Making the MBone Real [1]

Ajit S. Thyagarajan
Department of Electrical Engineering,
University of Delaware,
Newark, Delaware 19716
ajit@ee.udel.edu

Stephen L. Casner
Information Sciences Institute,
University of Southern California,
Marina del Rey, California 90292
casner@isi.edu

Stephen E. Deering
Palo Alto Research Center,
Xerox Corporation,
Palo Alto, California 94304
deering@parc.xerox.com

Abstract

IP multicasting is an extension of the Internet Protocol that efficiently delivers a single datagram to multiple hosts instead of a single host. Its benefits for applications such as live audio and video conferencing among Internet sites around the world have been clearly demonstrated over the past three years in an experimental deployment called the "Multicast Backbone", or MBone. The time has now come for the next step—bringing IP multicast service up to production quality and extending deployment to the entire Internet so that the MBone ceases to exist as a separate entity. This paper discusses some of the enhancements that are underway to make the MBone and IP multicast real: scalable routing, congestion management, and diagnostic tools.

1 Introduction

Over the last few years, the use of Internet multicasting has transitioned from being an experimental deployment over private research testbeds, to a widely accepted means of data communication over the Internet. This has resulted from the creation of the MBone—a virtual network layered on top of the Internet, consisting of routers and hosts that support IP multicasting. Ever since its inception in 1992, the MBone has clearly demonstrated its potential by multicasting numerous conferences and events using a variety of presentation formats. Many of the multicast applications that have been developed are capable of generating large quantities of real-time data resulting in uncontrolled consumption of valuable resources (e.g. bandwidth). The public availability and the immense popularity of these applications has resulted in an exponential growth rate as shown in Figure 1. With newer applications of multicast and multicast-based services appearing regularly, and with the population of potential IP multicast users increasing as the Internet itself grows, the existing MBone infrastructure has to accommodate this growth. In order to make the MBone real, we must address issues in the following main dimensions:

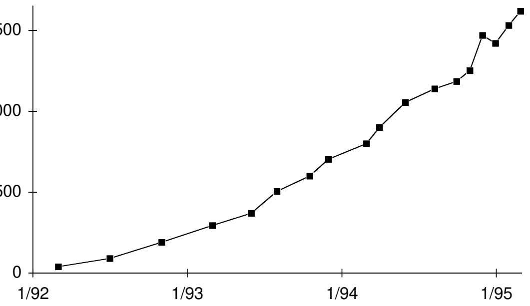

Figure 1: Growth of the MBone—Number of Subnets

Scalability: One of the most pressing problems is the growth of the number of subnets in the MBone such that it can no longer be managed as a single, flat routing domain. In order to reduce the amount of routing information exchanged between routers and stored at every router, modifications to the routing architecture are necessary. Another cause for concern is the capability of the current MBone infrastructure to support a large number of simultaneously-active multicast sessions. Two problems need to be

[1] This research was supported in part by the Advanced Research Projects Agency under contracts DABT63-94-C-0073, DABT63-91-C-0001 and NAG 2-638, and National Science Foundation grant NCR-93-01002. The views expressed here do not reflect the position or policy of the U.S. government.

addressed in this regard: reducing the incidence of unnecessary traffic along links, and reducing the amount of active state maintained at each router. Alternate routing algorithms which address these problems are being developed.

Congestion Control: With the anticipated increase in multicast traffic, the non-applicability of traditional flow control mechanisms, and the current shortage of bandwidth, a number of new congestion-control mechanisms appear to be necessary. Short term solutions to prevent saturation of links have been implemented, but as the MBone grows, other congestion control mechanisms will be necessary.

Diagnosability: Proper diagnostic tools are required in order to aid in the monitoring of the network, and to provide feedback on the operation of the routing and congestion management algorithms. Several on-line and off-line utilities have been developed to accomplish this, and other mechanisms are being implemented to provide feedback from individual receivers.

This paper describes a number of steps that must be taken to evolve the MBone into a reliable and manageable multicast service that is available everywhere in the Internet. It begins with a brief description of the MBone and its applications, followed by some of the proposed solutions to the problems currently outlined above.

2 Overview of the MBone

The MBone grew out of an experiment in March 1992 to enable remote participants to take part in an Internet Engineering Task Force (IETF) Meeting, and to demonstrate technology initially deployed and tested in the ARPA-sponsored network testbed called DARTnet [2]. Live audio was multicast from the meeting site and from many of the remote participants to destinations around the world. That initial MBone topology interconnected approximately 40 subnets in 4 countries. Since then, the MBone has proved to be remarkably successful, having multicast over 250 conferences and international events, with the recent Rolling Stones concert in Dallas being one of the more publicized events. Currently, the MBone connects over 1500 subnets (see Figure 1) in over 25 countries.

This impressive growth has been propelled by the development of several new multicast applications for real-time audio and video conferencing (such as LBL's *vat* and *vic*, Xerox's *nv*, INRIA's *ivs*, and UMass's *nevot*), shared presentation tools (such as LBL's *wb*), and conference managers (such as LBL's *sd* and USC-ISI's *mmcc*). In addition to teleconferencing, other applications include data dissemination utilities (such as multicast netnews, University of Hawaii's *imm* and the Woods Hole Oceanographic Institution's telemetry software), and distributed interactive simulators (such as the Naval Postgraduate School's *NPSnet*)[2]. Today, the MBone carries the audio and video traffic of numerous academic and technical conferences, working group meetings, University seminars, public radio and television programs, Space Shuttle missions and the occasional live music performance.

Since IP multicast is a relatively recent addition to IP, the majority of the installed Internet routers currently do not support multicast routing; most of the MBone routing and forwarding is being performed by general-purpose hosts, such as Unix workstations running multicast routing software (e.g. Sun, DEC, HP, IBM, Silicon Graphics, and PC's running a variant of Unix). Even though a number of router vendors have begun to produce multicast-capable routers (e.g. Cisco, Proteon, Alantec, and Bay Networks), there still exists a very large installed base of Internet routers which do not have this capability. These routers need to be either phased out or upgraded in order to support multicast routing.

Figure 2 shows an example where host H1 is acting as a multicast router between subnets (e.g., LANs) S1 and S2 because router R1 is not yet multicast capable. H1 forwards IP multicast packets as required between S1 and S2, including those relayed to and from S3 by multicast-capable router R2. In this paper, we use the term *multicast router* (or just *router*, when the context is clear) to refer to either a multicast-capable Internet router or a host running multicast routing software.

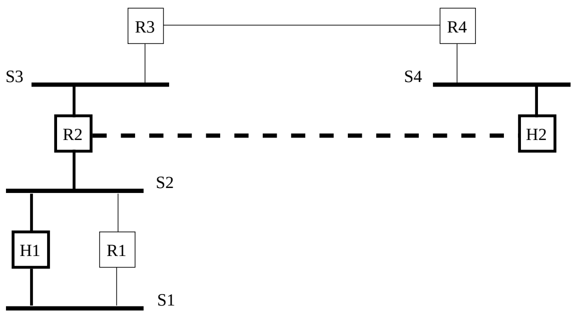

Figure 2: MBone Components

[2] Additional information on the MBone and its applications can be obtained from the following URL *http://www.eit.com/techinfo/mbone*

"Islands" of multicast routing capability, such as the S1–H1–S2–R2–S3 conglomerate in Figure 2, are connected to other islands via *tunnels* through non-multicast-capable parts of the Internet. A tunnel is a path between a pair of multicast routers over which a multicast packet is forwarded by encapsulating it inside a conventional unicast packet addressed to the router at the far end of the tunnel. Tunnels are "installed" simply by manually configuring the router at each end with the IP address of the other end. The figure shows a tunnel configured between multicast routers R2 and H2. The tunnel acts as a virtual point-to-point link between R2 and H2 (shown as a dashed line in the illustration) connecting the pair of multicast routers, and thus the islands in which they reside. Due to the high-bandwidth nature of most multicast traffic, tunnels have been carefully engineered to go over high-bandwidth links only. The set of multicast routers, the subnets to which they are directly attached, and the interconnecting tunnels constitute what is known as the MBone.

The MBone may be thought of as a virtual network (or an overlay network) layered on top of the Internet. As more of the Internet routers are upgraded to support multicast routing, hosts will be retired from performing multicast routing, and tunnels will be eliminated. Only then will the MBone become "less virtual" and eventually become one with the (unicast) Internet. Until that time however, the topology of the MBone will differ from the topology of the Internet. In Figure 2, the unicast topology consists of all the boxes and solid lines, while the multicast topology is made up of only the thick-outlined boxes and the thick lines, including the dashed line.

Due to difference in the MBone and the Internet topology, the multicast routers must run their own topologydiscovery protocol, i.e., routing protocol, in order to decide where to forward multicast packets. In most of the current MBone, the protocol is the Distance-Vector Multicast Routing Protocol, DVMRP [4]. In some parts of the MBone, routers use different multicast routing protocols, in particular MOSPF [8] and PIM [5]; these routers interoperate with DVMRP routers by implementing a subset of DVMRP – just enough to "fool" a DVMRP router into "doing the right thing".

The current implementation of DVMRP treats the entire MBone as a single, flat routing domain. This results in a very large routing table being maintained at each router. Routers using MOSPF and PIM reduce routing information passed to DVMRP routers by replacing a set of subnet routes with a single network route, thus exploiting the hierarchy encoded in IP addresses. An examination of the current set of MBone routes however reveals that this reduction is still not sufficient to match the MBone's growth rate. As with the unicast routes, some sort of hierarchy in managing the routing tables is necessary to reduce the amount of routing data stored at each router. As the MBone becomes an integral part of the Internet, it may be necessary to merge unicast and multicast routing information to avoid having to maintain separate routing tables. The next section describes some of the approaches designed to deal with these problems.

3 Improving Routing Scalability and Traffic Distribution

As the number of subnets in the MBone continues to increase, routers have to maintain large amounts of routing information. In addition, the current traffic distribution model causes the amount of state maintained at each router to increase linearly with the product of the number of active groups and the average number of senders per group. In order to scale to arbitrarily large networks, the amount of routing and traffic state overhead must be reduced significantly. Approaches to achieving more scalable routing and efficient traffic distribution are being studied under the auspices of the IETF.

3.1 Hierarchical Partitioning of the MBone

As with large-scale unicast routing, hierarchically partitioning the topology into multiple, independent routing domains will reduce the amount of routing information that each router must maintain [7]. The primary motivation for deploying hierarchical routing in the MBone is a reduction in the amount of topological information that multicast routers must store and exchange with other routers. Additional benefits which are derived from hierarchical routing are the following:

- Partitioning the MBone into regions allows the simultaneous use of different multicast routing protocols in the various regions of the MBone. Region boundaries isolate the damage that can be done by a misbehaving routing protocol, and allow new protocols to be tested and deployed on a region-by-region basis.
- The effects of low-level topology changes (such as link or router failures and recoveries) are isolated to only those routers in the same topological region as the affected components. Similarly, the effects of changes to

the high-level interconnection topology of regions are limited only to the inter-region routers. Thus, the normal low degree of volatility within any region of the MBone does not sum to a high degree of volatility across the entire MBone, as is currently the case. This is particularly beneficial when using a distance-vector style routing protocol, such as DVMRP, which can incur long convergence times after some topology changes (due to the so-called "counting-to-infinity" problem [6]).

- Limits on the maximum diameter of the topology imposed by some routing protocols (including DVMRP, to ameliorate the counting-to-infinity problem) can be relaxed by using hierarchical routing. As different instances of routing protocols are used for different parts of the MBone, these limits apply only to those parts and not to the whole.

- Some of the new multicast routing protocols [1, 5] do not have their own topology-determining component, but rather depend on the information provided by the underlying unicast routing protocols. This approach works correctly only when the multicast and unicast topologies are identical, which is not the case for the MBone now. Hierarchical routing thus enables the use of these new protocols within regions that satisfy the topological constraint, without necessitating that it be satisfied everywhere.

A scheme for hierarchical multicast routing has been developed [10] and is currently being implemented. Figure 3 shows an example of a network partitioned into separate regions and how routing takes place between the regions. Intra-region routers forward multicast traffic to all destinations within the region and to all "boundary routers", i.e., inter-region routers directly attached to the region. The boundary routers in turn forward this traffic to boundary routers in other regions using a separate distribution tree computed between the boundary routers. If a boundary router determines that any of its attached regions have members of the destination group, the traffic is injected into the region for delivery to the members. Intra-region routers forward multicast traffic to all destinations within the region while the boundary routers forward traffic between the different regions.

Thus, boundary routers maintain only inter-region routing information and intra-region routers maintain routing information pertinent to the region they are in. This leads to a significant

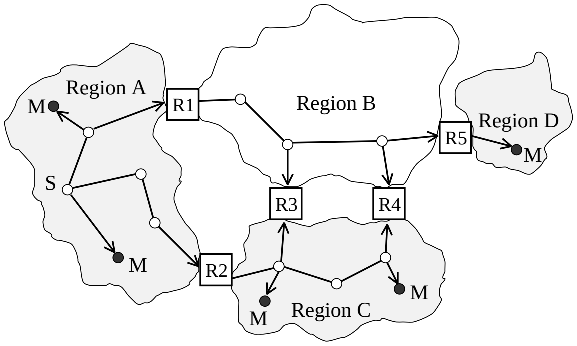

Figure 3: Routing of packets between Regions. S is a source in Region A. Regions A, C and D have members (M) of the destination multicast group. Region B is a transit region.

reduction in the amount of routing information maintained at each router. This routing approach can also be extended to multiple levels of hierarchy to reduce routing overhead.

3.2 Improving the Scaling of Multicast Routing Algorithms

Existing multicast routing algorithms are not well-suited to large-scale multicasting to "sparse" groups, that is, groups whose members are sparsely distributed across a routing domain. For example, DVMRP [4] uses a "broadcast-and-prune" technique in which multicast packets are occasionally broadcast to all routers in the domain, and state becomes established in routers that are not on the path to members of active multicast groups. Similarly, MOSPF [8] periodically broadcasts group membership information to all routers in the MOSPF domain, and requires all routers to remember that information, whether or not they are involved in forwarding to those groups. In domains where bandwidth or router memory is a scarce resource, and where there are many sparse multicast groups active, the overhead of the occasional broadcast traffic and the storage requirement imposed on all routers is very undesirable. To address this problem, two new multicast routing algorithms—CBT and PIM—have recently been designed and are being deployed and tested in parts of the MBone.

CBT stands for Core Based Trees [1]. It solves the overhead problems of the previous algorithms by establishing, for each multicast group, a single multicast delivery tree rooted at a *core* router. Rather than using broadcast mechanisms, multicast senders and multicast receivers send data packets and group membership information, respectively, towards the group's core. CBT allows for a group to have multiple cores for robustness reasons, i.e., to avoid the core becoming a single point of failure.

The single (per-group) delivery tree of CBT is shared among all senders to the group, as illustrated in Figure 4. A benefit of maintaining only one tree per group is that routers need not keep per-source state information. For groups with a large number of active senders, this can be a significant saving over the previous algorithms, DVMRP and MOSPF, that construct a separate delivery tree rooted at each subnet that is a source of multicast packets.

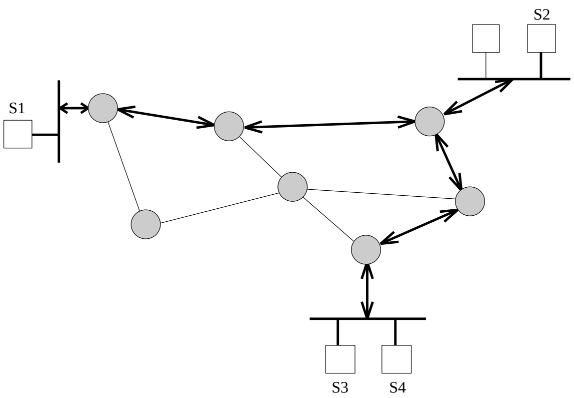

Figure 4: Shared tree for wide-area multicasting

On the other hand, there are a couple of drawbacks to using shared trees. First, for groups in which more than one sender is active, the single tree causes the traffic to be concentrated over fewer links than the per-source tree algorithms, potentially reducing the bandwidth available to each source. Second, for groups that are sensitive to delivery delay, delivering multicast packets over a shared tree usually increases the delivery delay to at least some group members, relative to the per-source, shortest-path tree approach.

The PIM (Protocol Independent Multicast) algorithm [5] is a hybrid algorithm that combines the benefits of CBT and the previous algorithms. Like CBT, it eliminates the occasional broadcast behavior of the previous algorithms by designating one or more core routers (called *rendezvous points* in PIM) for each group, from which a shared tree is built connecting all active senders to all group members. However, PIM also supports the automatic establishment of shortest-path, per-source delivery trees for high-data-rate sources, for which the data-concentration property of a shared tree would be a significant performance limitation.

3.3 Improving the Efficiency of Multicast Traffic Distribution

Significant improvements have been made over the last couple of years to restrict the distribution of multicast traffic to just those destinations where it is needed and wanted. The publicly available multicast software has been enhanced to prune the multicast distribution tree effectively. In addition, modifications to the host-router interface are currently being implemented to provide finer granularity in selecting traffic at a receiver. These techniques thus help reduce the incidence of multicast traffic on unnecessary links. In some cases it is also desirable to restrict delivery of certain multicast traffic to within an organization for privacy, or because it is known that adequate bandwidth is available there. An administrative scoping mechanism has been implemented to achieve this.

3.3.1 Pruning mechanisms

Initially, the MBone software was released with an experimental version of the multicast traffic distribution mechanism that was not intended for full-scale deployment. In that version of the software, multicast traffic originating from a source travelled to all multicast-capable routers subject only to a limited scoping mechanism based on packet hop-count thresholds.

The multicast implementation has since then been completed to detect the absence of group members on a downstream link and prune such links from the multicast distribution tree. Although this mechanism has been made available for quite sometime now, a large fraction of the multicast routers have not been upgraded to support this capability. Part of the problem lies in the fact that the deployment of Internet multicast has been a voluntary group effort so far, and getting individual routers upgraded periodically involves participation from the respective vendors and router administrators. With the traffic on the MBone increasing, the presence of these routers can cause a situation where low-bandwidth links become saturated by traffic that has no downstream receivers. It is hoped that with more widespread use of the MBone and the porting of the MBone routing software to different platforms, this problem will be eliminated in its entirety.

3.3.2 Per-Source Joining and Leaving

The Internet Group Membership Protocol (IGMP) is the protocol used by hosts to join and leave multicast groups [3]. Modifications are being made to the protocol to provide finer granularity by supporting per-source joins and leaves. This will enable a receiver to identify the specific senders that it wishes to hear, or not to hear, thereby reducing the incidence of traffic from unwanted sources. The exact details of the implementation are currently being finalized.

3.3.3 Administrative Scoping

Many organizations felt it necessary to implement some sort of bounding mechanism by which certain multicast traffic was limited to being within the organizational boundaries. An administrative scoping mechanism which limits traffic of certain multicast groups by enforcing strict scoping rules has been implemented. The boundary routers of a particular region are explicitly configured to limit traffic destined for certain multicast groups to within that region. For this purpose, a certain subset of the multicast group addresses (239.x.x.x) have been set aside as scoped addresses. Network administrators now have the flexibility to control the distribution of multicast traffic across their organizational boundaries by simply configuring the boundary routers to deny certain multicast traffic from passing in either direction.

A multicast address used by an application within a scoped region may be used by another application in another scoped region without traffic from each interfering with one another. This leads to efficient re-use of addresses in different regions. For example, the routers R1 through R5 could be configured to scope multicast traffic sent to group address 239.1.1.1. Thus, multicast traffic sent to this group address in Region A would not be forwarded to any of the other Regions. The scoping mechanism also allows traffic to be restricted to regions that have sufficient bandwidth available. This allows applications to make full use of the available bandwidth without worrying about the possibility of this traffic spilling over to low-bandwidth links.

4 Congestion Management

The high-bandwidth, long-lived, non-flow-controlled traffic typical of the MBone, such as real-time audio and video streams, can cause severe congestion in the Internet. Much of the infrastructure of the Internet was primarily designed to support TCP (which has its own congestion control mechanism and has been engineered to be "well-behaved") and short-transaction UDP traffic. The current architecture therefore lacks additional traffic congestion control mechanisms. Multicast traffic is much harder to manage. Traditional flow control methods are ill-suited to this traffic, due to the problem of acknowledgements from many receivers "imploding" on a single sender simultaneously.

With the use of the MBone gaining popularity, suitable mechanisms to avoid overwhelming network links due to multicast traffic are necessary. On the other hand, none of these mechanisms are a substitute for the provision of adequate bandwidth. Fortunately, bandwidth is becoming cheaper and more available in many parts of the Internet, and applications other than the MBone, such as the World Wide Web, are driving the installation of that bandwidth.

4.1 Rate-Limiting Mechanisms

Many network providers are reluctant to have multicast traffic pass through their networks for fear of affecting the performance of non-multicast applications. However it is not only multicast applications that are bandwidth-intensive; a number of non-multicast applications, such as the World Wide Web, also devour precious bandwidth. Hence mechanisms to enable proper management and use of network resources are universally required.

As a temporary measure, a rate-limiting mechanism that provides an option to set a ceiling for multicast traffic has been implemented in current versions of the MBone software. This mechanism is particularly useful in preventing low-bandwidth links from being saturated by multicast traffic. It has been implemented using a simple token-bucket filter as shown in Figure 5. An incoming packet on the link is first queued and then forwarded if there are enough tokens present in the bucket. If the packet queue is full, the incoming packets are discarded. In the long term, this mechanism is expected to be replaced by other traffic management and resource control mechanisms.

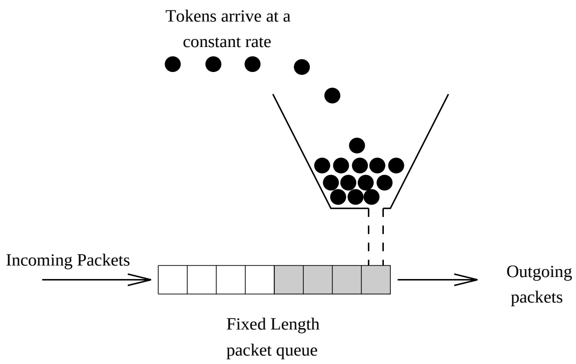

Figure 5: Token bucket rate-limiting mechanism to set an upper limit on multicast traffic

4.2 Priority Dropping Mechanism

Along with the rate-limiting mechanism, a simple mechanism for allocating bandwidth to the different traffic streams has been implemented. Using this mechanism, a router assigns different priorities to multicast packets based on the source and destination addresses and possibly other fields. If the packet queues on a link were to overflow, the lower priority packets present in the queue are dropped in preference to the higher priority packets. This ensures that the quality of service of high priority traffic is not compromised as a result of congestion.

4.3 Fair-Queuing Mechanisms

The non-flow-controlled nature of multicast traffic can be hazardous for a network if several high-bandwidth sources are active simultaneously. Low-rate multicast traffic is then forced to compete with the high-rate traffic. While this may not affect the high-rate traffic, the performance of low data-rate applications may be reduced significantly. In order to prevent any one source from consuming all of the capacity of a bottleneck link, a fairness criterion is required to moderate between the active sources. One way to accomplish this would be to divide the available bandwidth among the active sources and not allow any source to exceed its quota. A fair-queuing mechanism which regulates both multicast and unicast traffic per source would accomplish this effectively. Suitable algorithms which would perform "fair" queuing are currently being researched.

4.4 Explicit Reservation Mechanisms

Another approach to achieving "controlled" use of the available bandwidth is to use explicit reservation mechanisms to guarantee or deny bandwidth for multicast traffic. Each receiver explicitly identifies the traffic desired and the routers reserve the requisite bandwidth for that purpose. Admission control algorithms can then make use of the reservation information in making decisions about forwarding multicast traffic. This approach is being pursued by the ReSerVation Protocol (RSVP) [11] group in the IETF.

4.5 Layered Media Encoding

With any of the preceding mechanisms, a receiver would like to adjust its reception rate according to the bandwidth available. A source variably encodes high bandwidth streams into multiple substreams with possibly variable bandwidths, distinguished by different priorities and/or multicast groups. The receiver then joins only those groups which it can receive clearly. For example, a high-rate video stream could be hierarchically encoded into a thumbnail picture as one group, a low-resolution stream as another group and a high-resolution video stream as yet another group. A receiver attempts to join each stream, and depending on the packet loss rate experienced in each group, adjusts its group membership to obtain the best quality of service possible. This is illustrated in Figure 6, which shows a source sending data in two streams. Receiver R1 is able to receive both the streams, whereas receiver R2 is able to receive only one stream due to a low bandwidth link on its subnet.

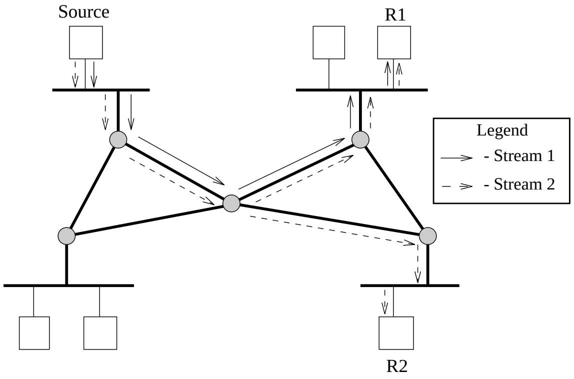

Figure 6: Encoding a single stream into two for reception on low bandwidth links

5 Diagnostic Mechanisms and Tools

As with any large network, proper diagnostic tools are required to monitor the state of the network and to provide feedback on the operation of the routing and congestion management algorithms. The nature of the multicast traffic and lack of effective tools has made it difficult to identify problems quickly. Only recently has significant interest been focused on developing tools for the diagnosis of multicast traffic. Some of the tools and mechanisms developed are described below.

5.1 Multicast Traceroute

A multicast traceroute utility (Mtrace) similar to the unicast traceroute tool has been developed for monitoring the state of the network and providing feedback on the operation of the routing and congestion management algorithms. Because the techniques used by the unicast traceroute tool are not applicable to multicast, a new multicast traceroute function has been defined through an extension to the IGMP protocol which is implemented in the multicast routers and in the Mtrace utility. Mtrace thus aids in the quick diagnosis of faults among routers and provides coarse-grain traffic statistics useful for determining "hot spots" in the network.

The basic-level operation of the Mtrace tool is shown in Figure 7, and can be described as follows:

(1) A requester (H) sends a traceroute query to the destination router (Rd) requesting a trace from the destination (D) towards the source (S).

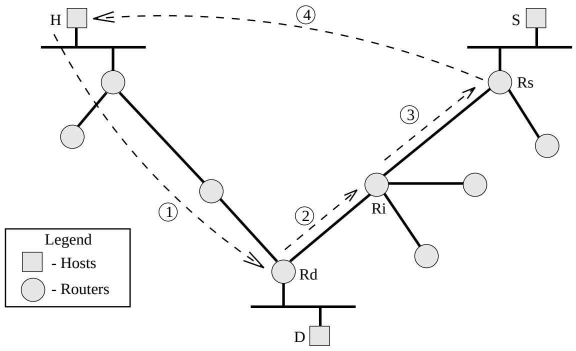

Figure 7: Multicast traceroute between S and D performed by remote host H

(2) The destination router (Rd) adds its response report to the packet and forwards it to the next hop router (Ri) towards the source.

(3) Each router in the path then adds its response to the packet and forwards it towards the source. Each response includes packet and interface counts, routing, and error information. This information can then be used to identify lossy links and misbehaving routers for corrective action.

(4) The router closest to the source (Rs) adds its response and completes the trace by sending the completed response packet back to the requester (H).

It is quite possible that the destination router (Rd) is unknown to the requester (H), or (Rd) and/or (H) may be unreachable via unicast. This problem can often be overcome by multicasting the query or response.

5.2 Other Utilities

The widespread use of multicast has prompted the development of several utilities for monitoring the health of the routers and collective state of the MBone. Some of these include a multicast router information utility (*mrinfo*), an off-line route debugger (*mrdebug*) and an MBone mapping utility (*mapper*), which builds the complete connectivity diagram of the MBone. Support for Simple Network Management Protocol (SNMP) has also been added to the multicast routers to permit monitoring of multicast traffic using standard network management techniques.

5.3 RTP reports

The most effective means to diagnose multicast traffic distribution problems and to monitor the quality of the distribution on a large scale is to collect reception quality feedback from the receivers themselves. By this technique, the real data serves as the test traffic. Distribution quality is monitored continuously in real time for all receivers.

This reception quality feedback mechanism is an integral part of the Real-time Transport Protocol (RTP) recently developed by the Audio/Video Transport Working Group of the IETF [9]. This function is related to the flow and congestion control functions of other transport protocols. The feedback may be directly useful for control of adaptive encodings as well in monitoring the quality of the distribution. Each receiver sends reception feedback reports to all participants so that one who is observing problems can evaluate whether those problems are local or global. If all of the receivers in one area report poor reception, then the common links upstream from that area would be suspect. It is expected that most of the IP multicast audio and video tools will incorporate RTP.

6 Conclusions

The MBone has been regarded as one of the Internet's "success disasters": an experiment that has rapidly outgrown the confines of the lab or the testbed, using prototype software that was never meant to operate at the scale that is being demanded by its users. Its growth has been enabled in part by significant improvements in multicast traffic distribution, the use of temporary congestion control mechanisms and the availability of useful diagnostic tools. In order to foster the continued growth, additional mechanisms to restrict traffic to only where required are being tested, alternate traffic routing strategies to reduce routing overhead are being deployed, and feedback mechanisms are being implemented provide useful information on the reception quality.

As the MBone continues to grow, additional measures will be required to accommodate the effects of the increased traffic. The availability of higher capacity network links and high-performance routers will play a crucial part in the widespread use of this technology. This, along with the implementation of good congestion control and feedback mechanisms will enable the MBone to merge with the global Internet, thus making Internet multicasting a powerful mechanism for distributing data efficiently.

7 Author Information

Ajit S. Thyagarajan is a doctoral student in the Department of Electrical Engineering at the University of Delaware. Since 1993, he has been actively involved in the design, implementation and support of the MBone infrastructure.

His current research interests include routing algorithms and network protocol design. He received his B.Tech in Electronics and Communications from the Indian Institute of Technology, Madras in 1990, and M.S. in Computer Engineering from Villanova University in 1992.

Stephen L. Casner received his M.S. in Computer Science from the University of Southern California in 1976. Since 1973, he has worked at USC's Information Sciences Institute on network protocols and systems for real-time communication over packet-switched networks. He is currently Project Leader for Multimedia Conferencing, working on teleconference session management architectures and protocols. He is also chairman of the Audio/Video Transport working group of the Internet Engineering Task Force which is developing the Real-time Transport Protocol (RTP) for packet audio, video and other real-time applications. He was the primary organizer for the establishment of the worldwide Internet Multicast Backbone (MBone).

Stephen E. Deering is a member of the research staff at Xerox PARC, engaged in research on advanced internetwork technologies, including multicast routing, mobile internetworking, scalable addressing, and support for multimedia applications over the Internet. He is present or past chair of numerous Working Groups of the Internet Engineering Task Force (IETF), and a co-founder of the Internet Multicast Backbone (MBone). He received his B.Sc.(1973) and M.Sc.(1982) from the University of British Columbia, and his PhD (1991) from Stanford University.

References

[1] Ballardie, A. J., P. F. Francis, and J. Crowcroft, "Core Based Trees", In *Proceedings of the ACM SIGCOMM '93*, August 1993.

[2] Casner, S. and S. Deering, "First IETF Internet Audiocast", *ACM SIGCOMM Computer Communications Review*, Vol.22, No.3, July 1992.

[3] Deering, S., "Host Extensions for IP Multicasting", RFC 1112, August 1989.

[4] Deering, S., *Multicast Routing in a Datagram Internetwork*, Ph.D. Thesis, Stanford University, December 1991.

[5] Deering, S., D. Estrin, D. Farinacci, V. Jacobson, C. Liu, and L. Wei, "An Architecture for Wide-Area Multicast Routing", *Proceedings of the ACM SIGCOMM '94*, August 1994.

[6] Hedrick, C., "Routing Information Protocol", RFC 1058, June 1988.

[7] Kleinrock, L. and F. Kamoun, "Hierarchical Routing for large networks; performance evaluation and optimization". *Computer Networks*, Volume 1, pages 155-174, 1977.

[8] Moy, J., "Multicast Extensions to OSPF", RFC 1584, March 1994.

[9] Schulzrinne, H., S. Casner, R. Frederick, V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", Internet Draft, March 1995.

[10] Thyagarajan, A., S. E. Deering, "Hierarchical DVMRP for the MBone", *To appear in Proceedings of the ACM SIGCOMM '95*, August 1995.

[11] Zhang L., S. E. Deering, D. Estrin, S. Shenker, and D. Zappala, "RSVP: A Resource Reservation Protocol", *IEEE Network*, September 1993