

Traffic Management in an ATM Multi-Service Switch for Workgroups

*Stefan Bodamer and Thomas Renger
University of Stuttgart
Institute of Communication Networks and Computer
Engineering, Pfaffenwaldring 47, D-70569 Stuttgart
e-mail: {bodamer, renger}@ind.uni-stuttgart.de*

*Georg Rößler
Bosch Telecom
UC-PN/EGH, Kleyerstraße 94, D-60326 Frankfurt am Main
e-mail: Georg.Roessler@pcm.bosch.de*

Abstract

ATM switches have to provide traffic management functions to meet the QoS requirements of different service categories. Among the traffic management functions we will focus on connection admission control (CAC) and priority control in this paper. Besides the mechanisms and algorithms behind these functions the issue of application to a concrete architecture of an ATM multi-service workgroup switch is emphasized.

For CAC a new method is proposed that is based on a simple approximation of the effective bandwidth. The difficulties occurring if connections with different parameters are mixed together are solved by handling CBR traffic separately and dividing the parameter space for VBR connections into several regions. Within the VBR regions the linear method is applied while for connections associated to different regions a reduced service rate is considered.

The performance of both the priority control mechanisms and the CAC method is evaluated using analysis and simulation.

Keywords

ATM, Traffic Management, Connection Admission Control, Priority Control

1 INTRODUCTION

Traffic management in ATM switches comprises several functions which act at all levels from cell level to connection level in order to make sure that the Quality of Service (QoS) requirements are met for all connections carried by the switch. Traffic management becomes especially demanding if traffic ranging from video conferencing over interactive applications like the World Wide Web (WWW) to bulk data transfer has to be handled by a switch with minimal interference between different traffic streams.

Video conference connections have stringent requirements in terms of delay and cell loss ratio which can only be fulfilled if the service rate closely matches the rate of the offered traffic. Data traffic, however, is very elastic in that it is not sensitive to delay and can be served at any rate the switch can provide. Of course, large buffers are needed to store bursts until they obtain service. Interactive applications like WWW generate traffic that is somewhat elastic but users naturally prefer short response times. Hence, the service rate for these connections should approximately follow the demand. The large spectrum of traffic characteristics and QoS requirements has led to the definition of five ATM service categories (CBR, rt-VBR, nrt-VBR, ABR, UBR) by the ATM Forum (1996) which largely correspond to the transfer capabilities specified by the ITU (1996), and the discussion about new categories is still ongoing.

The key areas of traffic management are connection admission control (CAC), usage parameter control (UPC), and priority control in the switching hardware, all of which have been extensively studied since many years (Kröner, 1995). Quite often, however, only one particular aspect is investigated and described in detail, whereas other parts of the system are neglected or approximated by some simplifying assumptions. In this paper, we concentrate on the ensemble of traffic management functions of an ATM Multi-Service Switch (MSS) which supports all traffic categories defined by the ATM Forum.

The next chapter presents the switching architecture of the ATM MSS and the priority control functions implemented in the system. Section 3 focuses on the CAC function of the MSS which takes into account that multiple service categories have to be supported. Results of both cell level simulations of the system and analytical as well as simulative investigations of the CAC are presented in Section 4.

2 SWITCHING ARCHITECTURE

2.1 Overview

The ATM MSS uses a folded bus as its switching fabric to which different types of units are connected (Figure 1). The bus is used for transmitting both user data and switch-internal control information. The main call processing functions are located on a Control Unit (CU), whereas terminal equipment is connected to Port Units (PU). Two types of PUs are currently defined: the PU25 contains twelve 25 Mbit/s ports, the PU155 has four 155 Mbit/s ports.

An ATM MSS consists of up to 16 units which share the bus with a total capacity of 1.6 Gbit/s (switch internal cells have a length of 56 octets).

All units are connected to the bus through the Bus Access Controller (BAC) which is realised as an ASIC. The BACs together with the bus implement the ATM layer in the

MSS. Physical layer chips and AAL chips exchange cells with the BAC via a standard Utopia interface (ATM Forum, 1995). The total capacity of this interface is 400 Mbit/s.

The input part of the BAC contains a small input buffer for 20 cells as well as control logic for a much larger External Input Buffer (EIB). An output buffer holds approximately 100 cells after the transfer via the bus until they can be forwarded to their destination ports. Rate matching buffers with small capacity for every port complete the picture of the queuing model. The resulting model of the ATM MSS, as it has been used for simulations, can be found in Figure 2.

The ATM MSS belongs to the class of ATM switches with buffers on both the input and output sides. The bus which implements the switching fabric is a shared resource, which means that blocking may occur when it is occupied by another unit.

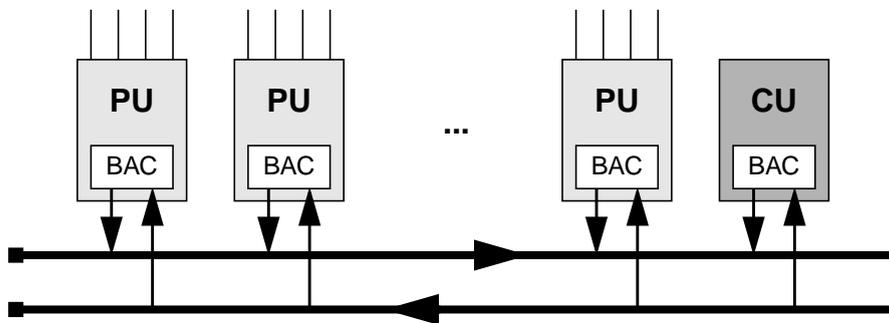


Figure 1 Internal switch architecture

2.2 Bus Access

Access to the bus is determined by a distributed bus access protocol which is implemented by the BAC. The bus consists of a data bus and a separate control bus used parallel to the data transmission in the current slot to determine the PU which will obtain the next slot. All BACs with a cell to send compete for the next slot by sending their unit identifier on the control bus. The PU with the dominant identifier receives back its own identifier and hence knows that it is allowed to use the next slot.

Fairness is ensured by a credit mechanism. Each PU uses one credit per transmitted cell. When the credits are exhausted, it may no longer compete for time slots. The unit continues to observe the control bus and regenerates a number of credits when it detects that no unit competes for the next slot which means that the other PUs either have no cells to send, or their credits are exhausted.

A compromise has to be found for the number of credits to regenerate. On one hand, the number of credits has to be small to limit the cell delay variation as well as to minimise the risk of cell loss. On the other hand, a larger number of credits reduces the overhead caused by the slot which is lost when all PUs regenerate credits. A good compromise is to generate in the order of 10 credits per PU for each cycle. Note that different credit values may be assigned to the PUs in order to prioritise some of them.

The BAC also implements a backpressure (BP) mechanism which avoids cell loss if the output buffer is occupied. In this case, the receiving BAC signals that it could not store the cell. The sending BAC then keeps the cell and schedules it for retransmission. Head-of-line blocking is avoided by a service discipline which is not strictly first-in, first out. Up to six cells at the head of the buffer can be considered for transmission on the bus, such that cells waiting behind blocked cells can be served in the meantime. Blocked cells are discarded after a number of failed transmission attempts.

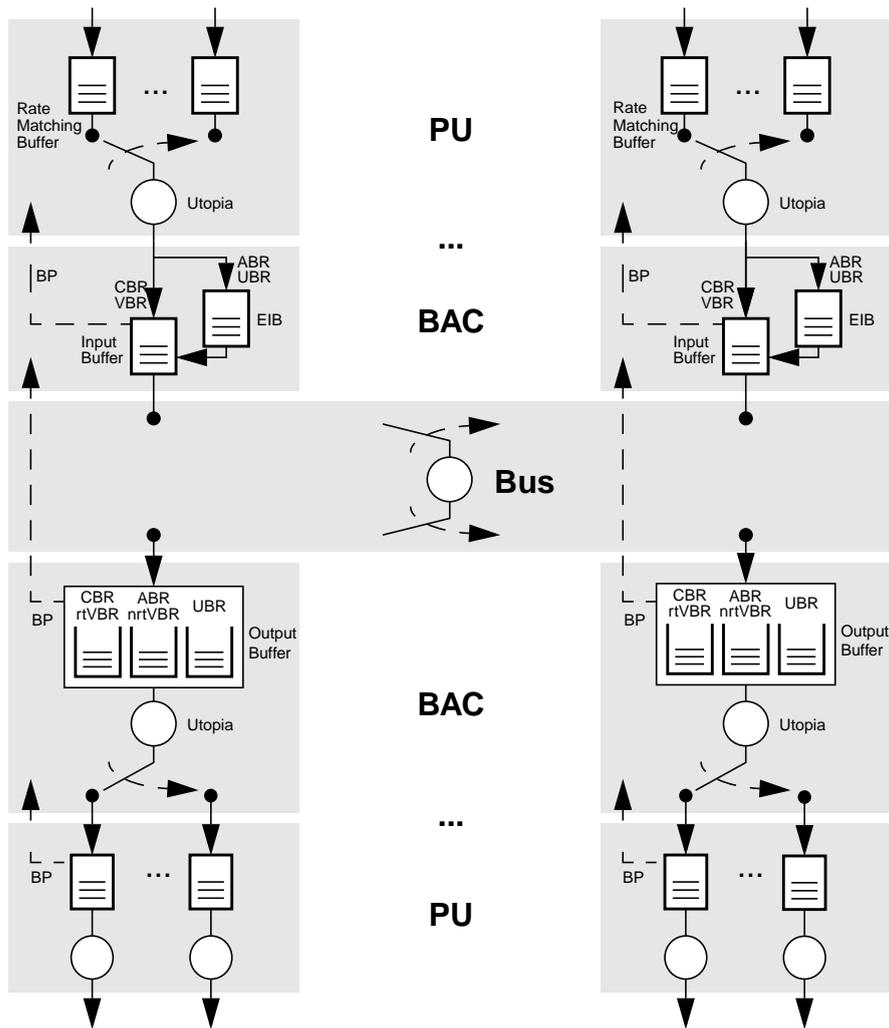


Figure 2 Queuing model of the switch

2.3 Priority Control

The purpose of priority control is to ensure that each service category obtains its required QoS. In the ATM MSS, priority control applies to both cell loss ratio and cell transfer delay.

On the input side of the BAC, UBR and ABR cells are immediately forwarded to the EIB which can store at least 2048 cells. The control logic for the EIB implements partial buffer sharing (Kröner et al., 1991) between ABR and UBR, such that UBR cells are discarded if the occupancy of the EIB becomes high while ABR cells can use the remaining buffer space. The control logic retrieves cells from the EIB only if the occupancy level of the input buffer drops below an adjustable threshold, hence VBR and CBR cells are prioritised.

The output buffer of the BAC implements three logical queues with statically assigned service priorities. CBR and rt-VBR cells are always served first in order to minimize their transit delay, then nrt-VBR and ABR cells, and finally UBR cells. Moreover, a pushout mechanism is implemented which discards a UBR cell if a cell having higher priority arrives from the bus. Thus the backpressure mechanism across the bus is activated only if the output buffer is completely occupied by VBR, CBR and ABR cells. The number of transmission attempts is configurable per service category, and the number is usually limited to one attempt for UBR cells because it doesn't make sense to keep UBR cells for retransmission if the load is already that high that the output buffer is completely occupied.

While the mechanisms described above are sufficient for prioritising VBR and CBR within a PU, they cannot effectively limit the amount of bandwidth on the bus that ABR and UBR traffic from other PUs can occupy. An enhancement of the bus access protocol ensures that the CAC can allocate a significant share of the bus capacity to VBR and CBR connections and that the capacity for this traffic can be guaranteed. Of course the administrator can limit the share of the bus capacity that the CAC allocates to VBR and CBR traffic.

The BAC implements two logical queues in the input buffer, one for VBR and CBR cells, the other for UBR and ABR. VBR and CBR cells always have higher priority, and the BAC competes for a slot if the credits are not exhausted. UBR and ABR cells are considered for transmission only if at least a certain number of credits are left, since part of the credits per cycle are reserved for the transmission of VBR or CBR cells. The number of credits per cycle that UBR or ABR cells may use can be limited to one by setting the limit to the same value as the number of credits to regenerate per cycle. In this case at least 80 percent of the bus capacity can be guaranteed for VBR and CBR traffic.

3 CONNECTION ADMISSION CONTROL

3.1 Requirements

Connection admission control (CAC) has to determine whether a new connection setup request can be accepted or should be rejected. This decision is based on the constraint to meet the negotiated QoS requirements of all existing connections as well as of the new connection. Another desirable goal is to allow for a statistical multiplexing gain, i.e. an

efficient CAC method should accept as many connections as possible without violating any QoS guarantees.

CAC is performed as part of the connection setup process and generally runs as a software module on a switch control processor. Therefore, further evaluation criteria of CAC algorithms are processing time and memory requirements as well as the implementation and integration effort. As setup times and control processor load must be limited, time consuming and complex CAC algorithms are not appropriate for real systems.

The main focus of CAC is on guaranteed traffic, namely CBR and VBR connections and the minimum cell rate (MCR) share of ABR connections. Among that type of traffic statistical multiplexing is only possible for VBR connections. UBR connections are not necessarily controlled by CAC because they cannot influence the performance of other service categories due to priority-based buffer control mechanisms of the switch. The presented system, which is designed for a workgroup environment, needs a CAC algorithm that is reasonably efficient but inexpensive with regard to processor resources.

Finally, a CAC algorithm has to be robust and should therefore not rely on a description given by the user but make a worst-case assumption according to the parameters negotiated during connection setup and monitored by the UPC function. Additionally the CAC may use measurements taken in the switch (e.g. for the actual mean rate of a connection) to improve effectiveness (Kröner et al., 1994; Antunes et al., 1997).

3.2 Classification and Characteristics of CAC Methods

Almost all CAC methods concentrate on the cell loss ratio (CLR) as the most crucial QoS parameter. Some algorithms also take into account delay requirements, e.g. the one presented by Antunes et al. (1997). However, this is not necessary here, as a limitation of transfer delay is guaranteed by the switch architecture with small buffers for real-time traffic and appropriate priority control mechanisms.

The CLR-based CAC methods make use of the results for the cell loss probability in an ATM multiplexer. Cell losses may occur due to congestion on cell level or on burst level. The cell level effects can be quite well described by an M/D/1-(S+1) model (Hui, 1988). The evaluation of this model shows that for a buffer size in the order of 50-100 cells the load should not exceed a value of about 85%. This result can be taken as a simple approximation to characterize the cell level effects.

For the characterization of the burst level effects two types of multiplexing strategies can be distinguished: rate envelope multiplexing (REM) considering the multiplexer to be bufferless and rate sharing where buffering is used to share the available rate among the connections (Roberts et al., 1996).

The switch described in Section 2 does not provide large buffers for VBR traffic which would be mandatory for queueing bursts. Therefore, in the following we will focus on the REM paradigm only. For this strategy the cell loss probability B is given by only considering the overload states where the aggregate rate $R(t) = R_1(t) + \dots + R_N(t)$ of all N connections exceeds the service rate C :

$$B = \frac{E[(R(t) - C)^+]}{E[R(t)]}. \quad (1)$$

This leads to a convolution of the rate distributions of the single connections. We assume on-off traffic sources with parameters m_i (mean rate) and h_i (peak rate). In the

homogeneous case where all sources have the same parameters m and h the convolution simplifies to a binomial distribution:

$$B = \frac{1}{m \cdot N} \cdot \sum_{i=\lfloor C/h \rfloor + 1}^N \binom{N}{i} \cdot \left(\frac{m}{h}\right)^i \cdot \left(1 - \frac{m}{h}\right)^{N-i} \cdot (i \cdot h - C). \quad (2)$$

CAC methods which directly use the cell loss probability result have to evaluate the convolution or binomial distribution in real-time. Each time a new connection request occurs, the cell loss probability including the new connection must be computed and compared with the CLR objective. As this is a very time-consuming task for larger values of N , using effective bandwidths appears to be an appropriate method.

The idea is to calculate an effective bandwidth c_i for each connection, so that the CLR requirements are met as long as the following equation holds:

$$\sum_i c_i \leq C \quad (3)$$

The difficulty here is, however, to find an appropriate expression for the effective bandwidth. In the case of homogeneous on-off sources a solution can be obtained by numerically inverting equation (2) to find the largest number of connections \hat{N} maintaining a cell loss probability still below the CLR objective. The effective bandwidth is then given by $c = C/\hat{N}$.

Various approaches have been proposed that give an approximation of the effective bandwidth (Gu erin, 1991; Kelly, 1991; Lindberger, 1991; Lindberger 1994, Sykas et al., 1992). Among them the solution presented by Lindberger (1994) which is also described in (Roberts et al., 1996) is promising as it yields a simple formula for computing the effective bandwidth based on the source parameters m and h , the service rate C and the CLR objective \hat{B} :

$$c = \begin{cases} a \cdot m \cdot (1 + 3 \cdot z \cdot (1 - m/h)), & z \leq \min(1, h/3m) \\ a \cdot m \cdot (1 + 3 \cdot z^2 \cdot (1 - m/h)), & 1 \leq z^2 \leq h/3m \\ a \cdot h, & \text{otherwise} \end{cases} \quad (4)$$

$$\text{with } a = 1 - \frac{\log \hat{B}}{50} \text{ and } z = \frac{-2 \cdot \log \hat{B}}{C/h}.$$

In Figure 3 the results of this approximate formula are compared with those obtained by numerically inverting the binomial distribution. The approximation matches quite well at least in the region of large values of C/h . Figure 3 also reveals that the effective bandwidth approaches the mean rate with increasing service rate. The multiplexing gain which can be expressed by $G = h/c$ increases with the service rate C up to the burstiness h/m of the source.

While in the homogeneous case the maximum number of connections that can be accepted only depends on the source parameters and the service rate, it is also influenced by the actual traffic mix in the heterogeneous case. This can be described by a multi-dimensional acceptance boundary. The typical shape of such a curve is concave. That means the effective bandwidths are higher compared to the homogeneous case if connections of different types are mixed together.

The acceptance boundary may be linearly approximated by assuming that the effective bandwidths are independent of the traffic mix. For this approximation the acceptance region is specified by the condition

$$N_1 \cdot c_1(m_1, h_1, C, \hat{B}) + N_2 \cdot c_2(m_2, h_2, C, \hat{B}) \leq C \quad (5)$$

where N_1 and N_2 are the number of type 1 and type 2 connections, respectively. The effective bandwidths c_1 and c_2 correspond to the values at the intersections with the axes. They can be computed either exactly or using an approximation, e.g. Lindberger's formula.

The linear approximation is quite good if the parameters of the source types do not differ too much. However, if for example CBR connections are mixed together with very bursty VBR connections the divergence of the acceptance boundaries is considerable (Figure 4). That means the linear approach underestimates the cell loss probability in a non-negligible way.

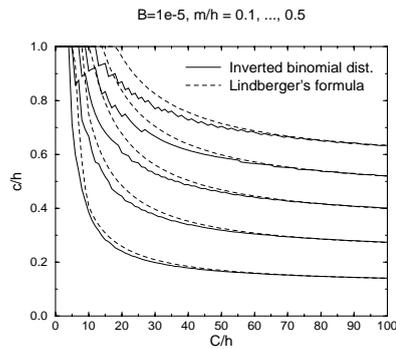


Figure 3 Effective bandwidth

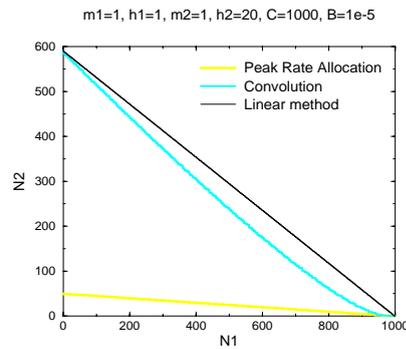


Figure 4 Boundary for two source types

3.3 Effective Bandwidths for Heterogeneous Traffic

As described in the previous section the linear CAC method is a simple and reasonably exact approximation as long as the multiplexed sources are not too different. Therefore the first idea is to separately handle CBR traffic (including also the MCR share of ABR connections), which causes most of the difficulties when put together with bursty VBR traffic, and to apply the linear method only to VBR connections.

For the computation of the effective bandwidth of a VBR connection with mean rate m_i and peak rate h_i a slightly modified version of Lindberger's formula is used where the effective bandwidth is limited to h_i . Instead of the full service rate C a reduced rate $C_{red} = C - C_{CBR}$ is given as parameter with C_{CBR} being the aggregate bandwidth of all CBR connections. Simply reducing the capacity available for VBR connections leads to an overestimation of the VBR loss probability as then the CBR traffic is assumed to be loss-free. This can be corrected by a modified CLR objective

$$\hat{B}^* = \hat{B} \cdot \frac{M_{VBR} + M_{CBR}}{M_{VBR}} \quad (6)$$

as parameter in the effective bandwidth formula. M_{CBR} and M_{VBR} are the aggregate mean rate of the CBR and VBR connections, respectively.

Partitioning the service rate into a VBR and a CBR share allows to handle the extremely concave behaviour in the acceptance boundary caused by CBR traffic. However, among the VBR connections there might still be differences which cause similar effects. Therefore, partitioning is extended to VBR connections by dividing the VBR source parameter space into K regions. An example with four VBR regions distinguished according to different values of burstiness and mean rate is depicted in Figure 5.

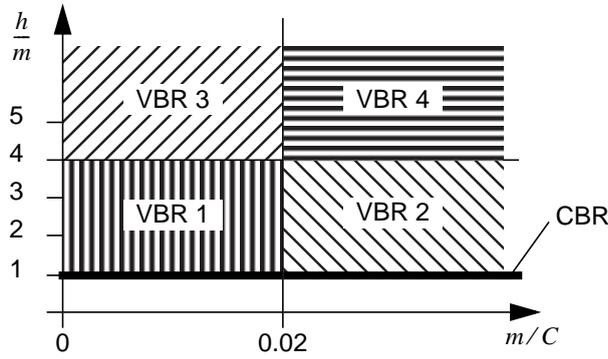


Figure 5 Example of VBR partitioning with four regions

The linear CAC method is now applied to connections of different types belonging to the same region. On the other hand, if two source types are located in different regions i and j , a reduced service rate $C_{red,i}$ is taken as a parameter for the effective bandwidth formula for connections in region i :

$$C_{red,i} = C - C_{CBR} - \sum_{\substack{k=1 \\ k \neq i}}^K C_{VBR,k}. \quad (7)$$

The service rate is reduced by the aggregate CBR bandwidth C_{CBR} and by the aggregate effective bandwidths of all VBR connections associated to regions different from i . A correction of the CLR objective corresponding to equation (6) is also performed:

$$\hat{B}_i^* = \hat{B} \cdot \frac{\sum_k M_{VBR,k} + M_{CBR}}{M_{VBR,i}}. \quad (8)$$

An alternative approach defined by Villen-Altamirano and Sanchez-Canabate (1997) is to subtract only the aggregate mean rates of the VBR connections in the other regions:

$$C_{red,i} = C - C_{CBR} - \sum_{\substack{k=1 \\ k \neq i}}^K M_{VBR,k}. \quad (9)$$

The partitioning of VBR sources into different regions depends among other parameters on the link rate and the application types using the system. Therefore it is a difficult task to find an optimal partitioning strategy. In general, it is important to keep the number of regions as small as possible to maintain efficiency on the one hand and to assign very different VBR sources to different regions on the other hand.

3.4 Proposed CAC Algorithm

If equation (7) is applied, variables C_{CBR} , M_{CBR} , $C_{VBR,k}$ and $M_{VBR,k}$ ($k = 1, \dots, K$) have to be provided by the CAC. When a VBR connection request with parameters m_{new} and h_{new} arrives the following procedure is performed:

1. Determine the region k_{new} for the new connection by evaluating m_{new} and h_{new} .
2. Compute the reduced service rate $C_{red,k_{new}}$ for region k_{new} according to equation (7).
3. Calculate $\hat{B}_{k_{new}}^*$ applying equation (8).
4. Compute the effective bandwidth $c_{new} = c(m_{new}, h_{new}, C_{red,k_{new}}, \hat{B}_{k_{new}}^*)$ using Lindberger's formula.
5. If $c_{new} + C_{CBR} + \sum_k C_{VBR,k} \leq C$ accept the connection, refuse it otherwise.
6. If the connection could be accepted, update $C_{VBR,k}$ and $M_{VBR,k}$ ($k = 1, \dots, K$).

The parameters h_{new} and m_{new} of a VBR connection correspond to the negotiated traffic contract parameters PCR and SCR, respectively. This rather conservative approach has been chosen to guarantee robustness.

The update at the end of the acceptance procedure (6.) comprises the incrementation of $M_{VBR,k_{new}}$ and $C_{VBR,k_{new}}$ by m_{new} and c_{new} , respectively, but it is also necessary for the corresponding variables of the other regions. This is because $C_{red,k}$ and \hat{B}_k^* change if a connection in any other region is added or removed. So the effective bandwidths of all existing connections have to be recalculated and summed up for each region to obtain $C_{VBR,k}$. This procedure has to be done iteratively until the variables $C_{VBR,k}$ converge to their exact values. In practice, however, it is sufficient to perform only a few iteration steps. If a CBR instead of a VBR connection is set up, an update in the same way is necessary.

If a connection with parameters m_{rem} and h_{rem} is removed, a procedure similar to the one described above has to be performed. The effective bandwidth for the removed connection is calculated and subtracted from the $C_{VBR,k_{rem}}$ value for the corresponding region. $M_{VBR,k_{rem}}$ is reduced by m_{rem} . Afterwards an update of all $C_{VBR,k}$ variables is performed in the same way as for a connection setup.

As the effective bandwidth can be computed very fast with Lindberger's formula, the update is not too time-consuming for a reasonable number of connections if an effective implementation is used. Nevertheless, there is a trade-off between necessary processing power and accuracy of the effective bandwidth calculation. It is justified to leave out updates if the traffic mix has not changed too much since the last update, i.e. if the $C_{VBR,k}$ show only moderate fluctuations. An appropriate solution is to carry out the update calculations not for every connection arrival or departure but in regular intervals, either timer-triggered (e.g. every 30 seconds) or on a counter basis (e.g. after every 10 setup or release requests).

3.5 Application to the Switch Architecture

The CAC algorithm presented in the previous section is applied to several locations in the switch. The local bus is shared by all connections. Here the largest multiplexing gain can be achieved due to the high bus service rate. Furthermore, CAC is necessary at the 25 Mbit/s and 155 Mbit/s output links. For each output link in the system a dedicated CAC entity using the presented algorithm has to be provided.

Additionally, an overload can occur at the Utopia interface on the input as well as on the output side of the switch if PUs with four 155 Mbit/s ports are used. However, here the maximum aggregate input rate exceeds the service rate of the Utopia interface only by a factor of about 1.6 so that statistical multiplexing is very limited. Therefore peak rate allocation is more appropriate than a more complex CAC algorithm at this location.

The CAC controlling the traffic over the bus must also take care of UBR traffic. As shown in Section 2 the influence of UBR traffic may be reduced by adapting the credit mechanism but it cannot be completely avoided. UBR cells consume a share of the bus bandwidth which heavily depends on the configuration and the traffic pattern. This effect is considered by a limit for the bus bandwidth that can be allocated by the CAC for CBR/VBR traffic. The limit depends on the number of units in the ATM MSS and the number of credits that are regenerated in each cycle. The limit can be further reduced if the operator wishes to reserve a larger proportion of the bus bandwidth for UBR/ABR traffic.

4 PERFORMANCE EVALUATION

4.1 Switch Architecture

The switch architecture presented in Section 2 has been extensively evaluated using a simulation program based on the queueing model depicted in Figure 2. Only some issues can be addressed in this paper. We will concentrate on effects related to the credit mechanism controlling bus access while investigations about buffer dimensioning and validating the buffer control strategies shall be omitted here.

As described in Section 2.3, the credit mechanism had to be changed to prioritise CBR/VBR cells at the bus and to reduce the influence of ABR/UBR traffic. If 10 is chosen to be the maximum number of credits allocated to each PU, on average 10 CBR/VBR cells per PU can be transmitted instead of one ABR/UBR cell. For a configuration with 14 PUs half of them sending VBR traffic contributing 50% to the total offered load while the other PUs only carry UBR traffic, the results depicted in Figure 6 have been obtained. The total load offered to the bus is increased while maintaining the proportion of VBR and UBR traffic unchanged. As soon as the offered load exceeds 1, the UBR share of the bus bandwidth decreases linearly until it reaches saturation at a value of $1/11$, which is directly determined by the number of credits each PU receives during the regeneration cycle.

Obviously, limiting the bandwidth that ABR/UBR traffic may occupy to an even lower value can be achieved by further increasing the maximum number of credits. However, large values of the credit per PU may potentially lead to unfairness. Assuming symmetric traffic from sources that are uniformly distributed over 14 PUs, we get the results shown in Figure 7. For large credit numbers one can see that PUs with lower

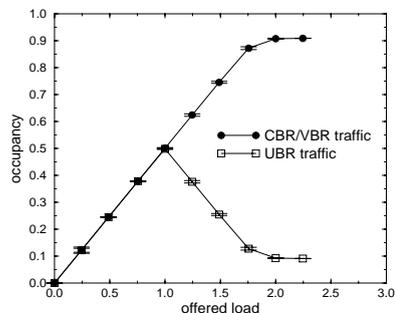


Figure 6 Influence of UBR traffic

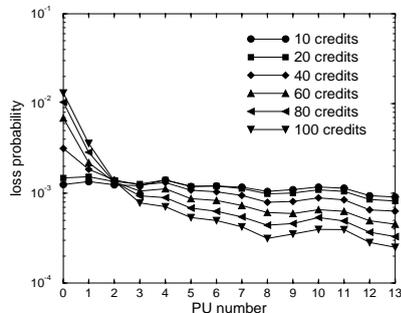


Figure 7 Influence of number of credits

identifiers suffer much more from losses. This effect is due to the fact that a credit regeneration in all PUs is performed if there is no BAC that has credits as well as cells in its input buffer. If the maximum credit value is greater than the size of the input buffer, PUs with lower identifiers cannot use up their credits because of the limited number of cells that they can accumulate until they receive the right to send. PUs with higher identifiers, however, can make better use of their credits because they can immediately send a newly arrived cell provided they still have credits. Thus, when credits are regenerated the PUs with higher numbers usually have exhausted their credits whereas the PUs with lower numbers still have credits but currently no cells to send. Hence the high priority PUs receive a larger share of the bus bandwidth which in fact reduces their cell losses. As a conclusion from that the number of credits is chosen to be in range of 10 to 20.

4.2 CAC Algorithm

The performance of the CAC algorithm presented in Section 3 heavily depends on the traffic pattern, the service rate, the CLR objective and the definition of the VBR regions. In the following we will concentrate on a traffic mix that consists of connections of two different source types:

- type 1 with parameters $m_1 = 1$ Mbit/s, $h_1 = 2$ Mbit/s and
- type 2 with parameters $m_2 = 1$ Mbit/s, $h_2 = 20$ Mbit/s.

The CLR objective is chosen to be $B = 10^{-5}$ for all case studies.

If the parameters of the source types belong to different VBR regions, partitioning is applied. The reduced capacity is calculated by either using effective bandwidths (equation (7)) or mean rates (equation (9)). In Figure 8 and Figure 9 the acceptance boundaries of both approaches are compared to those obtained by the convolution method and by peak rate allocation. For a service rate of $C = 1.37$ Gbit/s referring to the net transmission rate of the bus, the partitioning curves are rather closely tied to the ideal boundary. Here an enormous improvement compared to peak rate allocation can be achieved. For a smaller service rate of $C = 140$ Mbit/s that results as a net rate of a 155 Mbit/s output link the results are not that good. Partitioning yields acceptance boundaries that are still much better than peak rate allocation but quite far away from the ideal curve. However, this is not only the effect of partitioning but also of using Lindberger's effective

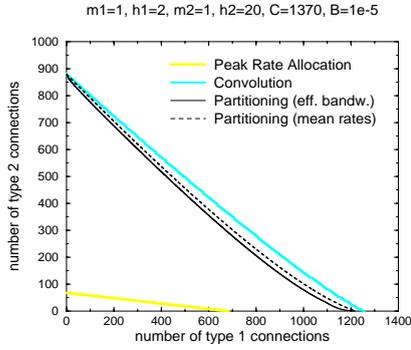


Figure 8 Boundary for $C = 1.37$ Gbit/s

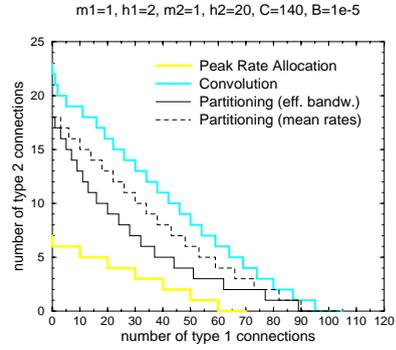


Figure 9 Boundary for $C = 140$ Mbit/s

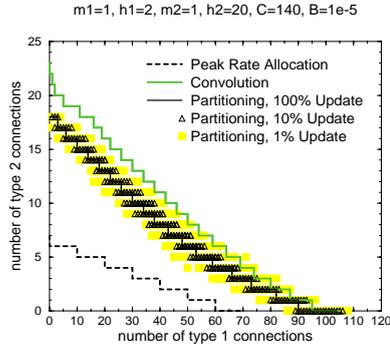


Figure 10 Update policies for partitioning using mean rates

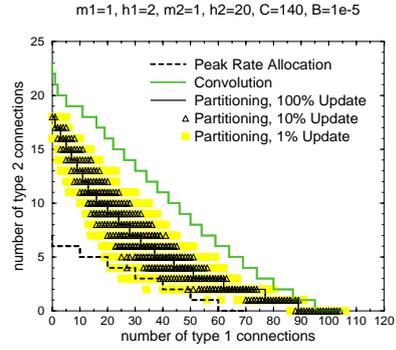


Figure 11 Update policies for partitioning using effective bandwidths

bandwidth approximation formula, as can be observed from the deviation at the axes intersections.

Comparing the two partitioning approaches one can see that the acceptance boundary for the method using aggregate mean rates for reducing the entire capacity is always closer to the ideal curve, especially for smaller service rates. However, this algorithm may be too optimistic if the ratio C/h is further reduced (Villen-Altamirano and Sanchez-Canabate, 1997). Furthermore, the effect of being quite conservative by using equation (7) may be compensated by the linear approach applied within the VBR regions.

As mentioned in Section 3.4, the update at the end of the acceptance procedure may be delayed in order to reduce the processor load caused by CAC. Two cases of performing an update only after every 10 and every 100 CAC procedure calls have been investigated. Thereby the processor load which is mainly due to the update is reduced to 10% and 1%, respectively. Simulation results in Figure 10 and Figure 11 show that the acceptance boundary curve changes to an area consisting of states in which a connection request either of a type 1 or a type 2 connection may be rejected. The rejection states

have been collected from several simulations with different loads based on Markovian arrival and holding-time processes. The number of observed rejection states increases if the update rate is reduced. This is due to the fact that the effective bandwidth calculation for a new connection has to rely on possibly out-of-date parameters for the reduced bandwidth. Therefore the acceptance decision more and more depends on the random walk produced by the connection arrival and release processes.

The effect of “spreading” rejection states is much heavier for the approach using effective bandwidths for the reduced capacity calculation. An update rate of 1% would be not acceptable as it may be even less effective than peak rate allocation. For the method based on mean rates a very low update frequency may extend the acceptance region beyond the ideal curve leading to illegal states. An update frequency of 10% seems to be a good compromise for both CAC methods.

5 CONCLUSIONS

The design of traffic management functions for an ATM switch must precisely take into account the internal switch architecture as well as QoS requirements, utilization targets and implementation effort. Here an ATM Multi-Service Switch (MSS) for workgroups is considered which is built of several port units that are attached to an internal bus. In order to support all five ATM service categories, appropriate control functions are mandatory which guarantee QoS objectives and can be realized with acceptable implementation effort.

Priority control has to ensure that access to switch resources is given according to QoS requirements of different service categories. Especially the performance of the credit mechanism for the bus access has been investigated by simulation. The results confirm the ability of the protocol to limit the impact of ABR and UBR traffic on the guaranteed services as CBR and VBR to an acceptable level. Furthermore, it has been shown that fairness between different port units can be ensured by a proper setting of the number of credits regenerated per cycle.

While CAC for CBR services becomes trivial (peak rate allocation), for VBR type of traffic the effects of statistical multiplexing have to be taken into account. We found a simple formula for the calculation of an effective bandwidth for each VBR connection to be a good compromise between accuracy and implementation complexity. To overcome the typical problems when applying the concept of effective bandwidths in heterogeneous traffic environments, a classification into VBR traffic regions is proposed. This allows to consider the mutual influences between these regions by replacing the total link capacity with the capacity which is effectively available for each region. Obviously, this is a conservative approach but the performance study reveals that the deviation from the ideal boundary is acceptable. Finally, the MCR portion of an ABR connection is treated similarly to CBR, while UBR requests are never rejected.

The control framework comprising UPC, CAC and various priority control schemes has been tailored to the internal MSS architecture. Together with some other control functions which are not described in detail (e.g., backpressure mechanisms), the MSS is equipped with powerful traffic management capabilities allowing to effectively support a broad spectrum of applications and services.

REFERENCES

- Antunes, N., Rocha, R., Pinto, P. (1997) "Analysis and Simulation of a Traffic Management Control Scheme for ATM Switches with Loose Commitments", *Conference on Communication Networks and Distributed Systems Modeling and Simulation CNDIS '97*, Phoenix, Arizona.
- ATM Forum - Technical Committee (1996) *Traffic Management Specification*, Version 4.0, af-tm-0056.000.
- ATM Forum - Technical Committee (1995) *UTOPIA Level 2*, af-phy-0039.000.
- Guérin, R., Ahmadi, H., Nagshineh, M. (1991) "Equivalent Capacity and Its Application to Bandwidth Allocation in High-Speed Networks", *IEEE Journal on Selected Areas in Communications*, Vol. 9, No. 7, pp. 968-981.
- Hui, J.Y. (1988) "Resource allocation for broadband networks", *IEEE Journal on Selected Areas in Communications*, Vol. 6, No. 9, pp. 1598-1608.
- ITU-T (1996) *Recommendation I.371: Traffic Control and Congestion Control in B-ISDN*, International Telecommunication Union - Telecommunication Standardization Sector, Geneva.
- Kelly, F.P. (1991) "Effective bandwidths at multi-class queues", *Queueing Systems*, Vol. 9, pp. 5-15.
- Kröner, H., Hébuterne, G., Boyer, P., Gravey, A. (1991) "Priority Management in ATM Switching Nodes", *IEEE Journal on Selected Areas in Communications*, Vol. 9, No. 3, pp. 418 - 427.
- Kröner, H. (1995) *Verkehrssteuerung in ATM-Netzen - Verfahren und verkehrstheoretische Analysen zur Zellpriorisierung und Verbindungsannahme*, Ph. D. Thesis, University of Stuttgart.
- Kröner, H., Renger, T., Knobling, R. (1994) "Performance Modelling of an adaptive CAC strategy for ATM networks", *Proceedings of the 14th ITC*, Antibes Juan-les-Pins, pp. 1077-1088.
- Lindberger, K. (1991) "Analytical models for the traffical problems with statistical multiplexing in ATM networks", *Proceedings of the 13th ITC*, Copenhagen, pp. 807-813.
- Lindberger, K. (1994) "Dimensioning and design methods for integrated ATM networks", *Proceedings of the 14th ITC*, Antibes Juan-les-Pins, pp. 897-906.
- Roberts, J., Mocci, U., Virtamo J. (1996), *Broadband Network Teletraffic - Final Report of Action COST 242*, Springer, Berlin.
- Sykas, E.D., Paschalidis, I.C., Mourtzinou, G.K., Vlakos, K.M. (1992) "Congestion Avoidance in ATM Networks", *Proceedings of IEEE INFOCOM '92*, Florence, pp. 904-914.
- Villen-Altamirano, M., Sanchez-Canabate, M.F. (1997) "Effective Bandwidth dependent of the actual traffic mix: an approach for bufferless CAC", *Proceedings of the 15th ITC*, Washington D. C., pp. 47-57.