A Comparison of ATM Stream Merging Techniques

M. Baldi, D. Bergamasco, S. Gai, D. Malagrino

Dipartimento di Automatica e Informatica — Politecnico di Torino

Corso Duca degli Abruzzi, 24 — I-10129 Torino (Italy)

Tel. +39-11-5647087 — Fax +39-11-5647099

Abstract

Multi-layer forwarding approaches (a.k.a. multi-layer switching or routing) which use ATM as transport technology, have proven not to scale enough unless route aggregation is performed. In ATM networks route aggregation implies stream merging: cells from different incoming streams are switched to the same outgoing link and labeled with the same stream identifier. This identifier could be either the whole VPI/VC pair (VC merging) or only the VPI (VP merging). Stream merging approaches are quite often referred to as VC merging approaches and in this paper we follow this naming convention. The standard way of carrying IP over ATM exploits the ATM Adaptation Layer 5 (AAL5) which does not provide native support for VC merging.

This paper provides an overview of the VC merging problem and presents a review of the most common solutions proposed so far. It also presents CLIMAX, a solution that could fit in different scenarios to solve the VC merging problem.

Keywords

Multi-layer routing, stream merging, VC merging, ATM, Multi Protocol Label Switching.

1 INTRODUCTION

Multi-layer routing techniques are quite often conceived for generic layer 2 and layer 3 protocols; however, their most natural application seems to be the combination of ATM (*Asynchronous Transfer Mode*) and the TCP/IP protocol suite in order to benefit from the performance of the former and the well-known

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-0-387-35388-3_42](https://doi.org/10.1007/978-0-387-35388-3_42)

properties of the latter. Multi-layer routing techniques have proven not to be sufficiently scaleable [SCALE] if the ATM network does not allow Virtual Connections (VCs) to be *merged*, i.e., cells from different incoming VCs to be switched to the same outgoing link using the same Virtual Path Identifier/Virtual Channel Identifier (VPI/VCI) pair. This capability, known as *VC merging*, allows multipoint-to-point VCs to be implemented. VC merging is crucial also to the implementation of scaleable group multicast over ATM, since it requires multipoint-to-point VCs too. This paper focuses on the application of VC merging to multi-layer routing, but most of the drawn considerations apply also to the exploitation of VC merging in ATM multicast.

In fact, VC merging is not the only possible solution to improve scalability in multi-layer routed (or switched) networks. In this paper many solutions that improve network scalability performing VP merging, instead of VC merging, are described. Thus, *Stream Merging* would be the best term to address the whole set of techniques; nevertheless, since the term VC merging has been traditionally used, we follow this naming convention in the paper.

VC merging cannot be performed by common ATM switches when higher layer packets are transmitted using the services provided by the *ATM Adaptation Layer 5* (AAL5) [ITU_AAL], as recommended by most of the proposals for carrying both data and multimedia traffic over ATM networks. In fact, AAL5 relies on the ATM layer delivering all the cells over a VC in the same order they were sent and without misinserted cells. Instead, when VCs are merged by a switch, cells belonging to different VCs get mixed together and are not distinguishable any more. Many different approaches for supporting VC merging have been proposed so far, but none of them has still proven to be the best in every situation. Each of them is particularly suitable for a specific network environment and for specific needs. In Section 2 the proposals appeared so far are briefly described and then compared. Conclusions are drawn in Section 3.

2. REVIEW

After providing a framework for classifying VC merging approaches, their advantages and drawbacks are described highlighting:

- hardware and software changes required in both core ATM switches and devices at the edge of the network;
- performance in terms of delay, jitter, throughput, and buffer requirements;
- specific problems.

A comparison of the different approaches is presented in Section 2.11.

2.1 Classification of VC merging approaches

Various VC merging approaches have been proposed in the recent past; some of them present many similarities and they can be broadly classified in three

categories according to the philosophy adopted to solve the problem, as shown in Figure 1:

- approaches based on avoiding cell interleaving;
- approaches based on VP switching;
- approaches based on AAL5 modification.

Approaches based on avoiding cell interleaving cause intermediate switches not to forward cells belonging to different packets simultaneously on the same output VC. All the cells belonging to the same packet are gathered and then forwarded all together. Approaches based on VP switching adopt the VCI to identify the packet to which a cell belongs. Finally, approaches based on AAL5 modification introduce an identifier in the cell payload and use it in order to discriminate among cells carrying different packets and traveling on the same VC. As far as the two last approaches are considered, they could be further subdivided into two categories, according to whether the identifier is associated to a packet or a sender.

Figure 1 - A classification of VC merging approaches.

2.2 MPLS Proposal

The IETF's *MultiProtocol Label Switching* (MPLS) working group proposes to avoid cell interleaving [MPLSARCH]. ATM switches are modified to implement a special queuing policy for incoming cells traveling on merged VCs. Each switch queues all the cells belonging to a packet until a cell with the *End Of Message* (EOM) bit¹ set (an EOM cell, for short) is received. This indicates that a whole packet has been completely received and buffered. Then, all the cells are transferred to the output queue for transmission. This mechanism avoids that cells belonging to different packets get interleaved on the output link.

¹ The EOM bit is set by a transmitting AAL5 entity to identify the last cell of a packet.

(a)

(b)

Figure 2- Cell buffering in the MPLS approach.

Figure 2 schematically shows the behavior of a switch implementing the MPLS approach. In Figure 2(a) no packet is being forwarded on the merged (outbound) VC, because none of the input buffered packets has been completely received. Cells belonging to incoming packets are being queued until their EOM cell is received. When the EOM cell of the gray packet is received, all the cells of the gray packet are transferred to the output queue at once, as shown in Figure 2(b). Note that even if some cells of the white and black packet have reached the switch before the gray ones, they will wait in input queues until their EOM cell is received, i.e., the whole white packet will be transmitted after the gray one.

AAL5 is not modified and ATM switches are not required to parse the cell payload. Even though connection endpoints do not need any change, this approach modifies the forwarding paradigm of switches and this, in turn, implies hardware modifications in ATM switches. Messages are not forwarded cell by cell and thus switches do not feature the latency properties characterizing ATM. However, since packets are not required to be completely reassembled, the MPLS approach demands less processing and introduces shorter latency than packet forwarding at intermediate switches. The extra buffer capacity and the per packet queuing needed in ATM switches could limit scalability.

2.3 Simple and Efficient ATM Multicast

Simple and Efficient ATM Multicast (SEAM) [SEAM] is very similar to the MPLS solution in buffering incoming cells until the EOM cell is received. Nevertheless, it aims at increasing throughput by forwarding cells immediately (before arrival of the EOM cell) when the output link is idle (*cut-through*). Figure 3 shows how SEAM works; the output queue being empty, the switch immediately forwards the cells of the first packet it began receiving, i.e., the cells of the white packet in Figure 3(a). This prevents cells belonging to other packets from being forwarded; as shown in Figure 3(b), even if the EOM cell of the black packet is received, it waits in the input queue until the EOM cell of the gray packet has been moved to the output queue.

Figure 3 - SEAM approach to VC merging

If the EOM cell of the packet being switched gets lost, cells belonging to other packets are blocked waiting for that EOM cell. A timer is used to overcome this problem; its duration is crucial and impacts significantly buffer requirement into switches. It should be determined on the basis of the bandwidth of the merged VC, the capacity of the links, and the load in the network.

SEAM shares most of the MPLS characteristics and it is not clear if performance is really improved. Short-cutting packets can imply longer latency for short packets (such as TCP acknowledgment messages), when a long packet is being

forwarded. Furthermore, it is hard to determine a suitable timer value since it depends on many parameters; if it is too short, good packets could happen to be discarded during congestion, and if it is too long, it could seriously affect buffer requirement, latency and jitter.

2.4 Cell Re-labeling at Merge-points

CRAM (*Cell Re-labeling at Merge-points*) is essentially an optimization of SEAM. It introduces a new type of *Resource Management* (RM) cell, which carries the multiplexing information. When cells are received from more than one input link for the same merged VC, switches suspend forwarding on the merged VC until a reasonable number of cells has been gathered. Then, the cells arrived from different sources are sent on the outgoing link without interleaving them, i.e., grouped according to their source². The trains of cells are preceded by an RM cell which contains a list of *Source IDentifiers* (SIDs) uniquely identifying the sender of each following group of cells. SIDs are ordered as the corresponding groups of cells on the merged VC.

Two mechanisms to guarantee SID's uniqueness are envisioned:

- Globally Unique SID Allocation: SIDs are uniquely assigned by a central server or through a distributed mechanism.
- Dynamic SID Allocation: network nodes dynamically remap the SIDs so that only local uniqueness is required.

CRAM is not compatible with current core and edge devices, even if it does not require any actual modification to the AAL. It requires some minimal changes into switches in order to cope with the new type of RM cell and implement its specific queuing mechanisms. Moreover, some mechanism must be used to cope with the assignment of unique SIDs. Finally, the support to Early Packet Discard (EPD) must be re-implemented because it requires to parse the RM cell payload.

2.5 Improved VP Switching and Merging

The improved (or extended) VP switching [VPMERGE] has been proposed by the ATM Forum and it can be categorized among VP switching based approaches. It consists in merging ATM Virtual Paths (VPs). Cells belonging to packets coming from different sources are discriminated through a VCI uniquely assigned by each source. Improved VP switching features all the characteristics of ATM cell switching, thus allowing resource reservation and cell scheduling policies to be kept unchanged, not introducing additional delay in the (VP) merging points.

² The boundary of the groups does not have to coincide with a packet; a group can either contain cells belonging to one or more packets.

Sources must be provided with a method for identifying a unique VCI value which is chosen at connection setup. At least two different categories of VCI assignments could be identified:

- *Server-based*: a central server in the network is responsible of the assignment of unique VCIs.
- *Signaling-based*: VCIs are negotiated by neighbor nodes.

VP merging presents the disadvantage of using a scarce resource, namely VP space, which limits the maximum number of merged connections on the same link. To overcome this problem, the improved VP switching approach proposes to enlarge the VPI field (18 bits) at the expense of a smaller VCI field (10 bits); the total cell label length is kept unchanged. This is not compatible with the standard operation of ATM switches and every cell must contain an indication of whether the switches should use the long or the short VPI field. The most significant bit of the VPI field is used to provide such an indication, thus halving the available VPI space. Moreover, implementation of improved VP switching requires ATM switches to be modified in order to cope with the new partitioning of the VPI/VCI field.

2.6 Dynamic Identifier Assignment

The *Dynamic Identifier Assignment* (DIDA) approach [DIDA] is similar to Improved VP Switching technique and it also comes from the ATM Forum. DIDA does not require packet reassembly at intermediate switches or usage of globally coordinated identifiers. DIDA assigns to each message a locally unique identifier which is inserted in the VCI field. Cells are routed according to their VPI, and the VCI is changed by each switch. The switch identifies any new VCI on incoming cells as the beginning of a new message and assigns a new locally unique VCI to the cell when it is transmitted on the outgoing port.

There are two differences between DIDA and Improved VP Switching:

- In the DIDA approach the VPI space, remains unmodified and is consequently smaller than in improved VP switching.
- VCI semantics and assignment are different. In Improved VP Switching the VCI identifies the source of the cell, while in DIDA it identifies a packet (i.e., packets generated by the same source can have different identifiers).

According to DIDA each identifier is assigned to a message only while it is traveling, thus requiring a small identifier space and no global uniqueness of VCIs. As well as Improved VP Switching, DIDA requires some modification to ATM switches which must modify the VCI in each cell, even though they do not use it for routing the cell. The number of merged connections across each port is limited to 4096, because the VPI field is not extended.

2.7 Double Identification Label Swapping

Similarly to DIDA, the *Double Identification Label Swapping* (DILS) approach from IETF [DILS] uses a double level of identification for each packet. The first level identifies the destination and the second the source. DILS envisages three options for the location of the identifiers:

1. The VPI identifies the destination and the VCI the source; the network performs VP switching.
2. One half of the VPI/VCI space is used to identify the source and the other half the destination; switches route cells based on the second half.
3. The VPI/VCI identifies the destination and the source identifier is placed in the cell payload; switches do not require any modification since routing is based on VPI/VCI.

DILS needs an auxiliary protocol to assign source identifiers; hardware changes are needed only with options 2 and 3, listed above. Software changes will be needed, when implementing DILS according to option 1. Options 2 and 3 show higher scalability than option 1, because of the larger labeling space available.

Performances are quite similar to those of Improved VP Switching and DIDA; cell switching is performed with neither extra delay introduced nor extra buffer capacity required.

2.8 The Sink Tree Paradigm

The *sink tree paradigm* [SINKTREE] is an innovative approach for ATM Local Area Networks (LANs) which is strongly based on VC Merging. Every switch in the LAN is the root of a multipoint-to-point VC (a *sink tree*) connecting it to all the other switches. A set of sink trees provides full connectivity among switches. Special cells called *connectionless cells* are transmitted over sink trees; they are differentiated according to a bit in the VPI field and are handled differently than regular ATM cells traveling over ordinary VCs. When a source host transmits connectionless cells carrying a packet to a destination host, the source switch places these cells on the sink tree associated with the switch of the destination host. The VPI/VCI fields of connectionless cells carry (1) the source and destination switch identifiers in order to identify the sink tree over which cells must travel, (2) the destination host identifier in order to allow the destination switch to deliver the cells to the proper host, and (3) the source host identifier. The latter enables the destination host to distinguish the cells coming from different sources and properly reassemble them even if many sources simultaneously transmit cells to the same destination host and they get interleaved while traversing the sink tree.

Storing all this information in the VPI/VCI fields limits the scalability of the approach. In fact, the length of the source and destination switch identifiers is 8 bits, and the length of source and destination host identifiers is 5 bits. This means that the largest LAN can span up to 256 switches, each having up to 32 hosts

directly connected, i.e., the maximum number of hosts allowed in a LAN is 8192. These numbers sound quite reasonable in a LAN environment, but prevent the scheme to be exploited in a wide area network.

The Sink Tree Paradigm requires switches to be modified to route cells based on the portion of the VPI/VCI field which identifies the destination switch (i.e., the sink tree on which the cell must travel). Moreover, a protocol for building sink trees and accordingly configuring the forwarding tables of switches is necessary. Edge devices need modifications too since the basic principles for VC creation and management have changed.

Even if the Sink Tree approach keeps the cell based forwarding paradigm typical of ATM, it is not suitable to the provision of service guarantees to applications in terms of controlled delay and jitter. In fact, switches cannot discriminate and properly handle the traffic of a specific application in order to provide it with the required quality of service. The finest possible granularity of traffic segregation into switches is the source-destination pair.

2.9 AAL5⁺

In [AAL5+] the VC merging problem is solved through a new AAL slightly differing from AAL5. AAL5⁺ overcomes the problems due to cell interleaving by marking all the cells belonging to the same packet with a *Message Identifier* (MID). Its value is assigned by sources on a per packet basis and it is randomly chosen in the range [0, 65535] with a uniform probability distribution. Destinations distinguish cells belonging to different packets thanks to the MID field and properly reassemble incoming packets, even if their cells got interleaved. Since the MID is chosen randomly, two or more messages may have the same MID at the same time. If their cells get interleaved the messages are lost because the destination cannot discriminate the cells belonging to the various messages. This phenomenon is called a *MID conflict* or a *MID collision*. MID conflicts are shown to be really rare and thus they are not explicitly handled. The upper layers reveal the incorrectness of packets affected by MID collision and discard them.

Figure 4 - AAL5⁺ SAR-PDU

The MID field, which is 16 bit long, is placed in the ATM cell payload by the *Segmentation And Reassembly* (SAR) sub-layer of the AAL, as shown in Figure 4. Since AAL5⁺ uses two octets out of the 48 of the ATM cell payload to carry the MID, its efficiency is lower than AAL5's one (i.e., $46/53=86.8\%$ versus $48/53=90.57\%$, respectively). AAL5⁺'s efficiency is anyway higher than AAL3/4's one (83%) which inserts a MID in each cell as well and could in principle

represent an alternative to support VC merging. Nevertheless, it is not used to this purpose because the ALL3/4 MID is intended for multiplexing on the same VC different kinds of traffic from the same source (not from different sources). Being AAL3/4 MID shorter (10 bits) than AAL5⁺ one, it is not suited to a random assignment, because the probability of MID collisions would be significantly higher.

Even though the efficiency loss introduced by AAL5⁺ with respect to AAL5 is not a major issue, the 46 byte payload of the SAR-PDU is not large enough to allow a TCP control message (e.g., an acknowledgment message) to be fully contained into a single cell³. This halves the efficiency in the transmission of TCP control messages (e.g., ACK segments) since two ATM cells must be transferred instead of one. Actually, the default encapsulation method of IP packets over ATM networks, requires an LLC/SNAP (*Logical Link Control/SubNetwork Attachment Point*) header (8 bytes) to be put in front of each IP packet in order to allow for multiplexing of different upper layer protocols [RFC1577] on the same VC. In this case TCP control messages do not fit anyway into the cell payload. Moreover, if IPv6 packets [RFC1883] are transmitted using AAL5, TCP control messages do not fit in a single ATM cell since the IPv6 header is 40 bytes by itself.

2.10 CLIMAX

The CLIMAX (*Cell-Interleaved Merged ATM conneXions*) approach [CLIMAX], analogously to AAL5⁺, proposes the exploitation of randomly chosen 16 bit Message Identifiers (MIDs) to allow cell interleaving at VC merging points. CLIMAX encompasses two possible implementations which basically differ in the way the MID is carried into cells.

AAL5⁺ Based CLIMAX inserts the MID in the first two bytes of the cell payload using the same format proposed in [AAL5+].

VP Switched CLIMAX inserts the MID in the VCI field of the cell header. This requires a software modification at the transmitting side of end systems in order to randomly choose a VCI value for the cells resulting from the segmentation of the same packet. VP Switched CLIMAX does not require any modification to the hardware of both network nodes and end systems (or edge devices). ATM switches perform VP switching on CLIMAX merged connections and VC switching on other VCs. This solution has a clear scalability limit due to the small dimension of the VPI field. If switches must support both traditional ATM VCs and CLIMAX merged VPs, a bit in the VPI must be used to differentiate between the two kind of VCs and the space of the merged VP identifiers is consequently reduced. VP Switched CLIMAX is completely transparent to the destination, which will distinguish cells belonging to different packets in the same way AAL5 usually

³ The TCP header (20 bytes), the IP header (20 bytes) and the AAL5 CS-PDU trailer (8 bytes) fit exactly in the ATM cell payload.

does. In fact, cells belonging to different packets arrive on different VCs, unless two different sources transmitting on the same VP have chosen the same MID to identify their packets (i.e., a MID collision takes place).

As well as AAL5⁺ (see Section 2.9), CLIMAX does not try to avoid MID collision since, in reasonable operating conditions, the MID collision probability is low and the consequent loss is acceptable [CLIMAX-TR], especially if EPD is implemented into intermediate switches as briefly discussed below.

The SAR sublayer in the receiver (of AAL5⁺ or AAL5, depending on the specific CLIMAX implementation) gathers payloads of cells with different MID values in different packet reassembly buffers. When an EOM cell is received, the SAR sublayer delivers the corresponding packet to the upper layer and releases the reassembly buffer associated with it. The memory required by the buffers concurrently used by the receiving entity to reassemble messages can limit the scalability of the approach.

Figure 5 - State-transition diagram of a receiving entity SAR sublayer in CLIMAX

If an EOM cell gets lost, the reassembly buffer associated with one of the packets being reassembled will not be released anymore. If switches implement EPD or a similar packet discarding mechanism, this phenomenon is rare. Since these techniques try to discard only entire packets instead of cells belonging to different packets, they limit the number of incomplete packets delivered to the destination, thus lowering the number of unreleased buffers. Anyway, packet discarding techniques like EPD have not yet reached a large diffusion.

The loss of EOM cells increases the probability of having MID collisions since each unreleased buffer is equivalent to keeping a MID in use until it generates a collision. CLIMAX strictly limits extra buffer allocation exploiting a *buffer release timer*. This reduces MID collision frequency and memory requirements in edge devices. Figure 5 shows the state-transition diagram of a SAR sublayer receiving

entity exploiting the buffer release timer. Each state represents the number of packets currently being reassembled or, equivalently, the number of reassembly buffers simultaneously allocated. A state change occurs when either a cell with a new MID or an EOM cell are received, or the timer associated to a particular MID expires. `CURR_MID_SET` is the set of the MIDs associated with packets being reassembled.

The choice of the duration of the buffer release time-out is critical, as shown by preliminary results of ongoing simulation work [CLIMAX-TR]. A too conservative (too long) timer might result in the need for a large amount of buffer memory and an increased MID collision frequency. This is the reason why the traditional AAL5 time out mechanism which is too loose, is considered not effective. At the opposite end, if the buffer release timer is set too short (i.e., shorter than the maximum cell delay variation experienced in the network), partially reassembled messages may be discarded due to even a single cell which has experienced the maximum delay.

Since buffering requirements at the destinations can affect the scalability of the approach, it is worth comparing CLIMAX with other approaches from this point of view. Destinations implementing CLIMAX allocate a buffer for each message being received on a multipoint-to-point VC. Assuming no loss of EOM cells, the maximum number of allocated buffers equals the number of sources transmitting concurrently on the merged VC. In a real scenario, EOM cells can get lost and reassembly buffer left open, but the exploitation of an effective buffer release timer can keep the number of buffers in use very close to the lossless case. Notice that when multipoint-to-point communications or multi-layer forwarding are performed without exploiting VC merging (i.e., group communications are implemented through a mesh of point-to-multipoint VCs and point-to-point VCs are used with multi-layer forwarding schemes) the total buffering capacity required in each receiver equals the number of sources, i.e., the upper bound for CLIMAX. This is because in each receiving node a different AAL5 entity must be instantiated to terminate each VC, with the consequent allocation of a reassembly buffer. Alternatively, when VC merging is performed by avoiding cell interleaving in merging points (e.g., like in the MPLS approach described in Section 2.2), the buffer space used by CLIMAX receiving entities is needed into switches.

2.11 Comparison

In Table 1 a comparison among the three classes of approaches discussed in Section 2 is outlined. The comparison is based on issues relevant to the production and deployment of these schemes (e.g., need for hardware modification).

Hardware changes are needed in either edge or core devices, but most of the approaches do not require both of them. The approaches based on AAL5 modifications require hardware changes in edge devices, while the others usually impact on the core of networks. Notice that usually in wide area networks the ratio

between core and edge devices is 1:20, making it simpler and preferable to change the former. AAL5 compatibility is obviously not granted by approaches based on AAL5 modifications, while it is generally preserved in the others.

Table 1 - Comparison of the VC merging approaches

<i>Category</i>	<i>No cell-interleave</i>	<i>VP switching</i>	<i>Modified AAL5</i>
Examples	MPLS, SEAM, CRAM	Impr.VPswitching, DIDA, DILS (opt.1 & 2), Sink Tree, CLIMAX	DILS (opt.3), AAL5+
Hardware changes in edge devices	No	No	Yes
Hardware changes in the switches	Yes	No (if VPI/VCI partitioning is not changed)	No (if EPD is not needed)
AAL5 compatibility	Yes	Yes	No
Label space for destination	VPI/VCI (28 bits)	VPI (12 bits)	VPI/VCI (28 bits)
EPD Compatibility	Yes (changes needed)	Yes (no changes)	Yes (changes needed)
Buffering required	High	Low	Low
Latency	High	Low	Low
Switching	Pseudo-packet-switching	VP level cell-switching	Pure cell-switching (not for CRAM)
QoS capabilities	Low	Medium (VP based)	High (connection based), lower for CRAM

The label space per destination is an indicator of the scalability of the approach because, if limited, it can reduce the maximum number of edge devices that could be connected to the network.

EPD compatibility indicates if any changes are needed in order to support packet discarding techniques, like EPD. Approaches based on avoiding cell interleave can support EPD, but they will need changes in ATM switches hardware. This is not an added limitation, since hardware changes are needed anyway in this case. Approaches based on usage of a packet identifier - either carried in the VCI field (VP switching) or in cells (AAL5 modifications) - could easily interoperate with current implementations of EPD, but it would be more effective to base packet discarding on the identifier used to support VC merging.

Buffering, latency and switching method are considered significant performance indicators. The first one impacts on cost and complexity of switches while the last two affect the suitability of the approach for controlling delay and jitter. Approaches based on VP switching and those requiring modification of AAL5 present better performances, while approaches based on the avoidance of cell-interleaving could have some limitations, especially when handling traffic other than best-effort.

The Quality of Service (QoS) capability row expresses the suitability of the category of VC merging approach for guaranteeing QoS. Of course, the more cell switching and its properties are preserved, the higher the suitability for providing QoS guarantees.

3. CONCLUSIONS

A network using the standard IP over ATM protocol stack in intermediate and end systems does not allow Virtual Connections (VCs) to be merged. This feature is essential to allow for transmission of packets on multipoint-to-point VCs to either solve scalability problems in multi-layer forwarding or group multicast communications. This paper presents a survey of the most common approaches proposed so far to solve the ATM *VC merging* problem. The approaches are grouped into three categories which are compared according to issues relevant to the production and deployment of the required equipment.

Currently, the mainstream approach to solve the VC merging problem in the context of the MultiProtocol Label Switching (MPLS) IETF's working group is based on avoiding cell interleaving in merging points. (Modified) ATM switches buffer all the cells of a packet before starting to forward them; this represents a step away from cell switching towards packet switching.

We consider CLIMAX a very promising approach due to its properties. It is easy to implement and operate, and since it implements traditional cell switching, it is suitable to the provision of Quality of Service (QoS) guarantees. Two CLIMAX implementations are possible: one based on usage of VP switching, the other based on a modification of AAL5 named AAL5⁺. The latter has higher scalability, but

requires hardware changes in edge devices and thus, due to the large number of such devices, it is not an attractive solution for immediate deployment.

We envision a migration towards the massive adoption of VC merging in ATM networks, where the most suitable short term solutions are VP Switched CLIMAX implementation (in small networks) and the MPLS approach (in large networks with a high ratio between the number of edge and core devices and with no QoS requirements). For the long term, the solution which will best combine scalability and cell switching performance is AAL5⁺ based CLIMAX.

4. REFERENCES

- [AAL5+] F.Hoymany, D.Mossé, "More Powerful ATM Connectionless Support Using AAL5", Proc. Of IASTED Networks '96, Orlando, Florida, January 1996.
- [CLIMAX] D.Bergamasco, S.Gai, D.Malagrino, "CLIMAX, Cell Interleaved Merged ATM connections", to appear in Proc of 6th Int'l. Conference on Telecommunication Systems, March 1998, Nashville, TN
- [CLIMAX-TR] M.Baldi, D.Bergamasco, D.Malagrino, "ATM VC Merging Techniques: Comparison and Simulative Study", Technical Report, TR-DAI-NET-980320, Politecnico di Torino, Torino, March 1998
- [DIDA] J.Calvignac, P.Droz, C.Basso, D.Dykeman, "Dynamic Identifier Assignment (DIDA) for Merged ATM Connections", ATM Forum Contribution 97-0504, July 1997.
- [DILS] G.Goren, I.Iliadis, P.Droz, "Double Identification Label Swapping (DILS) for Merged ATM Connections", Internet-Draft <draft-droz-dils-arch-00.txt>, July 1997.
- [ITU_AAL] International Telecommunications Union - Telecommunications Sector, "ATM Adaptation Layer", Recommendation I.363, March 1993.
- [MPLSARCH] E.Rosen et al. - "A proposed Architecture for MPLS", Internet-Draft, <draft-ietf-mpls-arch-00.txt>, August 1997.
- [VPMERGE] R.Venkateswaran, C.S.Raghavendra, X.Chen, V.P.Kumar, "Support for Multiway Communications in ATM Networks", ATM Forum Contribution 97-0316, May 1997.
- [RFC1577] M. Laubach, "Classical IP and ARP over ATM", RFC 1577, January 1994.
- [RFC1883] S. Deering, R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, December 1995.
- [SEAM] M.Grossglauser, K.K.Ramkrishnan, "SEAM: Scalable and Efficient ATM Multicast", INFOCOM '97, Kobe (Japan), April 1997.
- [SINKTREE] R.Cohen, B.V.Patel, F.Schaffa, M.Willebeek-LeMair, "The Sink Tree Paradigm: Connectionless Traffic Support on ATM LANs", IEEE Transactions on Networking, V. 4 N. 3, June 1996.

[SCALE] Z.Wang, G.Armitage, "Scalability Issues in Label Switching over ATM", Internet-Draft, <draft-wang-mpls-scaling-atm-00.txt>, July 1997.

5 BIOGRAPHY

Born in Cuneo, Italy, on November 9th, 1968, **Mario Baldi** is Assistant Professor of Computer Engineering at Computer and Control Engineering Department of Politecnico di Torino. From April to May 1997 and from July to September 1997, as part of his Ph.D. program, he visited the Computer Science Department at Columbia University in New York, New York, so as to work with Dr. Y. Ofek and Dr. B. Yener on synchronous packet switching techniques for real-time services.

Davide Bergamasco was born in Turin (Italy) in 1968. He received his Master Degree in Computer Engineering from Politecnico di Torino, Turin (Italy), in 1995. Currently he is a Ph.D student with the Computer Science Department of the same University. In November 1997 he joined Cisco Systems Inc., San Jose, CA (USA), where he will be working as a visiting researcher until March 1999. His current research interests include, but are not limited to, QoS and performance evaluation of packet switched networks based on computer simulation techniques.

Silvano Gai was born in Asti, Italy, on January 20th, 1957. He received the Dr. Eng. degree in electronic engineering from Politecnico di Torino. From 1990 he has been a professor of Computer Network at Politecnico di Torino. His current research topics include: high speed Ethernet, virtual LANs, multilayer switching, IPv4, IP over ATM, IPv6 and Multicast applications. He wrote a book titled "Computer Network: from cabling to internetworking" that is currently used as the classical course book in the Italian Universities. In 1997 he wrote for McGraw-Hill Italia the book "Guide to IPv6". From February 1st, 1997 to October 30th, 1997 he was in sabbatical at Cisco System, San Jose, CA. In 1998 he wrote the book "Internetworking IPv6 with Cisco Routers", published by McGraw-Hill USA.

Dante Malagrino (born in 1971) is a Ph.D. student at Politecnico di Torino. He received his Master Degree in 1997, defending a Thesis on "Comparative Analysis of Internetworking Solutions". Since March 1995, he has been working as International Project Manager for the IRISI, a European Project on the Information Society, and he participated to many other European projects. He will carry out a Ph.D. Thesis on "Gigabit Routing" and his main interests include network simulation, high speed network architecture, routing and switching. He published some papers on scientific Italian magazines and on international conferences proceedings, and he participated to the MPLS Working Group of the IETF.