

A Framework Model for Packet Loss Metrics Based on Loss Runlengths *

H. Sanneck and G. Carle
GMD Fokus, Kaiserin-Augusta-Allee 31, D-10589 Berlin, Germany

ABSTRACT

For the same long-term loss ratio, different loss patterns lead to different application-level Quality of Service (QoS) perceived by the users (short-term QoS). While basic packet loss measures like the mean loss rate are widely used in the literature, much less work has been devoted to capturing a more detailed characterization of the loss process. In this paper, we provide means for a comprehensive characterization of loss processes by employing a model that captures loss burstiness and distances between loss bursts. Model parameters can be approximated based on run-lengths of received/lost packets. We show how the model serves as a framework in which packet loss metrics existing in the literature can be described as model parameters and thus integrated into the loss process characterization. Variations of the model with different complexity are introduced, including the well-known Gilbert model as a special case. Finally we show how our loss characterization can be used by applying it to actual Internet loss traces.

Keywords: Packet Loss, Short-Term QoS Metrics, Loss Burstiness, Internet Real-Time Services

1. INTRODUCTION

Existing work on Internet Quality of Service (QoS) identified delay, jitter and packet loss as metrics of interest for users and operators ⁽¹⁾. Among these, packet loss is an important metric in the statistically shared environment of the Internet, where the demand for a resource (such as bandwidth or buffer) may exceed the available capacity. For multimedia applications such as coded audio and video (live or stored), loss may also result from inordinate delay in the network which causes the play-out deadline of a sample or a frame to be missed. While existing work focuses on capturing the mean loss (long-term QoS), less emphasis is put on modelling loss distribution (short-term QoS, cf. ^(2,3) and the references therein). It is important to note that for the same mean loss, different loss patterns can produce different perceptions of QoS, as described in ⁽⁴⁻⁸⁾. Also, as shown by Cidon et al. ⁽⁶⁾ and Shacham/McKenney ⁽⁹⁾, many forward error recovery approaches become less efficient as the loss burstiness increases. Thus, it is important to be able to capture the actual loss process with suitable (and simple) metrics. As an example, Fig. 1 shows mean loss rates $p_m(s)$ for a voice stream versus its sequence number s averaged with a sliding window of five and 100 packets respectively ($p_5(s)$, $p_{100}(s)$). It can be seen that the distribution of loss rates (and thus the perceptual impact) over a small window size ($p_5(s)$) varies strongly. At the same time, the mean loss rate evaluated over the larger window size ($p_{100}(s)$) varies within a much smaller interval. The mean loss rate calculated over a large time interval is thus not suitable to reveal differences of perceived QoS.

Thus, our aim in this paper is to develop a model with the following properties:

- possibility of expressing well-known long- and short-term QoS metrics (like e.g. those of the Gilbert model);
- adjustable complexity dependent on specific application/network requirements, supporting additional metrics for detailed short-term QoS analysis;
- coverage of both loss burstiness, as well as distances between losses.

In the following section, we describe a characterization of the loss process, and develop an analytical framework using a model based on loss and no-loss run-lengths. In section 3, we show the applicability of the developed measures to actual Internet loss traces. Finally, we conclude in Section 4 with a summary and point out directions for future research.

* To appear in *Proceedings of the SPIE/ACM SIGMM Multimedia Computing and Networking Conference 2000 (MMCN 2000)*, San Jose, CA, January 2000




Figure 1. Mean Loss Rates for a voice stream averaged over 5 and 100 packets respectively

2. LOSS PROCESS CHARACTERIZATION AND SHORT-TERM QoS METRICS

In order to capture and control the loss process itself, short-term QoS metrics are needed. Such metrics are useful in evaluating the network performance as perceived by the end users. Existing work on QoS metrics (^{10–13}) does not provide a framework that integrates these metrics, provides a common notation and allows for their simple computation. In the following we define a model that serves as such a framework. We begin with needed definitions:

The *loss indicator function* for a certain flow[†] at a certain node dependent on the packet sequence number s is:

$$l(s) = \begin{cases} 0: & \text{packet } s \text{ is not lost} \\ 1: & \text{packet } s \text{ is lost} \end{cases}$$

We define a *loss run length* k for a sequence of k consecutively lost packets detected at s_j ($s_j > k > 0$) with $l(s_j - k - 1) = 0, l(s_j) = 0$ and $l(s_j - k + i) = 1 \forall i \in [0, k - 1]$, j being the j -th “burst loss event”. The occurrence of a loss run length k is given by o_k . Thus for a given number of packets a of a flow that experience $d = \sum_{k=1}^{\infty} k o_k$ packet drops, we have the *relative frequency* $p_{L,k} = \frac{o_k}{a}$ for the occurrence of a loss burst of length k and the mean loss rate $p_L = \sum_{k=1}^{\infty} k p_{L,k}$.

2.1. Loss run-length model with unlimited state space

We define the random variable X as follows: $X = 0$: “no packet lost”, $X = k$: “*exactly*[‡] k consecutive packets lost”, and $X \geq k$: “*at least*[§] k consecutive packets lost”. With this definition, we establish a loss run-length model (Fig. 2) with an unlimited (possibly infinite) number of states, which gives loss probabilities dependent on the burst length[¶]. In the model, for every additional lost packet which adds to the length of a loss burst a state transition takes place. If a packet is successfully received, the state returns to $X = 0$. Thus the state probability of the system for $k > 0$ is $P(X \geq k)$. Given the case of a finite number of arrivals for a flow a , introduced in the previous paragraph, we can approximate the state probabilities of the model for $k > 0$ by the cumulative loss rate $p_{L,cum}(k) = \sum_{n=k}^{\infty} p_{L,n}$ (Table 1).

[†]In our definition, a flow is an application-layer data stream. For IPv4 it can be identified by the tuple (*source address, destination address, protocol ID, source port, destination port*).

[‡]“Exactly” means that the two packets immediately preceding and following the k lost packets are not lost with probability 1.

[§]“At least” means that the packet immediately preceding the k lost packets is not lost with probability 1.

[¶]The basic model is similar to the one employed by Varma (¹⁴) and Hsu et al. (¹⁵).




Figure 2. Loss run-length model with unlimited state space ($m \rightarrow \infty$)

Loss run-length model (unlimited states)	a arrivals	$a \rightarrow \infty$
burst loss ($k > 0$)	$p_{L,k} = \frac{o_k}{a}$	$P(X = k)$
mean loss	$p_L = \sum_{k=1}^{\infty} k p_{L,k}$	$E[X]$
cumulative loss ($k > 0$)	$p_{L,cum}(k) = \sum_{n=k}^{\infty} p_{L,n}$	$P(X \geq k)$ (state prob.)
conditional loss ($k > 1$)	$p_{L,cond}(k) = \frac{p_{L,cum}(k)}{p_{L,cum}(k-1)} = \frac{\sum_{n=k}^{\infty} o_n}{\sum_{n=k-1}^{\infty} o_n}$	$P(X \geq k X \geq k-1)$ (state transition prob. $p_{(k-1)(k)}$)
burst loss length ($k > 0$)	$g_k = \frac{o_k}{\sum_{n=1}^{\infty} o_n}$	$P(Y = k)$
mean burst loss length	$g = \frac{d}{\sum_{k=1}^{\infty} o_k} = \frac{\sum_{k=1}^{\infty} k o_k}{\sum_{k=1}^{\infty} o_k} = \sum_{k=1}^{\infty} k g_k$	$E[Y]$

Table 1. QoS measures for the loss run-length model with unlimited state space

The matrix of state transition probabilities for this model is given by

$$\begin{bmatrix} p_{00} & p_{01} & 0 & \cdots & 0 \\ p_{10} & 0 & p_{12} & & \vdots \\ p_{20} & 0 & 0 & \ddots & \vdots \\ \vdots & \vdots & \vdots & \ddots & 0 \\ p_{(n-2)0} & 0 & \cdots & 0 & p_{(n-2)(n-1)} \\ p_{(n-1)0} & 0 & 0 & \cdots & 0 \end{bmatrix}$$

The transition probabilities for $k > 1$ which can also be described as conditional loss probabilities can be computed easily as:

$$p_{(k-1)(k)} = P(X \geq k | X \geq k-1) = \frac{P(X \geq k \cap X \geq k-1)}{P(X \geq k-1)} = \frac{P(X \geq k)}{P(X \geq k-1)}$$

Again, if the burst loss occurrences o_k constitute a statistically relevant dataset, we can compute approximations for the conditional loss probabilities as given in Table 1.

Additionally, we also define a random variable Y which describes the distribution of burst loss lengths with respect to the burst loss events j (and not to packet events like in the definition of X). $E[Y]$ then is the expected mean burst loss length (loss gap). Table 1 shows the performance measures of the loss run-length model for a finite number of




Figure 3. Loss run-length model with limited state space $[(m + 1)$ states]

arrivals a using the loss run length occurrences o_k , as well as the relation to the transition/state probabilities of the model ($a \rightarrow \infty$) with the random variables X and Y . The cumulative loss rate for $k = 0$ is defined as the “no loss” case (corresponding to $P(X = 0)$):

$$p_{L,cum}(k = 0) = 1 - \sum_{k=1}^{\infty} p_{L,cum}(k) = 1 - \sum_{k=1}^{\infty} \frac{\sum_{n=k}^{\infty} o_n}{a} = 1 - \sum_{k=1}^{\infty} \frac{k o_k}{a} = 1 - p_L$$

2.2. Loss run-length model with limited state space

To assess the performance of a network with respect to real-time audio and video applications, a model with a limited number of states is sufficient. This is due to the fact that real-time audio and video applications have strict requirements and cannot use a network service with a significant number of “long” loss bursts. For these applications it is desirable to use only few model parameters, and to focus on key aspects of the loss process. In addition, memory and computational capabilities of the system that performs modelling have to be taken into account (see also paragraph 2.6).

For these reasons we derive from the basic model a loss run-length model with a finite number of states $(m + 1)$. Fig. 3 shows the Markov chain for the model. Table 2 gives performance measures similar to those in Table 1^{||}, however the state probability for the final state m ($P(X = m)$) and the probability for a transition from state m to state m are added. For $0 < k < m$, $X = k$ represents as before “*exactly* k consecutive packets lost”. Due to the limited memory of the system, the last state $X = m$ is just defined as “ m consecutive packets lost”. Thus $P(X = m)$ can be seen as a measure for the “loss over a window of size m ” (independently of actually larger loss run lengths).

Figure 4 shows the base measures used to compute the loss run-length based metrics. In Figure 4 (a), each point indicates whether there was a loss (1) or not (0), representing the loss indicator function. Figure 4 (b) shows the loss run lengths. Figure 4 also shows some of the state transitions when a given loss trace is applied to a model of either $m = 2$ or $m \geq 4$. With $m = 2$ for a loss burst of length $k = 4$, the system is three times ($k - m + 1$, see Fig. 4 (c)) in state 2, and thus two $(k - m)$ transitions $m \rightarrow m$ occur. This leads to the computation of $p_{L,m}$ and $p_{L,cond}(m)$ (as approximations for $P(X = m)$ and $P(X = m|X = m)$ respectively) as given in Table 2.

Interestingly, Miyata et al. ⁽¹²⁾ propose precisely $p_{L,m}$ as a performance measure for FEC-based audio applications. This “sliding window” of m consecutively lost packets allows to reflect specific applications’ constraints, e.g. m can be set to the lowest number of consecutively lost packets for which a complete audio “dropout” is perceived by a user. Then, larger loss bursts do not have a higher impact and thus do not need to be taken into account with their exact size. We extend the above approach by looking at the occurrence of a certain number of packets lost within the window of length m . This allows e.g. to assess how effective FEC protection applied to groups of packets would be without keeping track of the actual Application Data Unit (ADU) association of the individual packets. In section 1 we introduced the mean loss rate over a sliding window of length m ($p_m(s)$) which can be defined as the convolution of the analysis window with the loss indicator function (Fig. 5):

^{||}The burst loss length measures are computed as in Table 1 and are therefore not shown.

Loss run-length model ($m + 1$ states)	a arrivals	$a \rightarrow \infty$
burst loss ($0 < k < m$)	$p_{L,k} = \frac{o_k}{a}$	$P(X = k)$
burst loss ($k = m$) loss over window m	$p_{L,m} = \sum_{n=m}^{\infty} \frac{(n - m + 1)o_n}{a - m + 1}$	$P(X = m)$ (state probability)
mean loss	$p_L = \sum_{k=1}^{\infty} \frac{k o_k}{a}$	$E[X]$
cumulative loss ($0 < k < m$)	$p_{L,cum}(k) = \sum_{n=k}^{\infty} \frac{o_k}{a}$	$P(X \geq k)$ (state probability)
conditional loss ($1 < k < m$)	$p_{L,cond}(k) = \frac{p_{L,cum}(k)}{p_{L,cum}(k-1)} = \frac{\sum_{n=k}^{\infty} o_n}{\sum_{n=k-1}^{\infty} o_n}$	$P(X \geq k X \geq k-1)$ (state transition prob. $p_{(k-1)(k)}$)
conditional loss ($k = m$)	$p_{L,cond}(m) = \sum_{n=m}^{\infty} \frac{(n - m)o_n}{d - m}$	$P(X = m X = m)$ (state transition prob. p_{mm})

Table 2. QoS measures for loss run-length model with limited state space [($m + 1$) states]

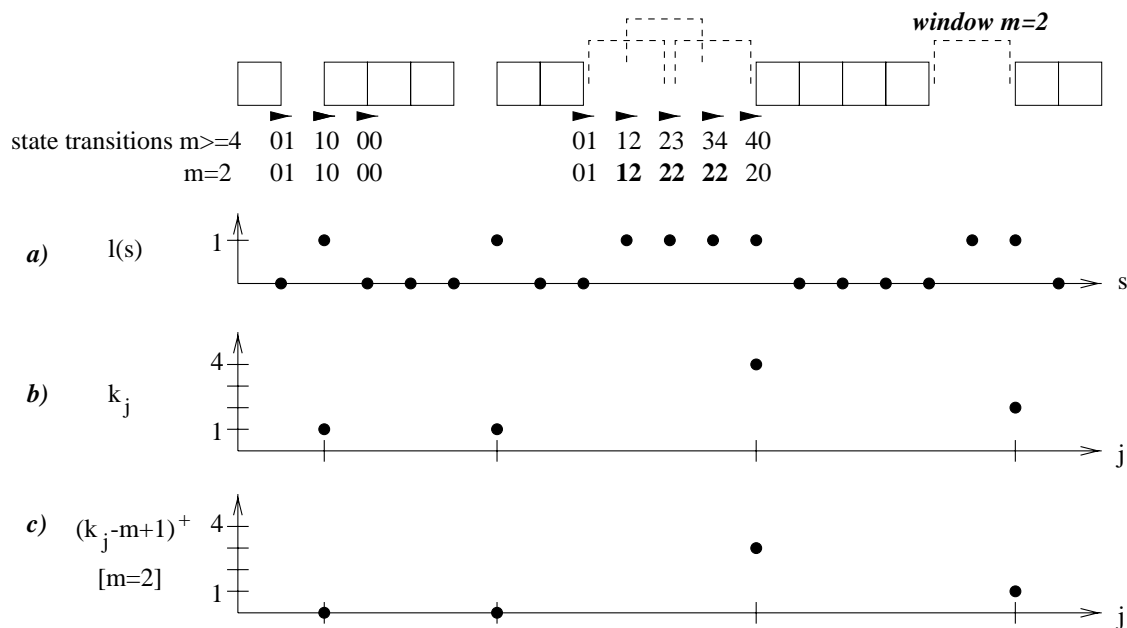


Figure 4. Basic loss measures




Figure 5. $p_m(s)$: mean loss rate over a sliding window of length m

$$p_m(s) = \frac{l(s) * w_m(s)}{m} = \frac{\sum_{\nu=0}^a l(\nu)w_m(s - \nu)}{m}$$

The tradeoff following from the above formula can be described as follows: computing the actual histogram of $p_m(s)$ (solid line in Fig. 5 c) captures the accurate sequential relation of the loss bursts, however makes the measure still depend on s (this is the approach taken in ⁽¹⁶⁾). When using only the sum of histograms calculated separately over each individual loss burst (dashed line Fig. 5 c) some information is lost, however only the tracking of loss bursts k is needed. To describe the latter approach we use $o_{w,m}(k)$ which describes the occurrence of k consecutive packets lost within the window of length m

$$o_{w,m}(k) = \begin{cases} (m - k + 1)o_k + 2 \sum_{n=k+1}^{\infty} o_n & 0 < k < m \\ \sum_{n=m}^{\infty} (n - m + 1)o_n & k = m \end{cases}$$

Summing over the weighted $o_{w,m}(k)$ we get in fact the overall mean loss rate (see the Appendix for the proof):

$$\sum_{k=1}^m \frac{k o_{w,m}(k)}{m a} = \sum_{k=1}^{\infty} \frac{k o_k}{a} = p_L \quad (1)$$

Similar window-based measures were proposed also in ⁽¹⁷⁻¹⁹⁾.

2.3. Gilbert Model

For the special case of a system with a memory of only the previous packet ($m = 1$), we can use the runlength distribution for a simple computation of the parameters of the commonly-used Gilbert model (Fig. 6) to characterize the loss process (X being the associated random variable with $X = 0$: “no packet lost”, $X = 1$ “a packet lost”). Then the “loss over window m ” is equal to the mean loss or unconditional loss probability, and only one conditional loss probability (transition $1 \rightarrow 1$) is defined.




Figure 6. Loss run-length model with two states (Gilbert model)

Gilbert	a arrivals	$a \rightarrow \infty$
burst loss ($k = 1$)	$p_L = \sum_{k=1}^{\infty} \frac{k o_k}{a}$	$P(X = 1)$
loss over window 1	mean loss rate	unconditional loss prob.
conditional loss ($k = 1$)	$p_{L,cond} = \frac{\sum_{n=1}^{\infty} (n-1) o_n}{d-1}$	$P(X = 1 X = 1)$
burst loss length ($k > 0$)	$g_k = \frac{o_k}{\sum_{n=1}^{\infty} o_n}$	$P(Y = k)$
mean burst loss length	$g = \frac{d}{\sum_{k=1}^{\infty} k o_k} = \frac{1}{1 - p_{L,cond}}$	$E[Y] = \frac{1}{1 - P(X=1 X=1)}$

Table 3. QoS measures for the loss run-length model with two states (Gilbert model)

The matrix of transition probabilities of the Gilbert model is:

$$\begin{bmatrix} p_{00} & p_{01} \\ 1 - P(X = 1|X = 1) & P(X = 1|X = 1) \end{bmatrix}$$

The Gilbert model implies a geometric distribution for residing in state $X = 1$. For the probability of a burst loss length of k packets we have $P(Y = k) = P(X = 1|X = 1)^{k-1}(1 - P(X = 1|X = 1))$, $k > 0$. The mean burst loss length $E[Y]$ is therefore:

$$E[Y] = \sum_{k=0}^{\infty} k P(X = 1|X = 1)^{k-1}(1 - P(X = 1|X = 1)) = \frac{1}{1 - P(X = 1|X = 1)}$$

For a finite number of arrivals we then can define a packet loss gap g that corresponds to the values for the higher order loss run-length model:

$$g = \frac{1}{1 - p_{L,cond}} = \frac{\sum_{k=1}^{\infty} k o_k}{\sum_{k=1}^{\infty} k o_k - \sum_{k=1}^{\infty} (k-1) o_k} = \frac{d}{\sum_{k=1}^{\infty} o_k}$$

These values can be compared to the respective values for the higher order loss run-length models to see how well the actual loss process is approximated by the simple two state model. Table 3 shows the performance measures for the Gilbert model.

2.4. No-loss run-length model with limited state space

User perception is not only affected by the length of burst losses ($k \in [1, m]$), but also by the length of the intervals between consecutive losses. We therefore define a *no-loss run length* K detected at s_J ($s_J > K > 0$) with $l(s_J - K - 1) = 1, l(s_J) = 1$ and $l(s_J - K + i) = 0 \forall i \in [0, K - 1]$, J being the J -th “no-loss burst event”. As in paragraph 2.2, we limit K to an interval $[1, M]$ dependent on the application. For audio, M could e.g. be set to the lowest value for which consecutive loss events are perceived by the user as being separate rather than a single

distortion of the signal. The occurrence of a loss distance K is given by o_K .

By defining a random variable X' as: $X' = 0$ if a packet was lost, $X' \geq k$ if at least k packets have *not* been lost, we can derive the same state model as the loss run-length model with finite state space for the “no-loss” case. Once m consecutive packets have been served (meaning not lost), the following packet arrivals (state transition: $m \rightarrow m$) are not taken into account in terms of the distance to the previously lost packet. Similarly to Tables 1-3 we can define model parameters and QoS measures for the no-loss run-length model. Additionally, we also have a random variable Y' which describes the distribution of no-loss lengths with respect to the no-loss events J . Of particular interest here is the relative frequency of a no-loss length K : $G_K = \frac{o_K}{\sum_{N=1}^{\infty} o_N}$ ($P(Y' = K)$ for $a \rightarrow \infty$).

2.5. Composite metrics

Obviously, both no-loss and loss models of any order can be combined to form a single model. Additionally it is possible to define metrics based on both no-loss and loss events. An example is a measure which already exists in the literature (^{10,18}) called the *noticeable loss rate (NLR)*. *NLR* defines a *loss distance constraint* (which is the no-loss runlength model order M) above which losses are excluded from the measure (are said to be not “noticeable”). Since the loss distance constraint must be at least one, all the losses in a loss run-length (except the first dependent on the distance to the previous loss) are said to be noticeable. Thus, using the previously introduced variables the *NLR* can be defined as follows:

$$NLR_M = \frac{d - \sum_{K=1}^{M-1} o_K}{d} = 1 - \frac{\sum_{K=1}^{M-1} o_K}{\sum_{k=1}^{\infty} k o_k}$$

2.6. Parameter computation

In this section we have demonstrated how to use loss and no-loss runlengths to compute models ranging from two states (Gilbert model) over $m + 1$ states to a potentially infinite state space. However, our formulas used the assumption that all (no-)loss burst lengths up to potentially infinite length are tracked. In a real system however, we clearly need to limit the maximum tracked burst length as a tradeoff between needed model complexity to assess the network performance with regard to specific applications, and memory or computational limitations.

Therefore we can limit the tracing of runlengths up to a length γ , with $\gamma \geq m$. Typically γ will be set according to the highest model order required ($\gamma = m$). This results in $p_L = \sum_{k=1}^{\gamma} \frac{k o_k}{a} + \frac{d_{\gamma}}{a}$, where d_{γ} are the packet drops which occur in bursts with higher lengths than γ . $p_{L,cum}(k) = \sum_{n=k}^{\infty} \frac{o_n}{a} = \sum_{n=k}^{\gamma} \frac{o_n}{a} + \frac{e_{\gamma}}{a}$, where e_{γ} are burst loss events with bursts larger than length γ . Thus essentially two additional counters are necessary, which keep track of $e_{\gamma} = \sum_{k=\gamma+1}^{\infty} o_k$ as well as $d_{\gamma} = \sum_{k=\gamma+1}^{\infty} k o_k$, rather than the individual o_k values.

3. APPLICATION OF THE METRICS

To demonstrate the previously introduced metrics, we performed measurements on a congested 22 hop path between Germany and France. Voice streams were generated by an audio tool and measured (RTP encapsulation, 160 bytes payload). The number of packets sent was between 170000-230000 for each of the measurements. We do not claim that these measurements are “typical”, however the results seem to tie in well with extensive measurement studies like (^{20,21}) and the measurement results of (²²). By visual inspection of a sliding window average (as in section 1, however with a window size of 1000) we checked the traces for non-stationarity (abrupt changes in the smoothed loss rate, linear increase or decay seen over the entire trace) before applying our models with limited state spaces.

Figure 7 shows an exemplary measurement with $P(X = 1) = 0.0418$ and $P(X = 1|X = 1) = 0.4694$ giving values for the raw data $P(X = k)$, the two-state Gilbert model $P(X = 1|X = 1)^{k-1}(1 - P(X = 1|X = 1))$ and the




Figure 7. Performance parameters for an example measurement 4

state and state transition probabilities for the model with limited ($k < m$) and unlimited state space. Additionally the state probability $P(X = m)$ for the model with limited states is given. We see that the probability $P(Y = k)$ to lose exactly k consecutive packets in a burst loss event drops geometrically fast in an interval of approximately $[1, 10]$. In this case, a loss run-length model confirms that the loss process is approximated well by the Gilbert model $P(X = 1|X = 1)^{k-1}(1 - P(X = 1|X = 1))$. For larger bursts ($k > 10$), the loss burst probabilities are significantly larger than for the Gilbert model. Thus the loss process in that area is underestimated. However, the absolute values of the loss probabilities are several orders of magnitude smaller than for the singleton loss ($k = 1$) case and do not seem to follow a specific distribution. Therefore it is not necessary that this area is considered by a model. The conditional loss probabilities $P(X \geq k|X \geq k - 1)$ increase with increasing loss burst length k , however their values are already very close to 1 for $k \gtrsim 10$ and stay there (in the shown area). This means that (as mentioned above) only few burst loss events larger than 10 packets take place and thus models with a higher number of states do not give much additional information. This can also be seen from the state probability curve $P(X \geq k)$, where values in the mentioned region do not change significantly. Fig. 7 also shows $P(X = m)$, the final state probability for models of order $m + 1$, which gives the probability that all packets within a window of length m are lost. This curve is very close to the "raw data" burst loss probabilities $P(X = k)$ for $k > 35$, which makes clear that no significant probability for burst losses larger than that exist, i.e. the data trace contains no "outages".

4. CONCLUSIONS

In this paper, we have presented a framework model which integrates both known and novel loss metrics and thus allows to capture the effect of loss distribution on continuous media applications. We constructed and parametrized models of different complexity that capture loss characteristics fully or partially, with the well-known Gilbert model

being a special case of these models. We showed how to directly relate model parameters to runlength-based data tracing. By applying the run-length model to measurement traces of IP voice flows, we demonstrated the tradeoffs between accurate multi-parameter modelling and simple but coarse modelling. We conclude that for Internet QoS support within applications and gateways an “intermediate model” is needed, which has to be more complex than the simple Gilbert model. The needed complexity of such an intermediate model is determined by the application requirements.

For future work, the autocorrelation of the loss indicator function, as well as the autocorrelation and composite metrics like the crosscorrelation function of the loss/no-loss run-lengths⁽²¹⁾ remain to be evaluated with regard to their practical usability and their relation to our model. Furthermore, it seems promising to combine the model for loss characterisation presented in this paper with results on user perception of multimedia applications like e.g. in^(23,24). This allows a precise and relevant characterisation of application-level QoS based on network level performance metrics, and opens a way for novel approaches to QoS support, which enforce certain short-term loss characteristics (patterns) on the stream rather than looking exclusively at a long-term loss rate^(10,25,26).

ACKNOWLEDGMENTS

We would like to thank Rajeev Koodli, Nokia Research Center, Boston, for many valuable discussions.

REFERENCES

1. V. Paxson, G. Almes, J. Mahdavi, and M. Mathis, “Framework for IP performance metrics,” RFC, IETF, May 1998. <ftp://ftp.ietf.org/rfc/rfc2330.txt>.
2. H. W. Chu and D. H. K. Tsang, “Dynamic bandwidth allocation for VBR video traffic in ATM networks,” in *Proceedings of ICCCN*, pp. 306–312, 1997.
3. H. Zhu and V. S. Frost, “In-service monitoring for cell loss Quality of Service violations in ATM networks,” *IEEE ACM Transactions on Networking* 4(2), pp. 240–248, 1996.
4. M. Bjorkman, A. Latour-Henner, U. Hansson, and A. Miah, “Controllability and impact of cell loss process in ATM networks,” in *Proceedings of IEEE GLOBECOM*, pp. 916–920, 1995.
5. J. C. Bolot and A. Vega Garcia, “The case for FEC-based error control for packet audio in the Internet,” *ACM Multimedia Systems*, 1997.
6. I. Cidon, A. Khamisy, and M. Sidi, “Analysis of packet loss processes in high-speed networks,” *IEEE Transactions on Information Theory* 39, pp. 98–108, January 1993.
7. M. Meky and T. N. Saadawi, “Degradation effect of cell loss on speech quality over atm networks,” in *Broadband Communications*, IFIP, Chapman and Hall, pp. 259–271, 1996.
8. I. E. G. Richardson and M. J. Riley, “Usage parameter control cell loss effects on MPEG video,” in *IEEE ICC*, pp. 970–974, 1995.
9. N. Shacham and P. McKenney, “Packet recovery in high-speed networks using coding and buffer management,” in *Proceedings ACM SIGCOMM '90*, pp. 124–131, (San Francisco, CA), June 1990.
10. R. Koodli and C. Krishna, “Noticeable loss: A metric for capturing loss pattern for continuous media applications,” in *Internet Routing and Quality of Service*, S. Civanlar, P. Doolan, J. Luciani, R. Onvural, Editors, *Proceedings SPIE Vol.3529A*, (Boston, MA), November 1998.
11. Z. Liu, P. Nain, and D. Towsley, “Bounds on finite horizon QoS metrics with application to call admission,” in *Proceedings IEEE INFOCOM '96*, (San Francisco, CA, USA), April 1996.
12. T. Miyata, H. Fukuda, and S. Ono, “New network QoS measures for FEC-based audio applications on the Internet,” in *Proceedings IEEE IPCCC 1998*, pp. 355–362, (Tempe/Phoenix, AZ, USA), February 1998.
13. C. Partridge, “A proposed flow specification,” RFC 1363, IETF, September 1992. <ftp://ftp.ietf.org/rfc/rfc1363.txt>.
14. V. Varma, “Testing speech coders for usage in wireless communication system,” in *Proceedings of IEEE Speech Coding Workshop*, pp. 93–94, (Montreal), 1993.
15. C. Hsu, A. Ortega, and M. Khansari, “Rate control for robust video transmission over wireless channels,” in *Proceedings of Visual Communications and Image Processing (VCIP)*, pp. 1200–1211, (San Jose, CA), February 1997.

16. M. Borella, D. Swider, S. Uludag, and G. Brewster, "Internet packet loss: Measurement and implications for end-to-end QoS," in *Proceedings of the International Conference on Parallel Processing*, August 1998.
17. S. Ono, T. Miyata, and H. Fukuda, "Loss metrics of grouped packets for IPPM," Internet Draft, IETF IPPM Working Group, August 1998. <ftp://ftp.ietf.org/internet-drafts/draft-ono-group-loss-00.txt>.
18. R. Koodli and R. Ravikanth, "One-way loss pattern sample metrics," Internet Draft, IETF IPPM Working Group, June 1999. <ftp://ftp.ietf.org/internet-drafts/draft-ietf-ippm-loss-pattern-01.txt>.
19. R. Nagarajan, J. Kurose, and D. Towsley, "Finite-horizon statistical Quality-of-Service measures for high speed networks," *J. High Speed Networks*, December 1994.
20. M. Yajnik, J. Kurose, and D. Towsley, "Packet loss correlation in the Mbone multicast network," in *Proceedings IEEE Global Internet 1996 (Jon Crowcroft and Henning Schulzrinne, eds.)*, pp. 94–99, (London, England), November 1996.
21. M. Yajnik, S. Moon, J. Kurose, and D. Towsley, "Measurement and modelling of the temporal dependence in packet loss," Technical Report 98-78, Department of Computer Science, University of Massachusetts, Amherst, 1998.
22. J.-C. Bolot, H. Crépin, and A. Garcia, "Analysis of audio packet loss in the Internet," in *Proceedings of the 5th International Workshop on Network and Operating System Support for Digital Audio and Video*, pp. 163–174, (Durham, NH), April 1995.
23. D. Reininger and R. Izmailov, "Soft quality of service with VBR+ video," in *Proceedings of 8th International Workshop on Packet Video (AVSPN97)*, (Aberdeen, Scotland), September 1997.
24. H. Sanneck and N. Le, "Speech property-based FEC for Internet Telephony applications," in *Proceedings of the SPIE/ACM SIGMM Multimedia Computing and Networking Conference 2000 (MMCN 2000)*, (San Jose, CA), January 2000. <ftp://ftp.fokus.gmd.de/pub/glone/papers/Sann0001:Speech-FEC.ps.gz>.
25. H. Sanneck and G. Carle, "Predictive loss pattern queue management for Internet routers," in *Internet Routing and Quality of Service, S. Civanlar, P. Doolan, J. Luciani, R. Onvural, Editors, Proceedings SPIE Vol.3529A*, (Boston, MA), November 1998. <ftp://ftp.fokus.gmd.de/pub/glone/papers/Sann9811:Predictive.ps.gz>.
26. H. Sanneck and G. Carle, "A queue management algorithm for intra-flow service differentiation in the "best effort" Internet," in *Proceedings of the Eighth Conference on Computer Communications and Networks (ICCCN '99)*, (Natick, MA), October 1999. <ftp://ftp.fokus.gmd.de/pub/glone/papers/Sann9910:Intra-Flow.ps.gz>.

Appendix

Proof of formula 1:

$$\begin{aligned}
& \sum_{k=1}^m \frac{k o_{w,m}(k)}{ma} = \\
&= \frac{1}{ma} \left[\sum_{k=1}^{m-1} \left(k(m-k+1) o_k + 2k \sum_{n=k+1}^{\infty} o_n \right) + m \sum_{n=m}^{\infty} (n-m+1) o_n \right] \\
&= \frac{1}{ma} \left[\sum_{k=1}^{m-1} k(m-k+1) o_k + 2 \left(\sum_{k=1}^{m-1} k \sum_{n=k+1}^{m-1} o_n + \sum_{k=1}^{m-1} k \sum_{n=m}^{\infty} o_n \right) + m \sum_{n=m}^{\infty} (n-m+1) o_n \right] \\
&= \frac{1}{ma} \left[\sum_{k=1}^{m-1} k m o_k - \sum_{k=1}^{m-1} k(k-1) o_k + 2 \left(\sum_{k=1}^{m-1} \frac{k(k-1)}{2} o_k + \frac{m(m-1)}{2} \sum_{n=m}^{\infty} o_n \right) + \sum_{n=m}^{\infty} n m o_n - m(m-1) \sum_{n=m}^{\infty} o_n \right] = \\
&= \frac{m \sum_{k=1}^{\infty} k o_k}{ma} = p_L
\end{aligned}$$