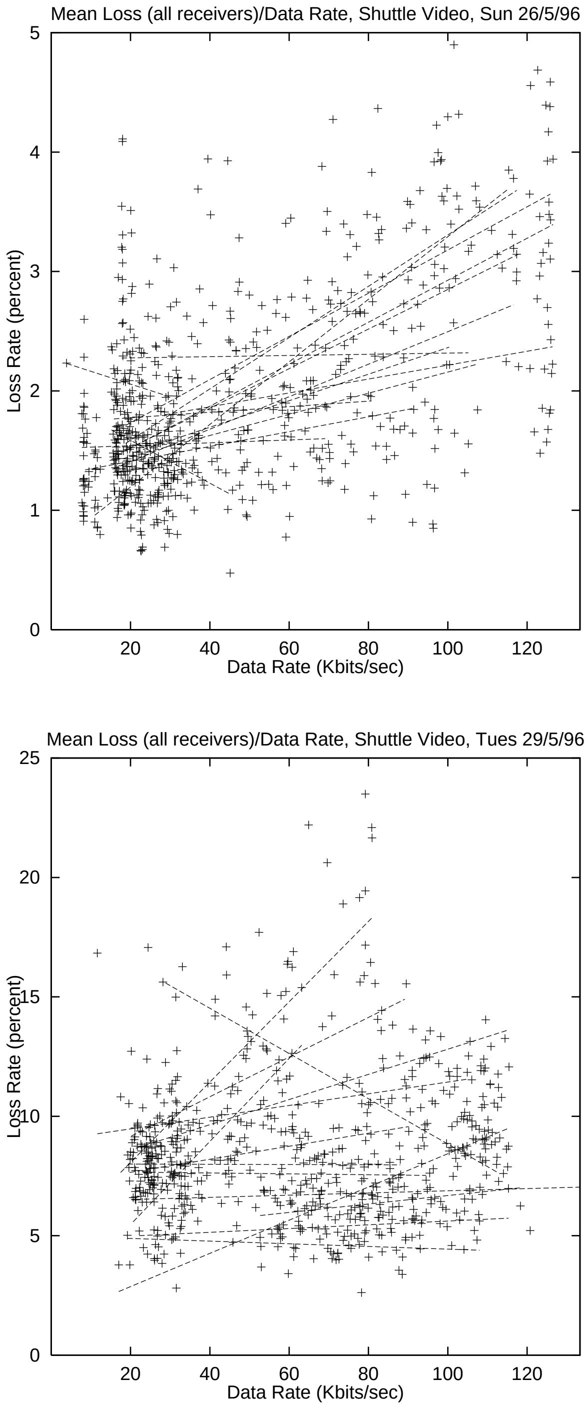

Figure 9: Loss Against Sent Data Rate

anism for multicast traffic that must adapt on significantly longer time-scales than TCP does. In particular, in the data for Tuesday 29th May, there are long periods of time where there is no discernible correlation between sent data rate and mean receiver loss rate. If even a fairly small proportion of receivers were showing such a correlation, then such a trend should be visible in this graph.

If this data is extrapolated, it shows that multicast applications that attempt to adjust so that no receiver is suffering significant loss will almost inevitably adjust so that they send no data (or whatever their minimum data rate is configured to). Thus schemes such as [2] are unlikely to be useful in the wide area mbone as we see it today.

Consecutive Losses

The graphs in figure 10 (A-C) show the frequency of RTP sequence number increments as observed at UCL on each of several days. We know from figures 6 and 7 that at least two processes combine to produce loss, and this is reflected in figure 10. The base losses appear to obey a simple rule where the probability p of a sequence number increase of $n, n >= 1$, is approximated by $p(n) = an^b$ where for the graphs shown, b is in the range -8.7 to -3.5. In addition there is a peak in all the graphs in figure 10 at around 5 to 15 consecutive losses which does not follow the simple rule. These peaks are consistent with the thirty-second periodic loss bursts observed in figures 6 and 7.

Figure 10D shows the conditional loss probability; that is given that packets x to $x + n - 1$ were lost, the probability that packet n is also lost[1]. As these curves show, packet losses are not independent, and tend to occur in longer, bursts than would be the case if they were independent. However, with the exception of the thirty-second periodic losses, the excess of bursts of 2-5 packet losses compared with what would be expected from random loss, although statistically significant, is not sufficient to greatly influence the design of most applications, and single packet losses still dominate.

In the case of applications such as UCL's Robust-Audio Tool[11] and INRIA's FreePhone audio tool that provide redundancy by carrying a highly compressed redundant copy of data one or more packets behind the primary encoding, these curves show that there is some gain from adding the redundant encoding more than one packet behind the primary. For example, from the data shown in 10C (16:00-18:00), 77243 packets were sent and 18640 (24.1%) were lost. Adding a single redundant payload one packet behind the primary encoding allows the reconstruction of audio from 7082 missing packets resulting in an effective loss rate of 15.0%. Adding the same redundant payload two packets behind the primary allows the

[1] n=0 gives the probability that a packet is lost if the previous packet was not lost