Connectionless Service for Public ATM Networks

Brett J. Vickers and Tatsuya Suda
bvickers@ics.uci.edu and suda@ics.uci.edu

Department of Information and Computer Science,
University of California, Irvine,
Irvine, CA 92717-3425
(714) 856-4105 (voice)
(714) 856-4056 (facsimile)

Abstract

The advent of broadband switching and communications technology (*e.g.,* ATM) has made possible the efficient transmission of multimedia information like voice and video. ATM, however, is connection oriented in nature. Transmission occurs on virtual channels established prior to transmission. In order to accommodate the internetworking of more traditional connectionless LANs and MANs across public wide area networks like B-ISDN, a means by which connectionless service can be provided on top of connection oriented technologies must be devised. This paper surveys architectural and conceptual alternatives that have been proposed and explored as ways to provide connectionless services in public ATM environments.

Suggested Summary Sentence: Providing connectionless service for the interconnection of LANs and MANs in a public ATM environment will likely entail the use of connectionless servers.

1 Introduction

A large number of existing networks are connectionless by nature of their medium access schemes. Broadcast or shared media networks (*e.g.,* Ethernet) and token or slot passing networks (*e.g.,* Token Ring, FDDI, DQDB) are inherently connectionless; no connection setup is required, and the network simply broadcasts or forwards packets using common MAC, logical link, or bridging protocols. Connectionless wide area networks (*e.g.,* Internet) are also inherently conducive to connectionless service as a result of their utilization of datagram routers. Connectionless networks place few burdens on end systems; they are typically not responsible for connection management or routing decisions.

ATM networks, on the other hand, are connection oriented. Virtual circuits are established between end systems, and packets (or *cells*) are switched according to connection identifiers. Because network resources are statistically allocated on a per-connection basis, connection oriented networks like ATM allow for strict quality of service guarantees. Furthermore, because routing decisions are made only at connection establishment time, cell sequence integrity is maintained.

In order to obtain the benefits of connection oriented networks in connectionless networks, connection oriented services are often offered. For instance, the IEEE 802.2 logical link control (LLC) layer has the ability to provide connectionless networks with services that emulate the functionality of connection oriented networks on a hop-to-hop basis. It contains primitives for establishing logical connections, for providing guaranteed data delivery, and for releasing connections. Higher layers also have the ability to provide connection oriented services. Transport layers such as TCP commonly provide the user with connection oriented services that keep track of transmission state on an end-to-end basis and thereby emulate the behavior of connections. Connection oriented services are typically implemented in protocols, because they are not inherently supported by connectionless networks.

Network technologies like ATM are the reverse situation. They are inherently connection oriented but provide connectionless services to obtain the benefits of connectionless networks and to allow interoperability. There are essentially two suggested methods for delivering connectionless service in ATM networks. One is to map connectionless medium access control (MAC) or IP protocol functionalities directly onto ATM connections. This path is being followed by the ATM Forum and the Internet Engineering Task Force's IP over ATM Working Group. An alternative approach is provision of connectionless service at the user-network interface. This approach is popular with pub-

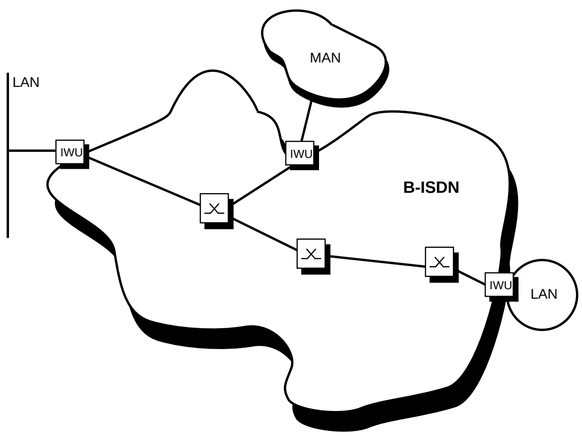

Figure 1: The indirect approach to connectionless service

lic network service providers as a means for supporting the interconnection of connectionless LANs and MANs, and has been prescribed by the ITU-T (formerly the CCITT) for B-ISDN. This latter approach will be the main focus of this paper.

2 The ITU-T's Recommendations

The ITU-T recommends two general approaches for providing connectionless service in an ATM network: the *indirect* and *direct* approaches (See ITU-T recommendations I.211 and I.364). Within each of the approaches, there are a multitude of architectural options available to the network designer, many of which are discussed in later sections.

2.1 The Indirect Approach

In the indirect approach, connectionless service is offered by the network through the use of virtual connections established between each pair of ATM interworking units (IWUs). An interworking unit forms the interface between the connectionless local area network and the wide area ATM network. Virtual connections (which may be viewed as "pipes") are simply built from one interworking unit to another, creating a dense mesh of connections within the network. See Figure 1. Such connections may be established all at one time using permanent or semi-permanent virtual connections (PVCs), or they may be established when necessary and torn down during idle periods if switched virtual connections (SVCs) are utilized. The choice of SVCs or PVCs depends on such factors as the size of the network and the quality of service desired.

There is an important tradeoff between the use of PVCs and the use of SVCs to provide indirect connectionless service. The use of SVCs is more scalable since connections are not maintained when there is no data to transmit. Thus, the use of PVCs may only be applicable in ATM environments where the number of connected LANs is small enough to allow for a fully interconnected mesh. On the other hand, SVCs may suffer from significant connection establishment overheads, especially in a wide area environment. Interworking units become responsible for buffering packets while connections are being established, and delays in transmission may result. PVCs have the advantage that such delays are completely avoided.

Regardless of whether PVCs or SVCs are used, the indirect approach suffers from its poor utilization of network resources, particularly bandwidth. If, when a connection is established for connectionless service, the amount of requested bandwidth is too large, it will remain underutilized by connectionless service users for extended periods of time. Similarly, if too small an amount of bandwidth is reserved for the connection, large delays and frame loss rates may result. Furthermore, the indirect approach does not scale well. As the number of interconnected end systems grows, so too does the number of connections required to support them all.

Schemes to avoid such problems have been suggested. One such scheme is best effort or available bit rate (ABR) service [1]. ABR service attempts to accommodate connectionless traffic by allowing it to utilize unused bandwidth within the network. ABR is often refered to as best effort service because it attempts to accommodate connectionless traffic with whatever available bandwidth the network has at the time. The network uses feedback to notify interworking units of congestion, and the interworking units adjust their output rate accordingly. This type of service may be extremely useful in an ATM LAN. However, in wide area environments where propagation delays are significant, feedback-based services such as ABR may not be feasible without significant supporting hardware. (Best effort service for wide area ATM networks is the subject of section 4.)

Variants of the indirect approach have been and are being adopted by both the ATM Forum and the Internet Engineering Task Force's IP over ATM working group. While these schemes are different from the indirect approach in that the network itself is not offering connectionless service to the user, they are similar to the indirect approach in that only ATM connections are used to facilitate connectionless service. No additional hardware is necessary.

The ATM Forum has focused its attention on the provision of connectionless service to end systems through LAN emulation [1]. The basic idea of LAN emulation is to make an ATM end system appear to the network as a typical IEEE 802 LAN end system. This is achieved through the implementation of a specialized MAC sublayer tailored

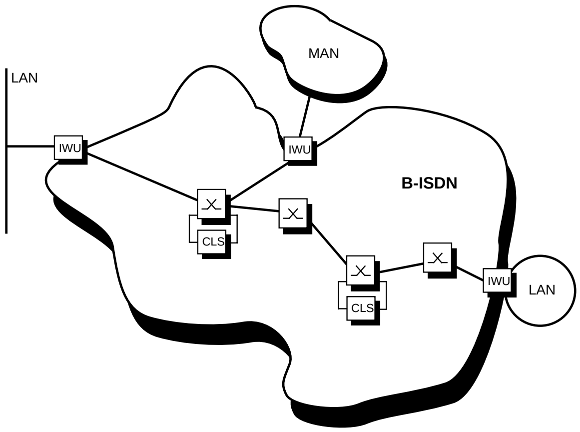

Figure 2: The direct approach to connectionless service

specifically for ATM. Within the ATM MAC sublayer proposed by the ATM Forum exists such functionality as connection management, signalling, address resolution, and MAC-layer encapsulation. By emulating the IEEE 802 interface, the LAN emulation approach allows ATM end systems to interconnect with other networks through bridges or routers in much the same way that systems in Ethernets and Token Rings do now.

The IETF has taken a slightly different approach from the ATM Forum, choosing instead to support connectionless service for ATM end systems through existing connectionless network layers (i.e., IP) rather than through the MAC sublayer. The IETF solution consists of encapsulating IP frames with IEEE 802.2 LLC and IEEE 802.1a SNAP headers [2]. The encapsulated IP frames are then stored in the payload of an ATM Adaptation Layer 5 (AAL 5) convergence sublayer PDU for subsequent segmentation and transmission via the ATM network. Address resolution in the IETF model of connectionless service is achieved through the use of ARP servers located in each IP subnet [3]. Interconnection of isolated ATM networks is realized through the use of IP routers.

The IETF and ATM Forum solutions appear as if they will be adopted in ATM LANs and within the Internet. Furthermore, they will most likely be used to provide connectionless services to ATM end systems as opposed to providing them for LANs and MANs. The public carriers, due to their desire to control the impact of bursty LAN/MAN data traffic and to support LAN/MAN interconnection in wide area environments, are focusing on providing connectionless service using another approach: the direct approach.

2.2 The Direct Approach

The direct approach alleviates many of the scalability and bandwidth utilization concerns raised by the indirect

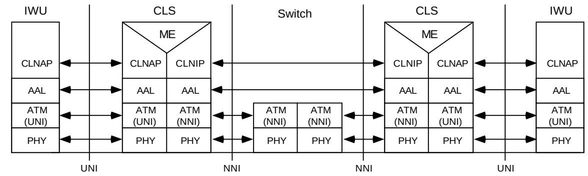

Figure 3: A direct approach protocol architecture

approach. In the direct approach, a connectionless service function (CLSF) is realized through the use of connectionless servers (CLSs) and interworking units. See Figure 2. Interworking units connect connectionless networks to the ATM network and simply forward connectionless data after segmenting or reassembling it. Connectionless servers operate within the ATM network and are usually connected by PVCs to reduce connection setup delays. Each server is responsible for making routing decisions that determine the next-hop server or interworking unit to which a packet or cell must be delivered in order to reach its final destination.

A sample protocol architecture for the direct approach to connectionless service is depicted in Figure 3. At the interworking units, the Connectionless Network Access Protocol (CLNAP) encapsulates datagrams before handing them to the ATM Adaptation Layer. The CLNAP frame format (see ITU-T recommendation I.364) is almost identical to the DQDB Initial MAC PDU format and may contain up to 9188 octets of user information. Once the CLNAP frame has been generated, it is encapsulated in an AAL 3/4 convergence sublayer PDU and segmented by AAL 3/4 into a number of cells. These cells are transmitted to the next hop connectionless server where they are either reassembled or transmitted one cell at a time to the next hop in a cut-through manner. At the network node interface (NNI), CLNAP frames are encapsulated with a 4 octet header by the Connectionless Network Interface Protocol (CLNIP). The Mapping Entity (ME) is responsible for performing the necessary encapsulation and decapsulation. Eventually, the connectionless datagrams arrive at the destination interworking unit where they are reassembled and delivered to the appropriate end system.

The direct approach has a number of advantages over the indirect approach that make it attractive for public wide area environments. First, each connectionless LAN requires only one connection at the edge of the network in order to forward connectionless data. Interworking units are therefore not responsible for routing decisions; the network is. Furthermore, interworking units are only required to manage a single connection. Second, all connectionless traffic is aggregated onto a small number of connections

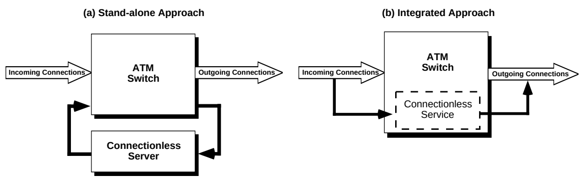

Figure 4: Stand-alone and integrated connectionless servers

between connectionless servers. This provides a statistical multiplexing gain, reduces burstiness of the traffic somewhat, and makes traffic easier to manage. Third, and perhaps most importantly, the degree of connectivity imposed by the indirect approach is reduced. Only connectionless servers are interconnected, diminishing the number of connections required to realize connectionless service. This makes the direct approach inherently scalable.

Of course, the direct approach also introduces some complications to connectionless service. For example, connectionless servers or connections between servers may become performance bottlenecks. The ability to renegotiate bandwidth allocated on connections between connectionless servers may improve performance during periods of congestion, but the servers still have to operate at fast speeds to keep up with ATM switches. Furthermore, the direct approach fails to eliminate the complexity of making routing decisions; it simply pushes this responsibility onto the network. Also, the direct approach has been slated to utilize AAL 3/4, whereas the ATM LAN community has settled on the use of AAL 5. This may complicate interoperability between ATM LANs and public ATM WANs.

3 Providing the Connectionless Service

If connectionless service is to be provided in an efficient and effective manner in public ATM environments, a number of issues must be dealt with. These issues include: where to locate connectionless servers, how to interconnect them, how to forward packets, how to manage bandwidth, and how to resolve destination addresses.

3.1 Locating and Connecting Connectionless Servers

There are two approaches to adding connectionless servers to an ATM network. One is to integrate the connectionless service function into the ATM switch (called Option A in recommendation I.364). The other is to append a connectionless server unit externally to the ATM switch (Option B). Option A will be referred to as the *integrated* approach, while Option B will be referred to as the *stand-alone* approach.

In the stand-alone approach, the connectionless server acts as a self-functioning unit with connectionless data streams arriving from one or more of a switch's outputs. The processed data streams exit the connectionless server and return to one or more of the same switch's inputs. A model of the stand-alone connectionless server is illustrated in Figure 4a. The stand-alone approach is flexible, because it allows the network designer to easily add connectionless servers to switching nodes in the ATM network. It does, however, impose slightly longer delays on connectionless traffic since it is routed through each switch twice, possibly contributing to congestion at ATM switches.

Alternatively, the connectionless service function may be placed inside the switch, resulting in the integrated approach. In this approach (see Figure 4b), cells containing connectionless data are delivered to the connectionless service component of the switch where they are processed and transmitted on the appropriate outgoing switch port. This approach is attractive but requires that wide area ATM equipment manufacturers implement the connectionless service function in their switches.

It is not likely that connectionless servers will be placed at every switch in the ATM network. Therefore, the determination of where to place them and how to interconnect them has important performance consequences. One recommendation is to place connectionless servers at each interworking unit's first-hop switch [4]. This configuration eliminates switch-to-switch connections between an interworking unit and the first-hop connectionless server, improving the bandwidth efficiency and subscriber cost of first-hop connections. It also makes management of connectionless service simpler, since address screening and tariff calculation will most likely occur at the local switch. However, placing connectionless servers more centrally may reduce the number of servers required.

In addition to the location of connectionless servers, their interconnection also determines the performance and flexibility of the offered connectionless service. The topology of a virtual connectionless network ultimately determines how quickly connectionless cells are routed from source to destination. Due to the connection oriented nature of ATM, virtually any topology for the interconnection of connectionless servers is possible. However, five basic topologies have been proposed as feasible alternatives: *complete mesh* [5], *arbitrary connections*, *hierarchical tree topology* [5], *partitioned hierarchical tree topology* [5], and *ring topology* [6]. Figure 5 illustrates each of these alternatives.

The complete mesh is perhaps the most intuitive topology and entails the connection of each connectionless server to every other server. For relatively smaller pub-

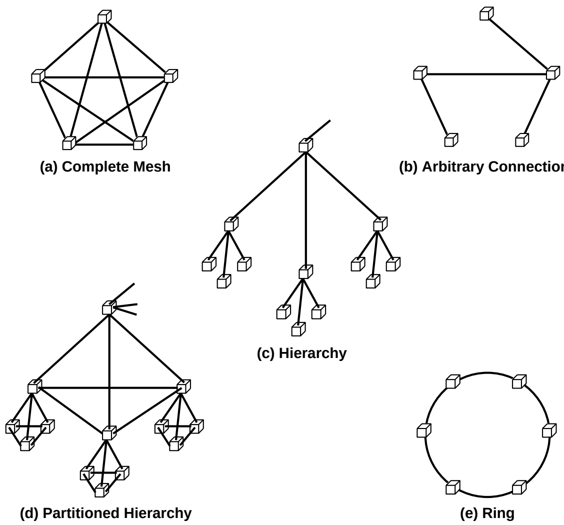

Figure 5: Connectionless server network topologies

lic networks, this method may prevail since the number of servers required is small and end-to-end delay is minimal. However, for public wide area ATM offerings, the mesh is problematic for the same reason that the indirect approach is problematic. If n is the number of connectionless servers in the network, $O(n^2)$ connections, each of which may use a non-negligible amount of network bandwidth, become necessary to complete the mesh.

The arbitrary interconnection of connectionless servers reduces the number of connections between servers required to support connectionless service. Each connectionless server is connected arbitrarily to one or more other connectionless servers in the network, and routing tables in each server are updated to reflect their interconnections. This topology has the advantage that the routing tables in each server are smaller than those for the mesh.

An alternative to the rather intuitive server interconnection schemes presented above is to connect connectionless servers hierarchically into a tree topology. Figure 5c illustrates such a virtual connectionless network. This topology may be useful for hierarchical addressing schemes like E.164. For instance, each level of the tree may perform its routing operation on a different chunk of the destination E.164 address. The problem with the strict hierarchical topology, however, is that servers near the root of the tree may become congested as communication across sub-trees increases.

One way to alleviate this congestion is to organize connectionless servers into a partitioned hierarchical tree topology as shown in Figure 5d. Each level of the hierarchy is partitioned into several groups of servers, each of which is capable of communicating directly with any other server in its group. If the number of servers in a partition is kept low, routing tables remain relatively small. This scheme has the advantage that it reduces congestion through parent servers. The size of the routing table is slightly larger than it is in the case of a non-partitioned hierarchy, but only by an amount proportional to the size of the partition mesh.

Also possible for the interconnection of connectionless servers is the ring topology as shown in Figure 5e. One of the primary benefits of the ring topology is that it reduces the complexity of routing algorithms that are otherwise required by mesh topologies. Cells are merely forwarded along fixed routes between connectionless servers and copied to destination interworking units when appropriate. Another advantage of a ring topology is that it simplifies broadcast and multicast transmission of connectionless data. Since data circulates on a closed ring of connectionless servers, it can be copied and sent to any number of interworking units in the network. However, one of the problems associated with this topology is that the ring may grow substantially in larger networks, resulting in transmission and processing delays.

3.2 Forwarding Schemes

Connectionless servers may transmit connectionless data in one of two basic ways [7]. They may send cells to the next-hop connectionless server or interworking unit as soon as they are received. This is called *cell-based forwarding*. Alternatively they may reconstruct frames at each server before resegmenting and transmitting them. This method is known as *frame-based forwarding*.

The cell-based forwarding approach requires the use of AAL 3/4. When the first cell of a frame arrives at the connectionless server, a unique outgoing connection identifier (VCI/VPI) and message identifier (MID) combination is looked up in an internal routing table using the destination address stored in that cell. This information is stored in the server so that subsequent cells with the same incoming identifiers can be forwarded as soon as they arrive. In order to keep pace with the flow of traffic, a cell-based connectionless server must perform the reception and transmission of a cell and the look-up of connection identifiers within a single cell transmission time. For an OC-3 SONET interface rate of 155.52 Mb/s, cells must be processed and transmitted within 2.7μs. For the OC-12 622.08 Mb/s service, the maximum time is reduced to 680 ns. This required performance impacts design choices, forcing the use of high-speed and parallelized digital circuitry.

On the other hand, connectionless servers using frame-based forwarding internally buffer a frame's cells until the frame is entirely received. The frame is then segmented and forwarded to the next hop. Because frames are processed one at a time, AAL 5, which lacks a multiplexing identifier, may be used. This makes frame-based forwarding a feasible candidate for the interconnection of ATM LANs. Frame-based forwarding alleviates the time restrictions that plague cell-based forwarding servers. It also enables the connectionless server to drop an entire frame if a single cell is damaged, thereby reducing useless traffic in the network. However, it makes buffer size and buffer management more important concerns and imposes several types of delays on frames [8]:

- *Reception delay:* caused by delaying the transmission of the first cell of a frame until the frame's last cell arrives at the connectionless server. This delay is at least as long as the transmission time of the frame.

- *Queueing delay:* caused by delaying the transmission of a frame until cells from other frames contending for the same outgoing port are transmitted.

- *Processing delay:* caused by processing that takes place in the server due to table lookups, buffer management, address resolution and error checking. This delay is dependent on the processing speed of the server.

So, while frame-based forwarding has the advantage that it eases timing requirements and allows frame-level control, it can impose considerable end-to-end delays on connectionless traffic.

3.3 Bandwidth Management for Connectionless Service

If the connections between connectionless servers are to be provided with some degree of guaranteed service, i.e., if they are to be allocated bandwidth, it may become necessary during their lifetime to adjust bandwidth allocation due to changes in connectionless traffic flux. One strategy is to have connectionless servers (or interworking units) directly renegotiate bandwidth with the network when traffic conditions change [9]. Output buffers in the rate control devices of connectionless servers or interworking units are monitored. When the buffer size exceeds an upper threshold, additional bandwidth is requested from the network by using the appropriate signalling protocol. When the buffer size decreases below a lower threshold (or, alternatively, when the buffer remains empty for a period of time), bandwidth is released back into the network using a similar signalling procedure.

Due to the fact that the request and release of bandwidth entails network signalling, renegotiation is less than immediate. The signalling overhead must be taken into account. Since wide area ATM signalling protocols have not yet been fully standardized, it is difficult to draw detailed conclusions regarding signalling overhead. However, it can safely be concluded that, at a minimum, signalling cells must be sent from one connectionless server to another in order to perform renegotiation on the link between them. Thus, the propagation delay becomes an important factor in evaluating the applicability of renegotiation, particularly in wide area ATM networks. If the time required to renegotiate bandwidth is longer than the interval over which traffic load may fluctuate severely, renegotiation may not be a worthwhile option.

Another problem with bandwidth renegotiation is that choosing thresholds for the bandwidth tracking algorithm is difficult. If, for example, upper and lower thresholds are chosen too close together, a hysteresis effect may result, creating an inordinate amount of signalling traffic. If, however, the thresholds are set too far apart, the algorithm may not be sensitive enough to changes in traffic conditions, and unwanted buffering delays or bandwidth inefficiency may result.

A bandwidth management strategy that ameliorates some of renegotiation's less desirable consequences is bandwidth advertising [10][11]. Bandwidth advertising is a hybrid between best-effort service and bandwidth renegotiation. Like best-effort service, unused bandwidth is "borrowed" from other connections to send connectionless traffic. A connectionless server or interworking unit knows when it can borrow bandwidth because the network periodically "advertises" the amount that is available. On the other hand, like bandwidth renegotiation, bandwidth can be requested if the amount of advertised bandwidth remains consistently insufficient to support connectionless traffic.

Bandwidth advertisement is beneficial because it reduces the number of renegotiations (and therefore the amount of signalling) that would otherwise have to be performed whenever a large burst of connectionless traffic arrived. It also reduces the amount of buffering that is required since large bursts of traffic can be immediately sent into the network when unused bandwidth is available. On the other hand, it introduces traffic in the form of available bandwidth notifications and requires the implementation of a traffic monitoring function within the network. Furthermore, it can have a deleterious effect on other types of traffic in the network, since bursts are sent on unreserved bandwidth.

3.4 Address Resolution

One other issue complicating the provision of connectionless service for wide area ATM networks is address resolution. Suppose host a in LAN A wishes to transmit connectionless data through the wide area ATM network to host b in LAN B. Assume the interworking unit connecting LAN A to the ATM network knows that host b resides in LAN B. Communication should then take place according to the route $\{a, \ldots, \text{IWU}_A, \ldots, \text{IWU}_B, \ldots, b\}$, where IWU_X is the interworking unit connecting LAN X to the ATM network. In this simple example, complex address resolution is not required since LAN A is aware of host b's location. However, in most cases, interworking units are not initially aware of the LAN in which another host resides. Address resolution must therefore take place in order to discover the address of the destination interworking unit so that communication with the destination host can occur.

There are essentially two classes of address resolution techniques that an interworking unit may utilize in order to discover a destination interworking unit address. One technique consists of broadcasting the destination host address to each interworking unit in the ATM network and waiting for the appropriate interworking unit to respond with its ATM address. This technique is similar to the Address Resolution Protocol (ARP) [12] commonly used in local area networks and requires the existence of ATM multicast service. The other technique consists of querying centralized or lightly distributed address resolution databases, each of which is updated with a list of host addresses whenever a new interworking unit is added to the network. This is the approach that the IETF IP over ATM working group has adopted [3].

Although ARP was originally intended for use in networks using shared media, it can be extended to ATM networks if they have broadcast or multicast capability. One possible approach to using multicast ARP for ATM connectionless service is to model it after the multi-LAN address resolution technique [13]. When an interworking unit wishes to transmit a connectionless frame to a host in a LAN connected to the network at some unknown location, it broadcasts an address resolution request to every other interworking unit in the ATM network. As interworking units receive the request, they compare the requested destination address against an internal cache containing a list of other recently requested destination addresses. If found, and the destination host exists in the interworking unit's local network, then the interworking unit returns its own ATM address to the source of the ARP request so that it may serve as a proxy for the destination host. If the address is not found in the cache, the interworking unit performs a local ARP request on each of its connected LANs and determines if the specified host is connected.

An alternative to multicast ARP is to use address resolution entities or ARP servers [1] [3]. These ARP servers are placed in the ATM network and respond to address resolution requests from interworking units. Whenever an interworking unit joins the ATM network, it updates the nearest ARP server with a list of host addresses in its attached local area network(s). Address resolution proceeds by querying ARP servers with a host address and receiving the destination interworking unit address in return. As in the case of the multicast address resolution protocol, interworking units serve as proxies for host data. Note that the functionality of this type of address resolution requires that each interworking unit know at least one ARP server's address.

For even larger interconnected environments, where networks are interconnected to networks which are in turn interconnected to other networks, similar server-based solutions have been proposed for address resolution. This has been the primary area of work for the IETF's Routing over Large Clouds Working Group [14].

There are tradeoffs involved with these two address resolution techniques. The multicast ARP method requires the broadcast of ARP packets to a large number of interworking units. This broadcasting may not be feasible in wide area ATM networks where network resources are scarce. Although interworking units can reduce their need to broadcast ARP requests by locally caching the results of previous requests, the broadcast of address resolution packets to a large set of interworking units of which only one contains the actual destination host is inefficient to say the least. The ARP server technique reduces the amount of traffic necessary to perform address resolution but transfers complexity and storage requirements to one or more central servers. It is also much less fault tolerant than the multicast ARP method. If a single ARP server crashes, its entire database may become unavailable unless there are multiple ARP servers.

4 Best Effort in the Wide Area

Among ATM researchers, there is growing consensus that connectionless service should be provided on a best effort basis. As it stands, the direct approach to connectionless service, while practical and attractive for public network service providers, does not support the best effort transport of connectionless data. Instead, it requires the guaranteed allocation of network resources on connections between connectionless servers.

On the other hand, the best effort service schemes suggested thus far utilize feedback-based control and do not adequately scale to wide area offerings. Feedback can take substantial time to propagate in a wide area environment. In order to deal with large bandwidths and feedback propagation delays that ATM can support, wide area switches offering best effort service end up requiring a significant amount of buffer space.

What is needed for wide area public ATM networks is a best effort connectionless service that maintains the scalability and resource manageability of the direct approach but allows for the available bit rate (ABR) transmission of connectionless data.

4.1 Credit-based and rate-based flow control

The two classes of flow control for best effort service that have been proposed to the ATM Forum are credit-based and rate-based schemes. In credit-based schemes (e.g., Kung and Chapman's FCVC scheme [15] and DEC's AN2 flow control), each virtual channel requires a credit before a data cell can be sent. When a data cell is transmitted, a credit is consumed. As data cells are transmitted on the VC, credits are intermittently sent to the upstream node in order to maintain a continuous flow of data. In rate-based schemes (e.g., Peter Newman's BECN scheme [16]), ATM switches notify traffic sources of congestion when an output buffer's occupancy exceeds a threshold. Only those sources contributing to the congestion are notified. Upon receiving a congestion notification, a traffic source slows the rate at which it transmits connectionless data into the network.

Credit-based schemes have the advantage that they apply link-by-link backpressure on sources generating congestion. This allows them to achieve excellent bandwidth utilization with relatively little loss of connectionless traffic in a fully loaded ATM network. However, one of the drawbacks of these schemes is that they require per-VC buffers in every switch. In the case of a wide area network where the propagation delay and bandwidths are substantial and the number of VCs is large, the required buffer space becomes prohibitively high. Large switch buffers are undesirable because they result in longer end-to-end delays and significantly complicate the development of switch hardware. Another drawback of the credit-based schemes is that they often require that a significant fraction of the bandwidth be used for the transport of credit cells.

Rate-based schemes fail to achieve optimal link utilization, but they ease buffering requirements in switches and generate fewer control cells on average. Unlike credit-based schemes, rate-based flow control avoids per-VC buffering. Instead, switch output buffers are monitored. However, like credit-based flow control, rate-based flow control does not scale well to the wide area. This is due largely to the fact that congestion notifications must travel from the congested node back to the originating traffic source. Wide area switches require significant buffer space to support such large propagation delays.

Thus, rate-based and credit-based schemes, while popular in the ATM LAN environment, do not scale well to the wide area ATM network.

4.2 Flow control in the wide area

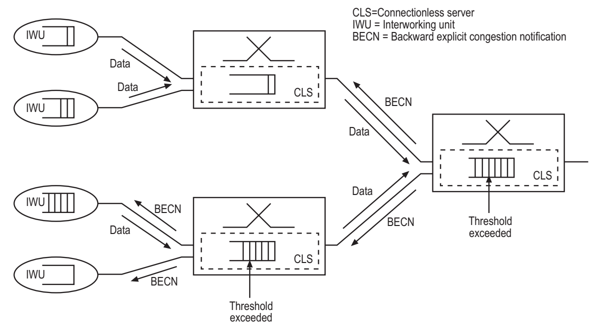

Figure 6: Best effort connectionless service for public ATM

One way to support best effort connectionless service in the wide area is to perform link-by-link congestion notification between integrated frame-based connectionless servers. Instead of allocating bandwidth to connections between connectionless servers, ABR connections are used. When a connectionless server's buffer occupancy exceeds a threshold, congestion notifications are sent to the last hop connectionless servers as shown in Figure 6. As connectionless servers choke their transmission rates, the congestion effect propagates to the source interworking units, creating a backpressure effect similar to that achieved in credit-based schemes. In effect, this scheme is a hybrid of rate-based, credit-based, and connectionless server approaches to connectionless service.

This form of feedback-based best effort service has a number of characteristics that make it useful in a wide area ATM network. First, it does not require the per-VC buffering that credit-based schemes do. One shared buffer space is used for the reassembly of all frames arriving from other connectionless servers or interworking units. Second, buffer space for connectionless data is decoupled from buffer space for guaranteed traffic. Buffer space in switches is preserved for the use of guaranteed bit rate data, while buffer space in connectionless servers is shared by the less delay-sensitive connectionless datagrams passing through the switch. Third, feedback propagation delays are ameliorated since congestion notifications are transmitted from node to node instead of from node to source, as is the case with rate-based schemes. Finally, interoperability with AAL5-based ATM LANs is guaranteed since AAL5-capable frame-based servers can be utilized.

Given that public ATM service providers appear reluctant to support feedback-based best effort service in the switches themselves, moving this functionality to the connectionless server may be an attractive option.

5 Conclusion

The interconnection of connectionless LANs and MANs over a connection oriented B-ISDN presents a dilemma with regard to efficient interoperability. In order to provide solutions to this dilemma, there have been many techniques proposed. We have attempted to describe and evaluate the architectural and conceptual alternatives available for the implementation of a connectionless service for public ATM networks. We have also suggested a scheme through which best effort connectionless service may be provided in public ATM networks using connectionless servers and hop-by-hop flow control. There is as yet no consensus on issues such as bandwidth management, address resolution, connectionless server forwarding modes, or connectionless server topology, and therefore much of the area remains open for research. These research issues must be addressed, however, in order to realize the interoperability, and thereby the acceptance, of ATM.

References

[1] P. Newman, "ATM Local Area Networks," *IEEE Communications Magazine*, pages 86–98, March 1994.

[2] J. Heinanen, "Multiprotocol Encapsulation over ATM Adaptation Layer 5," Request for Comments 1483, July 1993.

[3] M. Laubach, "Classical IP and ARP over ATM," Request for Comments 1577, Jan. 1994.

[4] T. Kawasaki, S. Kuroyanagi, R. Takechi, and K. Hajikano, "A Study on High Speed Data Communication System Using ATM," *Proceedings of IEEE Globecom*, vol. 3, pages 2105–2109, Phoenix, Ariz., Dec. 1991.

[5] K. Murakami, K. Hajikano, S. Abe, and Y. Kato, "Communication Service and Media Control Using ATM," *IEICE Trans. Comm.*, E 74(4):772–779, Apr. 1991.

[6] K. Yamazaki, Y. Wakahara, and Y. Ikeda, "Networks and Switching for B-ISDN Connectionless Communications," *IEICE Trans. Comm.*, E76-B(3):229–236, March 1993.

[7] G. Fioretto, T. Demaria, R. Vaglio, A. Forcina, and T. Moro, "Connectionless Service Handling Within B-ISDN," *Proceedings of IEEE Globecom*, vol. 1, pages 217–222, Phoenix, Ariz., Dec. 1991.

[8] D. F. Box, D. P. Hong, and T. Suda, "Architecture and Design of Connectionless Data Service for a Public ATM Network," *Proceedings of IEEE Infocom*, vol. 2, pages 722–731, San Francisco, Calif., March 1993.

[9] L. Mongiovi, M. Farrell, and V. Trecordi, "A Proposal for Interconnecting FDDI Networks through BISDN," *Proceedings of IEEE Infocom*, vol. 3, pages 1160–1167, Bal Harbour, Flor., April 1991.

[10] P. Crocetti, G. Gallassi, and M. Gerla, "Bandwidth Advertising for MAN/ATM Connectionless Internetting," *Proceedings of IEEE Infocom*, vol. 3, pages 1145–1150, Bal Harbour, Flor., April 1991.

[11] M. Gerla, T.-Y. C. Tai, and G. Gallassi, "Internetting LAN's and MAN's to B-ISDN's for Connectionless Traffic Support," *J. Select. Areas of Comm.*, 11(8):1145–1159, Oct. 1993.

[12] D. C. Plummer, "An Ethernet Address Resolution Protocol or Converting Network Protocol Addresses to 48-bit Ethernet Address for Transmission on Ethernet Hardware," Request for Comments 826, Nov. 1982.

[13] J. Postel, "Multi-LAN Address Resolution," Request for Comments 925, Oct. 1984.

[14] J. Heinanen and R. Govindan, "NBMA Next Hop Resolution Protocol (NHRP)," Internet Draft, Oct. 1993.

[15] H.T. Kung and A. Chapman, "The FCVC (Flow Controlled Virtual Channels) Proposal for ATM Networks," available through anonymous FTP to *virtual.harvard.edu:/pub/htk/atm/atm-forum/fcvc.ps*.

[16] P. Newman, "Backward Explicit Congestion Notification for ATM Local Area Networks," *Proceedings of IEEE Globecom*, vol. 2, pp. 719–723, Houston, Texas, Nov. 1993.