One-way Loss Pattern Sample Metrics
<**draft-ietf-ippm-loss-pattern-01.txt**>

Abstract

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. Currently, the focus has been on specifying
base metrics such as delay, loss and connectivity under the
framework described in [frame-work]. It is useful to capture
specific Internet behaviors under the umbrella of IPPM framework,
specifying new concepts while reusing existing guidelines as much as
possible. This draft proposes the use of "derived metrics" to
accomplish this, specifically providing means for capturing the loss
pattern on the Internet.

1. Introduction
In certain real-time applications (such as packet voice and video),

R. Koodli, R. Ravikanth   [Page 1]

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. There is ample evidence in the
literature indicating the importance and existence of loss burstiness
and its effect on packet voice and video applications
[Bolot], [Borella], [Handley], [Yajnik].

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. The Approach

This document closely follows the guidelines specified in [frame-work].
Specifically, the concepts of "singleton, sample, statistic",
measurement principles, Type-P packets, as well as standard-formed
packets all apply. However, since the draft proposes to capture
specific Internet behaviors, modifications to the sampling process
may be needed. Indeed, this is mentioned in [AKZ], where it is noted
that alternate sampling procedures may be useful depending on specific
circumstances. This draft 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. Otherwise, the metrics would have to be fundamentally
   different)
 - 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 [frame-work], this
translates to deriving *sample* metrics from the respective
singletons. The process of deriving sample metrics from the singletons
is specified in [frame-work], [AKZ], and others.

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

[3](#). Basic Definitions:

[3.1](#). Bursty loss:

The loss involving consecutive packets of a stream.

[3.2](#). Loss Distance:

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

   Example. Let packet with sequence number 50 be considered lost
           immediately after packet with sequence number 20 was
           considered lost. The loss distance is 30.

{Comment: note that this definition does not specify exactly how to
associate sequence numbers with test packets. In other words, from
a timeseries sample of test packets, one may derive the sequence
numbers. For example, assign consecutive integers to each packet in
the time series.}

[3.3](#). 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
                                1 1 1 1 1 1
      i: 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5

   $f(P_i)$ is,

   $f(P_i)$:   0 0 0 1 0 1 0 0 0 1 0 1 1 0 0 0

           and there are four loss periods in the above sequence
           begining at P_3, P_6, P_{10}, and P_{13}.

[4](#). Definitions for Samples of One-way Loss Distance,
   and One-way Loss Period.

[4.1](#) Metric Name:

[4.1.1](#)   Type-P-One-Way-Loss-Distance-Stream

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

4.2 Metric Parameters

 +   Src, the IP address of a host
 +   Dst, the IP address of a host
 +   T0, a time
 +   Tf, a time
 +   lambda, a rate in reciprocal of seconds
 +   Path, the path from Src to Dst (See [AKZ] for comments)

4.3 Metric Units

4.3.1 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 [AKZ], and loss
distance is either zero or a positive integer.

4.3.2 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 [AKZ], and loss period
an integer.

4.4. Definitions:

4.4.1 Type-P-One-Way-Loss-Distance-Stream

When a packet is considered lost (using the definition in [AKZ]), 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 [AKZ]. 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.

{Packet loss may also be considered as a result of exceeding some delay
threshold. This is particularly applicable to delay-sensitive audio
(or video) applications.
}

4.4.2 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 3.3.
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 [AKZ], 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.

{Note: 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.}

4.4.3 Example:

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 Definition 3.3. However, for
packet 10, the counter remains at 4 satisfying Definition 3.3 again.
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>}

4.5. Methodologies:

The same methodology outlined in [AKZ] can be used to conduct the
sample experiments.

4.6 **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. For
example, for a given packet loss rate, the proposed metric
indicates how the losses are spread. On the other hand,
the Loss-Period-Stream metric allows the study of loss burstiness
for each occurrence of loss. Note that a single loss period of
length 'n' can account for a significant portion of the overall
loss rate.

4.7 **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 [Sriram] model, in which an ON period
starts with certain probability 'p', during which 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
[Padhye1]. See [Padhye2] for an application of the model.

5. **Statistics:**

5.1 **Type-P-One-Way-Loss-Noticeable-Rate**

Define loss of a packet to be "noticeable" [RK97] 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 = 100. Let us assume that packet 50 is lost
           followed by a bursty loss of length 3 starting from
           packet 125.
           All the *four* losses 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. (Alternately, it
can also be defined as the number of "noticeable losses" to the number
of successfully received packets).

5.2 **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.

5.3 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 first entry equal
to "loss period."

{Note: This statistic represents the number of packets lost in each
loss period.}

5.4 Example

We continue with the same example as in Section 4.4.3. 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>}

6. Security Considerations

Since this draft proposes sample metrics based on the base loss metric

defined in [AKZ], it inherits the security considerations mentioned in
[AKZ].

7. Acknowledgements

Many thanks to Matt Zekauskas for the constructive feedback on the draft.
Thanks to Guy Almes for encouraging the work, and Vern Paxson for all the
comments during the IETF meetings.

8. References

[AKZ] G. Almes and S. Kalindindi and M. Zekauskas, "A One-way Packet
Loss Metric for IPPM", Internet Draft <draft-ietf-ippm-loss-07.txt>,
May 1999

[Bolot] J.-C. Bolot and A. vega Garcia, "The case for FEC-based
error control for Packet Audio in the Internet", ACM Multimedia
Systems, 1997.

[Borella] 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.

[Handley] M. Handley, "An examination of MBONE performance",
Technical Report, USC/ISI, ISI/RR-97-450, January 1997

[RK97] 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.
ftp://aurelius.ecs.umass.edu/pub/koodli/thesis.ps.Z

[Padhye1] 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.

[Padhye2] 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

[Paxson] V. Paxson, "End-to-end Internet packet dynamics", Computer
Communication review, Proceedings of ACM SIGCOMM'97 Conference,
Cannes, France, September 1997, 27(4), pages 139-152, October 1997

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

    [Sriram] K. Sriram and W. Whitt, "Characterizing superposition
    arrival processes in packet multiplexers for voice and data", IEEE
    Journal on Selected Areas of Communication, September 1986, pages
    833-846

    [Yajnik] M. Yajnik, J. Kurose and D. Towsley, "Packet loss
    correlation in the MBONE multicast network", Proceedings of IEEE
    Global Internet, London, UK, November 1996.

Author's Addresses

Rajeev Koodli
Nokia Research Center
3, Burlington Woods Drive, #250
Burlington, MA 01803
Phone: +1 781-359-5136
Email: rajeev.koodli@nokia.com

Rayadurgam Ravikanth
Nokia Research Center
3, Burlington Woods Drive, #250
Burlington, MA 01803
Phone: +1 781-238-4905
Email: rayadurgam.ravikanth@nokia.com