Multicast Group Behavior in the Internet's Multicast Backbone (MBone)*

Kevin C. Almeroth
Mostafa H. Ammar

Networking and Telecommunications Group
College of Computing
Georgia Institute of Technology
Atlanta, Georgia 30332-0280
{kevin, ammar}@cc.gatech.edu
(404)894-3292
(404)894-0272 (fax)

November 5, 1999

Abstract

The Multicast Backbone (MBone) is a network overlaying the global Internet designed to support multipoint applications. In this paper, we first give an overview of the development and architecture of the current MBone. One important characteristic of the MBone is its reliance on IP multicast which allows receivers to join and leave groups asynchronously. We describe The *Mlisten* data collection tool that was created to provide a mechanism for capturing information about when members join and leave a multicast group. Using data collected with *Mlisten*, we present statistics about some of the MBone sessions we monitored. Results are provided for key parameters including multicast participant inter-arrival times and participant durations, multicast tree size and characteristics, and inter-session relationships. By collecting data about MBone usage we can improve our understanding of how multicast communication and real-time protocols are being used today.

*The tool to collect group behavior, *Mlisten*, is at http://www.cc.gatech.edu/computing/Telecomm/mbone/

1 Introduction

The Multicast Backbone (MBone) is a network integrated into the global Internet[1]. It provides one-to-many, best effort delivery of real-time multimedia data between MBone-connected hosts. The MBone has served as a testbed for the development of multicast protocols, real-time protocols, and group conferencing tools. As the MBone has evolved over time, it has come to provide a relatively stable platform for the world-wide delivery of seminars, conferences, working group meetings, and even entertainment. As such, the MBone represents a good first estimate of how large-scale multicast services will be used when the Internet becomes entirely multicast capable.

Of particular interest, and the focus of this paper, is the real use of multicast communication (over the MBone infrastructure) to broadcast events on a very large scale. Insight is gained by observing and characterizing the behavior of MBone session members. Collecting and analyzing the join/leave behavior is of great benefit in understanding how any future networking infrastructure with multicast and real-time capabilities will be used. More fundamentally, basic statistics and an understanding of changes in group membership and characteristics are useful for developing and testing more efficient multicast routing and resource reservation algorithms and protocols. Furthermore, the ability to use collected data to build workload traces and models is a useful result of this research.

The MBone behavior we are most interested in is based on the scheduling of world-wide MBone "events" which are typically announced ahead of time in a global session directory. Using information in this session directory, interested people can join multicast groups, and receive real-time audio, video and/or other media streams. Our primary interest lies in monitoring: (1) the rate at which users join a session and the duration of their membership (temporal statistics), and (2) the geographical location of group members and their distance from the source (spatial statistics). To this end we have developed *Mlisten*, a tool to monitor MBone sessions and collect behavior data.

This paper documents our specific efforts in two areas. First is the capture of accurate MBone usage data. This involves building a collection tool and pre-processing raw data to eliminate errors and inconsistencies. Second is our effort to characterize and analyze MBone sessions both temporally and spatially.

We are not aware of any similar work aimed at collecting MBone multicast group membership or using this data to generate statistics or models. Nevertheless, our work is related to a body of work aimed at understanding the operation of the MBone. Such efforts include debugging tools[2, 3] that have been developed to analyze network problems. These tools try to isolate problems by determining routes and the status of links in the MBone. One of these tools[3] was used in our study to approximate the set of MBone links connecting a source to a set of destinations. In the area of traffic analysis, one effort evaluates the volume and types of traffic carried on the MBone[4].

The remainder of this paper is organized as follows: Section 2 presents an overview of the evolution of the MBone and its operation. Section 3 describes the *Mlisten* tool, and the collection and processing steps. Section 4 is a description of the types of analysis conducted, and reports statistics and analysis of data collected for some MBone sessions. The paper is concluded in Section 6.

2 Conferencing on the MBone

The MBone has been in a constant state of development since its inception in 1992 when 20 sites were connected together to receive the March meeting of the Internet Engineering Task Force

(IETF)[1]. That first audio conference, carried over the Internet, allowed a few members from all over the world to hear what was being said in San Diego. In addition to the conferencing software, the most significant achievement was the deployment of the first multicast-capable routers. These first 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 via one logical link but may have crossed several Internet routers. Once a multicast packet was received at a tunnel endpoint, it could be broadcast on a local network and received by other members. Routing between multicast-capable routers was provided using the Distance Vector Multicast Routing Protocol (DVMRP)[5].

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)[6] are also being used. Improvements are not limited to just multicast routing but also include protocols like the Real-Time Protocol (RTP)[7], which make the Internet much better at handling delay sensitive applications. The trend and eventual goal of the MBone is to facilitate the efficient and effective transmission of real-time multimedia data over the Internet.

Since the first audio conference in 1992 the MBone has seen the development of a wide range of new applications using audio, video, whiteboard, and text as media. MBone *sessions* are a combination of these media. Information about these sessions are periodically transmitted across the MBone on a well-known multicast address. Using the Session DiRectory (SDR) tool [1], MBone users can obtain a list of advertised sessions. Through SDR, an MBone user can choose from this list and launch the MBone tools (including audio, video, whiteboard, or test) required to receive the component streams of a session. For each of these tools there is a multicast group which the user *joins* when the tool is started. Once part of the group, members will receive group transmissions and they can actively participate or simply listen. Joining a group means that a user must be *grafted* into the multicast tree. The existing multicast routing protocols are capable of seamlessly providing both join (graft) and leave (prune) functions.

Communication in an MBone session takes place using an advertised multicast IP address and two pairs of UDP ports. The first pair of UDP ports, the *data ports*, are used to transmit audio data; the second pair, the *control ports*, are used to transmit *group membership information*. Each pair of ports consist of an advertised port number used to receive packets (either control or data), and a local port number used to transmit packets. Each of the MBone tools uses a session's control address to periodically broadcast identifying information about itself and the user. This information is used by the tools to know about the size of the group and characteristics about the other members.

Our efforts are based on the desire to know when group members join and leave any of the MBone sessions. We concentrate specifically on collecting data for the audio and video components of globally and near-globally transmitted sessions for two reasons: (1) we have only attempted to collect data from a site at Georgia Tech and so are limited to the observation of sessions that reach us, and (2) we believe group membership data from these sessions is sufficiently representative to achieve a basic understanding of behavior.

[1] Available from ftp://cs.ucl.ac.uk/mice/sdr/

3 Collecting Multicast Group Membership Data

Accurate data collection involves two steps: (1) capturing member join/leave data, and (2) processing to correct timing and other inaccuracies. Collection is obviously done in real-time, while processing is typically done before the data is analyzed. An overview of each step is described next.

3.1 Capturing Group Membership

Our collection tool, *Mlisten*, continuously monitors the well-known multicast address used to advertise MBone sessions. For each session, *Mlisten* can join the audio and video groups and collect information about control and data packets. When a packet arrives, *Mlisten* identifies the packet's sender, the MBone session, and the time of arrival. This information is used to maintain a list of active group members in each session.

At periodic intervals, *Mlisten* identifies any group members who either stopped transmitting data packets or who left the group and stopped transmitting completely (including control packets). If no *data* packets are received for a threshold period of two minutes, a "talk period" is assumed to have ended. Similarly, if a group member stops sending packets completely, the member is assumed to have left the group. For MBone sessions, if there are no group members or session advertisements during a four hour period, the session is assumed to have ended.

For each group member who leaves a particular group, a record is written to a *log file* and includes the following information: (1) the type of record (receive, transmit, or session), (2) the session multicast address and port number, (3) the host IP address and UDP port number, (4) the date and time the first packet was received, (5) the membership duration for receivers or talk-spurt duration for transmitters, and (6) the number of packets received.

3.2 Log File Processing

The log file generated by *Mlisten* may not accurately reflect true group membership or behavior. Two reasons are: (1) unreliable packet delivery, and (2) anomalous MBone behavior. These are briefly described below. For a complete description of why these inaccuracies occur and how they are corrected, see [8].

Unreliability. Unreliability causes packet loss and delivery jitter which can result in some group members appearing to leave a session, but not actually doing so. Multiple records for the same user showing several group join/leave actions may actually be only one join/leave action that was split because of packet loss. Combining split records can be accomplished using the sender's UDP port number which changes between leaving one group and joining another.

Abnormal MBone Behavior. Abnormal behavior describes a set of MBone phenomenon which are caused by software bugs, performance experiments, or any of a variety of infrequently occurring anomalies. Since the goal of this research is to model normal use, we try to eliminate obvious cases where usage is atypical and might skew quantitative analysis. Using a combination of inspection and automated processing we can process the log file to eliminate abnormal behavior. A complete description with examples can be found in [8].

4 Analysis Methodology

MBone data has been collected since the beginning of 1995. Over the course of these nearly two years we have improved the range of data collected. The first data sets consisted of only the membership information for a single audio session. We can now collect information about all audio and video sessions including information about the frequency and duration of session transmissions. Our statistical results and analysis are based on as much data as was collected at the time. Additional results continue to become available as ongoing data sets become available. In this paper we only present a very small subset of available results. Other results are reported in [8, 9].

Results of our analysis are divided into four categories: (1) temporal, (2) spatial, (3) inter-session, and (4) intra-session. A description of each is given below.

- **Temporal results** consist of group join inter-arrival time and membership duration. These two parameters are sufficient to describe activity in one or more groups. Observed results are then fitted to well known distributions to create multicast group activity models.
- **Spatial results** focus on the size of the multicast tree, average distance from the set of receivers to the source, and a geographical breakdown of group members. We also compare the number of packet hops needed to reach the set of destinations using both unicast and multicast.
- **Inter-session results** use data from all concurrent MBone sessions. These results capture the "surfing" behavior of MBone users as they move among sessions in a way similar to how TV watchers surf between channels.
- **Intra-session results** show the relationship between media within a specific session. In particular we compare the behavior of members in audio-based vs video-based multicast groups.

4.1 Temporal Results

In this section we present the temporal results for three data sets: (1) the NASA STS-63 shuttle audio session from February 3, 1995, to February 11, 1995, (2) a UCB Multimedia Lecture Series audio session on February 17, 1995, and (3) all audio sessions collected from February 26, 1996 to March 10, 1996. Each of these data sets represent a class of sessions, and as such, are similar to the behavior exhibited in many of the MBone sessions. Figure 1 contains two sets of graphs with three parts: (1) the number of receivers as measured at five minute intervals, (2) the inter-arrival time between group member arrivals, and (3) the membership duration plotted (on a log scale) at the time of the join action. The all audio session data is not plotted because of its similarity to the graphs in Figure 1. However, the average inter-arrival and duration statistics are given in Table 1.

Inter-Arrival Data Analysis. For long running sessions (like any of the shuttle sessions) or continuous broadcasts (like the old IMS World Radio Network or CBC Newsworld), daily and weekly variations in turnover and participation are quite pronounced. The highest periods of activity occur during the day. Day time hours are based on time zones in the United States because more than half of the receivers are in the United States (see Section 4.2). At night, activity (as measured by the inter-arrival time between new group members) and group size are much lower. Weekends usually see the least amount of activity. For example, the average inter-arrival time for the entire

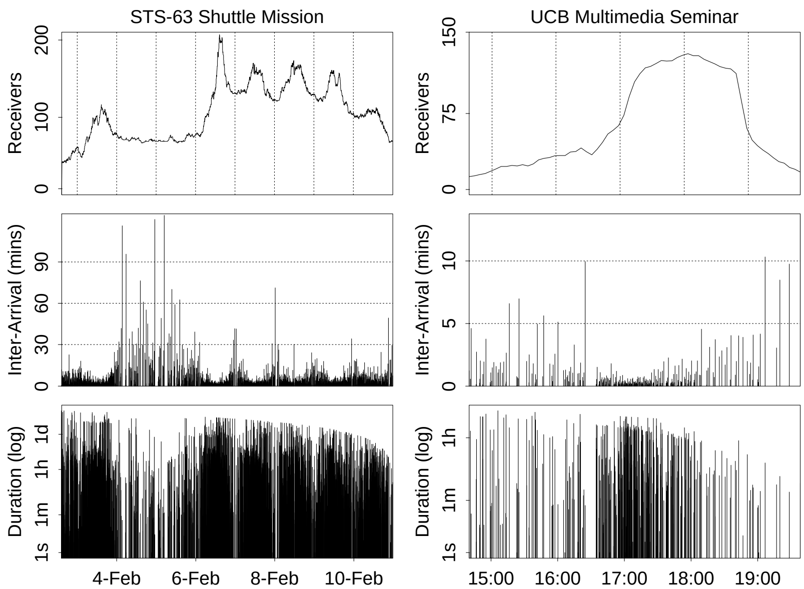

Figure 1: STS-63 Shuttle and UCB Seminar temporal statistics.

STS-63 session was 2.5 minutes, but the weekday average was 22.1 seconds; the weeknight average was 3.9 minutes; and the weekend average was 12.6 minutes. While these conclusions may seem qualitatively obvious the statistics are of great use in accurate multicast group models.

For shorter sessions lasting less than a day or two, turnover and participation are strongly correlated to the advertised start and end time of transmissions. The UCB seminar shown in the second part of Figure 1 was a talk that started at 17:30, lasted about an hour, plus another half hour for questions. Most short sessions exhibit an increasing amount of group join activity up to just after a program's scheduled start time.

Duration Data Analysis. The membership duration was relatively consistent and well behaved over the lifetime of observed sessions. One interesting observation is that group members typically either join for a very short period of time or stay for the entire session even if the session lasts for many days. For the STS-63 session, the average membership duration was 5 hours, but the median duration was only 6.5 minutes. Furthermore, during the STS-63 session, each host averaged 4.2 join requests suggesting that users would join for a couple minutes several times over the session lifetime.

For a shorter session like the UCB seminar, the average membership duration was much shorter (46 minutes) but the median duration was similar to long sessions (7 minutes). This suggests that there is a much higher turnover for shorter sessions. Finally, there appears to be a relationship between session content and membership duration. During a session, if there is no one transmitting, there is a noticeable decrease in average duration.

Temporal Data Fitting. One of the most important aspects of our analysis is fitting collected data to well known distributions. The most significant challenge in trying to model behavior is finding periods of stable activity, i.e. periods without significant increases or decreases in group membership. Since behavior is greatly affected by session content, time of day, and time of week variations, we chose to develop models based on activity during 8 to 10 hours of a typical week day.

Session Names	Join Requests		Average	Average
	Number	Percent	Inter-Arrival	Duration
NASA: STS-75	4230	27.55%	7.51 mins	4.3 hrs
IMS: World Radio Network	1998	13.01%	22.22 mins	1.7 hrs
IETF 1 - Audio	1343	8.75%	22.71 mins	2.6 hrs
IETF 2 - Audio	1234	8.04%	38.87 mins	2.4 hrs
Radio Free Vat	1179	7.68%	38.94 mins	49 mins
MBone Audio	839	5.46%	53.13 mins	4.4 hrs
FreeBSD Lounge	712	5.33%	67.57 mins	35 mins
Other 3 IMS Sessions	1314	8.56%	56.00 mins	2.0 hrs
9 Short-Lived Sessions	2507	16.33%	18.11 mins	42 mins
Total	15356	100.0%	1.24 mins	2.5 hrs

Table 1: Statistical breakdown for all sessions.

The results in Figures 2 for the STS-63 data show that inter-arrival data fits extremely well except for one or two points at the high end. For membership durations an exponential function does not fit well because the tail of the duration data is very long. While a majority of group members may join for only a few minutes or hours, some members are part of the group for many days. For membership duration data, we use a Zipf distribution which works when a large percentage of (duration) samples are concentrated at the beginning of the range while the remaining percentage are widely dispersed over the remainder of the spectrum[10]. This is exactly the type of behavior exhibited.

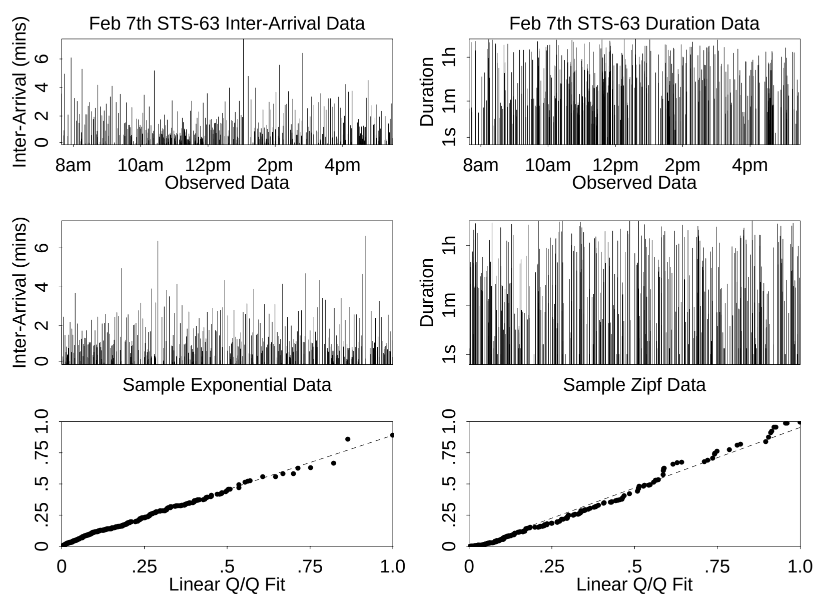

Figure 2: Fit of STS-63 data.

We have used the fitting technique described above to model a significant number of MBone sessions. In general, an exponential function works well for the inter-arrival data regardless of the type of session (short vs long vs an aggregate of sessions). When modeling duration data, an exponential function works well for short sessions while a Zipf distribution works well for longer sessions.

4.2 Spatial Results

In this section we present the spatial results for only one data set: the NASA STS-63 shuttle audio session. However, we also include analysis results from other sessions. In general, we find that many sessions exhibit similar characteristics or slightly different but predictable results.

Domain Breakdown. Using a DNS server, a domain name can be determined for each group member's IP address. Figure 3 shows a breakdown of the STS-63 session into three groups: (1) all hosts in North America, (2) all hosts connected to the MBone via trans-Atlantic links, and (3) all hosts connected via trans-Pacific links. Figure 3 shows that more than 50% of the receivers were in North America. The shuttle missions fall in the middle of the spectrum in terms of geographical distribution. Other, more general interest programs like the IMS World Radio Network have a much more even distribution while some programs are broadcast world-wide but have only local interest with respect to the source. A final observation is that activity in each of the three groups occurs during the day time hours but because each group is in a different part of the world, the daily activity is offset slightly.

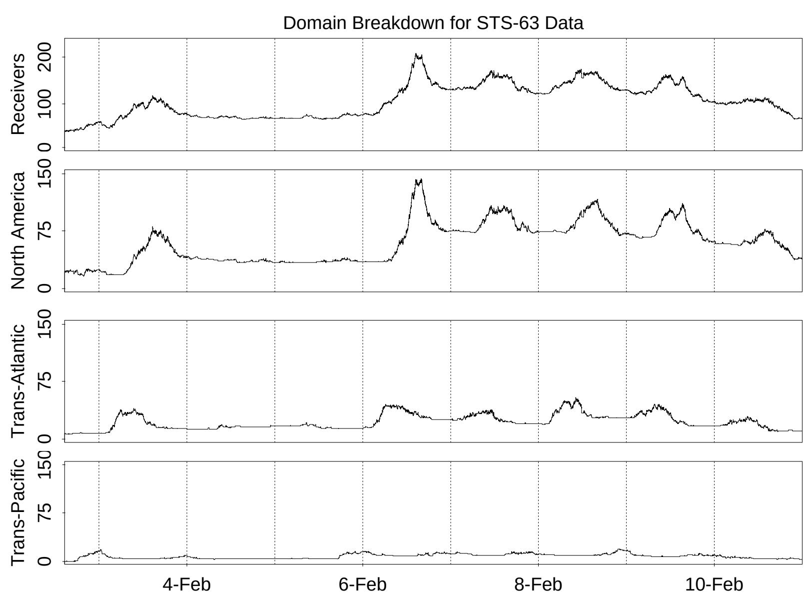

Figure 3: STS-63 domain breakdown.

Multicast Tree Hop Count. In computing the cost of reaching a set of destinations, we consider several types of trees. Our cost metric is that of packet-hops (first used in [11]). For example, a tree from a source to multiple destinations will have a cost in packet-hops equal to the number of links in the tree. The types of trees we consider are:

1. **MBone Tree with Observed Source**: We use the MBone topology and create a source to destination set tree. The routes used are roughly approximated using the MWatch tool[3].

The multicast hop count is compared to a unicast hop count which is computed assuming that the source transmits to each receiver individually (also known as the *separately-addressed packets* scheme [11]).

2. **MBone Tree with Georgia Tech Source**: In order to assess the sensitivity of tree costs to source location we repeat the above analysis but assume a host at Georgia Tech is the source. The tree is changed to reflect the change in source, with the group members remaining the same.
3. **Internet Tree with Georgia Tech Source**: MBone hops are really tunnels which can span multiple Internet hops. In order to estimate the tree cost in terms of Internet hops we use the *traceroute*[2] utility and assume a host at Georgia Tech to be the source[3]. To the extent that the results from Item 2 above show little sensitivity to the location of the source, then these estimates should be close to what we would have obtained using paths from the true source.

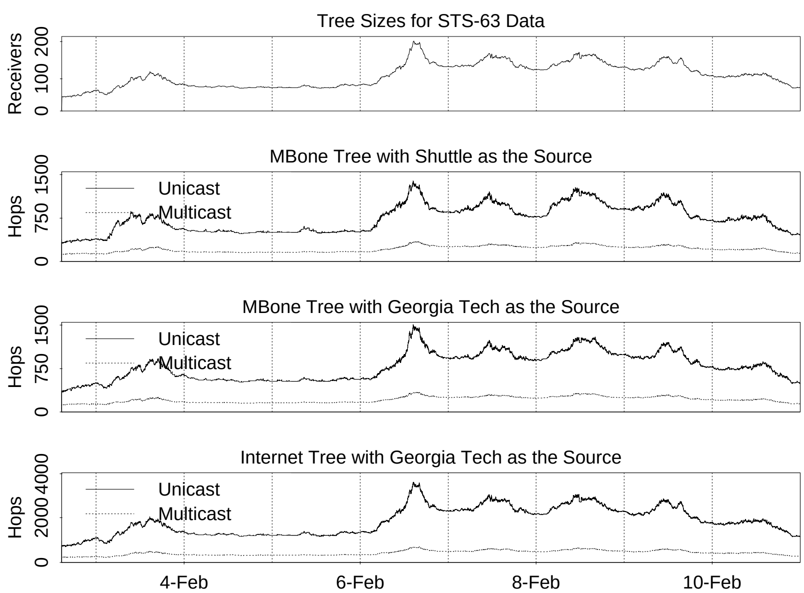

Figure 4: STS-63 multicast tree costs.

Figure 4 shows the tree costs for the STS-63 session. The results show that while there is a direct relationship between the number of group members and the number of unicast packet-hops, the number of multicast packet-hops remains nearly flat. Even when the number of group members increases, the number of packet hops increases only slightly. The distance between source and the set of destinations averages 7.1 hops but the cost in the multicast tree is, on average, only 30.1% of the unicast cost. Changing the source for the STS-63 session does not significantly affect the tree cost or the average distance. Using the Internet topology, the results are similar but the Internet multicast tree is larger, containing 2.3 times more hops then the MBone tree.

[2] Available from ftp://ee.lbl.gov/traceroute.tar.Z

[3] We are not aware of any technique that could have been used locally to determine Internet routes between two remote hosts.

4.3 Inter-Session Results

In one of our data sets we collected data from all audio sessions for a period of two weeks. Using this data, we analyzed the behavior of users in multiple groups, including what they do when they leave a session. Using the log data we can determine if, within a time interval of two minutes, a user leaving one session either joins another session or becomes *inactive*. For the data set summarized in Table 1, 70% of the hosts "surfed" to another session at least once. Furthermore, more than 50% of all join actions occurred within two minutes of the host leaving a previous group. These results suggest "session surfing" is a typical behavior. Table 2 shows exactly where members leaving one group go. The rows represent the group that a member leaves from and the columns represent where the member goes.

To → From ↓	STS-75	World Radio Network	IETF Chan 1	IETF Chan 2	Radio Free Vat	MBone Audio	Free BSD	Other IMS	Other Sessions	Leave MBone
New	27.2	12.8	18.1	11.0	6.8	6.0	3.8	5.4	8.9	N/A
STS-75	12.6	6.2	6.0	2.8	5.1	2.2	1.8	2.9	2.6	58.2
WRN	8.9	12.9	3.3	2.2	6.2	4.0	2.4	7.9	4.0	48.2
IETF-1	4.0	2.1	10.6	37.2	1.3	1.3	0.9	1.9	1.0	39.8
IETF-2	5.8	2.3	21.3	14.5	0.8	0.9	0.8	1.9	0.5	51.2
RFV	6.5	11.8	2.5	1.8	8.1	7.3	2.1	4.8	6.7	48.6
MBone Aud	7.6	8.0	4.5	2.3	10.7	9.9	2.9	3.9	5.8	44.4
Free BSD	8.3	4.3	6.7	0.3	2.4	1.6	12.1	9.9	6.7	47.7
Other IMS	6.6	13.7	2.6	2.2	5.4	3.2	2.4	29.6	5.1	29.2
Others	3.7	5.4	1.1	0.9	3.6	2.2	7.4	4.7	22.1	49.0

Table 2: Entries represent the percentage of members transitioning between sessions.

Two interesting observations are (1) some sessions are more tightly coupled and so see a much higher transition rate, i.e. IETF sessions, WRN and other IMS sessions, etc., and (2) sometimes users leave one group and within two minutes return to the same group; for example, 12.6% of STS-75 group members leave the STS-75 only to re-join it. Likely reasons for this include accidently closing then re-starting an audio tool, software crashes, or accidently starting multiple instances of a tool and then quitting all but one.

In addition to inter-session transition statistics we have also analyzed information about users who run more than one audio session simultaneously. Because of hardware limitations, it is not possible to actually listen to multiple sessions at the same time but a user can still be a member of multiple groups. Results show that for the all-session data 45% of all hosts were members of more than one group at one time or another. Furthermore, 17% of all group memberships occurred when the host was already a member of another group. These results suggest simultaneous sessions occur, but not frequently.

4.4 Intra-Session Results

Figure 5 shows the number of receivers in audio-based and video-based multicast groups for two data sets collected in 1996. The results show that there is a strong correlation between the behavior of members in audio and video groups. The large jump in receivers starting on September 16, 1996, is a result of the start of a new NASA shuttle mission. While there is significantly more audio-group members before the 16th, after the 16th, group membership was almost evenly distributed. This suggests that some people watching the shuttle didn't listen to the audio but simply watched the video. In the second data set there is significantly less activity and no major/popular session. Participation in video sessions dropped to almost half that of the audio sessions. In summary, when there is interesting content, there is significant correlation between audio and video sessions, but during other times audio seems to be the preferred media.

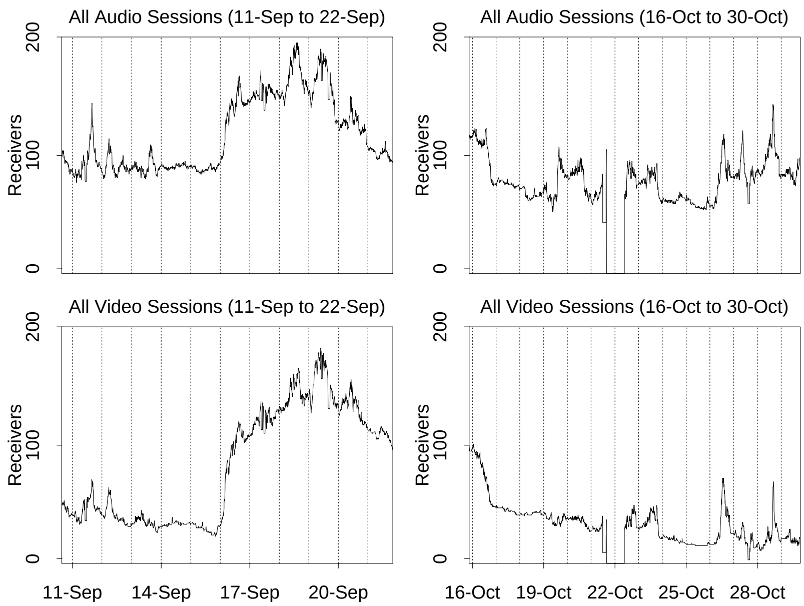

Figure 5: Comparison of behavior in audio and video groups.

5 Conclusions

There has been a significant amount of work geared towards developing mechanisms to provide network support for multicast communication. The work reported in this paper complements these efforts by providing a tool that can be used to collect and analyze MBone multicast sessions. The tool, called *Mlisten*, collects the join/leave times for multicast group members in MBone audio sessions. Using data collected for a set of single MBone sessions and also data collected for all sessions we perform analysis to produce (1) temporal results including group size, participant inter-arrival times and participant durations, (2) spatial results including multicast tree size, distance from receivers to the source, and a geographical breakdown of group members, (3) inter-session results in the form of "surfing" statistics, and (4) intra-session results comparing behavior and audio-based and video-based multicast groups. These results are extremely helpful in understanding multicast group behavior using both qualitative and quantitative measures. In addition to the conclusions drawn in this paper, results can be used to develop and realistically test the efficiency of multicast routing and resource reservation algorithms and protocols.

References

[1] S. Casner and S. Deering, "First IETF internet audiocast," *ACM Computer Communication Review*, pp. 92–97, Jul 1992.

[2] D. Agarwal and S. Floyd, "Tool for debugging internet multicast routing," in *22nd ACM Computer Science Conference*, (Phoenix, AZ), Mar 1994.

[3] A. Ghosh and P. Brooks, *MWatch 3.6.2*. University College Lonndon, Jun 1994. Available from http://www.cl.cam.ac.uk/mbone/index.html#Mrouted.

[4] B. Mah, "Measurements and observations of IP multicast traffic," Tech. Rep. UCB/CSD-93-735, University of California at Berkeley, Mar 1993.

[5] S. E. Deering and D. R. Cheriton, "Multicast routing in datagram internetworks and extended LANs," *ACM Transactions on Computer Systems*, vol. 8, pp. 85–110, May 1990.

[6] 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.

[7] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: A transport protocol for real-time applications," Tech. Rep. RFC 1889, IETF RFC, Jan 1996.

[8] K. Almeroth and M. Ammar, "Collecting and modeling the join/leave behavior of multicast group members in the MBone," in *IEEE Symposium on High Performance Distributed Computing*, (Syracuse, NY), pp. 209–216, Aug 1996.

[9] K. Almeroth and M. Ammar, "Characterization of MBone session dynamics: Developing and applying a measurement tool," Tech. Rep. GIT-CC-95-22, Georgia Tech, Jun 1995.

[10] G. Zipf, *Human Behavior and the Principle of Least Effort*. Reading, MA: Addison-Wesley, 1949.

[11] Y. Dalal and R. Metcalfe, "Reverse path forwarding of broadcast packets," *Communications of the ACM*, pp. 1040–1048, Dec 1978.