Network Working Group C.Demichelis CSELT

Internet Draft March 1998

expires 16 August 1998

Instantaneous Packet Delay Variation Metric for IPPM

<draft-ietf-ippm-ipdv-00.txt>

1. Status of this Memo

2. Abstract

This memo refers to a metric for variation in delay of packets across Internet paths. The metric is based on statistics of the difference in One-way Delay of consecutive packets. This particular definition of variation is called ''Instantaneous Packet Delay Variation (ipdv)''.

The metric is valid for measurements between two hosts both in the case that they have synchronized clocks and in the case that they are not synchronized. In the second case it allows an evaluation of the relative skew. Measurements performed on both directions (Two-ways measurements) allow a better estimation of clock differences. The precision that can be obtained is evaluated.

This memo is intended to have, as much as possible, the structure of the ippm draft on one-way delay metric.

3. Introduction

This memo defines a metric for variation in delay of packets that go from one host to another one through an IP path. Since the metric is related to a variation, different definitions are possible according to what the variation is measured against.

[Editor's Note: This memo refers to the Draft-ietf "One-way-delay metric for IPPM" that supposes as known. For sake of readability, some text is directly taken from that Draft. Text taken without modification is marked with trailing "TTTTTTT" and ending "EEEEEEE". These marks will be removed in next versions]

3.1. Definition

The Instantaneous Packet Delay Variation of an IP packet, inside a stream of packets, going from the measurement point MP1 to the measurement point MP2, is the difference of the One-Way Delay of that packet and the One-Way Delay of preceding packet in the stream.

3.2. Motivation

A number of services that can be supported by IP are sensitive to the regular delivery of packets and can be disturbed by instantaneous variations in delay, while they are not disturbed by slow variations, that can last a relatively long time. A specific metric for quick variations is therefore desirable.

In addition, this type of metric is particularly robust with respect differences and variations of the clocks of the two hosts. This allow the use of the metric even if the two hosts that support the measurement points are not synchronized. The related precision is comparable with the one that can be achieved with synchronized clocks. This will be discussed below.

3.3. General Issues Regarding Time

All what is contained in the paragraph 2.2. of the Draft ippm on one-way delay metric (2.2. General Issues Regarding Time) applies also in this case.

In addition, it is here considered that the relative skew of the two clocks can be decomposed into two parts:

* A fixed one, called in this context "skew", given, for example, by tolerances in physical dimension of crystals.

* A variable one, called in this context "drift", given, for example, by changes in temperature or other conditions of operation.

Both of this components are part of the term "skew" as defined in the referenced Draft and in the Framework document.

4. Structure of this memo

The metric will be defined as applicable to a stream of packets that flow from a source host to a destination host (one-way ipdv). The ini -tial assumption is that source and destination hosts have synchroni- zed clocks.

The definition of a singleton of one-way ipdv metric is first consi- dered, and then a definition of samples for ipdv will be given.

Then the case of application to not synchronized hosts will be dis- cussed, and the precision will be compared with the one of the pre- vious case.

A bidirectional ipdv metric will be defined, and the methodology for error corrections. This will not be a two-ways metric, but a "paired" one-way in opposite directions. Some statistics describing the IP path behavior will be proposed.

5. A singleton definition of a One-way ipdv metric

This definition makes use of the corresponding definition of type-P-One-way-delay, that is supposed to be known. This section makes use of those parts of the One-way-delay Draft that directly apply to the One-way-ipdv metric, or makes direct references to that Draft.

5.1. Metric name

Type-P-One-way-ipdv

5.2. Metric parameters

+ Scr, the IP address of a host
+ Dst, the IP address of a host
+ T1, a time
+ T2, a time

TTTTTTT

+ Path, the path from Src to Dst; in cases where there is only one path from Src to Dst, this optional parameter can be omitted.

{Comment: the presence of path is motivated by cases such as with Merit's NetNow setup, in which a Src on one NAP can reach a Dst on another NAP by either of several different backbone networks. Generally, this optional parameter is useful only when several different routes are possible from Src to Dst. Using the loose source route IP option is avoided since it would often artificially worsen the performance observed, and since it might not be supported along some paths.}

EEEEEEE

5.2. Metric unit

The value of a Type-P-One-way-ipdv is either a real number of seconds or an undefined number of seconds

5.3. Definition

Type-P-One-way-ipdv is defined for two consecutive packets from Src to Dst, as the difference between the value of the type-P-One-way-delay from Src to Dst at T2 [via path] and the value of the type-P-One-way-delay from Src to Dst at T1 [via path]. T1 is the wire-time at which Scr sent the first bit of the first packet, and T2 is the wire-time at which Src sent the first bit of the second packet.

Therefore, for a real number ddT "The type-P-one-way-ipdv from Src to Dst at T1, T2 [via path] is ddT" means that Src sent two consecutive packets whose the first at wire-time T1 (first bit), and the second wire-time T2 (first bit) and the packets were received by Dst at wire -time dT1+T1 (last bit of the first packet), and at wire-time dT2+T2 (last bit of the second packet), and that dT2-dT1=ddT.

"The type-P-one-way-ipdv from Src to Dst at T1,T2 [via path] is undefined" means that Src sent the first bit of a packet at T1 and the first bit of a second packet at T2 and that Dst did not receive one or both packets.

5.4. Discussion

Type-P-One-way-ipdv is a metric that makes use of the same measurement methods provided for delay metrics.

The following practical issues have to be considered:

+ Being a differential measurement, this metric is less sensitive to clock synchronization problems. This issue will be more carefully examined in section 6. of this memo. It is pointed out that, if the reciprocal clock conditions change in time, the accuracy of the measurement will depend on the time interval T2-T1 and the amount of possible errors will be discussed below.

TTTTTTT

+ A given methodology will have to include a way to determine whether a delay value is infinite or whether it is merely very large (and the packet is yet to arrive at Dst). As noted by Mahdavi and Paxson, simple upper bounds (such as the 255 seconds theoretical upper bound on the lifetimes of IP packets [Postel: RFC 791]) could be used, but good engineering, including an understanding of packet lifetimes, will be needed in practice. {Comment: Note that, for many applications of these metrics, the harm in treating a large delay as infinite might be zero or very small. A TCP data packet, for example, that arrives only after several multiples of the RTT may as well have been lost.}

+ As with other 'type-P' metrics, the value of the metric may depend on such properties of the packet as protocol,(UDP or TCP) port number, size, and arrangement for special treatment (as with IP precedence or with RSVP).

+ If the packet is duplicated along the path (or paths!) so that multiple non-corrupt copies arrive at the destination, then the packet is counted as received, and the first copy to arrive determines the packet's one-way delay.

+ If the packet is fragmented and if, for whatever reason, reassembly does not occur, then the packet will be deemed lost.

EEEEEEE

5.5. Methodologies

TTTTTTT

As with other Type-P-* metrics, the detailed methodology will depend on the Type-P (e.g., protocol number, UDP/TCP port number, size, precedence).

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

EEEEEEE

+ The need of synchronized clocks for Src and Dst will be discussed later. Here a methodology is supposed that is bases on synchronized clocks.

TTTTTTT

+ At the Src host, select Src and Dst IP addresses, and form a test packet of Type-P with these addresses. Any 'padding' portion of the packet needed only to make the test packet a given size should be filled with randomized bits to avoid a situation in which the measured delay is lower than it would otherwise be due to compression techniques along the path.

+ Optionally, select a specific path and arrange for Src to send the packet to that path. {Comment: This could be done, for example, by installing a temporary host-route for Dst in Src's routing table.}

+ 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 [via path].

+ If the packet arrives within a reasonable period of time, take a timestamp as soon as possible upon the receipt of the packet. By subtracting the two timestamps, an estimate of one-way delay can be computed.

EEEEEEE

Record this first delay value.

+ Repeat the procedure with the same parameters and record the second delay value. By subtracting the second value from the first the ipdv value is obtained.

+ If one or both packets fail to arrive within a reasonable period ot time, the ipdv is taken to be undefined.

5.6. Errors and Uncertainties

In the singleton metric of ipdv, factors that affect the measurement are the same that can affect the one-way delay measurement, even if, in this case, the influence is different.

TTTTTTT

The Framework document provides general guidance on this point, but we note here the following specifics related to delay metrics:

+ Errors/uncertainties due to uncertainties in the clocks of the Src and Dst hosts.

+ Errors/uncertainties due to the difference between 'wire time' and 'host time'.

Each of these are discussed in more detail below.

EEEEEEE

5.6.1. Errors/Uncertainties related to Clocks

If, as a first approximation, the error that affects the first measurement of one-way delay were the same of the one affecting the second measurement, they will cancel each other when calculating ipdv. The residual error related to clocks is the difference of the said errors that are supposed to change from the time T1, at which the first measurement is performed, to the time T2 at which the second measurement is performed. Synchronization, skew, accuracy and resolution are here considered with the following notes:

+ Errors in synchronization between source and destination clocks contribute to errors in both of the delay measurements required for calculating ipdv.

+ If the synchronization error is Tsync, and it is a linear function of time, through the skew value, at time T1 the error will be Tsync1 and at time T2 the error will be Tsync2. The ipdv measurement will be affected by the error Tsync2-Tsync1, depending from skew and T2-T1. To minimize this error it is possible to reduce the time interval T2-T1, but this could limit the generality of the metric. Methods for evaluating the synchronization error will be discussed below, since they come from a statistic of a significant sample.

+ As far as accuracy and resolution are concerned, what is noted in the above referenced Draft on one-way delay at section 3.7.1, applies also in this case, with the further consideration, about resolution, that in this case the uncertainty introduced is two times the one of a single delay measurement.

5.6.2. Errors/uncertainties related to Wire-time vs Host-time

The content of sec. 3.7.2 of the above referenced Draft applies also in this case, with the following further consideration:

The difference between Host-time and Wire-time can be in general decomposed into two components, whose one is constant and the other is variable aroud zero. Only the variable components will produce measurement errors, while the constant one will be cancelled while calculating ipdv.

6. Definitions for Samples of One-way ipdv

Starting from the definition of the singleton metric of one-way ipdv, some ways of building a sample of such singletons are here described that have to be further anlyzed in order to find the best way of considering all the related problems. In the following, the two packets needed for a singleton measurement will be called a "pair".

6.1. A "discontinuous" definition

A general definition can be the following:

Given particular binding of the parameters Src, Dst, path, and Type-P, a sample of values of parameters T1 and T2 is defined.

The means for defining the values of T1 is to select a beginning
time T0, a final time Tf, and an average rate lambda, then
define a pseudo-random Poisson arrival process of rate lambda,
whose values fall between T0 and Tf. The time interval between
successive values of T1 will then average 1/lambda. Another si-

milar, but independent, pseudo-random Poisson arrival process based on T0', Tf' and lambda', for each T1 value that has been obtained by the first process, is used for obtaining the interval T2-T1, that falls between T0' and Tf' with an average of 1/lambda'

This general definition is likely go give problems, if no limits are considered for the values T0, Tf, T0', Tf'. For example, the emission time of the first packet of a pair, could fall before the emission time of the second packet of the preceding pair. Probably this could be acceptable (provided that there are means to recognize pairs -e.g. use of sequence numbers-), but the concept itself of ipdv would be,at least, slightly changed. A way for avoiding problems can be to give some rules on the values T0, Tf, lambda, T0', Tf', lambda', without changing the meaning of the metric. For example it can be required that Tf'<T0 in order to assure that pairs of packets consist of two consecutive packets.

6.2. A "continuous" definition

A continuous stream of test packets can be supposed, where the second packet of a pair is, at the same time, the first packet of the next pair. Therefore the preceding definition becomes:

+ Given particular binding of the parameters Src, Dst, path, and Type-P, a sample of values of parameter T1 is defined.

The means for defining the values of T1 is to select a beginning time T0, a final time Tf, and an average rate lambda, then

define a pseudo-random Poisson arrival process of rate lambda,

whose values fall between T0 and Tf. The time interval between

successive values of T1 will then average 1/lambda. From the

second value on, T1 value of the pair n coincides with T2 of the

pair n-1, and the first packet of pair n coincides with the se-

cond packet of the pair n-1.

For the moment, in the following, this second definition will be considered. Further refinement is required and is for further discussion

6.3. Metric name

Type-P-One-way-ipdv-stream

6.4. Parameters

+ Src, the IP address of a host
+ Dst, the IP address of a host
+ Path, the path* from Src to Dst; in cases where there is only one path from Src to Dst, this optional parameter can be omitted
+ T0, a time
+ Tf, a time
+ lambda, a rate in reciprocal seconds

6.5. Metric Units:

A sequence of triads whose elements are:

+ T, a time
+ Ti, a time interval.
+ dT a real number or an undefined number of seconds

6.6. Definition

A pseudo-random Poisson process is defined such that ir begins at or before T0, with average arrival rate lambda, and ends at or after Tf. Those time values Ti greater than or equal to T0 and less than or equal to Tf are then selected. Starting from time T, at each time Ti of this process a value of Type-P-One-way-ipdv is obtained, and the time T becomes T = T+Ti. The value of the sample is the sequence made up of the resulting <time, time interval, ipdv> triad.

6.7. Discussion

TTTTTTT

Note first that, since a pseudo-random number sequence is employed, the sequence of times, and hence the value of the sample, is not fully specified. Pseudo-random number generators of good quality will be needed to achieve the desired qualities.

The sample is defined in terms of a Poisson process both to avoid the effects of self-synchronization and also capture a sample that is statistically as unbiased as possible. {Comment: there is, of course, no claim that real Internet traffic arrives according to a Poisson arrival process.}

EEEEEEE

6.8. Methodology

Since packets can be lost or duplicated or can arrive in a different order with respect the one of emission, in order to recognize the pairs of test packets, they should be marked with a Sequence Number or make use of any other tool suitable to the scope. For duplicated packets only the first received copy should be considered. If a packet is lost, two values of ipdv will be undefined, since each packet, in the supposed "continuous" definition, is common to two pairs.

Steps for measurement can be the following:

+ Starting from a given time T, Src generates a test packet as for a singleton metrics, adding also a Sequence Number, and sorts a Ti interval to determine the time at wich the next packet has to be sent.

+ On reception of the packet, Dst verifies the Sequence Number SN, and records SN and Tx timestamp that are contained in the packet and the Rx timestamp.

+ if the packet is not the first received and the SN is correct, ipdv is computed and Ti is recorded. Then Dst records SN, Tx and Rx timestamps as "old" values.

6.9. Errors and uncertainties

The same considerations apply that have been made about the singleton metric. An additional error can be introduced by the pseudo-random Poisson process as focushed in the above referenced Draft. Further considerations will be made in section 7.

6.10 Some statistics for One-way-ipdv

Some statistics are here considered, that can provide useful information in analyzing the behavior of the packets flowing from Src to Dst These statistics are given having in mind a practical use of them. Other statistics can be defined if needed.

6.10.1. Type-P-One-way-ipdv-inverse-percentile

Given a Type-P-One-way-ipdv-Stream and a time threshold, that can be either positive or negative, the fraction of all the dT values in the Stream less than or equal to the threshold, if the threshold is positive, or greater or equal to the threshold if the threshold is negative.

For many real-time services that require a regular delivery of the packets, this statistics can give the amount of packets beyond acceptable limits.

6.10.2 Type-P-One-way-ipdv-standard-deviation

Given a Type-P-One-way-ipdv-Stream, the distribution of ipdv values is considered and the Standard Deviation can be calculated as an indication of regularity of delivery. For practical purposes it can useful to define a total standard deviation, computed over the complete set of value, and a standard deviation computed over the subset of those values that do not exceede given positive and negative thresholds. This allows a more accurate description of the performance experienced by packets.

6.10.3 Type-P-One-way-ipdv-average

This statistic should tend to a value of ZERO for a number of ipdv values that tend to infinite. The behavior of Type-P-One-way-ipdv-average is an issue for the next section 7.

7. Discussion on clock synchronization

This section gives some considerations about the need of having synchronized clocks at Src and Dst. These considerations are given as a basis for discussion, they require further investigation. We start from the analysis of the mean value of the ipdv distribution related to a "continuous" sample.

7.1. Mean value of ipdv distribution.

If D(i) is the delay of packet "i", and ipdv(i) is the i-th value of ipdv in the distribution of a sample of "n" values, collected with the described methodology, we can write:

```
ipdv(1) = D1 - D0
..........
ipdv(i) = D(i) - D(i-1)
..........
ipdv(n) = D(n) - D(n-1)
```

The mean value of ipdv distribution will result in

E(ipdv) = (D(n) - D(0))/n

If an actual measurement is performed, that lasts a period of time long enough to contain a number "n" sufficiently large and, supposing synchronized clocks, such that the network conditions (traffic) allow to find a D(n) not too diferent from D(0), e.g. a time of 24 hours, E(ipdv) will tend to zero.

7.2. Effects of a varying traffic

If the mean values of delay D are changing inside a given period of time, for example they are increasing due to an increment of traffic, we can consider, as a first approximation, the ipdv values as decomposed into two components, one being instantaneous and another one as having a constant rate dD and corresponding to the increment "per interval" of the mean value of D. The mean value of the distribution will be shifted of the value dD corresponding to the mean value of the interval between test packets. When the conditions will come back to the initial ones, the distribution will resume a mean value around zero. At any time the distribution will correctly describe the behavior of the packet flow.

7.3. Effects of syncronization errors

We refer here to the two components that can generate this type of errors that are the relative "skew" and "drift" of the Src and Dst clocks. It is first of all noted that the variable component "drift" is physically limited and its effects can be interpreted by saying that the total skew of the two clocks can vary, ranging from a min to a max value in the time. This type of variation takes place very slowly being most connected to variations in temperature.

We suppose to perform a measurement between a Src and a Dst that have a reciprocal, initial skew of "ts1" and a reciprocal drift such that, after the time T the total skew is "ts2". It is not here a limitation to consider that at the beginning of time T the two clocks indicate the same time T0. In order to analyze the effects produced by this situation we suppose that packets are transferred, from Src to Dst, with a consatant delay D. In this conditions the measured ipdv should always be zero, and what is actually measured is the error.

An ipdv value is measured at the beginning of time T with two packets having an interval of Ti(1).Another ipdv value is measured at the end of T with two packets having a time interval Ti(2).

On our purposes other errors (like wire-time vs host-time) are not considered since they are not relevant in this analysis.

It is then possible to calculate the values of the Tx and Rx time-stamps as they are seen by the two clocks, and the related values of the two ipdv values.

The first ipdv value will be: ipdv1 = ts1*Ti(1) + ((ts2-ts1)/T)*Ti(1)

The second ipdv value will be: ipdv2 = ts2*Ti(2) +((ts2-ts1)/T)*Ti(2)

The error is given by the amount of variation during the time interval Ti(i) between the two packets of the pair, and a second orded term due to the variation of that variation in the same interval.

7.4. Related precision

This means that:

1) + If the skew is constant and is = ts all the ipdv(i) values are increased by the quantity Ti(i)*ts with respect the actual value.

2) + Considering the total skew as subdivided into a fixed part and a variable part (skew and drift),respectively, ts and + or - td, and a minimum time T in which the drift can go from -td to +td or vice -versa, each ipdv(i) value will be increased of the fixed quantity Ti(i)*ts plus or minus, as a maximum, the quantity 2*td*Ti(i)/T

3) + If the duration of the measurement is such that it is possible to consider that the effect of the items at points 7.1 and 7.2, and the effect of the drift are negligeables (related average tending to zero), the mean value of the ipdv distribution will have the value of the skew multiplied by the mean value of the emission interval.

4) + We observe that the displacemet due to the skew does not change the shape of the distribution, and, for example the Standard Deviation remains the same. What introduces a distortion is the effect of the drift, even if the mean value of this effect is zero at the end of the measurement. This is, anyway, a "second order" effect. If, for example, a drift of 30 parts per million (ppm) takes place along a time of 4 hours, and the used Ti(i) interval ranges from 200 ms to 1200 ms, with an average of 700 ms, the maximum error on ipdv(i) values will be in the order of:

Ti(i)*td/(4*3600)= 2.25 E -9 seconds

8. Definition for a bidirectional ipdv metric

We now consider that the action of the skew on one direction is the same, with opposite sign, of the action on the other direction. The idea of performing at the same time two independent measurements in the two directions is suggested by this fact.

If, after a long measurement, the variable conditions of the system under test have reached the situation of a contribution close to zero to the mean value of the ipdv distribution, it is expected that only the fixed action of the skew has modified the measured mean value. It is therefore expected that on one direction that value is equal and opposite to the one measured in the other direction.

A bidirectional measurement can be defined not only as twin one-way independent metrics that take place (nearly) at the same time, but also as a two-ways metric making use of packets looped back at one end. This metric, that can be object of further study/Draft, would be able to measure also the Round Trip Delay and its variations.

9. References

V.Paxon, G.Almes, J.Mahdavi, M.Mathis - "Framework for IP Performance Metrics", Internet Draft <draft-ietf-ippm-framework-01.txt> Feb. 1998

G.Almes, S.Kalidindi - "A One-way Delay Metric for IPPM", Internet Draft <draft-ietf-ippm-delay-01.txt> Nov. 1997

10. Author's Address

Carlo Demichelis <carlo.demichelis@cselt.it>

CSELT - Centro Studi E Laburatori Telecomunicazioni S.p.A

Via G. Reiss Romoli 274

10148 - TORINO

Italy

Phone +39 11 228 5057

Fax. +39 11 228 5069