Simple Integrated Media Access - a Novel Service Concept for Internet

Kalevi Kilkki

Nokia Research Center
e-mail: kalevi.kilkki@research.nokia.com

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.

The current approaches for meeting these requirements consist of several service specifications for different basic communication needs. There seems to be demand for three elementary services: first one for very reliable and high quality connections, second for connections with less stringent quality requirements, and third one for data connections which can smoothly adapt their bit rate. As the requirements of these elementary classes differ significantly, an obvious approach is to have different service specifications like the services specified by ITU [1] and ATM Forum [2], and recently also by Internet Engineering Task Force (IETF) [3].

The expected advantage of this approach is that by dividing the service specification task into several smaller parts the specification process is easier than if the all the service types were included in a single specification. However, this advantage is somewhat questionable because the whole service concept (with all the different service types) is what the network operator should manage and sell to customers and what the customer should buy and use. In particular, most customers of internet service providers (ISP) will certainly be reluctant to learn and understand several complicated services which may have very different structures, traffic parameters, charging schemes, etc. For real marketing purposes, the Internet service package must be simple, much simpler than what the current service models directly make possible. There are two main approaches to meet this requirement: to hide the complexity of network service from the end-user or to design an entirely new integrated service which is able to satisfy all the primary customer needs. The approach used in this document belongs to the later category. This paper presents a new simple approach for traffic management for integrated broadband networks: Simple Integrated Media Access (SIMA), see also [4].

2. Use of priorities

The basis of our proposal is the application of priority bits. Of course, priorities are used widely in telecommunication networks, and there are several different ways to apply them. Firstly, there is important division into two priority categories:

1. Priorities can be used to determine which packets (or cells[1]) shall be discarded under congestion situation. In this paper we use term “drop preference” to describe this priority scheme. Moreover, we assume that the technical realisation of the drop preference scheme is based on scheduling units before actual buffering of packets. The scheduling unit either accepts or discards a packet based on the information of load level in the buffer. It is also possible to discard packets that already are in the buffer, but that kind of case is not considered here.

[1] although the term “packet” is widely used in this paper, it can be replaced by “cell” in most cases.

2. Priorities can be used to determine that some packets shall be transmitted forward before some other ones. In this paper we use term “delay priority” to describe this priority scheme. The main purpose in applying delay priorities is that it makes possible an integration of connections with significantly different delay requirements into the same network link.

Another important issue is how the information of the priorities is transmitted in the network:

1. The necessary information can be transmitted into every network node (belonging to the path of the connection) before transmission of user packets. This means actually some kind of signaling system (or control plane) which transmits the necessary information. The basic drawback of this scheme is that all priorities are permanent and independent of the traffic sent by the user. In addition, the scaling of the network nodes is complicated due to the need of storing information.
2. The priority information can also be situated in every packet. An advantage of this scheme is that it makes possible a flexible change of priorities. Furthermore, the implementation could be much simpler if all the necessary information is in each packet, and there is no need to have a database for the priorities of every connection.

Figure 1 gives a framework in which the above presented priority schemes can be illustrated. The following notation is used: N = the number drop preference bits, M = the number of delay priority bits, L = other information related to the connection, S = scheduling unit, P = traffic profile measurement.

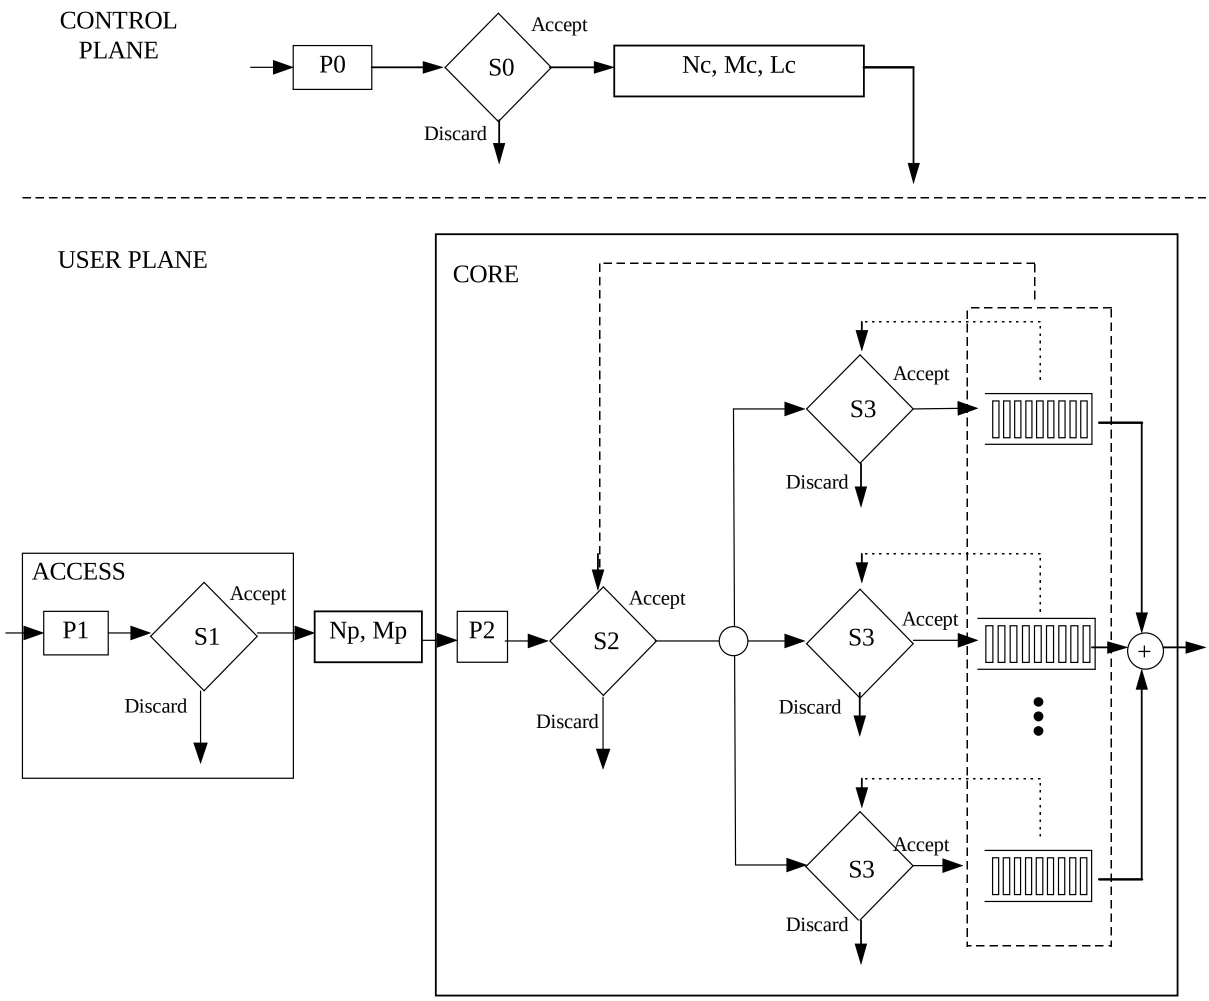

Fig. 1. A framework for services with drop preference and/or delay priorities.

On control plane the network does not handle packets or cells but connections. Therefore the meaning of profile measurement and scheduling is different: profile consists of information of traffic parameters of connections, and in some cases also, of traffic measurement results (P0). This information is used to either accept or reject the connection request in all network nodes (S0). If the connection is accepted the necessary priority information of the connection will be saved in the database of the node.

User plane functions can be divided into access node and core node functions. Access node means here the first point in which the network operator receives user packets. That is usually the point where traffic streams are controlled, i.e., packets are prioritised based on a traffic measurement (P1), and some packets can be discarded (S1). In some schemes the drop preference bits (Np) and/or delay priority bits (Mp) are transmitted in every packet.

In core network nodes, or actually in every multiplexing or switching unit, packets can be handled either based on priority bits in packet (Np, Mp) or permanent information (Nc, Mc). In addition, it is possible to make some kind of profile measurement inside the node (P2), and use that information for the scheduling. Packet scheduling can be done either before the queue selection (S2) or after the selection (S3). We assume in this paper that all accepted packets will be served and that each buffer uses a simple FIFO discipline. The queues are served typically in a way that a low priority buffer is served only if all the higher priority buffers are empty. The properties of some current network services and some preliminary proposals are presented in Table 1.

Table 1. Some network services using priorities.

		in Fig. 1:	CBR +UBR 1)	VBR with marking 2)	≈ Clark & Wro. 3)	UBR with FBA 4)	**SIMA**
Control plane	Admission control	S0	x	x	x	-	-
	Drop preference bits	Nc	-	-	-	-	-
	Service class bits	Mc	1	-	1 (?)	-	-
	Traffic parameters	Lc	PCR	several	several	-	-
Access node	Profile measurement	P1	x	x	x	o	x
	Number of profile levels		2	3	2	2	8
	Scheduling (discarding)	S1	x	x	-	o	-
Packets/cells	Drop preference bits	Np	-	1	1	-	3
	Service class bits	Mp	-	-	- (?)	-	1
Core node	Profile measurement	P2	-	-	-	x	-
	Scheduling based on the total load situation	S2	-	-	-	-	x
	Scheduling based on the load of each service class	S3	-	-	x	x	-
	Number of buffers		2	1	2 (?)	1	2
	Intra-queue discipline		DP	-	ns	-	DP

Legend x: mandatory function, o: optional function
ns: not specified, DP: delay priority

Notes:

1) The basic idea of this scheme is that the network operator may offer as simple ATM service packet as possible, while still meeting the requirements of both interactive and data applications. Small packet loss ratio and short delay for CBR (Constant Bit Rate) connections can be guaranteed by using connection admission control and prioritised buffer, while UBR (Unspecified Bit Rate) service class makes possible an effective statistical multiplexing but without any quality guarantees. All cells in CBR buffer are transmitted before cells in UBR buffer.

2) In this scheme VBR (Variable Bit Rate) service class with priority marking is used for all services. However, a more realistic scheme is that the previous scheme is supplemented by this VBR service. In that case the numbers of the two columns are combined: 2 service class bits are needed, the number of traffic profile levels is 3, and the number of buffers will be 2 or 3 (CBR and VBR services may use the same buffer). In case of VBR service with marking cells are classified at the access into three classes: high priority, low priority, and cells to be discarded immediately. High priority part of the traffic has to get high quality while the network does not usually give any guarantees for low priority traffic. Admission control is necessary for the high priority traffic.

3) This an interpretation of the scheme proposed by D. Clark and J. Wroclawski in IETF's Integrated Service Working Group [3]. The main sources of information are IETF draft "An Approach to Service Allocation in the Internet" (July, 1997), and some Integrated Service mailing list discussions on the topic, in particular D. Clarks mail "RE: differential services bof", 15th August, 1997. The main idea of the proposal is to mark all packets either "in" or "out" of a pre-defined traffic profile. All "out" packets are transmitted in the network but during congestion (that is, when the buffer occupancy level is high), they are dropped. Some kind of connection admission control is needed in order to guarantee that the probability of dropping "in" packets will be low enough. We have supposed that service class bits are transmitted on the control plane, because it is mandatory to transmit some other traffic parameters for connection admission control purposes. However, it is also possible to transmit service class information in each packet, e.g., by using Type of Service field in an IP packet. Number of buffers could be larger than 2 and different intra-queue disciplines can be applied.

4) This scheme is presented in an paper to be published in near future [5]. Fair Buffer Allocation (FBA) scheme is based on UBR or best effort service. In FBA all incoming packets or cells are accepted in the buffer if the occupancy level is not too high, say 75 % of the total capacity. Above that boundary the discarding decision of a cell depends on the number of cells of the connection which are already in the buffer (the profile measurement task is to keep a record of the cells of each connection in the buffer). At first, cells are discarded only if there are a lot of cells of that connection in the buffer, but if the buffer load level still rises, finally all cells which are the first one of a packet will be discarded. In practice, under congestion situation the FBA scheme tries to divide the available capacity evenly among all connections in progress. Packets or cells can be discarded at network access if the traffic exceeds a certain high bit rate, but that kind of control is not necessary.

Any one of these approaches could, at least in principle, form the basis traffic control of internet. However, all of them have some disadvantages. One of the key problems in internet, or in any other future integrated broadband network, will be the enormous number of simultaneous data flows. Therefore, if any complicated procedure is necessary, it shall be done preferable at the network access, not in the core network. Unfortunately, the first three schemes require a connection admission algorithm in every node, which also means that the state of every connection shall be maintained in the node. Another drawback is that most connection admission control methods use traffic parameters which shall be determined (by the user or application) before the connection establishment - this task seems to be very difficult to perform in case of web-surfing.

The fourth approach is essentially simpler than the other ones, as it does not require any connection admission control. However, a pure FBA cannot offer appropriate service for real-time services. This fact limits seriously it's applicability as a common internet service. Then, if CBR or VBR service is used with FBA, the overall service needs, again, complicated control methods.

In order to solve this problem, we propose a novel traffic control principle for internet: Simple Integrated Media Access (SIMA). The SIMA concept can be characterised by the following properties:

1. Control plane functions are not necessary (although they can be used if necessary),

2. drop preference is determined by 3 bits,
3. both drop preference and delay priority information is transmitted in every packet, and
4. scheduling is based on the total load situation of all buffers.

In the following chapter the SIMA service specification is presented in detail.

3. Service specification

The primary idea of the SIMA service is to maximise the exploitation of network resources with a simple control scheme while keeping the ratios of QoS levels offered to different flows unchanged under changeable traffic conditions. The maximisation is based on three key features: all flows with different QoS requirements share the total capacity of every link, the network attempts to avoid any unnecessary packet discarding, and flow (or call) level blocking can be avoided. The approximate constancy of QoS ratios and simplicity are achieved by using 8 priority levels which make possible a fair packet discarding scheme inside the network without keeping track on the traffic of every flow.

The SIMA specification covers the whole Internet service including charging, QoS and performance aspects, and traffic control functions in the network. As opposed to most service specifications, charging is the starting point of the SIMA concept. The prevalent charging scheme applied by Internet operators is a flat rate one with a constant monthly fee. Although this scheme is most reasonable when the network service is based on the best-effort principle, many network operators may still be willing to apply this scheme even with more complicated service models. The SIMA service model attempts to meet this demand.

3.1 Nominal Bit Rate

When the network operator offers the SIMA service, a customer first pays for some Nominal Bit Rate (NBR, kbit/s) and then he/she can trade the speed for QoS. Let us assume that a user pays X Ecu/month. This charge is translated to a NBR using an arbitrary function. Depending on the available information and the network capabilities, there are three basic approaches to manage NBRs. The simplest approach is to assign the NBR only to an interface, which means that the network measures the whole traffic going through the interface and handles this traffic as an indivisible entity. The users and flows that share the NBR obtain approximately the same QoS. In the second approach each user (identified by an IP address) has his/her own NBR. Now the network measures the total traffic generated by a user, and different flows compete with each other on a best-effort basis.

Both these approaches have the drawback that they do not separate different applications properly: a high-speed file transfer may disturb other flows, e.g., real time video connections, although the user may consider the file transfer a background process which uses only the capacity left by other more demanding applications. Therefore, as regards the performance and QoS of the SIMA service the most useful approach is the one where every flow has its own NBR. Later in this document we suppose that the network is capable to identify and measure every flow, and that every flow has its own NBR. The question how these NBRs are determined and managed can be left for network operators, and is, therefore, out of the scope of this document.

3.2 Real-time vs. Non-real-time service

The other part of the SIMA service concept is the possibility to request a real-time service. The user is entitled to him/herself determine whether the flow is a real-time (rt) or non-real-time (nrt) one. In practice, this decision can be made usually at the application level: a real-time service is requested only for interactive audio or video applications. If a real-time service is requested, the SIMA network attempts to offer as short delay and small delay variation as possible by using small buffers reserved only for real-time connections. The expense of this choice is that, if there are traffic variations of time scale from 0.1 ms to 10 ms, small real-time buffers cannot filter these variations (see illustration in chapter 4.1). Therefore, the measurement for the priority determination shall be more sensible as regards the traffic variations in case of real-time service than with non-real-time service.

If the user changes a VBR connection from nrt-service to rt-service without changing NBR or traffic process, the delay will decrease, but the cell loss ratio may increase because real-time measurement gives worse priorities during peak rates. If the user wants to obtain the same quality, this impairment of loss ratio should be compensated by increasing NBR. Real-time-service could, in this respect, be more expensive than non-real-time service although there is no difference in the actual tariffs. In consequence, 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. Furthermore, if the traffic variations are small enough, the user may always select a real-time service, because there is no difference in cell loss ratio between rt and nrt-services. In contrast, if there are significant traffic variations as with typical data applications, non-real-time service gives better quality, that is, smaller packet loss ratio.

3.3 Quality of Service expectations

The total SIMA service requested by a user consists of a nominal bit rate and of a possible real-time service request. This half of the service is as clear and reasonable as possible. The other half of the service is the expected QoS of the flow, or actually, the expected QoS of the application that the customer uses over the SIMA network. An essential issue for the success of the SIMA service is how reasonable and acceptable this part of the service concept will be.

Most customers have experience of circuit switched networks (like telephone networks) and packet networks with best-effort service (like the current Internet). In a circuit switched network a busy period means that the call blocking probability increases. In packet networks the packet loss ratio increases during busy periods, and effectively, the available capacity for a flow decreases if a TCP/IP type of protocol is used. In a SIMA environment, when a user buys a NBR for a flow and then sends traffic into a SIMA network, there is usually no flow level blocking. The quality of the flow depends on two issues: the NBR to actual bit rate ratio, and total load in the network. Therefore, a potential difficulty is that the customer cannot precisely know what the QoS of a flow will be because rapid traffic variations may bring about unexpected changes of QoS. Note, however, that even in the case of services using resource reservation the actual quality of flows using certain quality class may vary significantly because the quality can only be depicted by statistical parameters.

Because the quality of existing flows is not in the same way predictable as with services using complicated resource reservation mechanism, the SIMA network shall be implemented in a way that the users can rely on the fairness of the service. The fairness of the SIMA service is based on the fact that all flows with the same actual bit rate to NBR ratio perceives similar QoS. Thus, a home user with 10 kbit/s NBR receives the same QoS as a large company with NBR of 100 Mbit/s provided that both are transmitting at their own NBR.

Another aspect of fairness is the possibility to obtain more quality with higher price or lower price with less quality by changing the actual bit rate or NBR. This means that each customer is entitled to change the NBR to actual bit rate ratio and by that means to optimise his/her quality to charge ratio. If the ratio increases, the quality of the flow is enhanced. If the user sends traffic by using a constant bit rate, the SIMA service offers different quality levels (for variable bit rate traffic the levels are less distinct but basically the same). Although the absolute quality of each priority level depends on the network dimensioning and on actual traffic process, the quality levels can be described approximately as follows:

7 = reserved for non-SIMA services with resource reservation

6 = excellent quality: negligible packet loss ratio, and high availability even during network failures

5 = high quality: packet losses only during exceptional traffic peaks or during network failures

4 = good quality: small packet loss ratio even during busy hour

3 = moderate quality: usually small packet loss ratio except during busy hours

2 = satisfactory quality: from time to time very high packet loss ratio

1 = suitable for best-effort traffic during busy hour

0 = unusable during busy hour, but suitable for best-effort traffic during non-busy hours

Note that even very high load of the low quality levels has no significant effect on the packet loss ratio of the highest levels. It is reasonable to assume that the most intense traffic variations occur at the lowest quality levels, whereas the charging may dampen the variations at the highest quality levels. Thus, for most of the time the highest priority levels can be considered as insulated from the lower levels having more varying packet loss ratio.

4. Implementation of SIMA

There are two main alternatives for the realisation of the SIMA service: the first one based purely on packet network and the second one based on the use of ATM for the switching and transportation. As the basic implementation of these two alternatives does not considerable differ from each other, in the following both versions are presented in parallel. The main difference is that the ATM makes possible to realise more easily a satisfactory real-time service.

The implementation of the SIMA service consists of two main parts: access nodes and core network nodes presented in Fig. 2. There is a fundamental difference between these node types: the traffic measurement of every flow 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 flows.

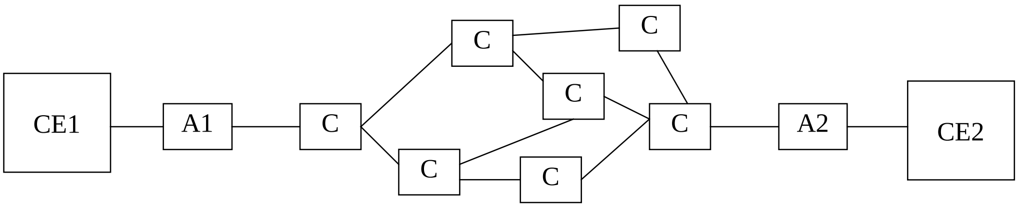

Fig. 2. Customer equipment (CE1) connected to an other customer equipment (CE2) through a SIMA network with access nodes (A) and core nodes (C).

4.1. Access node functions

Let us suppose that there is an IP flow (i) at an access node. A nominal bit rate, NBR(i), is associated to the flow and the user is transmitting IP packets (which may be converted into ATM cells) into the network according to an arbitrary traffic process. At the user/network interface there is a measuring device which measures the momentary bit rate of the flow at the arrival of the j:th packet (or cell). This rate is denoted by MBR(i,j). The device gives every packet (or cell) a priority, PL(i,j), based on the MBR(i,j) to NBR(i) ratio:

$$x = 4.5 - \ln\left(\frac{MBR(i,j)}{NBR(i)}\right) / \ln(2)$$

$$PL(i,j) = \begin{cases} 6 & \text{if } x \geq 6 \\ \mathrm{Int}(x) & \text{if } 0 < x < 6 \\ 0 & \text{if } x \leq 0 \end{cases} \qquad (1)$$

where Int(x) is the integer part of x.

Consequently, if MBR(i,j) = NBR(i) the packet (or cell) gets priority 4, if MBR(i,j) > 5.66 NBR(i) the packet (or cell) gets the lowest priority (0), and if MBR(i,j) < 0.17 NBR(i) the packet (or cell) gets the highest NBR-priority (6). Priority 7 is reserved for those connections that use a network service with guaranteed bandwidth and quality. The accepting and discarding of packets (or cells) inside a SIMA network is entirely based on the priorities.

Since the bit rate of every connection may vary significantly in several time scales, the operator must apply an averaging measuring principle to determine the instantaneous bit rate of each connection. 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. The measuring approach is based on the well-known principle of exponential moving average. In addition, we assume that ATM is used at the transport layer (ATM is also used in all the simulations presented later). If we suppose that the moving average is calculated at every time slot, 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.

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). The difference between actual and measured bit rates is illustrated in Fig. 3.

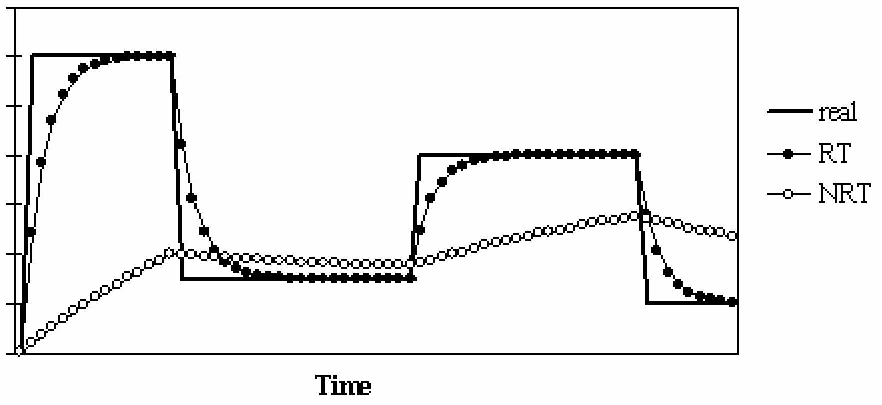

Fig. 3. The difference between actual bit rate, measured bit rate for a real-time flow and measured bit rate for a non-real-time flow.

4.2. Scheduling and buffering unit

The key issue in the implementation of the SIMA service in a high capacity core network is the packet or cell discarding system before the actual buffering shown in Fig. 4. At any instant there is an accepted level of priority (PL_a): if an incoming packet or cell has the same or higher 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}). In practice, the calculation of PL_a can be based on the use of a pre-calculated table (and therefore, there is no significant real-time processing effort).

All the packets or cells which have been accepted in the scheduling unit 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 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., 10 kbyte). All packets (or cells) in the real-time buffer shall be transmitted before any packet (or cell) in the non-real-time buffer. It should be emphasised that the delay priority of real-time flows has no effect on the packet loss ratios. The non-real-time buffer should be much larger (e.g., 1 Mbytes) because of the packet scale fluctuations in typical non-real-time traffic processes. Moreover, large buffers make it possible to offer reasonable service for those flows that are capable of adjusting their bit rate.

It should be emphasised that the function of each scheduling and buffering unit (SBU) is independent of all other SBU's; all the tasks of SBU are performed based on the information of incoming cells (or packets), and moreover, all the necessary function for the implementation are described in Fig. 4. Thus, due to the autonomous property of switching units and the unnecessity of resource reservation, the management of the SIMA network is very straightforward.

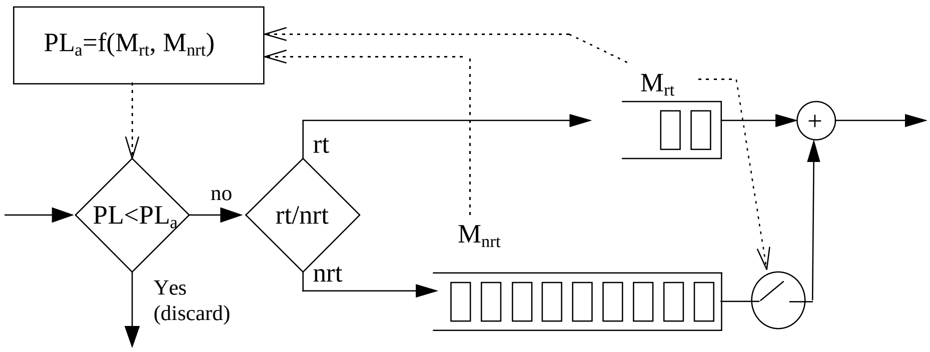

Fig. 4. A packet (or cell) scheduling and buffering unit (SBU) for a SIMA network node

5. Performance aspects

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 large 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 large amounts of users are willing to pay the additional charge.

Let us assume 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 used: 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 α we obtain: $\alpha_{rt} = 0.025$ and $\alpha_{nrt} = 0.001$.

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 effects 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
6	14.0	0

5	24.3	25
4	23.5	25
3	21.5	25
2	16.8	25

In Fig. 5, there is shown a graph illustrating the average cell loss ratio, P_{loss}, as a function of priority level for four specific load levels, r. 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.

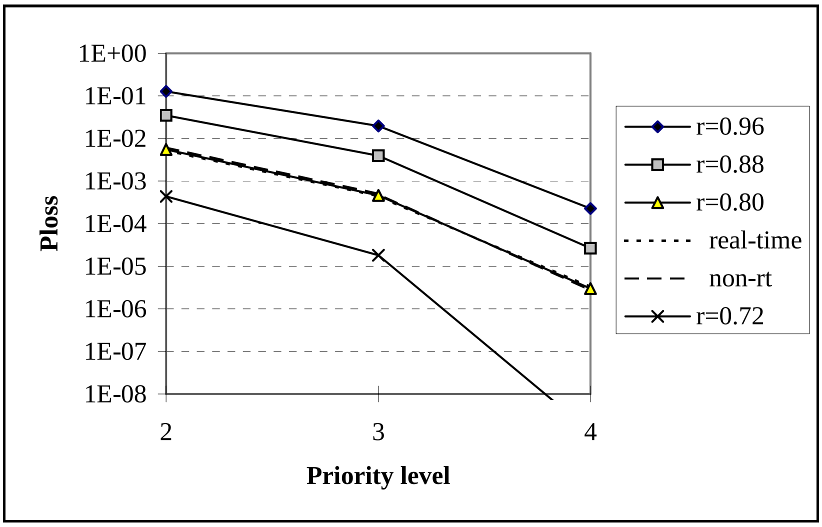

Fig. 5. Average cell loss ratio vs. priority level for load levels r = 0.72, 0.80, 0.88, 0.96. In case of load = 0.80 the cell loss ratios for real-time and non-real-time cells are presented by dotted and broken lines.

It should be emphasised that the difference in cell loss ratio between adjacent priorities depends strongly on the offered traffic process and, in particular, the inherent control loops. 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 behaviour 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 the 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 provides that:

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

$$P_{loss,7} = P^*_{loss,7}$$

$$P_{loss,k} = \frac{\rho^*_k P^*_{loss,k} - \rho^*_{k+1} P^*_{loss,k+1}}{\rho^*_k - \rho^*_{k+1}} \qquad for \quad k = 6 \ldots 0 \tag{4}$$

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. 6, there is illustrated in graphical form a relationship between cell loss ratio as a function of priority level for different load levels, r. Fig. 6 shows the cell loss ratios obtained by application of Equation (5) for different priorities. It is assumed in Fig. 6 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.

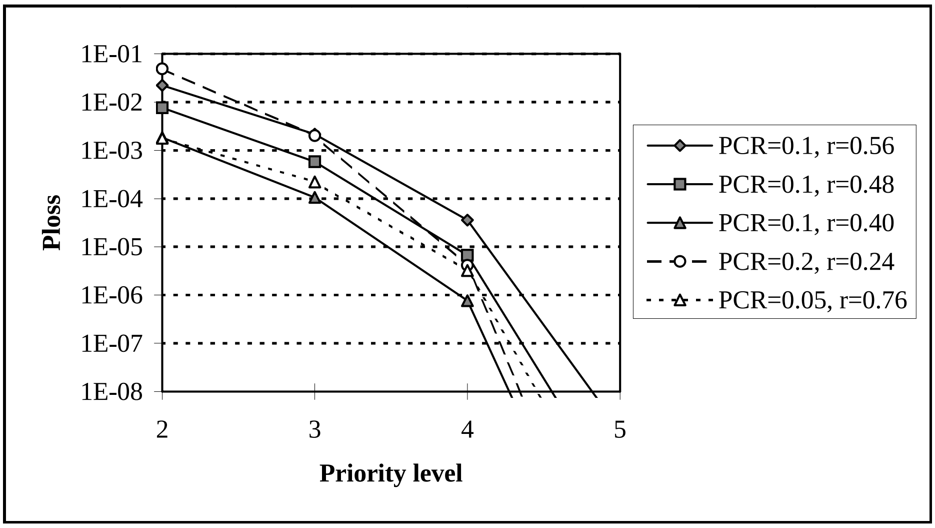

Fig. 6. Cell loss ratio vs. priority level for different load levels (r);
solid lines: peak cell rate of each connection = 0.1; broken line: peak cell rate of each connection = 0.2;
dotted line: peak cell rate of each connection = 0.05.

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. 5 and 6, 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 3 depending on the overall traffic variations. The cell loss ratio with priority 5 can be supposed to be less than 10^{-9} unless the traffic variations are very pronounced.

6. Conclusions

Notwithstanding the complexity of conventional traffic management schemes, the current ATM and Internet 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 paper is capable to fulfil 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 or packet, and separation of real-time and non-real-time connections at the buffer level. If a user needs a connection over an 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 to 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 high speed core network. In consequence, the SIMA concept is a very promising scheme for solving the most acute traffic control problems in Internet.

References

[1] ITU-T Recommendation I.371, Traffic control and congestion control in B-ISDN. International Telecommunication Union, 1996, 90 p.

[2] Traffic management Specifications v. 4.0, ATM Forum, af-tm-0056.000, 1996.

[3] http://www.ietf.org/html.charters/intserv-charter.html

[4] http://www-nrc.nokia.com/sima/

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