

Copyright © 2008 IEEE.

Reprinted from the Proceedings of the Australasian Telecommunication
Networks and Applications Conference, 2008:pp.22-27

This material is posted here with permission of the IEEE. Such permission of the IEEE does not in any way imply IEEE endorsement of any of the University of Adelaide's products or services. Internal or personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution must be obtained from the IEEE by writing to pubs-permissions@ieee.org.

By choosing to view this document, you agree to all provisions of the copyright laws protecting it.

On the Correlation of Internet Packet Losses

Hung X. Nguyen, Matthew Roughan

School of Mathematical Sciences, Adelaide University, Australia
{hung.nguyen, matthew.roughan}@adelaide.edu.au

Abstract—In this paper we analyze more than 100 hours of packet traces from Planet-Lab measurements to study the correlation of Internet packet losses. We first apply statistical tests to identify the correlation timescale of the binary loss data. We find that in half of the traces packet losses are far from independent. More significantly, the correlation timescale of packet losses is correlated with the network load. We then examine the loss runs and the success runs of packets. The loss runs are typically short, regardless of the network load. We find that the success runs in the majority of our traces are also uncorrelated. Furthermore, their correlation timescale also does not depend on the network load. All of these results show that the impact of network load on the correlation of packet losses is non-trivial and that loss runs and success runs are better modelled as being independent than the binary losses themselves.

I. INTRODUCTION

Packet losses severely deteriorate the performance of Internet applications. It is not only the average loss rate that is important but the correlation of these losses can also have profound effects. For example, highly correlated packet losses even at a low rate can severely degrade the performance of multi-media applications such as VoIP and Video streaming [7]. Correlation can also reduce the accuracy of network measurements [11]. Knowledge of the higher order statistics of Internet losses is therefore extremely important for application developers and network operators. In the last 15 years, several measurements test-beds such as the NIMI test-bed [10] have been built to measure Internet packet performance including packet losses. Thanks to this infrastructure, significant progress has been made in our understanding of Internet packet losses.

Recent studies of packet losses [1], [7], [13], [14] unanimously agree that the loss process is well represented as an alternating on-off process. When the process is in the “on” state, all packets are dropped; whereas when it is in the “off” state, no packet is dropped. The popular Gilbert model [6] is well suited to capture this typical behavior of Internet packet losses. The Gilbert model is essentially a k -state Markovian model. The number of states k in the Gilbert model is used to capture the amount of correlation in the packet losses. Even though the Gilbert model is commonly used, the literature has diverged on the number of states k needed in the Gilbert model. Indeed, each packet trace requires a different value of k and little is known about the factors that influence this parameter.

In this paper, we present a first step toward understanding the impact of network conditions on the correlation timescale of Internet losses. The answer of course is not simple as

there are many different aspects of the network that can have influence on k . Background traffic, buffer sizes and the queuing policies all can significantly change the correlation timescale of packet losses. All of these network entities however are difficult to study using active measurements. We instead concentrate on studying the impact of network load on the correlation scale of packet losses. Obviously, using active measurement we do not know the real load on the network path, so we use the loss rate itself as an indicator of the load.

We perform measurements on multiple Internet paths and apply statistical tests to study the correlations of Internet losses. We extend the findings in [1], [7], [13], [14] by examining these properties relatively to the load. We study in details the following properties on our traces:

- Correlation between binary losses,
- Correlation between runs of lost packets (*loss run*),
- Correlation between runs of successful packets (*success run*),
- The distributions of the lengths of the loss and success runs.

We explicitly concentrate on the impact of network load on these properties of the loss process and our contributions are as follows.

- We show that binary packet losses are correlated in most of our traces and their correlation timescales are strongly dependent on the network load.
- On the contrary, loss runs and success runs are uncorrelated in most of our traces regardless of the loss rates.
- We find that loss runs are short whereas success runs are geometrically distributed in the majority of the traces.
- We also find that even though the loss process can be modelled as an alternating renewal process regardless of the network load, its Markovian properties do depend on the network load.

The rest of the paper is organized as follows. In Section II we review existing work on characterizing loss behaviors on the Internet. In Section III, we define the network performance model and the problem of estimating correlation coefficient of a sequence of random variable. We introduce, in Section IV, a statistical test for the correlation timescale of a time series. In Section V we present the analysis of real data from Planet-Lab. We conclude the paper in Section VI.

II. RELATED WORK

The early work on characterizing Internet packet losses dated back to 1993 [2]. This paper, together with [9] a

few years later, studies packet correlation through the conditional loss probabilities of UDP and TCP data/ACK packets. Bolot [2] finds that for packets sent with spacing of less than 200ms, this conditional probability is much higher than the marginal probability. Paxson [9] confirms this finding and shows that the second packet in two consecutive TCP packets is much more likely to be lost if the first is lost.

Andr en *et al.* [1] show that packet losses in most of their traces have a simple structure: losses have weak correlation and are clustered into small groups. They also note that in some of the traces packet losses show evidence of long range dependence. That is, in these traces significant correlations exist even at large time lags.

Yajnik *et al.* [13] examine the correlation timescale of Internet losses using measurements from the University of Massachusetts at Amherst to several academic destinations around the world. They find that the correlation timescale is around 1000ms. That is, two samples that are separated in time by 1000ms or more are independent but samples closer than 1000ms show significant correlations. In the same paper, the authors fit a Markovian model of order k (with 2^k states and 2^{k+1} possible transitions) to the traffic traces and find that for most traces, 2^6 states are enough to fully capture the correlation of the packet losses. However, some traces require a Markovian model of order up to 40. The authors, however, did not provide explanations for the different values of k in their models.

Jiang *et al.* [7] introduce the extended Gilbert model that uses many fewer states to capture loss correlations. The main idea is to base the future loss probability on the previous k losses but not the previous no-loss data. The extended Gilbert model essentially exploits the on-off nature of Internet losses. They again find that different numbers of states are needed to represent their traces. For example, for most traces, an extended Gilbert model with 5 states is enough to accurately capture the correlation in the packet losses. However for some traces more than 10 states are needed.

Zhang *et al.* [14] perform similar studies with longer and more extensive data sets. Confirming the earlier results of [13], they find that loss correlation only happens on the timescale of less than 1000ms. Zhang *et al.* [14] also consider the runs of lost packets and find that in most of their traces, the loss runs are IID. Furthermore, the arrivals of the loss runs are well modeled as Poisson. They conclude that losses are mainly caused by trains of consecutive losses, rather than intervals over which loss rates become elevated and “nearby” but non-consecutive packets are lost. Losses are therefore better thought of as spikes during which there is a short-term outage, rather than epochs over which a congested buffer remains perilously full. This result confirms the appropriateness of the extended Gilbert model for Internet losses. Regarding the size of the loss runs, they find that loss run lengths are short, uncorrelated and follow a geometric distribution. The authors of [14] also note that the above observations are highly dependent of the network load.

III. PERFORMANCE MODEL AND PROBLEM DEFINITION

We consider the measurement of packet losses on an Internet path. The loss process on the path is a binary stochastic process $I(t)$: $I(t) = 1$ if at time t the buffer is full so that an arriving packet would be dropped and $I(t) = 0$ otherwise. In practice, $I(t)$ cannot be observed directly. The most common method to obtain samples from $I(t)$ is to use active probes: IP packets that are sent on the path to determine its quality. Using N of these active probes, we obtain N samples I_1, \dots, I_N , one for each probe at times t_1, \dots, t_N .

In this paper, we assume that the loss process is stationary. The statistical properties of loss processes are then well-defined and can be estimated from the samples I_1, \dots, I_N . For example, the average loss rate is defined as $p = E[I(t)]$ and p is estimated as the average statistics of I_1, \dots, I_N :

$$p^* = \frac{I_1 + \dots + I_N}{N}. \quad (1)$$

Note that p^* is an unbiased estimator of the true mean of the stochastic process I_1, \dots, I_N given that this process is stationary and ergodic. Furthermore, when the samples satisfy the weak lack of anticipation (WLA) requirement [8] (this technical requirement essentially means that the background traffic cannot predict when the samples are taken) the estimator p^* converges to the mean p of the process $I(t)$ as $N \rightarrow \infty$.

Higher order statistics such as correlations at different time scales can also be estimated from I_1, \dots, I_N . If the variance $\mathbb{E}[(I_i - p)^2]$ is finite, the autocorrelation function for $\{I_i\}$ is defined as

$$r_h = \mathbb{E}[(I_{i+h} - p)(I_i - p)] / \mathbb{E}[(I_i - p)^2], \quad (2)$$

where h is the time lag. Note that the time lag h in (2) is defined in term of the number of sample packets. To get the lag in time units, we need to multiply this number by the sample interval. Furthermore, in active measurements we do not know the mean p . Instead, we use the sample mean p^* in (1) for the true value of p in (2). Thus, r_h is estimated from the samples I_1, \dots, I_N as

$$r_h^* = \frac{\sum_{i=1}^{N-h} I_{i+h} I_i - p^{*2}}{\sum_{i=1}^N I_i^2 - p^{*2}}. \quad (3)$$

One way of looking at the losses is to consider the losses as periods of consecutive losses (loss run) and consecutive successful packets (success run). We denote the length of a loss run by Y_i . Y_i is measured by the number of consecutive sample packets that observe losses. Similarly, the length of a success run is denoted by Z_i and is measured by the number of consecutive no-loss packets. Y_i and Z_i can be easily obtained from the binary outputs I_1, \dots, I_N . Consider a simple example with 10 samples of the binary loss process $\{0, 1, 0, 0, 1, 1, 0, 0, 0, 1\}$. On this trace, we can identify two loss runs and two success runs. We therefore obtain two pairs $(Y_1, Z_1) = (1, 2)$ and $(Y_2, Z_2) = (2, 3)$ from this trace. Note here that we ignore the first $I_1 = 0$ and last sample $I_{10} = 1$ as the length of the corresponding success and loss runs cannot not be fully specified.

The correlation timescale k of a time series is the smallest time lag beyond which the value of the sample autocorrelation function becomes statistically insignificant. That is, $r_h = 0$ for all $h > k$.

We are interested in the statistical properties of the binary process $\{I_i\}$ and the alternating view $\{(Y_i, Z_i)\}$. Denote by \mathcal{Y} the distribution function of Y_i and \mathcal{Z} the distribution function of Z_i . In addition to testing for statistical properties of $I(t)$, we also investigate the distributions \mathcal{Y} and \mathcal{Z} .

IV. TESTING FOR CORRELATION

In this paper, we are primarily interested in finding the correlation timescale of $\{I_i\}$, $\{Y_i\}$, $\{Z_i\}$. To find the correlation timescale k of a time series, we apply statistical tests to determine if the time series is correlated at all time lags $h < k$. We start with $h = 1$ and keep increasing h until we reach the first time lag $h = k$ such that no significant correlation can be detected at a significant level α . This test could potentially underestimate the correlation timescale because the autocorrelation function could show dependence at some higher lag that has not been tested. We assume this error is small and use this method in this paper.

Finding the correlation timescale of a time series therefore boils down to testing whether the time series are independent at a given lag. Testing for independence is a classical problem in statistics. There are many statistical tests for independent properties of a sequence of random variables [4]. In this paper, we apply the popular autocorrelation function (ACF) test.

Let $\{X_i\}_{i=1}^{\infty}$ be a stationary sequence of random variables. The autocorrelation function of $\{X_i\}$ can be computed using (2). For a stochastic process that is independent at lag h , the autocorrelation function r_h at that lag is zero. Note however that due to randomness, a particular realization of such process may not have zero sampled r_h , even though this value is small. The following Proposition gives an idea of what constitutes a significant value for the sample autocorrelation function for a given number of samples N [3].

Proposition 1: For large N , the sample autocorrelations of an independent and identically distributed (IID) sequence with finite variance are approximately IID with a normal distribution $\mathcal{N}(0, 1/N)$.

If X_i are IID then approximately 95% of the sample autocorrelations should fall between the confidence bounds of $1.96/\sqrt{N}$. Using Proposition 1 we can then test the following hypothesis:

Hypothesis: $\{X_i\}$ is independent at lag h

The test statistic is calculated from the sample autocorrelation function at lag h as follow.

$$S = r_h * \sqrt{N}. \quad (4)$$

If the hypothesis is true then S has a limiting normal distribution $\mathcal{N}(0, 1)$. Thus, if $|S| > 1.96$ the hypothesis is rejected at the significant level 0.05. In this case, there is not enough evidence to prove that the samples are independent at lag h . Note however that when rejecting the hypothesis that

$\{X_i\}$ is independent at lag h , we still cannot claim that the samples are correlated at that time lag.

V. DATA ANALYSIS

In this section, we apply the test in Section IV to real measurement data collected from nodes on Planet-Lab [12]. Our main objective is to understand the correlation properties of the loss process $I(t)$.

The traces are obtained by sending probes between a pair of nodes randomly selected from 100 Planet-Lab hosts. These nodes are located mainly at universities around the world. 50% of them (50 nodes) are in the US, 40% (40 nodes) are in Europe and the other 10% are in South America, Asia and Australia.

Once the source and destination pair is selected, probes are sent from the source to the destination as a Poisson process with rate λ . We vary the sampling rate in our measurements. The probes are UDP packets with a payload of 40 bytes. This payload is used to store the packet sequence numbers, from which the receiver can identify the lost packets. We store the binary output for each probe.

We performed our measurements in a 3-day period from June 25, 2008 to June 28, 2008. In total, we successfully obtained 95 end-to-end traces. Out of these traces, 79 traces (80%) experience zero loss and 2 have loss rates of 1. The latter two clearly represent some problems with the measurement hosts, whereas the traces with zero losses do not provide any information for our analysis. We therefore exclude these traces from our data set. The remaining traces are summarized in Table I.

A. Stationarity Test

We first test the traces for stationarity. We use a simple test that is similar to the one used in [13]. We first divide the traces into segments of two-hour long and check each segment for stationarity. A time series is said to be stationary in the wide-sense if the mean and covariance function remain constant over time. Since there is no rigorous way to test for stationarity, we use the simple heuristic of checking whether the average loss rate varies significantly in the trace segments. We smooth the trace using a finite moving average filter with a window size of 1000 packets. We then apply two tests to the smoothed trace to test for stationarity. First, if there is any abrupt increases or decrease of greater than 0.05, the trace is then classified as non-stationary. Second, to test the gradual trend in the average loss rate, we fit a linear straight line to the data using the least-squares straight line fit. If there is a total change in the average loss rate of 0.015 or greater, over a two hour trace segment, the trace then is considered to be non-stationary.

Out of the 61 trace segments in our experiment, 48 are found to be stationary according to our tests. The details of these traces are again given in Table I. Among these 48 traces, 30 have less than 100 packet losses out of approximately 70000 packets over a two hour period. We again eliminate these trace segments because they do not provide enough losses for meaningful statistical analysis. In the end we have 18

TABLE I
TRACE DESCRIPTIONS.

Traces	Source	Destination	Sampling Rate	Duration	# Segments	# Stationary Segments
1	UIUC,USA	Intel-Berkeley,USA	5 packets/sec.	2 hours	1	1
2	Fraunhofer, Germany	UC-Berkeley	10 packets/sec.	3 hours	2	1
3	University of Washington,USA	UC-Berkeley,USA	10 packets/sec.	25 hours	13	11
4	Polytechnic,USA	University of Bern, Switzerland	10 packets/sec.	10 hours	5	1
5	JAIST,Japan	University of Bern, Switzerland	10 packets/sec.	6 hours	3	3
6	Fraunhofer,Germany	UC-Berkeley, USA	15 packets/sec.	26 hours	13	12
7	MPG-SWS,Germany	Perdue University, USA	15 packets/sec.	26 hours	13	13
8	Universidad Republicá,Uruguay	NAIST,Japan	15 packets/sec.	26 hours	13	11
9	Umass Amherst,USA	SICS,Sweden	6.25 packets/sec.	48 hours	24	19
10	Umass Amherst,USA	Los Angeles,USA	50 packets/sec.	2.5 hours	1	1
11	Umass Amherst,USA	Atlanta,USA	25 packets/sec.	6 hours	3	3

trace segments with meaningful losses. Since in our traces, the average loss rates are small (below 0.015), we add to our traces, the three traces (20Nov97, 20Dec97, and 26Jun98) of [13] that have significant loss rates (above 1.5%) to study lossy traces. These traces are measured using unicast probes, similarly to our traces. Altogether we have 39 stationary trace segments with meaningful losses.

We then apply the statistical test in Section IV to these traces to learn their correlation properties.

B. Results

1) *Correlation Timescale of The Binary Losses:* We first look at the binary losses $\{I_i\}$. We apply the ACF test to compute the correlation timescale of the binary data. In Figure 1, we plot a histogram of the correlation timescale for the 39 trace segments in our data. About 50% of the traces have lag 0, which means that the samples are uncorrelated. The samples in the other traces exhibit correlation at different timescales. Some are correlated up to a lag of 25 packets. These results are consistent with the findings in [13], [14].




Fig. 1. Correlation timescale of the stationary trace segments. Half of the traces have uncorrelated samples. Other traces have different correlation timescales.

We now investigate the dependency of correlation timescale on the network load. Recall here that in active measurements we do not know the real load on a network path, so we use the average loss rate as an indicator of the load. Our traces have different sampling rates. To compare the traces, we look at the correlation timescale in time unit rather than in packets. We

plot in Figure 2 the correlation timescale versus the average loss rates for our Planet-Lab traces and for the traces in [13] (Umass traces).

Figure 2 shows that the correlation timescale of the binary loss data is positively correlated with the average loss rates for Planet-Lab traces (left) but negatively correlated for Umass traces (right). The difference between the two sets of traces can come from multiple reasons. First, in the Planet-Lab traces probes are sent following a Poisson process whereas in the Umass traces probes are sent periodically. The opposite trends in the two sets of traces could also be caused by the different measurement platforms, background traffic (for example, losses caused by TCP versus those caused by UDP cross traffic), etc.

The results in this section suggest that the binary loss process is far from independent and that the average loss rate (which is representative of the traffic load) cannot be used to reliably predict the correlation timescale of the binary loss process. Other factors such as background traffic, probing process need to be taken into account when computing the correlation timescale. The above observations help explaining the many different numbers of states needed for the extended Gilbert model in the literature [7], [13], [14].

2) *Correlation Timescale of The Successful and Lost Packet Runs:* In this section, we view the loss process $\{I_i\}$ as an alternating process of successful and loss runs $\{(Y_i, Z_i)\}$. We investigate the statistical properties of the loss run lengths Y_i and success run lengths Z_i . These properties include: (i) their correlation timescales-especially their dependence on the network load, and (ii) the distributions of Y_i and Z_i . Note here that the success run lengths also represent the inter-arrival time of the loss runs and vice-versa.

We first look at the loss runs. We plot in Fig. 3 the correlation timescale of the lengths of the loss runs $\{Y_i\}$. We observe that loss run lengths are independent in all but 2 of our traces. Therefore, even though the binary packet losses are not independent their runs are independent regardless of the network conditions.

We next look at the successful packet runs. We plot in Figure 4 the correlation timescale of the success run lengths $\{Z_i\}$.

Lengths of the successful packet runs are independent in

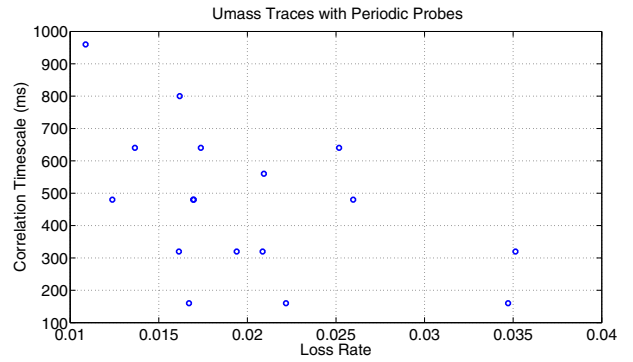
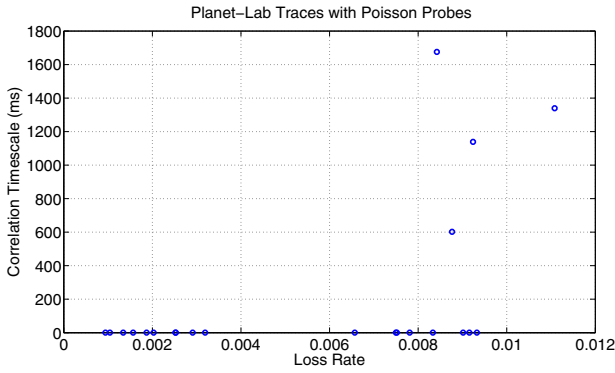


Fig. 2. Correlation function of the loss process under different loads. For the Planet-Lab traces, the higher the load, the less correlated the loss process becomes. The opposite is true for the Umass traces.




Fig. 3. Correlation of the loss runs under different loads. Contrary to the binary loss data, the loss runs are independent at all loads.




Fig. 4. Correlation of the successful packet runs under different loads. Contrary to the binary loss data, the correlation timescale of the success runs are not dependent on the network load.

about two third of the traces. For those traces that have significant correlation among the success runs, we observe that the correlation timescale of the success runs does not depend on the network load (reflected in the loss rate). This result indicates that the success runs are affected more by other network conditions than the actual load of the network.

Finally, we look at the correlation between loss runs $\{Y_i\}$ and success runs $\{Z_i\}$. The correlation of the pair (Y_i, Z_i) in

all of our traces is very weak (the cross correlation coefficient between $\{Y_i\}$ and $\{Z_i\}$ is below 0.07 in all of the traces), indicating that the lengths of adjacent loss run and success run bear little relationship to each other.

The finding in this section suggests that even though the binary losses are far from independent. The loss process can be approximated as an alternating renewal process with periods of independent loss runs and independent success runs.

3) *Distribution of Loss and Success Run Length:* In all of our traces, the loss run lengths are short. Most of them are 1. The length of the loss interval of the background traffic are therefore less than the sampling intervals, which in all traces are more than 2ms. To study the length of the loss runs beyond this time scale, more fine-grain measurements with shorter sampling intervals are needed. We plan to extend our measurements to higher probing rates for this purpose in our future work.

We consider next the distribution of the success run lengths. As an exploratory test, we plot in Figure 5 the log of the complementary cumulative distribution of the success run lengths for two traces with different loss rates. We observe that the success run lengths are geometrically distributed under different loads. Similar curves are obtained for other traces.

To test the hypothesis that these run lengths are indeed geometrically distributed we apply the Anderson-Darling test [5] to test whether to reject the hypothesis that the successful-run lengths are IID random variable with an exponential distribution (which is the continuous valued equivalent of the geometric distribution). This test has been found to have higher probability of rejecting the null hypothesis (that the distribution is exponential) than the Kolmogorov-Smimov or the χ^2 test. This is, in part, due to the fact that it employs the full empirical distribution rather than binning as in the χ^2 test, allowing it to give more weight to larger sample values.

We find that for Planet-Lab traces with small loss rates (below 1.5%) the exponential distribution fits well our data. Indeed, in 18 out of 22 traces exponential distribution is a good fit for the inter-arrival of the loss runs at the significant level 0.05. But for those with higher loss rates the exponential distribution is not a good fit. Thus for small loads, the arrivals

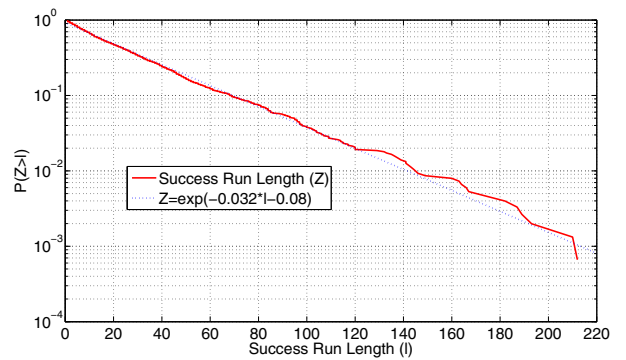
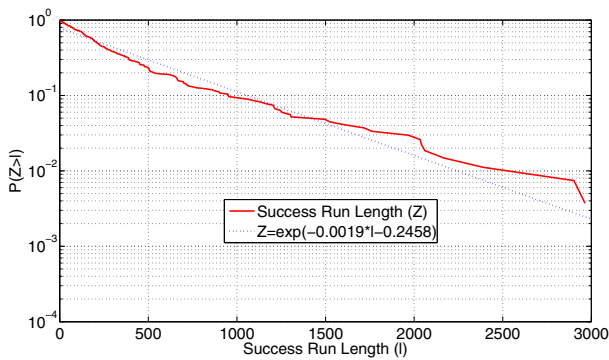


Fig. 5. Empirical distribution of the success run lengths under different loads: (left) Loss Rate of 0.025, (right) Loss Rate of 0.035. Clearly, under both loads, the success run lengths is geometrically distributed. The slopes in the two curves are different, suggesting that by looking at the slope we can estimate the parameters of the corresponding distribution.

of loss runs can be considered as Poisson, consistent with the results in [14]. However, for traces with higher loads, the above conclusion does not hold.

Thus, even though the loss process can be modelled as an alternating renewal process regardless of the network load, its Markovian properties do depend on the network load.

VI. CONCLUSION

In this paper, we study the correlation of Internet packet losses in a limited set of PlanetLab measurements. We found that for half of the traces packet losses are far from independent. More significantly, we found that the correlation time scale of packet losses is highly dependent on the network load. Our traces show little variation in the loss run lengths. The majority of them are 1. However, the success run lengths are highly variable. We found that the success runs in a majority of our traces are independent and that their correlation timescales are not related with the network load. All of these results point to the significant role the network load has on the correlation of binary packet losses. However, the loss process can be viewed as an alternating renewal process regardless of the network load.

An immediate future work is to extend measurements to higher probing rates to capture better the loss runs and to obtain more lossy traces. With these traces, we can then fully specify the correlation of packet losses at all network loads-not just the light loads in the current data sets.

REFERENCES

- [1] Jonas Andren, Magnus Hilding, and Darryl Veitch. Understanding end-to-end internet traffic dynamics. In *Proc. of IEEE Global Telecommunications Conference (Globecom'98)*, Sydney, Australia, 1998.
- [2] J. C. Bolot. End-to-end packet delay and loss behaviour in the internet. In *Proceedings of the ACM Sigcomm*, 1993.
- [3] P. Brockwell and R. Davis. *Introduction to Time Series and Forecasting*. Springer-Verlag, 1996.
- [4] C. Chatfield. Statistical inference regarding Markov chain models. *Appl. Statist.*, 22:7–20, 1973.
- [5] R. B. D'Agostino and M. A. Stephens. *Goodness-of-fit Techniques*. Marcel Dekker, 1986.
- [6] E. Gilbert. Capacity of a burst-noise channel. *Bell Systems Technical Journal*, pages 1253–1265, September 1960.
- [7] Wenyu Jiang and Henning Schulzrinne. Modeling of packet loss and delay and their effect on real-time multimedia service quality. In *Proc. of NOSSDAV 2000*, June, 2000.
- [8] B. Melamed and D. Yao. *Advances in Queueing: Theory, Methods and Open Problems*. CRC Press, 1995, ch. The ASTA Property.
- [9] V. Paxson. End-to-end internet packet dynamics. In *Proceedings of the ACM SIGCOMM*, Sep 1997.
- [10] V. Paxson, A. Adams, and M. Mathis. Experiences with NIMI. In *Proceedings of PAM 2000*, 2000.
- [11] Matthew Roughan. Fundamental bounds on the accuracy of network performance measurements. In *Proceedings of ACM SIGMETRICS*, June 2005.
- [12] www.planet lab.org.
- [13] Maya Yajnik, Sue Moon, Jim Kurose, and Don Towsley. Measurement and modelling of the temporal dependence in packet loss. In *Proc. of Infocom 99*, 1999.
- [14] Y. Zhang, N. Duffield, V. Paxson, and S. Shenker. On the constancy of Internet path properties. In *Proceedings of ACM SIGCOMM Internet Measurement Workshop*, San Francisco, 2001.