Multi-Protocol Over ATM (MPOA): Performance Tests

Heikki Vatiainen, Jarmo Harju, Hannu Koivisto, Sampo Saaristo, Juha Vihervaara
hessu@cs.tut.fi, harju@cs.tut.fi, hak@pori.tut.fi, s156953@cs.tut.fi, vihis@pori.tut.fi

Tampere University of Technology
Department of Information Technology
Telecommunications Laboratory
P.O. Box 553, FIN-33101 Tampere, Finland

Abstract

MPOA is a solution for routed networks to benefit more from the underlying ATM network and its QoS properties. MPOA offers more scalability, performance and lower latency than it is possible get from a router based network.

MPOA builds on many existing protocols and combines them to a scalable layer 3 switched routing solution. In this paper the performance and delay characteristics of an MPOA based solution are studied and compared to a “traditional” router based network where ATM shortcuts are not used

1 ATM on Linux

ATM support for the Linux operating system has been under active development since 1995. The project was started by M.Sc Werner Almesberger and is currently at its 56th release. The project still continues with steady pace and the ATM support will probably be included in the Linux distribution kernels in the near future. [Linux-ATM]

The ATM on Linux distribution supports a wide variety of ATM related protocols and utilities. The features include support for a number of ATM network cards, UNI 3.0, 3.1 and 4.0 signalling, UBR and CBR traffic categories, support for PVCs and SVCs through a Berkeley sockets based API, Classical IP, LANE and MPOA.

Tampere University of Technology has been active contributor to the ATM on Linux distribution. The work at TUT has mainly focused in IP over ATM techniques and protocols. Highlights include LANE service and client support which were developed in earlier projects with the latest addition being support for MPOA clients. [Kiiskilä]

2 What is MPOA?

Multi-Protocol Over ATM is a method for efficient transfer of inter-subnet unicast data in LANE environment. The existing functionality of LANE is preserved while allowing direct communication over ATM VCCs across network layer subnet boundaries. [LUNIv2]

These ATM VCCs, often called as MPOA shortcuts, bypass the routers decreasing the router load and lowering the possibility for network bottlenecks resulting from traffic congestion on often used links.

MPOA introduces MPOA Clients (MPCs) and MPOA Servers (MPSs) which co-operate with the existing LANE clients and LANE service. The MPCs query their MPSs for shortcut information and with that information create and utilize the shortcut connections. The query mechanism is based on IETF’s Next Hop Resolution Protocol (NHRP). [RFC 2332]

MPOA is specified by the ATM Forum. The version 1.0 of the MPOA specification was released in July 1997. [MPOA]

2.1 Problems with the LANE method

One of the goals of Multi-Protocol Over ATM is to resolve the problems with the current LAN Emulation model. The LANE model, which is effective within a network layer subnet, creates bottlenecks when packets need to be routed between subnets. The problems arise from the fact that the physical network topology needs not to correspond to the logical network layer topology which is used in routing.

Figure 1 shows what an ATM based network might physically look like. It is worth noting that the routers are attached to the rest of the network with a single link. This implies that all the routed packets have to pass twice through the same link.

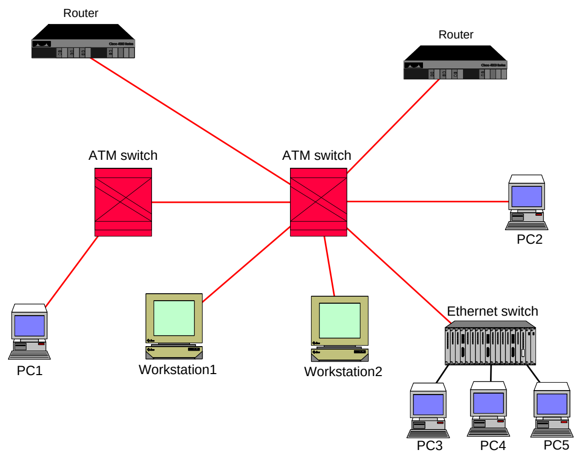

Figure 1: Network layout at the physical layer

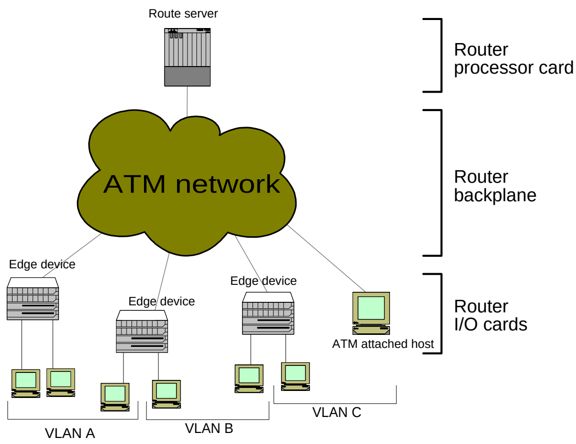

Figure 3: Virtual router

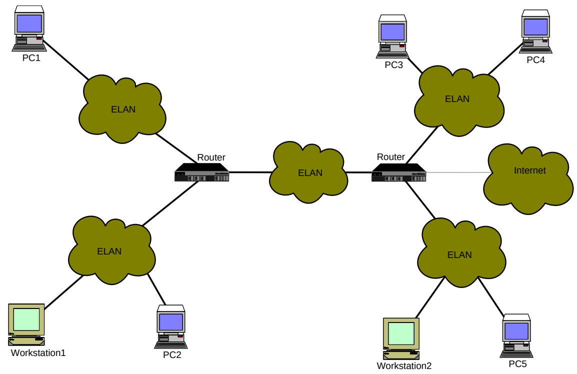

Figure 2: Network layout at the network layer

Figure 2 shows the same network seen from the network layer. The ATM and Ethernet switches are not visible and the routers have a central position in the network.

Even if the physical topology is hidden from the network layer, it still dictates how the actual flow of packets passes through the network. By comparing Figure 1 and Figure 2, it is clear that a high amount of traffic is passing through the routers and the routers' connections to the network can start hindering the network performance.

2.2 Virtual routing

The solution provided by MPOA to the LANE performance problems is the better usage of the underlying ATM network. Two LANE clients, which reside in the same network layer subnet, can establish direct VCCs between each other. MPOA extends direct VCCs across network layer subnets by distributing the routing functions across the MPOA enabled ATM network.

Since the routers are not needed on the data path between subnets, MPOA effectively separates the two functions of router: route calculation and packet forwarding. The route calculation in the MPOA model is done by the MPOA servers, also called route servers, and the packet forwarding is the responsibility of the MPOA clients which reside in ATM edge devices.

This technique, in which routing functions are distributed across the network, is often called virtual routing. The benefits virtual routing has over legacy router based networks include *(a)* effective communication across subnet boundaries and *(b)* increased scalability and manageability since the number of routers can be reduced.

Figure 3 illustrates how a traditional router can be divided to components and how these components correspond to the components of an MPOA system.

3 TUT MPOA test network

A number of different tests and measurements were done with an MPOA test network at TUT. The network consisted of three ATM switches, Fore ASX-200BX, FORE ASX-1000 and Cisco LS1010. The first tests with MPOA enabled equipment were done during the summer of 1998 with one router, Fore ASN-9000, and four ATM enabled hosts. The hosts were two Sun workstations and two Linux PCs. The network was expanded with a Cisco router and additional Linux PCs before the end of 1998.

Both routers and most of the edge devices were connected to the Cisco ATM switch. The ATM switches had multiple connections between them and were using PNNI protocol for call routing. Figure 4 shows the lay-

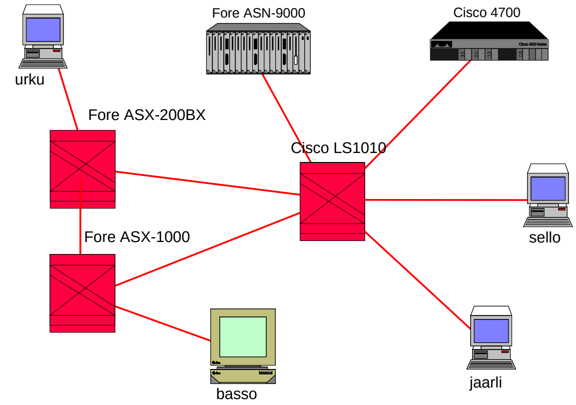

Figure 4: Laboratory test network: physical topology

out of the laboratory network. All the connections between the devices were 155Mbits fiber or twisted pair links.

The MPOA test network is a part of the TUT core campus network but only the devices shown in the Figure 4 were used in the testing.

The hosts were assigned into different IP subnetworks as shown in Figure 5. All the ELANs had LANEv2 service and MPOA enabled LANE clients in them. The ELAN between the two routers was dedicated for routing only.

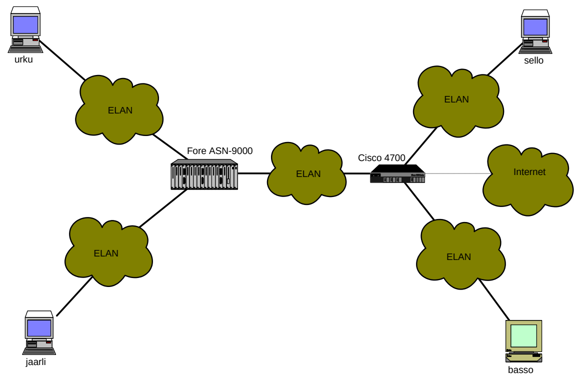

Figure 5: Laboratory test network: IP level topology

In the beginning of 1999 the Pori School of Technology and Economics, which is part of TUT, joined the MPOA test network. The network connection to Pori is an ATM virtual path tunnel with end points in Pori and TUT. The tunnel peak cell rate was unlimited and the total number of ATM switches between Pori and the test network MPOA servers was at least seven. Five switches are under TUT control and at least two are in the telecom operator's public network. The geographical distance between Pori and Tampere is about 180 kilometers (110 miles).

The equipment in Pori consisted of two MPOA clients which were using the test network's MPOA servers. This configuration made it possible to experiment how well MPOA functions with geographically wide VLANs.

4 MPOA interoperability and initial testing

Internetworking between Fore, Cisco and Linux implementations was surprisingly smooth. All three implementations had initially small problems which were reported and fixed during the testing period.

The experiments indicate that multivendor testing in heterogeneous environment is beneficial to testing the software for hidden bugs. Both Fore and Cisco had fatal bugs and some of the bugs were only seen when the other vendor's software did something unexpected.

The Linux implementation also had its problems but the decision to validate and parse packets in user space did not cause the whole system to become unstable even if the MPOA client daemon process, `mpcd(8)`, dies. The design of Linux MPOA client also ensures that the system will still function as a normal LANE client even if the MPOA client parts is not operational or dies.

5 Results of performance testing

The performance tests clearly show that MPOA can increase network throughput by bypassing routers. The Cisco 4700 router was a real bottleneck since its ATM module was only capable of routing about 35Mbits per second. When the 4700 was bypassed, the high end Linux PCs (urku, PentiumPro200 MHz and jaarli, PentiumII 350MHz) got much higher readings from `netperf(1)` benchmarking program.

The Fore ASN-9000 router is capable of routing much higher amount of traffic but its 155Mbits link can easily be saturated with the current PCs and workstations. The problem is common with any single armed router configuration including Fast Ethernet.

All the tests were done using the `netperf(1)` program. Netperf is freely available, comes with extensive documentation, supports a number of different tests and benchmarks and can be compiled on many different platforms. [Netperf]

The netperf runs were done between two hosts using TCP streams. Each stream lasted 60 seconds. Each test was repeated five times in a row and the average

Hosts	LANE Mbit/s	MPOA Mbit/s
urku – jaarli	114.57 ± 1.58	48.06 ± 4.96
jaarli – urku	119.12 ± 0.30	98.20 ± 0.58
sello – basso	30.59 ± 0.08	68.72 ± 0.30
basso – sello	20.19 ± 0.61	49.34 ± 0.17

Table 1: One stream over one router

Hosts	LANE Mbit/s	MPOA Mbit/s
urku – basso	5.33 ± 0.15	36.89 ± 0.64
basso – urku	15.62 ± 0.21	65.60 ± 0.40
sello – jaarli	4.35 ± 0.19	24.42 ± 0.28
jaarli – sello	11.11 ± 0.58	34.14 ± 0.45
Total	36.14	161.05

Table 2: Four streams over two routers

TCP throughput and average difference from average was recorded.

5.1 One stream over one router

The test was conducted by running a series of `netperf(1)` tests across one router. Only one stream was active at a time. The test results are shown in Table 1. Hosts see much better performance over the Cisco router but little lower over ASN-9000. The possible reasons for the lower performance of MPOA are discussed later in Section 5.3.

5.2 Four streams over two routers

The test was conducted by running four series of `netperf(1)` tests across two routers. Four streams were active at the same time. The test results are shown in Table 2. In this case, the combined rate of the shortcut streams is more than the theoretical rate of an 155 Mbit/s ATM link. Even if the routers had been capable of routing at high speed, they would have been outperformed by MPOA shortcuts.

5.3 Comparison between LANE and MPOA performance

The test results show that almost all the tests involving a router saw better network throughput when MPOA was used. The only exception was in Table 1 where the tests over Fore ASN-9000 saw better throughput over LANE than MPOA shortcut. Some of the possible and real reasons for this are given below:

- The Fore router is capable of routing close to wire speed and causes no bottleneck for single streams
- MPOA must do egress cache lookup when receiving packets over shortcuts. LANE has no egress cache and no cache lookup is required when receiving packets
- Even the longest LLC/SNAP header used along the shortcut (12 bytes) is shorter than the data link layer header, usually Ethernet header (14 bytes). The incoming packet must be copied to a bigger buffer before the shortcut encapsulation can be replaced with the data link layer header.

However, the benefits of MPOA are obvious when there are more than one high speed streams involved or routers are not capable of routing at wire speed. Just as Section 5.2 shows, the use of MPOA boosts the TCP throughput by almost 125Mbit/s when compared to throughput seen by the routed LANE streams.

MPOA can also help to extend the usable life time of existing routers. Cisco 4700 serves as an example of a router that does not offer very good routing performance even if the router has a high speed ATM module installed. However, the routing performance does not matter very much if it can function as a route server and let the MPOA clients do the data forwarding.

6 Results from delay variation measurements

A number of experiments were done to measure the delay variation, also known as “jitter”, experienced by packets when using LANE or MPOA. The tests were done by using two streams where one stream was used to generate load on the routers while the other stream was measured for jitter experienced by the packets.

The program used for generating the streams was `mgen` which is part of the MGEN-3.0 package. The accompanying program is called `drec`, which receives the flows generated by `mgen`. With the programs in the MGEN package it is possible to create and measure UDP/IP flows which use predefined amounts of bandwidth and have different traffic patterns. [MGEN]

6.1 7.2Mbit/s streams

An example of a 7.2Mbit/s packet flow generated by `mgen` is shown in Figure 6. The x-axis shows the packet sequence number and the y-axis shows the time delay (delta time) between this and the preceding packet. The main body of the points in the plot are located along two lines, the line in the middle of the figure and the line near the x-axis.

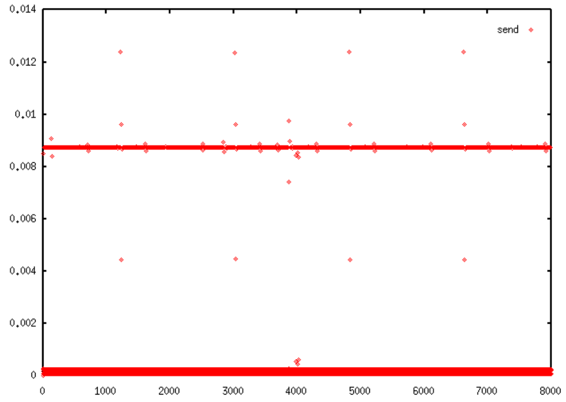

Figure 6: 7.2Mbit/s `mgen` generated packet stream

Figure 7 shows what the packet stream from Figure 6 looks like just before received by `drec` over an MPOA shortcut. The delta times in send and receipt have been plotted together and the plots overlap almost perfectly. Thus the delay variation introduced by MPOA is almost non-existent even if the shortcut VCC crossed two ATM switches.

Figure 8 shows the transmit and receive delta time plots for the packets which were routed over both routers while the 23.3Mbit/s Poisson distributed load was active. The delta time plots of sent and received packets overlap less and the amount of jitter has increased.

As Figure 9 shows, a bit different behaviour was seen with 23.3Mbit/s periodic background load.

As expected, the delay caused by routers is well visible. However, with the periodic load the router behavior

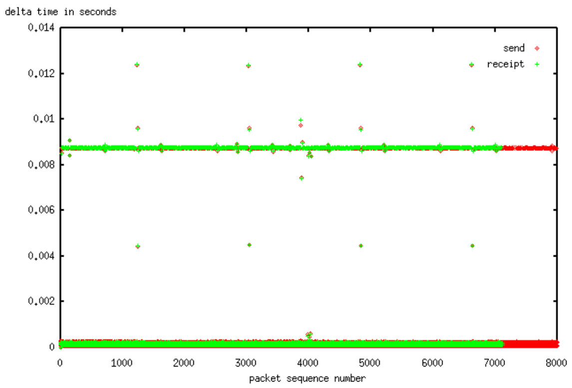

Figure 7: 7.2Mbit/s `mgen` generated stream as received over a shortcut

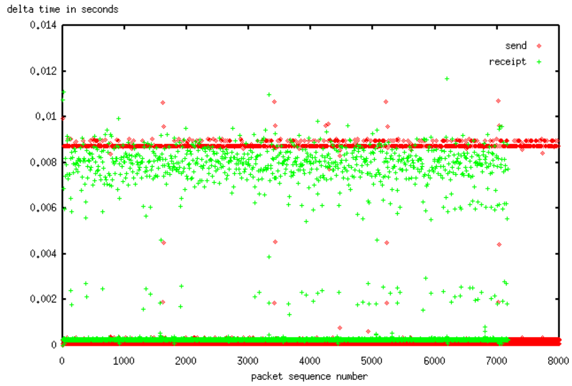

Figure 8: 7.2Mbit/s `mgen` generated stream over Poisson loaded routers

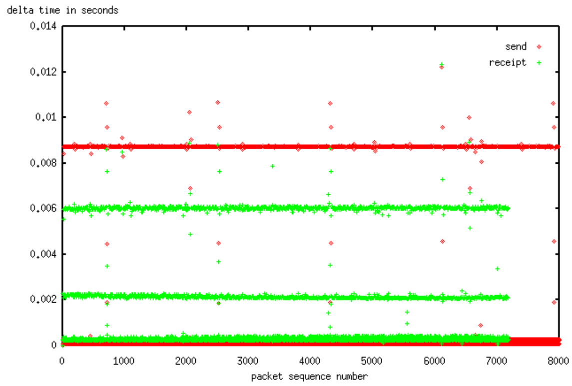

Figure 9: 7.2Mbit/s `mgen` generated stream over periodic loaded routers

differs considerably from the case where the Poisson distributed load was used. Since the testing was done when the network was almost free from other traffic, one possible explanation is some form of flow detection or buffering mechanism in one or both of the routers. This kind of buffering might be advantageous with non-delay sensitive applications but disadvantageous with applications that require constant delay from the network.

A closer examination also showed that some packets were lost or reordered during the transport over the routed path. As a result, some packets arrived before or after the main flow and some packets were completely lost. The number of lost or reordered packets was quite low, with the highest number being 226 lost packets of about 7200 sent packets. None of the packets sent over MPOA shortcuts were lost or reordered.

6.2 Delay comparison between LANE and MPOA

The tests verify that routing can introduce non-predictable behavior which can cause problems for applications that are delay sensitive or demand that packets arrive in correct sequence. The direct shortcut VCCs ensure that the delay is as low as the ATM network can provide and the higher layer packets will arrive in the correct order.

With all the routed LANE streams there were a number of anomalities introduced by the routers. Packets were sometimes lost and received out of sequence. The small variance and mean jitter experienced by shortcut streams indicate that the packets were received with almost constant delay. The constant delay makes MPOA advantageous for jitter sensitive applications.

7 Conclusions

The project was successful and multiple results were achieved. The ATM on Linux distribution was expanded to include MPOA client support and a number of different tests were done with TUT's experimental MPOA network.

Even though the positive results from interoperability, performance and jitter testing demonstrate the benefits of MPOA, it is uncertain how much MPOA will be deployed in production networks. MPOA based products from major vendors have been out for a while now (summer 1999), so the next year will probably show how well MPOA is accepted by the networking professionals.

References

[Kiiskilä] Marko Kiiskilä. *Implementation of LAN Emulation Over ATM in Linux.* Master's Thesis. Tampere University of Technology. Tampere, Finland October 1996. `ftp://sunsite.tut.fi/pub/Local/linux-atm/misc/`

[Linux-ATM] Werner Almesberger. *ATM on Linux homepage.* `http://ica1www.epfl.ch/linux-atm/`

[LUNIv2] ATM Forum Technical Committee. *Lan Emulation Over ATM Version 2 – LUNI Specification.* July 1997.

[MGEN] Brian Adamson. *The MGEN Toolset* `ftp://manimac.itd.nrl.navy.mil/Pub/MGEN/dist/`

[MPOA] ATM Forum Technical Committee. *Multi-Protocol Over ATM – Version 1.0* July 1997.

[Netperf] Rick Jones. *The Public Netperf Homepage.* `http://www.netperf.org/`

[RFC 2332] J. Luciani, D. Katz, D. Piscitello, B. Cole, N. Doraswamy. *NBMA Next Hop Resolution Protocol (NHRP)* April 1998.