

Modelling Packet Loss in RTP-based Streaming Video for Residential Users

Martin Ellis[‡], Dimitrios P. Pezaros[‡], Theodore Kypraios^{*} and Colin Perkins[‡]

[‡] School of Computing Science, University of Glasgow

^{*} School of Mathematical Sciences, University of Nottingham

Abstract—Packet loss is a major problem for real-time Internet applications. Markov models of packet loss are often used to develop and evaluate the performance of these applications. Despite their wide use, these models have not been validated in terms of how well they capture the loss conditions experienced by residential Internet users. We evaluate the accuracy of common packet loss models using traces of IPTV-like traffic measured on residential ADSL and Cable links, and find that these models are insufficient to capture the observed packet loss patterns. We introduce a new type of model, incorporating packet delay information, and show improved accuracy over previous models.

I. INTRODUCTION

Packet loss in residential broadband networks can be highly variable, and disrupts the user experience for real-time network applications such as Internet video and conferencing. Using models that capture the packet loss processes allows analytic and simulation studies to evaluate the performance of new applications and services, before full-scale deployment.

Previous studies have used Markov-chain models, such as the classical Gilbert model [3], to generate packet loss processes for evaluation of video streaming applications [15] and error repair [6]. More complex Markov-chain models [16], [5], and Hidden Markov Models (HMMs) [10], [11] have been proposed to describe packet loss on academic networks. However, the accuracy of these models for characterising packet loss on residential broadband networks remains unstudied. There is evidence that packet loss characteristics of residential networks differs from those of academic networks [2], so it is important to study if existing models are applicable in this context.

In this paper, we fit well-known Markov models to measurements of packet loss from residential Internet links [2]. We show that these models capture the measured loss conditions in many cases, but there are times when behaviour that they cannot capture is present. We introduce a new two-level model that uses both packet loss and delay information to better understand the state of the network, and show that this new model captures packet loss processes seen on residential ADSL and Cable links more accurately than previous models.

II. MODELLING RESIDENTIAL PACKET LOSS

We use the packet loss traces described in [2], taken from end-to-end measurements of streaming synthetic RTP traffic between 1–8.5Mb/s from a well-connected server to residential ADSL and Cable hosts. The dataset captures a range of different loss characteristics, with ~ 3800 traces varying between one and ten minutes in duration (between 6×10^3 and 6×10^5 packets

per trace). Sequence numbers and timestamps are attached to each packet, giving an observation sequence of packet losses, Z_i ($Z_i = 0$ for received packets, $Z_i = 1$ for lost) and queueing delay (DQ) estimates, that can be used as input to the models.

The first model we study is a two-state Markov chain, which we refer to as the Simple Gilbert Model (SGM) [3]. The states in the SGM, GOOD and BAD, directly represent Z_i (BAD always produces packet loss; GOOD never does). The SGM has been widely used to evaluate application performance since early work used it to model loss processes seen on academic networks [16]. Recent evaluations of video streaming [15] and video quality estimation tools [13] have used SGM models.

Since the SGM does not capture packet loss burstiness, the Extended Gilbert Model (EGM) was proposed [12]. Here, m states are used to represent losses, modelling loss bursts of up to m packets. State i ($i < m$) represents a loss burst of i packets, while state m represents a burst of *at least* m . Received packets are represented by state 0. We use $m = 5$, since $> 99\%$ of loss bursts in the traces are ≤ 5 packets [2].

The Gilbert-Elliott model [1] lets *both* states produce errors: GOOD representing isolated loss events; BAD modelling loss bursts. This is implemented as an HMM, using loss observations without knowing the state of the model. In HMMs, transition probabilities between hidden states and loss probabilities for each are estimated from observed data using Expectation-Maximisation algorithms [9]. HMMs can have more states, improving accuracy, but this is limited by estimation complexity. We focus on two- and three-state HMMs, as they are computationally feasible, and the states have a physical interpretation (i.e., network congestion state). In [10], HMMs of varying numbers of states were applied to loss traces, finding two- or three-state HMMs sufficient in most cases.

III. EVALUATING MODELS USING LOSS TRACES

We evaluate the SGM and EGM, which are widely used to model Internet packet loss [5]; and two- and three-state HMMs (2HMM and 3HMM), which aim to capture state changes, suggesting they are well-suited to model the variable loss patterns seen on residential links. To assess model accuracy, we test how well they capture the original data. We first estimate model parameters from each trace, using the process described in [5] for the SGM and EGM, and the R package *hmm.discnp* for the HMMs. Then, we use these parameters to simulate 1000 synthetic sequences, and compare these with the original data to assess goodness-of-fit. For each sequence, we calculate

a set of statistics, S_i , producing a distribution, S_i^{synth} , that is compared to S_i^{raw} (the value of S_i obtained from the raw data). As this technique generates new sequences using parameterised models, we refer to it as *parametric bootstrap*, in contrast to traditional bootstrap, which involves resampling within existing data. We then test the null hypothesis H_0 that the observed value of S_i^{raw} is a typical draw from the distribution S_i^{synth} , by calculating a central 95% confidence interval and checking if S_i^{raw} falls into that interval. If H_0 is not rejected, this suggests the model offers a good fit to the data, since realisations of the fitted model are similar to the observed data (in terms of S_i).

The performance metrics (S_i) we study are *mean packet loss fraction*, which describes the level of loss experienced in the traces, and the *loss and receive run-length distributions* which describe the loss patterns. We calculate the 5th, 25th, 50th, 75th, and 95th percentiles of the receive run-lengths (across the range of the distribution), and the mean, median, and max loss run-lengths (since these are less variable).

IV. PERFORMANCE OF SGM, EGM, AND HMMs

A. Examples of Loss Behaviour

Initial examination of the loss traces shows they can be divided into three categories: i) zero or very low loss; ii) non-bursty loss, or iii) bursty loss. Separating low loss traces allows further analysis to be more meaningful, focusing on bursty or non-bursty loss patterns that do not make sense with very low rate packet loss. Over 90% of traces which visually show very low loss had 15 or fewer lost packets, so we set the threshold at ≤ 15 losses. Traces with higher loss were classified by whether the loss is spread out in the trace, as in the top panel of Figure 1a; or confined in bursts separated by longer receive runs (Figure 1b). Here, we study 486 non-bursty and 433 bursty traces; of the rest, 1211 are very low loss, 1679 are loss free.

Figure 1 shows representative example traces of “non-bursty” and “bursty” loss behaviour. The top panel of each plot shows a measured loss trace, and the lower panels show example synthetic sequences generated by the models. Figure 1a shows that the synthetic sequences are comparable to the non-bursty loss in the trace, indicating that these models are suitable for non-bursty loss behaviour. Figure 1b shows a trace with bursty loss periods, separated by long receive runs. In this case, the SGM, EGM, and HMMs generate sequences that are quite different to the original data.

B. Parametric Bootstrap

Goodness-of-fit results from applying parametric bootstrap to the SGM, EGM, and HMMs are shown in Figure 2. These show, for each statistic S_i (y -axis), the number of traces where the model had “good fit”, in terms of S_i ; visually, longer bars mean that the model fits more traces. All models capture the mean loss rate, for both bursty and non-bursty traces. In terms of receive run-lengths, the SGM, EGM, and HMMs perform poorly for bursty traces, and a little better for non-bursty traces. Recall from Figure 1b that bursty loss was not well captured by these models. The HMMs aim to capture the changes in loss states and perform slightly better in bursty traces; this

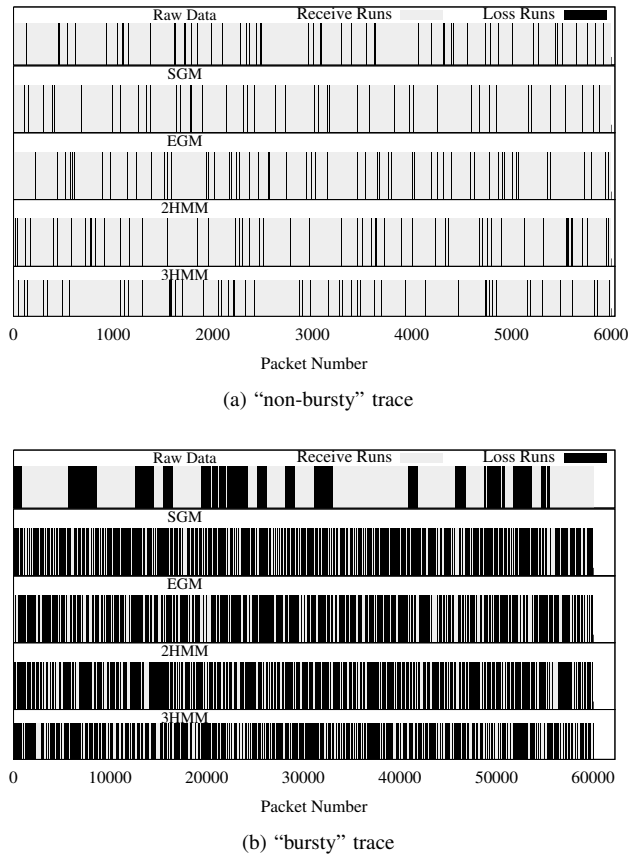


Figure 1. Original / synthetic loss sequences (loss models)

extra complexity does not help when the loss is not bursty. Patterns in loss run-length results are similar for both types of traces (with better performance on non-bursty traces). However, receive run-length statistics are not, leading to loss patterns that are quite different to the bursty traces (e.g., Figure 1b).

To summarise, these results show that although the SGM, EGM, and HMMs perform adequately in certain cases, they do not accurately capture the loss patterns of the most bursty traces (as demonstrated in Figure 1b). Since over 10% of traces were classified as bursty, and this has an impact on video performance, it must be accurately captured by the models.

V. A NEW TWO-LEVEL MODEL FOR PACKET LOSS

Bursty loss is often associated with congestion. We propose a new model that explicitly considers network load (similar to [14]), by incorporating measured queuing delay information [2]. We derive a *two-level* hierarchical model, with “outer” states representing network congestion state, and “inner” models to capture packet loss. The trace being modelled is split into fixed windows of one second, and the outer state for each window, Ω , is explicitly chosen using a simple classifier on loss and delay data. The Z_i observations from all windows classified as Ω are used to calculate the parameters of the inner packet loss model for Ω . Transitions between outer states are modelled as a Markov chain, with transition probabilities estimated by counting the transitions between windows of each outer state.

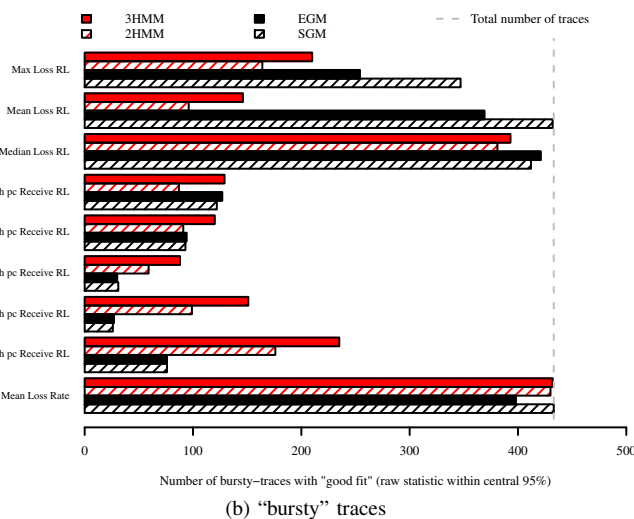



Figure 2. Parametric bootstrap results (loss models)

A. Pre-Classifying Network State

Figure 3 shows an example of the two-level model. Here, we chose three outer states, corresponding to different sources of packet loss along the measured end-to-end path [2]:

- electrical noise on access links causing higher-layer loss;
- edge congestion, where queue overflows at the ISP edge (i.e., DSLAM or CMTS) cause packet loss, and;
- congestive queue overflow within “core” networks.

Access link noise will cause low levels of uncongested loss, regardless of the levels of delay. Edge congestion causes higher levels of loss, associated with higher delay (since the building queues at edge routers will noticeably increase queuing delay DQ). Finally, core congestion also causes higher loss, but without noticeable effect in DQ (since higher statistical multiplexing at core routers means that the effects of queuing on DQ will be less obvious at the receiver). The outer states of the model correspond to these network conditions, modelled as a three-state Markov chain. In each outer state, packet loss is modelled by either the SGM or 2HMM as described in

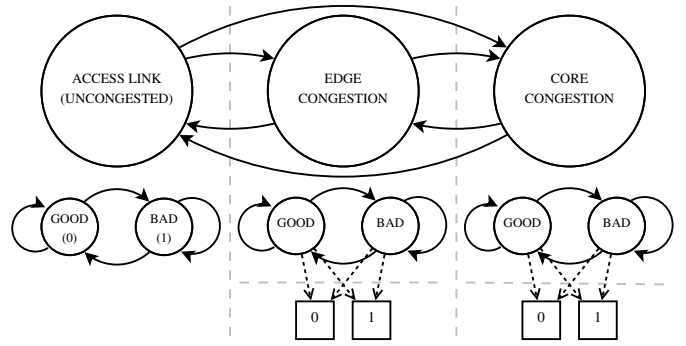


Figure 3. Two-Level Model, SGM/2HMM/2HMM configuration

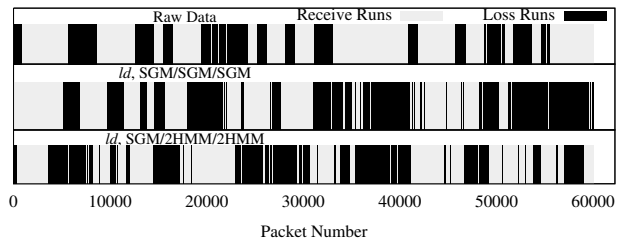


Figure 4. Example “bursty” trace (two-level loss/delay models)

Section II. The two configurations used for the inner models are SGM/SGM/SGM, where packet loss is modelled by an SGM in each outer state, and SGM/2HMM/2HMM, where the “congested” states are modelled by a 2HMM. Since Section IV showed that non-bursty loss is well-modelled by the SGM, it is suitable for uncongested loss due to access link noise.

B. A Threshold-Based Classification Scheme

We use a simple classification scheme (ld), based on loss and delay thresholds. This considers one-second time windows, examining the number of losses (N), number of loss bursts (M), and median DQ (\overline{DQ}) in each window. We use thresholds for N and M to identify periods of high loss (indicating congestion), and another threshold for \overline{DQ} to distinguish “core” and “edge” congestion. The choice of loss thresholds ($N > 2$ or $M > 2$ indicating congestion) is based on the assumption that non-congestive loss is unlikely to create more than two separate loss events with a one second window, and that loss bursts longer than two packets are likely due to congestion. The \overline{DQ} threshold (5ms) is also based on examination of the trace data; traces with non-bursty loss typically exhibit $DQ < 5$ ms. These thresholds are based on the dataset in [2]; further study is needed to confirm their generality.

VI. PERFORMANCE OF THE TWO-LEVEL MODEL

Figure 4 shows results from the two-level model on the loss trace from Figure 1b. This shows that the sequences generated by the new model are closer to the original data than those from the previous models. Figure 5 shows the results of applying parametric bootstrap to the non-bursty and bursty traces identified earlier. The two-level model has consistently

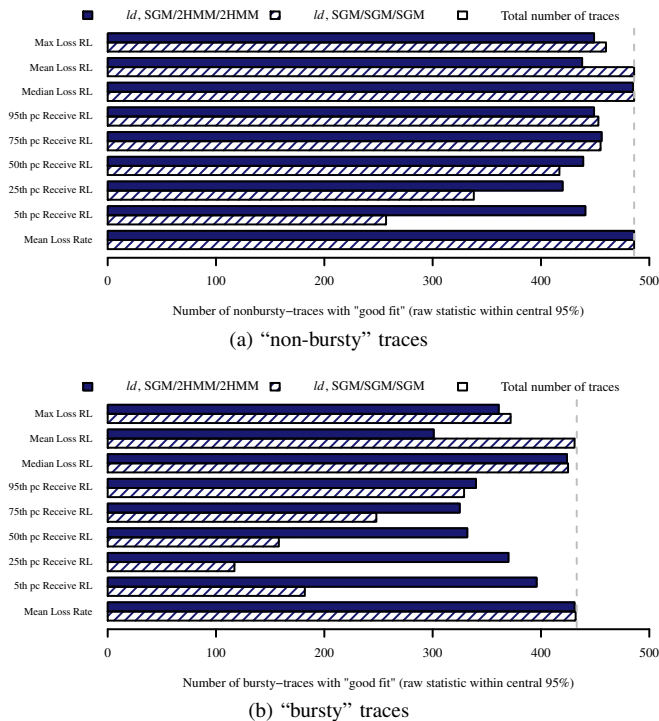


Figure 5. Parametric bootstrap results (two-level loss/delay models)

better performance for both classes, showing good fit for more traces than Figure 2. As before, the performance of the models on the non-bursty traces is better than for the bursty traces.

The SGM/SGM/SGM configuration improves slightly on the SGM, EGM, and HMMs using only loss data, but with receive run-length distributions still not well-modelled in many cases. However, the SGM/2HMM/2HMM configuration shows much improved performance over the previous models, with most traces being well-modelled in terms under every metric. This accurately captures the different modes of packet loss, using the most appropriate model for each.

These results show that the two-level models (using the SGM/2HMM/2HMM configuration) are more accurate than the previous SGM, EGM, or HMMs. For the non-bursty traces, the two-level models with the SGM/2HMM/2HMM configuration are suitable in almost all cases. In terms of the bursty traces, most of which were poorly modelled by the previous models, the two-level models show a clear improvement, as illustrated in Figure 4.

VII. RELATED WORK

Related models have been used for errors in wireless networks. These include a four-state Markov chain model, with separate states for long and short loss and receive run-lengths [17]; run-length thresholds were defined in terms of physical-layer characteristics of the channel, so the approach is not applicable to Internet losses. Markov-based trace analysis (MTA) [7] uses a data-preconditioning approach, similar to our pre-classification, classifying traces from GSM networks into lossy or loss-free sub-traces, then modelling these separately.

MTA and other Markov models were evaluated against a new alternative in [4], modelling loss and receive runs with lengths derived from mixtures of geometric distributions. This improved modelling accuracy, but only one trace was used in the evaluation, making generality of results unclear. Performance of the SGM, the four-state Markov-chain model, and MTA were compared in [8], using DVB-H traces, concluding that the four-state model is well-suited for DVB-H; however, they rely on manual parameter estimation, limiting applicability.

VIII. CONCLUSIONS AND FUTURE WORK

We evaluated the accuracy of commonly used packet loss models using traces from residential access networks, and found them insufficient to capture the bursty loss conditions present. We introduced a new two-level model to better capture these loss conditions, improving performance across all types of traces. Combining the simplicity of an SGM for non-congestive loss, and the power of HMMs for capturing bursty, congestive loss, the SGM/2HMM/2HMM configuration performs well.

Further work will compare the performance of the models for a real application (e.g., FEC performance), and investigate whether better accuracy can be obtained by improving on the simple classification algorithm in Section V-B.

Acknowledgements

This work was supported by Cisco Research and EPSRC. Thanks to Maurizio Filippone for his suggestion to use parametric bootstrap.

REFERENCES

- [1] E. O. Elliott, "Estimates of Error Rates for Codes on Burst-Noise Channels," *Bell Syst. Tech. J.*, vol. 42, no. 5, 1963.
- [2] M. Ellis *et al.*, "End-to-End and Network-Internal Measurements of Real-Time Traffic to Residential Users," in *Proc. ACM MMSys*, 2011.
- [3] E. Gilbert, "Capacity of a Burst-Noise Channel," *Bell Syst. Tech. J.*, vol. 39, no. 5, 1960.
- [4] P. Ji *et al.*, "Modeling frame-level errors in GSM wireless channels," *Perform. Evaluation*, vol. 55, no. 1-2, 2004.
- [5] W. Jiang and H. Schulzrinne, "Modeling of Packet Loss and Delay and Their Effect on Real-Time Multimedia Service Quality," in *Proc. NOSSDAV*, 2000.
- [6] S.-R. Kang and D. Loguinov, "Modeling Best-Effort and FEC Streaming of Scalable Video in Lossy Network Channels," *IEEE/ACM Trans. Netw.*, vol. 15, no. 1, 2007.
- [7] A. Konrad *et al.*, "A Markov-Based Channel Model Algorithm for Wireless Networks," *Wireless Networks*, vol. 9, no. 3, 2003.
- [8] J. Poikonen and J. Paavola, "Error Models for the Transport Stream Packet Channel in the DVB-H Link Layer," in *Proc. IEEE ICC*, 2006.
- [9] L. R. Rabiner, "A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition," *Proc. IEEE*, vol. 77, no. 2, 1989.
- [10] K. Salamatian and S. Vaton, "Hidden Markov Modeling for network communication channels," in *Proc. ACM SIGMETRICS*, 2001.
- [11] P. Salvo Rossi *et al.*, "Joint End-to-End Loss-Delay Hidden Markov Model for Periodic UDP Traffic Over the Internet," *IEEE Trans. Sig. Proc.*, vol. 54, no. 2, 2006.
- [12] H. Sanneck and G. Carle, "A Framework Model for Packet Loss Metrics Based on Loss Runlengths," in *Proc. SPIE/ACM MMCN*, 2000.
- [13] S. Tao *et al.*, "Real-Time Monitoring of Video Quality in IP Networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 5, 2008.
- [14] S. Tao and R. Guérin, "On-Line Estimation of Internet Path Performance: An Application Perspective," in *Proc. IEEE INFOCOM*, 2004.
- [15] P. U. Tournoux *et al.*, "On-the-Fly Erasure Coding for Real-Time Video Applications," *IEEE Trans. Multimedia*, vol. 13, no. 4, 2011.
- [16] M. Yajnik *et al.*, "Measurement and Modelling of the Temporal Dependence in Packet Loss," in *Proc. IEEE INFOCOM*, 1999.
- [17] Y. Yu and S. L. Miller, "A Four-State Markov Frame Error Model for the Wireless Physical Layer," in *Proc. IEEE WCNC*, 2007.