ASYNCHRONOUS TRANSFER MODE NETWORKS

Edited by

Yannis Viniotis
North Carolina State University
Raleigh, North Carolina

and

Raif O. Onvural
IBM
Research Triangle Park, North Carolina

Plenum Press • New York and London

ASYNCHRONOUS TRANSFER MODE NETWORKS

Edited by

Yannis Viniotis

North Carolina State University
Raleigh, North Carolina

and

Raif O. Onvural

IBM
Research Triangle Park, North Carolina

Plenum Press • New York and London

Proceedings of TRICOMM '93, held April 26-27, 1993,
in Raleigh, North Carolina

ISBN 0-306-44486-0

A Division of Plenum Publishing Corporation
233 Spring Street, New York, N.Y. 10013

Printed in the United States of America

ASYNCHRONOUS TRANSFER MODE NETWORKS

PREFACE

Broadband Integrated Services Digital Network (B-ISDN) is conceived as an all-purpose digital network supporting interactive and distributive services, bursty and continuous traffic, connection-oriented and connectionless services, all in the same network. The concepts of ISDN in general and B-ISDN in particular have been evolving since CCITT adopted the first set of ISDN recommendations in 1984. Thirteen recommendations outlining the fundamental principles and initial specifications for B-ISDN were approved in 1990, with Asynchronous Transfer Mode (ATM) being the transfer mode of choice for B-ISDN.

It seems fair to say that B-ISDN concepts have changed the face of networking. The expertise we have developed for a century on telephone systems and over a number of decades on packet networks is proving to be insufficient to deploy and operate the envisioned B-ISDNs. Much more needs to be understood and satisfactorily addressed before ATM networks can become a reality.

Tricomm'93 is dedicated to ATM networks. The technical program consists of invited papers addressing a large subset of issues of practical importance in the deployment of ATM networks. This is the sixth in a series of Research Triangle Conferences on Computer Communications, which emerged through the efforts of the local chapter of IEEE Communications Society.

We would like to thank all speakers who participated in the technical program, and, Mr. Len Felton, IBM, for his keynote speech. Ms. Margaret Hudacko, Center for Communications and Signal Processing, North Carolina State University, has patiently worked with us in the organization of the conference, and we are grateful to her for making it all possible. We also would like to thank Ms. Mary Safford, our editor, and Mr. John Matzka, both at Plenum Press, for publication of the proceedings.

Yannis Viniotis
Raif O. Onvural

COMMITTEES

Conference Chairs

Yannis Viniotis (NCSU)
Raif O. Onvural (IBM, Research Triangle Park)

Program Committee

Harry G. Perros (NCSU)
Arne A. Nilsson (NCSU)
Gerald A. Marin (IBM, Research Triangle Park)
Dan Stevenson (MCNC, Research Triangle Park)
Enis Erkel (BNR, Research Triangle Park)

Local Arrangements

Margaret Hudacko (NCSU)

In cooperation with

IBM, Research Triangle Park
Center for Communications and Signal Processing, NCSU
IEEE Communications Society, Eastern NC Chapter

CONTENTS

SOME OBSTACLES ON THE ROAD TO ATM

Daniel Abensour, Jean Calvignac(*), and Len Felton

RTP Laboratory
IBM Corporation
Research Triangle Park, N.C.

* La Gaude Laboratory
IBM France

Abstract

The introduction of ATM will take place in the local environment first. ATM will make very high speed LANs a reality. These future LANs will be based on switches, unlike their low speed predecessors which used shared media. ATM technology will contribute to making the distinctions between LANs, MANs and WANs eventually disappear. ATM will also be the transport mechanism of choice for multimedia. However there are some alternatives to ATM. One of them is FCS which can play the role of a very high speed LAN and is positioned with respect to ATM in this paper. When it is realized, the vision of a "world wired with fibers" where the videophones have replaced the conventional telephone sets and where video-conferencing are part of the normal working environment, will draw heavily upon ATM. We are discussing some of the significant problems which remain to be solved such as the distributed support of video-conferencing. The likely rollout of ATM products and the different approaches considered to address these problems are addressed here.

Introduction

Two new magic words have recently appeared in the vocabulary associated with high speed networking: ATM (Asynchronous Transfer Mode) and multi-media. The latter refers to a new family of applications involving voice, video, still images and data. The former is one of the technologies that satisfies some of the requirements of multi-media applications, namely support of isochronous traffic. From a network standpoint, two of the key characteristics from these applications are a considerable demand in bandwidth and, in some cases, a requirement for minimal delay. Today the existing multi-media applications use *separate networks* for voice, video and

Asynchronous Transfer Mode Networks, Edited by Y. Viniotis
and R.O. Onvural, Plenum Press, New York, 1993

data and demonstrate a limited integration of functions. ATM has been conceived and produced by the telecommunications carriers to accommodate the support of voice, video and data in the same network. ATM, belonging to the packet switching technologies[1] can offer a complete granularity of bandwidth and yet accommodate isochronous traffic requirements.

It is worthwhile mentioning that when the first CCITT documents [1] were produced, ATM was only presented as *the transfer mode for implementing B-ISDN*. Since that time two things have become clearer:

- ATM has gained incredible momentum and the technology is now being "borrowed" by a wide combination of industries with computer manufacturers at their forefront. As a consequence it is now evident that ATM technology will first appear in the Customer Premises Environment (CPE).
- In the carrier world ATM is definitely making more inroads than B-ISDN (Broadband Integrated Services Digital Network). It is not clear when B-ISDN will be offered by the carriers as a collection of services while ATM based services (advertised that way) can be expected within a few years.

In this paper we describe the rationale for a roll out of ATM in the CPE. We also describe and discuss some possible alternatives. We are focusing on some of the technical problems created by the advent of ATM in the local environment, namely the coexistence of ATM switches with PBXs and the evolution of the support of classical voice functions with the introduction of video-conferences. The reader is referred to [2] for a detailed description of ATM.

Alternative Configurations in the local environment

Connectivity in machine rooms

In modern machine rooms, connectivity is changing. Both in the main-frame environment and in the super computer environment we are facing an increased and an almost "any to any" type of connectivity. In the main-frame world, the advent of a new and faster channel operating in conjunction with a dynamic non blocking switching mechanism [3] allows any CPU to be connected to any other CPU using a Channel To Channel (CTC) protocol and also allows any CPU to access any control unit thru the switch. The corresponding paths are established in a dynamic fashion and the set up of each connection is performed in a remarkably short amount of time.

In the super computer world, a similar process has taken place. The HIPPI (High Performance Parallel Interface) standard [4] includes a "switch fabric" which allows a dynamic connection of the various end systems. These systems can be heterogeneous supercomputers, archival storage devices, array DASDs, main-frame com-

---

[1] "Fast Packet Switching" is more and more often associated with ATM, even though other technologies qualify to the name.

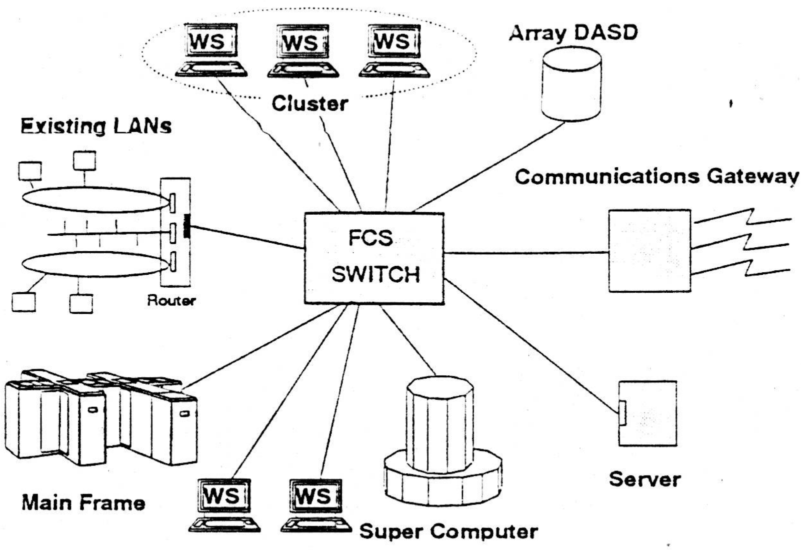

Figure 1. Potential FCS based configuration.

puters or even workstations. The utilization of a HIPPI switch results in a quite powerful but also quite expensive LAN model.

We will now see that the follow on to HIPPI, FCS (Fiber Channel Standard) [5] [6] [7] can and will play the same role, but in a much more affordable manner, and that FCS could be a real alternative to ATM in some cases. FCS was originally conceived by "channel" people as a replacement for HIPPI, IPI3 (Intelligent Peripheral Interface) and SCSI (Small Computer System Interface) types of interfaces. However, it rapidly became clear that FCS had real network attributes, at least in the local area. FCS has a wide range of speeds, up to 1.0625 Gbps or full speed, but lower speeds are also defined (half speed and quarter speed and even 1/8 speed). It covers distances[2] up to $2 \times 10Km$. FCS initially had 3 modes of operation or classes (class 1 is a point to point connection with guaranteed bandwidth, i.e. "circuit switched" type service, class 2 is connection-less service with notification of delivery or failure to deliver and class 3 is also a connection-less service with best effort delivery which is equivalent to a datagram mode). It is interesting to note that a class 4 has been recently proposed to support isochronous traffic, with the support of multi-media applications in mind.

Figure 1 illustrates a potential configuration based on FCS. Computers and workstations are directly attached to the FCS switch and communicate through the switch. As in the HIPPI case, it is expected that super computers and main-frames will also be connected using standard protocols and using paths going through the FCS switch. Again as in the HIPPI case, various devices can be directly attached to

---

[2] 10 km between the FCS switch and an attached device results in a potential "diameter" of 20 km.

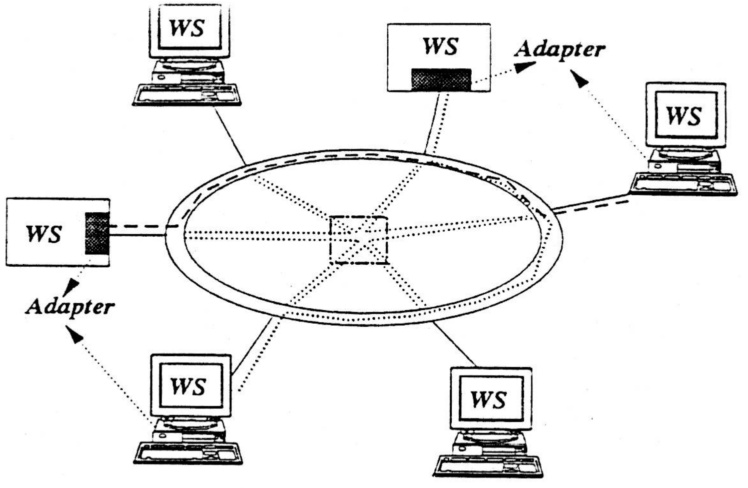

Figure 2. Shared medium based LAN

the switch. Therefore, these devices are accessible by the various processors represented. These devices may include disks arrays, frame buffers as well as print and file servers. Machines like routers may be used to allow the "old" workstations which do not have FCS adapters to communicate with the new systems and the role of the router will be to perform the required routing and protocol conversion functions. Not only will the workstations be able to communicate with each other but they can also form a *cluster*. The group of *clustered* workstations can behave like a single logical processor. This type of closely coupled processors has been developed in the past but typically each processor (or workstation) was using direct links in a point to point fashion to connect to its neighbors (using proprietary protocols). With a central FCS switch, these logical high speed links are defined through the switch and provide an equivalent fully meshed configuration with FCS connections (and using standard FC protocols). The FCS speed is a key ingredient for this type of application[3].

An important aspect of the utilization of FCS as a local network is that all the communications with the outside are funneled through a single or through multiple "gateways". Many different systems can provide the role of these communications platforms.

Switch based LAN model

We are suggesting that a switched based LAN model will be used for high speed connectivity. The classical LANs using a shared medium approach seem to be at a disadvantage when speed increases beyond 50-100 Mbps. Figure 2 depicts a clas-

[3] Obviously one of the key problems is the software required for the clustering (control programs and sub-systems). FCS role is to provide the necessary high speed connectivity and FCS does that, but only that.

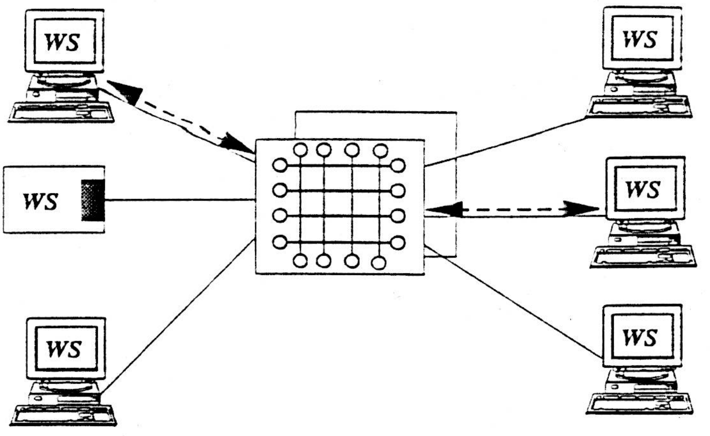

Figure 3. Switch based LAN

sical shared medium based LAN where all the attached elements (workstations, servers) compete for a common resource: the bandwidth of the shared medium. Each element has an adapter which allows the attachment to the shared medium. Each adapter must work at medium speed, even if the connection between every pair of attached elements offers only a fraction of the bandwidth of the medium (shared resource). For example if the LAN is an FDDI ring at 100 Mbps and if two workstations exchange data on the ring at, say 10 Mbps, the adapters must still work at 100 Mbps. With higher speeds, it seems that there is a significant cost difference. In addition the delay resulting from the processing by each intermediate workstation can also be significant when the two communicating nodes are separated by many of these intermediate nodes. From a topology stand point, there are no real differences between a "ring" and a switch based LAN given that rings are almost always using "wiring cabinets" and that their topology is actually more of a star than anything else. The aggregate throughput of the shared medium type of LAN has an upper limit determined by the bandwidth of the shared medium. The situation is going to be quite different with a switch based LAN where it is possible to extend the aggregate bandwidth in a virtually unlimited manner.

Figure 3 represents a model for a switch based LAN. Its major disadvantage is that it requires a switch... and the means to control the switch. This is an added cost and an added complexity. However, there are also advantages. With this approach the complexity of each end node adapter is significantly decreased since it only has to match the speed of the actual connection through the switch. Also, when a connection is established between a pair of nodes, the corresponding "bandwidth" resource is not shared. The aggregate throughput of the system can be increased by the addition of other switches in cascade or in parallel. The FCS approach follows this model and, as we shall see, so does ATM[4]. The debate comparing

[4] For ATM there is another advantage. The ports attached to the switch do not have to have the same speed. A server could use a much faster connection than a workstation. The control mechanisms will make sure that the information transfer rates are matched (leaky bucket for ATM).

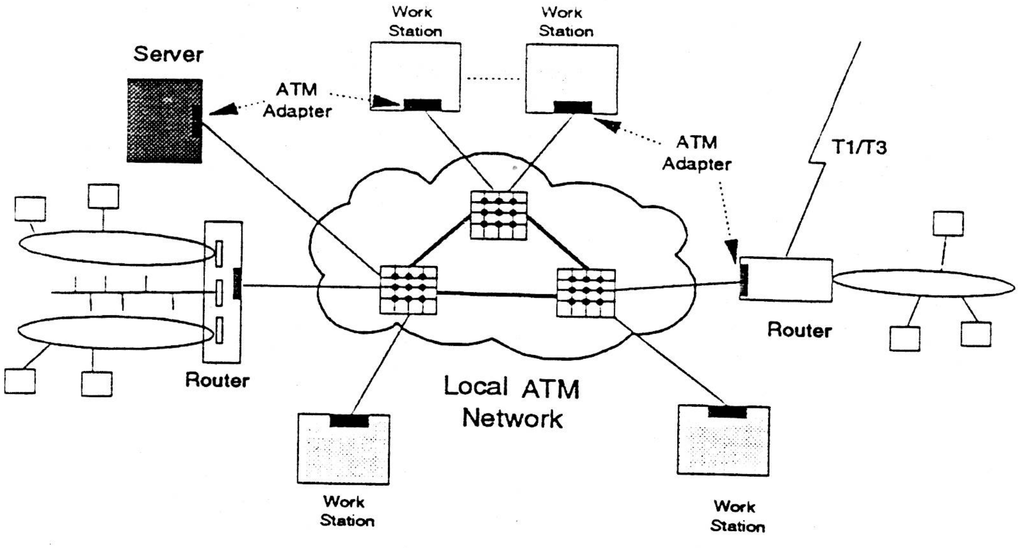

Figure 4. First phase of ATM in the local environment.

switch vs shared medium is not a new debate for LANs. There were similar discussions between PBXs proponents and LANs proponents at the time of the creation of the earlier LANs. It seems that from a technical stand point the advantages of one approach vs an other one go by cycles following the technology. But today for very high speeds and with the advent of new switching technologies and new switch architectures, all the indicators show an advantage to the switch based LAN.

ATM Rollout

Products using ATM will soon appear in the private local environment. Their introduction will be staged and the following paragraphs describe a likely scenario of the corresponding staging.

Stage 1

There are many motivations for ATM in the local environment. The key requirements that ATM address are the support of high bandwidth connectivity and the support of multi-media applications. We show a possible system configuration in Figure 4. This configuration is entirely local.

The ATM switches are systems which follow the new machine room model centered around switches [8] but are extending it to make it a new LAN. Multi-media capable workstations are attached to the ATM switches and communicate with each other through the switch. Also attached to the ATM switch are servers (like multi-media servers). Multiple workstations can have concurrent access to the server. We also find routers attached to the switch, their main purpose is to accommodate the existing LANs (Token Ring and Ethernet or even FDDI) and to allow connectivity between the existing, non ATM capable, workstations with the new ones and the new servers. The routers will have ATM adapters and will perform the necessary conversions between the existing workstations and the new ATM capable workstations.

There may be several ATM switches in the same local area and they will be interconnected by high throughput ATM links as shown in the picture. Initially speeds around 100 Mbps to 155 Mbps will be sufficient to interconnect the local ATM switches. For the workstations speeds between 25 and 45 Mbps will be satisfactory for several reasons. First a 45 Mbps link between a local ATM switch and a workstation can provide the workstation with a non shared 45 Mbps connection and there are few applications today requiring that much bandwidth. Second, the cost of the adapter is strongly influenced by the physical front end part and quite obviously a fiber interface requiring opto-electronics converters is more expensive than a lower speed front end using only copper. Third, lower speeds can be accommodated with Unshielded Twisted Pairs (UTP) at reasonable distances (at least in the local environment) and UTP saves costs of rewiring and installing new fibers. Connections to servers and inter-switch links require higher speeds because of the multiplexing of several simultaneous connections over those links.

Although this ATM model is conceptually very close to the FCS model described above, some major differences exist. ATM and FCS are intrinsically different, ATM uses a fixed 53B cell size while FCS, when operating in its "packet switching" mode, uses variable frame sizes. FCS can also operate in an almost "circuit switching" mode. More significantly, FCS has no networking background and no "communication carriers" support. Its design point was different. For example the optimization of the data transfer between disks and processors at the record level may involve very different algorithms than the ones used in the protocols associated with the transfer of data across a wide area network. ATM and the complementary architectures associated with it [9] are paying a lot of attention to the control of the network. Functions like topology service, directory service and route computation are required in a widely spread network. If FCS is indeed used as a LAN or a collection of interconnected LANs then the same functions will have to be provided[5].

What are the reasons for the utilization of ATM as a LAN in the local environment? We can consider the following:

- Natural evolution to a very high speed connectivity
- Isochronous traffic support
- Potential adapter cost advantage
- Common interface (eventually) for LANs, MANs and WANs...

High speed and isochronous traffic support are indeed functions inherent to ATM. This gives ATM a real advantage over FDDI since ATM can operate at gigabit speeds. To accommodate isochronous traffic, FDDI-II is required, but FDDI-II does not really exist. The adapter cost advantage can result from a) the significant ATM

[5] The same control mechanisms could conceivably also be adapted to, and used with FCS.

volumes and b) the relative simplicity of the adapters. The utilization of the same technology to interface a LAN or a MAN is indeed desirable but may not materialize for quite some time. The ATM interfaces on Figure 4 are all in the local environment and are not standardized by the carriers. On the other hand, the interfaces toward the public network must follow the various carriers' standards. Within the local environment, all speeds and physical interfaces will result initially from agreements between various CPE manufacturers. Currently the ATM forum is contemplating several different physical interfaces:

- 45 Mbps with DS3 framing
- 100 Mbps with 4b/5b encoding (using some of FDDI physical front end)
- 155 Mbps SONET OC3
- 155 Mbps with 8b/10b encoding (à la FCS but with different speed)

It is almost certain that this list will be extended and that some of the above interfaces will not be very successful. For example it is clear that a transmission over Unshielded Twisted Pairs (UTP category 3) at high speed rates (45 Mbps) is highly desirable and extremely likely to happen, while the intricacies of using DS3 framing and formatting (with all the wrap-tests mechanisms and the various levels of alerts) are not really required in the local environment.

In the carrier world, SONET OC3 and OC12 are the interfaces considered at this time. To accommodate the existing carrier networks it is also possible that DS3[6] ATM access be offered.

Note that the speeds considered in the local environment are not very different than FDDI. Yet it is reasonable to expect the support of very high speeds, in the gigabit range, for the applications justifying those transfer speeds, like some of the sought after FCS applications.

Stage 2

The local networks described above will be geographically dispersed and most often will need to get interconnected. Several possible approaches can be considered. Figure 5 shows the utilization of routers. These routers are using high speed private links for their interconnections. These links will normally be dedicated T1 or T3 links, but it is also conceivable to use public switched networks services for these interconnections. The choices offered today in the United States are High Speed Circuit Switched (HSCS) for switched T1 and switched T3, Frame Relay and SMDS. If the nature of the application demands isochronous traffic support then the choice is more limited and only HSCS meets the requirements, given that neither Frame Relay nor SMDS supports isochronous traffic[7].

---

[6] E3 in Europe

[7] IEEE 802.6 describes two modes of operations: pre-arbitrated and queued-arbitrated. The pre-arbitrated mode could lend itself to the support of isochronous traffic but Bellcore had elected not to support it in SMDS.

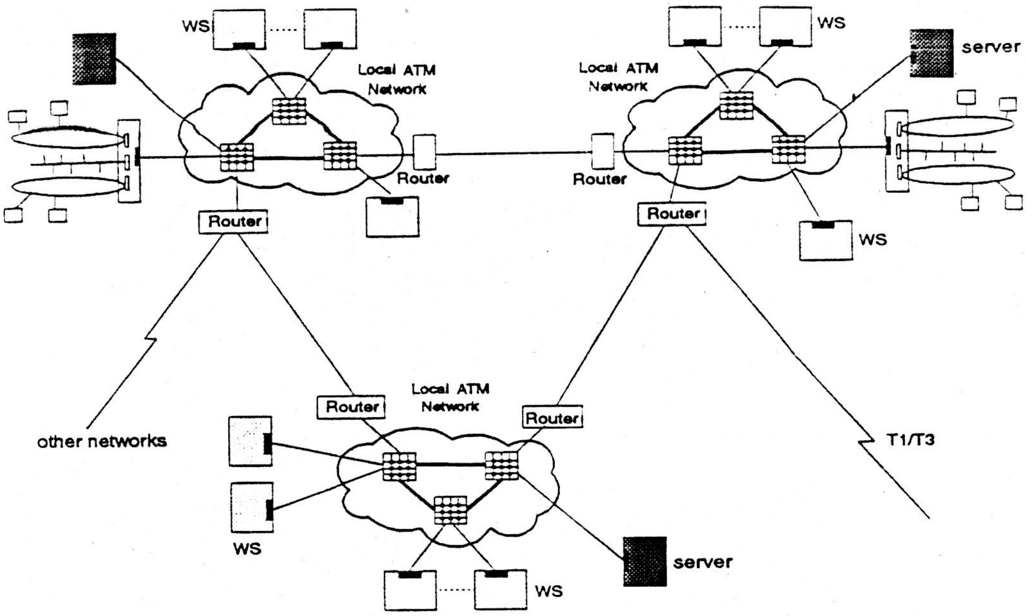

Figure 5. Interconnection of local ATM networks with routers.

For some low bandwidth applications Narrow band ISDN could also be used since it is circuit switching and, by definition, supports isochronous traffic. But then the related bandwidth would be limited to a maximum of ISDN Primary. In all cases, with this approach, the routers will have to adapt the ATM traffic to the wide area networks links used. Using routers to provide support for a wide area network has many limitations. A richer and more complete approach is shown in Figure 6

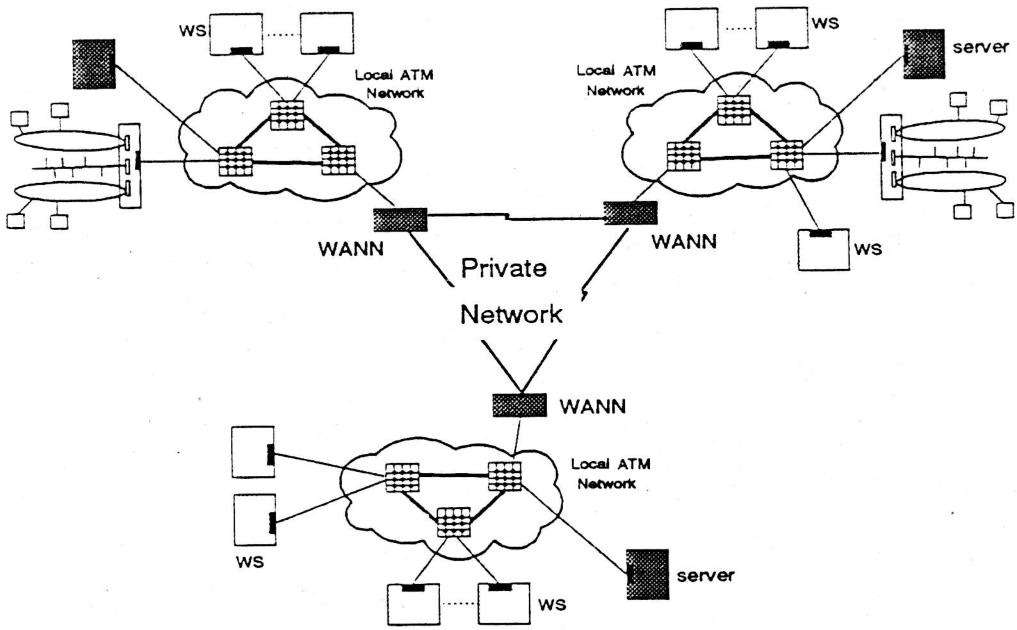

Figure 6. Interconnection of local ATM networks with bandwidth managers.