Packet Loss Correlation in the MBone Multicast Network*

Maya Yajnik, Jim Kurose, and Don Towsley
Department of Computer Science
University of Massachusetts at Amherst
Amherst MA 01003
{yajnik,kurose,towsley}@cs.umass.edu

Abstract

The recent success of multicast applications such as Internet teleconferencing illustrates the tremendous potential of applications built upon wide-area multicast communication services. A critical issue for such multicast applications and the higher layer protocols required to support them is the manner in which packet losses occur within the multicast network. In this paper we present and analyze packet loss data collected on multicast-capable hosts at 17 geographically distinct locations in Europe and the US and connected via the MBone. We experimentally and quantitatively examine the *spatial and temporal correlation in packet loss* among participants in a multicast session. Our results show that there is some spatial correlation in loss among the multicast sites. However, the shared loss in the backbone of the MBone is, for the most part, low. We find a fairly significant amount of of burst loss (consecutive losses) at most sites. In every dataset, at least one receiver experienced a long loss burst greater than 8 seconds (100 consecutive packets). A predominance of solitary loss was observed in all cases, but periodic losses of length approximately 0.6 seconds and at 30 second intervals were seen by some receivers.

1 Introduction

The recent success of multicast applications such as Internet teleconferencing tools [3, 9] for audio, video, and whiteboard, and distributed interactive simulation illustrates the tremendous potential of applications built upon wide-area multicast communication services. A critical issue for such multicast applications and the higher layer protocols that support them is the manner in which packet losses occur within the multicast network.

In this paper, we present and analyze packet loss data collected simultaneously at up to 12 hosts at geographically distinct locations in Europe and the US. These hosts are connected via the Multicast Backbone (MBone) network [3]. The primary goal of this work is to examine the *spatial and temporal correlation in packet loss* among participants in a multicast session. (Informally, by "spatially" correlated loss, we mean the loss, i.e., lack of reception, of the same packet at many sites; by "temporally" correlated loss, we mean the loss of consecutive packets at a given receiver.) Our results show that:

- For most of the traces, the loss on the backbone links of the MBone multicast network is observed to be small (2% or less), as compared to the average loss seen by a receiver. However, due to occasional outages lasting from few seconds to few minutes, in some backbone links, the spatially correlated loss between receivers does go up to 20%, in a few datasets.
- There is a significant amount of burst loss (consecutive losses) at each site. One or more extremely long loss bursts, lasting from a few seconds up to 3 minutes (around 2000 consecutive packets), occur in almost every trace. Such long loss bursts have been reported in [7] for the case of point-to-point connections.
- Most of the loss bursts consist of isolated single losses, but the few very long loss bursts contribute heavily to the total packet loss.
- Some receivers see periodic packet loss lasting for approximately 0.6sec. (8 consecutive packets) and occurring at 30 sec. intervals. This is possibly due to the routing updates as reported in [5].

The underlying packet loss process is of tremendous importance to error control protocols. This is particularly so with multicast communication, since many of the proposed error control protocols recover from packet loss by having receivers interact with other receivers rather than with the data source itself. Thus, the spatial correlation of loss is of particular importance. Although there has been a considerable amount of research on multicast error control protocols (see [6, 8] and the references therein) these works have either not examined or considered the underlying loss process, or have assumed that packet losses are both spatially and temporally independent; the two exceptions are [1, 2]. The work by Bhagwat *et al.* [1] describes a recursive analytic method for computing the probability that a packet is not received at one or more receivers given a specific multicast tree and known, independent loss probabilities on each link. The work by Bolot *et al.* [2] is the work most closely related to our present work. In that work, packet loss measurements are presented from a 10,000-packet trace between MBone sites in France and England. With respect to temporally-correlated loss, they find that "losses appear to be isolated"; they do not address the issue of spatially correlated losses. Interesting experimental observations on routing behavior in the Internet are presented in [7] which discusses a variety of observed routing pathologies and reports outages lasting longer than 30 secs. and up to 5 minutes long, due to changes in routing connectivity.

The remainder of this paper is structured as follows. In the following section we describe the measurement tools we constructed and how the data was collected. In section 3, we examine the spatial correlation of loss in the packet traces. In section 4, we examine the temporal correlation in loss. Section 5 concludes this paper.

2 Data Collection Background

Our measurements were performed by simultaneously monitoring and recording the received multicast packets during audio multicast sessions on the MBone at the 17 different MBone sites listed in Table 1. At some sites, two machines were used. Three different audio sources were used: the "World Radio Network" (WRN) transmitting from Washington DC, the "UC Berkeley Multimedia Seminar" (UCB) transmitting from California, and "Radio Free Vat" (RFV) also transmitting from California. These

*This work was supported in part by the Defense Advanced Research Projects Agency under contract F19628-95-C-0146, and the National Science Foundation under grant NCR-9508274

Machine Name	Location
alps	Georgia Institute of Technology
anhur, spiff	SICS, Sweden
artemis, atlas	I.B.P., Paris, France
bagpipe, ocarina	Univ. of Kentucky, Lexington
cedar	Univ. of Texas, Austin
collage, zip	E.I.T., California
dixie	Univ. of California, Irvine
edgar	Univ. of Washington, Seattle
erlang, trantor	Univ. of Massachusetts, Amherst
excalibur	Univ. of Southern California
float	Univ. of Virginia, Charlottesville
ganef	Univ. of California, Los Angeles
law	Univ. of California, Berkeley
pax	INRIA, France
tove	Univ. of Maryland, College Park
ursa, lupus	GMD Fokus, Berlin, Germany
willow	Univ. of Arizona, Tuscon

Table 1: MBone Sites

	Date	Source	# of Recs.	Time (mins.)	# of packets
1.	Sep 19,'95	WRN	8	23	17,000
2.	Sep 20,'95	UCB	9	14	21,000
3.	Oct 30,'95	WRN	10	76	57,000
4.	Nov 1,'95	WRN	9	55	41,000
5.	Nov 13,'95	WRN	9	53	40,000
6.	Nov 14,'95	WRN	8	40	30,000
7.	Nov 28,'95	WRN	7	27	20,000
8.	Dec 4,'95	WRN	8	60	45,000
9.	Dec 11,'95	WRN	9	93	70,000
10.	Dec 16,'95	WRN	7	45	50,000
11.	Dec 18,'95	WRN	7	92	69,000
12.	Apr 19,'96	RFV	11	60	45,000
13.	Apr 24,'96	UCB	12	62	93,000
14.	May 8,'96	RFV	10	99	148,000

Table 2: Datasets

audio sources transmit packets over the MBone at regular intervals. The WRN source transmitted packets at $80ms$ intervals each of which contained approximately 5Kbits of audio data within a vat audio packet. The UCB source transmitted at double the rate, at $40ms$ intervals and each packet contained 2.5Kbits worth of audio data. For the Apr 19th, 1996 trace, RFV transmitted at $80ms$. intervals, and for the May 8th, 1996 trace, it transmitted at $40ms$. intervals. By listening to the session multicast address at each site, it is possible to determine which packets arrive and which are lost. Note that while these packets contain audio data, our results are not tied to this specific application. We ignore the actual contents of these packets, essentially considering them as periodic test packets that are sent into the multicast network.

At each receiver, a process was run that listened to the multicast address and recorded and timestamped the vat headers of the arriving packets. The packet header contained a sequence number which uniquely identified each multicast packet sent by the source. These data collection daemons were remotely controlled by commands sent from a central control program to start, stop, and otherwise control them. Once the data was collected, the control program instructed the daemons to send the trace files via ftp to our centralized site.

14 different sets of traces have been collected, each lasting 15 to 99 minutes. Table 2 chronologically lists the datasets giving the source and the lengths of the traces. Not all receivers were able to receive data on a given day, either because the daemon was not set up at that time or because the site was disconnected from the MBone. All datasets can be obtained from our web-site http://www.cs.umass.edu/~yajnik/datasets.html or our ftp site ftp://gaia.cs.umass.edu/pub/yajnik.

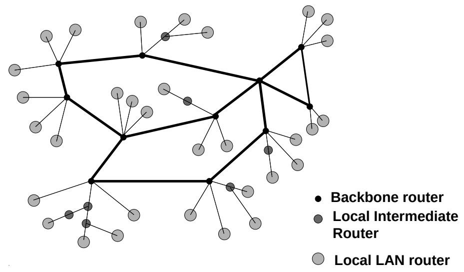

Figure 1: Hierarchical Topology of the MBone

3 Spatial Correlation of Loss

This section discusses the distribution of packet loss in the multicast transmission tree. Subsection 3.1 describes how the loss rates on the different segments of the transmission tree are determined. The backbone loss versus the average loss seen by the receivers for all the datasets is summarized later in the subsection.

We consider two ways of assessing the extent of spatial correlation of loss among receivers

- In subsection 3.2 we plot the distribution of M, the number of receivers that simultaneously lose a given packet. The measured distribution is compared with three computed distributions, each assuming different transmission topologies.
- The correlation coefficient of loss for a pair of receivers gives a measure of the spatial association of loss between them. The average of the correlations over all pairs of receivers is a measure of the overall spatial association for the dataset. Subsection 3.3 describes this method of measuring correlation.

Both analyses show that, for most datasets, the overall spatial correlation in loss in the network is small and does not have a major impact, *except* for loss occurring close to the source. This follows from our observation that backbone loss in the MBone is generally very low. Occasionally, there are extremely long periods of loss lasting for a few seconds or even a few minutes (as described in section 4) on the shared segments of the transmission tree. These long bursts of loss, when they occur, do contribute heavily to the spatially correlated loss.

3.1 Where Does Loss Occur?

The topology of the MBone is as follows. The MBone is a virtual multicast network built on top of the physical Internet to support routing of IP multicast packets. The design of the MBone is described in the MBone FAQ [4]. The nodes in the MBone are multicast-capable routers, logically connected to each other via IP routes known as "tunnels". That is, multicast packets are sent, in encapsulated form, over routers which are not multicast-capable, through point-to-point connections, called tunnels. The MBone has a "mesh-star" topology using two-tiered hierarchical routing as shown in Figure 1. There is a base of backbone multicast routers maintained by the service providers, interconnected by a mesh of tunnels, which forms the higher level of long-distance multicast connectivity (shown by bold lines in the figure). There are alternate routes between the main backbone routers, giving the MBone sufficient robustness to handle network failures. The "backbone router" provides multicast connectivity to its region by a star hierarchy of tunnels which fan out and connect to local multicast routers at organizations that wish to receive MBone packets. These in turn may branch out further to other local routers. Finally, there are multicast routers on the LAN of the intended receivers, each providing multicast connectivity to the machines on its LAN. The three different kinds of multicast routers shown

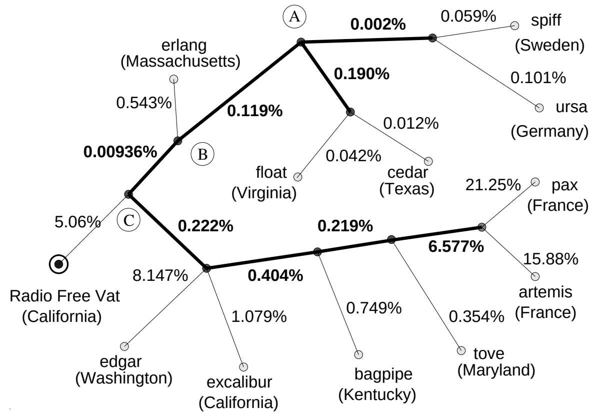

Figure 2: Transmission Tree: source - RFV, on Apr 19, 1996

in the figure are the backbone routers, the local multicast routers on the LAN of the intended receivers, and a few intermediate local routers which connect a backbone router to a local router and may also provide multicast connectivity to their own LANs.

Figure 2 provides a *logical* view the multicast transmission tree for the 11 receivers in the dataset of Apr. 19th, 1996, with the estimated probability of loss on each segment. The source of the packets was "Radio Free Vat" in California transmitting a packet every $80ms$. The receivers are shown, as are *selected* MBone routers between the receivers and the source. Every MBone router shown is the nearest common ancestor of all downstream receivers on the multicast tree. The multicast tree itself was constructed by joining together the multicast paths from each of the receivers to the source. Thus, a single tree segment is a *virtual* link and could include a series of tunnels and multicast connections. The routes taken by the multicast packets were determined by using the "mtrace" utility, the multicast "ping" program (with record-route option) and the "mrinfo" utility.

The bold lines in Figure 2 indicate the connections between the "backbone" routers, which form the base of the multicast tree and traverse much of the distance in the tree. The other branches of the tree are on the "edge" of the network. In some cases, these other branches may cross over backbone routers before reaching the local routers.

The data traces contain information that indicates which packets were lost by each of the receivers. For a given packet, examining which receivers received the packet and which did not can provide a valuable clue as to *where* in the multicast tree the packet was lost. For example, looking at Figure 2, if a packet is lost by `spiff`, `ursa`, `float` and `cedar` but received correctly at `erlang` it is likely that it was dropped between the multicast routers `A` and `B`. (It should be noted, however, that this need not be the case, as the packet could have been simultaneously and independently lost on the downstream paths from A, although we consider this latter scenario to be much less likely.) That is, the estimated number of packets lost on link from `B` to `A` is the difference between the number of packets lost by all receivers downstream from A and the number of packets lost by all receivers downstream from `B`. Let N_A be the number of packets lost by all receivers downstream from `A` and let N_B be the number of packets lost by all receivers downstream from `B`. Then the estimated probability of loss along link `AB`, p_A, is given by the following formula.

$$p_A = \frac{N_A - N_B}{N - N_B} \tag{1}$$

where N is the total number of packets sent by the source. Using this reasoning, we can determine the approximate percentage of packets lost on each of the links in Figure 2.

It is obvious in Figure 2 that the backbone loss, except for one segment between the USA and France, is rather low, ranging from 0.002% to 0.4%. Also, there is a major bottleneck, very close to the source which contributes 5% packet loss. Once the packets are past this bottleneck in California, there is very little loss, across the continent and even into Sweden and Germany. In general, looking at all datasets, we observed low loss rates (2% or less) along the MBone backbone. Occasionally, there are blackout periods or very long loss bursts, on the backbone, as discussed in section 4. However, the base loss rate, excluding extremely long burst loss, has been consistently low. This has important implications in the context of reliable multicast. When a receiver loses a packet, it may be able to recover the packet from a nearby receiver which correctly received it, instead of directly from the sender, as discussed in [6]. Such local recovery from loss would often be possible, due to the low backbone loss.

Another set of measurements we made regarding the spatial locality of loss was to determine whether any packet was being dropped at the receiving hosts themselves. To do so, we monitored the multicast session at two different workstations on the same end local area network at six sites: `anhur` (Sweden), `artemis` (in France), `bagpipe` (in Kentucky), `collage` (in California), `erlang` (in Massachusetts) and `ursa` (in Germany). We measured the percentage of all packets sent by the source that were lost by one receiver and not by the other. Surprisingly, the end-host loss was found to be negligible. It was zero in most cases and never exceeded 0.001%. We conclude that packets are almost never dropped between the network interface on the LAN of the receiver and the receiving daemon.

Table 3 shows the backbone loss rates vs. receiver loss rates for every dataset. The backbone loss rates were, in general, rather low (around 1%). However, some backbone links do occasionally show high loss of up to 20% due to the presence of a small number of extremely long loss periods extending from several seconds to several minutes. These long loss bursts are discussed in detail in section 4. The table gives the number of backbone links that experience long loss bursts versus the total number of backbone links in the transmission tree. A long loss burst is defined as a loss burst that affects 100 or more consecutive packets, when the sampling interval is $80ms$. For a sampling interval of $40ms.$, the threshold is 200 consecutive packets. The average backbone loss is the average over all the backbone links in the tree, and the average receiver loss is the average of the loss rates seen by each receiver in the dataset. In order to assess the backbone loss excluding these extremely long loss bursts, the table also shows the average "trimmed" backbone loss rates versus the average trimmed receiver loss rates. The trimmed loss rate for a backbone link is determined by computing the loss for the portions of the trace that do not show the long loss bursts. From the results in table 3, we conclude that average backbone loss is less than 2% for most datasets. The average *trimmed* backbone loss rate is always 2% or less in every case.

3.2 Distribution of the number of receivers that simultaneously lose a packet

From the point of view of a reliable multicast protocol, it is important to know the statistics of the number of receivers that simultaneously lose a given packet.

For the dataset described of Apr. 19th 1996, Figure 3 illustrates the distribution of M, the number of receivers that simultaneously lost a given packet. For this dataset, 47% of the packets sent by the source were lost by at least one receiver. In the context of reliable multicast, this implies that retransmission would have been necessary for 47% of the packets. The actual measured distribution is compared to three computed distributions, each based on a different model of the transmission tree. Note that temporal independence of loss is assumed in every model. The models are:

Data-Set	backbone links: # with long loss / total #	average backbone loss	average average recv. loss	average trimmed backbone loss	trimmed recv. loss
1.	2/ 5	6.94%	13.94%	2.07%	5.99%
2.	1/ 5	3.74%	18.08%	2.17%	10.87%
3.	1/ 6	2.02%	15.06%	1.51%	13.14%
4.	1/ 6	5.05%	20.99%	2.02%	14.75%
5.	0/ 5	0.29%	9.03%	0.29%	7.50%
6.	0/ 4	1.50%	22.50%	1.50%	13.95%
7.	1/ 4	5.00%	13.01%	0.95%	6.52%
8.	0/ 5	1.31%	13.69%	1.31%	13.52%
9.	1/ 4	1.19%	9.97%	0.32%	8.83%
10.	0/ 2	0.01%	4.90%	0.01%	4.83%
11.	0/ 2	0.12%	9.74%	0.12%	9.42%
12.	0/ 8	0.97%	9.28%	0.97%	8.78%
13.	1/ 8	1.36%	14.18%	1.06%	12.26%
14.	2/ 6	3.08%	14.61%	1.66%	9.35%

Table 3: Summary of Backbone Loss for all Datasets

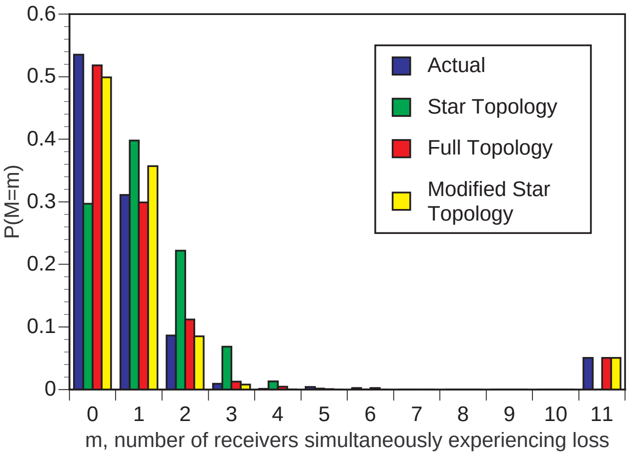

Figure 3: Distribution of M: source-RFV, on Apr 19, 1996

1. Star Topology: The packet loss is assumed to be spatially and temporally independent and measured probabilities of loss at the receivers are used to recursively compute the effective distribution of M. That is, the topology is assumed to be a "star" as shown in figure 4.
2. Full Topology: The packet loss is assumed to be spatially correlated as in the transmission tree of figure 2. The estimated probabilities of loss on each link are used to recursively compute the effective distribution of M, in a bottom-up fashion. That is, the distribution of M for a node is calculated using the calculated distributions for the downstream nodes.
3. Modified Star Topology: The distribution of M is computed based on a "modified star" topology shown in Figure 4. The probability of loss on the link from the source to node C is the fraction of packets lost by all the receivers. The rest of the loss is assumed to be spatially independent.

The histograms of Figure 3 show that the computed distribution using the first model based on a "star" topology is significantly different from the actual distribution of M. However, both the

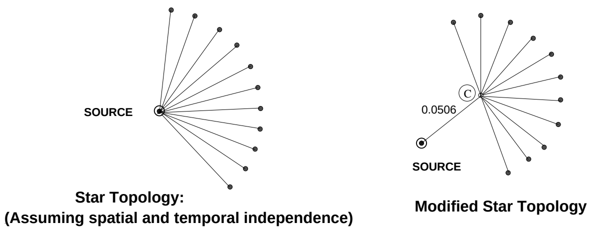

Figure 4: The Star and Modified Star Topologies

Dataset	% lost by 1 or more recs.	% lost by 1 rec.	% lost by 2 recs.	% lost by more than 2 recs.
1.	42.5%	10.8%	10.2%	21.5%
2.	64.3%	31.9%	14.9%	17.5%
3.	72.2%	33.8%	23.7%	14.7%
4.	55.1%	17.0%	3.8%	34.3%
5.	38.1%	27.4%	5.5%	5.2%
6.	69.5%	33.1%	18.1%	18.3%
7.	45.3%	20.3%	18.9%	6.1%
8.	63.2%	46.1%	8.9%	8.2%
9.	38.6%	24.7%	7.2%	6.7%
10.	69.0%	28.9%	4.2%	5.9%
11.	37.9%	29.9%	3.2%	4.8%
12.	46.5%	31.1%	8.7%	6.8%
13.	64.3%	34.1%	15.8%	14.4%
14.	62.6%	32.4%	11.9%	18.3%

Table 4: Summary of Distribution of M for all Datasets

distributions computed using the full topology and the modified star topology are close to the original distribution. This means that the topology is effectively that of a modified star, and the spatially correlated loss in the network is low except for the loss next to the source.

In general, for every dataset, the distribution computed using the full topology model with the transmission tree loss rates is close to the actual distribution. The distribution computed using the modified star topology model is close to the actual distribution for 9 out of the 14 datasets. The exceptions are the 1st, 2nd, 4th, 7th and 14th datasets (refer to table 3).

Table 4 contains a summary of statistics over a range of datasets taken on different days, and for different sources. The percentage of packets lost by more than two receivers ranges from 4.8% to 34.3%.

3.3 Correlation Between Pairs of Receivers

The correlation of loss for a pair of receivers is a measure of the association between them. The average correlation for all pairs of receivers in a dataset gives an overall measure of the spatial association in the dataset as a whole.

Let X_i be a binary random variable taking on the value 1 if the packet is lost at receiver i, and value 0 if the packet is correctly received by receiver i. Let $\overline{X_i}$ be the mean of variable, X_i. The correlation coefficient between any two receivers i and j is defined as

$$corr(X_i, X_j) \quad = \quad \frac{E[(X_i - \overline{X_i})(X_j - \overline{X_j})]}{\sqrt{E[(X_i - \overline{X_i})^2]}\sqrt{E[(X_j - \overline{X_j})^2]}} \qquad (2)$$

If X_i and X_j are independent, then $corr(X_i, X_j)$ is zero. The correlation coefficient varies between -1 and 1. Table 5 shows the average correlation coefficient between pairs of receivers for each dataset. The average is taken over all receiver

Dataset	Average Correlation Coefficient	Average Correlation without loss next to the source
1.	0.271	0.162
2.	0.351	0.320
3.	0.299	0.042
4.	0.556	0.099
5.	0.606	0.019
6.	0.506	0.099
7.	0.355	0.026
8.	0.458	0.174
9.	0.556	0.024
10.	0.488	0.004
11.	0.637	0.024
12.	0.666	0.114
13.	0.425	0.131
14.	0.341	0.114

Table 5: Average Correlation between Pairs of Receivers

Mach. Name	loss rate (%)	% of bursts	Avg. length	Coef. of Var.	99-%	long-est burst	% loss in long bursts
alps	5.93	3427	1.21	2.91	3	179	4.3%
anhur	5.15	3387	1.07	0.25	2	4	0.0%
cedar	14.22	7463	1.33	0.83	8	14	0.0%
collage	9.08	5508	1.16	2.07	3	175	2.75%
erlang	10.41	3793	1.92	21.30	4	2518	34.6%
float	10.44	6470	1.13	0.37	3	7	0.0%
law	12.09	4983	1.70	21.00	3	2518	29.8%
pax	16.98	7633	1.56	19.13	3	2603	21.9%
tove	5.46	3486	1.10	0.41	3	10	0.0%

Table 6: Burstiness of Loss: source-WRN, on Dec 11,1995

pairs in a dataset. We observed that, as indicated in subsection 3.1, that much of the shared loss, for these datasets, occurs on the shared link next to the source. So, in Table 5 we also tabulate the average correlation computed by excluding the loss next to the source, that is, the packets lost by all the receivers.

The average correlation varies from 0.271 to 0.666. When the loss that is common to all receivers is deleted from the traces, the average correlation drops by almost an order of magnitude, in most cases. Thus, much of the spatially correlated loss is due to the loss close to the source. An exception to this, is the 2nd dataset, for which the average correlation remains greater than 0.3, despite ignoring loss close to the source. This is because of the presence of a lossy backbone link which experienced a long loss burst affecting most but not all the receivers.

From the results in Table 5 we can conclude that there is, on average, little pair-wise spatially correlated loss in almost all datasets, except for the spatial correlation due to the loss occurring next to the source.

4 Temporal Correlation of Loss at a Single Receiver

This section describes our findings regarding the burstiness of the packet loss. We discuss the extent to which packets are lost consecutively (in long loss bursts) and the extent to which there are solitary losses (a single lost packet preceded and followed by successful reception).

We notice a predominance of solitary losses in the distributions of the loss burst length seen by each of the receivers in our traces, as reported in [2]. It is also apparent that the lengths of the bursts span different timescales. The distribution of loss burst length can be divided into three regions: lengths of 1 to 6 packets, 7 to 100 packets and greater than 100. Most loss bursts affect just 1 to 6 consecutive packets (equivalent to 0.08 sec. to 0.48 sec.). This is the dominant mode in the distribution. A different mode affecting 7 to 10 packets (around 0.6sec.) is observed at some receivers. And, most significantly, we observe loss periods, 100 to 1000 packets long (equivalent to 8sec. to 3 minutes), at various receivers in every dataset.

First, we discuss the burstiness of loss for a single dataset in detail, describing the patterns observed. Then, we generalize our observations by showing summary statistics for all of the datasets.

Table 6 shows statistics for data collected on Dec 11, 1995. The source was "World Radio Network" which transmitted packets at $80ms.$ intervals. The loss rate, number of loss bursts, average loss burst length and coefficient of variation of burst length are given for each of the nine receivers. Burst length is defined as the number of consecutive packets lost. The coefficient of variation of the burst length is defined as

$$c = \frac{\sqrt{E[(b-\overline{b})^2]}}{\overline{b}} \qquad (3)$$

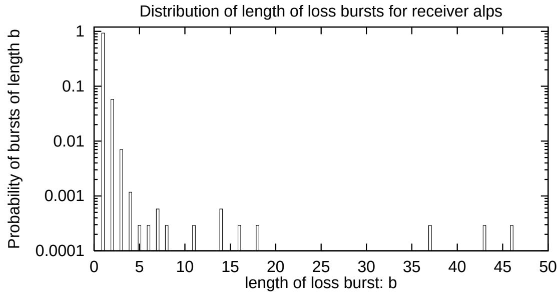

Figure 5: Distribution of loss burst length: `alps`(Georgia)

where b is the burst length or the number of consecutive losses and $\overline{b}$ is the mean burst length.

Table 6 also partially describes the distributions of the burst length by including the 99-percentile and the maximum burst length, for every receiver. For all receivers, both the median and the 75-percentile burst length were found to be 1, indicating the predominance of solitary losses. The table shows what percentage of the total loss is in bursts of length greater than 100. The 99-percentile burst length is low ranging from 2 to 8 consecutive packets. The length of the longest loss burst, on the other hand, is very high for five of the nine receivers. For example, `erlang` shows loss burst consisting of 2518 consecutive packets (equivalent to 3 minutes). There are thus a few extreme outliers, reflected in a coefficient of variation that is very high for some receivers. The extreme outliers, though infrequent, can contribute heavily to the total packet loss. For example, the burst of length 2518 seen by `erlang` accounts for 35% of its total loss. Receivers `alps`, `float` and `tove` received many duplicate packets. That is, almost half the packets received by each of them were duplicates. The other receivers received no duplicates. Duplication of packets was also noticed in some of the other datasets. In all cases, a packet is assumed to have been correctly received at a receiver if at least one copy of it is received.

Figures 5 and 6, show the distribution of the loss burst length for the receivers `alps` (in Georgia) and `cedar` (in Texas). These figures show the number of bursts with a given burst length. The y axis is a log scale so it is obvious that, in all three cases, the solitary bursts are the most frequent and that the probability drops sharply from burst length 1 to 6 , approximately as in a geometric distribution. In figure 5 for receiver `alps` there are stray bursts of length 7, 12, 37, 43, 46 etc.

In figure 6 for receiver `cedar` in Texas, there is an additional cluster of loss bursts of length 7 to 10 (around 0.6 sec.). Similar modes were observed in the burst length distributions of one or more receivers in many datasets. For example, in the 12th dataset two receivers in France saw the same concentration of bursts in the distribution, centered at 0.6sec. Upon taking a closer look at the timing of the losses, it becomes clear that this cluster of bursts was due to periodic loss occurring at 30 sec intervals, a phenomenon also discussed in [5]. This periodicity also shows

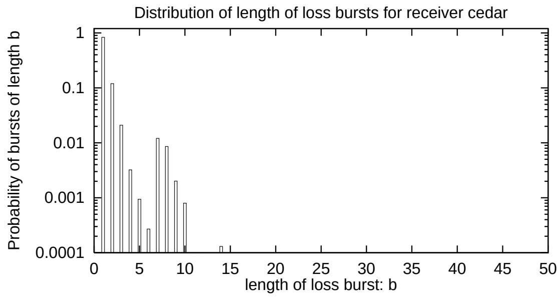

Figure 6: Distribution of loss burst length: `cedar`(Texas)

Data-set	median length	75-perc. length	95-perc. length	99-perc. length	maximum burst length
1.	1	1	2	3	758
worst			3	757	anhur,ursa 1161
3.	1	1	2	4	203
worst		3	8	12	collage 3697
4.	1	1	3	4	64
worst		2	4	12	erlang 1963
5.	1	1	2	3	9
worst			3	92	ursa 877
6.	1	1	2	9	1289
worst	2	3	6	19	cedar 1333
7.	1	1	2	3	515
worst			4	8	anhur,ursa 3342
8.	1	1	4	10	34
worst	2	3	9	20	ursa 662
9.	1	1	2	3	175
worst				8	pax 2603
10.	1	1	1	2	5
worst			7	24	float 260
11.	1	1	2	4	10
worst		2	3	8	pax 458
12.	1	1	2	2	26
worst		2	3	7	edgar 2050
2.	1	1	2	4	1648
worst		2	5	16	ursa 3434
13.	1	1	2	4	22
worst		2	5	9	erlang 18623
14.	1	2	6	10	3700
worst	2	4	14	49	ganef 3704

Table 7: Summary of Burstiness of Loss for all Datasets

up as an abrupt jump at 30 sec., in the autocorrelation function.

Table 7 summarizes the distribution of the length of the loss bursts over all datasets. The median loss burst length, the 75, 95 and 99-percentiles and maximum burst length are shown. For each dataset, the first line gives the median of the statistic over all receivers in the set and the second line gives the maximum value of the statistic and the names of the receivers which saw that maximum value. The first 11 datasets had sampling intervals of $80ms$., and the last three datasets had sampling intervals of $40ms$. In all our datasets, at least one receiver experienced a loss period of length greater than 200 (equivalent to 16 sec). In many cases, bursts of length greater than 1000 (equivalent to 1.3 minutes) were seen.

5 Conclusions and Future Work

This paper described the results of measurements of packet loss in the Mbone, a multicast network in widespread use. Measurements were taken for three sources in 14 data sets each collected on a different day. The data was collected simultaneously at up to 12 locations.

We presented a method for estimating the loss rates on the segments of the multicast transmission tree. We also presented two methods of judging the extent of spatial correlation between receivers: first, by plotting the distribution of the number of receivers that simultaneously lose a packet, making different assumptions about independence and topology and secondly, by computing the correlation coefficient between the loss at pairs of receivers. It was found that, in most datasets, the loss on the backbone links of the multicast transmission tree was small relative to the total loss seen by the receivers. The spatially correlated loss was small, on the average, *except* for the loss due to the link next to the source. A negligible number of packets were lost at the receiving hosts themselves.

With respect to temporally correlated losses, we found that a majority of the loss bursts were solitary losses. A few extremely long loss bursts greater than 8 sec. (or 100 packets) were also observed. At least one receiver saw one of these long loss bursts, in every dataset. Periodic bursts of length approximately 0.6sec.(8 consecutive packets) were observed for some receivers in some of the datasets.

A more thorough study of the loss in the different parts of the MBone by recording packets sent by sources in a greater variety of locations would indicate how widespread the loss patterns that we have observed are. The long loss bursts lasting for several seconds and minutes are of particular concern. It would be useful to pinpoint the reasons for such long outages and possibly find ways to remove them. Our traces can also be used directly in a simulation of a multicast network, to assess the performance of reliable multicast protocols. This would indicate which kinds of error-recovery methods are useful and in which situations. It would also show which aspects of the loss strongly affect the performance of reliable multicast protocols.

6 Acknowledgments

We would like to thank Henning Schulzrinne and Vinay Kumar for their immense help in the data collection, our colleague Sue Bok Moon for her support and encouragement, and to Andres Klemets on whose multicast recording tools, our data collection software is partly based.

We are grateful to Chuck Cranor of Washington Univ., St.Louis, Jon Crowcroft of University College London, Ron Erlich of Columbia Univ., Deborah Estrin of Univ. of Southern California , Serge Fdida of Institut Blaise Pascal, Jay Glicksman of Enterprise Integration Technologies, John Jendro of Portland State, Edward Knightly of Univ. of California at Berkeley, Simon Lam of Univ. of Texas, Jorg Liebeherr of University of Virginia at Charlottesville, Phillipe Nain of INRIA, Stephen Pink of Swedish Institute of Computer Science, Larry Peterson of Univ. of Arizona, Linda Prime of Univ. of Washington, Kishor Trivedi of Duke Univ., Tatsuya Suda of Univ. of California at Irwine, Jon Reid of University of Illinois at Urbana-Champaign, Satish Tripathi of Univ. of Maryland at College Park, Peter Wan of Georgia Institute of Technology, Raj Yavatkar of Univ. of Kentucky at Lexington and Lixia Zhang of Univ. of California, Los Angeles, for providing MBone-capable computer accounts which allowed us to take the measurements.

References

[1] P. Bhagwat, P. Misra, S. Tripathi, “Effect of Topology on Performance of Reliable Multicast Communication,” *Proc. IEEE Infocom 94,* (Toronto, June 1994), pp. 602 – 609.

[2] J. Bolot, H. Crepin, and A Vega Garcia, “Analysis of Audio Packet Loss in the Internet,” *Proc. 1995 Workshop on Network and Operaing System Support for Audio and Video,* pp. 163 - 174.

[3] S. Casner and S. Deering, “First IETF Internet Audiocast,” *ACM Computer Communication Review,* Vol. 22, No. 3 (July 1992), pp. 92 - 97.

[4] S. Casner “Frequently Asked Questions (FAQ) on the Multicast Backbone (MBONE)” available via ftp from ftp://ftp.isi.edu:mbone/faq.txt .

[5] S. Floyd and V. Jacobson, “The Synchronization of Periodic Routing Messages”, *ACM Trans. on Networking*, Vol. 2, No. 2 (April 1994), pp. 122 - 136

[6] S. Floyd, V. Jacobson and S. McCanne “A Reliable Multicast Framework for Light-weight Sessions and Application Level Framing” *Proc. of ACM SIGCOMM ’95*, Vol. 25, No. 4 (October 1995), pp. 342 - 356

[7] V. Paxson, “End-to-End Routing Behavior in the Internet” *to appear in Proc. 1996 ACM SIGCOMM Conf.*

[8] S. Pingali, D. Towsley, J. Kurose, “A Comparison of Sender-Initiated and Receiver-Initiated Reliable Multicast Protocols”, *Proc. 1994 ACM SIGMETRICS Conf.*, 221 – 230, May 1994.

[9] V. Jacobson, “Multimedia Conferencing on the Internet,” *Tutorial Notes - ACM Sigcomm94,* (London, Sept. 1994).