

Network Working Group
Request for Comments: 3357
Category: Informational

R. Koodli
Nokia Research Center
R. Ravikanth
Axiowave
August 2002

One-way Loss Pattern Sample Metrics

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

Using the base loss metric defined in RFC 2680, this document defines two derived metrics "loss distance" and "loss period", and the associated statistics that together capture loss patterns experienced by packet streams on the Internet. The Internet exhibits certain specific types of behavior (e.g., bursty packet loss) that can affect the performance seen by the users as well as the operators. The loss pattern or loss distribution is a key parameter that determines the performance observed by the users for certain real-time applications such as packet voice and video. For the same loss rate, two different loss distributions could potentially produce widely different perceptions of performance.

Table of Contents

| | |
|--|----|
| 1. Introduction | 3 |
| 2. Terminology | 3 |
| 3. The Approach | 3 |
| 4. Basic Definitions | 4 |
| 5. Definitions for Samples of One-way Loss Distance, and One-way Loss Period | 5 |
| 5.1. Metric Names | 5 |
| 5.1.1. Type-P-One-Way-Loss-Distance-Stream | 5 |
| 5.1.2. Type-P-One-Way-Loss-Period-Stream | 5 |
| 5.2. Metric Parameters | 5 |
| 5.3. Metric Units | 5 |
| 5.3.1. Type-P-One-Way-Loss-Distance-Stream | 5 |
| 5.3.2. Type-P-One-Way-Loss-Period-Stream | 5 |
| 5.4. Definitions | 6 |
| 5.4.1. Type-P-One-Way-Loss-Distance-Stream | 6 |
| 5.4.2. Type-P-One-Way-Loss-Period-Stream | 6 |
| 5.4.3. Examples | 6 |
| 5.5. Methodologies | 7 |
| 5.6. Discussion | 8 |
| 5.7. Sampling Considerations | 8 |
| 5.8. Errors and Uncertainties | 8 |
| 6. Statistics | 9 |
| 6.1. Type-P-One-Way-Loss-Noticeable-Rate | 9 |
| 6.2. Type-P-One-Way-Loss-Period-Total | 9 |
| 6.3. Type-P-One-Way-Loss-Period-Lengths | 10 |
| 6.4. Type-P-One-Way-Inter-Loss-Period-Lengths | 10 |
| 6.5. Examples | 10 |
| 7. Security Considerations | 11 |
| 7.1. Denial of Service Attacks | 12 |
| 7.2. Privacy / Confidentiality | 12 |
| 7.3. Integrity | 12 |
| 8. IANA Considerations | 12 |
| 9. Acknowledgements | 12 |
| 10. Normative References | 12 |
| 11. Informative References | 13 |
| Authors' Addresses | 14 |
| Full Copyright Statement | 15 |

1. Introduction

In certain real-time applications (such as packet voice and video), the loss pattern or loss distribution is a key parameter that determines the performance observed by the users. For the same loss rate, two different loss distributions could potentially produce widely different perceptions of performance. The impact of loss pattern is also extremely important for non-real-time applications that use an adaptive protocol such as TCP. Refer to [4], [5], [6], [11] for evidence as to the importance and existence of loss burstiness and its effect on packet voice and video applications.

Previously, the focus of the IPPM had been on specifying base metrics such as delay, loss and connectivity under the framework described in RFC 2330. However, specific Internet behaviors can also be captured under the umbrella of the IPPM framework, specifying new concepts while reusing existing guidelines as much as possible. In this document, we propose two derived metrics, called "loss distance" and "loss period", with associated statistics, to capture packet loss patterns. The loss period metric captures the frequency and length (burstiness) of loss once it starts, and the loss distance metric captures the spacing between the loss periods. It is important to note that these metrics are derived based on the base metric Type-P-One-Way-packet-Loss.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", "OPTIONAL", and "silently ignore" in this document are to be interpreted as described in BCP 14, RFC 2119 [2].

3. The Approach

This document closely follows the guidelines specified in [3]. Specifically, the concepts of singleton, sample, statistic, measurement principles, Type-P packets, as well as standard-formed packets all apply. However, since the document proposes to capture specific Internet behaviors, modifications to the sampling process MAY be needed. Indeed, this is mentioned in [1], where it is noted that alternate sampling procedures may be useful depending on specific circumstances. This document proposes that the specific behaviors be captured as "derived" metrics from the base metrics the behaviors are related to. The reasons for adopting this position are the following:

- it provides consistent usage of singleton metric definition for different behaviors (e.g., a single definition of packet loss is needed for capturing burst of losses, 'm out of n' losses etc.)
- it allows re-use of the methodologies specified for the singleton metric with modifications whenever necessary
- it clearly separates few base metrics from many Internet behaviors

Following the guidelines in [3], this translates to deriving sample metrics from the respective singletons. The process of deriving sample metrics from the singletons is specified in [3], [1], and others.

In the following sections, we apply this approach to a particular Internet behavior, namely the packet loss process.

4. Basic Definitions

Sequence number: Consecutive packets in a time series sample are given sequence numbers that are consecutive integers. This document does not specify exactly how to associate sequence numbers with packets. The sequence numbers could be contained within test packets themselves, or they could be derived through post-processing of the sample.

Bursty loss: The loss involving consecutive packets of a stream.

Loss Distance: The difference in sequence numbers of two successively lost packets which may or may not be separated by successfully received packets.

Example: In a packet stream, the packet with sequence number 20 is considered lost, followed by the packet with sequence number 50. The loss distance is 30.

Loss period: Let P_i be the i 'th packet. Define $f(P_i) = 1$ if P_i is lost, 0 otherwise. Then, a loss period begins if $f(P_i) = 1$ and $f(P_{(i-1)}) = 0$

Example: Consider the following sequence of lost (denoted by x) and received (denoted by r) packets.

r r r x r r x x x r x r r x x x

Then, with 'i' assigned as follows,

```

i:      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
          1 1 1 1 1 1

```

f(P_i) is,

```

f(Pi): 0 0 0 1 0 0 1 1 1 0 1 0 0 1 1 1

```

and there are four loss periods in the above sequence beginning at P₃, P₆, P₁₀, and P₁₃.

5. Definitions for Samples of One-way Loss Distance, and One-way Loss Period

5.1. Metric Names

5.1.1. Type-P-One-Way-Loss-Distance-Stream

5.1.2. Type-P-One-Way-Loss-Period-Stream

5.2. Metric Parameters

Src, the IP address of a host

Dst, the IP address of a host

T₀, a time

T_f, a time

lambda, a rate of any sampling method chosen in reciprocal of seconds

5.3. Metric Units

5.3.1. Type-P-One-Way-Loss-Distance-Stream

A sequence of pairs of the form <loss distance, loss>, where loss is derived from the sequence of <time, loss> in [1], and loss distance is either zero or a positive integer.

5.3.2. Type-P-One-Way-Loss-Period-Stream

A sequence of pairs of the form <loss period, loss>, where loss is derived from the sequence of <time, loss> in [1], and loss period is an integer.

5.4. Definitions

5.4.1. Type-P-One-Way-Loss-Distance-Stream

When a packet is considered lost (using the definition in [1]), we look at its sequence number and compare it with that of the previously lost packet. The difference is the loss distance between the lost packet and the previously lost packet. The sample would consist of <loss distance, loss> pairs. This definition assumes that sequence numbers of successive test packets increase monotonically by one. The loss distance associated with the very first packet loss is considered to be zero.

The sequence number of a test packet can be derived from the timeseries sample collected by performing the loss measurement according to the methodology in [1]. For example, if a loss sample consists of <T0,0>, <T1,0>, <T2,1>, <T3,0>, <T4,0>, the sequence numbers of the five test packets sent at T0, T1, T2, T3, and T4 can be 0, 1, 2, 3 and 4 respectively, or 100, 101, 102, 103 and 104 respectively, etc.

5.4.2. Type-P-One-Way-Loss-Period-Stream

We start a counter 'n' at an initial value of zero. This counter is incremented by one each time a lost packet satisfies the definition outlined in 4. The metric is defined as <loss period, loss> where "loss" is derived from the sequence of <time, loss> in Type-P-One-Way-Loss-Stream [1], and loss period is set to zero when "loss" is zero in Type-P-One-Way-Loss-Stream, and loss period is set to 'n' (above) when "loss" is one in Type-P-One-Way-Loss-Stream.

Essentially, when a packet is lost, the current value of "n" indicates the loss period to which this packet belongs. For a packet that is received successfully, the loss period is defined to be zero.

5.4.3. Examples

Let the following set of pairs represent a Type-P-One-Way-Loss-Stream.

```
{<T1,0>,<T2,1>,<T3,0>,<T4,0>,<T5,1>,<T6,0>,<T7,1>,<T8,0>,<T9,1>,<T10,1>}
```

where T1, T2,...,T10 are in increasing order.

Packets sent at T2, T5, T7, T9, T10 are lost. The two derived metrics can be obtained from this sample as follows.

(i) Type-P-One-Way-Loss-Distance-Stream:

Since packet 2 is the first lost packet, the associated loss distance is zero. For the next lost packet (packet 5), loss distance is 5-2 or 3. Similarly, for the remaining lost packets (packets 7, 9, and 10) their loss distances are 2, 2, and 1 respectively. Therefore, the Type-P-One-Way-Loss-Distance-Stream is:

```
{<0,0>,<0,1>,<0,0>,<0,0>,<3,1>,<0,0>,<2,1>,<0,0>,<2,1>,<1,1>}
```

(ii) The Type-P-One-Way-Loss-Period-Stream:

The packet 2 sets the counter 'n' to 1, which is incremented by one for packets 5, 7 and 9 according to the definition in 4. However, for packet 10, the counter remains at 4, again satisfying the definition in 4. Thus, the Type-P-One-Way-Loss-Period-Stream is:

```
{<0,0>,<1,1>,<0,0>,<0,0>,<2,1>,<0,0>,<3,1>,<0,0>,<4,1>,<4,1>}
```

5.5. Methodologies

The same methodology outlined in [1] can be used to conduct the sample experiments. A synopsis is listed below.

Generally, for a given Type-P, one possible methodology would proceed as follows:

- Assume that Src and Dst have clocks that are synchronized with each other. The degree of synchronization is a parameter of the methodology, and depends on the threshold used to determine loss (see below).
- At the Src host, select Src and Dst IP addresses, and form a test packet of Type-P with these addresses.
- At the Dst host, arrange to receive the packet.
- At the Src host, place a timestamp in the prepared Type-P packet, and send it towards Dst.
- If the packet arrives within a reasonable period of time, the one-way packet-loss is taken to be zero.

- If the packet fails to arrive within a reasonable period of time, the one-way packet-loss is taken to be one. Note that the threshold of "reasonable" here is a parameter of the methodology.

5.6. Discussion

The Loss-Distance-Stream metric allows one to study the separation between packet losses. This could be useful in determining a "spread factor" associated with the packet loss rate. In conjunction, the Loss-Period-Stream metric allows the study of loss burstiness for each occurrence of loss. A single loss period of length 'n' can account for a significant portion of the overall loss rate. Note that it is possible to measure distance between loss bursts separated by one or more successfully received packets. (Refer to Sections 6.4 and 6.5).

5.7. Sampling Considerations

The proposed metrics can be used independent of the particular sampling method used. We note that Poisson sampling may not yield appropriate values for these metrics for certain real-time applications such as voice over IP, as well as to TCP-based applications. For real-time applications, it may be more appropriate to use the ON-OFF [10] model, in which an ON period starts with a certain probability 'p', during which a certain number of packets are transmitted with mean 'lambda-on' according to geometric distribution and an OFF period starts with probability '1-p' and lasts for a period of time based on exponential distribution with rate 'lambda-off'.

For TCP-based applications, one may use the model proposed in [8]. See [9] for an application of the model.

5.8. Errors and Uncertainties

The measurement aspects, including the packet size, loss threshold, type of the test machine chosen etc, invariably influence the packet loss metric itself and hence the derived metrics described in this document. Thus, when making an assessment of the results pertaining to the metrics outlined in this document, attention must be paid to these matters. See [1] for a detailed consideration of errors and uncertainties regarding the measurement of base packet loss metric.

6. Statistics

6.1. Type-P-One-Way-Loss-Noticeable-Rate

Define loss of a packet to be "noticeable" [7] if the distance between the lost packet and the previously lost packet is no greater than delta, a positive integer, where delta is the "loss constraint".

Example: Let delta = 99. Let us assume that packet 50 is lost followed by a bursty loss of length 3 starting from packet 125. All the three losses starting from packet 125 are noticeable.

Given a Type-P-One-Way-Loss-Distance-Stream, this statistic can be computed simply as the number of losses that violate some constraint delta, divided by the number of losses. (Alternatively, it can also be defined as the number of "noticeable losses" to the number of successfully received packets). This statistic is useful when the actual distance between successive losses is important. For example, many multimedia codecs can sustain losses by "concealing" the effect of loss by making use of past history information. Their ability to do so degrades with poor history resulting from losses separated by close distances. By choosing delta based on this sensitivity, one can measure how "noticeable" a loss might be for quality purposes. The noticeable loss requires a certain "spread factor" for losses in the timeseries. In the above example where loss constraint is equal to 99, a loss rate of one percent with a spread of 100 between losses (e.g., 100, 200, 300, 400, 500 out of 500 packets) may be more desirable for some applications compared to the same loss rate with a spread that violates the loss constraint (e.g., 100, 175, 275, 290, 400: losses occurring at 175 and 290 violate delta = 99).

6.2. Type-P-One-Way-Loss-Period-Total

This represents the total number of loss periods, and can be derived from the loss period metric Type-P-One-Way-Loss-Period-Stream as follows:

Type-P-One-Way-Loss-Period-Total = maximum value of the first entry of the set of pairs, <loss period, loss>, representing the loss metric Type-P-One-Way-Loss-Period-Stream.

Note that this statistic does not describe the duration of each loss period itself. If this statistic is large, it does not mean that the losses are more spread out than they are otherwise; one or more loss periods may include bursty losses. This statistic is generally useful in gathering first order approximation of loss spread.

6.3. Type-P-One-Way-Loss-Period-Lengths

This statistic is a sequence of pairs <loss period, length>, with the "loss period" entry ranging from 1 - Type-P-One-Way-Loss-Period-Total. Thus the total number of pairs in this statistic equals Type-P-One-Way-Loss-Period-Total. In each pair, the "length" is obtained by counting the number of pairs, <loss period, loss>, in the metric Type-P-One-Way-Loss-Period-Stream which have their first entry equal to "loss period."

Since this statistic represents the number of packets lost in each loss period, it is an indicator of burstiness of each loss period. In conjunction with loss-period-total statistic, this statistic is generally useful in observing which loss periods are potentially more influential than others from a quality perspective.

6.4. Type-P-One-Way-Inter-Loss-Period-Lengths

This statistic measures distance between successive loss periods. It takes the form of a set of pairs <loss period, inter-loss-period-length>, with the "loss period" entry ranging from 1 - Type-P-One-Way-Loss-Period-Total, and "inter-loss-period-length" is the loss distance between the last packet considered lost in "loss period" 'i-1', and the first packet considered lost in "loss period" 'i', where 'i' ranges from 2 to Type-P-One-Way-Loss-Period-Total. The "inter-loss-period-length" associated with the first "loss period" is defined to be zero.

This statistic allows one to consider, for example, two loss periods each of length greater than one (implying loss burst), but separated by a distance of 2 to belong to the same loss burst if such a consideration is deemed useful. When the Inter-Loss-Period-Length between two bursty loss periods is smaller, it could affect the loss concealing ability of multimedia codecs since there is relatively smaller history. When it is larger, an application may be able to rebuild its history which could dampen the effect of an impending loss (period).

6.5. Examples

We continue with the same example as in Section 5.4.3. The three statistics defined above will have the following values.

- Let delta = 2. In Type-P-One-Way-Loss-Distance-Stream

{<0,0>,<0,1>,<0,0>,<0,0>,<3,1>,<0,0>,<2,1>,<0,0>,<2,1>,<1,1>},

there are 3 loss distances that violate the delta of 2. Thus, Type-P-One-Way-Loss-Noticeable-Rate = 3/5 ((number of noticeable losses)/(number of total losses))

- In Type-P-One-Way-Loss-Period-Stream

{<0,0>,<1,1>,<0,0>,<0,0>,<2,1>,<0,0>,<3,1>,<0,0>,<4,1>,<4,1>},

the largest of the first entry in the sequence of <loss period,loss> pairs is 4. Thus,

Type-P-One-Way-Loss-Period-Total = 4

- In Type-P-One-Way-Loss-Period-Stream

{<0,0>,<1,1>,<0,0>,<0,0>,<2,1>,<0,0>,<3,1>,<0,0>,<4,1>,<4,1>},

the lengths of individual loss periods are 1, 1, 1 and 2 respectively. Thus,

Type-P-One-Way-Loss-Period-Lengths =

{<1,1>,<2,1>,<3,1>,<4,2>}

- In Type-P-One-Way-Loss-Period-Stream

{<0,0>,<1,1>,<0,0>,<0,0>,<2,1>,<0,0>,<3,1>,<0,0>,<4,1>,<4,1>},

the loss periods 1 and 2 are separated by 3 (5-2), loss periods 2 and 3 are separated by 2 (7-5), and 3 and 4 are separated by 2 (9-7). Thus, Type-P-One-Way-Inter-Loss-Period-Lengths =

{<1,0>,<2,3>,<3,2>,<4,2>}

7. Security Considerations

Conducting Internet measurements raises both security and privacy concerns. This document does not specify a particular implementation of metrics, so it does not directly affect the security of the Internet nor of applications which run on the Internet. However, implementations of these metrics must be mindful of security and privacy concerns.

The derived sample metrics in this document are based on the loss metric defined in RFC 2680 [1], and thus they inherit the security considerations of that document. The reader should consult [1] for a more detailed treatment of security considerations. Nevertheless, there are a few things to highlight.

7.1. Denial of Service Attacks

The lambda specified in the Type-P-Loss-Distance-Stream and Type-P-Loss-Period-Stream controls the rate at which test packets are sent, and therefore if it is set inappropriately large, it could perturb the network under test, cause congestion, or at worst be a denial-of-service attack to the network under test. Legitimate measurements must have their parameters selected carefully in order to avoid interfering with normal traffic in the network.

7.2. Privacy / Confidentiality

Privacy of user data is not a concern, since the underlying metric is intended to be implemented using test packets that contain no user information. Even if packets contained user information, the derived metrics do not release data sent by the user.

7.3. Integrity

Results could be perturbed by attempting to corrupt or disrupt the underlying stream, for example adding extra packets that look just like test packets. To ensure that test packets are valid and have not been altered during transit, packet authentication and integrity checks, such as a signed cryptographic hash, MAY be used.

8. IANA Considerations

Since this document does not define a specific protocol, nor does it define any well-known values, there are no IANA considerations for this document.

9. Acknowledgements

Matt Zekauskas provided insightful feedback and the text for the Security Considerations section. Merike Kao helped revising the Security Considerations and the Abstract to conform with RFC guidelines. We thank both of them. Thanks to Guy Almes for encouraging the work, and Vern Paxson for the comments during the IETF meetings. Thanks to Steve Glass for making the presentation at the Oslo meeting.

10. Normative References

- [1] Almes, G., Kalindindi, S. and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

- [3] Paxson, V., Almes, G., Mahdavi, J. and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, May 1998.

11. Informative References

- [4] J.-C. Bolot and A. vega Garcia, "The case for FEC-based error control for Packet Audio in the Internet", ACM Multimedia Systems, 1997.
- [5] M. S. Borella, D. Swider, S. Uludag, and G. B. Brewster, "Internet Packet Loss: Measurement and Implications for End-to-End QoS," Proceedings, International Conference on Parallel Processing, August 1998.
- [6] M. Handley, "An examination of MBONE performance", Technical Report, USC/ISI, ISI/RR-97-450, July 1997
- [7] R. Koodli, "Scheduling Support for Multi-tier Quality of Service in Continuous Media Applications", PhD dissertation, Electrical and Computer Engineering Department, University of Massachusetts, Amherst, MA 01003, September 1997.
- [8] J. Padhye, V. Firoiu, J. Kurose and D. Towsley, "Modeling TCP throughput: a simple model and its empirical validation", in Proceedings of SIGCOMM'98, 1998.
- [9] J. Padhye, J. Kurose, D. Towsley and R. Koodli, "A TCP-friendly rate adjustment protocol for continuous media flows over best-effort networks", short paper presentation in ACM SIGMETRICS'99. Available as Umass Computer Science tech report from <ftp://gaia.cs.umass.edu/pub/Padhye98-tcp-friendly-TR.ps.gz>
- [10] K. Sriram and W. Whitt, "Characterizing superposition arrival processes in packet multiplexers for voice and data", IEEE Journal on Selected Areas of Communication, pages 833-846, September 1986,
- [11] M. Yajnik, J. Kurose and D. Towsley, "Packet loss correlation in the MBONE multicast network", Proceedings of IEEE Global Internet, London, UK, November 1996.

Authors' Addresses

Rajeev Koodli
Communications Systems Lab
Nokia Research Center
313 Fairchild Drive
Mountain View, CA 94043
USA

Phone: +1-650 625-2359
Fax: +1 650 625-2502
EMail: rajeev.koodli@nokia.com

Rayadurgam Ravikanth
Axiowave Networks Inc.
200 Nickerson Road
Marlborough, MA 01752
USA

EMail: rravikanth@axiowave.com

Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

