Network Working group
<draft-kalevi-simple-media-access-00.ps>
Internet-Draft
Expire in 20th September

K. Kilkki
Nokia Research Center
March 1997

Simple Integrated Media Access (SIMA)

Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

To learn the current status of any Internet-Draft, please check the "1id-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), ftp.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

Distribution of this memo is unlimited. Please send comments to kalevi.kilkki@research.nokia.com.

Abstract

The basic objectives of future Internet are to increase the network capacity, to offer a practical real-time service, and to develop a feasible charging scheme. These objectives introduce very strict requirements for the traffic control system. This paper presents a new simple approach for traffic management: Simple Integrated Media Access (SIMA) service. According to the SIMA concept each customer shall define only two issues before a connection establishment: a nominal bit rate (*NBR*) and the selection between real-time and non-real-time service classes. *NBR* has two purposes: it forms the basis of charging, and it defines how the network capacity is divided among different connections during overload situations. Simplicity means that, on the one hand, the network operator

does not guarantee the continuous availability of nominal bit rate, and on the other hand, the user is allowed to send data with any bit rate independently of the *NBR*. However, the bit rate of transmission defines the cell loss probability in the case of network congestion. In this way a simple, but effective, self-regulation of traffic can be realised.

Table of Contents

Abstract 1
1. Introduction 3
2. Simple Integrated Media Access 4
3. Viewpoints 7
3.1 Network operator viewpoint 7
3.2. Customer viewpoint 8
3.3. Manufacturer viewpoint 9
4. Advanced features 9
4.1 Interoperability with ordinary ATM networks and services 9
4.2 A measuring scheme using exponential moving average 10
4.3 A cell scheduling scheme 12
4.4 A packet discarding scheme 12
4.5 Priority feedback for controllable connections 13
4.6 Priority gain for connections with low average rate 14
4.7 A charging scheme 15
5. Performance evaluation 16
5.1. Performance evaluation with independent traffic process 16
5.2. User reactions to quality differences 20
5.3. Performance of controllable connections 21
6. Conclusions 23
Author’s address 24
Expiration 24

1. Introduction

The Internet is at a phase of great changes. There are several stringent new requirements for the network because of two reasons: the invasion of new users, and the rapid development of new applications. These requirements mean that network capacity must rapidly be increased, real-time service has to be fundamentally improved, and a feasible charging scheme must be introduced. Asynchronous transfer mode (ATM) technology may give some useful answers: it enables flexible increase of capacity and it supports the integration of real-time and non-real-time applications. However, there are fundamental problems with the conventional ATM technology mainly related to it's complexity.

Firstly, a simple and clear charging scheme is a mandatory component of a feasible network service. Unfortunately, as the total system of ATM traffic management is extremely complicated, a feasible and understandable charging scheme that takes into account all the necessary things (service classes, traffic parameters, QoS parameters, etc.) seems to be almost impossible to construct. This is especially evident if we take into account the reluctance of an ordinary customer to learn all the technical subtleties of telecommunication networks.

Secondly, any ATM network will be very difficult to construct, manage and control if all the properties specified in different recommendations are implemented and offered to customers. A complex traffic management system will also increase the probability of malfunctions. In particular, the traffic management of an advanced ATM network calls for high quality expertise and practical experience in that particular area, and we cannot suppose that all new network operators are willing to acquire the human resources needed for that purpose. In the reverse, it is much more likely that a lot of carriers would prefer to throw bandwidth to the problem of managing QoS, as John McQuillan has expressed it in Broadband Networking News, Nov. 26, 1996. But if that is the reality, what is the significance of all the traffic management functions defined in numerous ATM specifications?

Consequently, it will be very difficult to satisfy the requirements for low and reasonable tariffs expected by customers and high incomes expected by operators due to high operation costs. With Internet where customers are accustomed with low costs, this aspect is of the utmost importance. Therefore, there seems to be an actual demand, and opportunity, for a total re-thinking of the traffic management scheme of Internet even if we apply some parts of ATM specifications.

2. Simple Integrated Media Access

In order to obtain a practical traffic management system, we shall keep in mind the real, indispensable requirements for the future Internet:

1. Network operators should be able to offer simple and understandable service for ordinary customers.
2. The network should fulfil the requirements of both real-time and non-real-time applications.
3. A fair charging and accounting scheme should be simple to implement.
4. Network, including traffic management should be simple, in order to keep managing costs on a reasonable level.
5. Network nodes should not be too complicated in order to obtain high capacity, especially in the core network.

As stated previously, ATM might be the key technology to be used for fulfilling these requirements. The primary strengths of ATM are the possibilities to make very fast and large switches, and to guarantee short delays for real-time, interactive connections. If we take seriously the above 5 requirements for Internet, we shall, definitely, not construct such a complex system as the current ATM traffic management. This document shows that both of the basic advantages of ATM can be obtained without the greater part of complicated ATM specifications by using a new service specification: Simple Integrated Media Access (SIMA).

The starting point for the development of SIMA service model has the basic properties of ATM with additions of 8 priority levels, and a new concept called Nominal Bit Rate (NBR). NBR has two main tasks: the charging is based entirely on it, and the division of network capacity between different connections is determined using the NBR of each connection.

The charging part of the service is very simple: in the simplest model each user has a permanent NBR, e.g., 100 kbit/s, and the monthly fee is proportional to that value. If the user has several simultaneous connections, this NBR is divided between the connections. In addition, there could be a time-dependent charge directly proportional to the NBR as well. This issue is discussed later in this document.

The implementation of SIMA service consists of two main parts: access nodes and core network nodes presented in Fig. 1. There is a fundamental difference between these node types: the traffic measuring of every connection is performed at access nodes whereas at the core network nodes the traffic control functions do not need to know anything about the properties of separate connections.

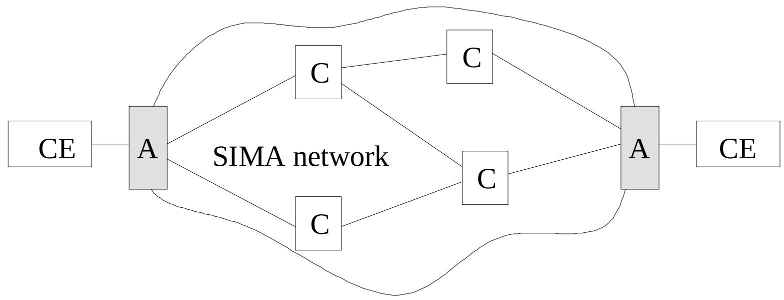

Fig. 1. Customer equipment (CE) connected in a SIMA network
with access nodes (A) and core nodes (C).

Let us suppose that an ATM connection (*i*) is terminated at an access node (i.e. there is a user/network interface based on ATM). A nominal bit rate (NBR_i) is associated to the connection and the user is transmitting cells into the network according to an arbitrary traffic process. At the user/network interface there is a device which measures the momentary bit rate of the connection. This rate is denoted by $MBR_{i,j}$. The effective measuring period is short for real-time connections and longer for non-real-time connections (the details of the measuring device are discussed in chapter 4.2).

The device gives every cell (*j*) a priority ($PL_{i,j}$) based on the $MBR_{i,j}$ to NBR_i ratio:

$$x = 4.5 + \frac{\ln(MBR_{i,j} / NBR_i)}{\ln(2)}$$

$$PL_{i,j} = \begin{cases} 7 & if \quad x \geq 7 \\ \lfloor x \rfloor & if \quad 1 < x < 7 \\ 1 & if \quad x \leq 1 \end{cases}, \qquad (1)$$

where $\lfloor x \rfloor$ is the integer part of *x*.

Consequently, if $MBR_{i,j} = NBR_i$ the cell gets priority 4, if $MBR_{i,j} > 5.66\, NBR_i$ the cell gets the lowest priority (7), and if $MBR_{i,j} < 0.17\, NBR_i$ the cell gets the highest *NBR*-priority (1). Priority 0 is reserved for those connections that use an ordinary ATM service with guaranteed bandwidth and quality (see chapter 4.1). The accepting and discarding of cells inside a SIMA network is entirely based on these priorities.

Next we deal with the question how the network can guarantee small delay for real-time connections. For this purpose every network node (and every switching block) shall have two parallel buffers: one for real-time cells and another for non-real-time cells. Before the connection is established, the user shall select either the real-time or non-real-time service class. All cells belonging to a real-time connection go through the real-time buffer and all other cells are using the non-real-time buffer. This selection can be left freely to the customer and there is no need to take it into account when determining the charge of the connection.

The key issue in the implementation of the SIMA service is the cell discarding system before the actual buffering of the cells shown in Fig. 2. At any instant there is an accepted level of priority (PL_a): if an incoming cell has the same or higher priority (i.e., the same or lower value of priority), it is accepted, otherwise it is discarded. The calculation of PL_a is based on the buffer occupancy levels of the real-time buffer (M_{rt}) and non-real-time buffer (M_{nrt}). See also chapters 4.2 and 4.3.

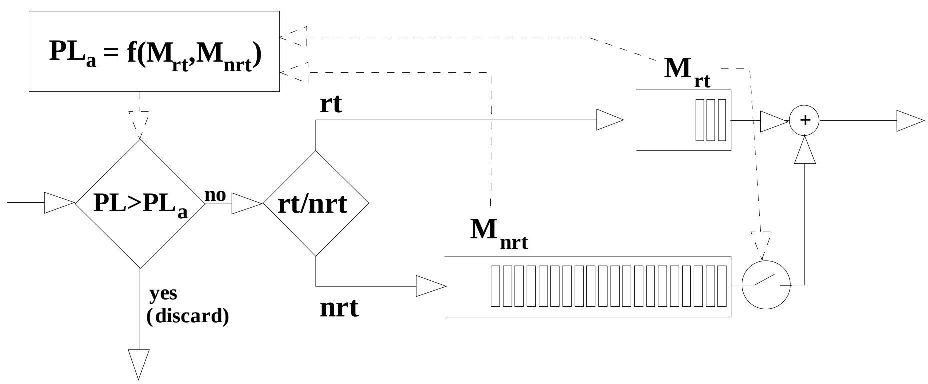

Fig. 2. A cell scheduling and buffering block (CSB) for SIMA service.

All the cells which have been accepted in the scheduling block are situated either in the real-time or non-real-time buffer (the scheduling algorithm can guarantee that there is no cell loss in actual buffers). Both buffers may apply the ordinary First In First Out (FIFO) principle. In order to obtain a small delay and delay variation, the real-time buffer should be relatively small (e.g., 200 cells). All cells in the real-time buffer are transmitted before any cell in the non-real-time buffer. It should be emphasised that this delay priority of real-time cells has no effect on the cell loss ratios of either real-time or non-real-time cells but only on the delays of different cells. The non-real-time buffer should be much larger (e.g., 20000 cells) because of the packet scale fluctuations in typical non-real-time traffic processes. Moreover, large buffers make possible to offer reasonable service for those connections that are capable to adjust their bit rate (see chapter 4.5).

3. Viewpoints

3.1 Network operator viewpoint

From a network operator viewpoint the two main properties of the SIMA service are simple traffic management and simple charging scheme.

Simple traffic management means that the operator offers in principle only one service with two components: a real-time service class and a non-real-time service class. Notwithstanding the simplicity, this one basic service is able to offer different quality levels with an automatic charging structure. A simple charging scheme is very advantageous for a network service since most customers are quite reluctant to become acquainted with complicated services and complicated charging systems. In many cases a simple flat rate scheme is the most desirable one and SIMA service is well suited to this scheme. However, if time-dependent charging is needed, there is available a clear solution: a needed NBR is negotiated between user and network and the charging of a connection is proportional to NBR and the duration of the connection (see chapter 4.7).

The technical basis of SIMA service lays on principles of Best Effort or Unspecified Bit Rate (UBR) service: on the one hand, users do not inform in advance the network on the needed bit rate or any other traffic parameters, and on the other hand, the network operator does not give any precise guarantees of the available bit rates or QoS (Quality of Service). The UBR principle with the aid of priorities makes possible a simple network structure and management and, at the same time, it results in good fairness among different connections and efficient statistical multiplexing. The basic version of SIMA service works without such ordinary management functions as Traffic Descriptor, QoS parameters, Service Classes, Connection Admission Control (CAC), or Usage Parameter Control (UPC). All these functions are replaced by two autonomous units: the measuring block and the cell scheduling and buffering block (CSB). In addition, it is possible to build an informative network service for those connections that are able to adapt their bit rate (see chapter 4.5).

The most difficult and crucial issue with the traffic management of SIMA service is the dimensioning of the network because it is the only available tool to keep customers satisfied with the service. One possible approach is that the operator attempts to offer satisfactory QoS to nominal connections. In practice this may mean that the operator measures the cell loss ratio of cells with priority level 4. This ratio should remain on a reasonable level, for instance less than 10^{-6}. If this cell loss ratio is exceeded continuously, the operator shall firstly identify the bottlenecks in the network and then increase the network capacity in those points. It should be noted that this capacity increase is a quite straightforward task because there is no need to make any new plans concerning switching structure, the capacity division between service classes or virtual paths, etc. (operator simply throws bandwidth, and the SIMA service manages QoS).

The lowest priorities (6 and 7) are used to divide the remaining capacity among highly variable connections. The cell loss ratio of cells with priority 7 could be temporarily very high, even 100%, while the cell loss ratio of cells with priority less than 3 shall be negligible ($<10^{-9}$). This phenomenon is elaborated later in chapter 5. The highest priorities (1 and 2) offer reliable, high quality service even at exceptionally busy hours and just before capacity updates. It should be noted, however, that the nominal bit rate of the connection shall be 16 times higher if the user wants to get priority 2 instead of priority 6 without changing the offered traffic stream.

3.2. Customer viewpoint

SIMA service can be clear and understandable even for an ordinary non-technical customer, because the charging is based purely on the nominal bit rate, and there is no pre-defined traffic or quality parameters for each connection.

Although it is possible that the NBR is determined according to a monthly fee, let us suppose that the customer may select different *NBR* for each connection. This choice is always a compromise between price and quality of service. The quality detected by the user depends on four issues: *NBR*, the average bit rate, the amount of traffic variations of the connection, and the current load situation in the network. If the user is not satisfied with the QoS of the connection, he or she has several alternatives: to keep the average bit rate unchanging but reduce the variation of traffic process, to decrease the average bit rate, or to increase the nominal bit rate (this means increased price, as well). In some cases the quality of an application might be improved by changing the service class from real-time to non-real time or vice versa, or by postponing the transmission until the network load decreases. As a last option, the user may change the network operator.

The selection of service class is left for the user or in many cases for the application (default value might be the non-real-time class). If the application is a real-time one, it is advantageous for the user to select the real-time class, because it is the only way to attain small delay and delay variation. However, if the user wants to obtain small cell loss ratio, the cell stream process should be smooth, because the effective measuring period is short for real-time connections (approximately 0.1 - 0.3 ms). If the user sends large bursts of cells, some of the cells may be marked with the lowest priority (7) and, consequently, they encounter very high cell loss ratio. If the application does not need small delay or delay variation, it is more useful for the user to apply the non-real-time class because it allows much larger bursts without significantly affecting the cell’s priority.

It should be stressed that every user is allowed to send data with any bit rate, higher or lower than *NBR*. In one extreme case with *NBR* = 0 all cells get the lowest priority, and in other extremity *NBR* can be higher than the transmission capacity at user/network interface (all cells may get even the highest priority). In all cases, every user which is transmitting data with his or her *NBR* encounters similar QoS independent of the real bit rate or service class.

3.3. Manufacturer viewpoint

Simplicity is the main property of SIMA service for a manufacturer as well. It should be noted that ATM switches or crossconnects can be build by using the CSB-blocks presented in Fig. 2, switching fabrics and routing functions. By using ATM virtual paths or IP switching technology the routing task can be kept small and efficient as well. In addition, Packet Discarding and Priority Feedback (see chapters 4.4 and 4.5) may be included in the CSB blocks without weakening their automaticity. Simple implementation of network nodes may result in inexpensive network infrastructure with high capacity.

The more complicated part of the network infrastructure is access nodes. The additional functions needed is a measuring unit which shall be able to measure in real time the traffic stream of every connection and to computation unit for determining the proper priority for every cell. The difficulty of these tasks is of the same order as the difficulty of the usage parameter control (UPC) in conventional ATM networks (note that UPC is a method to divide cells into two priority classes).

4. Advanced features

4.1 Interoperability with ordinary ATM networks and services

The main obstacle of SIMA service might be the incompatibility problems with the ordinary ATM technology. First of all, SIMA service requires 3 bits in each ATM cell for the determination of cell priority (or 2 additional bits if the current cell loss priority, CLP, bit in the cell header is used). In addition, one bit is needed due to the sorting of real-time and non-real-time connections. There are even some arguments to include the real-time/non-real-time bit in every cell. This approach may further simplify the implementation of CSB blocks and core network nodes as the blocks do not need to keep record of every connection. If the place of rt/nrt bit is known and stable, the realisation of CSB-block can be independent of the other part of node. In addition, this approach may alleviate some problems with connectionless traffic even though this type of traffic can be supposed to use non-real-time class.

One possibility is to use the current Generic Flow Control (GFC)-field with 4 bits in the cell header. In this case all 3 priority bits and rt/nrt-bit can be situated in the cell header. If this is not possible, the required bits (2, 3 or 4 depending on the use of CLP bit and the status of rt/nrt bit) shall be situated outside the current cell header.

Another compatibility question is how the SIMA service is working with ordinary ATM network using guaranteed connections. As the ordinary ATM services are offered in parallel with SIMA service, it shall be possible transmit ordinary ATM connections in a SIMA network. In this case the operator shall have an UPC device for each ordinary ATM connection (or possibly for each virtual path). All cells with CLP=0 in these connections are marked with the highest priority (0). If the operator wants to mark excessive cells as CLP=1 cells, those cells shall be marked with lower priority in SIMA network, for example, with priority 6.

There should be a simple connection admission control (CAC) method probably based on peak rate allocation. Peak rate allocation is supposed to be sufficient in practice because the charge of a conventional ATM connection with priority 0 will be essentially higher than the charge of SIMA connection with the same bit rate and with priority 3 or 4. In addition, the SIMA scheme exploits very efficiently the capacity between the allocated and used bit rates of VBR connections as there is no real capacity allocation for different connections or services. A possible interoperability scheme for transferring ATM services over a SIMA network is presented in Table 1.

Table 1. An interworking scheme between ATM services and SIMA.

ATM service	CLP	SIMA class rt/nrt	priority level	Connection acceptance in SIMA network
CBR	0	rt	0	Peak rate allocation
	1	rt	5 or 6	-
rt-VBR	0	rt	0	Peak rate allocation
	1	rt	5 or 6	-
nrt-VBR	0	nrt	0	Peak rate allocation
	1	nrt	5 or 6	-
ABR	0	nrt	?	?
	1	nrt	?	?
UBR	0	nrt	5 or 6	Not applicable
	1	nrt	6 or 7	

4.2 A measuring scheme using exponential moving average

Since the bit rate of every connection may change significantly in several time scales, the operator must apply an averaging measuring principle to determine the instantaneous cell rate of each connection. The time scale of the measurement shall depend on the service class (real-time or non-real-time) because the non-real-time buffer capacity can be 100 times larger than the real-time one. The approach presented in this chapter is applicable, but any measuring scheme which gives a feasible approximation of the instantaneous bit rate can be used, provided that it can be adjusted to the needed measuring period.

This measuring approach is based on the well-known principle of exponential moving average. If we suppose that the moving average is calculated at every time slot in the access node, the measured load generated by a connection (*i*) at the instant of transmission of *j*:th cell is:

$$\rho_{i,j} = (1-\alpha)^{N_{i,j}} \rho_{i,j-1} + \alpha \qquad (2)$$

where $N_{i,j}$ is the distant between *j*:th and (*j*-1):th cells in time slots and α is a parameter which defines the time scale of measurement. Formula (2) is obtained by assuming that the estimation for the instantaneous load is updated at every time slot, but all calculations are performed only at the arrival instant of a cell. The following starting values can be used: $\rho_{i,0} = 0$ and $N_{i,1} = C/NBR_i$.

In order to obtain an exact steady state value for constant bit rate connections the following conversion between load ($\rho_{i,j}$) and measured bit rate ($MBR_{i,j}$) shall be applied:

$$MBR_{i,j} = \frac{C\ln(1-\alpha)}{\ln\left(1-\dfrac{\alpha}{\rho_{i,j}}\right)} \qquad (3)$$

where *C* is the link capacity [bit/s] at the user/network interface. For numerical reasons (2) and (3) shall be replaced by

$$MBR_{i,j} = \frac{C}{N_{i,j}}$$

if $N_{i,j} > 10/\alpha$.

It should be noted that because α is usually a constant, formula (3) can be replaced by a table with a proper granularity. For the same reason, at least the term $(1-\alpha)^{N_{i,j}}$ in (2) can be tabulated.

The proper value for parameter α depends on the buffer capacity reserved for the service class used by the connection. With real-time services (with small delay variation) the buffer should be small, and thus the value of α must be quite high. On the contrary, when using a non-real-time service the user may want to send bursts of cells without high cell loss ratio. As a consequence α must be much smaller (or the averaging period should be much longer). As an interim approach the following approximation might be applicable:

$$\alpha \approx 5/K_n , \qquad (4)$$

where K_n is the buffer capacity in cells reserved for the service class n.

4.3 A cell scheduling scheme

The key point of the SIMA service lies in the function of the scheduling algorithm. The decision of the acceptance is based on two parameters: the priority level of the cell and the occupancy level of the two buffers. Let us use the following notations:

- M_{rt} = the number of cells in the rt-buffer
- K_{rt} = the number of buffer places in the rt-buffer
- M_{nrt} = the number of cells in the nrt-buffer
- K_{nrt} = the number of buffer places in the nrt-buffer

The average occupancy level of the total buffering system (x) might be determined in several ways, for instance:

$$
\begin{aligned}
x &= (x_{rt} + x_{nrt}) && (a) \\
x &= \sqrt{x_{rt}^2 + x_{nrt}^2} && (b) \\
x &= \max(x_{rt}, x_{nrt}) && (c)
\end{aligned}
\tag{5}
$$

where:

$$
\begin{aligned}
x_{rt} &= M_{rt}/K_{rt} \\
x_{nrt} &= M_{nrt}/K_{nrt}
\end{aligned}
$$

The cell is accepted if the following relation is valid

$$PL < a - bx \tag{6}$$

In reality formulae (5) and (6) can be implemented by using pre-calculated tables. The occupancy level of both buffers is divided into N levels, where N can be for instance 16. When a cell comes, the scheduling process inquires the current values of M_{rt} and M_{nrt}. By a simple calculation (especially if $K_{rt,}$ K_{nrt} and N are of the form 2^n) the process gets a rough estimate of the current occupancy level of both buffers. These two values determine the rows and columns of the table. The contents of each cell of table are the highest allowed priority level.

4.4 A packet discarding scheme

The basic SIMA service discards separate cells rather than whole frames or packets. If the cell loss probability is high, the useful throughput could be very low, because corrupted frames or packets should be sent again. A common solution to this problem is to discard all new packets when the network load or buffer occupancy exceeds a certain pre-defined limit. If the beginning of the packet is accepted in the buffer, the whole packet will be accepted (unless the whole buffer becomes full). The basic SIMA service does not comprise this type of property as it discards only individual cells.

A SIMA network operator may want to apply a scheme in which network nodes reject whole packets instead of individual cells. This property can be implemented in a SIMA network by using the priority levels. Each network node applies the normal accepting/discarding algorithm for the first cell of each packet as in the basic SIMA service. If the first cell of a packet is discarded, all the following cells of the packet are discarded as well. If the first cell of a packet is accepted then the algorithm gives a higher priority for all other cells belonging to that packet. Even a gain of one priority level seems to be sufficient to guarantee that there will only be very few partly transmitted packets. The priority gain should not be too large in order not to disturb the normal function of SIMA service.

4.5 Priority feedback for controllable connections

The conventional ATM traffic management seems to be too complicated for many practical purposes. However, it provides some useful properties, e.g., the *ABR* service might be quite helpful as it gives the users valuable information of the current load situation in the network. On the contrary, in the basic SIMA the user or application does not directly know what is the current load situation in the network, but the user has to send cells with different rates in order to achieve this information. This is a complicated process and it exploits excessively network resources.

In this chapter we describe a system which informs the user on the current practical priority level in the network. This service is called in this document Priority Feedback. The principle of the service is described in Fig. 3.

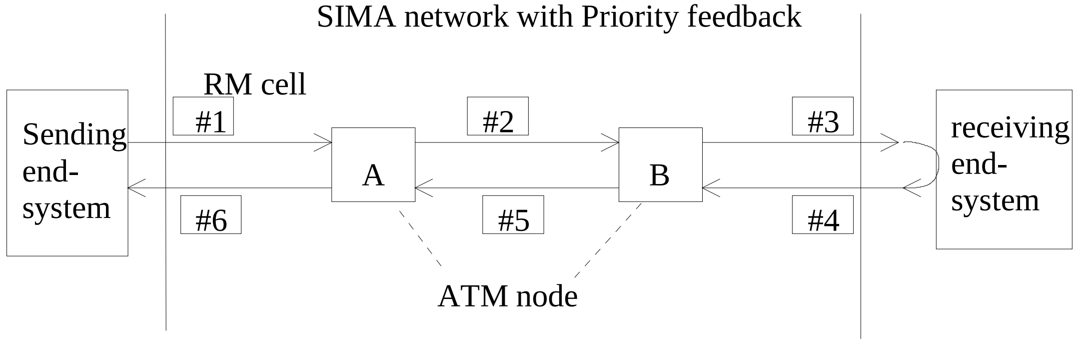

Fig. 3. The cell flow of a connection in a SIMA network with Priority Feedback.

The sending end-system sends ATM cells to the receiving end-system. Always when the sender wants to know what is the current load situation into network (in order to optimise its sending rate) it sends a special ATM cell (RM cell) into the network. This cell belongs to the normal connection flow in the sense that it is included in the measured bit rate (MBR) of the connection and the priority of the cell is determined in the same way as the priority of all other cells of the connection.

The RM cell contains an information field which determines a practical priority level for the connection, denoted by $PL_{fb,f}$. The sender end-system sets value 7 (the lowest priority) in this field (cell 1 in Fig. 3). Each network node or CSB block can detect the RM cell based on the information in the cell header. When the node or CSB block receives the RM cell, it examines the current load situation of the outgoing link of the connection. The load situation is defined as the lowest priority level which has been available for a typical connection, denoted by $PL_{fb,n}$ (priority level for feedback information in an ATM node). If $PL_{fb,n} < PL_{fb,f}$, then the node decreases the value $PL_{fb,f}$ field otherwise the node does not change the field. After this phase the RM cells go in the scheduler block which accepts and discards the cells according to the priority level of the cell and the current buffer occupancy level.

If the load levels of nodes A and B in Fig. 3 have been 5 and 6, respectively, the $PL_{fb,f}$ values will be 7, 5 and 5 in cells 1, 2 and 3, respectively. The receiver end-system turns the RM cells back. The $PL_{fb,f}$ value of the incoming (forward) RM is placed in another field in the RM cells, denoted by $PL_{fb,b}$. Then the receiver end-system sets $PL_{fb,f}$ the value 7 so that this field can be used for the same purpose as in the other direction.

$PL_{fb,b}$ of the backward RM cell (cell 4 in Fig. 3) might be the same as the $PL_{fb,f}$ in the income RM cell (cell 3 in Fig. 3). The receiver side is allowed to change the value of $PL_{fb,b}$, before it sends the RM backwards. The receiver side may use the information which it has of the received normal cells. If the receiver is supposing that some cells with priority level 5 have been lost, but all cells marked with PL=4 have been transmitted, the receiver end-system may set value 4 in the $PL_{fb,b}$ field. This value is usually not changed in the ATM network nodes in the backward direction. Therefore, in Fig. 3, the value of $PL_{fb,b}$ is 4 for both cell 5 and cell 6.

As option there could be several $PL_{fb,f}$ fields (and $PL_{fb,b}$ fields, respectively) in order to give the user information of the load situation during different time periods, for example, for during the last 100 ms, 10 s and 10 min. This could be advantageous because the capability of different services to adapt their bit rate may differ considerably.

4.6 Priority gain for connections with low average rate

The basic SIMA service determines the cell priority according to the instantaneous cell rate of each connection. In consequence, it does not give any advantage for those users which have early exploited less network capacity compared with connections that have used much network capacity continuously, if both types of connection have the same peak rate. For instance, there is only a small advantage for a customer to use variable bit rate video coding instead of a constant bit rate coding with the same peak rate. This is a somewhat unfair property because it is possible to transfer more connections of the first type with on/off nature than constant bit rate connections if the peak rates are equal because of the effect of statistical multiplexing.

Therefore, it could be advantageous for a network operator to give an extra profit for those connections that are exploiting the network capacity in long term less than the other connections. This priority gain can be realised by using two MBR's: one MBR_s using short measurement interval, and another MBR_l using long interval (e.g. some minutes). The only difference is that the parameter α is much smaller with MBR_l than with MBR_s. The final value for MBR is a weighted average of these two measurement results:

$$MBR = \beta MBR_s + (1-\beta) MBR_l \tag{7}$$

where β is a constant between values 0 and 1. If β = 0.5 the maximum profit for a connection due to low long term average bit rate is one priority level. However, because the application of this option complicates the traffic management and causes additional calculations, it is not clear whether it gives enough profit compared with its disadvantages.

4.7 A charging scheme

In this chapter we present a charging scheme for the SIMA service. The two main alternatives are to apply pure monthly fee or to use both monthly and time dependent charge. A possible approach is to use a monthly charge which determines a maximum allowed *NBR* of the user (denoted by NBR_{max}). In addition there can be a time dependent part of charge proportional to the *NBR* of each connection. If we take into account that both directions may have different *NBR*, the total charge of a customer during a month is as follows:

$$X = \beta_1 NBR_{\max} + \beta_2 \left(\sum_i \left(NBR_i^u + NBR_i^d \right) t_i \right) \tag{8}$$

where t_i is the holding time of connection *i* with an upstream nominal rate NBR_i^u and a downstream nominal rate NBR_i^d. The dimension for b_1 and b_2 can be $/(kbit/s) and $/(kbit/s)/min, respectively.

Even though more complicated charging functions can be used instead of the linear scheme presented above, the basic target of the concept is to apply as simple principles as possible.

The main reason to use NBR_{max} is that it alleviates the network dimensioning problem because the sum of NBR_{max}´s of different customers is known and relatively stable, the network operator knows the maximum load level for each priority level, and may even make feasible prediction of the QoS for different priority levels.

It should be noted that this charging scheme takes indirectly into account bit rate, QoS and traffic variations, the effect of network load, user expectations and willingness to pay (this issue is briefly addressed in the next chapter).

There is no technical obstacle to change the value of NBR during the connection because the change has an effect only at the user/network interfaces. However, this approach may result in more difficult network dimensioning, because the total sum of NBR's of different connections can vary significantly depending on the network load. In that respect, a pure monthly NBR seems to be the most practical approach.

5. Performance evaluation

The main difference in performance evaluation between SIMA and conventional ATM services is the priority levels. Therefore, the focus of this chapter is to illustrate the QoS and throughput at different priority levels. The first question is how big is the quality difference between adjacent priorities. It should be remembered that the price is doubled if the user wants to obtain one degree higher priority for every cell without changing the real bit rate. In consequence, the QoS should be improved so much that at least some users are willing to pay the additional charge.

5.1. Performance evaluation with independent traffic process

In this chapter it is assumed that there are many identical traffic sources which generate traffic independent of the current or previous load condition in the network. The following traffic parameters are assumed: the link capacity C = 1, peak bit rate = 0.1, the ON probability at the burst (or packet) scale = 0.2, and the average burst duration = 1000 time slots (i.e., the average packet size = 100 cells). In addition we are supposing that there is an upper on/off layer which is used to model the random process of connections. It is assumed that both the average on-period and off-period of this layer are 100 000 time slots. The real time buffer contains 200 cell locations and non-real-time buffer 5000 cell locations. By using the equation (4) for the time scale parameter a we obtain:

$$\alpha_{rt} = 0.025$$
$$\alpha_{nrt} = 0.001$$

In this example, eight different connection types are assumed: four connection types are real-time ones and four are non-real-time ones. Also, four different NBR values are assumed as: 0.2, 0.1, 0.05, and 0.025. The priorities corresponding to these NBR values, with maximum MBR = 0.1, are 3, 4, 5 and 6, respectively. It should be noted, however, that not all cells are assigned these exact priorities, and that especially with non-real-time connections, many cells obtain better priority values because of the affects of the averaging measuring principle. The distribution of cells having different priority levels, represented as percentages, is presented in Table 2.

Table 2. The percentage of cells of different priority levels

priority level	real (simulated) percentage of offered cells	percentage based on peak rates
1	6.1	0
2	7.9	0
3	24.3	25
4	23.5	25
5	21.5	25
6	16.8	25

In Fig. 4, there is shown a graph illustrating the relationship of average cell loss ratio, P_{loss}, as a function of priority level for four specific load levels, r. It is noted that, in the case of load = 0.80 the cell loss ratios for real-time and non-real-time cells are indicated by dotted and broken lines, respectively. The figure shows that the difference in cell loss ratio between real-time and non-real-time cells is insignificant.

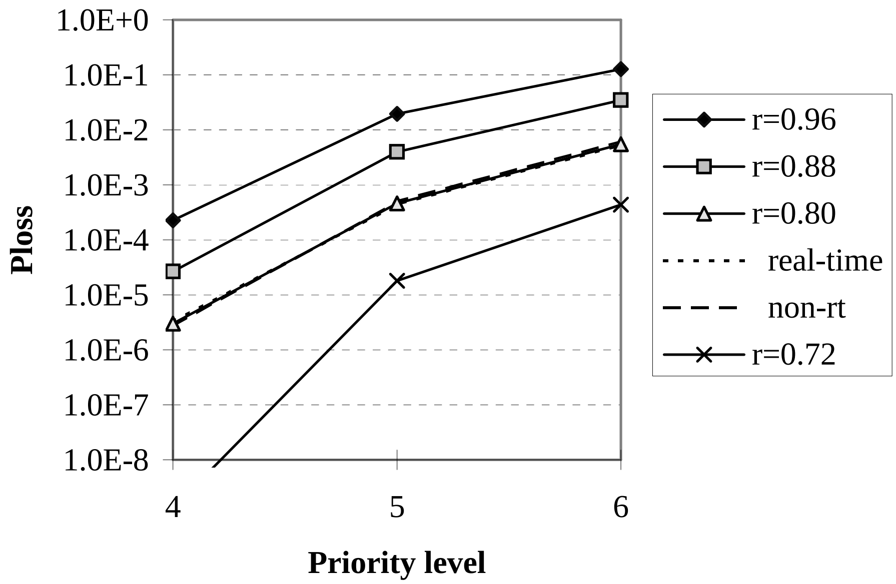

Fig. 4. Average cell loss ratio vs. priority level for load levels r = 0.72, 0.80, 0.88, 0.96. In case of load = 0.80 the cell loss ratios for real-time and non-real-time cells are presented by dotted and broken lines.

Given, for example, a traffic scenario where the operator wants to offer a cell loss ratio of 10^{-6} to cells with priority 4, the total load can be approximately 0.75. It can be assumed that this average cell loss ratio is sufficient for most video applications. Give the same traffic conditions, priority level 5, which corresponds to $P_{loss} \approx 10^{-4}$ can meet the requirements of many voice applications, and priority 6, which corresponds to $P_{loss} \approx 3 \cdot 10^{-3}$, is suitable for a TCP/IP type of file transfer, provided that there is an adequate packet discarding scheme in place.

It should be emphasized, however, that the difference in cell loss ratio between adjacent priorities depends strongly on the offered traffic process and, in particular, the inherent control loops of the SIMA service. When the user perceives an unsatisfactory QoS, the user can, and should, change either the actual bit rate or the nominal bit rate of the connection. In either case, the priority distribution changes as well. Nevertheless, if this phenomenon is temporarily ignored, the basic behavior of priority distribution may be further appreciated by making the following simplifying assumption. If it is assumed that all traffic variations are slow as compared to the measuring period and buffer size, then a well known, conventional ATM approach to approximating cell ratio may be used, with the additional requirement that the eight NBR priority levels are taken into account.

If the loss ratio of cells with priority k is denoted by $P_{loss,k}$ and the average loss ratio of cells with priority of 0 to k is denoted by $P^*_{loss,k}$, then the following equation, which ignores buffering effect, provides that:

$$P^*_{loss,k} = \frac{\sum_{j:\lambda_j > c} \Pr\left\{\lambda^*_k = \lambda_j\right\}\left(\lambda_j - C\right)}{\rho^*_k C}$$

$$P_{loss,0} = P^*_{loss,0}$$

$$P_{loss,k} = \frac{\rho^*_k P^*_{loss,k} - \rho^*_{k-1} P^*_{loss,k-1}}{\rho^*_k - \rho^*_{k-1}} \qquad for \quad k = 1\ldots7 \tag{9}$$

where, λ^*_k represents the momentary bit rate level of all cells with a priority of 0 to k, ρ^*_k represents the average offered load produced by these cells, and C represents the link capacity. The probability $\Pr\left\{\lambda^*_k = \lambda_j\right\}$ can be calculated in a straightforward manner by using known convolution techniques.

For purposes of further illustration, we assume the same sources described in the beginning of this chapter (except the long ON and OFF periods). Because of the long periods the peak rate always determines the cell priority. As in this case the buffers are not capable of filtering any traffic variations, the allowed load is much lower in this example than in the original case.

In Fig. 5, there is illustrated in graphical form a relationship between cell loss ratio as a function of priority level for different load levels, r. Fig. 5 shows the cell loss ratios obtained by application of Equation (9) for different priorities. It is assumed in Fig. 5 that the peak cell rate of each connection depicted by solid lines is 0.1. The peak cell rate of connection depicted by broken line is 0.2, which actually means that traffic variations have been doubled by changing both the peak cell rate and nominal bit rate. The peak rate cell rate of connection depicted by dotted line is 0.05. As the nominal bit rate is halved, as well, the traffic variations are decreased.

In a network that embraces the SIMA service concept, an increase of traffic variations has two main effects if the operator keeps the QoS of priority level 4 unchanged. First, the allowed load level is decreased in the same way as in conventional ATM, and second, the difference in cell loss ratio between adjacent priority levels decreases. For purposes of providing a rough estimate of QoS based on Fig. 4 and 5, it may be assumed that if priority level 4 offers a cell loss ratio of 10^{-6}, then the cell loss ratio will be approximately 10^{-4} to 10^{-3} with priority level 5 depending on the overall traffic variations. The cell loss ratio with priority 3 can be supposed to be less than 10^{-9} unless the traffic variations are very pronounced.

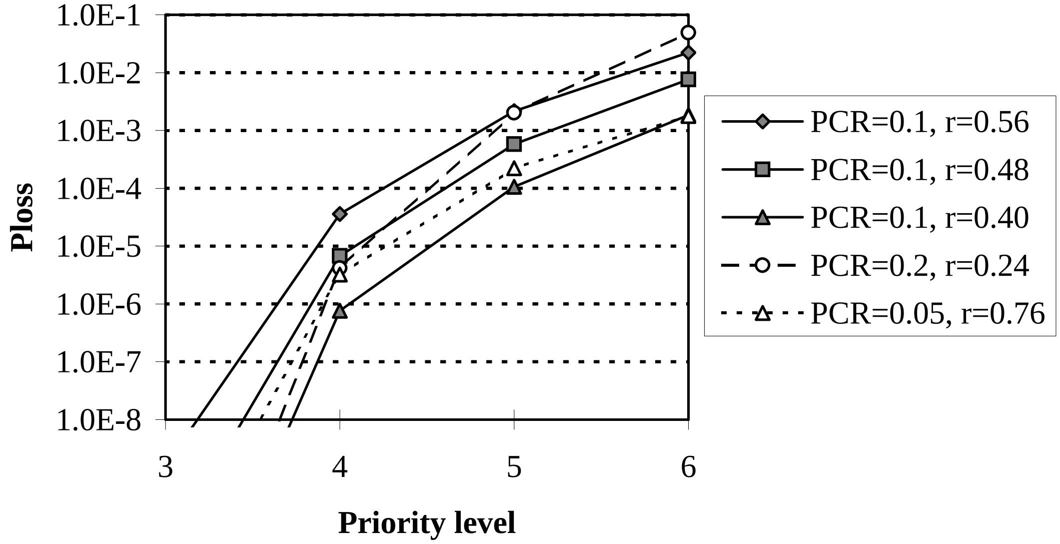

Fig. 5. Cell loss ratio vs. priority level for different load levels (r); solid lines: peak cell rate of each connection = 0.1; broken line: peak cell rate of each connection = 0.2; dotted line: peak cell rate of each connection = 0.05.

5.2. User reactions to quality differences

Although the previous examples provide illustrations between QoS and priority levels, it may be unfruitful to attempt to exactly determine the allowed load or the cell loss difference between adjacent priority level until user reactions to different QoS and usage charges are evaluated. In a SIMA service environment, a schedule of charges based on different QoS levels may be determined, in a certain sense, automatically. For example, if the difference in cell loss ratio between priority levels 4 and 5 is very small, it can be assumed that some of the connections will tend to move from priority level 4 to level 5 because of a lower assessed charge. This change indicates, apparently, that the cell loss ratio of priority level 4 decreases and the cell loss ratio of priority level 5 increases. It can be reasonably assumed that this type of movement continues until the QoS difference corresponds to the average user’s expectation of a reasonable charging structure.

Similar concerns are raised with regard to the differences when charging users during busy hours in contrast to idle hours. For example, it would appear reasonable to charge higher prices during low load periods for a certain QoS and bit rate. With SIMA service this difference is achieved automatically as the user has to use better priority my means of higher NBR during busy hours in order to obtain the same QoS as during idle hours. This "supply and demand" effect may tend to automatically even out the load between busy and idle hours.

In this respect there is a fundamental difference between SIMA and the conventional ATM, where the operator must itself plan a complicated but practical charging structure for several QoS classes.

5.3. Performance of controllable connections

Another important control loop is the adaptation of momentary bit rate. As presented in chapter 4.5 it is possible to give those sources that are able to change their bit rate information about the current load condition in the network. This chapter attempts to give an outlook of the performance to be obtained by these connections. Let us first take an example with a background traffic process generated by the sources determined in the beginning of chapter 5.1. In this example we have 10 connections of each type which means an average background load r = 0.80. There are also three sources (below called as feedback sources) that adjust their transmission rate according to the feedback information that they receive from the network. All these feedback sources have NBR = 0.01. The feedback sources are similar to each other except that the time period, used for determining the PL_{fb} information, is different, namely 10 000 (FB1), 30 000 (FB2) and 100 000 (FB3) time slots. For comparison purposes we also have three constant bit rate connections with the following parameters (link capacity = 1):

C4: bit rate/NBR = 1.4 , NBR = 0.01, priority of cells = 4

C5: bit rate/NBR = 2.8 , NBR = 0.01, priority of cells = 5

C6: bit rate/NBR = 5.5 , NBR = 0.01, priority of cells = 6

These sources transmit at a rate that is slightly less than the limit that introduces the next, worse priority level. For example bit rate/NBR = 1.5 would introduce already priority of cells 5 while bit rate/NBR = 1.4 yields priority level 4. The feedback sources used the same bit rate values in order to optimally utilize the network capacity.

Figure 6 shows the simulation results. The cell loss ratio is plotted as a function of accepted bit rate / NBR. The figure shows that constant bit rate sources are obtaining better loss vs. throughput when they are sending at some relatively good priority level. However, the results indicate that feedback sources are useful when the background traffic is changing slowly. In this case they can adapt to changes while constant bit rate sources cannot take advantage of the changing load of the network. With fast changing background a feedback source is not able to adjust to changes fast enough and its cell loss ratio increases.

Figure 6 also includes one source that roughly acts as some kind of TCP source with NBR = 0.01. This source halves its transmission rate when it gets information about a lost cell. It increases its transmission rate by 10% if it does not get back information about lost cells over a time period of 10 000. Comparison shows that this kind of source loses more cells than corresponding feedback source. This is expected since this kind of source has slower reacting to changes and it does determine the actual bit rate into the priority steps as the feedback sources.

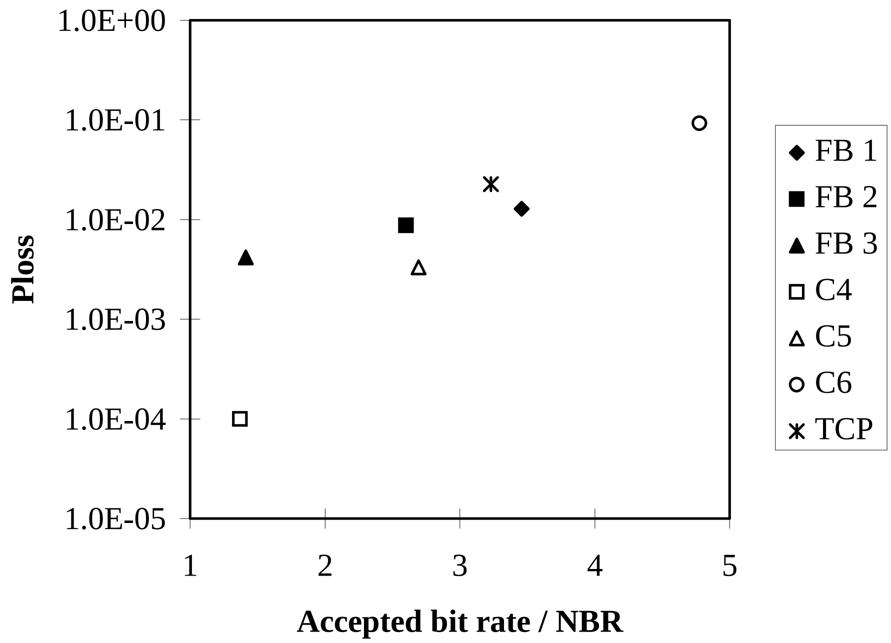

Fig. 6. Loss ratio of different sources as a function of accepted bit rate/NBR. Sources FB1, FB2 and FB3 are using feedback information for adjusting they rate. Constant bit rate sources C4, C5 and C6 are transmitting cells with priority levels 4,5 and 6, respectively. Source TCP halves its transmission rate when it receives a notice of a lost cell and slowly increases the rate in the absence of lost cells.

Other important questions are how well the connections can be adjusted into sudden capacity changes, and how fair will be the capacity division between different feedback connections. Fig. 7 shows a case where four feedback sources are first transmitting to the network node. Two of them have NBR = 0.25 (FB1 and FB2) and two have NBR = 0.0625. At time = 30 000 a uniform source starts to transmit with rate PCR = 0.333 and NBR = 0.333. At time 60 000 the source is switched off.

Figure 7 shows the throughput / capacity from each of the sources as a function of time. As figure shows, the feedback sources are able to adjust their transmission rate to the step wise change in the load. All the feedback sources have roughly the same decrease of throughput when the step source is switched on. After the constant source is switched off, the feedback sources restore their original throughput. No clear instabilities after the changes are seen.

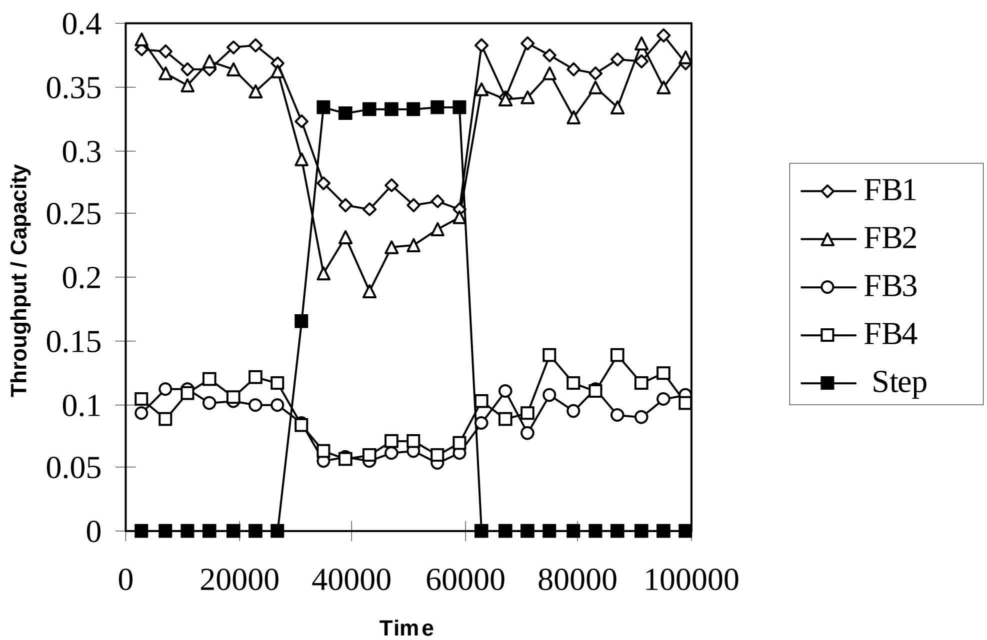

Fig. 7. Throughput/capacity as a function of time. Sources FB1 - FB4 are feedback sources. Step sources are switched on at time 30 000 and switched off at time 60 000.

6. Conclusions

Notwithstanding the complexity of conventional ATM traffic management schemes, the current ATM specifications fail to adequately address the need of simple management and feasible charging for future Internet and other networks with high capacity and quality requirements. Accordingly, there is a need in the communications industry for a network management architecture that is simple in concept and in its implementation, yet adequately addresses the quality of service requirements to support a variety of network services, including real-time and non-real-time services. There exists a further need for a system and methodology that provides for the implementation of a simple and effective charging capability that accounts for the use of network services. The present SIMA service introduced in this document is capable to fulfill these and other needs which remain unaddressed by current traffic management approaches.

The SIMA service is technically based on three key ideas: the use of nominal bit rate concept, the use of 8 priority levels for every cell, and separation of real-time and non-real-time connections at the buffer level. If a user needs a connection over a IP or ATM network, he should select a nominal bit rate which could be even a constant proportional to a monthly fee. The other decision needed before a connection establishment is that the user shall select either a real-time or a non-real-time service class. In addition to these two parameters the user does not need to give any information about the properties of the connection like required bit rate or quality of service. After the connection establishment the capacity division among different connections is based on a priority which is determined using a ratio of the measured bit rate and the nominal bit rate. This priority in addition to the real-time/non-real-time separation is sufficient information for every network node to properly manage the traffic in the network.

Because there is no need for various traffic classes, traffic parameters and network services, the SIMA service makes possible a simple and efficient implementation of network nodes, a simple and fair charging scheme, and very simple traffic management in the core, high speed network. In consequence, the SIMA concept is a very promising scheme for solving the most acute traffic control and management problems in Internet.

Author’s address

Kalevi Kilkki
Nokia Research Center
P.O.Box 422
FIN-00045 NOKIA GROUP
Finland

E-mail: kalevi.kilkki@research.nokia.com
Tel. + 358 9 4376 6817
Fax. + 358 9 4376 6851

Information about SIMA will also be available in near future from
http://www-nrc.nokia.com/sima/

Expiration

This document will expire in 20th September 1997.