Characterization of MBone Session Dynamics: Developing and Applying a Measurement Tool *

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)894-3293
GIT-CC-95/22

June 9, 1995
Revised: March 1, 1996

Abstract

Much of our current understanding of the operational aspects and the network requirements of multicast communication derives from the MBone. As an experimental network that overlays the Internet, the MBone has served as the testbed for multicast protocols. This work has led to the development of audio and video conferencing tools that have been used extensively recently. This paper describes our efforts in developing and using a measurement tool, called *Mlisten*, for characterizing the dynamic behavior in MBone sessions. Such a characterization will be of great benefit in understanding how any future networking infrastructure with multicast and real-time capabilities will be used. We discuss the challenges in monitoring MBone session participation and report on a methodology for pre-processing observed data to make it better reflect the true behavior of members in an MBone session. We also report on the results of using *Mlisten* to determine temporal, geographical and resource usage characteristics of several recent MBone sessions.

*A prototype of the tools are available at "http://www.cc.gatech.edu/computing/Telecomm/mbone/".

1 Introduction

The capability to efficiently support multicast communication is becoming an indispensable feature of any networking infrastructure. Multicast communication is necessary as a means to provide a variety of communication services such as video and audio conferencing, computer supported collaborative work and wide-scale information dissemination.

Much of our current understanding of the operational aspects and the network support requirements of multicast communication[1] derives from the MBone[2]. As an experimental network that overlays the Internet, the MBone has served as the testbed for many ideas. Most prominent is the development of a set of audio and video conferencing tools[3, 4, 5, 6, 7, 8] that have been used extensively in the last few years to multicast conferences, talks, radio broadcasts, concerts, Space Shuttle missions, and working group meetings. Despite it being an experimental network, many of the sessions using the MBone are real and as such represent well how multicast services will be used when the Internet becomes entirely multicast capable. For this reason we feel that a characterization of MBone sessions will be of great benefit in understanding how any future networking infrastructure with multicast and real-time capabilities will be used. Such characterization can also be used to build realistic simulation models which are useful for testing multicast techniques such as routing, congestion control and resource reservation.

MBone "events" are typically announced ahead of time in a "session directory"[3] and interested people join the appropriate multicast group to participate in the particular session. Session members may also leave a session at will and possibly rejoin. Our particular focus in this paper is the understanding of this *dynamic member behavior* and how it is influenced by the type of session being transmitted. Additionally, we are interested in how network resource demands (particularly bandwidth) change over time as a result of this dynamism. To this end we have developed a tool, called *Mlisten* to monitor an MBone session and extract member behavior. Using *Mlisten*, we analyze a number of recent MBone sessions.

In this paper, we document our efforts to develop *Mlisten* and analyze the collected data. First, we discuss the challenges in monitoring MBone session participation from a remote site. Part of this activity includes pre-processing collected raw data to make it better reflect the true behavior of members in an MBone session. We found this pre-processing to be necessary because various conditions in the MBone like jitter, loss, and experimental usage can bias the observed data. Second, we report on the results of using *Mlisten* to determine temporal, geographical and resource usage characteristics of several MBone sessions.

We are not aware of any similar work that has aimed at characterizing the dynamics of session participation based on actual MBone sessions. Nevertheless, our work is related to a body of work that has aimed at understanding the operation and usage of the MBone. Such efforts include debugging tools[9, 10, 11] that have been developed to analyze network problems. Most of these tools work by recording statistics about the routes used in the MBone, and using this information to isolate faulty routers or links. One of these tools[10] was used in our study to determine the set of MBone links used to connect a source to a set of destinations. In the area of traffic analysis, only preliminary work has been done. In [12], several traffic traces were taken, and the volume and distribution of MBone traffic was presented.

The remainder of this paper is organized as follows: Section 2 gives an overview of the MBone and the audio conferencing tools used. Section 3 describes the collection of session membership data. Section 4 details the steps used to make raw data more effectively represent operational behavior. Section 5 contains the results obtained by characterizing several MBone sessions and section 6 discusses an analysis of MBone session multicast tree costs. The paper is concluded in Section 7.

2 Conferencing on the MBone

MBone Overview. The MBone is an experimental virtual network that provides the means for multicasting data to any number of connected hosts[2]. One of the primary objectives of the MBone is to test the real-time, one-to-many delivery of audio- and video-conferencing data. A wide variety of multimedia applications can theoretically be supported, but the large-scale application most commonly demonstrated is that of lecture- or program-style sessions. These sessions typically have only one source for a majority of the time, and interactivity on the receiver side is usually only a small portion of the session.

The MBone is a virtual network that overlays the Internet. Some links of the MBone network are actually unicast tunnels between multicast-capable routers. Each tunnel may actually consist of several Internet links. Routing in the MBone is performed using the Distance Vector Multicast Routing Protocol (DVMRP) [11, 1].

Even though the MBone is still an unstable, experimental network, the transmission of programs or conferences has been a daily occurrence for some time[13]. Information about conferences or other events is announced by periodically transmitted messages over the MBone. Using the Session Directory(SD) tool[3], MBone users can see what sessions are active, and find out about other scheduled MBone broadcasts. An MBone user can choose from the list of sessions displayed in the SD menu, and launch all the tools (including audio, video, or whiteboard) necessary to receive all the component streams of a session. When one of these tools is started, the user *joins* the particular session and becomes a member of the multicast group. Members can either actively participate and transmit to the group, or they can simply listen. In either case, starting any of the tools means that the user will become part of the multicast group and will receive the transmissions of the group.

MBone Audio Tools. MBone conferences can have multiple components including audio, video and whiteboard communication. In practice most sessions have at least an audio component and the behavior of the other components closely follows the audio conference dynamics. We focus, therefore, on studying the audio conferencing component in this paper. The two most common audio tools used in the MBone are VAT[4] and NeVoT[8]. When one of these applications is started, the host on which the tool is started joins a multicast group dedicated to the session. One result of this action is that the tool begins transmitting *name information* to all other group members. This name information is a one line text identifier specified by the user and can be anything, but is typically a person's name, affiliation, and/or e-mail address. All group members of a particular session transmit and receive name information which is used by both VAT and NeVoT to create and display a list of group members, and the status of each[1]. Name information is transmitted by each user for as long as the audio tool is open. The frequency of transmission is dependent on the number of group members and has been observed to be 5 seconds for two members. The interval is made longer as the number of session members increases so that the name information traffic does not create unnecessary link congestion. When a member properly closes the audio tool, a final

[1]Collecting complete data for video conferences is likely impossible because name information is not always transmitted by the video tools.

message is transmitted informing the group that the member is leaving.

Both commonly used audio tools[4, 8] work using an advertised multicast IP address and a pair of UDP ports. One UDP port is used for transmitting audio data, and the second for member name information. It is this latter information that we are interested in collecting. This second UDP port, or *control port*, is typically the advertised data port number plus one [3, 8]. When an audio tool is started, two pairs of UDP port numbers are used; one pair each to receive and transmit control messages and data. The receive pair is set to the session address, but the transmit pair is locally determined for each each new instance of the tool. Different transmit port numbers can be used to distinguish between different *connections* by the same host. From an application layer viewpoint, a connection starts when the audio tool is opened and ends when the tool is closed. Over the lifetime of a session a single user may open many connections some even simultaneously.

3 Measurement Tool Development

The purpose of the measurement tool is to record information about the membership dynamics of MBone sessions. Join and estimated leave times for each member are collected by listening to a session's multicast address and recording the name information transmitted by the session's members. The operation of the measurement tool can be broken down into the following four major functions: (1) collecting packets, (2) processing packets, (3) estimating connection end times, and (4) maintaining active and archived records.

(1) Collecting Packets. Data is collected by opening a socket on the multicast IP address and UDP control port for a session. Packets are collected using the *recvfrom* instruction which allows *Mlisten* to determine the sender's IP address and UDP port number.

(2) Processing Packets. For an observed session, our tool maintains a record for each active connection. (Recall that a connection is identified by the receiver's UDP port number and IP address and that a single host can be involved in many connections over the lifetime of the session.) When a name information packet is received, it is processed by the measurement tool to either create a record, or update an existing one. For each active connection, the measurement tool maintains the following information:

1. IP address of the receiving host.
2. UDP port number on which the group member is receiving the session broadcast.
3. The date and time when the connection was established.
4. The status of the record which is active until the connection ends.
5. The latest date and time a packet was received.
6. The number of packets received for the connection.
7. The estimated date and time of when a connection ended.

(3) Estimating Connection End Times. In some cases, determining the time that a connection ends can be difficult. Because of network unreliability, neither the audio tools nor *Mlisten* can determine if a session member has really left. VAT deals with this problem by using two status levels in the display. All hosts which have been heard from in the last x seconds (observed to be approximately 45 seconds) are considered to be active. If after 45 seconds no packets have been received from a particular host, that host's entry in the audio tool's list of listeners is turned grey. Unless an explicit leave message is received from a member, their entry is not removed for a long time (observed to be on the order of 10 to 15 minutes).

Mlisten uses a threshold to decide if a group member has closed a connection and left a session. When the threshold is exceeded, *Mlisten* marks the connection record as being completed. A problem with any timeout value is that some MBone links, especially those that are transoceanic, can experience extremely high loss. In some cases, session members may not be heard from for several minutes and the measurement tool would believe they had left; only to return later with the same UDP port number. One solution is to increase the threshold to a value large enough to accommodate even very long response times. However, some multicast routing protocols require a member to inform the router that it is still part of the multicast group. A router would use a much shorter threshold in its pruning algorithm. Therefore it is more reasonable to use a short threshold and then use the pre-processing step to detect obvious cases of loss causing multiple records for one connection. The threshold used in collecting data reported in this paper was 1.5 minutes.

(4) Maintaining Active and Archived Records. The measurement tool periodically performs record maintenance which consists of writing completed connection records to log files, and freeing these records for re-use. For the tests reported in this paper, cycles are set to be 30 seconds long. At the end of a cycle, the measurement tool will scan the active connection records and look for session members who may have left the group. The measurement tool uses the timeout threshold on the latest date and time fields to determine if a connection is still active or if it has ended. If a connection is believed to have ended, a record with the following format is written to a log file:

```
-----------------------------------------------------------------------
Name:  '130.207.8.11.4909'.  Received:  74.  Status:  1.
Start Date:  02/02/95.  Start Time:  14:36:37.
Latest Date: 02/02/95.  Latest Time: 14:43:50.
End Date:    02/02/95.  End Time:    14:47:00.
-----------------------------------------------------------------------
```

4 Pre-Processing Raw Log File Data

Our interest in collecting session participation data is to analyze the dynamic characteristics of MBone sessions. Some characteristics of the MBone can, however, bias observed data and must be accounted for. These characteristics include:

1. High loss over some MBone links.
2. Unpredictable packet delays.
3. Downtime of machines running tunneling daemons.
4. Testing or experiments conducted on the MBone.

As was discussed in the previous section several of these characteristics make determining the real leave time of a multicast session member difficult. To remove these biases our pre-processing step looks for connection records with the same IP address and UDP port number and attempts to determine if multiple records are for the same connection and should be combined. Pre-processing also eliminates records generated as a result of experimental usage. These records are characterized by an extremely high rate of requests of very short duration from the same host.

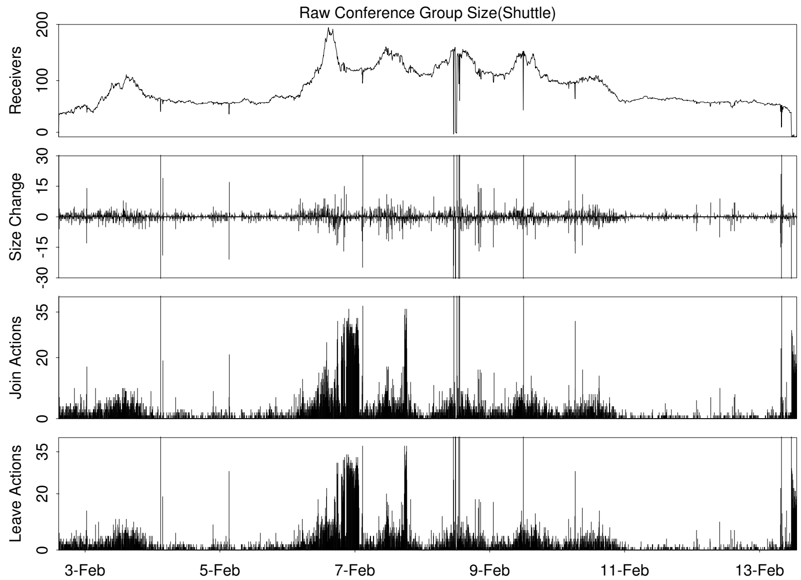

Figure 1: Group dynamics for the STS-63 Space Shuttle Session – Raw data.

Pre-Processing Example. Figures 1 and 2 show plots for the raw and pre-processed data from the STS-63 Space Shuttle multicast session. Each figure plots four different aspects of the data as a function of session time. Time is shown on the horizontal axis, and data points represent samples taken at 5 minute intervals. The vertical axis of each graph shows, in order from top to bottom: (1) the number of receivers, (2) change in group size between intervals (number joined - number left), (3) the number of new members, and (4) the number of members that left in a time interval. These numbers are discussed in more detail in the next section. For now we focus on the transformation from the raw data in Figure 1 to the processed data in Figure 2.

4.1 Verifying Connection End Times

Pre-processing consists of searching through the log file and identifying multiple connection records with the same host IP address and UDP port number. If such records are found and the time between the end of one and the start of another is small (on the order of 5 minutes or less) then the two records are probably for the same connection. So while the observed behavior is a leave and re-join, it is more likely that the connection lasted from the start of the first record to the end of the second. Multiple records with this characteristic are combined into single records. MBone links with frequent high losses may produce many connections for what was actually only a single

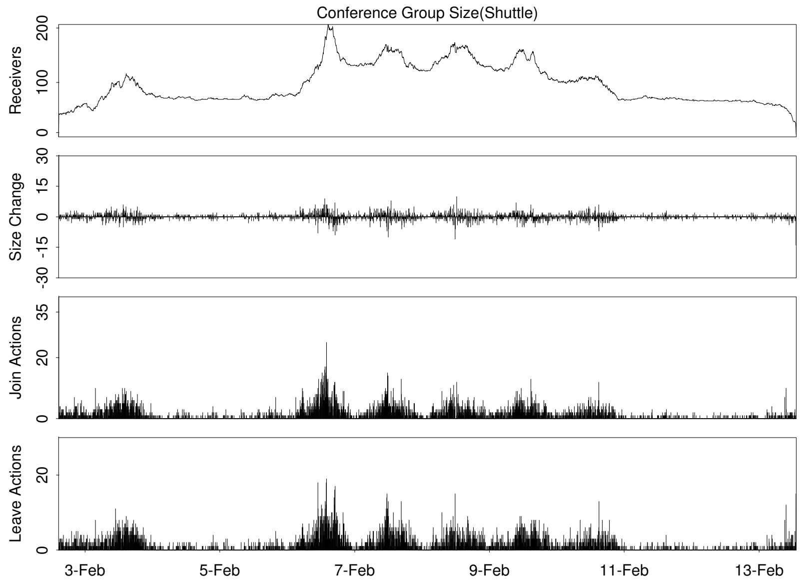

Figure 2: Group dynamics for the STS-63 Space Shuttle Session – After Pre-processing.

long connection. A comparison of Figures 1 and 2 shows that pre-processing resulted in a lower observed frequency of join and leave actions. In this particular instance pre-processing reduces 8229 observed connections to only 5055 "real" connections.

In addition to problems caused by high link losses, an occasional failure of a host running a multicast routing daemon or a link close to the measurement tool will result in the observation that all group members have left a session. This can be seen in Figure 1 on the evening of February 8th. Pre-processing will compensate by using connection records before and after the down-time to figure out what connections existed during that period. The one limitation is that the actual start or end time of a connection which joined or left during the down-time can not be measured. New connections are assumed to have started at the end of the down-time and ending connections are assumed to have ended at the start of the down-time.

4.2 Adjusting Data for Experimental Usage

Pre-processing also identifies unusually large numbers of short-lived, near-simultaneous connections by a single host. The graphs in Figure 1 show that on and slightly before February 7th, there was a period of extremely high activity in the group. A closer look at the log files showed that one host joined and left the group many hundreds of times over this 12 hour period. Each connection

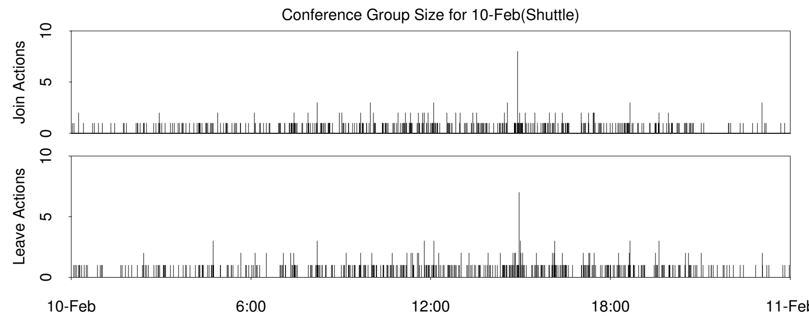

Figure 3: Enlargement of a 24 hour period of the Space Shuttle session

only lasted one or two seconds. After examining the log files, it was obvious someone was running membership benchmarking. For this reason, all the connections made by this host that lasted less than a couple of seconds and which occured during the 12 hour period were removed.

Experimental Usage vs High Turnover. Determining which observed data is the result of experimental usage and which is actual member behavior is difficult especially when member behavior is unknown. For this reason, removal of data suspected to be the result of experimental usage occurs infrequently in our pre-processing. Another reason for using caution when removing data is that certain types of user behavior look just like experimental usage. It has been observed that on occasion a host will start several near-simultaneous connections, some of which are short-lived. An example of this can be seen in the evening of February 10th in the graphs of Figure 2. There appears to be a spike in the the number of new connections and an almost simultaneous spike in the number of connections that ended. Figure 3 shows an enlargement of the time surrounding this event, and shows more clearly that an unusually large number of connections were started and ended at almost the same. Examination of the log file showed 7 connections were established and broken by the same host during a short period of time. This type of behavior may not represent typical usage of the application tools, but it does represent actual user behavior, and as such, it is part of the membership dynamics. One reason why this behavior might occur is that there is a delay when opening an audio tool. No indication is given in SD when a tool is successfully started so a user might continuing clicking on the "Open" button thereby opening multiple sessions by mistake. (This observation makes an important point about the effect that the application's user interface might have on multicast session dynamics.)

5 Analysis of MBone Session Dynamics

We now turn our attention to analyzing the data generated by *Mlisten*. In this paper we give an analysis of the sessions listed in Table 1. After pre-processing, the data was analyzed using another tool written by the authors called *Mprocess*. This tools generates graph data which can be plotted using most available graph tools.

Session Name	Measurement Start Day	Measurement Stop Day	Multicast IP Address	Multicast UDP Port Number	Source IP Address
STS-63 Space Shuttle Mission	Feb 3	Feb 13	224.2.248.153	59359	128.102.84.140
Info → http://www.ksc.nasa.gov/shuttle/missions/sts-63/mission-sts-63.html					
UCB Multimedia Seminar	Feb 17	Feb 17	224.2.238.75	65198	128.32.32.120
Info → http://roger-rabbit.cs.berkeley.edu/sp95-298/mss.html					
IPNG Working Group Meeting	Feb 9	Feb 11	224.2.255.152	65151	Unknown
IMS: World Radio Network	May 21	May 31	224.2.143.24	62253	204.62.246.73
Info → http://town.hall.org/radio/live.html					

Table 1: MBone sessions included in this paper.

5.1 Group Size and Join/Leave Data

A session's group size and the join/leave frequency of its members are important MBone session characteristics to examine. In addition to the Space Shuttle mission dynamics shown in Figures 2 and 4; Figures 5, 6, and 7 show statistics for several other MBone sessions. All the graphs use a 5 minute *sampling interval*. Data points are either sampled at the end of an interval (as with the number of receivers) or represent the number of actions occuring during the interval[2].

STS-63 Space Shuttle Mission. Figure 2 shows the results for one of the most popular type of MBone sessions. Like other Space Shuttle sessions, this one had a large number of receivers, exceeding 200 at one point in time. In addition to popularity, the shuttle missions have shared

[2]The sampling interval is simply a representation used in plotting results.

other similarities. The missions usually last on the order of 10 days with audio and video streams transmitting nearly continuously. And except for an occasional accidental transmission by someone else, the source transmits nearly constantly for the entire session. These similarities make it relatively straightforward to predict member behavior of future shuttle sessions.

An examination of Figure 2 reveals three noteworthy patterns:

- Activity is higher during normal working hours and seems to cycle every 24 hours with well defined peaks. Of course, because normal working hours vary around the world, some averaging occurs. (We examine the impact of location shortly.)
- Activity is low on the weekends. (February 4th is a Saturday.)
- The occurrence of significant events can have a pronounced effect on session activity. For shuttle missions, the most interesting events are the launch, and the landing. The primary mission objective is also popular and spikes in membership occur around this event. For example, an objective of the STS-63 mission was the practice approach for the rendezvous of the Space Shuttle and Space Station Mir[3].

One of the more surprising aspects of the graphs in Figure 2 is the constant "buzz" of join/leave activity even during times when there is not much happening. This has also been observed for other sessions. This can perhaps be attributed to the novelty (to some) of the MBone tools and/or to the fact that MBone usage is for free.

Another surprising aspect of the graphs is that join activity correlates well with leave activity. There are no periods when increased join activity does not result in an equal amount of leave activity. The cause is related to the connection duration which is discussed in Section 5.2.

We can further analyze group size information by grouping receivers based on their geographical location. Using a DNS server, a domain name can be determined for each receiver's IP address. Figure 4 shows session participation using 5 groups of domains: (1) "edu", (2) "com", (3) others in North America, (4) domains across the Atlantic, and (5) domains across the Pacific. One of the most interesting features of these graphs is that although the same peak and valley pattern is duplicated in all domains, the pattern is shifted according to time zone differences. Another observation is

[3]The actual docking of the two spacecraft during the STS-71 mission set a record for group members at 328.

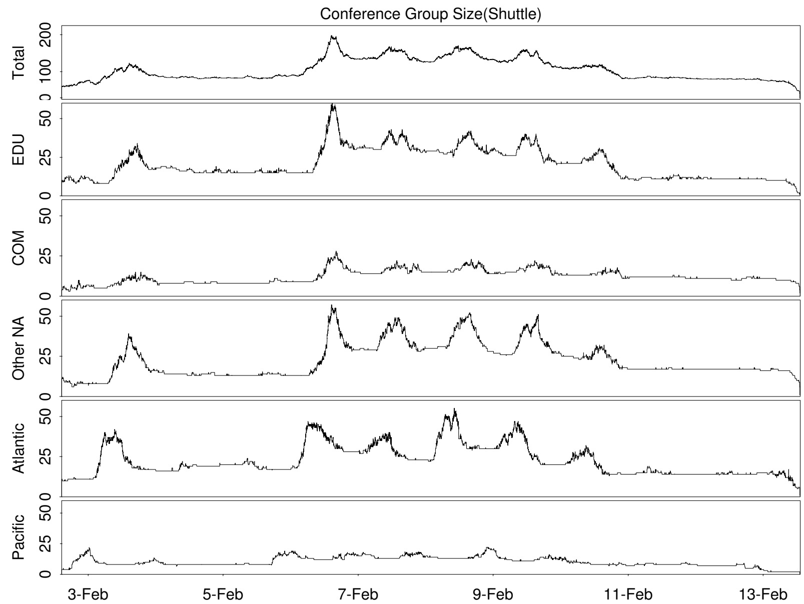

Figure 4: Breakdown of group size by domain for the Space Shuttle session.

that there is significant international (non-U.S.) interest in the space shuttle missions; on occasion accounting for more than 50% of the group members.

IMS World Radio Network. Figure 5 shows the membership dynamics for a less popular, but constantly transmitted session. The group size peaks around 30, but does not drop much below 15. On occasion there is some highly concentrated activity, but this can be traced to one member who started several sessions and then closed all but one of the sessions. One other interesting observation is that May 26th was a Friday and May 29th was a national holiday in the United States. Over the long weekend there was a noticeable drop in activity.

IPNG Working Group Meeting. Figure 6 is a different kind of session. Instead of continuously transmitting audio and video, this multicast session was a meeting that lasted for several hours on each of two days. Most of the group activity occurred during those two periods of time. One interesting result is that when the meeting was over at the end of the first day, a large number of members chose to remain in the group instead of leaving and coming back at the start of the next day's meeting. Even after the meeting ended some members did not leave. It is not known why these people did not leave, especially since nothing else was ever transmitted.

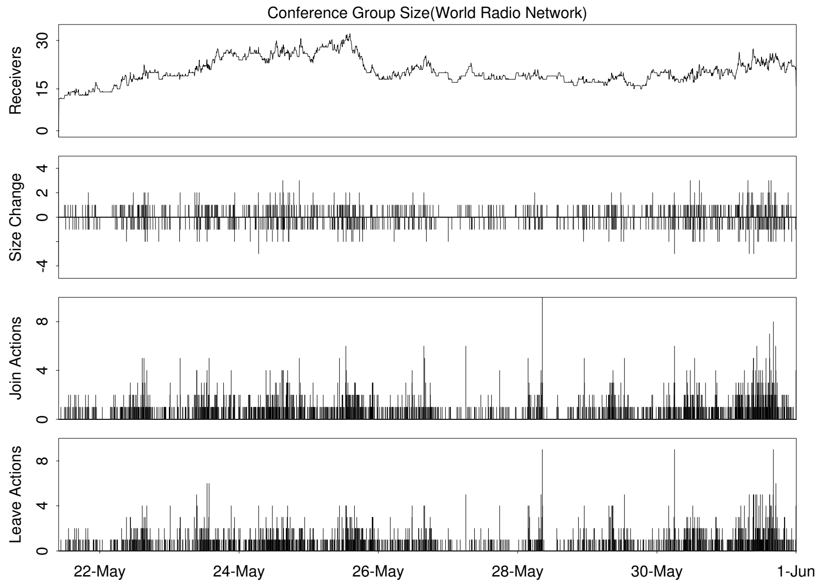

Figure 5: Group dynamics for the IMS World Radio Network.

UCB Multimedia Seminar. Figure 7 shows the results for a session that was scheduled to start at 5:00pm (EST) and last for approximately two hours. MBone behavior is similar to the behavior of patrons at a movie theater: members steadily arrive until just after the start time and most remain for the duration of the program. At the end of the session (most) everybody leaves. Both events have a few people coming and going in the middle of the program. In addition to the normal reasons (i.e. breaks, food, etc) activity in the MBone can be attributed to two other factors: (1) losses in the network were particularly high making the session nearly impossible to follow, and (2) because the MBone carries multiple sessions, some users are prone to "surf" among active sessions.

5.2 Per-Connection and Per-Member Durations

Another interesting measure that can be distilled from the information that *Mlisten* collects is the duration of a user's participation in a session. We distinguish between two measures:

- **Per-Connection Duration**: The time period from when a user connects to a session until the user quits from the session.
- **Per-Member Duration**: The sum of the durations for all connections that a member (identified by an IP address) makes to a particular session.

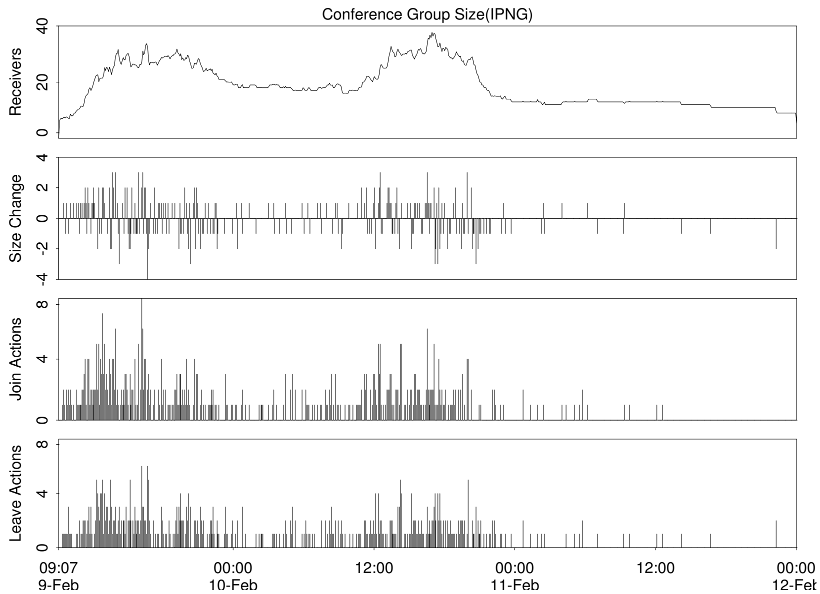

Figure 6: Group dynamics for an IPNG Working Group session.

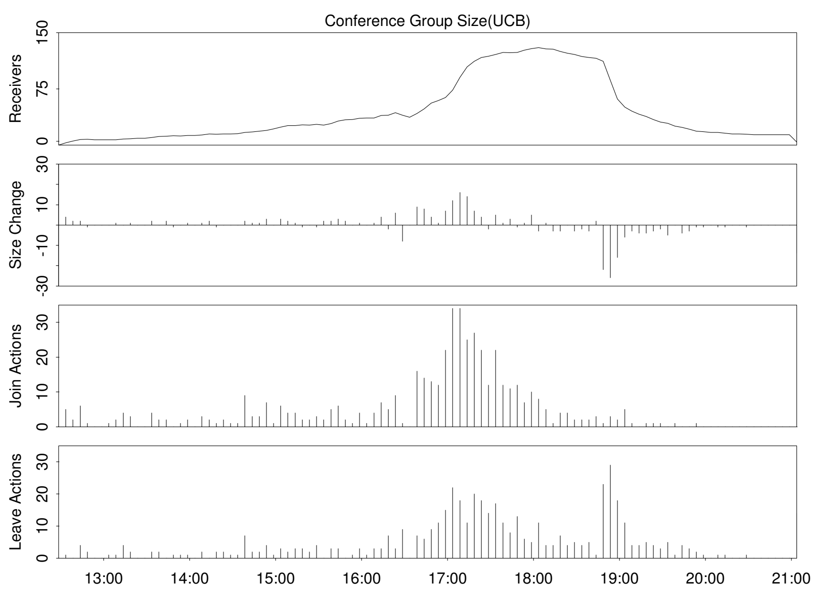

Figure 7: Group dynamics for the UCB Multimedia Seminar session.

Figure 8 shows a histograms for both types of measures for the STS-63 Shuttle multicast session. The histograms are shown with time scales from top to bottom of days, hours and minutes respectively. The y-axis is the number of connections on a logarithmic scale. It is perhaps surprising that most durations (per-connection and per-member) are relatively short with 70% of connections lasting less than one hour. It is interesting to note that there is a significant number of connections that last for the entire duration. The average and median connection durations were 5 hours and 6.1 minutes, respectively. The tremendous difference suggests that there are really two types of connections: many that are short-lived, and a few that are very long. The average and median per-member durations were 21 hours and 62 minutes, respectively. This suggests that on average a member made 4.2 connections. The difference between the per-connection and per-member histograms suggests that many hosts that join a session are likely to come back later.

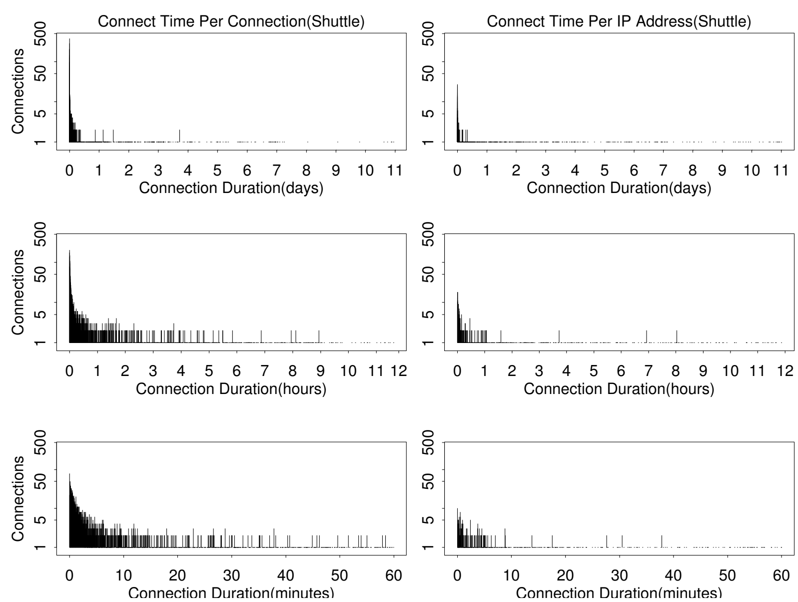

Figure 8: Per-connection and per-member duration histograms for the shuttle mission.

Figure 9 shows the connection duration histograms for the UCB session. Even though the session only lasted about two hours; some members had connections longer than six hours. However, most members only listened for approximately two hours. The average and median per-connection duration were 46 and 7 minutes respectively. The average and median per-member durations were 105 and 88 minutes. This translates to an average of 2.3 connections per member; a smaller number than for the shuttle session, but reasonable given the shorter duration of the UCB session.

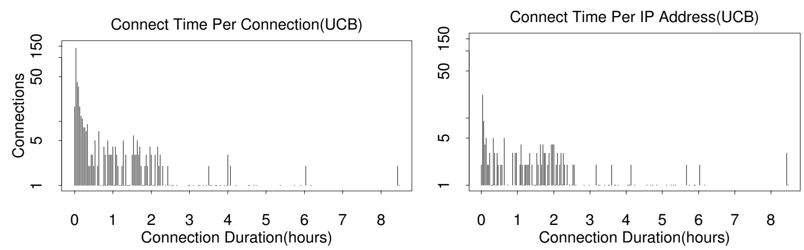

Figure 9: Per-connection and per-member duration histograms for the UCB session.

6 Analysis of MBone Session Multicast Tree Cost

With the information collected by *Mlisten* we can also derive and analyze information regarding the size of the multicast tree used by the session. A unique aspect of this analysis is that we observe how the tree characteristics change *over the lifetime* of a session.

In our analysis to follow we consider several types of trees. Our cost metric is that of packet-hops (first used in [14]) which is simply the sum, over all the links, of the number of packets carried over each link for each multicast. For example, a multicast from one source to multiple destinations will have a cost in packet-hops equal to the number of links in the tree. The following scenarios, summarized in Table 2, are:

1. **MBone Tree**: We obtain the topology of the actual MBone tree using the information we have about the IP address of the source and the IP addresses of the member. The routes in the MBone tree are roughly approximated using the MWatch tool [10]. The multicast hop count is compared to the number of packet hops in the same tree but assuming that the source transmits to each receiver individually (also known as the *separately-addressed packets* scheme [14]).

2. **MBone Tree with Georgia Tech host as the Source**: In order to assess the sensitivity of tree costs to the location of the source we repeat the calculations for both 1 and 2 above but assume a host at Georgia Tech was the source. The MBone tree is changed to reflect the change in source, with the session members remaining the same.

3. **Internet Tree with Georgia Tech as the Source**: MBone hops are really tunnels which can span multiple Internet hops. So in order to estimate an Internet hop cost we would like

be able to determine the actual cost of a tree as if it were running on a fully multicast-capable Internet. For this, we considered a Georgia Tech host as the source and use the *traceroute*[15] utility to determine the shortest path[4]. We use this tree to determine the appropriate unicast and multicast costs. To the extent that the results from item 3 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.

	Network	
Source	MBone	Internet
Observed Source	MWatch: Tree used by the actual session.	Not determinable from Georgia Tech
Georgia Tech	MWatch: Used to test sensitivity of source.	Traceroute: Approximates multicast-capable Internet.

Table 2: The different trees considered and the methods used to determine them.

Figure 10 shows the tree costs for the IMS World Radio Network, a small, continuous multicast session. Several observations are immediately apparent. First, and foremost, as the number of receivers varies, the cost of the multicast tree does not change significantly, but the cost of the unicast varies greatly. Second, even for a small conference, the advantages of using multicast communication are apparent. On average, the multicast tree cost is only 35.7% of the unicast cost. At the peak group size, this percentage is even better with the multicast tree cost only 32.5% of the unicast cost. A third observation is that the positioning of the source in the MBone does not have a significant impact on the cost. One reason why the difference is so small could be that the IMS source host and Georgia Tech host are separated by only a few hops. The final two graphs in Figure 10 shows that in the Internet the advantages of multicast (over unicast) are even greater. Because of tunneling on the MBone, the Internet multicast tree cost is 2.0 times more than the MBone tree cost.

Figure 11 shows the tree costs for the UCB Multimedia Seminar session. For this larger session, the advantages of using multicast communication become even more significant. Even as the number of receivers doubles and doubles again the cost of the multicast tree increases only slightly. For

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

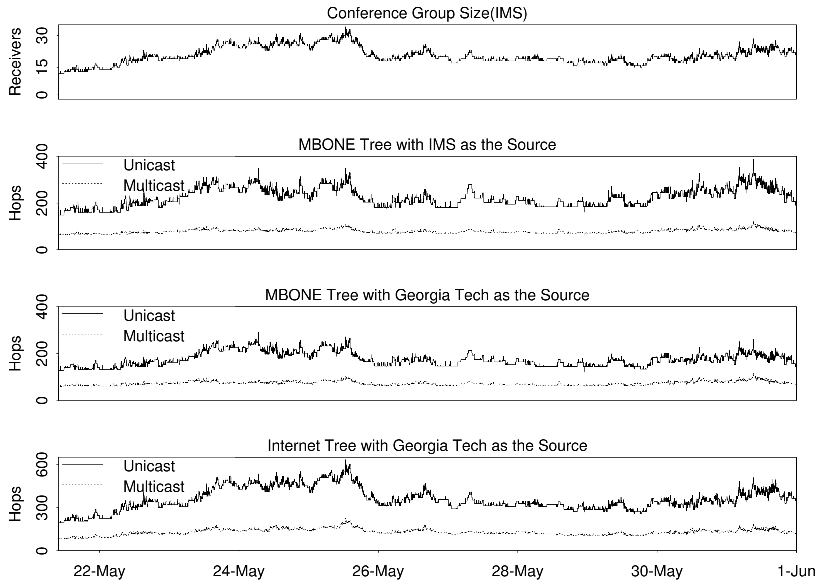

Figure 10: Multicast and unicast costs for the IMS World Radio Network.

almost the entire seminar the multicast tree cost is only 22.6% of the unicast cost. Again due to the effect of tunneling, the Internet multicast cost is 2.4 times the MBone multicast cost.

Figure 12 shows the tree costs for the STS-63 Space Shuttle session, the largest session for which data was collected. The results show that while there is a direct relationship between the the number of members and the number of unicast packet-hops, the number of multicast packet-hops remains nearly flat. The cost for the multicast tree is on average only 30.1% of the unicast cost. This is somewhat larger than the ratio for the UCB Multimedia Seminar. The reason for this is that the set of destinations for the shuttle mission is more geographically dispersed. Almost all of the receivers for the UCB seminar were in the United States while a larger percentage of the shuttle session members are outside of the United States(see Figure 4). Changing the source for the shuttle mission does not have a significant affect on the costs. Because of MBone tunneling, the Internet multicast tree costs 2.3 times the MBone tree; a number consistent with the ratios for the other sessions. The results for all three sessions are summarized in Table 3.

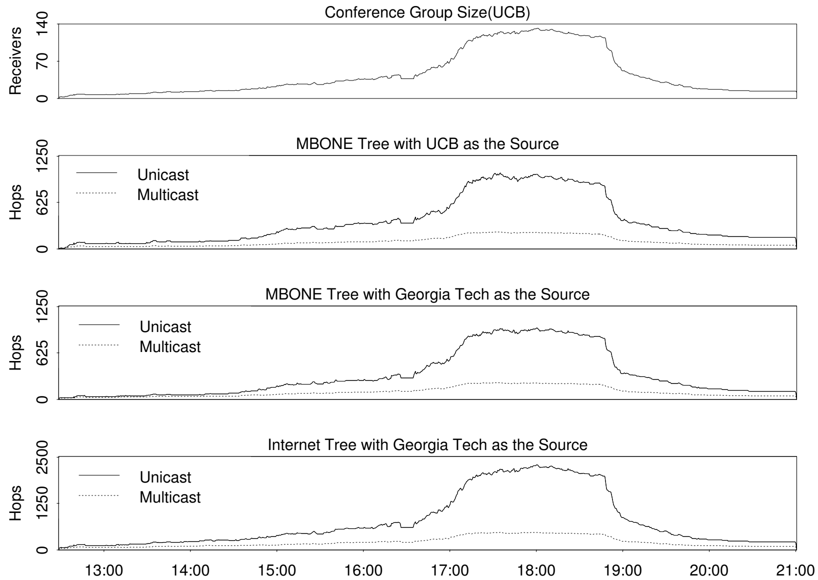

Figure 11: Multicast and unicast costs for the UCB Multimedia Seminar session.

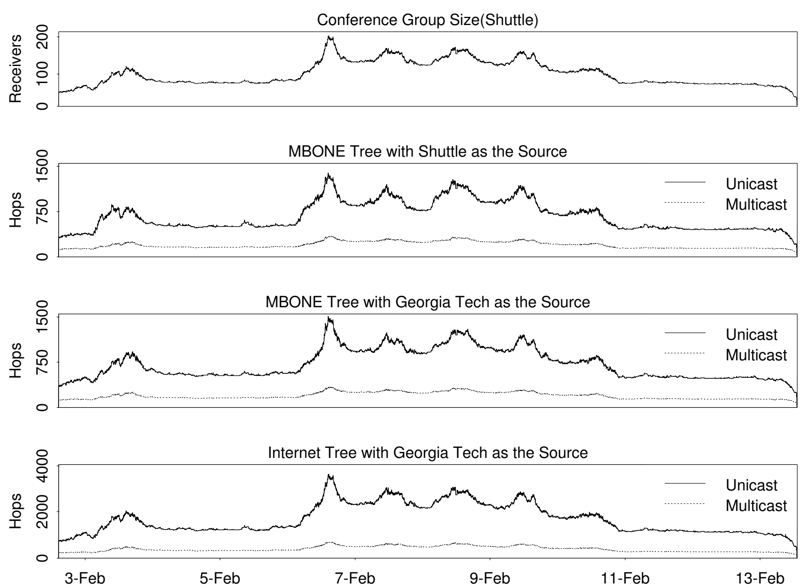

Figure 12: Multicast and unicast costs for the Space Shuttle session.

Session	MBone Topology Original Source		MBone Topology GA Tech Source		Internet Topology GA Tech Source	
	Average	Maximum	Average	Maximum	Average	Maximum
IMS World Radio Network	11.0	25	8.6	22	17.4	32
UCB Multimedia Seminar	8.8	21	7.1	14	16.8	23
STS-63 Space Shuttle Mission	7.1	17	7.5	15	17.5	36

Table 3: Average and maximum path lengths for multicast sessions.

7 Summary and Concluding Remarks

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 measure and analyze MBone multicast sessions. An analysis of this measured data can be extremely beneficial in understanding how multicast network support will be used. One of the challenging aspects of developing this tool was dealing with various MBone idiosyncrasies that can significantly bias observed behavior. We use our tool to analyze some recent MBone sessions. The analysis provides insights into the dynamics of member behavior in various types of multicast sessions. The data collected can also be used to understand the effect of the member behavior on network resource usage. We provide an example of this by analyzing multicast tree costs for the sessions we observed.

Other uses of our tool and the data that it produces need to be explored. This includes: (1) modeling data for use in simulations and as a basis for realistic traffic workloads, (2) using the data to produce visualizations of MBone session dynamics, and (3) extend *Mlisten* to collect data for all active sessions.

8 Acknowledgments

This research was conducted with assistance from several people. Atanu Ghosh, University College London, provided help with the MWatch tool. Bill Fenner at Xerox PARC and Steve Casner at Precept explained some of the mysteries and anomalies in the MBone. Amarnath Mukherjee, Naser Adas, and Steve Klivansky at Georgia Tech offered help with data analysis and S-Plus.

References

[1] 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.

[2] S. Casner, *Frequently Asked Questions(FAQ) on the Multicast Backbone(MBone)*. USC Information Sciences Institute(ISI), Dec 1994. Available from ftp://ftp.isi.edu/mbone/faq.txt.

[3] V. Jacobson and S. McCanne, *SD: Session Directory Tool*. Lawrence Berkeley Laboratory, February 1993. Available from ftp://ee.lbl.gov/conferencing/sd/.

[4] V. Jacobson and S. McCanne, *VAT: Visual Audio Tool*. Lawrence Berkeley Laboratory, February 1992. Available from ftp://ee.lbl.gov/conferencing/vat/.

[5] V. Jacobson and S. McCanne, *WB: Whiteboard Tool*. Lawrence Berkeley Laboratory, July 1994. Available from ftp://ee.lbl.gov/conferencing/wb/.

[6] S. McCanne and V. Jacobson, *VIC: Video Conference Tool*. Lawrence Berkeley Laboratory, November 1994. Available from ftp://ee.lbl.gov/conferencing/vic/.

[7] R. Frederick, *NV: Network Video Tool*. Xerox PARC, March 1993. Available from ftp://parcftp.xerox.com/pub/net-research/nv-3.3beta/.

[8] H. Schulzrinne, *Voice Communication Across the Internet: A Network Voice Terminal*. Department of Electrical Engineering, University of Massachussets at Amherst, Jul 1992. Available from ftp://gaia.cs.umass.edu/pub/hgschulz/nevot.

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

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

[11] S. Deering, A. Thyagarajan, S. Casner, and et al., *MROUTED 3.5 and Tools*, Jun 1995. Available from ftp://parcftp.xerox.com/pub/net-research/ipmulti/.

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

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

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

[15] V. Jacobson, *Traceroute*. Lawrence Berkeley Laboratory, February 1989. Available from ftp://ee.lbl.gov/traceroute.tar.Z.