

IP over ATM: Classical IP, NHRP, LANE, MPOA, PAR and I-PNNI

[Jun Xu](mailto:jun@cis.ohio-state.edu)<jun@cis.ohio-state.edu>

Due to the popularity and dominance of the TCP/IP, the success of ATM as a data networking technology lies largely in its support of IP on top of it. This report surveys the proposed approaches of transporing IP traffic over ATM networks, and examines their pros and cons.

[Other Reports on Recent Advances in Networking](#)

[Back to Raj Jain's Home Page](#)

Table of Contents

- [1. Introduction](#)
 - [2. Classical IP over ATM, NHRP, and IP Multicast over ATM](#)
 - [3. LAN Emulation](#)
 - [4. MPOA](#)
 - [5. PAR and I-PNNI](#)
 - [6. Summary](#)
 - [7. Annotated Bibliography](#)
-

[Other networking topics ...](#)

1. Introduction

1.1 The challenges of IP over ATM

The success of Asynchronous Transfer Mode (ATM) lies largely in its ability to transport legacy data traffic, mostly IP, over its network infrastructure. The complexity of interoperating IP with ATM originates from following two major differences between them.

- Connection-oriented vs. Connectionless

ATM is connection-oriented, that is, a connection need to established between two parties before they

can send data to each other. Once the connection is set up, all data between them is sent along the connection path. On the contrary, IP is connectionless so that no connection is needed and each IP packet is forwarded by routers independently on a hop-by-hop basis. When we need to transport IP traffic over an ATM network, we have two options. Either a new connection is established on demand between two parties or the data is forwarded through preconfigured connection or connections. With the first approach, when the amount of data to be transferred is small, the expensive cost of setting up and tearing down a connection is not justified. On the other hand, with the second approach the preconfigured path(s) may not be an optimal path and may become overwhelmed by the amount of data being transferred.

- QoS-aware vs. Best Effort

Quality of Service is an important concept in ATM networks. It includes the parameters like the bandwidth and delay requirements of a connection. Such requirements are included in the signaling messages used to establish a connection. Current IP (IPv4) has no such concepts and each packet is forwarded on a best effort basis by the routers. To take advantage of the QoS guarantees of the ATM networks, the IP protocol need to be modified to include that information.

1.2 Models of IP over ATM

- Peer Model

To run IP on top of ATM networks, we first need to figure out how to relate ATM protocol layers to TCP/IP protocol layers. Two models, one called peer model and the other overlay model, are proposed [Models]. Peer model considers the ATM layer a peer networking layer as IP and propose the use of the same addressing scheme as IP for ATM-attached end systems. ATM signaling requests will contain IP addresses and the intermediate switches will route the requests using existing routing protocols like OSPF. This scheme was rejected because although it simplifies the addressing scheme for end systems, it complicates the design of ATM switches by requiring them to have all the functions of an IP router. Moreover, if the ATM network will also support other networking layer protocols like IPX or Appletalk, the switch has to understand all their routing protocols.

- Overlay Model

The overlay model, which is finally adopted, views ATM as a data link layer protocol on top of which IP runs. In overlay model ATM networks will have its own addressing scheme and routing protocols. The ATM address space is not logically coupled with the IP addressing space and there will be no arithmetic mapping between the two. Each end system will typically has an ATM address and an unrelated IP address well. Since there is no nature mapping between the two addresses, the only way to figure out one from the other is through some addressing resolution protocol.

With overlay model, there are essentially two ways to run IP over ATM. One treats ATM as a LAN and partitions an ATM network into several logical subnets consisting of end systems with the same IP prefix. This is known as Classical IP over ATM. In Classical IP over ATM, end systems in the same logical subnet communicate with each other through end-to-end ATM connections, and like in LAN, ARP servers are used in logical subnets to resolve the IP addresses into ATM addresses. However, traffic between end systems in different logical subnets has to go through a router even though they are attached to the same ATM network. This is not desirable since routers introduce a high latency and become the bandwidth bottleneck. Next Hop Resolution Protocol (NHRP) steps in to solve this problem. Working in an ATM network partitioned into logical subnets, it allows an end system in one subnet to resolve the

ATM address (from the IP address) of an end system in another logical subnet and establish an end-to-end ATM connection, called a short-cut, between them.

The other approach uses an ATM network to simulate popular LAN protocols like Ethernet or token ring. IP runs on top of it in the same way it runs on top of Ethernet or token ring. This is known as LAN Emulation (LANE). LANE allows current IP applications run over an ATM network without modification. This will help accelerate the deployment of ATM networks. However, like in Classical IP over ATM, traffic between different emulated LANs (ELANs) still needs to travel through a router. As a combination of LANE and NHRP, Multiprotocol Over ATM (MPOA) solves the problem by creating shortcuts that bypasses routers between ELANs.

- PAR and I-PNNI

With above approaches, ATM and IP each run a separate routing protocol. For ATM it is P-NNI and for IP it is OSPF. With IP, the routers have no idea about the internal topology of the ATM network, and with ATM, the switches does not distinguish between an ATM-attached router and an ATM end system. Sometimes it is desirable for the routers to understand the routing protocols of ATM to figure out how to establish end-to-end ATM connections with other routers. This results in PNNI augmented routing (PAR), in which ATM-attached routers behave like an ATM switch and exchange topology and reachability information with switches and other routers. Another approach, called Integrated PNNI (I-PNNI), propose the use of PNNI as the single protocol to be used in a network of switches and routers. [Dorling]

1.3 IPv6 and IP Integrated Services over ATM

IP itself is an evolving technology; IPv6 improves the addressing capabilities of IPv4 and IP integrated services support real-time IP traffic. There are many issues concerning how ATM networks support both technologies. However, since each has a report (see [Other networking topics ...](#)) devoted to it, I will not duplicate the effort here to discuss them in detail.

2. Classical IP over ATM, NHRP, and IP Multicast over ATM

2.1 Classical IP over ATM

2.1.1 Logical Subnet

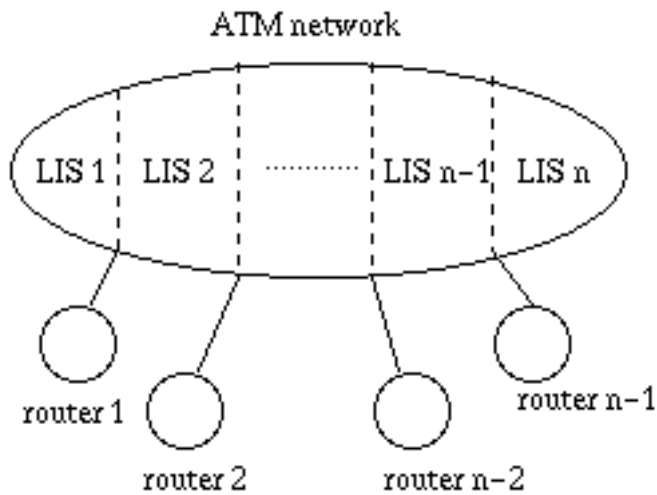


Fig. 1 Classical IP over ATM

Fig. 1 shows the configuration of Classical IP over ATM. As the name indicates, this model treats the ATM network as a number of separate IP subnets connected through routers. Such an IP subnet is called a logical subnet (LIS). A LIS has the following properties [RFC 1577]:

- End systems in a LIS share the same IP prefix and address mask.

In this way LIS is much similar to a traditional IP subnetwork over a broadcast LAN. However, traditional IP subnetworks are separated from each other by routers while LISs are actually connected to the same ATM network. This explains why it is called logical subnet: the membership to an LIS is defined by software configuration, not by hardware settings. Also this implies that inter-LIS communication need not necessarily go through a router.

- End systems in a LIS communicates with each other through end-to-end ATM connections.

When an end system A needs to communicate with an end system B, which is in the same LIS, it needs to establish a connection with B first. A has B's IP address but does not know its ATM address. To resolve the IP addresses into ATM addresses, like in traditional IP subnets each LIS contains an ARP server, called ATMARP server. A sends ARP query packet that contains B's IP address to the ATMARP server and the server will reply to it with B's ATM address. A then establishes a connection with B through P-NNI signaling.

- End systems in different LISs communicates with each other through a router.

Like in traditional IP subnets, a router is a member of multiple LISs and forwards IP traffic between them. Typically, each LIS contains a router and all IP packets that is not destined for an end system in the same LIS are forwarded to the router. If the router is in the same LIS as the destination end system, it forwards the packet to the destination end system using the scheme described above (intra-LIS). Otherwise it forwards the packet to another router and the packet is routed to the destination on a hop-by-hop basis. For example, in Fig. 1, router i is a member of both LIS i and LIS $i+1$ ($i=1, 2, \dots, n-1$). Then a packet from an end system in LIS 1 to an end system in LIS 2 will travel through router 1, router 2, ..., and router $n-1$. This is not desired since each router has to reassemble and disassemble the IP packet and this introduces a huge amount of delay. Since it is feasible to have direct connection between these two end systems as they are attached to the same ATM network, such hop-by-hop forwarding is definitely a waste of time and resource. NHRP fixes this problem by allowing direct

connection between the end systems that lie in different LISs.

2.1.2 ATMARP Server

The configuration and operation of an ATMARP server in a LIS is similar to an ARP server in a traditional IP subnet. An ARP_REQUEST packet is sent from an end system to the ATMARP server to resolve the ATM address of an IP address. ATMARP server contains a table consisting of <IP Address, ATM address> pairs. If a pair that matches the IP address is found, the corresponding ATM address will be returned to the inquiring end system in a ARP_REPLY packet. Otherwise, an ARP_NAK packet is returned. The table entries are either configured manually or learned through the registration of the end systems with the ATMARP server. Detailed description of LIS and ATMARP can be found in RFC1577.

2.1.2 Data Encapsulation [RFC 1483]

Data encapsulation for IP over ATM is specified in RFC 1483. AAL 5 is used to carry the IP packets end-to-end. Two modes of encapsulations, one called VC Based multiplexing and the other LLC encapsulation, are defined. The difference between the two is that the latter allows multiple protocols to be carried on a single VC while the former binds a VC to a single protocol.

- VC Based Multiplexing

With VC Based Multiplexing, the VC is bound to a specific protocol and the data is encapsulated into the CPCS-PDU field of AAL5 directly. So each connection can only carry one protocol. This approach is preferred when large number of VCs can be established in a fast and economical way.

- LLC Encapsulation

With LLC Encapsulation, a number of protocols (e.g., IP, IPX, AppleTalk.) can be carried over the same VC connection. In this approach, an IP packet will be prefixed with an IEEE 802.2 LLC header before it is encapsulated into the AAL5 frame. This approach is preferred when a separate VC for each carried protocol is either expensive or not possible, e.g., the only connection type supported is PVC or the charging is based on number of VC allocated. It is used as the default encapsulation method for all IP over ATM protocols.

2.2 NHRP

As mentioned before, with Classical IP over ATM, inter-LIS communication has to go through routers, which is not an optimal solution when both parties involved are attached to the same ATM network. Direct connection between them is desired and actually it is not difficult to achieve. All we need is a mechanism for an end system to resolve the IP address of another end system in a foreign LIS into its corresponding ATM address. NHRP is such a protocol that performs this task.

NHRP consists of two types of entities, NHRP server (NHSs) and NHRP client(NHC), and protocols between them. Each LIS contains at least one NHRP server and each end system is a NHRP client. When an end system needs to resolve an IP address, it sends a request to NHRP server in charge of his LIS. A NHS can serve more than one LISs and it keeps a table of <IP address, ATM address> pairs for all the hosts that belong to the LISs it is serving. If the IP address to be queried belongs to these LISs (known from the IP prefix), NHS expects to find an entry that matches the IP address and replies with the corresponding ATM address. Otherwise a negative reply will be returned.

So far NHS behaves exactly the same as ATMARP servers, and actually in LISs where NHC and ATMARP clients coexists, NHSs are coupled with the function of ATMARP server. However, a limitation of ATMARP server is that it can not resolve an IP address that belongs to another LIS while NHRP server can. When a query comes to a NHRP server regarding an IP address that belongs a LIS it does not serve, it will manage to forward the query to the NHS that serves the LIS. NHSs that serve LISs on an ATM network have preconfigured connections between them so that they form a routed network for NHRP queries. Thanks to the routing protocols like OSPF executed among these NHSs, like IP routers NHSs know which next-hop (another NHS) to forward the NHRP query in order to reach the destination NHS. This is exactly where the name of Next Hop Resolution Protocol (NHRP) is coming from. When the NHS that serves the LIS receives the query, it will reply with the corresponding ATM address to the end system that initiates the query. The reply will travel back through the NHSs and intermediate NHSs may cache the <IP address, ATM address> entry so that later NHRP queries with the same IP address will be intercepted and replied. This feature is expected to save a lot of NHRP traffic. Once the sender knows the ATM address of the receiver, it can establish an end-to-end connection with the receiver, called a shortcut, to transfer IP packets between them. Before the shortcut is established, the data will still be forward through the routers like in Classical IP over ATM.

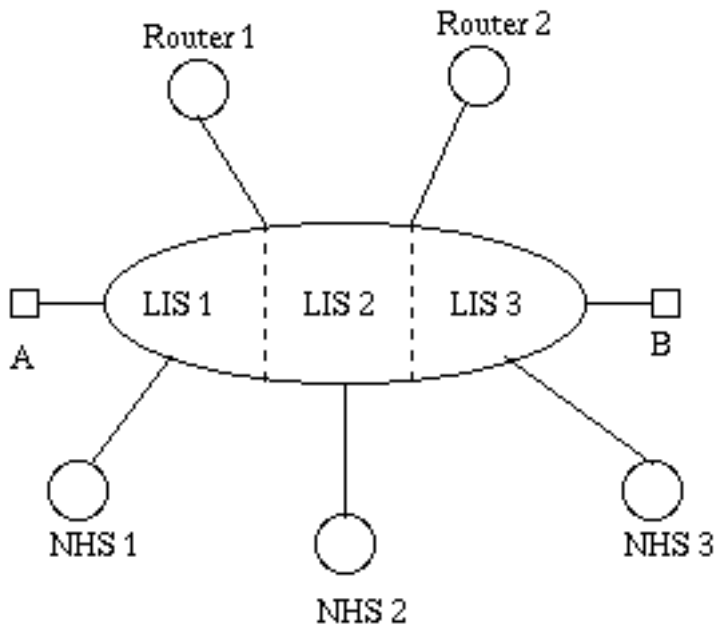


Fig. 2 NHRP

Fig. 2 shows an ATM network that is partitioned into three LISs, each of which is served by a NHS. Router 1 is connected to both LIS 1 and LIS 2, and router 2 is connected to both LIS 2 and LIS 3. End systems A and B are attached to LIS 1 and LIS 3 respectively. Now A wants to send some data to B. With Classical IP over ATM, the data will have to travel through router 1 and router 2. With NHRP, A sends a NHRP requests to NHS 1, which forwards it to NHS 2, and then to NHS 3. Since NHS 3 is serving the LIS 3, which contains the end system B, it will reply to A with the ATM address of B. When the reply travels back, it may go through NHS 2 and NHS 1. NHS 1 and NHS 2 will cache this information so that future NHRP requests will be replied by them directly without forwarding it to the serving NHS. When A gets the reply, it will establish an end-to-end connection with B to transfer the data between them.

2.3 IP Multicast over ATM [RFC 2022]

Classical IP over ATM and NHRP only support IP unicast over ATM. To support IP multicast, two issues need to be solved. First, we need an address resolution protocol to translate a multicast IP address into a list of ATM addresses, and this is solved by Multicast Address Resolution Server (MARS). Second, we need to specify how multicast data is transferred among the involved parties, and VC mesh and Multicast Server (MCS) are two possible solutions.

A Multicast Address Resolution Server (MARS) is introduced into each LIS to perform the multicast address resolution. It answers the queries for multicast addresses from the end systems in the same way as ATMARP server answers queries for unicast addresses. An end system joins or leaves a particular multicast group by sending Internet Group Multicast Protocol (IGMP) packets to the MARS.

When an multicast IP address is resolved into a list of end points, the data needs to be forwarded among the group members. from the sender to the recivers. One way to do this is to let each group member set up a point-to-multipoint connection with all other group members and this approach is called VC Mesh. The other way is to introduce a Multicst Server (MCS) into each LIS that supports multicast. When an end system query for a multicast address, MARS will reply to it with the ATM address of the MCS. The end system then sends the multicast packets to MCS. The MCS will build a point-to-multipoint connection or multiple point-to-point connection to the group members to forward the packet received from the end system to all the members of the group specified in the address field of the multicast packet.

VC mesh and MCS each its pros and cons. With MCS, if the membership of a multicast group changes, it only needs to modify the point-to-multipoint VC to the group members while with VC mesh, all connections in this "mesh" have to be modified. However, MCS needs to reassemble the cellified packets sent from the source and resend them to the group members so it may become the single point of conjection and introduce certain amount of latency. With VC mesh the reassembling is not needed so the latency is minimized.

3. LAN Emulation (LANE)

In LAN Emulation (LANE), an ATM network is configured to simulate an Ethernet or Token Ring LAN. The resulting LAN is called an Emulated LAN (ELAN). LANE defines a service interface that looks to the IP exactly the same as Ethernet or token ring. In this way the IP software that is running previously on Ethernet and token ring can be ported onto the ATM network without any modification. This helps accelerate the deployment of ATM as a LAN technology. [Ginsburg]

LANE specifies following four types of entities and protocols between them:

- **LAN Emulation Client (LEC):** An LEC runs on each ATM end system in an ELAN to simulate an Ethernet or Token Ring node. Each LEC has one or more MAC addresses associated with it. It contacts the LES (explained below) to resolve the MAC addresses into ATM addresses and performs certain control functions. It emulates Ethernet or Token Ring service interface to the IP layer by encapsulating the outgoing IP packets into ELAN frames or decapsulating the incoming ELAN frames into IP packets.
- **LAN Emulation Server (LES):** Each ELAN contains a LES, which acts as the coordinator. Each LEC will register with LES its <MAC address, ATM address> pair. Based on such information LES resolves

MAC addresses into corresponding ATM addresses in the same as ATMARP server in Classical IP over ATM.

- **Broadcast and Unknown Server (BUS):** Each ELAN includes a Broadcast and Unknown Server (BUS) to emulate the broadcast capability of Ethernet and Token Ring. A LEC who wants to broadcast a packet sends it to BUS, which forwards every ELAN member a copy. Before the direct data connection is built between two LECs, the data between them is also forwarded through BUS.
- **LAN Emulation Configuration Server (LECS):** There can be more than one ELAN running on an ATM network. LECS keeps the configuration information of each ELAN including the LECs, LES and BUS in each ELAN.

The operation of LANE consists of following four steps:

- **Configuration:** In this step, a LEC contacts the LECS to know each ELAN it belongs to and the address of LES and BUS to contact to join that ELAN. There are three ways for a LEC to access LECS. The first is to configure the ATM address of the LECS into the LEC. The second is to have a fixed VPI/VCI that directs to the LECS from every end system. The last is to get it through ILMI.
- **Registration:** Now that LEC knows the ATM address of the LES of the ELAN to join, it sends a registration message to the LES, which includes the MAC address and ATM address of the LEC. Upon receiving this message, the LES will record the information in its address resolution table and create a few connections to the LEC for the transfer of data and control information.
- **BUS Connection:** Using the ATM address of BUS obtained in the configuration stage, it establish a connection to BUS for the transfer of multicast data
- **Data Transfer:** This step is the major business of LANE. When a LEC wants to send some data to another LEC, it will query LES for the ATM address of that LEC. Using that address it establish a direct data connection with that LEC. Afterwards Ethernet or Token Ring frames between them are transported on this connection in AAL 5 frames with LLC encapsulation [RFC 1483]. If a LEC needs to broadcast a packet, it simply sends the packet to the BUS, which will forward every ELAN member a copy of the packet.

Although Intra-ELAN data are transferred through direct data connection between two LECs. LECs that belong to different ELANs may still have to communicate through routers. Described below, MPOA solves this problem by creating shortcuts between ELANs.

4. MPOA

MPOA is basically a combination of LANE and NHRP. MPOA improves LANE by allowing inter-ELAN traffic go through shortcut connections rather than routers. To build such a shortcut, NHRP is used to resolve destination IP address into ATM address. In this sense, MPOA is a combination of Layer 3 routing and Layer 2 bridging. [Ginsburg]

4.1 MPOA Environment

A typical MPOA environment is an ATM network with ATM hosts and edge devices attached to it. An edge device can be a bridge, which bridges a legacy LAN (e.g., Ethernet or Token Ring) to an ELAN, or a router, which connects a non-ATM IP subnet to the ATM network. On this ATM network defined a number of ELANs and LISs. Typically a LIS corresponds to an ELAN and legacy LANs that are bridged to it. The job of MPOA protocol is to find the best way for any two hosts in this environment to communicate efficiently with each other.

If one legacy system attached to an edge device needs to send data to another legacy system attached to a different edge device, obviously the best approach is to establish a direct connection between the two edge devices (assuming that they are not connected in certain ways outside of the ATM network) and transport traffic across this connection. In this case, the edge device that the sender is attached to is called an ingress endpoint and the edge device the receiver is attached to is called an egress endpoint. The major business of MPOA is to build end-to-end connection between an ingress endpoint and an egress endpoint for efficient communication. If we view communication between two legacy systems attached to different edge devices as the that between these two edge devices, we can simplify our discussion in the following by only considering the communications between end systems (hosts and edge devices).

4.2 MPOA Operations

MPOA is built on top of LANE. If two end systems are in the same ELAN, the traffic between them follows the LANE specification. With LANE, inter-ELAN traffic has to go through a router. MPOA allows shortcuts between ELANs by integrating NHRP functionalities into it. It allows an end system to obtain the ATM address of another end system that lies in a different ELAN and establish end-to-end connection with it. Such shortcuts are created automatically through a mechanism called flow detection. Before the direct connection between the end systems is built, inter-ELAN traffic is still forwarded to routers. However, an entity called MPOA Client (MPC) that runs on each end system is capable of detecting traffic. When it finds out that an end system has recently sent a significant amount of traffic destined for the same destination to the router, it will try to establish a connection to the egress point which is "nearest" to the destination. An entity called MPOA server runs in a MPOA environment to resolve the ATM address of the egress point given the IP address of the destination.

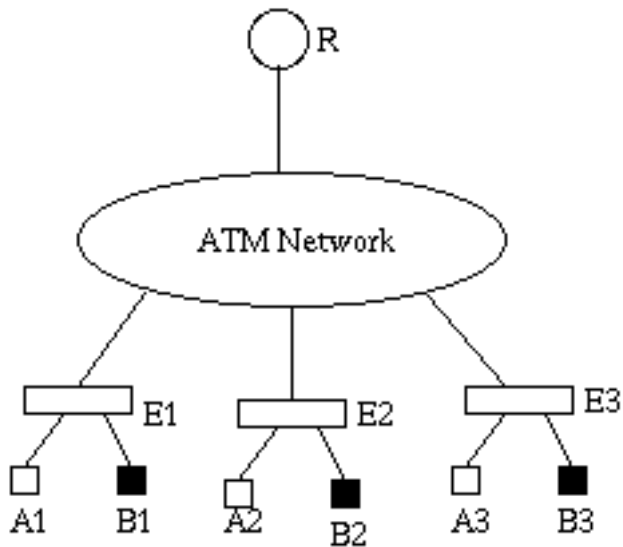


Fig. 3 MPOA

Let us look at an example. In Fig. 3 contains a MPOA environment. E1, E2 and E3 are edge devices connected to an ATM network. Each edge device connects two hosts to the ATM network (such edge devices are probably IP switches). Two ELANs are defined on the ATM network; LAN 1 consists of hosts A1, A2 and A3 and LAN 2 consists of hosts B1, B2 and B3. LIS 1 is defined on LAN 1 and LIS 2 is defined on LAN 2. The router R is a member of both LIS 1 and LIS 2.

If A1 wants to talk to A2, this is easily accomplished through LANE procedures as A1 and A2 belong to the same ELAN. However, if A1 wants to talk to B2, with LANE E1 has to send the packet from A1 to router R and R forwards it to E2, then to B2. This is obviously not a good approach. With MPOA, an entity called MPOA server (MPS) is running on R, which knows the edge device to contact for a given destination IP address. In this example, the MPS is running on R. If A1 wants to talk to B2, it first just sends the traffic to R as in LANE. However, E 1 is smart enough to detect this flow of traffic after a while. Then E1 will contact R to know which edge device is "nearest" to the destination B2. R then replies with the ATM address of E 2. Then E 1 builds a connection to E 2 and sends all traffic from A1 to B 2 along that connection.

In this example there is only one MPOA server. In real MPOA environment there can be multiple MPSs, each of which serves one or more ELANs. A MPS contains the function of a NHS, that is, if a MPS can not resolve an IP address it will forward the query to other MPSs.

MPOA supports the separation of route calculation and data forwarding. A traditional router has both functions. However, in MPOA environment if every edge device is a full-fledged router it will be too costly. MPOA removes the function of route calculation from the edge devices and a route server function is included in the MPS to direct these edge devices as to forward the data to correct destinations.

5. PAR and I-PNNI

In MPOA and NHRP, the connections between different routers and servers are preconfigured connections. When the network becomes large it is expected that such configuration can be complex and error-prone. PNNI augmented routing (PAR) tackles this problem by allowing the routers take part in the PNNI routing protocol and make connections to other routers on demand based on the reachability information it gets from the exchange of PNNI messages with other switches and routers. With PAR, the routers will still run the legacy IP routing protocols like OSPF and IGRP. The information it gets by taking part in the PNNI message exchange helps it establish connections with other routers. [Dorling]

I-PNNI proposes the use of PNNI as the only routing protocol in an environment of ATM network and ATM-attached routers. A peer group may include both routers and switches. Both signaling requests and IP packets are routed using the PNNI protocol. In this way, IP can benefit from the QoS guarantee and the scalability of the PNNI. In I-PNNI, both routers and switches will exchange PTSPs with their neighbours. Both IP addresses and ATM NSAP addresses can be advertised in PTSP packets. ATM switches will advertise their reachability to certain ATM addresses summarized as a prefix of 20-byte ATM NSAP address while IP routers advertise reachability to certain IP addresses summarized as a prefix of 16 bit IP address.

6. Summary

The ability to support legacy IP traffic is vital to the success of ATM networks. Treating ATM as yet another LAN technology, Classical IP over ATM is the simplest to implement. However, the drawback of it is that inter-LIS traffic has to travel through a router even though both parties are directly connected to the ATM network. NHRP fixes this problem by augmenting it with an address resolution protocol so that shortcut connections can be established between end systems that belong to different LISs. To accelerate the deployment of ATM technology, LANE emulates Ethernet and Token Ring LANs on an ATM network so that existing IP software running on such LANs can run on ELANs without modification. However, ELAN suffers the same drawback as Classical IP over ATM, that is, inter-ELAN traffic has to travel through a router. MPOA combines LANE and NHRP technology to support both IP routing and LAN bridging over an ATM network. In addition to these data models, two routing schemes, PAR and I-PNNI are proposed to be used in the environment of interconnected ATM network and routers. PAR allows the routers automatically discover each other and build ATM connections to exchange routing information. I-PNNI allows PNNI protocol to be used in a hybrid mesh consisting of ATM switch and routers. Actually these routing enhancements can be used in align with any data model.

7. Annotated Bibliography

[Cisco] A. Alles, "[ATM Internetworking](#)", Internet posting, Cisco Systems, Inc., May 1995.

[Ginsburg] D. Ginsburg, ATM: solutions for enterprise internetworking, Addison Wesley, 1996.

[Dorling] B. Dorling, "Internetworking over ATAM: An Introduction, Prentice Hall, 1996.

[Minoli] D. Minoli, et al., LAN, ATM, and LAN emulation technologies, Artech House Inc., 1996.

[MPOA] ATM Forum Multi-Protocol Over ATM Version 1.0.

[LANE] ATM Forum LAN Emulation 1.0.

[NHRP] IETF NBMA Next Hop Resolution Protocol (NHRP) Draft.

[I-PNNI1] Integrated PNNI (I-PNNI) v1.0 Specification.

[I-PNNI2] Issues and Approaches for Integrated PNNI.

[Models] ATM Forum/94-1015, On Models of Internetworking.

[RFC 1483] Multiprotocol Encapsulation over ATM Adaptation Layer 5.

[RFC 1577] Classical IP and ARP over ATM.

[RFC 2022] Support for Multicast over UNI 3.0/3.1 based ATM Networks.

[Other networking topics ...](#)

Last updated by: [Jun Xu](#) <jun@cis.ohio-state.edu>

August 14th, 1997