

Virtual LAN-WAN for Workgroup Multimedia Application¹

A. Azcorra, J.I. Moreno
Dpto Ingeniería de Sistemas Telemáticos
Technical University of Madrid
28040 MADRID (SPAIN)
E-mail: aazcorra@dit.upm.es, jmoreno@dit.upm.es

Abstract

Workgroup multimedia applications such as videoconference, teleteaching or CSCW demand multipoint service from Telecommunication Networks. This is not just a broadcast service but it is a service where N sites transmit traffic to N sites. This service requires that the network support a large number of workgroups with few (5-20) participants together with low delay, low delay jitter, high throughput and a connection-less service in a wide area. The service of required by each of these workgroups is what we call Virtual LAN-WAN (VLW).

At present, most of these multimedia applications have been developed over the IP protocol. Asynchronous Transfer Mode (ATM) will be the technology of the future Broadband Integrated Services Digital Network (B-ISDN), so ATM Networks must provide a framework for the development of multimedia applications based on IP protocols. In fact, the integration of IP over ATM Networks is as subject of intense work in standardization bodies. Within this effort, the ATM Forum and the Internet Engineering Task Force (IETF) have achieved some results in the provision of LAN service based on IP over ATM.

This paper shows the problems of supporting a VLW service based on IP protocol over ATM Networks. We summarize the ATM Forum model based on LAN Emulation Specification (LANE) and the work in progress of the IETF on this subject. Pros and cons of every solution are described, and a model considered to provide cost-effective performance is proposed. Finally, we present a case study that has been tested in a European scale. This case study is a distributed conference based on the ATM paneuropean pilot interconnecting ten sites in seven countries.

Keywords:

IP, ATM, multicast, groups, LANE, MARS, traffic, costs, management.

1. Introduction

Current multimedia applications require multicast service from communication networks. In the wide area networks, this service should be provided taking into account the relative high cost of the bandwidth consumed by multimedia traffic. The network service should allow the dynamic creation and destruction of groups, together with the entry or exiting of members during the group life.

Multicast service is provided today based on IP multicast and the shared media nature of LANs. In the wide area the solutions are based on multicast routers, which can be mixed with conventional routers using tunneling, being the MBONE the largest network supporting world-wide IP multicast service. The shared resources, and the layer 3 processing, impose some drawbacks for the usage of this approach in a widespread production (non-research) environment.

Solutions based on switched networks are very frequently found today in the local area [1]. Commercial products offer a various range of performance figures, but all of them suffer from proprietarieness and poor treatment of multicast. Moreover, most of them cannot scale beyond the local area.

¹ This work has been partly supported by the Comission of the EU under projects BRAIN, IBER and ISABEL.

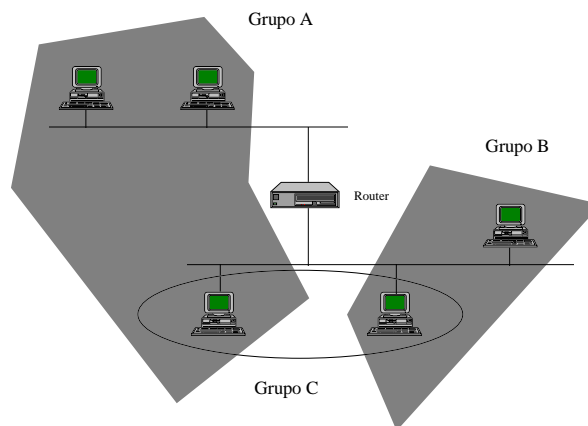
Within the wide area, ATM is consolidating as the leader technology for broadband networks. Statistical traffic multiplexing, transmission speed scalability, compatibility with LAN technologies and suitable multicast support are some of the features considered more relevant. This clearly suggests that solutions for the VLW service should be based on ATM.

In the ATM framework there are two approaches close to VLW being developed by relevant standard-bodies: LAN Emulation over ATM [2], by ATM Forum, and Virtual IP Networks over ATM, by IETF [3]. However, both approaches are mainly focused on the local area, and do not take into account scalability problems and resource optimization.

Next section presents the classical solution based on IP multicast routers and shared media LANs. Section 3 describes the advantages provided by using ATM, and the current state of the approaches from ATM Forum and IETF. In section 4 we propose a design better adapted to the requirements of workgroup multimedia applications over the WAN area. A case study thoroughly tested is presented in section 5. Finally, some conclusions are drawn in section 6.

2. Classical IP multicast routing

Multicast IP was developed based on shared LAN technologies (e.g. Ethernet, Token Ring, FDDI) that are inherently suitable for data broadcasting. Shared media allows that the resource usage, delay and delay jitter of a multicast frame be the same as the ones of a unicast frame. The interconnection of subnets is performed using multicast routers that, first, are capable of routing to all the necessary destinations, and, second, make the necessary copies along the distribution tree.



Addressing is based on a global, unassigned, flat numbering plan. Each class D address represents a global end-system group. Admission to the group is always granted. Collisions in group creation (i.e. two different workgroups attempting to use the same address) result in the inclusion of all end-systems in a single group. A given end-system may belong to as many groups as it wishes. An end-system sending data to a group does not know the identity of group members (just requires to know the group address), and does not have to belong to the group.

Group management is performed by using the IGMP protocol for member acquisition by the routers, and the explicit mechanisms embedded in the routing protocol for membership distribution between the routers:

- A group is created when an end-system sends an IGMP [4] message to subscribe to a previously free class D address. The group is deleted when the last member abandons the group, either explicitly or detected by time-out. Successive end-systems sending IGMP join/abandon messages are added/removed to the group by the routers.
- Within shared-media subnetworks there is an optimization consisting in that an end-system wanting to join a group, or that is already in a group, will not send an IGMP join message if another end-system in his subnetwork has already done so.

- Routers use specific multicast routing protocols (e.g. [5], [6], [7]) to maintain routing trees that reach all the members of a group, in order to send the datagrams only to those links/subnetworks that lead to members in zero or more routing hops.

The usage of classical IP multicast for workgroup multimedia applications presents several disadvantages:

- As there is no traffic separation mechanisms, QOS cannot be guaranteed. Some solutions based on modifying the network protocols [8], or designing independent virtual IP networks over ATM could be used to reduce this problem, but they are not considered as fully satisfactory. The problem of switched trunk capacity dimensioning over the WAN is therefore not applicable.
- SAR layer reassembly at every node and datagram copying at the IP layer is less cost-effective, for technological reasons, than end-to-end ATM cell transmission. Moreover, end-to-end delay (and delay jitter) is increased because of forwarding data datagram by datagram instead of cell by cell. Batch application such as multimedia broadcasting are not affected so much, but interactive groupware is.

3. Provision of Multicast Service in ATM Networks

Asynchronous Transfer Mode [9] has been selected by ITU-T and the majority of the industry to be the core network of the future Broadband-ISDN. Using ATM as the means to transport IP datagrams end-to-end across the WAN offers several advantages:

- Transmission is connection oriented with assigned resources. In ATM networks the connections are guaranteed the required QOS: required bandwidth and delay, and a small delay jitter.
- Point to multipoint capabilities. ATM switches permit to establish unidirectional point to multipoint connections. Cell copies are made in hardware inside the switch with minimum latency and implementation cost.
- Transmission is based on fixed size cells. Because of this, hardware based switching is used allowing a very low cell processing time. The reservation of resources already mentioned also causes low queuing time.
- The small cell size allows to minimize network latency caused by cell (re)transmission time at every node, and also to minimize delay jitter caused by cell multiplexing.

It is also true that ATM currently lacks some fundamental issues such as standardized routing, refined group addressing, and group management. However, there is much work being put on those aspects and most of them will be available in the short future. Two main standardization efforts are being carried out for the provision of a group multicast service over ATM. One of them by the ATM Forum through LAN emulation Service (LANE) and the other one by IETF through a collection of (draft) RFCs centered on the Multicast Service for Logical IP Networks (mainly the MARS protocol). In the remainder of this section we describe both approaches focusing on the multicast support for workgroup multimedia applications.

3.1 LAN Emulation Service (LANE)

LAN emulation service, defined by ATM Forum [2], was developed mainly to provide a fully compatible IEEE MAC service over the ATM protocol stack. These effort is aimed to support without any changes the vast amount of software and applications currently available on LANs, and simplify the configuration of an ATM local network to mimic the ease of use of LANs. advantages are service, called "LAN emulation" that emulates services of existing LANs across an ATM network. This service has been designed to allow interoperability between software applications residing on ATM-attached end systems and on traditional LAN end systems by means of specific components for current bridges and routers.

Each LANE consists of a set of so-called Lan Emulation Clients (LEC) and one LAN Emulation Service (LE Service). This LE Service consists of three elements with different functionalities: a LAN emulation configuration server (LECS), a LAN emulation server (LES) and a broadcast and unknown server (BUS). A detailed description of these elements can be found in [11] and [12]. Figure 1 shows the interactions between the LANE components.

In order to provide MAC multicasting service LANE specifies the usage of a specific server, the Broadcast and Unknown Server. LECs wishing to send a multicast MAC frame send it to the BUS using a

unicast VC. The BUS broadcasts the frame to **all** LECs in the LANE. This broadcasting is performed either using a point-to-multipoint VC (Multicast Forward VCC) if available, or a bidirectional use of the n VCs used by the n LECs to send data to the BUS. An effect of this approach is that a Client receives its own multicast traffic (in addition to the unknown), being thus necessary to explicitly filter it.

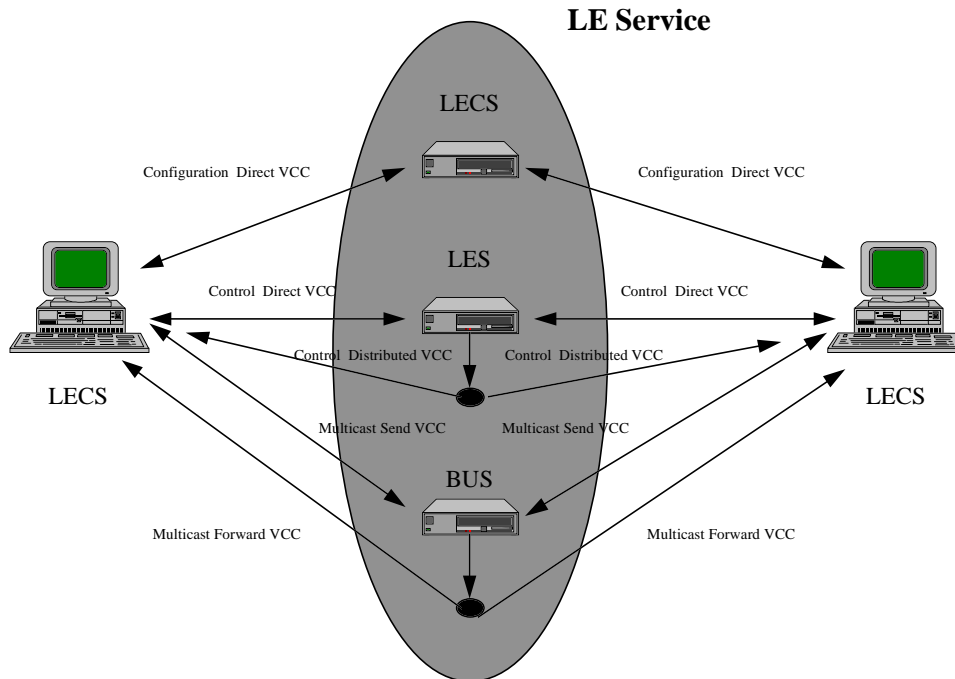


Figure 1: Components of the ATM Forum LANE

Basing the multicasting server on the BUS is a simple and clean solution, although it presents some potential problems:

- Scalability:
 - All multicast traffic is handled by the BUS. In a multimedia environment with large quantities of multicast traffic the BUS would become a bottleneck.
 - Every LEC receives **all** multicast traffic, being required to filter it at the MAC software-implemented layer. This would impose an unnecessary processing overhead at the clients to discard incoming frames corresponding to non-subscribed groups.
- Efficiency:
 - As frames are sent to all clients, indiscriminate transmission would result in unnecessary bandwidth consumption. While in a local environment with cheap bandwidth this is acceptable, in a WAN environment this would be a serious drawback.
 - As there is only a single BUS per LANE, clients require a double-hop to receive frames, resulting in an additional undesirable transit delay.

The ATM Forum document suggests that the BUS could be refined to send frames only to the destinations that are members of the group, and also that the BUS function could be distributed among several physical servers. However, no technical solutions are drafted in order to do it.

The reader should notice that the disadvantages described above are relevant in a WAN environment. This makes the LANE perfectly suitable for local area ATM networks, but not so much for the wide area.

3.2 Multicast Service for Logical IP Networks

As an alternative to the proposal of the ATM Forum, the IETF has proposed a design to build logical IP subnetworks (LIS) over the ATM networks ([13], [14]). Within this proposals, there is a specific document to provide multicasting service in logical IP networks [3].

A LIS is a set of systems (end-systems and/or routers) connected to an ATM network that share a common IP subnet prefix address. Communication between end-systems attached to different LIS must be made through a router(s) that connect both LIS. IP packets are directly encapsulated over ATM, and mapping between local IP addresses and ATM addresses is performed using a generalization of the ARP/RARP protocols.

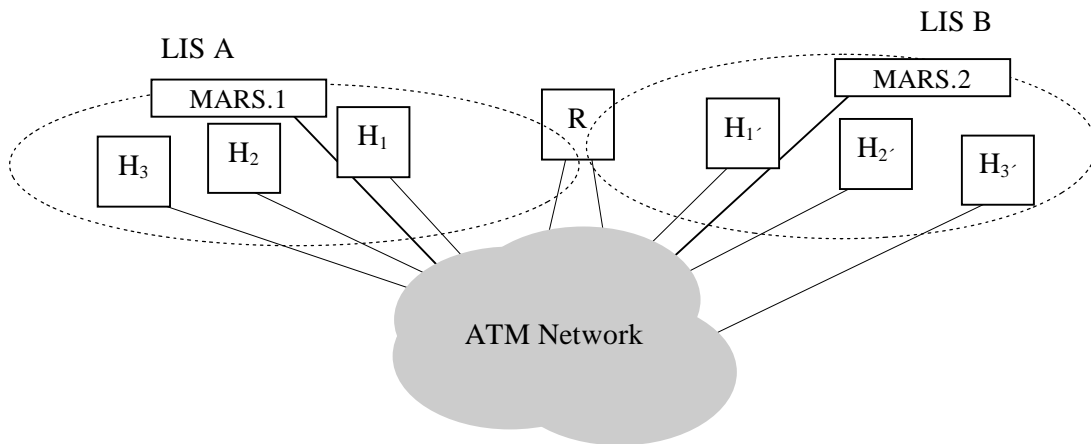


Figure 2: LIS and MARS

The specific mechanisms for providing multicast service consist in using one multicast address resolution server (MARS) per LIS. Clients are required to explicitly notify the server the entry/exit to a group, identified by a Class D address. For each active group the server maintains a list of the ATM addresses corresponding to all the systems that belong to the group, within its LIS. Upon request, the MARS broadcasts to all systems participating in any group within its LIS the list of members of the requested group.

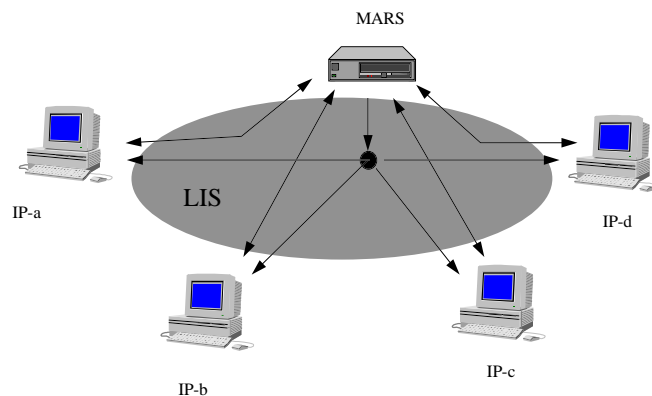
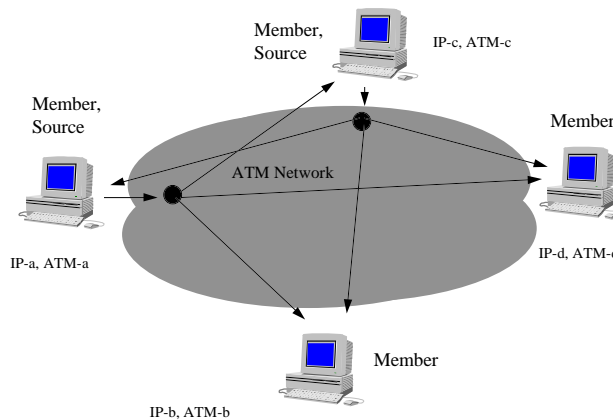


Figure 3: Communication between MARS and LIS clients

A client wishing to send information to a group uses the list of group members within the LIS received from the MARS to establish a point-to-multipoint VC connecting all destinations. In this way, datagrams are sent end-to-end using the ATM network, within the LIS. Multicast communication between LIS must be handled by multicast IP routers, that should register to the MARS in the corresponding LIS.



The IETF design does not present the scalability problems of the LANE. However, for workgroup multimedia applications there are some unaddressed issues. Notice that, for this approach to work properly, all workgroup members must be in the same LIS. Otherwise, transmission through routers would arise the problems described for Classical IP Multicast in Section 2. This implies some LIS management constraints such as the following.

As workgroups are dynamically created/destroyed in short time-frames (e.g. a four hour CSCW session), mechanisms are required for automatic LIS creation/destruction. Moreover, a given system can simultaneously participate in more than one workgroup, being then required mechanisms to allow that a system with a single ATM interface belongs simultaneously to several LIS. Finally, the IETF does not describe how resources should be allocated within the ATM network in order to optimize WAN transmission.

4. Proposed Network Design

4.1 Model of Workgroup Application Requirements

We consider that the computing trend is towards a highly distributed and cooperative environment. The bandwidth cost forecast suggests that in the very near future it will be fully cost-effective to form task-groups formed by persons physically located in distant places. Moreover, it will be frequent that a given person will be assigned to several task-groups, in order to contribute his highly specialized skills. The technical requirements for this organizational scenario suggest that the workgroup must be the central design point, in opposition to the current departmental (based mainly on a static physical proximity) approach to network structuring.

A workgroup approach implies that network resources should be so as to allow an efficient information sharing between task-group members. A group activity usually consists on batch and interactive periods. During batch periods members exchange/access to relevant information available to them. During interactive periods, the members cooperate by using full fledged multimedia applications. The duration of batch periods can be estimated in several weeks or months. The duration of interactive periods can be estimated in the range ten of minutes to ten hours. Technical solutions for batch periods are more or less well defined, but it is not the case for interactive periods, and that is the subject of study of this paper. The number of workgroup participants can be estimated in the range of 5 to 75 for batch work and 2 to 20 for interactive work.




Figure 4: Correlation of traffic sources

The large amount of multimedia traffic generated during the interactive periods suggests that a precise characterization would result in resource saving by tuning the network design. Experience in CSCW applications performed in several projects carried out in our organization show that traffic sources in an interactive session are strongly correlated. The sources have different traffic generation states which are mutually related. The typical case is a scenario in which sources have two traffic states, high and low, and only one of the sources is in high state at any given time (see Figure 4). In audioconference sessions, one source is speaking while the remaining ones are listening. In videoconference sessions the situation is very similar, having one active source speaking and being seen at top video quality, while the remaining sources are listening and being seen at low video quality.

The conclusion is that the peak traffic of the aggregation of the n sources is not n times the peak traffic of each source. The peak aggregated traffic is the high-state traffic plus $(n-1)$ times the low-state traffic. This traffic model has a very important impact in network design because the incoming bandwidth allocated to any member is much lower than n times the peak traffic of each source (outgoing bandwidth of the member).

In a LAN environment where abundant resources exist, this characterization is not very relevant. However, in the WAN environment, even in an scenario where line costs are reduced one order of magnitude, the economical saving obtained from reducing by n the required bandwidth (in one direction) should be taken into account.

Another characteristics of the workgroup traffic are variable, bounded, bit rate (in some cases it could be statistically modeled), and multicasting to all workgroup members.

4.2 Applicability of Existing Network Models

The current solutions to provide multicast support in a WAN environment based on Classical multicast IP (see section 2) does not satisfy the above requirements. In particular, resource sharing makes not possible to guarantee that the application will work during an interactive period because of delay jitter and bandwidth shortage. Another drawback is that software processing makes it inefficient for transport interactive multimedia applications.

Extending the LAN solutions based on LANE or MARS (see section 3) to the WAN environment is a better approach but is not straightforward.

In the case of ATM Forum approach, the adaptation to the workgroup WAN environment would require to build a LANE for each workgroup. In this way the scalability problems would be reduced because the number of clients in a LANE would be small. However, the efficiency problems remain: a double hop to the BUS is required, and much more important, multicast frames are sent to *all* clients. Even if this is acceptable for batch periods, assuming that most of the traffic is batch and unicast, it is not so for interactive periods. In interactive periods the delay increase may be unacceptable when network transit delay is high (e.g. satellite connections). Moreover, as not all members participate in every interactive period, the WAN bandwidth waste caused by indiscriminate broadcast would be unacceptable. A solution for this last problem could be creating a specific LANE for each interactive period, having as clients the participating members. Finally, it should be explicated the mechanisms that in the WAN case, perform the ATM resource allocation corresponding to the n multicast-send-VCCs taking into account the correlation of workgroup multimedia sources, as explained below.

In the case of using the IETF MARS approach, the adaptation to the workgroup WAN environment would require to build a LIS for each workgroup. The IETF approach does not expose the problems found in LANE. In this case there is no double hop and there is no indiscriminate broadcast, as long as interactive periods use a Class D address to which only the participating members have joined. However, the IETF should clarify the mechanisms that allow a system belong to different LIS, as one person can belong to different task-groups. Using IP subinterfaces over a single ATM interface would solve only part of the problem, as LIS is based on using one ATM layer address for each IP address. Multi-ATM-address cards should then be required, with the appropriate mapping software. Finally, it should be explicated the mechanisms that in the WAN case, perform the ATM resource allocation corresponding to a Class D address multicast circuit taking into account the correlation of workgroup multimedia sources, as explained below.

4.3 Network Design

In a WAN scenario, the usage of multicast servers is considered inefficient because of double-hop transmission, bottlenecks with the single advantage of solving the downstream resource allocation, but maintaining the upstream one. For this reason, the IETF approach based on per-group and source point-to-multipoint circuits is considered better suited, although more complex to manage. The mechanisms required to allow a system belong to several workgroups should be studied, and the resource allocation added.

The complete design proposed in this paper, describing the resource allocation optimization adapted to workgroup multimedia traffic, is described below. The description center on WAN design, and end-systems will be attached to the WAN. Obviously, the most frequent case will be that in which broadband LANs will be attached to the WAN access points, but this fact does not affect the WAN network design.

WAN transmission is performed using ATM encapsulation end-to-end: there are no IP intermediate systems. Transmission is performed directly from the source to the members: there is no multicast central server. ATM multicast, point-to-multipoint VCs, must be available in the switches. The resources are structured as a set of hop-by-hop VPs that form a shortest-path, depending on physical topology and resources, acyclic graph spanning all the end-systems. The bandwidth reserved for each VP is symmetric (save for source-only or member-only access-lines) and the characteristics correspond to the aggregated traffic of the interactive session. Traffic is sent using point-to-multipoint VCs, from each source, that are multiplexed together in the available VPs. Figure 5 illustrates this design.

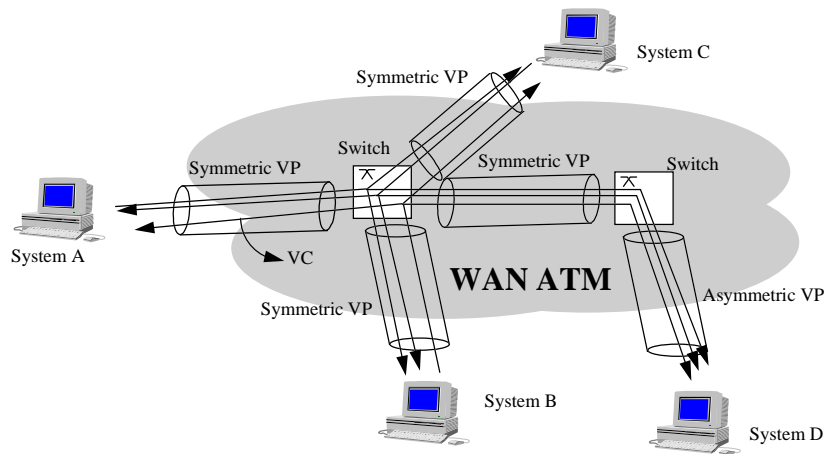


Figure 5: Proposed Network Design in the Wide Area.

Notice how the traffic is balanced, and how the network trunks VPs requested are dimensioned to support just the aggregated traffic, and not n times the traffic sent by one source, being n the number of sources. Following the model of multimedia traffic, the resources corresponding to each VP are the ones of the high-traffic state plus $(n-1)$ times the low-traffic state. For example, an application generating a peak of 1.5 Mb/s in high-traffic state and a peak of 150 Kb/s in low-traffic state, in a session with 4 source members and 4 watch-only members would require a peak 1.95 Mb/s for each VP, instead of 6 Mb/s.

This design imposes some state-of-the-art functions on the ATM equipment. First, switches must be able to combine traffic from different incoming point-to-multipoint VPI/VCI within a single output VPI within different VCIs. Moreover, the switches must have traffic-shaping capabilities on output ports per VP. It is not enough to have appropriate policing at the sources such that the addition of all the incoming peaks be less than the peak traffic reserved for the VP. For example, three incoming VCs with a source sending no more than 1 Mb/s cannot be directly mapped over a VP with a peak rate of 3 Mb/s, unless the aforementioned shaping on output is provided. The peak traffic is reflected in the inter-cell arrival time, not in per-second averages. Figure 6 illustrates this fact.




Figure 6: Peak addition requires shaping on output

It is not clear when the PNOs will offer multicast as a network service. In this case, the user should provide multicasting switches, ideally attached directly to the public switches that are nodes in the acyclic graph. The VPs would then be established, across the public network, between end-systems and these private switches, and between the private switches themselves. If private switches are not available at the required topological locations then the spanning acyclic graph will not be optimal, i.e. it will not be a shortest path, but the design criteria will still be valid.

In LANE approach, the topological situation of the BUS function is performance-dependent. Consider the LANE network in a WAN environment shown in Figure 7. Multicast traffic is sent from LEC_a , LEC_b and LEC_c to the BUS through three *Multicast Send VCCs*. By default, each of these LEC would independently signal the ATM network to establish a VC with the peak traffic corresponding to each source. If the application generates a peak of 1.5 Mb/s in high-traffic state and a peak of 150 Kb/s in low-traffic state, each VC would be established with a peak 1.5 Mb/s, consuming 4.5 Mb/s over the trunk line. It would be much wiser to ask for a shared 1.8 Mb/s VP to carry the VCs corresponding to a, b and c. The network resources used, and consequently the line cost paid, would be cut almost by three, making advantage of the traffic correlation. As already mentioned, Node A must make a per VP traffic shaping in order to accommodate the multiplexing of the three VCs to a VP without losing cells. In the downstream direction, from BUS to LECS using the *Multicast Forward VCC*, the resource and shaping problems do not appear due to serialization made by BUS.

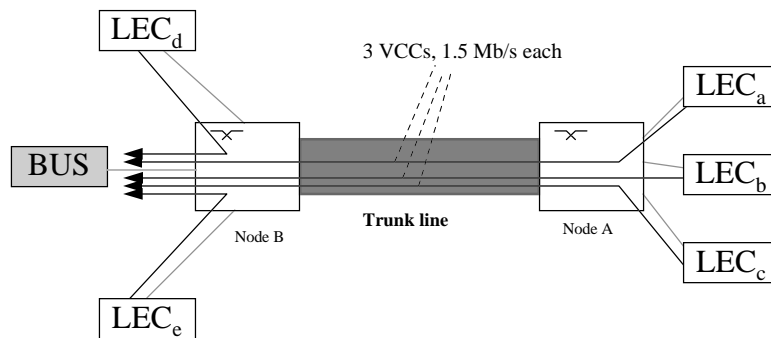


Figure 7: LANE aggregation problem

In the MARS case, the problem is exactly the same, with the difference that it appears only downstream by the combination of all sources. If sources allocate VCs in an independent manner, the resources demanded to the network would be much higher. In Figure 5, assuming the same traffic patterns just described for LANE, the resources required on the internodal trunk line would require 4.5 Mb/s for the independent allocation and 1.8 Mb/s in the allocation taking into account the traffic correlation.

The main disadvantage is the management complexity. As each source-member must establish a point-to-multipoint VC to all members in the workgroup, coordinating their requirements to demand a shared VP, appropriate member-to-member and member-to-network protocols should be designed. As members and sources may entry/exit during an interactive period, the protocol should include procedures for these cases. The MARS address server solves part of the problem. Group addressing and automatic configuration mechanisms based on signaling protocols will also contribute to the solution. Work in progress through UNI 4.0 from ATM Forum will support group addressing and a point-to-multipoint root and leaf initiated.

5. Case of Study: Summer School 95

During June 1995 “Third Summer School on Advanced Broadband Communications” [16] was held between 10 sites of 7 countries in a european-wide scenario. For this purpose it was used the ISABEL CSCW application that conforms to the workgroup traffic model presented in section 4.1. The network was built over the pan-European ATM Pilot that offers unicast peak-limited end-to-end PVP service.

The network design used to support the event was an approximation to the one proposed in this paper because of several restrictions. First, as the public network did not offer multicasting service, it was provided by private switches. These were available at the Madrid and Bassel nodes, and this implied the VP topology used. Second, as per-VP traffic shaping on output was not available, it was not possible to multiplex VPs by means of VCIs. Consequently, it was necessary to multiplex the VPs at the IP level, thus implying traffic serialization based on IP datagrams. This was performed by traffic aggregation at the tree nodes using a multicast server based on ATM workstations.

The network topology was the tree shown in Figure 8, with 6 international PVPs and 3 national PVPs. According to the multimedia application traffic correlation, each bidirectional VP bandwidth was dimensioned to 16,600 cells per second (~6.3 Mb/s). Trunk lines between countries of the ATM Pilot are typically one or two E3 34 Mb/s lines. Access lines between sites and the pilot nodes depend on the country and site and are typically STM-1 155 Mb/s or E3 34 Mb/s.

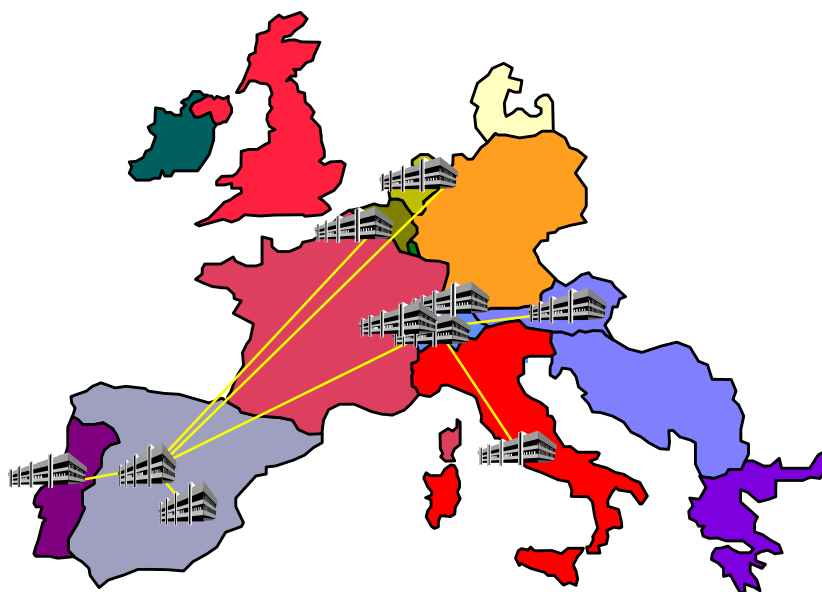


Figure 8: VP Network Topology for Summer School 95

In order to perform traffic aggregation and distribution a spanning tree was built having ASPA (Switzerland) as node and DIT-UPM (Spain) as the root. End systems connected to either the tree node or directly to the root as shown in Figure 9. Upstream traffic is aggregated using a specific Workstation placed at the tree nodes (ASPA and DIT-UPM). Downstream traffic is distributed using point-to-multipoint VCs generated at ATM switches placed at the tree nodes and carried over the pilot network VPs, as shown also in Figure 9 (VPI/VCI labels are shown for each connection, arrows show traffic flows).

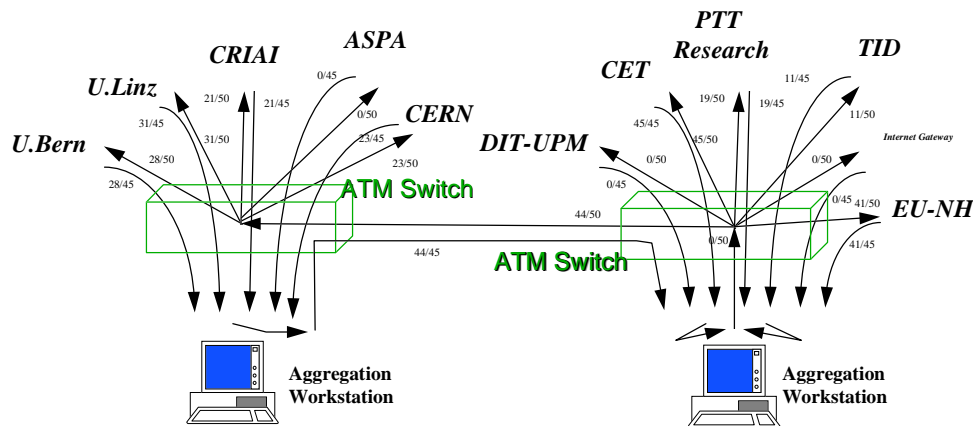


Figure 9: Network Design for Summer School 95

This design allowed a very simple management at the leaves, as it was only necessary to configure the end-system (ISABEL Workstation) to output traffic through VCI 45 and receive it through VCI 50. Simplification of configuration was important because sites participating in the Summer School changed during time and the configuration was carried out by persons belonging to different institutions, with part-time participation in the project.

An undesirable effect of centralized adirectional approaches is that a source+member receives a reflexion of its own traffic. It could be filtered based on the source IP address, as is suggested in a variety of the IETF MARS approach based on multicast servers, but in our case we performed it at the application layer.

This network model served to run the ISABEL application in various configurations. Audio configuration was always one speaking site at 16 bit linear with 22 KHz sampling. Typical video configurations were:

- One lecturing site at 768x567, true color, 8 frames/s and 9 listening sites at 120x90, true color, 3 frame/s each.
- One question site and one answering site, both at 480x360, true color, 13 frames/s and 8 listening sites at 120x90, true color, 3 frame/s each.
- Nine panneling sites at 360x270, true color, 6 frames/s.
- Three panneling sites at 400x300, true color, 10 frames/s and 7 listening sites at 120x90, true color, 3 frame/s each.

Although traffic sent is not linear with uncompressed image quality, it is clear that the aggregated traffic is much lower than n times the traffic of one source at top quality. Unless traffic optimization had been applied, we would have required around 8 times more bandwidth for each VP from the pilot. The reader may take a look to these real figures and evaluate the cost-effectiveness of trying to optimize the WAN.

6. Conclusions

State of the art multimedia and workgroup applications require communication services not directly solved with today's technology. Current efforts at leading engineering groups as ATM Forum and IETF are sketching excellent solutions in the LAN environment. These solutions present some problems when

scaling them to the WAN environment. In particular, neither the ATM Forum, nor the IETF have mechanisms that allow the adaptation of resources required to specific traffic patterns.

WAN network design for multimedia applications require a precise traffic characterization in order to save resources by tuning the network design. This paper shows an approach to structuring the network in order to reduce the required resources. This work is focused on multimedia applications in which traffic sources are correlated. CSCW, teleteaching or audioconference are examples of applications that conform to the traffic model described in this paper.

The VLW network design proposed in this paper is complementary to the approaches followed by ATM Forum and IETF. This design is a distributed, optimal performance solution that makes advantage of traffic correlation to minimize resources at the WAN level. The most relevant concepts of the design have been validated in a large scale experiment including 10 sites in seven countries.

An important issue still open in the proposed model is the configuration complexity. Mechanisms for automatizing resource allocation and configuration appropriate to each situation are still to be developed. These mecanism are greatly dependant on forthcoming standards on ATM multicast signalling and group addressing. The current trend on the development of UNI 4.0 seems to be very appropriate to greatly simplify the procedures to automatize VLW management.

Acknowledgements

We thank the development teams of DIT-UPM, TID and CET for all their support during the installation and development of Summer School'95. We also thank all the participants in IBER, BRAIN and ISABEL for suggesting, criticizing and helping to refine many of the ideas presented here.

References

- [1] S.King. "Switched Virtual Networks" Data Communications. September 1994
- [2] ATM Forum LAN Emulation Sub-working group. "LAN Emulation Over ATM Specification - Version 1.0" January 1995.
- [3] G.Armitage. "RFC Draft: Support for Multicast over UNI 3.1 based ATM Networks. Bellcore. May 1995.
- [4] S. Deering, "RFC 112: Host Extensions for IP Multicasting". Stanford University. August 1989.
- [5] J.Moy. "RFC 1585: MOSPF. Analysis and Experience" Proteon Inc. March 1994
- [6] D.Waitzman, C Patridge. S. Deering "RFC 1075: Distance Vector Multicast Routing Protocol" November 1988.
- [7] S.Deering & others. "RFC Draft: Protocol Independent Multicast (PIM): Motivation and Architecture". January 1995.
- [8] Zhang, L. "RFC Draft: Resource Reservations Protocol (RSVP)- Version 1 Functional Specification" March 1995
- [9] J. Le Boudec, "The Asynchronous Transfer Mode: a tutorial". Computer networks & ISDN Systems n. 24. 1992
- [10] ITU-T "Rec E.164. Numbering plan for the ISDN era" 1991.
- [11] P.Newman. "ATM Local Area Networks". IEEE Communications. March 1994.
- [12] Hong Linh, W. Ellington & others. "LAN Emulation on an ATM Network". IEEE Communications Magazine. May 95.
- [13] M. Laubach, "RFC 1577: Classical IP and ARP over ATM". Hewlett-Packard Laboratories
- [14] J. Heinanen, "RFC 1483: Multiprotocol Encapsulation over ATM Adaptation Layer 5". Telecom Finland, July 1993.
- [15] ATM Forum. "ATM User-Network Interface Specification Version 3.0" Septiembre 1993
- [16] T. de Migueis, J. Quemada, A. Azcorra, M.Duarte,R Aquiar, S. Plagemenn "ISABEL - Distributed Multimedia Application Experiments over Broadband Networks" BROADBAND ISLANDS '95 University College Dublin. September, 1995