Dynamic multiplexing and UPC renegotiation for MPEG video over ATM networks

M. T. Andrade, A. P. Alves

{mandrade, palves}@inescn.pt

INESC Porto, Praça da República, 93 R/C, Porto, Portugal

FEUP, DEEC, Rua dos Bragas, Porto, Portugal

telephone: +35122 2094238 fax number: +351 22 2084172

Abstract

This paper presents an experimental approach for QoS-aware and bandwidth-efficient transmission of multimedia sources over ATM networks. VBR video sources are statistically multiplexed at the application level. Usage Parameter Control (UPC) values of a single connection for the aggregate traffic are dynamically calculated and renegotiated with the network during the session lifetime. New sources join the system by identifying their desired quality of service from the user perspective. A statistical multiplexer dynamically assigns bandwidth to each source according to their requests and network resources availability. A dynamic UPC manager negotiates the traffic contract of the connection for the multiplexed stream. Based on feedback information sent by the network regarding resources availability and on the bandwidth requested by the statistical multiplexer, the UPC manager may initiate a renegotiation process with the network to adjust the traffic parameters to the present needs of the aggregate traffic. Performing the connection admission control procedures for the aggregate traffic, allows significant reductions in the value of the total peak rate when compared to the sum of peak rates of individual connections. It also reduces the burstiness of the flow submitted to the network, thus increasing the lifetime of each set of UPC parameters.

Keywords: Quality of Service, ATM, UPC, statistical multiplexing, renegotiation

List of abbreviations:

AAL, ATM Adaptation Layer
ABR, Available Bit Rate
ATM, Asynchronous Transfer Mode
CAC, Connection Admission Control
CBR, Constant Bit Rate
CDV, Cell Delay Variation
CDVT, Cell Delay Variation Tolerance
CTD, Cell Transfer Delay
EFCI, Explicit Forward Congestion Indication
FTTC, FTTN, FTTH, Fiber To The Curb/Node/Home
GOP, Group Of Pictures
MBS, Maximum Burst Size
MCR, Minimum Cell Rate
MPEG, Motion Picture Expert Group
NIC, Network Interface Card
PCR, Peak Cell Rate
QoS, Quality of Service
RM, Resource and Management
SCR, Sustainable Cell Rate
UNI, User Network Interface
UPC, Usage Parameter Control
VBR, Variable Bit Rate
VC, Virtual Circuit
VoD, Video on Demand

Introduction

Deployment of high-quality networked multimedia services remains an important challenge in the design of high-speed networks and applications. The continuing growth both in consumer demand (in variety and quality) and providers choice, requires efficient methods for resource allocation and quality of service guarantees. The MPEG video compression algorithms are able to provide high-quality services while significantly reducing the number of bits necessary to represent the information. In the case of video sources this reduction is obtained by exploiting spatial and temporal redundancy, which will vary in time as the content of video sequences changes. For that reason, constant quality of MPEG compressed video streams will only be possible with variable bit rate (VBR) encoding.

Transmission of constant quality compressed video streams and therefore variable bit rate streams, is possible in broadband networks using the Asynchronous Transfer Mode. Traditional transmission channels offer only constant bit rate connections which either leads to an inefficient bandwidth utilisation or to varying picture quality. For the compression of video sequences, high-activity scenes will require more bits to achieve the same quality level as low-activity scenes. The same applies when scene changes occur. The use of VBR channels provides a way to cope with these time-varying requirements together with efficient network resource allocation. However, the use of VBR connections, characterised in ATM networks by the traffic parameters PCR, SCR, MBS and CDVT [AF], it is not by itself an assurance that the requirements of the bit stream will be satisfied for the whole duration of the call together with efficient network bandwidth utilisation. Redundancy of video sequences is considerably variable especially between different scenes. Traffic parameters chosen at connection set-up to closely match the initial characteristics of the video stream, may no longer be suitable after a scene change. Also, within the same scene, there will be pictures exhibiting more activity/motion than other pictures. These fast-rate changes, *intra-scene variations*, present smaller fluctuations when compared to the less frequent *inter-scene variations* referred previously.

A dynamic UPC with adjustable renegotiable parameters is able to contribute to a more efficient use of bandwidth while meeting the requirements imposed by the low-frequency bit rate variations. Because the renegotiation process imposes a rather significant overhead to the network and because there is latency in the response from the network, the renegotiation frequency should be kept as small as possible. Considering that the UPC negotiates a joint connection, the minimum inter-renegotiation interval can be set to the duration of the smallest scene in the aggregate traffic or to the mean duration of scenes.

To deal with the intra-scene variations we use the statistical multiplexer which aggregates a collection of VBR video sources and dynamically assigns bandwidth taking advantage of the fact that independent video sources are highly uncorrelated. Bandwidth made available by one source in periods of low-activity can be re-assigned to other sources instantaneously demanding more bandwidth. The objective is to obtain increased overall statistical gain while improving the quality of service for an aggregation of video sources.

Statistical multiplexing algorithms have been extensively studied and used to allow savings in overall transmission bandwidth in packet networks. Literature in this field is quite representative of the work already developed [Ohta] … [Heym]. Statistical multiplexing has been traditionally applied to the use of a fixed-bandwidth channel by a group of VBR sources [Perk] and [CLI] and inside ATM networks to efficiently support simultaneously a number of different connections [Veci] and [Ross]. More recently, researchers have started to study the potential gain, both in effective bandwidth and quality improvement, in the transmission of video signals over packet networks using adjustable UPC and traffic parameters renegotiation [Rein] … [Giord: 2]. However the effects of using the two techniques combined have not yet been sufficiently explored. In [Fulp], [Noz] and [Teix], the authors present approaches and results which demonstrate the benefits of multiplexing a set of video sources together with renegotiation of the traffic parameters of a single network connection for the aggregate stream. In our work we try to jointly explore the benefits of performing statistical multiplexing at the user level and dynamic adjustable UPC based on the following premises:

- independent video sources are highly uncorrelated;
- constant quality of compressed video is only possible with VBR;

- constant quality compressed video presents two types of rate variability: low-amplitude/high-frequency intra-scene variations and high-amplitude/low-frequency inter-scene variations (respectively short-term and long-term in [Ohta]);
- scene changes occur independently [Ohta];
- renegotiation of traffic parameters provides a way of matching more closely the application requirements but it imposes a significant overhead to the network and consequent latency in the response of the network to those requests;
- a reduction in the number of renegotiation requests reduces the blocking probability of the network (or improves the network performance).
- a customer can combine in a single ATM connection several video channels and share at his will the overall requested bandwidth while having to control one single network access unit;

Equally important is the existence of applications that may actually benefit from the use of the proposed approaches. Figure 1 shows a possible application scenario, where new network architectures such as the combination of FTTC or FTTN with xDSL or complete FTTH extend the broadband capabilities closer to the end-user.

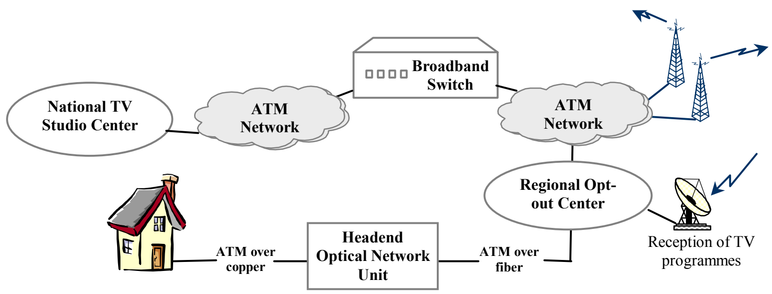

Figure 1 – Application scenario

Examples of areas of potential application for the combined use of statistical multiplexing and dynamic renegotiation of traffic parameters are:

- the TV broadcasting industry for programme contribution, distribution and TV studio pre and post-production [Atl], [Vid];
- infotainment applications where users combine and interact with distinct video sources and other types of media, in a completely unstructured way;
- VoD services where ATM backbones can be used to interconnect remote islands of users to the central server (where films are stored

Description of the experiments

Overview

In order to be able to devise a methodology for the estimation of values for aggregate statistical bandwidth and UPC + traffic parameters, we have used traces of a number of MPEG encoded video sequences. Compression video encoding algorithms, such as those described in MPEG standards [MPEG1,MPEG2], operate on a frame basis and use 3 different picture types: I (Intra), P (Predicted) and B (Bi-directional). Encoded pictures, belonging to one of the 3 possible types, succeed repeatedly in a structured way, the GOP structure. The amount of bits generated at the output of the encoder reflects this structure. Typically, within this GOP structure, I frames require more bits while B frames are those which require less bits. The traces we have used are publicly available by Oliver Rose from the University of Wuerzburg (*ftp site: ftp-info3.informatik.uni-wuerzburg.de/pub/MPEG*). Each encoded sequence presents the following characteristics: *40 000 frames long at 25 Hz, with spatial resolution 384x288; a GOP structure with N=12 and M=3 (**I***

*B B P B B P B B P B B **I** B…); VBR encoded with fixed quantisation step sizes of 10, 14 and 18 respectively for I, P and B frames*.

Conceptually our system comprises four hierarchical layers from the user/application level to the network. Figure 2 represents the system architecture, identifying the tasks to be performed in each of the first three levels.

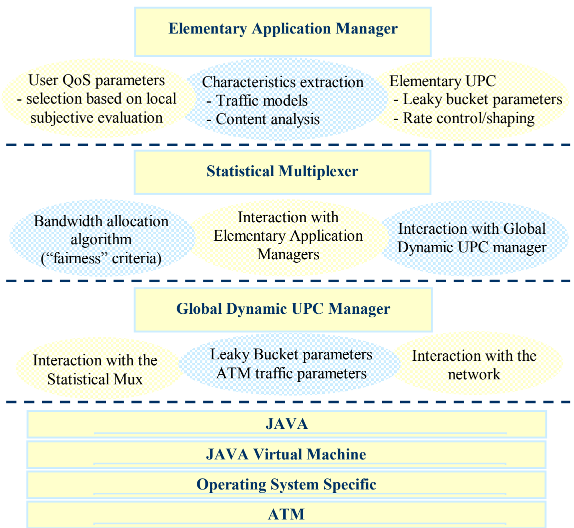

Figure 2 – Conceptual system architecture

All modules implement closed-loop control algorithms in the sense that they all react from information coming from another module. The *Statistical Multiplexer* and the *Global Dynamic UPC* can be seen as dual closed-loop control algorithms. Interaction among these components is performed in both directions: from the source of generation of bits down to the network and from the network towards the source. Using this approach, applications should connect themselves to the system via their application manager, supplying QoS parameters from the user perspective: target picture quality (TPQ), minimum picture quality (MPQ) and service type (ST). With these parameters and a sample analysis of the bit stream to extract associated peak and mean bit rates, the application manager selects the parameters of the elementary UPC and passes the bit rate information and the user MPQ, TPQ and ST parameters to the statistical multiplexer. This entity estimates necessary values of peak and mean bit rates to allow the transmission of the aggregate traffic from all active sources. It passes this information to the UNI, which then calculates the appropriate parameters of a global UPC capable of monitoring the aggregate traffic. It also specifies the ATM QoS and traffic parameters to request an ATM channel adequate to support the transmission of the multiplexed stream.

The *statistical multiplexer* implements a closed-loop bandwidth allocation algorithm reacting to requests from the sources and to feedback information from the global dynamic UPC. Likewise, the *UPC managers*, both the elementary and the global, use a dual closed-loop parameters adjustment algorithm. The *elementary UPC* interacts with the real-time VBR video encoder and with the stat mux. The *global UPC manager* interacts with the stat mux and with the network, receiving feedback information regarding network resources availability.

Real-time VBR video encoders also use a closed-loop rate control mechanism to increase or decrease the rate of generation of bits. The rate adaptation is performed by accordingly modifying the values of the quantisation step sizes. This is only possible for live video sources or through the use of trancoders for pre-recorded sources. VBR *on-line* video sources, i.e. video sources being encoded as they are being transmitted, have their rate of generation of bits constraint, such that the current picture quality is maintained within two pre-defined thresholds. Pre-encoded and stored material, *off-line* video sources may have not been subject to these constraints and therefore may present higher levels of rate variability.

Our experiments have been restricted to the *off-line* case and assuming that the requested network bandwidth was always available. We have only computed required values for traffic parameters and UPC re-negotiation frequency imposed by the bit rate requirements of the sources. Further experiments will be implemented once our system becomes complete and operational. In real implementations, the total bandwidth necessary to completely satisfy all requests will not always be available. In those cases it will be necessary to use rate conversion mechanisms able to interact with *off-line* sources to reduce the number of bits. For example, transcoders or filters may perform the necessary rate adaptation by modifying the quantisation step sizes or by selectively eliminating DCT coefficients.

Elementary Application managers

Elementary sources can be either real-time on-line encoded video or pre-encoded and stored video. In the later case the application manager extracts from the stored data, information concerning bit rates and service type. This information is directly used to select the parameters that better match the characteristics of the source. In the former case, there are two alternative methods to obtain the same information: manually, with the intervention of a human, supplying the desired user QoS parameters and bit rate information; automatically, with the application manager performing an analysis of a sample number of frames of the video source and using a parameterised set of values for the objective picture quality.

So, the first operation executed once a new source becomes active is the definition of the desired user QoS and bit rates. With this information, the application manager calculates the parameters for a flow rate monitoring mechanism. This monitoring is performed by an elementary dual Leaky Bucket mechanism. The dual LB implements the UPC for each elementary VBR source, the way as proposed by the ATM FORUM [AF]. It comprises two LB, each one parameterised by a drain rate and a buffer size. Each LB can be implemented through a finite-capacity buffer, the token buffer, with defined sizes and drain rates, B and R, according to the desired peak and sustainable rates, length of bursts and spacing between bursts of the stream. One of the LBs monitors the peak rate with parameters R_p and B_p and the other monitors the sustainable rate with parameters R_s and B_s. The values of the LBs parameters are related with the traffic parameters in the following way:

$$R_p = PCR \qquad B_p = 1$$
$$R_s = SCR \qquad B_s = R_p \times T - b$$
$$MBS = R_p \times T \quad ; \quad MBS = int[1 + \tau_s/(T_s - T_p)] \quad ; \quad MBS = 1 + int\,[(\,mín\,(T_1 - T_s\,,\,\tau_s)/(T_s - T_p)]$$

with,

b = number of tokens drained from the token buffer B_s in a frame period T;
T_1 = minimum spacing that bursts of size equal to MBS must observe to be in conformance;
τ_s = burst tolerance.

The dual LB constraints the source such that its maximum bit rate does not exceed the negotiated peak rate PCR and its average rate does not exceed the negotiated sustainable rate SCR and that the length of bursts (number of consecutive cells transmitted with the minimum inter-cell spacing $T_p = 1/R_p$) does not exceed $MBS = (R_p * B_s)/(R_p - R_s)$. The token buffers, initially set to their capacity B_p or B_s, are continuously drained at a constant rate R_p or R_s and filled at the rate of submission of cells to the network. A cell may only be submitted to the network if it is able to add a token into the token buffer. If the buffer is full, the cell is not in conformance with the UPC and therefore should not be transmitted. The traffic parameters PCR, SCR and MBS and consequently the UPC values, are continuously evaluated for each GOP and their values set to:

$$PCR \equiv \textit{bit rate of the largest picture in the GOP (in cell/s)}$$
$$SCR = (1 - k) * R_{avg} + k * PCR \quad \text{with } 0 < k < 1$$
$$R_{avg} = \textit{weighted sum of the bit rates of all but the largest frame in the GOP (in cell/s)}$$

If significant changes in these values are detected, the application manager contacts the statistical multiplexer, requesting a modification in the allocation of bandwidth.

For stored VBR video, we have performed experiments using accompanying metadata files with the following information: average and peak bit rates for each scene, initial and ending times of each scene (or the period during which the previous values are valid), frame period and GOP structure. If this metadata is not available, the application manager must calculate them throughout the session lifetime or completely at connection set up.

Statistical multiplexer

Statistical multiplexing, the ability of dynamically allocating bandwidth among sources according to picture activity and target quality, is based on the theory that independent video sources are likely to exhibit the same amount of activity simultaneously and therefore, the need for the same amount of bits. Statistical multiplexing aims at providing better channel utilisation and better picture quality across a number of video sources. The way it proposes to achieve this goal is to dynamically distribute the overall channel bandwidth, assigning to each source variable amounts of bit rates according to their present activity (complexity and motion).

We have made calculations using the video traces referred above to obtain values of required bandwidth, statistical gain and UPC + traffic parameters. Different transmission scenarios have been considered:

- sources transmitted separately with a fixed set of UPC parameters for the whole stream;
- as above, but using renegotiation of UPC parameters;
- as above, but using a GOP buffer to smooth the elementary bit rate prior to submit the stream to the network;
- association of different number of sources, assuming they are GOP-aligned and using a fixed set of UPC parameters;
- as above but assuming sources are not GOP-aligned;
- as above but using renegotiation of UPC parameters.

Table1 presents the elementary sources behaviour in terms of mean and peak bit rates and degree of variability in those bit rates. Statistics performed at the GOP level assume that each application manager has a buffer of sufficient capacity to be able to send all but the largest picture in the GOP at the sustainable cell rate. The picture with the highest bit rate observed in the GOP, is sent at that maximum rate, the peak cell rate. The quantity *avg bit rate* on a GOP basis is calculated using the bit rates of all but the largest picture in the GOP, while the *max bit rate in GOP* is set to the bit rate of the largest picture

	Frame basis		GOP basis		Frame basis		GOP basis			
	avg	max	avg	avg of max			in avg rate		in max rate	
sources					25%	50%	25%	50%	25%	50%
Asterix	0.559	3.68	0.448	1.78	27042	18954	1049	380	338	104
Dino	0.327	2.00	0.232	1.38	25886	18496	770	244	179	42
Simpons	0.588	3.52	0.487	1.7	27185	21238	683	219	275	69
Lambs	0.183	3.35	0.113	0.95	26058	18360	1050	411	179	47
Movie	0.357	4.31	0.259	1.44	28407	26970	1460	729	612	253
Bond	0.608	6.12	0.474	2.08	26980	26618	439	126	212	46
Starwars	0.233	3.12	0.154	1.10	26591	19930	1331	594	219	50
Terminator	0.273	1.99	0.935	1.99	28663	23322	1321	466	463	92
mTV-1	0.615	5.73	0.512	1.75	28195	26429	1213	505	653	204
mTV-2	0.495	6.29	0.4	1.54	27207	20124	1198	504	596	213
Race	0.769	5.06	0.659	1.98	25335	16256	511	119	272	61
ATP-tour	0.547	4.77	0.425	1.89	26827	19433	659	217	263	99
Soccer	0.678	4.68	0.56	1.98	27269	26752	427	126	273	86
Sbowl	0.588	3.52	0.487	1.7	27185	21238	683	219	275	69
	bit rate values (Mbit/s)				variations in bit rate					

Table 1 - Statistics of elementary sources

Table 2 characterises the behaviour of aggregate traffics in the same terms as for individual sources. It presents measures for combinations of different number of sources and assuming alignment or not at the GOP level. It also compares the bit rates required by those multiplexed bit streams with the bit rates required by the sum of corresponding individual sources From the statistics presented in those tables, it can be seen that using the same amount of rate variability as the threshold for initiating a renegotiation, the number of requests significantly decreases when more than 3 sources are multiplexed before accessing the network. It can be noted that the inclusion of extra sources in the multiplex contributes on average with a reduction of 5% in the number of renegotiation requests. The values of rate variability used to trigger the renegotiation are 25 and 50% of change relative to the values in use for the traffic parameters.

The estimation of the initial bandwidth required for the multiplexed stream uses an algorithm that assumes sources are not aligned at the GOP level and that a reduction from 10% to 5% of the sum of bit rates is achieved for each new source joining the multiplex. These values were obtained in the measures performed with the MPEG traces. Information about activity and bit rate requirements to achieve a certain quality level relative to each elementary application, is then used to effectively distribute the negotiated bandwidth. The statistical multiplexer continuously performs the bit rate calculations. If significant changes occur, on an n-GOP basis, it requests the dynamic UPC manager to alter the global UPC parameters. In that situation, the UPC manager initiates with the network, negotiations for a new traffic contract. After receiving a confirmation of the acceptance or rejection of the new traffic parameters, the UPC manager notifies the statistical multiplexer and adjusts its own policing parameters. If the new values were accepted, the statistical multiplexer distributes the bandwidth as requested. Otherwise it re-distributes the available bandwidth based on a fairness criterion, regarding values of bit rate requested and target quality level.

Number of sources			GOP period basis						
			Average rate			Maximum bit rate			
			t rate value	variations		g of peak rate	ax of peak ra	variations	
				25%	50%			25%	50%
all	Sum		5.42	12794	4858	22.40	57.44	5001	1434
	mux	A	5.42	22	1	22.87	30.97	0	0
		NA	6.56	8	1	8.35	14.35	29	3
12	Sum		5.06	10142	3799	20.17	54.03	4127	1293
	mux	A	5.06	31	1	20.17	27.99	3	1
		NA	5.945	20	1	7.641	13.66	108	7
8	Sum		2.82	7838	3116	12.81	37.25	3032	1008
	mux	A	2.82	247	17	12.81	19.252	23	1
		NA	3.54	86	3	5.345	11.104	144	29
6	Sum		1.949	1846	2500	9.522	25.234	1846	614
	mux	A	1.949	247	24	9.522	14.747	40	1
		NA	2.469	92	7	3.813	8.307	218	32
4	Sum		1.587	3938	1570	7.130	1756	1392	518
	mux	A	1.587	279	28	7.130	11.287	62	3
		NA	1.927	145	15	3.295	6.934	165	20
2	Sum		0.537	1819	624	3.863	9.8	517	146
	mux	A	0.537	538	98	3.863	8.386	154	15
		NA	1.038	427	60	2.579	6.549	173	21

Table 2 - Statistics of combined sources on a GOP period basis

When the bandwidth granted is not enough to completely satisfy all requests, the allocation algorithm starts by assuring that the minimum picture quality is maintained in all sources and then distributes the remaining bit rate according to a ratio between bandwidth requested and target picture quality.

Fairness allocation bandwidth algorithm

The statistical multiplexer uses a *fairness allocation bandwidth algorithm* based on the following definitions, measures and operations:

$$LB_{requested} = \sum_{i=1}^{N} \max_i (1) \times (1-k),$$

where:

Number of Sources	k	Number of Sources	k
2	0.30	8	0.70
4	0.50	12	0.75
6	0.60		

Table 3 – Experimental values of multiplexing gain for peak rate

$LB_{network}$: bandwidth actually granted by the network
$max_i(1)$: maximum bit rate requested by source i to achieve the target picture quality level, QoS(1)
$max_i(2)$: maximum bit rate requested by source i to achieve the minimum picture quality level, QoS(2)
k : defines the multiplexing gain and depends on the number of sources in the multiplex. Values for this parameter, obtained experimentally using the mentioned traces, are given in table 3

Pseudo C code describing the fairness allocation algorithm:

```
if (LBnetwork ≥ LBrequested)                                        /* assign bandwidth as requested
    for(i = 1 ; i < numberOfSources ; i++)        maxiassigned = maxi(1);
```

```
if (LB_network < LB_requested)                                          /* applies reduction factor
  {
    reduction_1 = (LB_network - LB_requested) / LB_requested;
    for(i = 1 ; i < numberOfSources ; i++)
        {
          if((1-reduction_1 × max_i(1)) < max_i(2))     max_i assigned = max_i(2);        /* assign at least to guarantee QoS(2)
          else  max_i assigned = (1-reduction_1) × max_i(1) ;
        }
    Σmax_i assigned = 0 ;
    for(i = 1 ; i < numberOfSources ; i++)              Σmax_i assigned += max_i assigned ;
    while((1-k) × Σmax_i assigned > LB_network )
    {
        reduction_2 = (Σmax_i assigned - LB_network) / LB_network;
        for(i = 1 ; i < numberOfSources ; i++)
        {
            if(max_i assigned > max_i(2))
            {
                if((1-reduction_2) × max_i assigned < max_i(2)) max_i assigned = max_i(2);        /* assign at least to guarantee QoS(2)
                else     max_i assigned = (1-reduction_2) × max_i assigned ;
            }
        }
        Σmax_i assigned = 0 ;
        for(i = 1 ; i < NumberOfSources ; i++)        Σmax_i assigned += max_i assigned ;
    }
  }
```

The set of calculations and operations are indicated for the distribution of the peak bit rate of the aggregate connection. It is also valid for the mean bit rate except that the multiplexing gain indicated would not be used. The case we are considering, assumes that each elementary source, if transmitted isolated, would pass through a GOP-buffer before accessing the network.

Dynamic renegotiable UPC

In ATM networks, the establishment of a new connection must be preceded by the negotiation of a service contract between the new application and the network. The contract must identify the ATM service attributes requested from the network, through specification of traffic parameters (PCR, CDVT, SCR and MBS) and QoS parameters (peak-to-peak CDV, max CTD and CLR).
If the network is not able to support stating of QoS in terms of individual numerical values for the negotiable QoS parameters, it will support its indication through QoS specified classes. QoS specified classes identify a set of ATM performance parameters (CTD, CDV and CLR) and objective values for those parameters. The user must only select the specified QoS classe he wants to negotiate. Once the connection is accepted, the network will provide the requested QoS for those cells that conform to the negotiated traffic values. If the user violates this traffic contract, the network no longer has the responsibility of maintaining the agreed QoS. It is therefore the interest of the user, to ensure that all cells submitted to the network do not exceed the negotiated values.

Conformance of the flow submitted to the network is monitored by an UPC mechanism implemented through a dual Leaky Bucket algorithm. The operation of the LB mechanism has already been described in the subsection "*Elementary Application manager*". The global UPC manager sets its parameters according to the request sent by the statistical multiplexer. At connection set up, it negotiates with the network a traffic contract specifying the ATM traffic parameters PCR, SCR and MBS for the aggregate stream and selecting the specified QoS class 2. This specified QoS class supports a QoS that meets the performance requirements of the ATM service class suitable for real-time VBR video and audio applications [AF]. The value of the PCR is the value of the maximum bit rate observed in a frame period for the current GOP period. The values of the SCR and MBS are estimated based on the average bit rate of the current GOP period, the duration of bursts and the spacing between bursts. Formulas for calculating these parameters have already been presented in subsection "*Elementary Application managers*".

For each GOP period consisting of 12 frame periods (T_{frame} = 40 ms), it is assumed that the PCR is used during a frame period and the SCR during the rest of the GOP. This means that we have a constant spacing of 11 frame periods between bursts. The values of the global UPC are therefore calculated in order to match the requirements of the multiplexed stream as indicated by the statistical multiplexer. When a request for altering those parameters is received, the dynamic UPC before actually updating them, interrogates the network about the availability of resources. Table 4 presents values of individual UPC parameters calculated according to the formulas and assumptions referred above. Table 5 presents the same kind of information for aggregate traffics.

	Ravg		PCR		SCR (k = 0.1)		MBS		Bs			Ravg		PCR		SCR (k = 0.1)		MBS		Bs	
	mean	stdev	mean	stdev	mean	stdev	mean	stdev	mean	stdev		mean	stdev	mean	stdev	mean	stdev	mean	stdev	mean	stdev
asterix	1174.3	667.1	4739.1	1170.4	1530.7	685.0	189.6	46.8	128.3	31.3	*mTV1*	1660.9	669.5	4544.7	1510.8	1949.3	723.0	181.8	60.4	115.3	44.0
atp	1566.1	632.6	5127.1	1543.6	1922.2	689.2	205.1	61.7	128.2	42.2	*mTV2*	1244.5	553.0	3859.0	1031.8	1506.0	572.4	154.4	41.3	104.6	31.0
bond	1160.6	449.4	3993.9	1357.5	1443.9	527.2	159.8	54.3	113.3	39.6	*race*	2022.9	891.0	5198.4	1425.1	2340.5	922.3	207.9	57.0	127.0	34.5
dino	839.0	318.8	3412.9	726.8	1096.4	334.8	136.5	29.1	103.0	23.9	*sbowl*	1361.5	553.0	4224.4	1322.0	1647.8	579.5	169.0	52.9	114.5	45.1
socc	1775.1	704.1	5022.4	1619.6	2099.8	764.7	200.9	64.8	129.9	46.9	*simps*	1374.9	592.2	4901.9	1416.2	1727.6	636.9	196.1	56.6	141.1	44.3
lamb	531.4	246.2	2705.8	791.7	748.9	280.2	108.2	31.7	87.0	26.2	*starwars*	834.7	456.2	3316.5	1242.9	1082.9	509.6	132.7	49.7	99.3	38.2
movie	954.2	443.2	3666.2	1243.3	1225.4	491.0	146.6	49.7	108.5	40.1	*termin*	722.4	274.3	2567.8	614.8	906.9	285.6	102.7	24.6	73.8	20.6
	Ravg, PCR, SCR in cell/s ; MBS and Bs in cells																				

Table 4 UPC parameters for individual sources

Our approach to implement the interaction between the dynamic UPC and the ATM network is to use RM cells in a similar way as specified for the ABR service category [AF]. The network access interface of the multiplexed source sends to the network RM cells with the indication of target bit rates. There are no values specified for the transmission frequency of those cells but typically, reported implementations send one RM cell for each group of 48 service cells. In response, the network sends feedback information in backward RM cells. This feedback consists of information about network bandwidth availability and state of congestion and can be explicitly stated in the fields *Explicit Rate*, *Congestion Indication* and *No Increase* of backward RM cells. It is not possible to report results in this field, namely of performance and feasibility of the proposed method and network blocking probability to renegotiation requests, because experiments have not yet been performed. It is however the intention of the authors to proceed the work in this direction.

Results and further work

We have evaluated our experimental system using traces of pre-encoded MPEG sequences. We have thus only implemented and off-line experiment of our proposed methodology, where the entire bit streams to be submitted to the network were known in advance. Although having an arguably restricted field of application, this experience has revealed itself very important:

- it provided a way of evaluating the effectiveness of the proposed system:
- it provided valuable information to infer the methodology to use with on-line cases, where video sequences are being real-time encoded as they are being submitted to the ATM network.

We are developing an API in Java to provide applications the kind of functionality we have described for the *Elementary application managers* combined with the dynamic renegotiation of UPC parameters. Presently, this API provides only basic functionality for native ATM applications to request an ATM connection and specify ATM QoS and traffic parameters. In our premises at INESC, we have installed a 155 Mbit/s private ATM network with two FORE switches and a number of end stations equipped with 25 and 155 Mbit/s Fore ATM NICs. Using this infrastructure, we have set up a platform to test several TV studio applications developed within the framework of research projects [Atl] and [Vid]. The platform allows the experimentation of the application scenario "*TV studio pre and post-production*". This scenario involves both pre-encoded material, thus the *off-line* case, as well as *on-line* real-time encoded sources. Our work will evolve towards the experimentation of the *on-line* case. For that matter we will use, once finalised, the full-featured API for native-ATM QoS-aware applications within the TV studio applications scenario.

The results we have obtained so far can be extracted from the presented tables. They show us that it is possible to obtain statistical gain ranging from nearly 75% to 15% for the same picture quality objectives, depending on the number of sources being multiplexed. Also, that the frequency of renegotiation, initiated by the occurrence of the same amount of bit rate variation, is significantly smaller when sources access the network in a multiplex rather than individually. We have also calculated initial values of traffic parameters for individual streams and aggregate flows using the information of tables 1 and 2 and the formulas presented.

We are also extending the statistical characterization of compressed video sequences to a content analysis. Our objective is to correlate statistical models with content classification to build more complete models that may describe and predict more accurately the behaviour of video sources. In the network side we will further investigate the possibility and usefulness of using RM cells to implement the UPC-network interaction.

References

[AF] ATM Forum Technical Committee: Traffic Management Specification Version 4.0 (1996)
[Ohta] Naohisa Ohta, "Packet Video: Modelling and Signal Processing", Artech House, Boston (1994)
[Atl] ACTS project 078, ATLANTIC, Working documents, http://www.bbc.co.uk/atlantic
[Vid] VIDION project, http://newton.inescn.pt/projects/vidion
[Rose] O. Rose, "Statistical properties of MPEG video traffic and their impact on traffic modelling in ATM systems", University of Wuerzburg, Institute of Computer Science (1995)
[Mich] H. Michiel, K. Laevens, "Teletraffic Engineering in a Broad-Band Era", Proc. IEEE, Vol. 85, No. 12 (1997)
[Krunz] M. Krunz, R. Sass and H. Hughes, "Statistical Characteristics and Multiplexing of MPEG streams", Michigan State University, Department of Electrical Engineering
[Magla] B. Maglaris et al., "Performance Models of Statistical Multiplexing in Packet Video Communications", IEEE JSAC, Vol. 36 (1988)
[Heym] D. Heyman, A. Tabatabai, T. Lashkman, "Statistical analysis and simulation study of video teleconferencing traffic in ATM networks", IEEE Transactions on Circuits and Systems for Video Technology, Vol. 2 (1992)

[Rein] D. Reininger, Y. Senda and H. Harasaki, "VBR MPEG-2 Encoder for ATM networks with UPC renegotiation", NEC USA, C&C Research Laboratories (1996)
[Rein:2] D. Reininger, D. Raychaudhuri and J. Hui, "Bandwidth Renegotiation for VBR Video Over ATM Networks", IEEE Journal on Sel. Areas in Comm., Vol. 14, No. 6 (August 1996)
[Rein:3] D. Reininger, D. Raychaudhuri and R. Siracuse, "Video Transport in ATM networks: a systems view", NEC USA, C&C Research Laboratories (1995)
[Rein:4] D. Reininger, G. Michelitsch, M. Ott, G. Welling, "An architecture for QoS aware browsing in distributed multimedia systems", NEC USA, C&C Research Laboratories (1996)
[Mark] B. Mark and G. Ramamurthy, "Joint source-channel control for real-time VBR over ATM via dynamic UPC renegotiation", NEC USA, C&C Research Laboratories (1996)
[Mark:2] B. Mark and G. Ramamurthy, "Real-time estimation and dynamic renegotiation of UPC parameters for arbitrary traffic sources in ATM networks", IEEE/ACM Transactions on networking, Vol. 6, No. 6 (Dec. 1998).
[Giord] S. Giordano and J. Le Boudec, "The renegotiable variable bit rate service: characterisation and prototyping", EPFL ICA, Lausanne, Switzerland (1998)
[Giord:2] S. Giordano and J. Le Boudec, "The renegotiable variable bit rate service", EPFL ICA, Lausanne, Switzerland (1998)
[Fulp] E. Fulp and D. Reeves, "Dynamic bandwidth allocation techniques", North Carolina State University, Departments of ECE and CSC (1997)
[Noz] Y. Nozawa, H. Harasaki, Y. Senda and K. Kobayashi, "Statistical multiplexing gateway for piecewise CBR transmission channel", Proceedings of the Packet Video'99 Conference, New York (1999)
[Teix] L. Teixeira and M. Andrade, "Joint Control of MPEG VBR video over ATM networks", Proceedings of the ICIP 97, International Conference on Image Processing, Sta Barbara, CA, USA (1997)
[Perk] M. Perkins and D. Arnstein, "Statistical Multiplexing of MPEG2 video channels", Proceedings of the SMPT 136th Technical Conference and World Media Expo, Los Angeles, USA, (1994)
[CLI] Compression Labs white paper, "The next step in Digital Broadcast: statistical multiplexer enhances compressed video MPEG2"
[Veci] G. de Veciana, G. Kesidis and J. Walrand, "Resource Management in wide area ATM networks using effective bandwidth, IEEE JSAC, Vol. 13 (1995)
[Ross] K. Ross, V. Veque, "Analytic Models for Separable Statistical Multiplexing", University of Pennsylvania, Department of Systems Engineering (1994)
[MPEG1] D. Le Gall, "MPEG: A video compression standard for multimedia applications", Trans. of ACM (1991)
[MPEG2] ISO/IEC 13818-2, ITU-T Rec. H.262, "Generic coding of moving pictures and associated audio: video".