A Retrospective View of ATM

Charles Kalmanek
AT&T Labs Research
180 Park Avenue, Room A113
Florham Park, NJ 07932
crk@research.att.com

ABSTRACT

ATM was the focus of active research and significant investment in the early to mid 1990's. This paper discusses several visions for ATM prevalent at the time, and analyzes how ATM evolved during this period. The paper also considers the implications of this history for current connection-oriented technologies, such as optical transport networks and MPLS.

Keywords

ATM, transport networks, flow switching, MPLS.

1. INTRODUCTION

Asynchronous Transfer Mode (ATM) networking had its origins as a switching and multiplexing technology suitable for the design of high capacity switches. The essential features of ATM are a fixed-length packet (called a *cell*), which is switched based on a virtual circuit identifier in the cell header. End-hosts request that the network set up a virtual circuit via a signaling (control) protocol that allows them to specify the desired quality of service. Quality of service per virtual circuit is provided through admission control and switch scheduling algorithms, allowing delay-constrained traffic, such as voice and circuit-emulated TDM traffic, to share a single network infrastructure with bursty data traffic. The cell size was kept small to support low delay for voice (although introducing enough delay that echo cancellation is needed.)

For a period of time in the early to mid 1990's, investment and research on ATM exploded, based on an expectation that ATM would revolutionize networking. For telecom providers, ATM promised to unify a number of disparate networks (voice, private line, data) on a single switching network. The fixed cell size fit well with designs for large self-routing switch fabrics suitable for the construction of very high-capacity switches. ATM's proponents anticipated that ATM would be ubiquitous, and that end-to-end quality of service would enable an entirely new class of network applications to be built.

The reality today is far different. ATM is used today to provide Virtual Private Network (VPN) services to businesses, consisting primarily of point-to-point virtual circuits connecting customer sites. ATM services represented a $ 2B business in 2001. ATM also provides the underpinnings of Digital Subscriber Loop (DSL) services, which are growing rapidly. In DSL access networks, ATM enables local exchange carriers to switch subscriber traffic to different Internet Service Providers. ATM is also used as the core network infrastructure for large Frame Relay networks and for some IP networks. While these uses of ATM are important and should be viewed as a mark of success for ATM technology, there is a perception in the network research community that ATM "failed." Indeed, when compared with the grandiose visions that many of its proponents had, ATM was not as successful as it might have been. This paper explores some of the visions for ATM that were pursued both by telecommunications service providers and the research community in the early to mid 1990's and presents some of the technical and business issues that drove the evolution of ATM.

2. MANY VISIONS FOR ATM

Starting from a small set of initial design principles, the development of ATM technology progressed in a number of different directions, based on the business and technical visions of the companies and individuals who were driving the technology. This section summarizes several of the early visions for ATM prevalent in the research community and industry. While attempting to give a balanced overview of the work that was going on, I have no illusions that this survey is exhaustive or that is does justice to any one of these visions. However, I hope that it gives some notion of the tremendous scope that the ATM community was trying to address.

One vision of ATM's role in telecom networks was that it would provide a *single multi-service network*. I distinguish between two variants of this vision. One is that it would serve as a multi-service core network supporting primarily data services such as Frame Relay, IP, and ATM service, possibly with some DS1, DS3 or higher rate private line and voice services. The other is that it would eventually replace the circuit-switched TDM hierarchy and provide a *next-generation transport network*, supporting long-term capacity (bandwidth) management for all services. In the first instance, an ATM core network would support multiple service edges, such as frame relay, IP, ATM, and possibly private line and voice. A high level view of this architecture is shown in Figure 1, which gives an example of multiple networks, each with its own "edge" switches connected over dedicated links. Figure 2 utilizes a single core network to connect edge switches for each service. This approach optimizes link utilization through statistical multiplexing of edge-to-edge traffic over the core network. It was also understood that the introduction of hierarchy could improve the scalability of the network from a routing perspective. Rather than scaling independent networks as the number of customers grew, the networks within each region could be scaled independently, and interconnected over a core network running the ATM PNNI routing and signaling protocol. The edge-core approach also had the potential to reduce network operations expense through the consolidation of individual service networks.

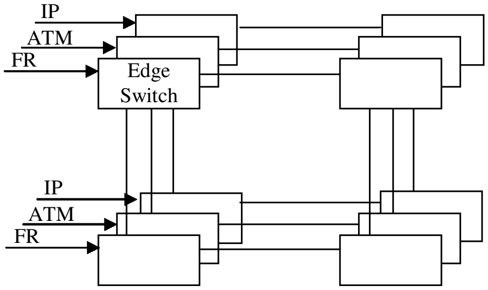

Figure 1: Multiple Edge Networks

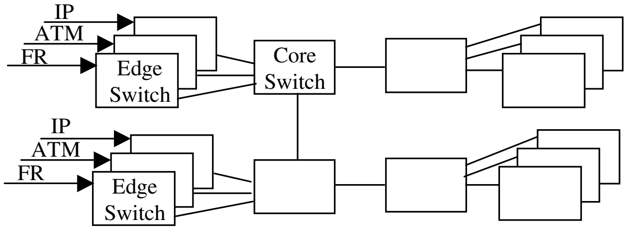

Figure 2: Edge-Core Architecture

To consider ATM as a next-generation transport network, we first need to understand how transport networks are built. Transport networks have traditionally been based on a hierarchy of time-division multiplex (TDM) switches. Figure 3 presents typical network architecture, simplified from [1]. DS1 private line or voice trunks are aggregated and/or switched in a Digital Cross-connect System with DS3 interfaces and a switching granularity of DS1, denoted as a DCS-3/1. DS3 private line or aggregated traffic demands are switched in a DCS-3/3 supporting DS3 interfaces and a switching granularity of DS3. These DS3s and other higher rate signals are carried over SONET rings[1] or linear chains that provide restoration in case of link or interface failure. SONET links are transported between central offices by an Optical Transmission System (OTS). Due to the existence of multiple layers of network hierarchy, this network architecture often results in inefficient overall network utilization, for a number of reasons. One is that the partial utilization at each network layer is compounded as you go up the hierarchy. Another reason is that circuits at a given layer may end up being routed inefficiently as a result of capacity planning processes that are designed to maximize the utilization of the layer that carries them. For example, a DS1 circuit may not follow the shortest path between two central offices because it is routed over a pre-existing DS3 circuit following a less efficient path.

From its earliest design, ATM was intended to support virtual circuits across a wide range of rates. With the promise of ATM switches with aggregate capacities of 10's of Gbps in the mid 1990's growing to 100's of Gbps by the end of the decade, network designers began to seriously consider using ATM in the transport layer, supporting multiple services and consolidating several layers of the TDM hierarchy. Figure 4 illustrates a network architecture in which all transport bandwidth management below OC-48 is done at the ATM layer. In the figure, ATM and IP layer demands are carried over an optical cross-connect (OXC) layer, which takes on the restoration function supported in the SONET layer in Figure 3. Since bandwidth management on the ATM network occurs on a slow time scale, ATM VC setup only occurs on provisioning time scales. By eliminating the hierarchical transport network architecture below OC-48, ATM promised to simplify the task of managing the transport network and to improve overall network efficiency.

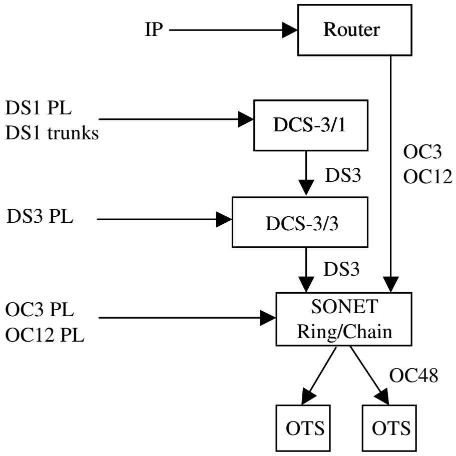

Figure 3: Traditional Transport Network

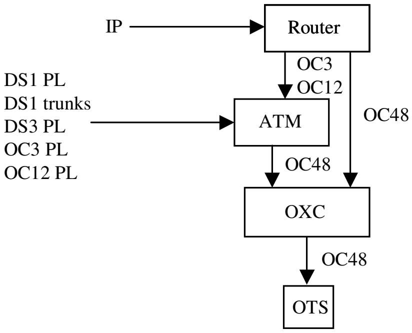

Figure 4: ATM-based Transport Network

A second vision of ATM, fostered by the research community, envisioned the use of ATM as a universal end-to-end packet service [2]. ATM's emphasis on end-to-end quality-of-service was an important part of this vision. In the end-to-end vision, desktop computers, networked appliances and large servers would all support ATM, and set up end-to-end connections with quality-of-service when needed. There was a tremendous amount of work on ATM host interfaces, low cost ATM interface chips, and ATM LAN switches. Protocol stacks were developed for end-hosts by extending the Berkeley socket layer to allow applications to directly establish ATM virtual circuits [3]. ATM's small cell size seemed particularly well suited to managing quality-of-service in access networks, where bandwidth is scarce, such as DSL and wireless networks [4], etc. Small cells implied that links could be

---

[1] While SONET ring technology does not scale very well to large networks, large SONET cross-connects were not yet available in the early to mid 1990's.

scheduled on a fine time scale, allowing delay-sensitive applications to be supported alongside elastic applications.

There was also a significant amount of research on ATM signaling. In addition to ATM standards, ATM signaling was adapted to support connection handoff for mobile wireless devices [5]. A lightweight ATM signaling protocol was proposed in [6] which established forwarding state for a connection very efficiently, while allowing additional signaling along the slow path to establish QoS for the connection. The “open signaling” community proposed that ATM switches should support a standard low-level control protocol [25], allowing service providers to customize their ATM control plane. The research community also pursued even more radical ideas, such as desk-area networks [7], in which ATM was used as the fabric in a desk-area distributed computing environment comprising processing resources, I/O devices, and storage. As a result of the investment by the vendor community, leading edge products were available causing some enterprises to start to use ATM in their server environment, and to gear up to push ATM to the desktop.

The above vision of ATM was of interest to ATM purists. At the same time, there began to be significant interest in IP flow switching concepts [8] that promised to better integrate ATM with IP. If one looks at IP traffic, a significant fraction of the “flows” are small transactions, such as DNS lookups, for which the overhead of ATM connection setup is on par with or larger than the duration of the transaction. Rather than setting up a connection for short transactions, IP flow switching proposed to setup ATM *shortcut* connections only for long-lived flows [9, 10], offloading slow IP routers and leveraging high-performance ATM switches. The basic idea is illustrated in Figure 5, which shows “default” connectivity via ATM permanent virtual circuits (PVCs) using solid lines, and a shortcut connection that has been set up dynamically between two routers using a dashed line. The flow switching architecture proposed to enhance ATM switches by putting smart algorithms for detecting IP “flows” on ATM line cards. Service providers began to investigate how they could use ATM flow switching as a way of more tightly integrating the IP layer with the ATM layer than was envisioned in the architecture shown in Figure 2.

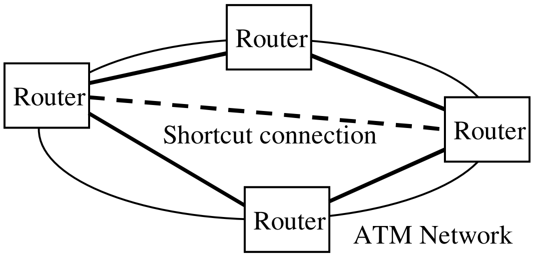

Figure 5: IP Flow Switching

The flow switching concepts were further developed in both the IETF and the ATM Forum. Using the Next-Hop Resolution Protocol (NHRP) [11], a router queries a Next-Hop Server to determine the ATM address of the next IP hop towards an IP destination. An alternative approach [10] utilized extensions to IP routing to carry the ATM address of the next-hop router. In addition, the Multicast Address Resolution Server (MARS) architecture [12] addressed the problem of mapping IP multicast forwarding onto ATM's connection-oriented services. MARS is based on point-to-multipoint VC's and uses either VC meshes or multicast servers to support the IP multicast service. A Multicast Address Resolution Server maintains a mapping from IP address to a set of ATM addresses in a Logical IP Subnet (LIS), and is updated by a host in the LIS when it joins or leaves an IP multicast group.

Another set of activities focused on enabling ATM to support the large embedded base of software built on bridged LANs. Initial implementations of LAN emulation on ATM were available around 1994, and the ATM Forum developed the LAN Emulation (LANE) specification [13]. LANE provides transparent support for Layer 2 Ethernet bridging services: the ATM LAN emulation protocol stack is shown in Figure 6. LANE's primary goal was to support common end-host drivers, such as Network Driver Interface Specification (NDIS) from Microsoft, which runs over bridged LANs. The function of the LAN emulation layer was to exactly mimic the MAC layer interface of a bridged LAN, so the higher layers would think they were running over a standard Ethernet or token ring network. An ATM LAN bridge would support fragmentation of MAC frames into ATM cells, emulating all of the functions supported by LAN bridges on top of ATM and inter-working with existing bridged networks.

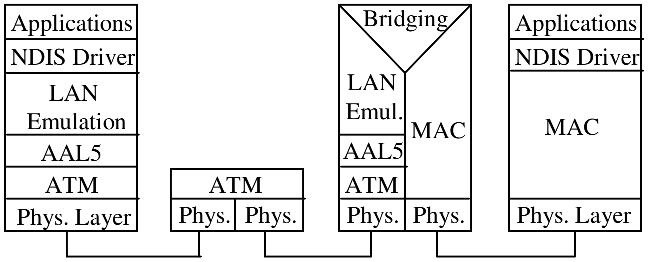

Figure 6: LAN Emulation

None of these visions of how ATM would evolve proved to be correct. The question is why?

3. HOW IT PLAYED OUT

The vision of end-to-end ATM with quality-of-service was extremely compelling to many people. Internet service was widely seen to be unpredictable, and ATM promised a solution. The ATM Forum developed a comprehensive Traffic Management framework, supporting five classes of service: CBR, VBR-rt, VBR-nrt, ABR and UBR[2]. This framework developed many of the key traffic management concepts that are in common use today: traffic descriptors, shapers, policers, priority and weighted fair scheduling, as well as signaling support for connection admission control and QoS routing. CBR defines mechanisms to support delay-constrained constant bit rate traffic, while ABR supports network feedback to sources allowing them to adapt their sending rate to the max-min fair rate of the bottleneck link along the path of the connection [14, 15]. This

---

[2] CBR = constant bit rate; VBR-rt = Variable Bit Rate (real time); VBR-nrt = Variable Bit Rate (non real time); ABR = Available Bit Rate; UBR = Unspecified Bit Rate.

work contributed a significant number of innovations to packet networking.

However, end-to-end ATM faced a deployment challenge due to the law of network externalities. Until ATM deployment reached a critical mass, end-to-end ATM QoS couldn't be realized. Since most existing applications were based on IP, use of ATM QoS for IP applications would need to be mediated by a (non-existent) IP QoS application programming interface. Moreover, ATM deployments would bear the burden of inter-working with existing applications and hosts that had not been upgraded to ATM. Another factor limiting the realization of end-to-end QoS is that, despite the variability of Internet performance, new network technologies are often initially deployed in enterprises. Here, the application drivers for end-to-end QoS never really materialized. And in this environment, high-speed Ethernet-based LANs began to dominate the desktop thereby reducing the perceived need for ATM's traffic management framework.

Nonetheless, the momentum behind ATM deployment was strong enough that for a period of time, it seemed possible that ATM to the desktop might succeed. Economic forces worked against it, however. In the early 1990's, an ATM host adaptor cost roughly \$3K. By the mid 1990's, this price had fallen to roughly \$1K. At that time, Ethernet adaptors cost about \$100. This price difference was a significant impediment to widespread adoption of ATM.

We can also consider the vision of ATM as a future core data or transport network. Since ATM provides a flexible bandwidth management capability, it seemed very well suited to the role of a multi-service core network or even core transport network, replacing SONET rings with a more efficient mesh structured bandwidth management layer. Unfortunately, there were a number of factors that made it difficult to realize this vision. One factor is that growth for data services through the mid 1990's dramatically exceeded expectations. This growth included Frame Relay, IP, ATM and private line services. When the total Frame Relay, ATM, IP and DS1 private line demands were considered, the switch capacities of commercially available ATM switches were not adequate to support the requirements of large central offices.

Another issue is that each service emphasized a slightly different set of requirements, such that the union of the service requirements was difficult for vendors and service providers to cope with. For example, private line has stringent requirements on reliability and restoration capabilities, and requires ATM line card support for TDM circuit emulation; voice has requirements for high switched virtual circuit (SVC) setup rates; and ATM PVC and SVC services for data likely require high port densities and reasonable reliability, at a significantly different target for price/performance than either voice or private line. If we focus specifically on transport network evolution, switch reliability was an issue. Digital Cross-connect Systems typically are designed to meet less than 3 minutes per year of downtime: relatively immature ATM switches could not meet this requirement. One way of dealing with some of these problems would have been to select different vendors for the edge and core switch nodes. However, due to relatively slow progress on ATM signaling standards, vendor inter-operability was delayed. As a result of these considerations, service providers gave up on the idea of a single core network, and took the more conservative approach of evolving separate IP, transport, and ATM networks. The one exception was Frame Relay networks, which evolved to run over ATM core networks since higher speed ATM interfaces and switch capacities were needed to support growing Frame Relay demand.

It is also important to recognize that ATM was being standardized and deployed just as IP was beginning to pick up steam. To be successful, ATM clearly needed to provide some inherent advantages (and few disadvantages) in carrying IP traffic. To some, ATM's high-speed and the emerging flow switching technologies seemed at the time like a winning combination. However, flow switching actually introduced uncertainty in how the IP layer would best utilize ATM. This uncertainty and the complexity of the underlying technical issues may have slowed rather than accelerated ATM deployments in large-scale data networks. Standardization in the IETF and ATM Forum's Multi-Protocol over ATM (MPOA) group [24] seemed likely to take a long time. In addition, flow switching introduced another layer of complexity into the architecture, requiring vendors to understand and be competitive in ATM switching, IP routing, *and* IP flow switching. To utilize flow switching effectively, service providers would need to provision IP flow detection algorithms on edge switches. Before developing software tools to do this, there needed to be strong evidence that flow switching would provide either significant cost savings or end-user performance improvements. Demonstrating cost savings in a large service provider network would depend on a combination of factors, including the ratio of router interface costs to ATM interface costs, the amount of traffic that would actually utilize shortcuts, etc. Demonstrating end-user performance improvements would depend on being able to deliver shortcut traffic with better end-to-end performance (e.g., throughput, delay) than routed traffic. While this might have been possible, the improvement would have to be significant to justify a new architecture.

Another issue with IP over ATM was support for IP multicast. The IP multicast community was extremely vocal about both the need for IP multicast and the difficulties of supporting it in ATM. As a result, ATM standards evolved to include support for point-to-multipoint unidirectional virtual circuits, and [12] defined a set of mechanisms to support IP layer multicast using them. An alternative multicast model using core-based trees was defined in [16]. Nonetheless, the debate about IP multicast over ATM contributed to the uncertainty about ATM, despite the fact that IP multicast has *still* not been widely deployed.

While these issues were being played out in the marketplace, there was also significant investment in IP router technology, due to the tremendous growth in IP traffic demands. Improvements in silicon technology and algorithms for IP forwarding table lookups [17, 18] resulted in a situation where commercially available ATM switches did not offer any speed advantage over commercially available IP routers. For a multiplexing layer to make sense, it needs to offer some speed advantage over the demands that it will be carrying. In fact, router vendors began to use ATM switch technology to increase their aggregate switching capacity, and interface rates on ATM switches often lagged behind those of IP routers.

It is also useful to consider the evolution of link technologies and their impact on ATM. Around 1992, when work on ATM was getting started, 100 Mbps shared FDDI rings were the fastest switching technology in widespread use. There was no other cost effective high-speed link technology for the LAN – Fast Ethernet was not yet available. ATM took advantage of the newly developed Fibre channel line coding chips, which promised relatively inexpensive 155 Mbps LAN links. In the WAN, ATM SONET interfaces at 155 Mbps and 622 Mbps were commonly used as the interface between backbone routers. Then, in 1994 - 1995 Fast Ethernet came out, and around 1997, Gigabit Ethernet appeared. In 1997, Packet over Sonet (POS) also appeared, supporting high-speed IP over HDLC over SONET without the overheads of ATM. When compared with IP router technology, ATM had a number of initial advantages: high capacity switching, high-speed links, etc., but lost advantage after advantage as IP router technologies advanced.

The end result was that ATM never reached critical mass for going to the desktop, given the cost of network adaptors, network externalities and the existence of competing Ethernet technology. ATM was also not ready to support the needs of a multi-service core network or core transport network. Finally, proposals to integrate ATM and IP, such as IP flow switching, were complex and were ultimately superceded by advances in IP router design, which incorporated many of the innovations that had been developed in ATM.

It is clear that ATM fell short of the technological vision that many people had. Where ATM *succeeded* was as a Layer 2 switching technology that is used in access/aggregation networks and as a core network for Frame Relay and native ATM services. Layer 2 VPNs consisting of point-to-point PVC's provide high reliability connectivity services to enterprises at a lower price point than private line services. ATM is also widely deployed as part of DSL access networks. The latter application is shown in Figure 7, where traffic from a DSL Access Multiplexer is carried over ATM to an ISP. It is important to note that ATM networks today carry a significant amount of Frame Relay traffic, DSL traffic, both enterprise and carrier voice traffic, and some backbone IP traffic.

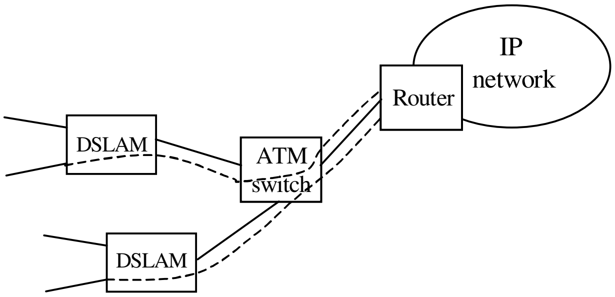

Figure 7: ATM-based DSL Access

ATM services continue to provide some technical advantages over IP services. ATM is more mature than IP in its ability to provide stringent quality of service guarantees. While IP differentiated services support class-based quality of service, IP differentiated services face a number of deployment challenges in large ISPs, including the difficulty getting good traffic data as an input to traffic engineering, and performance problems in legacy routers when quality of service features are enabled. In addition, ATM's guaranteed bandwidth on demand, fast re-route, and OAM features are important to many large customers. None of these features has been thoroughly integrated into IP as yet. Identifying and sectionalizing problems with end-to-end service in IP, even within a single ISP, is difficult. In time, MPLS may be able to support these features well, but ATM will be the technology of choice for some customers for many years to come.

4. LESSONS FOR TODAY

One natural question given this background is what lessons can be learned from ATM that might be applicable in the latest incarnations of connection-oriented technology: optical transport networks and MPLS. Large SONET cross-connects, typically with OC-48 or OC-192 ports and STS-1 switching granularity are now commercially available and are the basis of a new generation of transport infrastructure. These switches use variants of ATM's PNNI signaling or MPLS signaling (called Generalized MPLS) protocols [19] to set up and manage connections. For the foreseeable future, SONET technology appears to be the likely basis of the transport infrastructure, perhaps augmented by the use of transparent optical switching for managing large optical bit pipes at some point in the future.

For IP networks, MPLS is being touted as providing a routing and switching layer that can enable multiple types of traffic to share a common packet switched infrastructure. This sounds familiar, and it is worth understanding how MPLS is evolving in order to understand whether it will succeed where ATM failed. Like ATM, MPLS is a virtual circuit technology. In fact, MPLS has borrowed a number of the essential ideas of ATM: virtual circuit switching, fast re-route, and the notion of a single network infrastructure. However, MPLS does not attempt to solve an end-to-end problem, but rather focuses on a single administrative domain and is tightly integrated with the IP forwarding paradigm. To support quality of service, MPLS reuses IP differentiated services. MPLS does not support fast/dynamic connection setup like ATM SVC's.

One key application of MPLS is support for Layer 2 and Layer 3 virtual private networks (VPNs) on an MPLS label switched core network. As with ATM, the opportunity here is to reduce network operations expense through the consolidation of individual service networks. The IETF 'Martini" encapsulation [20] allows frame relay, ATM, Ethernet, and IP packets to be transported over an MPLS network. An MPLS tunnel between ingress and egress label-switched routers (LSRs) can carry packets associated with different services – the egress LSR uses the innermost label to distinguish the service to which packets are to be de-multiplexed. This approach is designed to support ATM permanent virtual circuits, although it does have some deficiencies. For example, the Martini encapsulations do not support service inter-working such as between Frame Relay and ATM. In addition, the related signaling extensions [21] were not designed to support switched virtual circuit setup. However, another approach to ATM over MPLS provisions an overlay ATM network on MPLS tunnels, and runs PNNI among the overlay network edges. This approach allows switched virtual circuits to be supported. In addition to supporting ATM, Frame Relay and Ethernet 'virtual circuits," MPLS protocols are also being developed to support Transparent (bridged) LAN service over MPLS [22].

Layer 3 VPNs are supported using extensions [23] of the IP Border Gateway Protocol (BGP) that provide support for private address spaces and virtual private IP networks on a shared MPLS core network. Given that many enterprises use private IP addressing and do not want mission-critical applications exposed to the Internet, the isolation provided by Layer 3 VPNs based on MPLS are likely to be important, and will compete with Layer 2 and IPsec-based VPN's. Note also that the ability to run virtual private IP networks on a shared core network allows large Internet service providers to resell IP backbone network capacity. Since the cost structure of an ISP depends on economies of scale, this drives down costs. The RFC2547 approach may eventually change the way that ISPs handle Internet routing. While IP forwarding and MPLS label switching currently co-exist in core network routers, in time IP forwarding tables could essentially disappear from core routers in an MPLS enabled network – Internet default-free routes would only exist as one of the virtual routing and forwarding tables in an MPLS provider edge router.

While it is impossible to predict the future, it is clear that MPLS has avoided a number of the pitfalls that plagued earlier visions for ATM. First, MPLS is not attempting to provide an end-to-end solution -- it is clearly targeted at service provider core networks. As a result, MPLS doesn't need to be ubiquitous to be successful. Second, MPLS protocols are an extension of existing IP protocols, while ATM's control plane evolved to be quite complex, including signaling, routing, MPOA, LANE, etc – all of which were needed *in addition to* IP protocols. This may simplify development and deployment. Third, MPLS is riding the same technology curve as IP routers, which suggests that switch capacity will not put MPLS at a disadvantage relative to IP.

There are still interesting questions about the role of MPLS in ISP networks. For example, the cost and performance tradeoffs among restoration alternatives at different layers (IP layer rerouting, MPLS fast re-route, and transport network restoration) are an active area of research. In addition, as mentioned earlier, transport networks and ATM networks are traditionally more reliable than IP networks – in part because routing problems in one ISP's network can affect other providers. RFC2547 isolates routing in Layer 3 VPNs from Internet routing, which directly addresses this problem for VPN customers.

5. CONCLUSION

This paper surveys several of the visions for ATM explored by the networking community in the early to mid 1990's. Many of the traffic management concepts developed in ATM have become part of networking practice, and ATM is widely used to support VPN's, DSL access networks, and as a core networking technology for Frame Relay and some voice and IP services. Nonetheless, ATM did not succeed in revolutionizing networking. Economic factors, network externalities, the complexity of emerging standards and implementations, and the rapid development of alternative technologies were all factors which made it difficult for ATM to take over the world as many people expected.

From an historical perspective, the debate between connection-oriented and connectionless networking technologies has existed since the early days of packet switching. Even with the explosive growth in IP communications over the last decade, it appears that the tension between the two *technologies* is alive and well. Connection-oriented technologies are the basis of the optical transport networks that underlie most data networks below OC-48 rates, while MPLS is now becoming mature enough to support VPN services in large ISP backbones. It seems likely that connection-oriented technologies will continue to play a significant, if largely invisible, role in data networks at Layers 1.5 and 2 for some time.

6. ACKNOWLEDGMENTS

The author gratefully acknowledges the contributions of Tom Afferton, Elie Francis, Han Nguyen and K.K. Ramakrishnan, for their careful feedback on early drafts of this paper. K.K. Ramakrishnan provided the perspective on the evolution of datalink technologies and ATM.

7. REFERENCES

[1] R. D. Doverspike, S. Phillips, and J. R. Westbrook, *Future Transport Network Architectures,* IEEE Communications Magazine, August 1999.

[2] I. M. Leslie, D. R. McAuley, and D. L. Tennenhouse, *ATM Everywhere?*, IEEE Network Magazine, March 1993.

[3] R. Sharma and S. Keshav, *Signaling and Operating Systems Support for a Native ATM Applications*, Proc. SIGCOMM '94.

[4] J. Condon, et al., *Rednet: A Wireless ATM Local Area Network using Infrared Links*, Proc. MOBICOM '95.

[5] R. Yuan et al., *A Signaling and Control Architecture for Mobility Support in Wireless ATM Networks*, Mobile Networks and Applications, Special Issue on Wireless ATM, Vol. 1, Issue 3, December 1996.

[6] G. Hjalmtysson and K.K. Ramakrishnan, *UNITE: An Architecture for Lightweight Signaling in ATM Networks*, Proc. Infocom '98.

[7] M. Hayter and D. McAuley, *The Desk Area Network*, Operating Systems Review, Vol. 25, No. 4, pp. 14-21, 1991.

[8] P. Newman, G. Minshall, and T. Lyon, *IP Switching: ATM under IP*, IEEE/ACM Transactions on Networking, Vol. 6, No. 2, April 1998.

[9] A. Shaik, J. Rexford, and K. G. Shin, *Load-Sensitive Routing of Long-Lived Flows*, Proc. SIGCOMM '99.

[10] C.R. Kalmanek, A. Lauck, and K.K. Ramakrishnan, *SUBMARINE: An Architecture for IP Routing over Large NBMA Subnetworks*, Proc. INFOCOM '99.

[11] J. Luciani et al., *NBMA Next-Hop Resolution Protocol*, IETF Request for Comments 2332, April 1998.

[12] G. Armitage, *Support for Multicast over UNI 3.0/3.1 based ATM Networks*, IETF Request for Comments 2022, November 1996.

[13] *LAN Emulation over ATM Version 2.0 – LUNI Specification*, ATM Forum, AF-LANE-0084.000, 1997.

[14] A. Charny, K.K. Ramakrishnan, and T. Lauck, *Time Scale Analysis and Scalability Issues for Explicit Rate Allocation in ATM Networks*, IEEE/ATM Transactions on Networks, Vol. 4, No. 4, August 1996.

[15] *Traffic Management Specification Version 4.0*, ATM Forum, AF-TM-0056.00, 1996, available at:ftp://ftp.atmforum.com/pub/approved-specs/af-tm-0056.000.pdf.

[16] M. Grossglauser, K.K. Ramakrishnan, *SEAM: A Simple and Efficient Architecture for ATM Multipoint-to-Multipoint Communication*, Proc. Infocom'97.

[17] M. Waldvogel et al., *Scalable High-Speed IP Routing Lookups*, Proc. Sigcomm '97.

[18] M. Degermark et al., *Small Forwarding Tables for Fast Routing Lookups*, Proc. Sigcomm '97.

[19] E. Mannie (ed.), *Generalized Multi-protocol Label Switching (GMPLS) Architecture*, IETF Work in Progress.

[20] L. Martini et al., *Encapsulation Methods for Transport of [Ethernet frame, ATM Cells, etc.] over MPLS Networks*, IETF Work in Progress.

[21] L. Martini et al., *Transport of Layer 2 Frames over MPLS*, IETF Work in Progress.

[22] L. Andersson et al.*, PPVPN L2 Framework*, IETF Work in Progress.

[23] E. Rosen and Y. Rekhter, *BGP/MPLS VPNs*, IETF Request for Comments 2547, March 1999.

[24] *Multi-Protocol over ATM Technical Specification Version 1.0*, ATM Forum, AF-MPOA-0087.000, 1997.

[25] Newman P., et al., *Ipsilon's General Switch Management Protocol Specification Version 2.0*, Request for Comments 2297, March 1998.