A Long-Term Analysis of Growth and Usage Patterns in the Multicast Backbone (MBone)

Kevin C. Almeroth

Department of Computer Science
University of California
Santa Barbara, California 93016-5110
almeroth@cs.ucsb.edu
(805)893-2777(office)
(404)893-8553(FAX)

July 1, 1999

Abstract

The Multicast Backbone (MBone), the Internet's multicast research infrastructure, has existed since the early 1990s. Since its inception, there have been few formal studies investigating the "state of multicast", i.e. the success of multicast deployment. Our work attempts to understand (1) how the MBone is used, (2) how multicast deployment has progressed, and (3) what barriers exist for the continued deployment of multicast. Future research and deployment efforts would benefit significantly with answers to these macro-scale questions. Our work is based on data sets collected over a 4.5 year period. These data sets include join/leave statistics for many of the MBone sessions advertised through the MBone's session directory tool. Using this data, we examine characteristics about the number and frequency of multicast groups and about the users participating in these groups. In addition, we attempt to qualify the accuracy of our results by examining other sources of multicast traffic statistics. Among our conclusions, we find that the use and deployment of multicast has continued to grow, but at a relatively slow rate. Furthermore, most of the activity in the MBone occurs among a relatively small group of core users. Finally, in an attempt to understand why these conditions exist, we offer several possible reasons why multicast deployment has not been more rapid.

1 Introduction

The Multicast Backbone (MBone), the research infrastructure for evaluating multicast communication in the Internet, has existed since the early 1990's. One focus of the MBone's early evolution was on the development and refinement of multicast routing. Another was on transport layer protocols to provide additional services like reliability and real-time streaming media support. Having achieved a relatively versatile and effective set of routing protocols, the MBone is at a new juncture in its evolution. Commercial interests are increasingly driving the requirement for multicast to evolve beyond a research concept. Pressure is mounting to make multicast a fully deployed, ubiquitous service on par with unicast. Some in the multicast community suggest the MBone should disappear as we no longer need a "research infrastructure". The thinking is that we should now be working toward a fully deployed, native multicast infrastructure. This process is underway and a number of Internet Service Providers (ISPs) are focused on deploying multicast[1].

Since the MBone's inception, relatively little formal research has been conducted on how the MBone is used, what performance gains multicast has provided, how multicast deployment has progressed, and what barriers exist for the continued deployment and success of multicast. While some of these issues have begun to be addressed, many of the efforts have been targeted at answering detailed questions about group performance and stream statistics. Understanding the bigger picture has proven to be a difficult task, primarily because there exist insufficient data and experience to fully describe the "state of multicast".

As the clamor for more rapid deployment has increased, the need for a qualitative examination of MBone deployment and performance has grown. There is a desire to know how multicast tools have been used, are being used, and will be used. This kind of understanding would offer insight into how to develop multicast protocols, especially from the point of view of identifying, anticipating, and correcting performance bottlenecks.

To address these needs, we present a long-term analysis of MBone group dynamics. This analysis is based on the collection of join/leave data for MBone sessions advertised through the session directory tool, *sdr*. The focus of this paper's analysis is on discovering long-term trends about the MBone. In particular, we first examine the characteristics of group members by reporting results about MBone group size and the number of joins and average duration for each unique IP address. Second, we look at the MBone population as measured by unique IP addresses and also examine a breakdown based on domain name information. The third set of results looks specifically at the lifetime (duration of participation in the MBone) for each IP address. Finally, we look at *sdr* sessions characteristics.

A basic conclusion of our analysis is that the MBone is growing slowly and that only a small percentage of MBone users are actively participating. We find similar results for *sdr* sessions. Given these conclusions, we offer a number of possible explanations and then describe some of the work to address the underlying problems.

The remainder of this paper is organized as follows. In Section 2 we offer a brief history of the MBone and relevant analysis work. Section 3 describes our data collection methodology. Section 4 describes the results of our analysis. Section 5 discusses problems identified by our analysis and potential solutions. The paper is concluded in Section 6.

2 The MBone and Related Analysis Work

2.1 A Brief History of the MBone

Interest in building a multicast-capable Internet began to achieve critical mass in the late 1980s. Part of Stephen Deering's dissertation[2] proposed an architecture and service model sufficient to begin deploying multicast[3]. The evolution of this work lead to the creation of the Multicast Backbone (MBone)[4]. In 1992, the MBone carried its first world-wide event when 20 sites connected together to receive the March meeting of the Internet Engineering Task Force (IETF)[4]. This first audio conference allowed a few members from all over the world to hear what was being said at the San Diego meeting. In addition to the conferencing software, the most significant achievement was the deployment of the first multicast-capable routers. These routers were actually workstations running a daemon to process and forward multicast packets. Connectivity between these machines was provided using point-to-point IP-encapsulated *tunnels*. Each tunnel connected two end-points as a logical link but physical connectivity may have included intermediate several routers. Once a multicast packet was received at a tunnel endpoint, it could be broadcast on a local network and/or sent through other tunnels. Routing between multicast-capable routers was provided using the Distance Vector Multicast Routing Protocol (DVMRP)[2].

Since 1992, the MBone has grown tremendously. It no longer is a simple virtual network sitting on top of the Internet, but is rapidly becoming integrated into the Internet itself. In addition to simple DVMRP tunnels between workstations, the MBone now has native multicast capability provided in the routers themselves. Other multicast routing protocols like Protocol Independent Multicasting (PIM)[5] and Multicast Open Shortest Path (MOSPF)[6] have been standardized and deployed. Improvements have also been made in transport layer services. The Real-time Transport Protocol (RTP)[7] assists loss- and delay-sensitive applications in adapting to the Internet's besteffort service model. The trend and eventual goal of the MBone is to facilitate the efficient and effective transmission of real-time multimedia data over the Internet.

With the MBone's growth, the demand for new, better, and different applications has also grown. Since the first audio conference in 1992, the MBone has seen the development of new applications based on an increasingly diverse set of media types. Originally, the MBone was considered a research effort and its evolution was overseen by members of the MBone community. Coordination of events was handled almost exclusively through the use of a global session directory tool, originally called *sd*, but now called *sdr*[1]. As multicast deployment has continued, and as it is has been integrated into the Internet as a native Internet service, the informal use agreements and guidelines have faded. Even though *sdr*-based sessions remain at the core of Internet multicast events, their percentage of the total is shrinking. Other applications are being deployed which do not coordinate sessions through *sdr* or use RTP. This potpourri of tools has enriched the types applications available, but has made tracking multicast use more difficult.

2.2 Relevant Analysis Work for the MBone

The subject of long-term growth and usage patterns in the MBone has been the focus of only a handful of research studies. Efforts have been more focused on debugging problems and developing new protocols and applications. The analysis work on usage patterns that has been conducted has primarily been focused on analyzing group member behavior. Relevant work in each of these areas is described below.

The measurement of MBone traffic was some of the earliest work done in multicast analysis[8]. Mah captured several traces of multicast traffic at UC Berkeley and presented various results including total traffic bandwidth, packet rates, and a breakdown based on media type. More recently, Yajnik et al., examined packet loss for a topology of 17 locations and analyzed the spatial and temporal characteristics of packet loss[9]. In other work, Handley has looked at the details of packet loss in a multicast tree and how it varies over time[10]. Massey and Fenner have characterized DVMRP implementation bugs and observed their impact on the stability of routing tables[11].

With respect to understanding group and tree dynamics, Almeroth and Ammar conducted an investigation of the join/leave behavior of various MBone sessions with the goal of modeling the inter-arrival time and duration of users into a group[12]. This work also looked at tree dynamics and how the size of the tree changed as the group membership changed. In a related vein, the

---

[1] Two repositories are: http://www-mice.cs.ucl.ac.uk/multimedia/software/ and http://www-mash.cs.berkeley.edu/mash/.

Cooperative Association for Internet Data Analysis (CAIDA) has visualized various topological characteristics of the MBone[13].

Finally, a discussion of relevant work would not be complete without mentioning some of the online efforts currently under way. AT&T, Merit Network, CAIDA, and UC–Santa Barbara all offer WWW pages detailing the status of the MBone using various metrics ranging from multicast group statistics to routing table trends.

3 Data Collection Methodology

The basis for the analysis in this paper is the collection of RTP packets for group members joining MBone sessions. Information about transmitters was gathered by listening to RTP data packets. Information about group participants (including transmitters) was gathered by listening to Real-time Transport Control Protocol (RTCP)[7] packets. The data sets used have been collected using a two-step process: (1) capturing RTP/RTCP packets, and (2) post-processing the raw data. Details about the data collection process are described in the following sections.

3.1 Capturing Group Membership

Data was collected using a modified version of the publicly available collection tool *Mlisten*[14]. *Mlisten* collects information about multicast group membership using the following method:

1. *Mlisten* joins the well-know *sdr* multicast group address and listens for session advertisements.
2. For advertised sessions heard at the collection site, *Mlisten* joins the multicast group and port number for media types that are audio, video, whiteboard, or text streams. Each media type uses two <*IP_mcast_address, UDP_port*> pairs, an even numbered port for RTP packets and an odd numbered port for RTCP packets.
3. For each <*IP_mcast_address, UDP_port*> pair, *Mlisten* collects data about arriving packets. Information includes the time the packet arrived, the packet's sender and local UDP port number, and the packet's multicast address.
4. At periodic intervals, *Mlisten* will attempt to identify transmitters who have stopped sending RTP packets, and receivers who have stopped participating. Information about group activity and membership is written to a log file.

Mlisten uses an abstraction called a "connection" to represent the duration of a source's transmission or the amount of time a receiver has been in the group. The connection is defined to start when a packet arrives with the new, unrecorded tuple: <*multicast_addr*, *sender_IP_addr*, *sender_local_UDP_port*>. The connection logically ends at the point when no more packets arrive for a specified tuple. Identifying the "last" packet requires use of a timeout value and additional post-processing (see below). If no packets are heard for a period longer than the timeout value, the connection is considered ended.

For each completed source/receiver connection for each stream in each group, a log record is written. Each record contains the record's unique tuple and information about the start time of the connection, the duration, and the total number of packets received. These records are the atomic elements of the collected data sets.

The data sets generated by *Mlisten* may not accurately reflect true group membership or behavior. The primary problem stems from lost RTP/RTCP packets. If loss between a group member and the collection location is high, several or even many RTP/RTCP packets may be lost. The result is that some group members are not heard from for a period of time longer than the end-of-connection timeout value. Multiple records showing several short connections may actually be one single, longer connection. Post processing the raw log files will combine split records with the same tuple. The uniqueness of a tuple implies a unique connection because each new instantiation of an MBone tool uses a new, local UDP port number.

3.2 Raw Data Sets

The data sets analyzed in this paper cover a 4.5 year period ranging from February 2, 1995, to November 15, 1999. All data was collected from a host located at Georgia Tech in Atlanta, Georgia. During this time, the functionality of the *Mlisten* tool has been continually expanded to collect more information. The data sets collected in 1995 include only 7 audio sessions covering time periods primarily in the middle of February, late May, late June, and early July. Beginning in January of 1996, data was collected for all audio sessions, and as of April 1996, data was collected for all video sessions. In November 1997, *Mlisten* was extended to collect data for all audio, video, whiteboard, and text sessions heard through *sdr*. A summary of the data sets is shown in Table 1.

3.3 Collected Data Sets in Perspective

One of the critical aspects of analyzing data is understanding its limitations. The data collection task for this paper has been particularly difficult. The primary reasons are the duration of the collection period and the difficulty of accurately collecting MBone-wide statistics. In an attempt to quantify the data accuracy, we have identified a number of critical limitations. Given these limitations, we have constrained our investigation to analyses which should not be affected. In particular, we focus on a long-term, macro-scale analysis of growth and usage. We consider a

Start Date	End Date	Type of Data Collected	Num Packets Collected
02-Jan-95	17-Jan-95	Shuttle launch, UCB session, IPNG meeting	46×10^6
22-Jun-95	31-May-95	IMS session, MBone Audio, Radio Free Vat	25×10^6
22-Jun-95	11-Jul-95	Shuttle launch	113×10^6
08-Jan-96	17-Jan-96	All audio	49×10^6
22-Feb-96	31-Mar-96	All audio	112×10^6
27-Aug-96	10-Dec-96	All audio,video	976×10^6
16-Jan-97	30-Oct-97	All audio,video	2665×10^6
31-Oct-97	31-Dec-97	All audio,video,wb,text	702×10^6
01-Jan-98	15-May-98	All audio,video,wb,text	1987×10^6
25-May-98	31-Dec-98	All audio,video,wb,text	5124×10^6
01-Jan-99	31-May-99	All audio,video,wb,text	2932×10^6
01-Jun-99	15-Nov-99	All audio,video,wb,text	1619×10^6
02-Feb-95	15-Nov-99	TOTAL	16.3×10^9

Table 1: Summary of collected data sets.

discussion of data set limitations to be independently useful because it provides insight into the difficulties of collecting this type of data on this kind of scale. The limitations include:

Single Collection Point. Not all *sdr*-based sessions are global and so a single collection site can limit the sessions seen by the collection tool. The ideal solution would be to install *Mlisten*-style tools at a couple, several, or even many different locations around the world. However, for the results in this paper, it is not critical to have a distributed, global view. The reason is that the focus of our analysis is to investigate the number and general characteristics of MBone participants. Even if users primarily participate in local sessions, it is likely that from time to time, they will join a global session. *Any* participation in a global session gives us information about a user's IP address. An important question, but one beyond the scope of this paper, is whether (and how) people participate in local sessions differently than in global sessions.

Dependence on RTCP Feedback. Our collection methodology is based on the capture of RTCP packets. If a single user, group of users, or entire organization is unwilling or unable to transmit RTCP feedback across an administrative boundary and into the MBone, the *Mlisten* tool will not be aware of those users' participation. Furthermore, recent deployment and monitoring efforts suggest that while a large number of sites are capable of receiving MBone traffic, some sites are incapable of sourcing traffic. This anomaly, for many sites, is possibly unintentional and may not even be a known problem. Our requirement for inclusion in this study is that participants must be able to successfully transmit multicast packets into the MBone.

Non-MBone-Style Applications. During the early development of the MBone, an informal, historical set of use and operating guidelines came to exist. "MBone-style" sessions use *sdr* for session advertisement and RTP/RTCP for real-time stream support. Even private sessions are advertised, but they are usually encrypted to provide privacy. The problem is that as multicast has evolved, users have not been using *sdr* or any other tool to advertise their sessions. Even more significant is the increasing number of multicast-capable tools that do not implement RTP and have no need, desire or means of tracking group membership. In particular, several popular commercial tools have added multicast support, but do not have any feedback mechanisms. Without a feedback mechanism, it becomes nearly impossible to capture group membership with an *Mlisten*-style tool.

One User $\neq$ One IP Address. Our data collection method can only collect the IP address of the host on which the session tools are run. We do not know if there is a single user or multiple users per machine. The solution to this problem is to not equate users with IP addresses and report results only for IP addresses.

In general, we attempt to minimize the impact of any of these limitations by focusing our analysis on tracking participation in global, MBone-style sessions. Furthermore, we now attempt to quantify the ability of *Mlisten* to collect data for all global multicast sessions. By taking a snapshot of multicast forwarding state at a router in the network we can compare how many sessions are seen by *Mlisten* versus how many groups have forwarding state in the network. We have taken snapshots of active multicast groups at a core router in the MBone (known as "fixw-mbone"). For a recent snapshot taken early in 1999, fixw-mbone reported the existence of forwarding state for 199 multicast addresses. The router's *sdr* cache showed 158 addresses in use. A snapshot taken for *Mlisten* showed 128 multicast addresses in use. Possible reasons for the differences in these numbers include:

- The difference in number of groups between fixw-mbone's *sdr* cache and fixw-mbone's forwarding table can be attributed to two reasons. First, any non-*sdr* sessions will appear in the forwarding table but will not appear in the *sdr* cache (increases number of forwarding entries). Second, any groups with no traffic, i.e. no RTP/RTCP transmitters or receivers, will not appear in the forwarding table (reduces number of forwarding entries).
- The difference in fixw-mbone's *sdr* cache and *Mlisten's* list of session can mostly be attributed to differences in the number of local sessions.

For the snapshots listed above, it is important to look at the actual multicast addresses. Comparing fixw-mbone's forwarding table with fixw-mbone's *sdr* cache, 20% of the *sdr* cache addresses had routing state. Also, we identified roughly 20 addresses used by multicast routing protocols[15, 5, 6], group management, and multicast tools like *mtrace*[16]. In addition, 30 groups were being announced by *sdr* but had no group members. More than 50 groups were identified as not being MBone-style groups. Informal observations suggest that this trend continues to grow. The differences between fixw-mbone's *sdr* cache and the *Mlisten* cache were due almost entirely to local groups. These results suggest there are some groups not being captured by *Mlisten*, but information for a large percentage of groups is being collected.

4 Analysis

In this section, we look at four different characteristics related to long-term MBone growth and use. We first examine the characteristics of group members and report results about overall MBone participation. For each unique IP address, we compute the total number of joins and total duration statistics. Next, we look at the MBone population and present a breakdown based on host domain names. The third set of results looks specifically at the lifetime (duration of participation in the MBone) for each IP address. Finally, we look at characteristics of *sdr* sessions.

4.1 Group Member Activity and Characteristics

Figure 1 shows the number of receivers and transmitters observed from January 1, 1997 to November 15, 1999. Significant drops in activity are typically the result of collection discontinuities. Analysis of individual sessions can be found in other work[12]. The results in the top row show that except for times with major events, there have been only about 100 to 200 active group members. These numbers are down significantly compared to data prior to 1997 when single events drew hundreds of participants. The results in Figure 1, along with data about *sdr* session types, suggests that the MBone is primarily used either for one-to-many, broadcast-style applications or for small-scale collaborations. An obvious trend is that the overall rate of participation is not increasing and can even be seen to be decreasing. Several reasons are discussed in Section 5, but worth noting here is the increased use of non-RTCP-based tools.

Turning to more specific characteristics about IP addresses, the results in Figure 2 show both time- and value-ordered graphs for average join and duration information for each unique IP address heard in the MBone. The three graphs, from top to bottom, show the number of joins, the average duration per connection, and the total MBone join time for each unique IP address. Results are calculated over the lifetime of each IP address. Each of these three characteristics are described below.

- A total of 33,545 unique IP addresses were heard during the 4.5 year collection period. Results show, that in general, only a small number of users frequently use the MBone. For total joins, 90% of all users joined multicast groups less than 74 times, and only the top 1% have joined an MBone group more than 350 times.
- For the average per-connection duration data, results show that 20% of IP addresses had an average duration of at least an hour, and 70% had an average of at least one minute. It is difficult to draw specific conclusions about use based solely on this data. For example, using data about MBone "session surfing"[12] suggests users join and leave several groups before finding an active group with worthwhile content. The distribution for an individual's duration data is likely to reflect many short and a few long connections.
- By multiplying a user's number of joins and average per-connection duration, we can compute the total amount of MBone join time. This result is shown in the last row of Figure 2. We see that only slightly more than 20% of all users joined the MBone for an aggregate period of more than one day over the course of the 4.5 years of the study. Furthermore, less than 5% of users have used the MBone for more than 300 hours.

4.2 MBone Population Statistics

We now present results that describe long-term characteristics of the MBone's population. The population is measured by counting the number of unique IP addresses. In Figure 3, the graph in the upper-left corner shows the "growth" in MBone membership. This graph was generated using a population counter set initially to zero and then parsing each data set in chronological order. Each time a new IP address was encountered, the counter was incremented. The flat parts early in the graph represent periods when data was not collected (see Table 1). The graph shows that the number of unique IP addresses seen in the MBone is steadily growing. However, further analysis shows the second-order growth rate is decreasing and the overall rate is slowing. From another perspective, even though the MBone has sustained growth, the percentage of Internet hosts who have participated in global MBone sessions is close to zero.

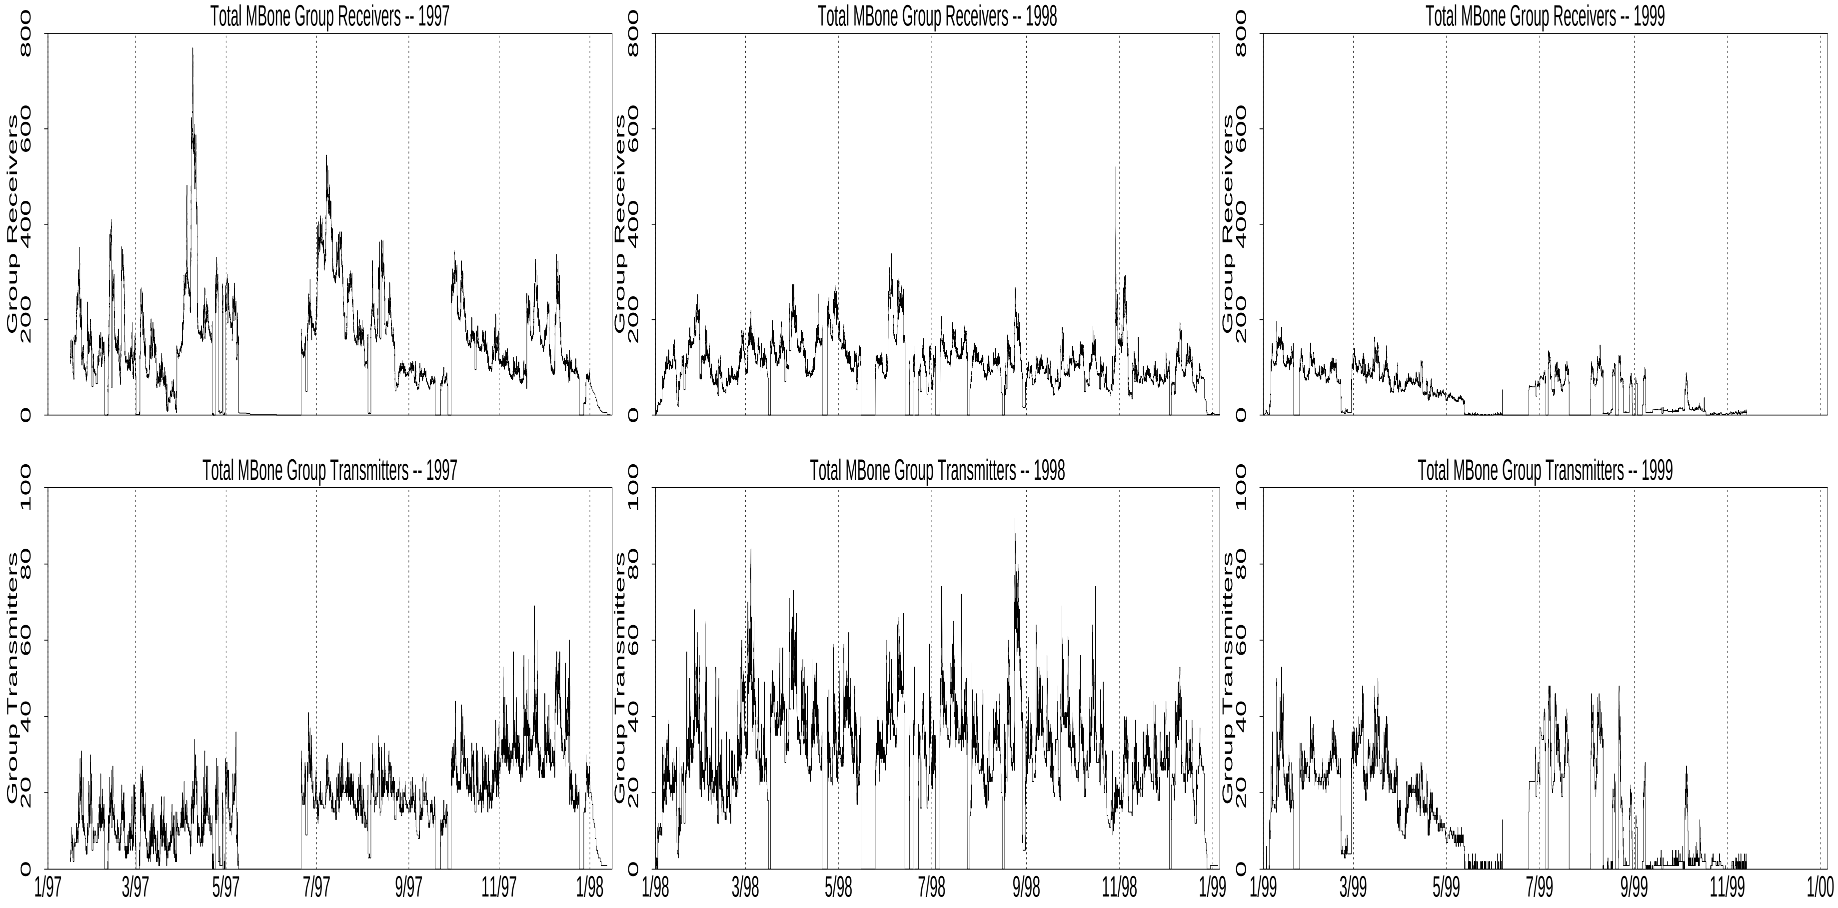

Figure 1: Receivers and transmitters from January 1, 1997 to June 30, 1999 (Y-axes have different values).

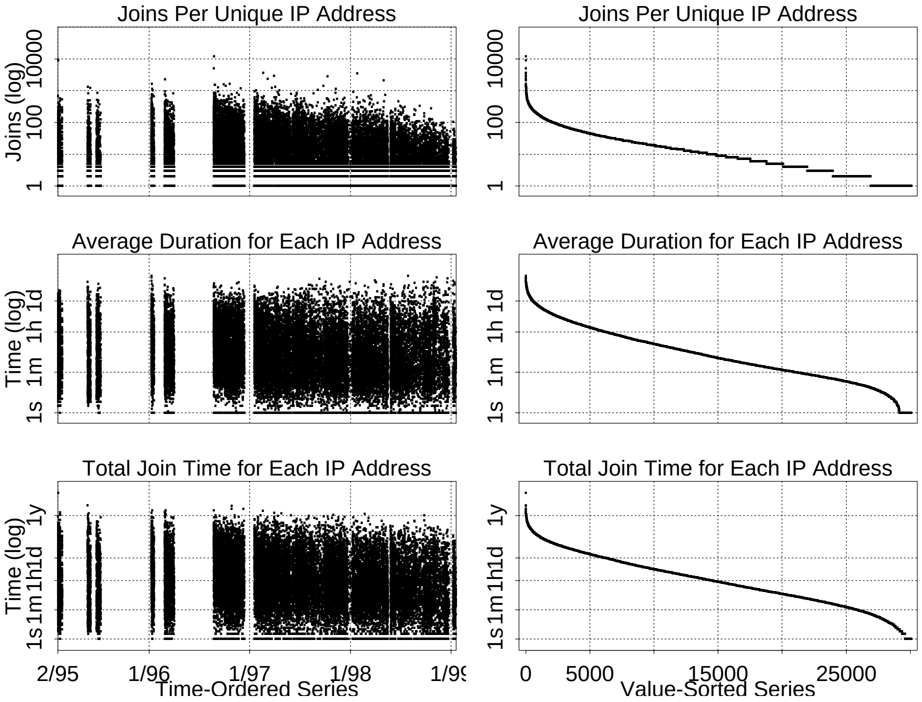

Figure 2: Join and duration results for each unique IP address observed in the MBone.

Not only is the MBone growing slowly, but many of the IP addresses that have been seen at one time or another are no longer participating. This “loss” of members is shown in the lower-left graph in Figure 3. This graph was generated by using the same method as the growth curve but data was processed in reverse chronological order. Instead of finding the *first* time an IP address was heard, the loss curve measures the *last* time an IP address was heard. The graph can be interpreted by noting that each point represents the number of unique IP addresses that have been heard since the date on the x-axis. For example, since the beginning of 1998, only 10,600 of the MBone’s 33,545 total IP addresses have been heard. This means that more than 20,000 IP addresses have not joined a global session in more than a year. If we can consider the size of the MBone to be the number of IP addresses participating in the last six months, the size of the MBone (as of the last data point in January 1999), is approximately 5000 IP addresses.

The two graphs in the right column of Figure 3 show growth and loss results broken into four categories based on the domain name of the DNS entry for each IP address. Hosts are grouped into one of four categories: (1) North American hosts (edu, com, ca, mx, etc.), (2) European hosts (connected via trans-Atlantic links), (3) Pacific Rim hosts (connected via trans-Pacific links), or (4) unknown. Results show the that North American hosts represent the largest number of MBone participants. However, more recent trends suggest many North American hosts have stopped participating. European hosts had significant growth during the middle of 1997. A likely hypothesis

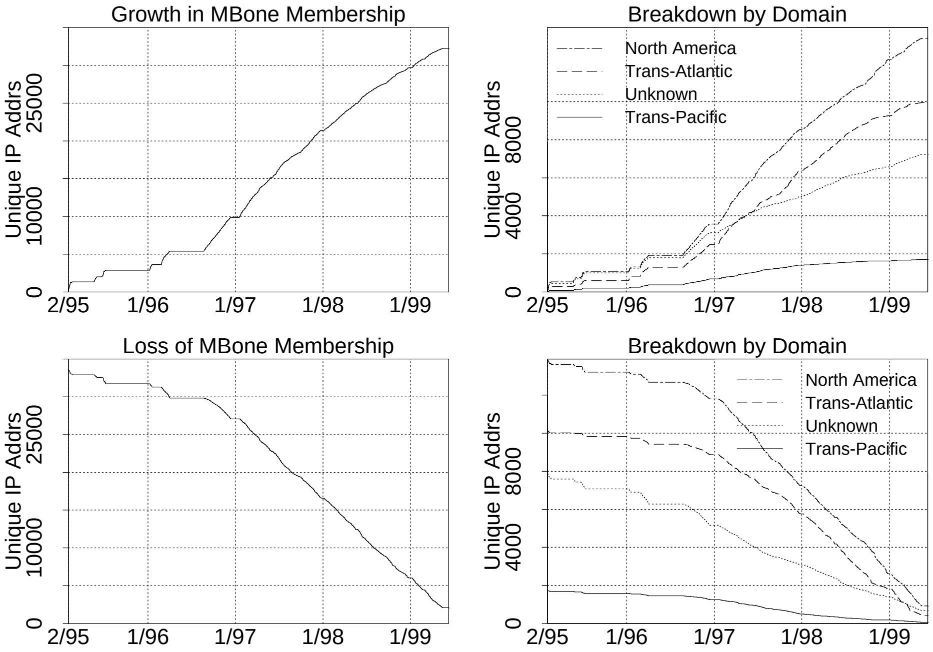

Figure 3: Growth and loss of MBone users.

is that much of this interest was generated by the IETF being held in Munich, Germany. Overall, the diversity of the MBone continues to grow at a relatively healthy rate.

Using these results to understand why the loss of membership has been so dramatic is the topic of Section 5. However, one reason, related to Figure 3, can be described now. Results show there are a large number of hosts whose IP address cannot be resolved into a name. Closer examination has revealed two possibilities. Either a host does not have an assigned name, or it was once assigned but no longer exists. By running a series of *ping* tests we have been able to verify that most of the hosts no longer exist. The question then becomes why these hosts constitute such a large percentage of MBone participants. Our reasoning is that some of the events broadcast over the MBone (including IETFs and most conferences) use temporary network address ranges. Such events tend to generate significant activity in the local, temporary network. During an event, these hosts will be reachable and have valid DNS entries but after the event they will no longer exist. Unless these addresses are later reused, these hosts will be considered lost.

4.3 MBone Lifetime Statistics

Our results with respect to MBone population lead us to investigate the lifetimes of multicast IP addresses more closely. Figure 4 shows a graph of the lifetime of each unique IP address plotted on a time scale. Each point represents the lifetime of the IP address and is positioned on the x-axis according to when the host was first heard. The left graph shows lifetime results for the entire collection period. The graph on the right shows an expanded view of the 1998 data. Both graphs have lines representing the maximum lifetime possible given when the IP address was first heard. For example, an address first heard on January 1, 1998 could only have a maximum lifetime of slightly more than one year. These two graphs show that IP address lifetimes vary widely, but both graphs also show a small concentration of addresses that have been active for close to the maximum time.

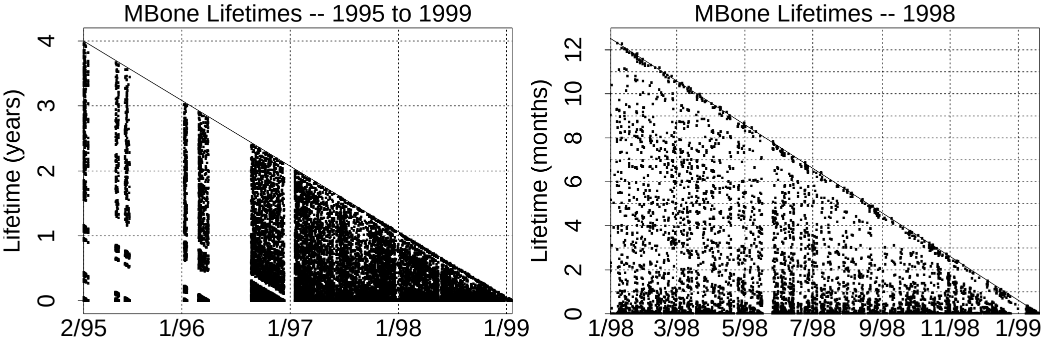

Figure 4: Lifetime for each MBone IP address.

A new question raised by these results is whether newer MBone participants are participating more or less than older participants. With these results, we cannot determine whether they are participating *more* but we can determine whether they are remaining active longer. Figure 5 shows the number of new IP addresses observed per day and the average lifetime of the set of new MBone IP addresses per day as a percentage of the maximum possible lifetime. One hypothesis is that if the quality of the MBone has improved, newer participants will be more likely to remain active for a longer period of time. However, such a hypothesis is not supported by the results in Figure 5. In fact, there is no discernible trend in the results. Old MBone participants are active in the MBone for the same percentage of their maximum lifetime as new participants.

Finally, one explanation for why average lifetime as a percentage of maximum lifetime is so low is because the distribution of lifetimes varies widely. The data set includes some very short-lived participants. Therefore, the average is low. Finally, while there are a number of IP addresses that participate on a continuous basis, they represent only a fraction of the total.

We now remove the time element and show a simple value-sorted series of IP address lifetimes. Figure 6 shows the distribution of these values both with a linear and logarithmic y-axis. What is important to observe is that the overall distribution has a very heavy tail. There are an unusually

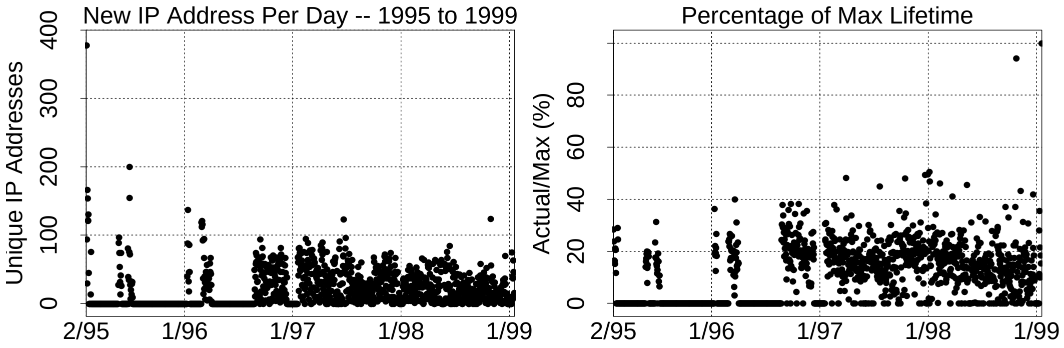

Figure 5: New IP addresses per day and the ratio of IP address lifetime to maximum lifetime.

large number of IP addresses that have been heard from once and only once. For more than 10% of the addresses we have received only one RTCP packet. This is unusual because it means a user joined a group, immediately quit, and then never joined again. For these hosts, there was hardly a chance to receive any content before quitting. Quite possibly there are other reasons, including routing problems or mis-configurations, which allowed one or a few packets to be received but no more.

Looking beyond the very short-lived lifetimes, results do not improve significantly for larger quantiles. For example, 25% of MBone IP addresses have a lifetime of less than 10 minutes, and 50% have a lifetime less than one week. Combined with other results, the generally short lifetimes suggest many users have tried multicast but found the performance and applications rather uninteresting.

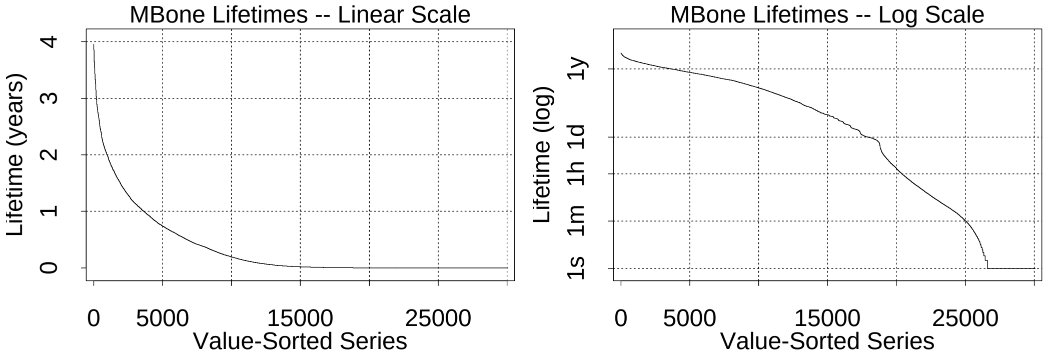

Figure 6: Lifetime values for IP addresses plotted on a linear and logarithmic y-axis.

4.4 SDR Session Characteristics

We now turn our attention to an analysis of *sdr* sessions. In processing the data files, we found that there were 1,611 different session names using 4,825 unique <*IP_mcast_address, UDP_port*> pairs. These pairs covered 4,407 unique IP multicast addresses (excluding the UDP port). Given that we do not have more detailed information about sessions, it is difficult to draw many quantitative conclusions. For example, with simple division, we estimate that if there are 1,611 sessions and 4,825 unique <*IP_mcast_address, UDP_port*> pairs, that each session used slightly more than 3 addresses. This tends to be an over-estimation since we cannot distinguish between sessions using the same name. For example, we cannot determine whether an audio session named "test" and a video session named "test" are two separate sessions or two components of a single session.

One noteworthy statistic occurs when the port number is removed. Over the 4.5 year period, 418 multicast addresses were used more than once. If two groups used the same address at the same time, regardless of port number, the two groups would have over-lapping router forwarding state. As a consequence, all sources would be received by both groups' receivers. However, these type of collisions tend not to be observable by users because of the different port numbers. Collisions should be unlikely, especially since *sdr* chooses a random address not already in use. Upon further investigation, we found only two instances where concurrent sessions used the same address. Since collisions are highly unlikely, we looked to explain why some addresses seemed to be used repeatedly but without colliding. The likely explanation is that some users tend to manually enter a "favorite" address instead of allowing *sdr* to randomly choose. Even though a particular address is used frequently, collisions are avoided because the user will not knowingly create them.

Turning to more qualitative results, Figure 7 shows various characteristics for each unique <*IP_mcast_address, UDP_port*> pair. The upper-left graph shows the growth of multicast addresses. This graph was constructed using the same technique as for computing growth of unique IP addresses. The upper-right graph shows the lifetime of each session. A session is considered to be active as long as the session is advertised through *sdr* or there are active group members. This tends to inflate lifetimes since many users do not leave a group once the advertised session is gone. This may seem like strange behavior, but it is quite common. One unusual example is when IETF groups have participants for several *weeks* after the event has ended.

Figure 7 also includes results for the total number of joins, talk periods, unique receivers, and unique transmitters. The results suggest that many groups are short-lived and are used only for test purposes. Almost 50% of all groups had no receivers, and almost 80% had less than 10 receivers. Results for transmitters show the same trend but the magnitude is lower given there are generally

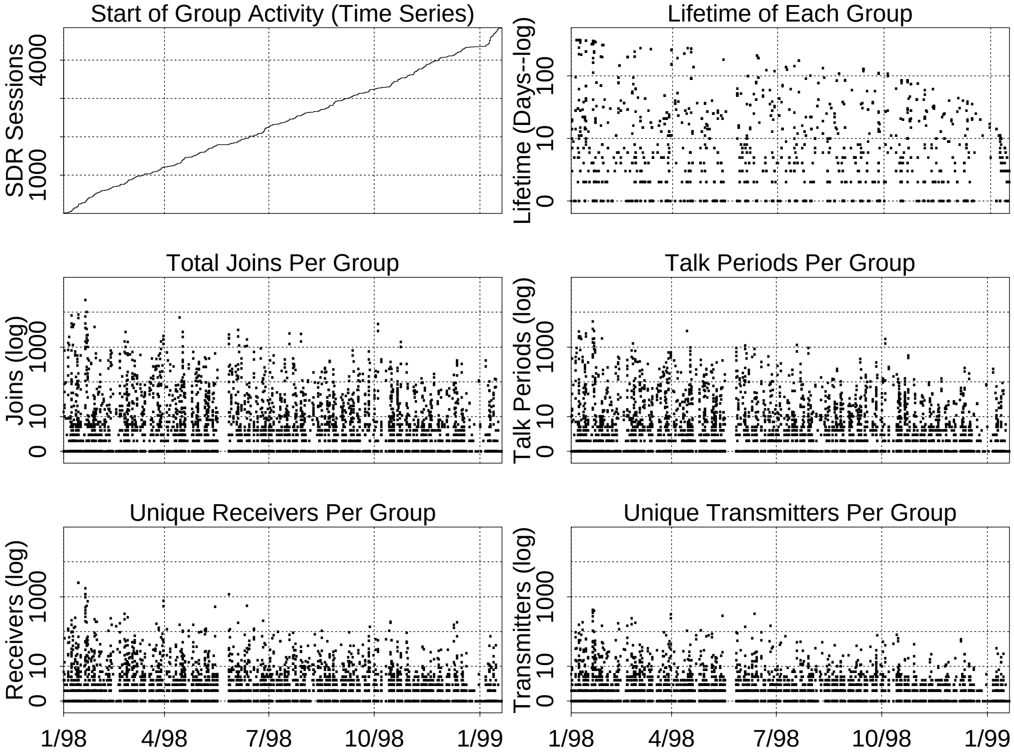

Figure 7: Statistics for each <*IP_mcast_address, UDP_port*> pair heard through *sdr*.

many fewer transmitters. At the other end of the spectrum, for the few groups that are popular, use statistics are very high. Some sessions have long lifetimes, thousands of participants and even see thousands of unique IP addresses. For these sessions there is a wide variety in the number of transmitters. Some of these popular groups, as with the case of the Interactive Multimedia Jukebox (IMJ)[17], have only a few unique, but frequent transmitters. Other sessions have large numbers of unique transmitters. For example, "Places Around the World" has had 421 different transmitters over its lifetime.

5 The "State of Multicast"

The results presented in this paper suggest that the "state of multicast" is not particularly good. The multicast community has worked hard to identify problems and develop solutions. In some cases, solutions are already being deployed. Based on our results, we believe the key problems include:

Overhead in Joining the MBone. Getting an MBone connection requires a significant amount of effort. Potential users have two choices. First, the traditional approach is to find someone willing to host a "tunnel". This process is not particularly easy and can be further complicated if a user does not have significant knowledge about multicast. The second choice is to find an ISP that offers multicast as a service. The problem is few ISPs currently offer a true multicast service.

Poor Session Quality. Many sessions today are simply unusable because of poor quality. Other researchers have reported on this problem[10, 9]. Our own measurements, using a tool called *mhealth*[18] confirm these results. One typical result is shown in Figure 8. These results, from an IETF meeting, show that on some occasions more than 50% of the receivers had loss rates greater than 25%.

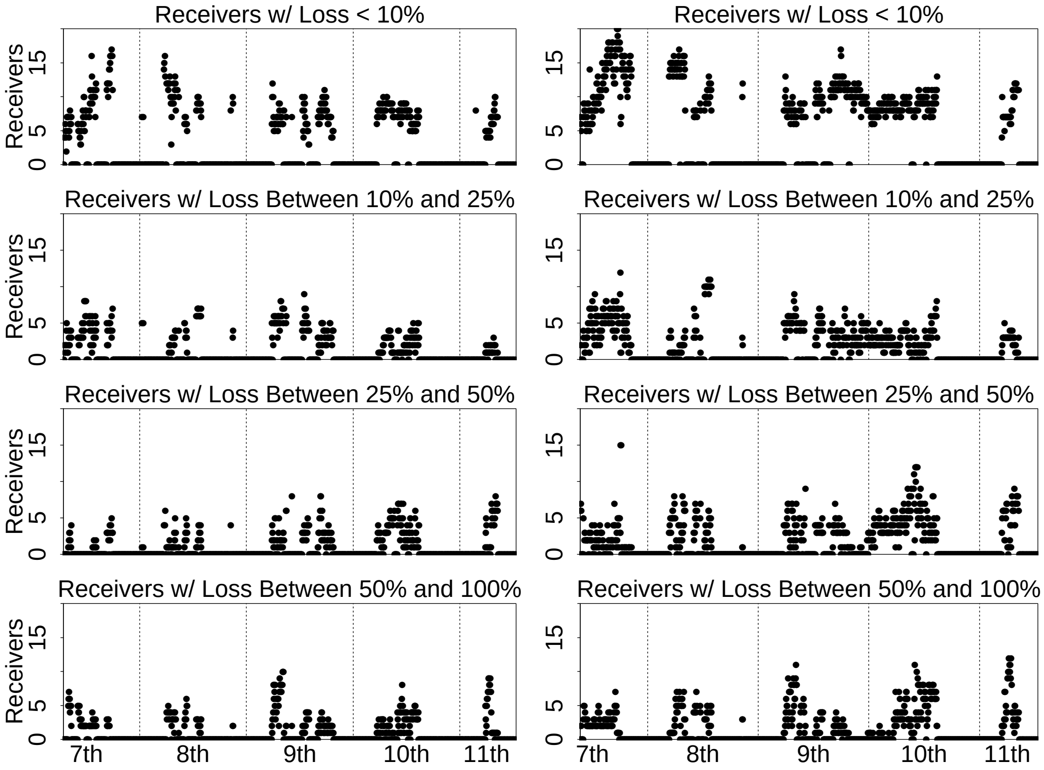

Figure 8: Loss quantiles for data collected during the 43rd IETF in Orlando, Florida.

Large Bandwidth Requirement for Users. The MBone was developed and used primarily by researchers who have large bandwidth capacities. This is especially the case when compared to users who connect to the Internet via a dial-up service. Because a large percentage of Internet users have dial-up connections, very little MBone content can be seen by most Internet users. A number of solutions are available including layered video encodings, and RTP/UDP/IP header compression.

Lack of Compelling Content. As our results show, the most popular applications are real-time broadcasts of IETF meetings, NASA programs, IMJ programming, and various radio stations. The reason why better content is not available is the classic chicken-and-egg problem. Without content, there is no push for deployment, and without a suitable infrastructure there is no customer market to motivate commercial content. This problem is less technical and is being addressed by groups like the IP Multicast Initiative (IPMI).

Lack of Motivation to Deploy Multicast. There are three groups who are impacted by multicast: content providers, network service providers, and end-users. Creating interest in multicast among members of any of these groups is likely to be difficult. Users do not care how they get content as long as they get it. Network providers have a difficult time understanding how to price multicast to create a profitable service. Content providers are the most motivated group but unless multicast makes financial sense, they will not use it. The problem is rooted in the fact that multicast has such a small installation base. A multicast-based service will only be usable by a small percentage of customers.

Of the items above, the recurring needs are more multicast-savvy engineers, better protocols, and more widespread deployment. There are both short-term and long-term solutions being discussed in the IETF. In the short-term the direction seems to be towards developing and deploying inter-domain multicast routing protocols. Under investigation is a version of the Border Gateway Protocol (BGP) that supports exchange of multicast routes[19], and a protocol, called the Multicast Source Discovery Protocol (MSDP)[20], to exchange information between domains running sparse mode protocols. One long-term solution is the Border Gateway Multicast Protocol (BGMP)[21] though other protocols are under consideration. Additional information on the evolution and direction of multicast can be found in [22].

6 Summary

In this paper we have presented the results of our study on the "state of multicast". We have attempted to answer questions about (1) how the MBone is used, (2) how multicast deployment has progressed, and (3) what barriers exist for the continued deployment of multicast. More specifically, we have analyzed four years of data and examined (1) characteristics about MBone group members, including group size and number of joins, (2) MBone population, as measured by unique IP addresses and domain name information, (3) statistics about the lifetimes of MBone users, and (4) the characteristics of *sdr* sessions. The key conclusions of our research are:

- The MBone is growing slowly. Not only is growth slow but there is a relatively high attrition among existing users.
- Use in the MBone is not widespread. A small percentage of users represent almost all of the MBone's activity.

- Only a small percentage of multicast groups contain significant levels of activity. There are only a handful of long-lived, popular groups.

References

[1] C. Diot, B. Lyles, B. Levine, and H. Kassem, “Requirements for the definition of new IP-multicast services,” *IEEE Network*, January/February 2000.

[2] S. Deering and D. Cheriton, “Multicast routing in datagram internetworks and extended LANs,” *ACM Transactions on Computer Systems*, pp. 85–111, May 1990.

[3] S. Deering, “Host extensions for IP multicasting.” Internet Engineering Task Force (IETF), RFC 1112, August 1989.

[4] S. Casner and S. Deering, “First IETF Internet audiocast,” *ACM Computer Communication Review*, pp. 92–97, July 1992.

[5] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, G. Liu, and L. Wei, “PIM architecture for wide-area multicast routing,” *IEEE/ACM Transactions on Networking*, pp. 153–162, Apr 1996.

[6] J. Moy, “Multicast extensions to OSPF.” Internet Engineering Task Force (IETF), RFC 1584, March 1994.

[7] H. Schulzrinne, S. Casner, R. Frederick, and J. V., “RTP: A transport protocol for real-time applications.” Internet Engineering Task Force (IETF), RFC 1889, January 1996.

[8] B. Mah, “Measurements and observations of IP multicast traffic,” Tech. Rep. UCB/CSD-93-735, University of California at Berkeley, March 1993.

[9] M. Yajnik, J. Kurose, and D. Towsley, “Packet loss correlation in the MBone multicast network,” in *IEEE Global Internet Conference*, (London, ENGLAND), November 1996.

[10] M. Handley, “An examination of MBone performance,” Tech. Rep. ISI/RR-97-450, Information Sciences Institute (ISI), University of Southern California (USC), January 1997.

[11] D. Massey and B. Fenner, “Fault detection in routing protocols,” in *International Conference on Network Protocols (ICNP)*, (Toronto, CANADA), November 1999.

[12] K. Almeroth and M. Ammar, “Multicast group behavior in the Internet’s multicast backbone (MBone),” *IEEE Communications*, vol. 35, pp. 224–229, June 1997.

[13] B. Huffaker, K. Claffy, and E. Nemeth, “Tools to visualize the internet multicast backbone,” in *Proceedings of INET ’99*, (San Jose, California, USA), June 1999.

[14] K. Almeroth, *Multicast Group Membership Collection Tool (mlisten)*. Georgia Institute of Technology, September 1996. Available from http://www.cc.gatech.edu/computing/Telecomm/mbone/.

[15] D. Waitzman, C. Partridge, and S. Deering, “Distance vector multicast routing protocol (DVMRP).” Internet Engineering Task Force (IETF), RFC 1075, November 1988.

[16] W. Fenner and S. Casner, “A ‘traceroute’ facility for IP multicast.” Internet Engineering Task Force (IETF), draft-ietf-idmr-traceroute-ipm-*.txt, August 1998.

[17] K. Almeroth and M. Ammar, “The interactive multimedia jukebox (IMJ): A new paradigm for the demand-driven delivery of audio/video,” in *World Wide Web Conference (WWW)*, (Brisbane, AUSTRALIA), April 1998.

[18] D. Makofske and K. Almeroth, “MHealth: A real-time graphical multicast monitoring tool for the MBone,” in *Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, (Basking Ridge, New Jersey, USA), June 1999.

[19] T. Bates, R. Chandra, D. Katz, and Y. Rekhter, “Multiprotocol extensions for BGP-4.” Internet Engineering Task Force (IETF), RFC 2283, February 1998.

[20] D. Farinacci, Y. Rekhter, P. Lothberg, H. Kilmer, and J. Hall, “Multicast source discovery protocol (MSDP).” Internet Engineering Task Force (IETF), draft-farinacci-msdp-*.txt, June 1998.

[21] S. Kumar, P. Radoslavov, D. Thaler, C. Alaettinoglu, D. Estrin, and M. Handley, “The MASC/BGMP architecture for inter-domain multicast routing,” in *ACM Sigcomm*, (Vancouver, CANADA), August 1998.

[22] K. Almeroth, “The evolution of multicast: From the MBone to inter-domain multicast to Internet2 deployment,” *IEEE Network*, January/February 2000.