Simplified management of ATM traffic

Marko Luoma, Mika Ilvesmäki

Helsinki University of Technology
Laboratory of Telecommunications Technology
Otakaari 5A, 02150 Espoo, Finland

ABSTRACT

ATM has been under a thorough standardization process for more than ten years. Looking at it now, what have we achieved during this time period? Originally ATM was meant to be an easy and an efficient protocol enabling varying services over a single network. What it is turning to be is 'yet another ISDN'-- network full of hopes and promises but too difficult to implement and expensive to market. The fact is that more and more 'nice features' are implemented on the cost of overloading network with hard management procedures. Therefore we need to adopt a new approach. This approach keeps a strong reminds on 'what is necessary.' This paper presents starting points for an alternative approach to the traffic management. We refer to this approach as 'the minimum management principle.' Choosing of the suitable service classes for the ATM network is made difficult by the fact that the more services one implements the more management he needs. This is especially true for the variable bit rate connections that are usually treated based on the stochastic models. Stochastic model, at its best, can only reveal momentary characteristics in the traffic stream not the long range behavior of it. Our assumption is that ATM will move towards Internet in the sense that strict values for quality make little or no sense in the future. Therefore stochastic modeling of variable bit rate connections seems to be useless. Nevertheless we see that some traffic needs to have strict guarantees and that only economic way of doing so is to use PCR allocation.
Keywords: ATM, CBR, Congestion Control, Resource Allocation, Traffic Measurement

1. INTRODUCTION

The key issue for successful implementation of a new network and a protocol is the fact that it should provide something better than earlier networks. In order to accomplish this, ATM must provide an easy and scaleable platform to support a variety of services. 'Easy' in this context means that the threshold for introducing new concepts and services should be low. 'Scaleable' in this context means that all services, ranging from few kilobits per second to terabits per second, may be introduced at the same time.

Designing a new network protocol, like ATM, is always inspiring work. Many people want to leave their mark in the history. This can be observed in the ATM Forum's contributions list which has had the tendency to double every three years. This 'enthusiasm' can lead to complex solutions resulting in inefficiency as real-world applications are implemented. Sometimes the mistake of having designed an unnecessarily complex protocol can easily be repaired; but tedious work and a strong effort are needed. One of the most commonly known mistakes of this kind has been the implementation of ATM adaptation layers (AAL) 3 and 4. As the work went further and further on, more and more features were introduced making the result intolerable for wide usage. This was soon realized by the community and a strong effort to simplify their structure started - leading to AAL5, which is ***the*** AAL nowadays.

Same kind of development can now be observed with traffic management. The ATM Forum Traffic Management Specification v.4[1] introduced many fine concepts, like available bit rate service (ABR), but nevertheless it is becoming a giant that might die due to its internal complexity and impossibility. Therefore a new approach needs to be taken. In this approach a strong reminder of 'what is necessary' needs to be kept in mind.

Communications services can be examined from different viewpoints: as the viewpoint changes, the appreciated values in a service change. Normally one can look from the viewpoint of a customer but hardly ever catch the viewpoint of the operator.

Further author information -
M.L: (correspondence): Email: marko.luoma@hut.fi; WWW: http://keskus.hut.fi/~mluoma; Telephone: + 358 9 451 2463;
Fax: +358 9 451 2474
M.I: Email: lynx@luuri.hut.fi; WWW: http://keskus.hut.fi/tutkimus/ipana/mika.shtml; Telephone: + 358 9 451 4784;
Fax: + 358 9 451 2474

A customer, who is always the one who decides whether a new service introduced by the operator will succeed, has a high value on issues like ease of use, price and availability. These values are heuristic -- they do not imply any real implementation neither they form mathematical formula to solve. More or less the operator offering the service has to have the capability to adjust the service based on the response from the customers. The customers can be divided roughly into two categories: business and residential. Residential users, the big mass, are purely interested in the values described before but also want something radically new and valuable from the new service in order to invest their money on it. Business customers on the other hand are more willing to experiment with new concepts if there is a potential benefit for them on the market; whether it be a strategic or an operative one.
The operator's viewpoint to the new service and the network is completely different. The operator is more interested in whether a new investment will extend the life cycle of previous investments. After that they also want to tune network so that they get as high revenue as possible with customer still satisfied. The satisfaction of the customer can be observed from the fluctuation in the number of customers. To maximize revenue the operator may overbook the network far too much which unquestionably leads to customer dissatisfaction.

2. TRAFFIC MANAGEMENT

Traffic management as a term is quite complex to formulate. Traffic management is the overall mechanism that makes it possible to offer reliable services on top of a transmission infrastructure. Reliable means that services are available whenever used and are offered on the controlled manner. Control functions time scale (average time between subsequent executed functions) ranges from microseconds to years.
The key issue in traffic management is the way in which the problem is approached. One can implement the so called socialistic approach where individual users are treated fairly and network relies on the good behavior of individual users. This means that individual users are responsible to adapt their transmission to the network conditions. The other opportunity is the capitalistic approach where users are treated on the basis of their financial capability. This means that the network executes all of the control functions, allocates resources unevenly to the users and charges the users based on the used capacity or some other policy. The fact has been that the Internet relies on good will of all users whereas telephone networks are totally self controlled entities; they do not presume anything from the users. This approach has also had implications to the services. The Internet services have developed to be rather flexible, since no actual guarantees about the parameters of transmission channel have existed. Telephone services, with full guarantees of the transmission channel, have developed to be more strict.

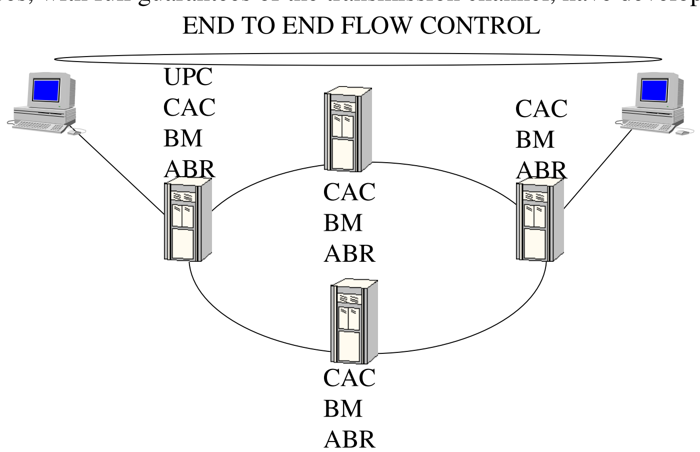

Figure 1: Different possibilities for traffic control.

Traffic control in ATM networks can be categorized to two classes: preventive and reactive methods. Preventive methods try to predict the possibility of congestion and take actions before congestion really occurs. Preventive methods presented before are: connection admission control (CAC), usage parameter control (UPC), buffer management (BM) and available bit rate service (ABR). The only reactive method is also buffer management (BM). CAC is the only method operating at the level of connection management and the others are basically link management procedures.

3. CONNECTION MANAGEMENT

Connection management deals with issues like connection admission control (CAC) and service classes. The basic idea behind packet based networks is the dynamic possibility to share resources e.g. buffers and bandwidth. Bandwidth is usually overbooked in the network by a factor from 2 to 20, depending how far in the core we are. This overbooking is done because many services are naturally variable bit rated and this kind of statistical multiplexing gives higher revenues for the operator.

3.1 Service classes

Traffic in ATM network is a combination of heterogeneous traffic from all of existing and forthcoming services. These services have widely different characteristics when considering tolerance of delay and cell losses or rate and variation of traffic they are emitting. These differing combinations of service demands and typical applications are presented in figure 2.

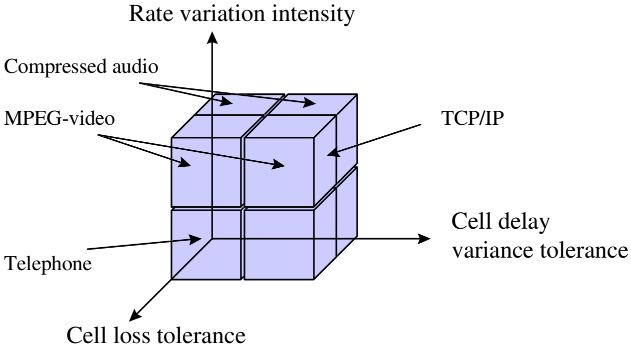

Figure 2: Characteristics of applications

Figure 3 shows that applying only one service type operator can not meet all the demands of the applications and a high utilization of the network at the same time in controlled manner. As a consequence, the network operator must either adopt more services or accept higher cell loss ratio, larger delays, lower utilization or acceptance of variable bit rate traffic.

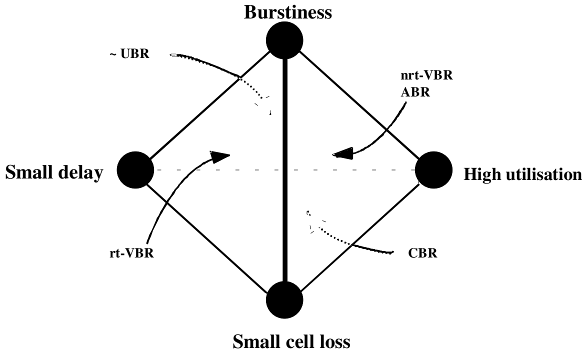

Figure 3: Different service components, their requirements and efficiency[2]

On the other hand, choosing the suitable service classes for an ATM network is made difficult by the fact that the more services one implements the more management he needs. This is especially true for the variable bit rate (VBR) services which are usually treated based on stochastic models. These stochastic models try to transfer the key characteristics of the actual information flow to two or three parameters. The other side of the coin is the constant bit rate (CBR) service, which is fully expressed with one parameter. Because there is no general stochastic model for every VBR source and because usually stochastic models can only reveal momentary characteristics in the traffic stream and not the long range behavior of, e.g. a video stream. We believe that it should be adequate to consider only cell losses and delays, because the rate variations can either be buffered or discarded depending on the way we want to treat them.

		Cell Loss Probability	
		Small	Large
Delay	Small	CBR with rate enforcement	NBR+
	Large	NBR	UBR

Figure 4: Proposed service classes and their qualities

To accommodate the most viable services in a network only two service classes are needed: NBR and CBR. Nominal bit rate (NBR) is for variable bit rate connections capable of handling cell losses and for non-real-time applications like Internet applications. Constant bit rate (CBR) is for real-time applications. The nominal bit rate service is divided in three subsections where each section has a slightly different goal.

1. CBR with rate enforcement; CBR with rate enforcement is for the applications with critical timing and coding. These applications can not tolerate cell losses nor delays. Therefore they have to be treated as separate as possible from the other traffic in the network. To guarantee a small cell loss, the sources using this service must have traffic shaping capabilities.
2. NBR+; NBR+ is for the applications that require a minimum capacity, small delay but can tolerate cell losses. NBR+ is one way to send variable bit rate traffic. One can make the contract for a rate that one expects to be suitable and then transmit even more, but facing the risk of a severe cell loss due to small buffers.

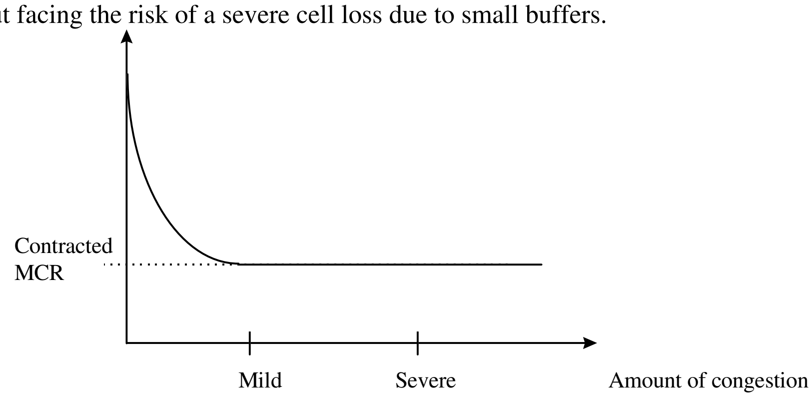

Figure 5: Received bandwidth as a function of congestion level

3. NBR; NBR is for the applications that require a minimum capacity and tolerate cell losses. NBR is also a way to send variable bit rate traffic: One can make the contract for a rate that one expects to be suitable and then transmit even more on the penalty of losses.

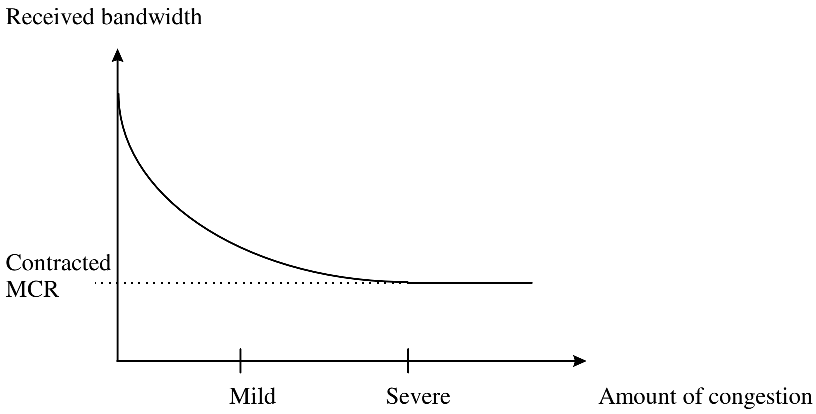

Figure 6: Received bandwidth as a function of congestion level

4. UBR; UBR equals to NBR when MCR is set to zero.

As seen from above, there is no variable bit rate category. If one wishes to use a variable bit rate connection, one has two options:

1. Long queuing delays, if relatively bursty (third case).
2. Shape the traffic so that it falls in the second or first category depending on the loss requirements.

3.2 Connection Admission Control

3.2.1 Meaning of connection admission control

The main task for the connection admission control (CAC) is to decide whether a new connection, if accepted, deteriorates ongoing connections. This decision is made based on the knowledge of ongoing connections and advertised parameters of the new connection. The simplicity of the implementation including parameter calculation and routing aspects is the main requirement for a practical CAC method. Parameter calculation means the tasks which the management computer needs to perform in order to transpose advertised parameters to suitable parameters for implemented CAC formula. These advertised parameters are called the traffic descriptor. A typical traffic descriptor includes following parameters: the sustainable (mean) cell rate (SCR), the intrinsic burst tolerance (IBT e.g. MBS) and the peak cell rate (PCR). The parameter, that is typically used by the CAC algorithm, is the effective bandwidth (EB). Effective bandwidth is an imaginary value which should approximate the required bandwidth for the connection. The main difficulty is related to the conversion of the traffic descriptor to the effective bandwidth. Traffic models used in the conversion are usually tied to a certain time scale and are not valid on the other scales. Therefore huge mistakes can be done in the approximation of the required bandwidth.

3.2.2 The problem of variable bit rate connections

Variability in the transmitted bit rate is natural for many applications in telecommunications. Variability comes either from the source of the information (packet segmentation process in data applications) or from the compression algorithm (video applications) used to cut down the expenses of the communication channel. From the network point of view if some guarantee needs to be set, smooth traffic is far easier to handle whereas bursty traffic requires more management resources. In the present methodology it is assumed that the user or the application determines suitable traffic parameters for the connection and signals them to the network. For smooth (CBR) connections this task is easy to accomplish but what then, is the case for the bursty applications. Following figures represent some traffic profiles taken from the MPEG-1 coded videos [3]. As these pictures show, video in the open loop coding, produces traffic that, when produced online, cannot be estimated beforehand by the values of PCR, SCR and IBT. If played out from the server it could be possible to include additional information about required traffic contract.

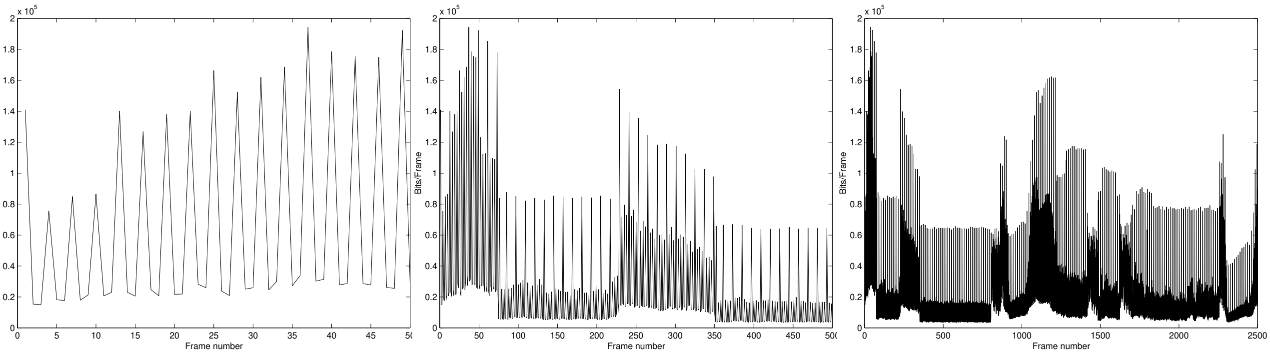

Figure 7: Amount of data in single I, P or B frame. Time of single frame is 33ms.

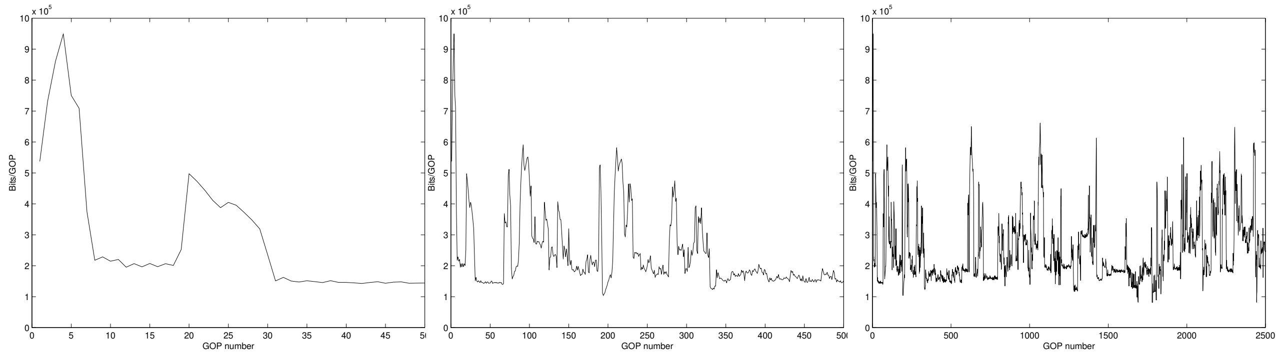

Figure 8: Amount of data in single group of pictures (GOP). GOP consists all frames (12) between two consecutive I-frames. Time of single GOP is 400ms.

How can this problem then be solved? We can either make a rough estimation that does not tie the user and do some measurement of the actual traffic distribution or dump the whole CAC and operate on pure congestion control.

3.2.3 Connection admission control based on traffic measurement

To cope with the uncertainty in parameter calculations other approaches have been introduced. Saito[4] has introduced approaches that use traffic measurements as a filter for the uncertainty. He has introduced dynamic and hybrid approaches for CAC both utilizing traffic measurements. The dynamic CAC model is a non-parametric CAC model that uses the distribution of the arriving cells tuned by the traffic measurement. This kind of CAC is suitable for high values of statistical multiplexing as in data traffic but not so suitable for the high quality traffic. Saito has proposed the hybrid CAC model to overcome this problem. The hybrid CAC uses peak rate allocation for the high quality VC's and dynamic CAC for the low quality VC's. Kilkki has proposed the SIMA[5] concept where no actual CAC is needed when resource allocation and congestion control are executed based on real time traffic measurements. A measurement scheme based on the exponential moving average is used to measure the actual bit rate of a connection. This measured bit rate (MBR) is compared to the contracted nominal bit rate (NBR) in order to set priority for each cell transmitted further from the access switch. Priorities 1 to 7 are set based on the MBR to NBR ratio. A cell is discarded as a function of the buffer occupancy and the priority that it has.

3.2.4 Proposed CAC technique

The viewpoint of the user to the received bandwidth is usually heavily dependent on the application the user is running. Some of the applications are able to provide better service when they are able to receive more bandwidth. Typically one can characterize applications by how elastic they are.

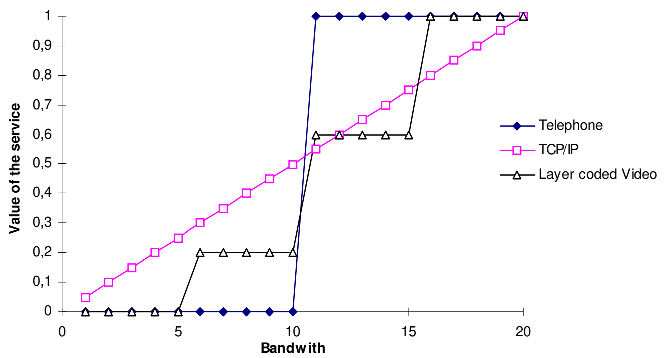

Figure 9: Value of the service for the application versus bandwidth available to the application

Figure 9 shows some typical applications and their value curves. Rigid applications (value curve equals to the step function) are typically interactive applications whose coding is based on the deterministic quantization. Elastic applications (value curve equals to the smooth line) are typically data applications with flow control. To be able to fulfill requirements of value curves of the applications, the connection admission principle should allow some initial bandwidth for each user. For excessive bandwidth, an elastic approach is applied. Our proposal has therefore equivalencies to both of the measurement based approaches presented before. We base our approach to the assumption that ATM should move towards Internet in the sense that strict values for the quality make little or no sense in the future as in the SIMA concept. Also we see that some traffic needs to have strict guarantees and that the only way of doing this economically is to use the PCR allocation.

4. LINK MANAGEMENT

Congestion control and resource allocation are linked by the fact that the usage of the other affects to the other. Theoretically it is possible to allocate network resources so that congestion does not exist but this way of operating is not, however, practicable. On the other extreme there is a scheme that has no resource allocation and the network operates purely on congestion control. In practical networks both resource allocation and congestion control are needed.

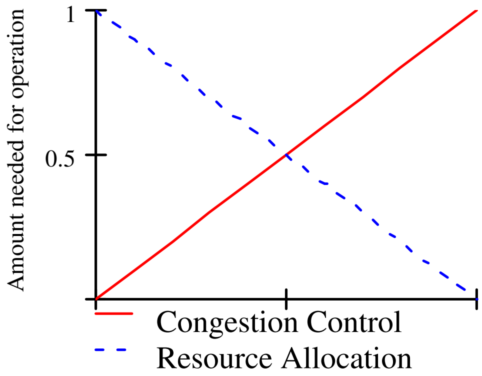

Figure 10 Resource allocation versus congestion control; trade off between two complexities.

Statistical multiplexing gives rise to the possibility of an overload in the link and can therefore, if persisting, cause switch buffers to overflow. This overflow reduces the performance of the network and therefore has to be done in a controlled way. Key issues in the buffer management are the adequate buffer size and an efficient control algorithm.
For the guaranteed traffic buffer requirements are relatively straightforward to determine. For the best effort traffic the requirements are more based on the control we are implementing. The issue of what is the best suited algorithm for the cell discarding is, is more or less a matter of an opinion. Several relatively equal algorithms exist. Some of them are based on the actual queue occupancy and some on the average occupancy. We believe that some of these methods suit quite well for one type of the traffic but to handle different services different algorithms should coexist.
Congestion control is a vital block in the network devices. The main task for the congestion control is to ensure that the network resources are used efficiently and fairly among different (competing) users. This fairness is the key issue in the Internet that has no guaranteed provisioning of the network resources. Also when we speak about ATM we must also deal with the fairness issue. Fairness in this case comes to the picture in the buffering systems of an ATM switch.

	FBA	*NBA*	*EPD*	*RED*
UBR	X			X
IP/Video over CBR			X	
NBR		X		

Figure 11: Combinations of the buffer management algorithms and the services that they are best suited for

Short descriptions of the previous algorithms are presented on the following chapters.

4.1 Early Packet Discard

Early packet discard (EPD) was proposed by Allyn Romanow[6] in 1994. Early packet discard is a scheme where the switch monitors the instantaneous queue length *[X]*. After a certain predetermined threshold level *[R]*, the switch starts to discard entire packets - not individual cells from multiple packets. This method is predictive in the manner of operation. It tries to avoid the congestion by starting the cell discarding before the buffer is full. In this way there are no corrupted packets in the buffer and no bandwidth is wasted. To make sure that the control algorithm does not oscillate or at least the frequency is lower some hysteresis *[r]* needs to be implemented.

$$X_{+} \geq R \vee X_{-} \geq R - r \quad (1)$$

4.2 Fair Buffer Allocation

Fair buffer allocation (FBA) was proposed by Juha Heinänen and Kalevi Kilkki[7] in 1995 [4]. Fair buffer allocation adapts the fundamental idea of EPD with the key issue of the best effort service - fairness. In the fair buffer allocation the switch monitors both the instantaneous queue lengths of an individual connection *[Y_i]* and the aggregate value of the buffer usage *[X]*. Decision of whether a cell should be discarded is based on two rules; the buffer occupancy is exceeding the threshold level *[R]* and the connection is using more than its fair share *[W_i]* of the resources. In the calculation of the fair share only the active connections *[N_a]* i.e. those having at least one cell in the buffer, are taken in to the account. To smooth operation in a boundary area a smoothing factor is calculated based on the buffer size *[K]*, the buffer occupancy and the threshold level.

$$X \geq R \wedge \begin{cases} W_i = \dfrac{Y_i N_a}{X} \\ W_i > Z\left(1 + \dfrac{K - X}{X - R}\right) \end{cases} \tag{2}$$

4.3 Random Early Detection

Random early detection (RED) was proposed by Sally Floyd and Van Jacobson[8] in 1993. RED implements a two parameter approach to congestion control. First it measures the average queue size by using an exponential decay filter and compares it to the fixed threshold values. If the threshold is smaller than the average queue size, a packet discard probability is calculated and marking is executed based on the calculated probability. If the upper threshold is exceeded all packets are marked.

$$avg_t = \alpha q_t + (1 - \alpha) avg_{t-1} \tag{3}$$

Random early detection is based on an assumption that the fairness among competing connections can be achieved if connections are treated randomly. By this assumption it can be claimed that the faster connections lose more cells than the slower ones.

4.4 Nominal Buffer Allocation

Our buffer management scheme for the NBR class is the nominal buffer allocation (NBA). Nominal buffer allocation operates fundamentally identically with the FBA but it relates the fair share of the resources to the contracted cell rate. A cell or a packet is discarded if and only if the queue length exceeds the threshold and the connection uses more that its fair share of the buffer capacity, based on the ratio of the used capacity versus the allocated nominal capacity.

5. CONCLUSIONS

ATM traffic management is facing a big dilemma, whether to pursue ahead or to turn around and see where it did go wrong. We believe that the time to turn around has come. To implement an efficient ABR, CAC, buffer management and all service classes that the TM4[1] expects is far too complex a task. The strength of ATM lies in its simplicity in the link layer not in the complexity of the traffic management. We have suggested a simplified approach to the traffic management and believe that further research under the goal of a simple and an efficient traffic management needs to be done.

Traffic management is the key feature when very high speed networks start to grow in a way they are expected to grow. The new services and their characteristics will very much determine which of the traffic management functions will be implemented and in what way. Our opinion is that the high speed networks will eventually evolve towards the Internet and more elastic applications. Therefore simpler and more straight forward mechanisms need to be implemented. Nominal bit rate service based on the minimum cell rate reservation seems to be the right goal. We have presented first ideas about these functions and will pursue in order to formulate these functions more precisely.

If harmonizing of traffic management functions is not done, would you really like to manage a nation wide network?

ACKNOWLEDGEMENTS

This work was supported by Finnish Academy of Research under contract 34097 (MITTA).

REFERENCES

1 ATMF, Traffic Management Specification Version 4.0, April 1996.

2 J. Virtamo, Inauguration speech, 18.9.1995, Otaniemi, Espoo, Finland.

[3] O. Rose. "Statistical properties of MPEG video traffic and their impact on traffic modelling in ATM systems". Proceedings of the 20th Annual Conference on Local Computer Networks, Minneapolis, MN, pp. 397-406,1995.

Traces: **URL:** ftp://ftp-info3.informatik.uni-wuerzburg.de /pub/MPEG/

[4] H. Saito, *Teletraffic technologies in ATM networks*, p.174, Artech House, Inc., Norwood, 1994.

[5] K. Kilkki, Simple Integrated Media Access (SIMA), **URL:** http://ds.internic.net/internet-drafts/draft-kalevi-simple-media-access-01.txt, 18.6.1997.

[6] A. Romanow, S. Floyd, “Dynamics of TCP Traffic over ATM Networks”, SIGCOMM '94, pp. 79-88, August 1994.

[7] J. Heinänen, K. Kilkki, A Fair Buffer Allocation Scheme, submitted for publication, 1995.

[8] S. Floyd, V. Jacobson, “Random Early Detection gateways for Congestion Avoidance”, IEEE/ACM Transactions on Networking, V.1 N.4, p. 397-413, August 1993.