

Analysis of the Requirements for ATM Multicasting Based on Per-PDU ID Assignment¹

Josep Mangues-Bafalluy and Jordi Domingo-Pascual

Advanced Broadband Communications Center, Computer Architecture Department.

Polytechnic University of Catalunya.

Campus Nord, Mòdul D6, Sala 008.

Jordi Girona 1-3 – 08034 Barcelona (Spain)

{jmangues, jordi.domingo}@ac.upc.es

<http://www.ccaba.upc.es>

Abstract. The main issue to solve for ATM multicast forwarding is the cell-interleaving problem using AAL5. In strategies using a multiplexing ID, there are two main options: to assign an ID per source or per PDU. This paper argues for the convenience of using PDU ID because it allows the share of IDs between all the senders, thus reducing the ID consumption and management complexity. Compound VC (CVC) goes a step beyond in proposing a variable-length ID to adapt to the group size and to reduce the overhead introduced by the multiplexing ID. The size of this ID is negotiated at connection establishment according to the traffic and group characteristics. This paper provides some hints towards obtaining some dimensioning rules for the number of IDs based on its dependencies on traffic and group parameters.

1 Introduction

When dealing with ATM multicast forwarding, though ATM is cell-based, one expects the information in higher layers to be generated as packets. AAL5 seems to have been widely accepted and used for transmitting most data and multimedia communications over ATM. But AAL5 has no fields in its SAR-PDU allowing the multiplexing of cells belonging to different CPCS-PDUs. Therefore, the cell-interleaving problem must be solved in order to allow the receivers to reassemble the original packet when merging occurs.

One way to offer multicasting over ATM is IP multicasting over ATM [1], though it doesn't fully exploit ATM possibilities. A more efficient option is to provide multicasting mechanisms at the ATM level, that is, Native ATM Multicasting mechanisms. This term refers to those mechanisms implemented at the switches to allow the correct ATM level forwarding of the information being interchanged by the members of a group. There is no AAL5 CPCS-PDU reassembly inside the network.

A classification of these mechanisms can be found in [2], though it is slightly modified in [3] according to our conception of the mechanisms and to introduce our

¹This work has been funded by the Spanish Ministry of Education (CICYT) under grant number TEL99-1117-C03-03

proposal. Thus, table 1 presents the classification as we think it should be with some more new mechanisms than in [2].

Techniques belonging to the first type solve the cell-interleaving problem by avoiding cells from different packets to be interleaved. VC Merging techniques ([4], [5]) reassemble all the cells of each PDU in separate buffers and forward them without mixing cells of different buffers (or PDUs). SMART [6] uses a token passing scheme to allow just one sender to put data in the multicast tree at any instant. In the second type, the VPI identifies the connection and the VCI is used as the multiplexing ID (identifying the source [7] or the PDU[8]). Compound VC switching [3] also uses multiplexing IDs per packet, but it allows the negotiation of its length (see Section 3). And the last type allows multiplexing inside the same VC either by adding overhead in the transmitted data ([9],[10]) or by using the GFC field in the header of the ATM cell [11].

Table 1. Classification of Native ATM Multicasting mechanisms

Avoid Cell-interleaving		VP switching		Compound VC switching	Allow Multiplexing inside a VC	
VC Merging (Buffering)	Token control	Source ID	Packet ID		Added overhead	GFC
SEAM	SMART	Standard VP switching	DIDA	CVC	SPAM CRAM	Subchannel (WUGS)
VC Merge (MPLS)		VP-VC switching				

Those strategies using multiplexing IDs may be classified in two main groups: Source IDs and Packet Data Unit (PDU) IDs. In strategies using Source ID, the ID is related to the source that transmitted the packet. On the other hand, PDU ID strategies assign an ID (independent of the source) to each packet. The next section explains both philosophies in further detail.

This paper argues for the convenience of using PDU IDs instead of Source IDs for multiplexing, due to the efficiency in terms of ID consumption. The advantages of variable-length PDU IDs and the flexibility they allow in group size by providing minimum overhead are also discussed. The next step is to determine the correct size of the PDU ID according to the traffic characteristics and the group size. For that purpose, we carried out some simulations. Thus, the goal of this paper is to find a methodology to study the dependence on traffic and group characteristics of the PDU losses due to running out of identifiers at a given switch.

A discussion on Source ID and PDU ID advantages and drawbacks is presented in the next section. Following that, there is a brief description of the Compound VC mechanism, as it is the mechanism for which calculating the size of the PDU ID makes sense. Determining the size of the ID is the issue of the following section.

Section 5 discusses the results of our simulations. Finally, conclusions and future work are presented in the last section.

2 Source ID and PDU ID

The classification presented in table 1 is based on the way each strategy employs to solve the cell-interleaving problem. Except for the mechanisms labeled as ‘Avoid Cell-Interleaving’, which use buffers or a token, all the rest use some kind of multiplexing ID (muxID) to deal with this problem. The purpose of the muxID is to identify each particular cell so as to make the end-system able to correctly reassemble all the PDUs it receives.

Multiplexing IDs allow cells belonging to different PDUs to be interleaved, and thus the traffic characteristics of all the sources in the group are respected, making these strategies suitable for multimedia communications. However, in VC merging strategies, the buffer requirements and the increase in CDV and burstiness due to buffering limits its application to multimedia communications. The connection management complexity of SMART may also limit its application to multimedia communications. Therefore, time-constrained traffic may be more suitably served by allowing multiplexing of cells belonging to different PDUs. The price paid by muxID strategies is the extra overhead to carry the ID, which adds to the intrinsic ATM overhead.

In Source ID mechanisms, the ID is related to the source that transmitted the packet. Therefore, there is a binding between the source and the ID at each switch. This binding must be unique at each switch so as to avoid ID collision at merge points. The ID collision problem may be solved either by globally assigning IDs for the group or by locally remapping the IDs at each switch [7]. The management required in the former option may limit its scalability. On the other hand, local remapping maintains a list of free IDs at each switch, where a local mapping of IDs is carried out.

Source ID mechanisms usually overdimension the size of the ID so as to solve the worst case in which there could be a lot of senders (usually up to 2^{15} or 2^{16}). However, not all groups will have such a huge number of senders, and most overhead will be unused, e.g. in the local area.

PDU ID strategies assign an ID to each packet, and this assignment is independent of the source this packet came from. A new incoming PDU to the switch is assigned an ID from a pool of free IDs. Thus, packets coming from all the sources in the group share the identifiers. In this way, ID consumption is smaller than with Source ID. However, the solutions proposed up to now also use fixed size identifiers except in one case, Compound VC [3]. DIDA uses a 16-bit field, which is also overdimensioned, even more than in the Source ID case, because these IDs are shared by all the senders. GFC, on the other hand, uses small IDs (the 4 bits of the GFC field), which may be insufficient for bigger groups. In this case, more than one such GFC-connection should be used and the group management is then increased.

The exception comes from Compound VC (CVC), which allows flexible ID size negotiation at connection establishment so as to adapt to the ID consumption required

by each group. In this way, the overhead is minimized for two reasons: the PDU ID philosophy and the ID size negotiation. The next section briefly explains how CVC works.

3 Compound VC (CVC)

The main goal of the CVC mechanism [3] is to solve some of the problems of the mechanisms that have been commented above, namely buffer requirements, alteration of traffic characteristics, overhead, and flexibility.

Like in DIDA, there is a multiplexing ID per PDU that is carried by each cell belonging to that packet. CVC is based on local remapping of multiplexing IDs at the switches. Thus, no global ID assignment mechanism is needed, which would limit its potential deployment in wider environments than the local area.

VP switching techniques use the VPI to identify the group. This limits scalability due to VPI exhaustion, and also the exclusive use of the VPI by the operator. The solution CVC proposes is to treat a group of VCs as a whole when switching cells. The number of VCs in this group could be a power of 2 between one and 2^{16} . Thus, by providing this flexibility, scalability is increased. But scalability must not limit the potential application of a mechanism to multimedia, which is growing in importance. Multimedia traffic imposes stringent QoS constraints that make mechanisms that interleave cells more suitable due to traffic and group interaction constraints.

In CVC, the size of the PDU ID is negotiated at connection establishment depending on the group and traffic characteristics. The number of bits used as multiplexing ID may be determined by means of a mask. As far as the architecture is concerned, CVC requires a dynamic updating of the tables at the switch each time the first cell of a new PDU arrives and each time the last cell of a PDU arrives. Therefore, a new column in the switching table is required to consider a mask, but table size is not globally increased as there is just one entry for the whole group. The mask determines the portion of the VCI that will identify the group of VCs and the portion that will contain the multiplexing IDs (see fig. 1). CVC masks allow a lot of smaller groups with diverse sizes where there was just one group connection inside a VPI.




Fig. 1. VCI field in Compound VC

4 Determining the Number of Identifiers

The two preceding sections served to argue for the convenience of having PDU IDs of variable length to solve the cell-interleaving problem. CVC allows such philosophy. Therefore, one of the main problems to solve for CVC is the dimensioning of such variable-length ID.

The dimension of the PDU ID at a given switch through which a CVC connection passes should be chosen according to the traffic characteristics and the group size. The characteristics of the aggregated traffic observed by a switch are related with the traffic characteristics of each particular sender, that is why in this paper we try to determine the dependence of the required number of IDs as a function of the parameters of one source. This makes sense because we will study a homogeneous environment, i.e. all the senders transmit traffic with the same characteristics. The traffic parameters considered in this initial study are the PCR, the mean traffic and the length of the PDU.

With respect to group characteristics we see that passive members of the group, i.e. hosts that only act as receivers, do not have any effect on this traffic. As a consequence, the group size is characterized by the number of senders (N) in the group, when the purpose is to dimension the ID.

The goal of this paper is to find a methodology to study the dependence on traffic and group characteristics of the PDU losses produced due to running out of identifiers at a given switch. The procedure we will follow will start by calculating the histogram representing the frequency of the number of slots in which a given number of simultaneous PDUs is being transmitted through a given output port, i.e. the discrete probability density function of the number of simultaneous PDUs. Next, a graph for the PDU loss probability due to running out of identifiers is calculated, i.e. the graph will represent the probability of loss of an arriving PDU as a function of the number of IDs (nID) used. This graph is obtained from the previous one by adding from nID up to N-1 simultaneous PDUs. This sum is carried out for nID values from 0 to N. The probability that a PDU passes through a switch as a function of nID may also be obtained from the initial graph. In this case, we add the values ranging from 0 to nID-1 simultaneous PDUs.

As one of the goals of the paper is to obtain a value for nID as a function of traffic and group characteristics, our study will focus on the PDU loss probability graph. The idea is to find some dimensioning rules that relate such probability with the parameters being considered. Such rules would allow dimensioning the PDU ID at CVC connection establishment given a PDU loss probability acceptable for the user.

An analytical expression for finding the probability that a burst of the GFC mechanism is lost is presented in [11].

$$\sum_{i=h}^{n-1} \binom{n-1}{i} p^i (1-p)^{(n-1)-i} \quad (1)$$

This equation depends on the number of sources (n), the number of subchannels (h), the average of busy sources (m), and p is the probability that any given source is transmitting a burst ($p=m/n$). When applied to CVC, a burst is taken to be a PDU.

However, parameter m is too generic to be useful at connection establishment when trying to determine the number of required IDs. Other more easily obtained parameters, which could be directly related to the source, should be used for this purpose.

5 Simulation Environment

The simulated scenario consists of some bursty sources sending traffic to the same output port of a switch. They are characterized by their average cell rate (R), their peak cell rate (P), and the number of cells per burst (B). The sources are homogeneous, i.e. they are all modeled by means of the same statistics with the same parameters. A MMDP (Markov Modulated Deterministic Process) with one active state and one silence state is used to model each source, which is a particular case of ON-OFF source (figure 2). For this model, the sojourn time at both states follows a geometric distribution. In the OFF state, the source does not send any cells. In the ON state, it sends cells at its peak cell rate (P).




Fig. 2. ON-OFF source model

From the analysis of this Markov chain we may obtain that the probabilities of being at the ON and OFF states are:

$$P_{\text{ON}} = \frac{p}{p+q} \quad P_{\text{OFF}} = \frac{q}{p+q} \quad (2)$$

and the transition probabilities are related to the source parameters as follows [12]:

$$p = \frac{1}{B(b-1)} \quad q = \frac{1}{B} \quad (3)$$

where $b=P/R$ is the burstiness of the source.

Each source sends just one PDU at any given time. In our simulator, the output queue is modeled as a counter of the number of simultaneous PDUs for each slot. During all the simulations, the number of sources is varied depending on the mean traffic introduced to the switch so as to assure that the losses only occur due to running out of identifiers and not due to overload. This assumption is possible if we consider that there is enough buffer space at the switch to absorb the burstiness due to the aggregation of sources. The simulated time is 10^{10} μs .

The reference source is characterized by the following parameters. At peak cell rate (PCR), the source transmits one cell out of fifteen, i.e. for an STM-1 link the PCR is 10 Mbps. The mean traffic introduced by each source is 0.5 Mbps. The number of sources ranges from 100 to 300 to study the behavior of the mechanism for mid and high loads without reaching instability. PDUs are composed of an average of 5 cells and the sojourn time at ON state follows a geometric distribution. The OFF state also follows a geometric distribution with a mean of 1425 cells, which was chosen to obtain an average of 0.5Mbps per source. Any variation with respect to these parameters will be noted when presenting the results.

6 Results

The distribution of the number of simultaneous PDUs for a given reference scenario is presented in figure 3. Logarithmic scale has been chosen for both axis to provide a further detail for the range of values of interest, i.e. between 16 (4 bits) and 128 (7 bits) identifiers. We focus on these values because few bits in the ID provide low losses due to running out of ID (see 4). Statistically non-significant values have been removed from the figure.

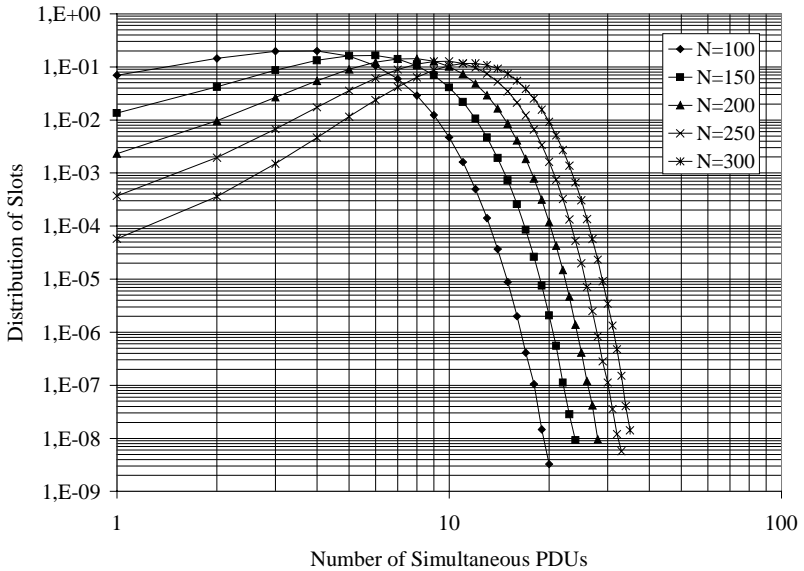


Fig. 3. Distribution of the number of simultaneous PDUs at the output port of a switch where merging occurs. The reference scenario is average=0.5Mbps per source, 5 cells per PDU, and PCR=150Mbps. Each curve corresponds to a different number of sources (N).

Figure 4 represents the probability that an arriving PDU does not find a free ID at the switch. Each curve corresponds to a different number of sources, and thus, different average loads at the switch, ranging from 50 Mbps for $N=100$ to 150Mbps for $N=300$. Only values that presented small 95% confidence intervals are represented.

A comparison with the analytical expression commented above (equation 1) is also presented. The goal of the comparison is to study the goodness of the fit between the results obtained through simulation and the theoretical expression for the simulated range of values. The PDU loss curves obtained by applying the analytical expression are labeled as N (*theor*) in figure 4. The parameters appearing in the expression were mapped to parameters in our simulation to obtain such curves. A burst is taken to be a PDU in our simulation, n corresponds to the number of sources (N in the figures), m

corresponds to the average number of simultaneous PDUs sent by the sources to the same output port, and h is the number of IDs. From the analysis of the expression it follows that each of the terms that are summed corresponds to the fraction of slots in which there are a given number of bursts being transmitted. In our case a burst corresponds to a PDU, therefore, figure 3 represents each of the terms being summed. And the sum from h up to $n-1$ is exactly how we obtain the graph of the PDU loss probability. It can be observed that the results of the simulation and those of the expression coincide for those values with small 95% confidence interval.

In a similar way, from the sum of the rest of the values, i.e. from 0 up to $h-1$, we obtain the probability that an arriving PDU to the switch is correctly multiplexed and forwarded through the output port because it was able to allocate a free multiplexing ID. These results give us an idea of the throughput of PDUs. These curves are not presented because those of figure 4 provide the complementary information. Just remark that, for this scenario, with 16 identifiers, PDU throughputs near 90% are obtained for the highest load case ($N=300$).




Fig. 4. PDU Loss Probability due to running out of identifiers. Reference scenario is: average per source = 0.5Mbps, mean PDU length = 5 cells, and PCR = 10 Mbps.

The main conclusion that may be drawn from these results is that with a few bits (between 4 and 6) and by using the PDU ID strategy, low PDU loss probabilities may be obtained even with a high number of sources. This scenario, which represents a possible scenario in future group communications, shows the advantages of PDU ID over Source ID multiplexing. Source ID requires a number of IDs assigned to the group equal to the number of sources. As a consequence, the overhead introduced by the multiplexing ID in these latter strategies is much higher than that of CVC. CVC allows the negotiation of the ID size to adapt to the traffic and group requirements.

Another characteristic that can be observed in figure 4 is the linearity (when using the logarithmic representation) of the PDU loss probability in the range of values of interest for the IDs. Therefore, the curve may be divided into three main linear regions. The first one starting at low ID values would be flat, which would tell that the number of IDs is not enough for such kind of traffic and group characteristics, as the PDU loss probability is 1. The second region, the one whose characterization is our main concern, would be a line going from the number of IDs where the curve starts to bend up to a number of IDs equal to $N-1$. The third region is characterized by a vertical line starting at $nID=N$, meaning that it is nonsense to use more IDs than sources in the group, because no losses occur due to running out of identifiers when $nID \geq N$.

Of course, such an approach is a rough approximation of the actual curve, but this characterization would allow a CVC switch to obtain, by means of a simple expression, the number of required IDs as a function of group and traffic characteristics and accepted PDU loss probability during connection establishment. The dependence of the parameters of these lines as a function of traffic characteristics is left for further study.

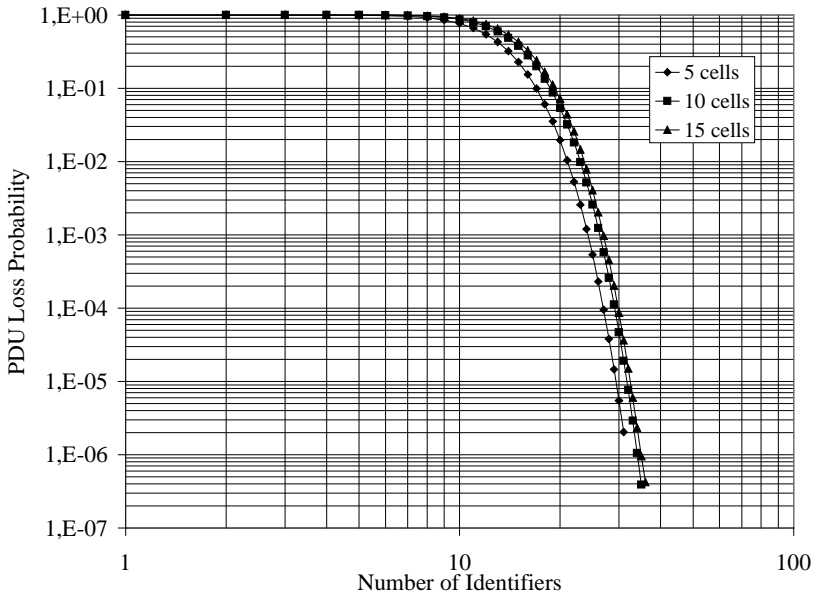


Fig. 5. PDU Loss Probability comparison varying the mean PDU length. Average per source=0.5 Mbps, PCR=10 Mbps, and $N=300$.

Other simulations have been carried out with different traffic and group parameters. For instance, PDU loss probability results were obtained for a number of sources ranging from 500 to 1500 with an average traffic per source of 0.1 Mbps. This average value is obtained by varying the mean sojourn time at OFF state. The same

observations stated above apply for this new scenario. The simulation results are also compared with the analytical ones. Both curves coincided for statistically significant values. However, simulations with high mean average per source (e.g. 5 Mbps) showed a difference between both curves.

Finally, the comparison between the results presented in figure 4 (reference scenario) and those obtained for the scenario with average per source=0.1 Mbps showed an almost imperceptible variation in the curves when the aggregated average load was the same.

The rest of the curves presented in this paper provide some results to study the dependence of the PDU loss probability curves on each of the considered parameters. For instance, figure 5 presents a comparison of the curves obtained for $N=300$ when varying the mean PDU length. The simulated values are 5, 10, and 15 cells per PDU, which correspond to reasonable mean values according to current Internet traffic, and in particular, to multimedia traffic. As it could be expected, the longer the PDU, the longer the IDs are occupied, and the more IDs are required. However, the variation between these curves is not very high. Thus, it may be observed that the curves show the same behavior (they all have the same shape). The only difference is a slight shift. Therefore, in the rough linear model we proposed to describe the behavior of these curves, it seems that the dependence of the equation of the line in the region of interest would be in the position of the point where the curve bends and not in the slope. Anyway, for reasonable mean values, such dependence would not be very strong.

Finally, figure 6 presents a comparison of the reference scenario described before with others in which one or two parameters are changed with respect to the reference one. These results were obtained for $N=250$ sources, which produce an aggregated traffic of 125Mbps, as the average traffic per source is 0.5Mbps.

We first focus on the curves that just vary the PCR while maintaining the rest of the reference parameters. They are labeled as $PCR=2Mbps$, $avg=0.5$ (which corresponds to $PCR=10$ Mbps), $PCR=30Mbps$, and $PCR=150Mbps$. In this case, the range of possible values is wider than in the PDU case. The effect of varying the PCR is to modify the burstiness of each source. This is so because the average per source remains unchanged while the cells are being transmitted in longer or shorter ON periods. $PCR=150$ Mbps corresponds to the most bursty sources, and $PCR=2$ Mbps corresponds to the smoothest simulated traffic.

For instance, to obtain a PLP of $1e-6$, the number of bits required for the PDU ID is respectively 7 (128 IDs), 5 (32 IDs), 4 (16 IDs), and 3 (8 IDs) for the 2, 10, 30, and 150 Mbps cases. This is due to the relationship between the PCR and the burstiness of the traffic introduced by the source. That is, if we maintain the same average traffic per source and we vary the PCR, the same number of cells per PDU is sent, but at a higher speed. Therefore, the PDU lasts less and as a consequence it is using an ID during less time, making it possible for other sources to get that ID.

These results show a strong dependence of the shift of the curves with respect to the PCR. However, the slope of the line in the region of interest seems to show only a slight dependence on the PCR. These dependencies will be studied in future work.

The curve labeled as $PCR=2Mbps/10$ cells per PDU serves to confirm that the PLP curve is more sensitive to PCR variations, i.e. to the burstiness of the traffic, than to

the length of the PDUs. Its values are more similar to those of the curved labeled as $PCR=2Mbps$ than to those obtained for the 15 cells per PDU case.




Fig. 6. PDU Loss Probability comparison for different parameters

After examining the results for this scenario, and without aiming to generalize, we could give some hints for the dimensioning of the ID size for a given CVC connection. They are mostly based on the bursty behavior of the sources, which is the characteristic responsible for the most important variations in the PLP values obtained for a given aggregated average load.

We calculated the burstiness of the sources for which their PLP curves are represented in figure 6. The burstiness (b) is 4 ($=10Mbps/0.5Mbps$) for the sources whose curves are labeled as $PCR=2Mbps$, and $PCR=2Mbps/10cellsPDU$, 20 for $avg=0.5$ and 15 cellsPDU , 60 for $PCR=30Mbps$, and 300 for $PCR=150Mbps$. For a PLP of $1e-6$, the number of bits required is respectively, 7, 5, 4, and 3. Therefore, in this scenario, we could use these ID sizes for values of burstiness around those calculated. Further work is required to derive rules that apply to scenarios with different average per source. They should also take into account the aggregated load, whose importance has been noted in comparisons presented above.

The diversity in scenarios and requirements for different groups also shows the advantages of having flexible ID size negotiation, such as the one offered by CVC. For instance, in multimedia group communications more losses could be accepted for video than for audio, and different number of sources would require different number of IDs.

7 Conclusions and Future Work

The results presented in this paper confirm that the efficiency in terms of multiplexing ID overhead could be highly reduced with PDU ID strategies when compared to Source ID strategies.

The convenience of flexible ID size negotiation, such the one offered by the Compound VC (CVC) mechanism, can be deduced from the diversity in requirements and group characteristics.

The PDU loss probability in the range of values of interest, i.e. from 16 IDs (4 bits) to 128 IDs (7 bits), shows a linear trend in logarithmic representation. Finding the expression of this line as a function of the traffic and group characteristics would allow easy ID negotiation during connection establishment. The derivation of this expression is left for future work.

The results have shown that the expression proposed in [11] and the PLP curve obtained through simulation coincide in most cases. However, the application of such an expression at connection establishment is limited because the user does not know the mean number of simultaneous PDUs the connection will have. Parameters that are available at connection establishment should be used instead.

The results also showed that the PLP is more sensitive to PCR variations than to PDU length. Some simple hints for the dimensioning of the ID size were given for a reference scenario. Once the aggregated load is fixed, the most important parameter is the burstiness.

The price paid when deploying CVC is the extra complexity in the switches. Implementation issues of CVC are left as future work to determine whether the benefit of introducing CVC is higher than its cost.

Other traffic parameters should be considered in future simulations so as to be able to draw more general conclusions on the behavior of PDU loss probability. Other kinds of traffic types may also be introduced to simulate more real environments. And scenarios of heterogeneous traffic sources may also be of interest in our future work.

References

- [1] Armitage, GJ 'IP multicasting over ATM Networks.' IEEE Journal on Selected Areas in Communications 15(3): 445-457, April 1997.
- [2] Baldi M, Bergamasco D, Gai S, and Malagrino D. 'A Comparison of ATM Stream Merging Techniques.' Proceedings of IFIP High Performance Networking (HPN'98): 212-227, Vienna, September 1998.
- [3] Manges-Bafalluy J and Domingo-Pascual J. 'Compound VC Mechanism for Native Multicast in ATM Networks.' Proceedings of the 2nd International Conference on ATM (ICATM'99): 115-124, Colmar (France), June 1999.
- [4] Grossglauser M and Ramakrishnan KK. 'SEAM: Scalable and Efficient ATM Multicast.' Proc. of IEEE Infocom'97: 867-875, Kobe (Japan), April 1997.
- [5] Rosen EC, Viswanathan A, and Callon R. 'Multiprotocol Label Switching Architecture.' IETF Draft, draft-ietf-mpls-arch-05.txt, April 1999.
- [6] Gauthier E, Le Boudec J-Y, and Oeschlin P. 'SMART: A Many-to-Many Multicast Protocol for ATM.' IEEE Journal on Selected Areas in Communications 15(3): 458-472, April 1997.

- [7] : Venkateswaran R, Raghavendra CS, Chen X, and Kumar VP. 'Support for Multiway Communications in ATM Networks.' ATM Forum/97-0316, April 1997.
- [8] Calvignac J, Droz P, Baso C, and Dykeman D. 'Dynamic Identifier Assignment (DIDA) for Merged ATM Connections.' ATM Forum/97-0504, July 1997.
- [9] Komandur S and Mossé D. 'SPAM: A Data Forwarding Model for Multipoint-to-Multipoint Connection Support in ATM Networks.' Proc. of the 6th International Conference on Computer Communications and Networks (IC3N). Las Vegas, September 1997.
- [10] Komandur S, Crowcroft J, and Mossé D. 'CRAM: Cell Re-labeling At Merge Points for ATM Multicast.' Proc. of IEEE International Conference on ATM (ICATM'98), Colmar (France), June 1998.
- [11] Turner J. 'Extending ATM Networks for Efficient Reliable Multicast.' Proc. of Workshop on Communication and Architectural Support for Network-Based Parallel Computing, Springer Verlag, February 1997.
- [12] Bolla R, Davoli F, and Marchese M. 'Evaluation of a Cell Loss Rate Computation Method in ATM Multiplexers with Multiple Bursty Sources and Different Traffic Classes.' Proceedings of IEEE Globecom, pp. 437-441, 1996.