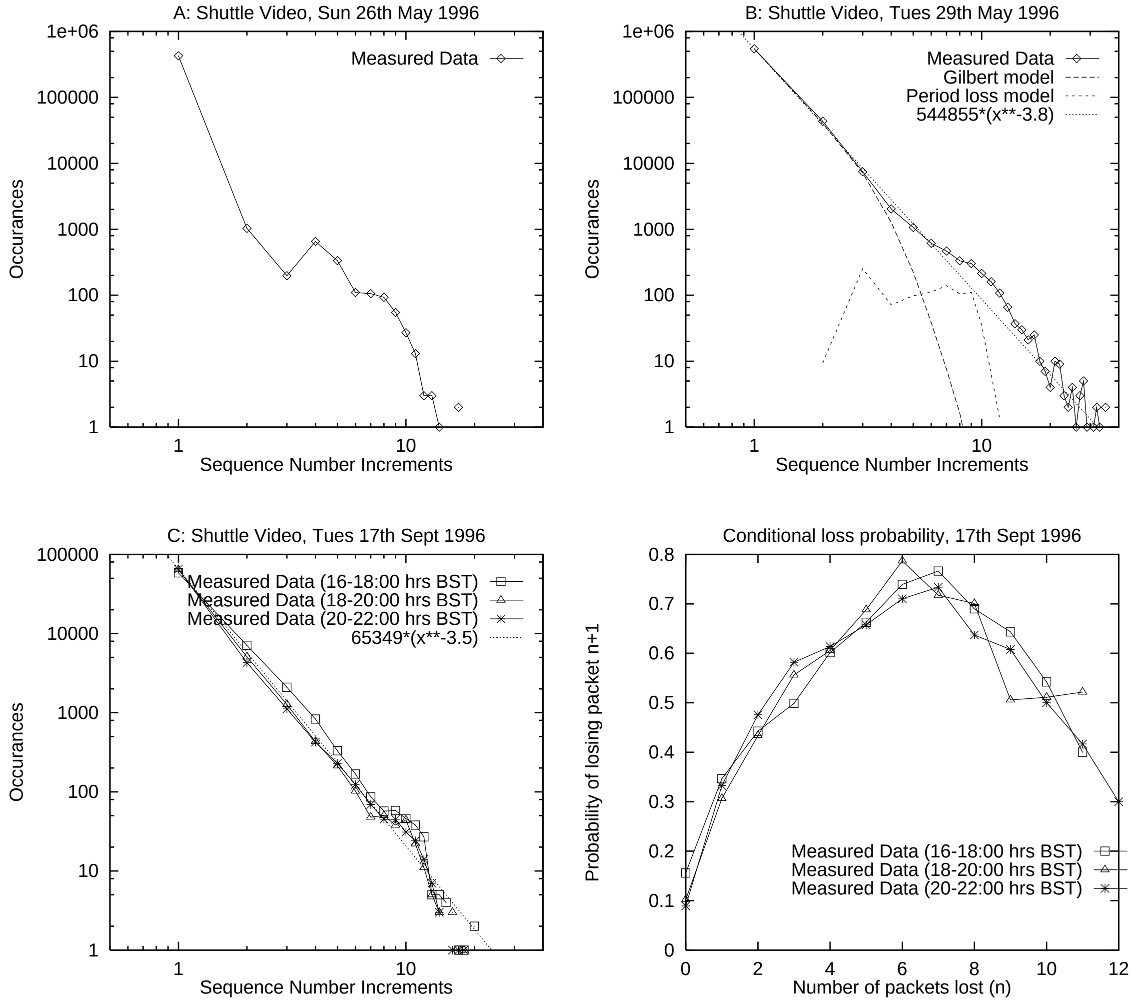

Figure 10: Frequency of Sequence Number Increments (Consecutive Losses - 1)

reconstruction 7885 packets resulting in a effective loss rate of 13.9%. Unfortunately this extra gain comes at the expense of extra delay in the case of this audio coding scheme so may or may not be worthwhile depending on the purpose of the stream itself.

Whilst these results agree with the conclusions of Bolot[3], the two-stage Gilbert model used only provides a first order approximation to these results. Figure 10B shows a two-state Gilbert model (figure 11, p=0.083, q=0.823).

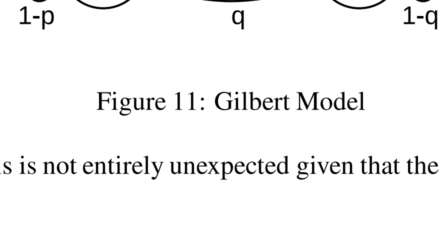

Figure 11: Gilbert Model

This is not entirely unexpected given that the loss process measured clearly consists of a congestive loss process which the Gilbert model tracks quite well and an additional periodic failure which the model does not take into account. Modeling this periodic failure using the traffic data rate distribution in figure 9B and assuming a 500ms failure every 30 seconds gives the distribution shown in 10B as Periodic model. This does go some way towards explaining the tail on these distributions, but does not do so entirely, so some other failure must be occurring fairly infrequently.

4 Conclusions

Drawing general conclusions from limited data is a process fraught with dangers. We believe that the RTCP data we have presented is reasonably representative of network conditions on the Mbone as a whole. The same general patterns show up in data gathered at different times and