NTS-13
Trondheim, 20. - 22. august 1996

Experiences with ATM Equipment and Network Operation

Marko Luoma, Kalevi Kilkki, Hanna Karjalainen, Kauko Rahko

**Helsinki University of Technology, Laboratory of Telecommunications Technology
Otakaari 5A, 02150 Espoo, Finland**

Marko.Luoma@hut.fi

Abstract

There are ATM nodes from several vendors in the market and some operators even are offering public ATM services. However, in order to really make a breakthrough of ATM we must develop ATM applications which are acceptable also by others than computer-wizards. Our paper gives an outline of the ATM evolution from separate nodes to real and practical services utilizing ATM technology. One of the most prominent experiences of current situation is continuously evolving standards. First adopters have faced a dilemma of never ending software and hardware updates. Although we may except that Internet will be upgraded to ATM because of increasing capacity and quality requirements, the definitive role of ATM in the whole network architecture is still shrouded in mystery, and a lot of practical and theoretical questions are open.

1. Introduction

In this paper we give an insight to the current situation of ATM technology. It seems, when reading magazines, that ATM has reached maturity. With maturity we mean the level on which you can build networks without fear of hidden pitfalls. When we first started our mission to build up test network, early in the year 1995, we were on the situation where there were many vendors who said that their equipment was ready to deliver all benefits of ATM. This was partly true; they were able to fulfill their promises but not in a standardized way. Nowadays things have changed, ATM has been standardized on the level which ensures multivendor interoperability.

The paper is organized as follows. First, we make a short review on the current opinion in the telecommunication area. Then we present our experience on the building of ATM infrastructure in a university environment including ATM nodes, ATM cards, measurement tools and video servers. Finally we attempt to sketch a picture on the near term future of broadband networks and the most imminent problems in the teletraffic area.

2. History and background

Let us first consider the main advantages of ATM technology. There must be solid ground for introducing a completely new technology in the existing telecommunication environment. Without this ground the future of the new technology will be, usually after a short phase of enthusiasm, very uncertain. It is easy to find several examples of this kind of life cycle in telecommunications: ISDN (Integrated Services Digital Network) as the most famous case, and DQDB (Distributed Queue Dual Bus) seems to be one of the recent downfalls.

So a vital question is if there is a real demand which is very difficult to satisfy with any other technology. Our opinion is that there really is although we have to be cautious when predicting the customer demand. The present condition in the telecommunication and information technology gives us some valuable indications. Firstly, we cannot pass over the Internet boom, and the problems entailed. These problems, especially the poor Quality of Service (QoS) of Internet, certainly need good solutions. Secondly, new video services are steadily emerging, especially if there were enough capacity in all parts of the network to transmit the bandwidth needed by video services with an adequate QoS. Thirdly, the huge investments needed for the realization of a (wide area) broadband network imply that the there must be a workable billing system with an appropriate tariff structure for different services. All these issues should be thoroughly considered independently of the technology selected for the implementation.

As a consequence, any broadband technology intended to public networks must have the following properties:

- scalability from some megabits per second to tens of gigabits per second,
- suitability for integrated real-time services with video and audio components, and
- a proper network management system including billing.

It seems to be almost a unanimous opinion that ATM is the most suitable, currently available technology to meet *all* these requirements. But why has ATM not yet been a real success? Let us make a short review on popular answers. J. McQuillan [1] has presented 10 reasons, most of which are connected with the customer or the market. Despite the technological excellence of ATM there are many obstacles before the benefits can be brought to the end-users. The primary obstacles are obviously related more to the software and protocols than to the ATM hardware: the ATM protocol infrastructure with about 200 000 lines of new code may prove to be too complex [2]. It will take years before all the protocol software components are acting together without problems. Even though it might be possible to build the network and offer a service like videoconferencing, the hidden costs of installation and management can be too high in order to justify the introduction of the new service, see e.g. [3].

One prominent experience of ATM technology is continuously evolving standards. First adopters have faced a dilemma of never ending software and hardware updates. Nowadays the hardware has reached its maturity but the software seems to be updating in steady speed - new software updates are coming out almost monthly. To operate a multivendor environment in such an unstable situation may drive the management people on the verge of despair.

A typical problem is connected with the fact that new applications often require native functions of ATM. Native functionality is brought by Application Programming Interface (API). Unfortunately, API is not yet standardized and therefore constantly changing. In this the interoperability cannot be guaranteed even in updates of single vendor. The worst case is that you have a switch from one vendor, Network Interface Cards (NIC) from an other and an application from a third. When a new UNI specification is finished vendors race for a first update. On that race updates come out on the time scale of half year. If you update your switch you usually lose all the connectivity to the old version of NICs. If you update your NICs you lose your connectivity to the switch and application (which is presumed to use API and which has changed as well). This lack of interoperability forces to wait until all of the vendors have finished their updates, after that you may step one step further and update your whole network. This total update would require time for equipment to be out of service and therefore it might be quite a harmful event.

From the operating point of view there seems to be only little problems as long as someone is willing to do the hard work of connection configurations. Signaled connections are easy to implement in small scale but large networks still need a lot of handwork. Setting up connections is relatively easy when you have one protocol to do that. In contrast, if you have more than one protocol in your hand - as it is with ATM, the situation is much more complicated. Standardized signaling protocols have been developing at a steady speed of one release per year. Interoperability of new releases has been poor and, moreover, the lack of sophisticated functionality in standardized protocols has impelled some vendors to create their own signaling protocols which are not so promising for multivendor environment.

Nevertheless, the market leader Fore has managed to conquer a big share of the market with their SPANS protocol. SPANS has features which are needed when more than one switch operates in an ATM network. This evolution of protocols where several parties are competing with each other, will be prevalent until the standardization process is completed (if that will ever happen). Network operators and users can only attempt to adapt to the situation. In consequence, many potential users of ATM may suspend their investments in ATM technology due to the uncertain future of current equipment.

3. Services

Network is worth of nothing if it has no services. Services are implemented through applications which are steadily but slowly moving towards native ATM. When talking about using applications in an ATM network we cannot yet expect them all to be plug-and-play. In this respect the applications supposed to use ATM network are still imperfect. Actually, in order to really make a breakthrough of ATM technology we must develop ATM applications which are acceptable also by others than computer-wizards.

As we explained earlier, one of the primary limitations is a frequent updating of the software and the hardware relating with ATM arising from constantly evolving standards. Compatibility problems are due to this fact. In order not to run to these problems one has to be aware of evolution process: for example what is the latest release of an ATM-driver in a workstation and what are the applications supporting it.

Another problem area is the applications using Proxy Signaling and Permanent Virtual Circuits. This means that a person who wants to use these applications has to change configurations in an ATM switch. To do that might not be an arduous task, but it certainly will be an obstacle or even a stumbling block to a person who is not used to doing these things - most people are not very familiar with the delicate characteristics of ATM.

Another key service, which is said to solve all problems, is LAN Emulation (LANE). LAN Emulation has been developed to assist easy adoption of ATM. LAN Emulation should emulate characteristics of legacy networks in a way that all of the legacy LAN programs and services would run on the top of ATM. LAN Emulation is simply conversion between legacy addressing and ATM addressing. In the case of Ethernet emulation we have MAC to ATM conversion. What could be so difficult with this? We have had bridges, routers and gateways which have done the same thing. Yes, this is true but in the case of LAN Emulation we have to emulate a connectionless network on top of a connection oriented network. This arrangement results in a tremendous signaling load to the network. In order to cope with this load LANE has been divided into three parts: a client and two servers. A client is a process residing in the workstation or PC which communicates with servers to find an appropriate ATM address to a certain MAC address. Servers are responsible for information gathering. They do all the raw work on the network to find out right conversion for addresses. To configure these functions properly in the case of more than one switch appears to be a work that needs a real expert or people who are willing to spend their nights beside switches.

4. Experiences

To study and exploit new ATM services Laboratory of Telecommunication Technology has driven a project to establish a test platform for new services. This platform has been built up in close co-operation with VTT and Laboratory of Communications. Our laboratory has 2.5Gbps switch (Fore ASX-200) which is equipped with 12 155Mbit/s OC-3c optical interfaces and 4 155Mbit/s UTP-5 twisted-pair interfaces. To have desktop connectivity we have two switches (Madge Collage 280) with 25Mbit/s twisted-pair interfaces. To be able to exploit different services we have acquired a video compressor Nemesys AVA-300 and a decompressor Nemesys ATV-300, network interface cards (Fore SBA-200, Fore PCA-200, Madge ATM 25) and a measurement tool (Adtech AX/4000). The environment is presented in Fig. 1.

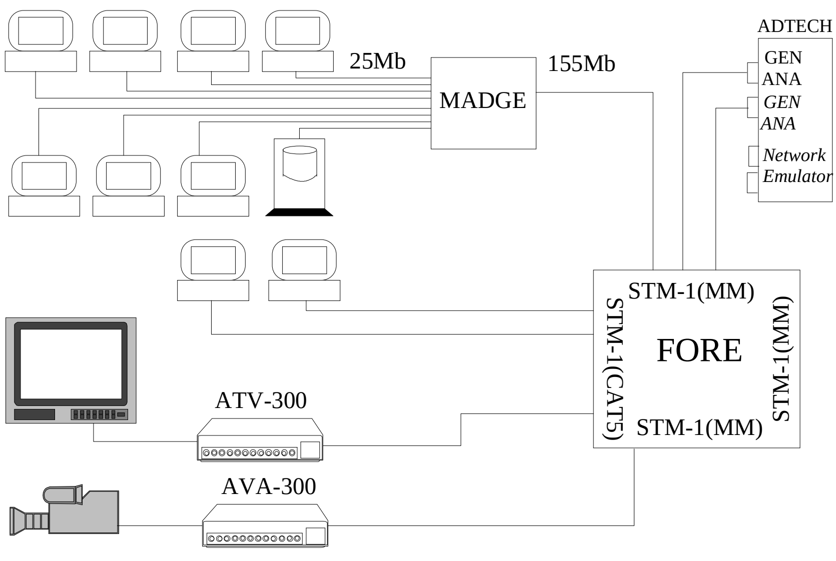

Figure 1 Laboratory environment

Good examples of complexity of current ATM environment are Nemesys AVA-300 and ATV-300. AVA-300 is a device which converts analogue video signal into a digital format, either raw or M-JPEG-compressed, and encapsulates it to ATM cells using AAL5 frames. ATV-300 decompresses M-JPEG-video streams and makes the digital-to-analogue conversion. AVA can work both in PVC mode or using multicast Switched Virtual Circuits but in the latter case it needs manager process, which handles signaling for the device. In order to implement this kind of proxy signaling one still needs to create a couple of PVC, the amount depending on the signaling protocol used.

Software updates have proved to be difficult. New software for the switch comes out every half a year and for NICs little bit more often. It has been common that API has changed with every update on NICs or even if it has not changed, there has been safety lock which prohibits older application to run over a new API. This has been rather annoying since it takes time from application vendors to change their code. Last problem came out when Fore moved to software release 4 and changed their API. Our Nemesys video converters use API stack for proxy signaling. This change in API was not documented and it took several weeks to find out were problem was.

Another curious problem has been signaling stacks. Although vendors say that their signaling uses ITU-T standardized protocol they do not specify whether it is Q.93B, Q.2931. Actually Q.93B and Q.2931 are the same document but they use different Service Specific Coordination Function (SSCF) and Service Specific Connection Oriented Protocol (SSCOP). UNI 3.0 is based on Q.93B and UNI 3.1 on Q.2931. If a vendor has come up with some own idea (like some have) there is an imminent conflict between two switches. They seem to recognize others but do not work correctly.

Operation of Madge and Fore switches has been mainly stable and satisfying. They have implemented LAN Emulation and connection management services inside switches. This on-box solution gives a feasible picture of their operability. Initial configuring is quite easy to accomplish as long as the network does not require usage of centralized LAN Emulation. Madge has rather alluring functionality to switch ethernet and ATM traffic. This bridging function allows easy migration to ATM.

We have accomplished some performance studies with the Fore switch. The detailed results are presented in a parallel paper [4]. The main goals of the evaluation have been to collect information on the performance measurements in ATM switching systems, and to verify their usability in a real testing environment. The results of the measurements met mainly with the requirements stated to an appropriate ATM switch. The switch performed traffic management functions quite well and consistently although there were some minor deficiencies. This enables the spread of ATM networks, because problems in traffic management have slowed down the widespread utilization of broadband networks. Now when problems with late collisions in Ethernet traffic have been solved, the ATM-network build in our laboratory seems to be a suitable solution for campus environment where budget is small and migration time could be quite long.

5. Prospects

In spite of the problems, it is very probable that Internet will be upgraded to ATM in one way or other because of increasing capacity and quality requirements. In contrast, it is not clear what the definitive role of ATM in the whole network architecture will be. One possibility is to exploit the current technology as widely as possibly and to use ATM only as a tool that makes possible to meet the three requirements mentioned above. The Internet protocol (IP) switching is a typical example of this approach [5]. The strength of this concept is that it uses ATM switching for speed, but addressing, routing and bandwidth allocation are accomplished via Internet protocols which are supposed to be less complicated. The real usefulness of this approach is still uncertain because there are fundamental management problems in all networks with both real-time and non-real time services.

Another possibility is to aim straightly at a pure ATM network similar to that in our laboratory. This approach requires, in addition to complete ATM network protocols, economical interfaces for workstations and personal computers, enhanced ATM APIs and general acceptance for this scheme among software developers. Our objective has been to investigate the usefulness of a network architecture which makes possible ATM connections from desktop to desktop. Because this network architecture is especially suitable for video applications, we have begun studies relating to the performance requirements of video services in ATM environment. According to our experiment, end-to-end ATM connections are technologically possible, even though it is somewhat difficult to recognize any commercial video application which positively requires a pure ATM connection.

ATM network may as well be applied for more efficient voice transmission because of its intrinsic on/off characteristic. This approach seems to have even now demand among some network operators, and with multiplying data transmission ATM will have an increasing importance in mobile networks. A standardization effort is going on in this area. Unfortunately, completing the standardization may take so long that several proprietary solutions arise making the wide employment of voice over ATM troublesome, see e.g. [6].

As a consequence of these prospects there is an impending danger of designing an extremely complicated ATM network with specific functions for all services with several network services having different priorities. Although this strategy may appear to be attractive since it promises lots of excellent properties, it will be very difficult to build the network and manage the services. One approach is to keep the network as simple as possible taking into account only the most significant service demands. This means that in the first phase the network may offer only two or three relatively simple services with which the operator may satisfy the three basic requirements: scalability, service integration and network management, see e.g. [7].

6. Conclusions

As a conclusion of experiences it can be said that software updates form the main bottleneck of ATM. It is probable that real breakthrough can be expected when software stabilizes and the API stack is standardized. After that we may suppose that a lot of new and fascinating multimedia applications will emerge.

What could be the conclusion from these insights as regards the teletraffic research? ATM nodes must support both real-time and non-real-time services at the same time, be appropriate to both pure ATM connections and IP-based connections, have an enhanced billing system, and be inexpensive. Every topic has something to do with traffic issues: the coexistence of real-time and non-real-time services induce several fundamental difficulties; the cooperation between IP and ATM worlds is still an unresolved puzzle; and, particularly, tariffs and billing need a lot of theoretical studies and practical experiments. Finally, we must keep the whole system as simple as possible in order to enable inexpensive services.

References

[1] “Internet Poses Opportunities for ATM as well as Threats”, Broadband Networking News, May 14, 1996.

[2] N. Lippis: “ATM’s Soft Side”, Data Communications, May 1996, pp. 31-32.

[3] “Video Still on Pause”, Communications International, February 1996, pp. 22-23.

[4] M. Luoma: “Performance Evaluation of ATM Switching Systems”, Trettende Nordiske Teletrafikkseminaret (NTS-13), Trondheim,1996.

[5] P. Newman, T.Lyon, and G. Minshall: “Flow Labelled IP: A Connectionless Approach to ATM, Proc. IEEE Infocom, San Francisco, March 1996.

[6] K. Taylor: “Voice over ATM: A Bad Connection”, Data Communications, February 1996, pp. 55-58.

[7] K. Kilkki: ATM Node with Real-Time and Non-Real-Time Services, ITC Specialist Seminar, Lund 1996.