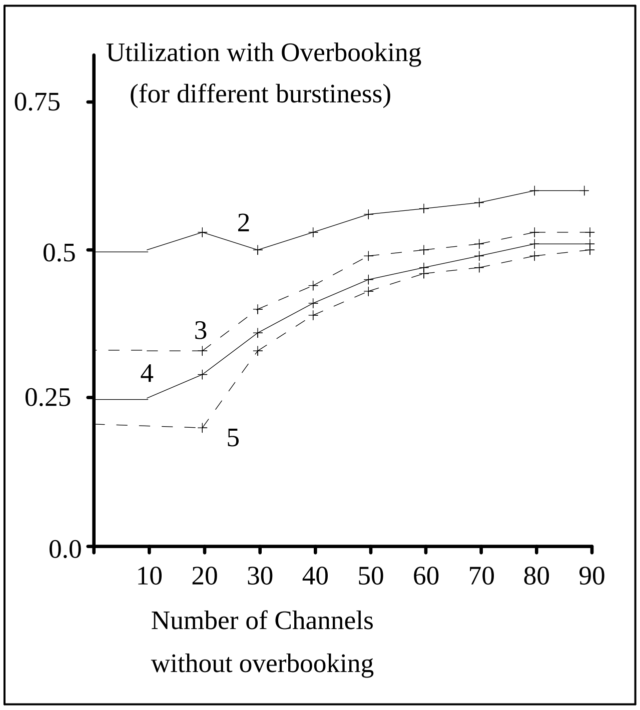

Figure 4