

Quality-of-Service (QoS) in Heterogeneous Networks: CLIP, LANE, and MPOA Performance Test

Kai-Oliver Detken
Head of wwl network Bremen

WWL Internet GmbH, Goebelstr. 46,
D-28865 Lilienthal/Bremen, Germany
detken@wwl.de,
Business: <http://wwl.de>
Private: <http://kai.nord.de>

Abstract. Interoperability is a most challenging issue in today's network infrastructures. Especially ATM and IP have to be ever more adaptive and integrative with each other. Additionally, the issue of QoS has not been adequately addressed in today's available heterogeneous platforms. This paper will be discussing the interoperability of different technologies such as the Internet Protocol (IP) and the Asynchronous Transfer Mode (ATM) in a heterogeneous environment, because they are considered the most important protocols in the near future. Furthermore, this paper will be discussing Classical IP (CLIP), LAN Emulation (LANE) and Multiprotocol-over-ATM (MPOA) solutions concerning the adaptation and integration of IP and ATM and their performance. As a manufacturer-independent company, the WWL Internet AG (WWL) tested different already available solutions of different manufactures and compared the results. WWL regularly evaluates new technologies for ones own better understanding and in order to build up special knowledge to support qualified customers. Following there will be, furthermore, discussed the effectiveness of the integration of IP and ATM and this will be accompanied by test results.

1 Introduction

In the Internet and in other networks, Quality-of-Service (QoS) is the idea that transmission rates, error rates, and other characteristics can be measured, improved, and, to some extent, guaranteed in advance. QoS is of particular concern for the continuous transmission of high-bandwidth video and multimedia information. Transmitting this kind of content dependably is difficult in public networks using ordinary best-effort protocols. QoS refers to the performance guarantees that a network can offer to its users, which in turn determines what the network can be used for. QoS is essential to the current and future markets for telecommunications.

Some examples that represent the most rudimentary attempts to provide QoS control can be found in the Internet's Resource Reservation Setup Protocol (RSVP) and in the ATM world. In the RSVP environment packets passing through a gateway host

can be expedited based on policy and reservation criteria arranged in advance. In an ATM network provision is made such that a company or user pre-selects a level of quality in terms of service. QoS can be measured and guaranteed in terms of the average delay at a gateway, the variation in delay in a group of cells, cell losses, and the transmission error rate. Latency, i.e. delay, is one of the basic parameters that determines a network's QoS. The main question to ask is then how latency and performance affects the Internet protocols? Especially the performance has been tested in this article.

2 Comparison of LANE and MPOA

For the adaptation or integration of IP into ATM, there are different methods existing developed by the ATM-Forum (LANE and MPOA) and IETF (CLIP and Multiprotocol Label Switching). For this paper, LANE and MPOA have been tested compared with CLIP, because they are based on each other, they are in ATM products currently available and support legacy LANs such as Ethernet and Token Ring.

LAN Emulation (LANE) is after CLIP the second solution for IP-over-ATM and can be best characterized as a service developed by the ATM Forum that will enable existing LAN applications to be run over an ATM network. In order to do so, this service has to amalgamated the characteristics and behaviors of traditional Ethernet, Token Ring and FDDI networks. Moreover, it has to support a connectionless service as current LAN stations send data without establishing a connection before the operation takes place. Therefore, LANE must support broadcast and multicast traffic such as the kind of traffic allowed over shared media LANs. It has to allow the interconnection of traditional LANs with the amalgamated LAN and at the same time to maintain the MAC address identity associated with each individual device that is attached to a LAN. Finally, it has to protect the vast installed basis of existing LAN applications and enable them to work in the same way as before over an ATM network. This may be put into practise over the OSI layer 2, whereas the LANE technology can be understood as a bridging technology.

MPOA clients established VCCs with the MPOA server components in order to forward data packets or to request information. Using these components, the client is able to establish a more direct path. MPOA supports various kinds of routable protocols (IP, IPX, AppleTalk, etc.), integrates existing internetworking protocols (RFC-1577, RFC-1483, NHRP, MARS, RSVP) and the IETF and ATM Forum solutions (LANE, P-NNI) into a virtual router environment. Yet actually, MPOA supports only IP and based on LANE, RFC-1483, and NHRP. MPOA is a service with layer 3 internetworking support for hosts attached to ELANs, ATM networks and legacy LANs. Thus, the real premise behind MPOA is to provide and deliver the functionality of a router and to take as much advantage of the underlying ATM network as possible. MPOA works as a virtual router on the OSI layer 3. [1]

2.1 TCP/IP-Over-ATM

IP is an important protocol used to achieve interoperability in a heterogeneous network. In contrast, the effectiveness of IP in high speed networks is not widely known. There is much doubt about the TCP performance, in particular, as the acknowledgment mechanisms, overhead size and parameter setting are considered to be obstacles. UDP is more appropriate for real-time data streams, but does not have any security mechanisms.

The TCP/IP protocol stacks were not designed for high speed performance networks in the first place. Several extensions of TCP protocols have been suggested in order to achieve higher performance over these networks and to improve the performance of connections with a high bandwidth delay. New bandwidth-intensive applications such as multimedia conferencing systems and characteristics of high speed networks have triggered research on advanced transport protocols. Therefore, the discovery of additional error sources is not surprising. Following, the reader will find factors identified by WWL that are responsible for inefficiencies of TCP protocols over ATM: [3]

- Send and receive socket buffer size
- Network: Maximum Transport Unit (MTU)
- Protocol: Maximum Segment Size (MSS)
- Transmitter: use of Nagle's algorithm
- Round Trip Time (RTT)
- Receiver: delayed acknowledgement mechanisms
- Transmitter: Silly Window Syndrome (SWS)
- Copy strategy at the socket interface
- Network congestion and lost notice

2.2 Scenarios for LANE and MPOA

The following ATM devices were used during the tests in order to measure the performance of MPOA and LANE:

1. ATM-Switch Centillion 100 (Bay Networks) with ATM module ATMSpeed MDA MCP, software C100 SpeedView 2.2/3.0.1
2. ATM-Switch Centillion 1600 (Bay Networks) with ATM OC-3 module 155M-SMFS, software C1000 7.0(1).2
3. ATM-Switch CS3000 (Newbridge Networks) with ATM155 module (OC-3), VIVID software version 3.0 and system manager
4. VIVID Route Server (Newbridge Networks) with SC Connector and ATM MMF OC-3 155 module, Routing Protocols RIP, OSPF, NHRP and support of IP, IPX, etc.
5. VIVID Orange Ridge (Newbridge Networks) with Ethernet interface 10BASE-T Orange Ridge MMF, and ATM uplink interface.

- 6. ATM NIC from ForeRunnerLE series with throughput of 155 Mbps, bus architecture PCI 2.0/2.1, 32-Bit, 33 MHz PCI Bus and Windows95/NT 3.51/4.0, NDIS 4.0 und NDIS 5.0
- 7. ATM NIC from Olicom Rapid Fire OC-6162: 155 Mbps over MMF (SC Connector), NDIS 3.0 NIC driver

In the first step, the measurement scenarios were defined. In doing so, WWL distinguished between LANE and MPOA tests. LANE was tested without any router as there has been just one Emulated LAN (ELAN) used. This can be considered an optimal way of using LANE, because there is no performance bottleneck. The MPOA scenario was different to LANE, because Logical IP Subnets (LIS) were needed in order to use the forwarding and routing functionality of MPOA.

Nonetheless, a direct comparison regarding the effectiveness and the performance was possible, yet it has to be stressed that LANE did not use a router in the test scenario. In order to prevent a falsification of the measurement functionality, switches like Spanning Tree and dynamical routing were not tested. On the other hand, measurements with various packet sizes were carried out to make sure that performance and latency were represented correctly. Fragmentation, a process carried out by bigger packet sizes, influenced the measurements considerably. Typical packet sizes of IP are 576 Byte (Internet), 1500 Byte (Ethernet), and 9180 Byte (CLIP).

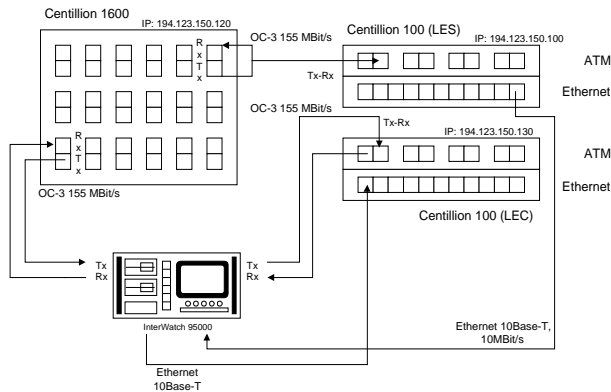


Fig. 1. LANE test scenario with components of Bay Networks: The BUS allows only unknown addresses to get through and these addresses were used for the measurements. The ARP-Request was not used because terminal equipment was missing at the time the tests were carried out. A problem of the benchmark tests was the sending of a ping. Finally, it was send as continuous broadcast.

LANE was tested by a pure Bay Networks scenario. The Centillion 100 has been configured as LAN Emulation Server (LES), whilst the second Centillion 100 has been configured as LAN Emulation Client (LEC). The Centillion 1600 was used as a backbone switch. In doing so, the Centillion 1600 worked in this scenario as a native ATM switch. LANE functionality is generally only to achieve by using an additional module, yet was not required for the tests.

The test device InterWatch 95000 with its ATM OC-3 interfaces was connected to the two Centillion 100 switches. As of the configuration of LES and LEC, the communication between the two has to be directly connected with the switch Centillion 1600 and the InterWatch 95000. The test device was, moreover, connected to the switches of Centillion 100 via Ethernet interfaces in order to produce data traffic and to measure the performance of the Ethernet. Using this scenario, the following tests were carried out:

- Performance Ethernet to Ethernet
- Performance ATM to ATM
- BUS performance

The measurement results differed at the very beginning, because the MAC addresses were not known right from the start of the test. A reply message was sent regularly every few minutes in order to briefly inform about the missing address solved this problem. In doing so, the performance of one link can exactly found out. Since none permanent broadcast has been sent, the results showed a better exploitation of the connection and, furthermore, only the real traffic was measured. The author would like to stress that this is a frequent mistake of benchmark tests.

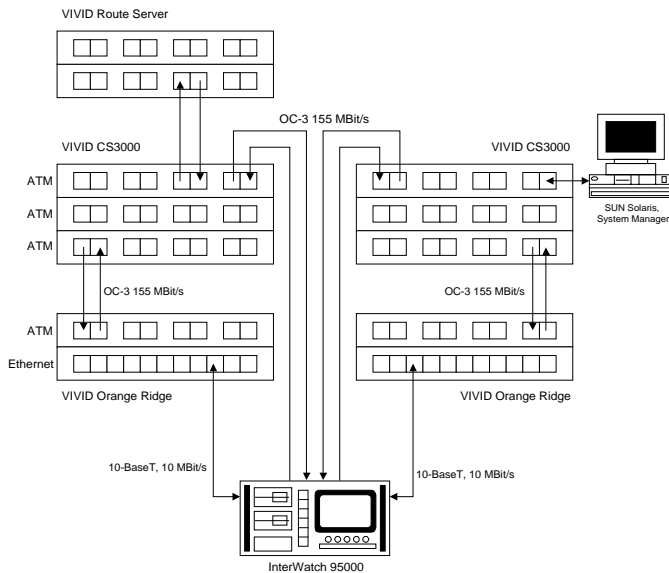


Fig. 2. MPOA test scenario with components of Newbridge Networks: The MPOA measurements differ to the tests with the Bay Networks equipment and LANE. The Newbridge Networks equipment was used with the VIVID series. Two ATM switches CS3000 were used for the backbone. The Newbridge Orange Ridge edge devices which convert Ethernet frames into ATM cells and the other way around have been used for the MPOA functionality. Furthermore, P-NNI was configured in a way to automatically establish dynamical routing. Newbridge devices work only with LANE version 1.0. Therefore, this software was the only one available in order to establish a heterogeneous network with Newbridge components.

During the tests, MPOA operated with its specific parameters. The InterWatch 95000 was similar to the LANE configuration connected via the Ethernet ports to the edge devices VIVID Orange Ridge. Additionally, the test device was connected between the two ATM switches in order to measure the ATM performance. Summarizing, the following measurements have been carried out in performance test Ethernet-to-Ethernet and ATM-to-ATM. Using this configuration, it was possible to conduct research under the same conditions on the Route Server performance similar to the LANE BUS performance. The test procedure was repeated to gain test results which could be compared for the different methods of LANE and MPOA. [2]

2.3 TCP/IP-Over-ATM Scenario

The tests were carried out over a 155 Mbps connection between two separate Pentium-II/350MHz and Pentium-II/266MHz workstations with operating systems Windows98 and Windows NT 4.0. Both clients had an ATM Network Interface Card (NIC) from Fore Systems (ForeRunnerLE155) and also the RapidFire6162 ATM 155 PCI Adapter from Olicom. The central ATM switch was Cabletron's SmartSwitch2000 with a single ATM module and four OC-3 interfaces. AAL-5 was used for all measurements to encapsulated IP packets. The sender buffer varied between 16-64 kbyte, whilst the receiver buffer was always of the same value. Classical IP (CLIP) was used for the scenario without any routing. Therefore, it was a point-to-point connection between both workstations under optimal conditions.

Netperf, developed by Hewlett Packard, was the software used for test-purposes on both clients. It is a benchmark program that to measure the performance of networks. Measurement opportunities for burst traffic and request/response performance are the main tasks of the program. Optional extensions of the software are available. In contrast to the full Netperf version that is only available for Unix operation systems, the used test-version was especially designed for Windows end systems.

2.4 LANE Measurement Results

The switch Centillion 100 was started immediately afterwards in order to reactivate LANE. It comes to an interruption of the signalling packets at full load. Therefore, the switch went down, which is not a good feature in real network environment. Yet, in a real situation this will hardly happen as the BUS normally has to forward just the unknown and broadcast packets. After the establishment of the connection, the link is directly existing between LEC to LEC. Therefore, the BUS does not need to have such a performance and is able to work with short bursts. Figure 4 shows the short measurements which increase to 53.000 cells per second and go down again after a few minutes. In the next step, the exploitation was set at 50% in order to allow a longer operation of the BUS and to measure the unknown packets. Moreover, it was considered a more real environment for a BUS operation. The measurements were stopped after a short residence time.

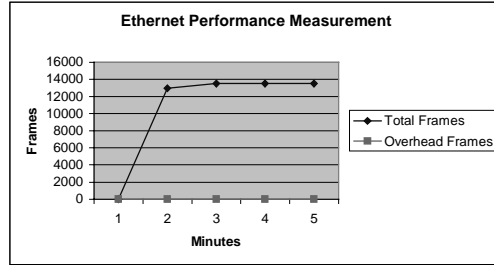


Fig. 3. In order to measure the performance of the Ethernet side, a PING was sent in both directions. In doing so, the measurements could be limited to one port. Furthermore, an exploitation of 100% was adjusted. Figure 3 clearly shows that the throughput of the Frames after a very short time climbs up to 13.500 frames/sec and remains at that level. This figures demonstrate an effective exploitation of 99,4%. Yet, the payload is included and that means you have to add some collisions (27) and the overhead.

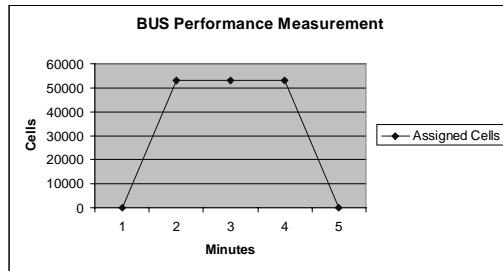


Fig. 4. In a next step, the BUS performance on the ATM side was tested. The maximum BUS throughput is approx. 53.000 cells per second. This represents a data rate of 22,4 Mbps. Full duplex was chosen as the full possible data rate plus overhead had been reached. This result has got independence from the data rate and can be hold only approx. one minute.

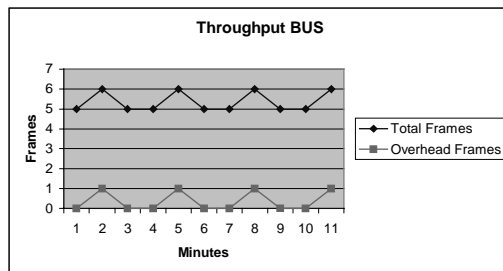


Fig. 5. This figure shows the unknown packet throughput that was responsible for the allocation of the addresses. The unknown packets did not load the BUS as figure 5 illustrates. The BUS sent only 5-6 frames with 0-1 frames overhead.

2.5 MPOA Measurement Results

The MPOA measurements could not be configured directly together with the Inter-Watch 95000, because MPOA had not been not implemented at that time. Therefore, the first tests were repeated as they were also valid for MPOA. That means, the encapsulation of the data traffic and the operation was similarly to LANE.

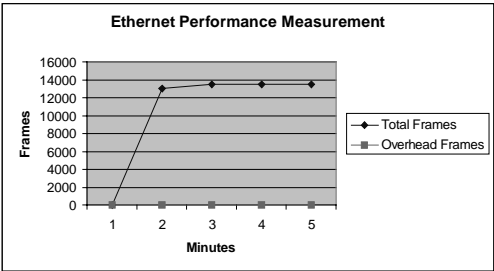


Fig. 6. Ethernet-to-Ethernet performance: The measurements showed the same result as the aforementioned with LANE. Changes of the packet sizes from 100 byte to 1000 byte and 64 kbyte did not influence the test results. The explanation to that is the 10 Mbps data rate. The exploitation was 99,4% with 1.100 collisions.

This test has been conducted with more than two subnets in order to force the router to route and to forward the packets. With MPOA the route server established a shortcut connection to directly pass the data traffic from one switch to the other. 27.000 cells as the maximum throughput was measured. This equals a data rate of 11,45 Mbps. The reason for this data rate was the unidirectional connection that was measured. Otherwise the performance was identical with LANE without router integration. In case of a network with different LIS, MPOA this may be very useful.

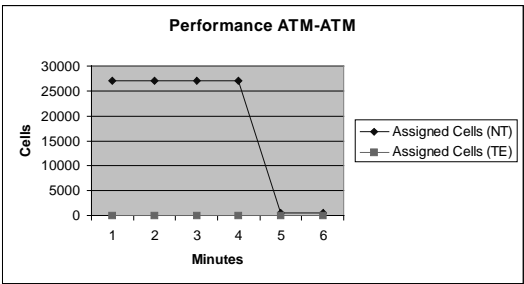


Fig. 7. ATM-to-ATM performance: The second measurement has been conducted with 100% exploitation. After a defined period of time, the route server also had to be restarted, similarly to the BUS at LANE.

2.6 TCP/IP-Over-ATM Performance

The measurement of TCP/IP-over-ATM was carried out by using various packet sizes in order to represent the effectiveness of IP-over-ATM. The throughput alone was not really interesting, because of the different bottlenecks that were already mentioned before. The fluctuations and throughput breaks were more important. Classical IP (CLIP) has been used for this measurements point-to-point, because of the highest effectiveness. The test phase had a time duration of 60 seconds for each measurement via Netperf.

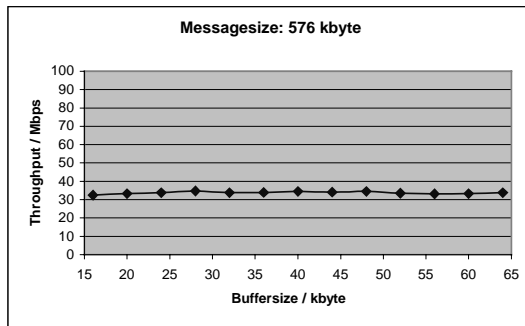


Fig. 8. IP-over-ATM, 576 byte: This figure shows the TCP/IP measurements with a packet size of 576 byte via CLIP. The buffer varies from 15-64 kbyte. The packet size of 576 byte represents a normal datagram in the Internet environment. Figure 8 demonstrates the effectiveness of the achieved throughput during a normal point-to-point session with minimal overhead. By the less datagram size the effectiveness went down to approx. 35 Mbps.

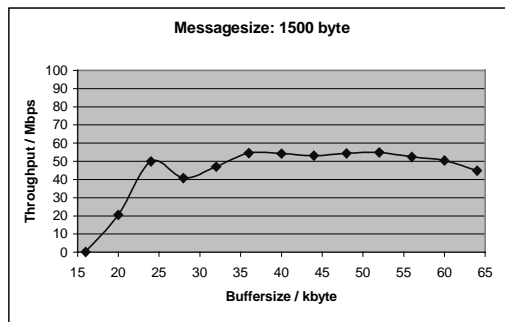


Fig. 9. IP-over-ATM, 1500 byte: Additional breaks and interruptions happened with small and big packet sizes. The MTU of 1.500 byte required a fragmentation of the packets and this has to be considered the reason for it. If the fragmentation was increased the performance went down. Figure 9 shows a differing result. In this case, the packet size was 1.500 byte which allows higher data rates then before. The results are only limited to approx. 60 Mbps. The measurements show some fluctuations, because of the different buffer sizes, which was used.

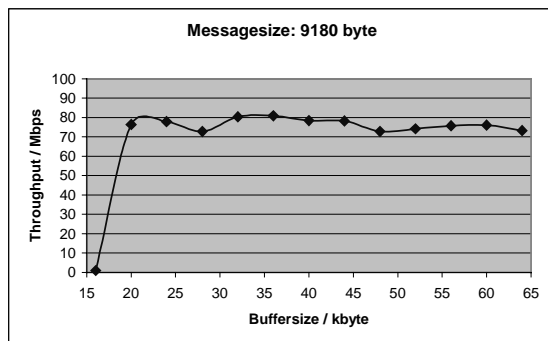


Fig. 10. Fluctuations happened, especially at a buffer size of 55 kbyte. The best results were achieved at 40 and 65 kbyte. A fundamentally improved performance was achieved with a packet size of 9.180 byte. The throughput increased up to 80 Mbps. Fluctuations happened, again (especially on less buffer sizes), like during the other measurements. The small fragmentation of the packets has to be considered the reason as the MTU size was 9.180 byte.

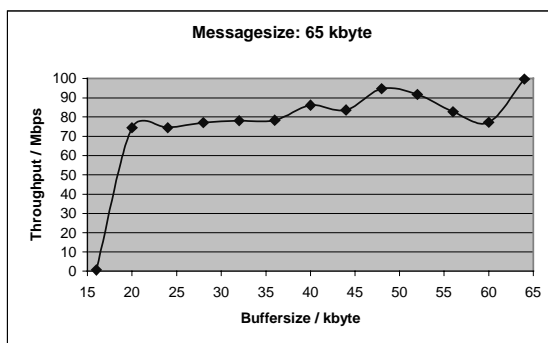


Fig. 11. IP-over-ATM, 65 kbyte: IP-over-ATM, 9180 byte: For the last measurement a packet size of 65 kbyte was chosen. It triggered several fluctuations as figure 11 shows. At first the throughput went down to approx. 5 Mbps. This performance has to be considered poor and was not efficient enough to be used in practice. Otherwise this results show the best data rates till 100 Mbps. Concluding, the big packets are useful for data networks with big buffer sizes and the maximum transport unit (MTU) is responsible for an effective data throughput.

The NIC RapidFire6162 from Olicom together with the LE155 from Fore Systems have been an efficient throughput and high Switched Virtual Circuit (SVC) performance. The LE155 has been also an active processor for direct traffic support without influence the CPU of the workstation. Both NICs need 128 Mbyte for a efficient operating. Additionally, both cards support in MPOA mode different service classes as Constant Bit Rate (CBR), Variable Bit Rate (VBR), and Unspecified Bit Rate (UBR). Available Bit Rate (ABR) was only available at the RapidFire6162. But, for the support of ABR you need this service inside the ATM switch, which was not integrated in every component yet.

3 Conclusions

There are basically three fundamental elements needed for a successful QoS implementation. At first, QoS within a single network element (for example, queuing, scheduling, and traffic shaping tools). Second, QoS signalling techniques for coordinating QoS from end-to-end between network elements. And third, QoS policy, management, and accounting functions to control and administer end-to-end traffic across a network. Three basic levels of end-to-end QoS can be provided today:

- a) *Best-effort service* – also known as lack of QoS, best-effort service is basic connectivity with no guarantees.
- b) *Differentiated service (statistical QoS)* – some traffic is treated better than the rest. This is a statistical preference, not a hard and fast guarantee.
- c) *Guaranteed service (guaranteed QoS)* – an absolute reservation of network resources for specific traffic.

The type of service that is needed depends on the application or problem the user is trying to overcome. There are applications where a best-effort service may be sufficient. There are also problems such as real-time control over a network where the guaranteed service is absolutely necessary. It may also happen that a best-effort customer will migrate to differentiated service and then to guaranteed service. In many applications there may exist an upgrade path from the technology needed to provide differentiated services to that needed to provide guaranteed services. Network performance is influence e.g. by traffic offered and thus network performance should be taken as a random variable and has been analysed from this point of view in this article.

LANE is a very stable standard and simply to configure. The components of Bay Networks showed a reliable and efficient performance. If there was no router bottleneck, LANE had the same performance as MPOA, because of the direct connections between the two LECs provided the BUS found the correct addresses. If there was a router in the LAN, MPOA was necessary. As LANE version 1.0 supports only Unspecified Bit Rates (UBR) no Quality-of-Service parameter could be used. MPOA will solve this problem, because the version 1.1 will support service classes.

The missing standard conformity in software 3.0 of the VIVID components limited the possibility of using all defined functionality of MPOA. Only the Next-Hop Resolution Protocol (NHRP) and LANEv1.0 were supported. Additionally the edge device Orange Ridge had only one ATM uplink that did not allow redundant connections. The management software was not easy to install and only available for Solaris 2.5.0. Newer Solaris-versions were not supported. Furthermore, the series of VIVID used a propriety protocol for the management with the name Control Packet Switching System (CPSS). Other components and devices by other manufactures were not supported and could not even be recognized by the management software. Summarizing, the VIVID components were highly reliable and showed a good performance.

TCP has been extended and further developed for better and more efficient mechanisms in high-speed networks. Yet in practice, TCP quite often does not have the op-

timal throughput. Several parameters, devices and mechanisms were responsible for this effect. This caused fluctuations and delays during the transmission of data traffic. Mainly TCP's own mechanisms were responsible for the small throughput such as acknowledgements, small window sizes and sequence number overloads. Additionally, the performance of the ATM boards and the end stations were important. The adapter boards from Fore Systems and Olicom showed only a maximum throughput of approx. 100 Mbps, if you use CLIP. Earlier measurements conducted in some European projects (e.g. European Information Exchange Services – EIES) got values of 117 Mbps at LANE and 134 Mbps at CLIP with high buffers and two Sun Solaris workstations. Therefore, there are three main issues regarding latency in Internet protocols, which are able to improve the performance:

1. The default “window size” in many TCP/IP protocol implementations acts as a bottleneck on communications over high-latency links. On many implementations, the standard window prevents sending enough data to fill a high-latency connection.
2. TCP includes two essential congestion control mechanisms called “slow start” and “congestion avoidance”. These mean that all Internet connections (such as viewing web pages and sending e-mail) start out at lower bandwidth and then throttle up to higher speed if no congestion is encountered. The problem is that each cycle of speed increase requires a full round-trip communication between sender and receiver, and dozens of such round-trips can be necessary to reach the full potential of a link.
3. There are research efforts to look at increasing the performance of TCP by tricking the connection on the other side into believing it is communicating over a low-latency link. Unfortunately, these schemes fundamentally alter the semantics of TCP communications, introducing the possibility of data corruption. Moreover, they are incompatible with the IP security protocols (IPsec), which promise to bring an unprecedented and badly needed degree of security to the Internet

In order to implement IP-over-ATM effectively, it is necessary to assign enough time to the configuration of the participating devices and software. ATM is a new technology which considerably extends the usable bandwidth with regard to QoS parameters. Nowadays, workstations are designed for work in traditional networks such as Ethernet or Token Ring. These deadlocks must be compensated in order to use efficiently IP-over-ATM.

References

1. Detken, K.-O.: Interworking in heterogeneous environment: Multiprotocol over ATM (MPOA); European conference Interworking98; Ottawa 1998
2. Detken, K.-O.: ATM Handbook; publishing house Hüthig; Heidelberg 1999
3. Detken, K.-O.: ATM in TCP/IP environment: adaptations and effectiveness; European conference ATM Traffic Symposium; Mykonos 1997