Network Working Group
Request for Comments: 3393
Category: Standards Track

C. Demichelis
Telecomitalia Lab
P. Chimento
Ericsson IPI

November 2002

IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

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

Abstract

This document refers to a metric for variation in delay of packets across Internet paths. The metric is based on the difference in the One-Way-Delay of selected packets. This difference in delay is called "IP 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. We discuss both in this document.

Table of Contents

1. Introduction

This memo defines a metric for the variation in delay of packets that flow from one host to another through an IP path. It is based on "A One-Way-Delay metric for IPPM", RFC 2679 [2] and part of the text in this memo is taken directly from that document; the reader is assumed to be familiar with that document.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY" and "OPTIONAL" in this document are to be interpreted as described in BCP 14, RFC 2119 [3]. Although BCP 14, RFC 2119 was written with protocols in mind, the key words are used in this document for similar reasons. They are used to ensure the results of measurements from two different implementations are comparable and to note instances where an implementation could perturb the network.

The structure of the memo is as follows:

+ A 'singleton' analytic metric, called Type-P-One-way-ipdv, will be introduced to define a single instance of an ipdv measurement.

+ Using this singleton metric, a 'sample', called Type-P-one-way-ipdv-Poisson-stream, will be introduced to make it possible to compute the statistics of sequences of ipdv measurements.

+ Using this sample, several 'statistics' of the sample will be defined and discussed

1.1. Terminology

The variation in packet delay is sometimes called "jitter". This term, however, causes confusion because it is used in different ways by different groups of people.

"Jitter" commonly has two meanings: The first meaning is the variation of a signal with respect to some clock signal, where the arrival time of the signal is expected to coincide with the arrival of the clock signal. This meaning is used with reference to synchronous signals and might be used to measure the quality of circuit emulation, for example. There is also a metric called "wander" used in this context.

The second meaning has to do with the variation of a metric (e.g., delay) with respect to some reference metric (e.g., average delay or minimum delay). This meaning is frequently used by computer scientists and frequently (but not always) refers to variation in delay.

In this document we will avoid the term "jitter" whenever possible and stick to delay variation which is more precise.

1.2. Definition

A definition of the IP Packet Delay Variation (ipdv) can be given for packets inside a stream of packets.

The ipdv of a pair of packets within a stream of packets is defined for a selected pair of packets in the stream going from measurement point MP1 to measurement point MP2.

The ipdv is the difference between the one-way-delay of the selected packets.

1.3. Motivation

One important use of delay variation is the sizing of play-out buffers for applications requiring the regular delivery of packets (for example, voice or video play-out). What is normally important in this case is the maximum delay variation, which is used to size play-out buffers for such applications [7]. Other uses of a delay variation metric are, for example, to determine the dynamics of queues within a network (or router) where the changes in delay variation can be linked to changes in the queue length process at a given link or a combination of links.

In addition, this type of metric is particularly robust with respect to differences and variations of the clocks of the two hosts. This allows the use of the metric even if the two hosts that support the measurement points are not synchronized. In the latter case indications of reciprocal skew of the clocks can be derived from the measurement and corrections are possible. The related precision is often comparable with the one that can be achieved with synchronized clocks, being of the same order of magnitude of synchronization errors. This will be discussed below.

The scope of this document is to provide a way to measure the ipdv delivered on a path. Our goal is to provide a metric which can be parameterized so that it can be used for various purposes. Any report of the metric MUST include all the parameters associated with it so that the conditions and meaning of the metric can be determined exactly. Since the metric does not represent a value judgment (i.e., define "good" and "bad"), we specifically do not specify particular values of the metrics that IP networks must meet.

The flexibility of the metric can be viewed as a disadvantage but there are some arguments for making it flexible. First, though there are some uses of ipdv mentioned above, to some degree the uses of ipdv are still a research topic and some room should be left for experimentation. Secondly, there are different views in the community of what precisely the definition should be (e.g., [8],[9],[10]). The idea here is to parameterize the definition, rather than write a different document for each proposed definition. As long as all the parameters are reported, it will be clear what is meant by a particular use of ipdv. All the remarks in the document hold, no matter which parameters are chosen.

1.4. General Issues Regarding Time

Everything contained in Section 2.2. of [2] applies also in this case.

To summarize: As in [1] we define "skew" as the first derivative of the offset of a clock with respect to "true time" and define "drift" as the second derivative of the offset of a clock with respect to "true time".

From there, we can construct "relative skew" and "relative drift" for two clocks C1 and C2 with respect to one another. These are natural extensions of the basic framework definitions of these quantities:

+ Relative offset = difference in clock times

+ Relative skew = first derivative of the difference in clock times

+ Relative drift = second derivative of the difference in clock times

NOTE: The drift of a clock, as it is above defined over a long period must have an average value that tends to zero while the period becomes large since the frequency of the clock has a finite (and small) range. In order to underline the order of magnitude of this effect,it is considered that the maximum range of drift for commercial crystals is about 50 part per million (ppm). Since it is mainly connected with variations in operating temperature (from 0 to 70 degrees Celsius), it is expected that a host will have a nearly constant temperature during its operation period, and variations in temperature, even if quick, could be less than one Celsius per second, and range in the order of a few degrees. The total range of the drift is usually related to variations from 0 to 70 Celsius. These are important points for evaluation of precision of ipdv measurements, as will be seen below.

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

The purpose of the singleton metric is to define what a single instance of an ipdv measurement is. Note that it can only be statistically significant in combination with other instances. It is not intended to be meaningful as a singleton, in the sense of being able to draw inferences from it.

This definition makes use of the corresponding definition of type-P-One-Way-Delay metric [2]. 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.

2.1. Metric name

Type-P-One-way-ipdv

2.2. Metric parameters

+ Src, the IP address of a host

+ Dst, the IP address of a host

+ T1, a time

+ T2, a time

+ L, a packet length in bits. The packets of a Type P packet stream from which the singleton ipdv metric is taken MUST all be of the same length.

+ F, a selection function defining unambiguously the two packets from the stream selected for the metric.

+ I1,I2, times which mark that beginning and ending of the interval in which the packet stream from which the singleton measurement is taken occurs.

+ P, the specification of the packet type, over and above the source and destination addresses

2.3. Metric unit

The value of a Type-P-One-way-ipdv is either a real number of seconds (positive, zero or negative) or an undefined number of seconds.

2.4. Definition

We are given a Type P packet stream and I1 and I2 such that the first Type P packet to pass measurement point MP1 after I1 is given index 0 and the last Type P packet to pass measurement point MP1 before I2 is given the highest index number.

Type-P-One-way-ipdv is defined for two packets from Src to Dst selected by the selection function F, as the difference between the value of the type-P-One-way-delay from Src to Dst at T2 and the value

of the type-P-One-Way-Delay from Src to Dst at T1. 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. This metric is derived from the One-Way-Delay metric.

Therefore, for a real number ddT "The type-P-one-way-ipdv from Src to Dst at T1, T2 is ddT" means that Src sent two packets, the first at wire-time T1 (first bit), and the second at 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 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.

Figure 1 illustrates this definition. Suppose that packets P(i) and P(k) are selected.

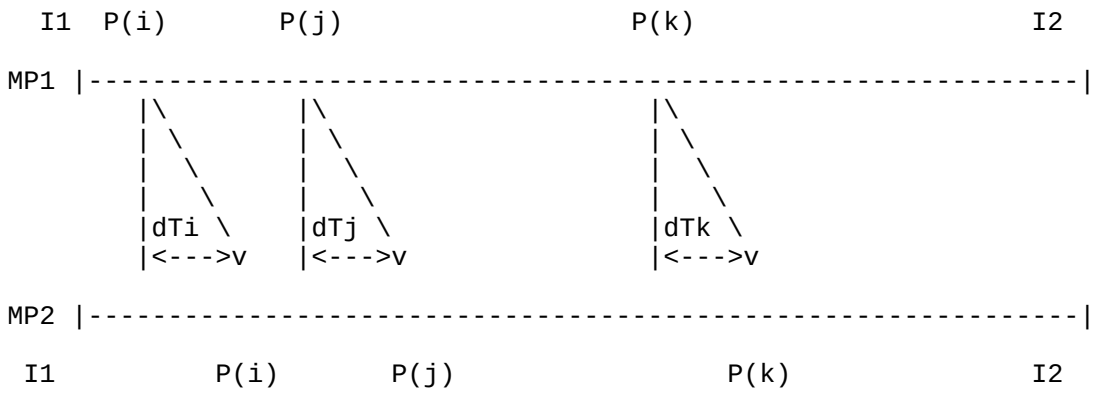

Figure 1: Illustration of the definition

Then ddT = dTk - dTi as defined above.

2.5. Discussion

This metric definition depends on a stream of Type-P-One-Way-Delay packets that have been measured. In general this can be a stream of two or more packets, delimited by the interval endpoints I1 and I2. There must be a stream of at least two packets in order for a singleton ipdv measurement to take place. The purpose of the selection function is to specify exactly which two packets from the stream are to be used for the singleton measurement. Note that the

selection function may involve observing the one-way-delay of all the Type P packets of the stream in the specified interval. Examples of a selection function are:

+ Consecutive Type-P packets within the specified interval

+ Type-P packets with specified indices within the specified interval

+ Type-P packets with the min and max one-way-delays within the specified interval

+ Type-P packets with specified indices from the set of all defined (i.e., non-infinite) one-way-delays Type-P packets within the specified interval.

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 5 of this memo. It is pointed out that, if the relative clock conditions change in time, the accuracy of the measurement will depend on the time interval I2-I1 and the magnitude of possible errors will be discussed below.

+ 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).

+ ddT is derived from the start of the first bit out from a packet sent out by Src to the reception of the last bit received by Dst. Delay is correlated to the size of the packet. For this reason, the packet size is a parameter of the measurement and must be reported along with the measurement.

+ 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, re-assembly does not occur, then the packet will be deemed lost.

In this document it is assumed that the Type-P packet stream is generated according to the Poisson sampling methodology described in [1].

The reason for Poisson sampling is that it ensures an unbiased and uniformly distributed sampling of times between I1 and I2. However, alternate sampling methodologies are possible. For example, continuous sampling of a constant bit rate stream (i.e., periodic packet transmission) is a possibility. However, in this case, one must be sure to avoid any "aliasing" effects that may occur with periodic samples.

2.6. Methodologies

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).

The measurement methodology described in this section assumes the measurement and determination of ipdv in real-time as part of an active measurement. Note that this can equally well be done a posteriori, i.e., after the one-way-delay measurement is completed.

Generally, for a given Type-P, the methodology would proceed as follows: Note that this methodology is based on synchronized clocks. The need for synchronized clocks for Src and Dst will be discussed later.

+ Start after time I1. At the Src host, select Src and Dst IP addresses, and form test packets of Type-P with these addresses according to a given technique (e.g., the Poisson sampling technique). 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.

+ At the Dst host, arrange to receive the packets.

+ At the Src host, place a time stamp in the Type-P packet, and send it towards Dst.

+ If the packet arrives within a reasonable period of time, take a time stamp as soon as possible upon the receipt of the packet. By subtracting the two time stamps, an estimate of One-Way-Delay can be computed.

+ If the packet meets the selection function criterion for the first packet, record this first delay value. Otherwise, continue generating the Type-P packet stream as above until the criterion is met or I2, whichever comes first.

+ At the Src host, packets continue to be generated according to the given methodology. The Src host places a time stamp in the Type-P packet, and send it towards Dst.

+ If the packet arrives within a reasonable period of time, take a time stamp as soon as possible upon the receipt of the packet. By subtracting the two time stamps, an estimate of One-Way-Delay can be computed.

+ If the packet meets the criterion for the second packet, then by subtracting the first value of One-Way-Delay from the second value the ipdv value of the pair of packets is obtained. Otherwise, packets continue to be generated until the criterion for the second packet is fulfilled or I2, whichever comes first.

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

2.7. Errors and Uncertainties

In the singleton metric of ipdv, factors that affect the measurement are the same as those affecting the One-Way-Delay measurement, even if, in this case, the influence is different.

The Framework document [1] 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 errors is discussed in more detail in the following paragraphs.

2.7.1. Errors/Uncertainties related to Clocks

If, as a first approximation, the error that affects the first measurement of One-Way-Delay were the same as 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 errors that are supposed to change from time T1, at which the first measurement is performed, to 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.

+ The effect of drift and skew errors on ipdv measurements can be quantified as follows: Suppose that the skew and drift functions are known. Assume first that the skew function is linear in time. Clock offset is then also a function of time and the error evolves as e(t) = K*t + O, where K is a constant and O is the offset at time 0. In this case, the error added to the subtraction of two different time stamps (t2 > t1) is e(t2)-e(t1) = K*(t2 - t1) which will be added to the time difference (t2 - t1). If the drift cannot be ignored, but we assume that the drift is a linear function of time, then the skew is given by s(t) = M*(t**2) + N*t + S0, where M and N are constants and S0 is the skew at time 0. The error added by the variable skew/drift process in this case becomes e(t) = O + s(t) and the error added to the difference in time stamps is e(t2)-e(t1) = N*(t2-t1) + M*{(t2-t1)**2}.

  It is the claim here (see remarks in section 1.3) that the effects of skew are rather small over the time scales that we are discussing here, since temperature variations in a system tend to be slow relative to packet inter-transmission times and the range of drift is so small.

+ As far as accuracy and resolution are concerned, what is noted in the one-way-delay document [2] in 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. Errors introduced by these effects are often larger than the ones introduced by the drift.

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

The content of sec. 3.7.2 of [2] 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, of which one is constant and the other is variable. Only the variable components will produce measurement errors, while the constant one will be canceled while calculating ipdv.

However, in most cases, the fixed and variable components are not known exactly.

3. Definitions for Samples of One-way-ipdv

The goal of the sample definition is to make it possible to compute the statistics of sequences of ipdv measurements. The singleton definition is applied to a stream of test packets generated according to a pseudo-random Poisson process with average arrival rate lambda. If necessary, the interval in which the stream is generated can be divided into sub-intervals on which the singleton definition of ipdv can be applied. The result of this is a sequence of ipdv measurements that can be analyzed by various statistical procedures.

Starting from the definition of the singleton metric of one-way-ipdv, we define a sample of such singletons. In the following, the two packets needed for a singleton measurement will be called a "pair".

3.1. Metric name

Type-P-One-way-ipdv-Poisson-stream

3.2. 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 seconds

+ L, a packet length in bits. The packets of a Type P packet stream from which the sample ipdv metric is taken MUST all be of the same length.

+ F, a selection function defining unambiguously the packets from the stream selected for the metric.

+ I(i),I(i+1), i >=0, pairs of times which mark the beginning and ending of the intervals in which the packet stream from which the measurement is taken occurs. I(0) >= T0 and assuming that n is the largest index, I(n) <= Tf.

+ P, the specification of the packet type, over and above the source and destination addresses

3.3. Metric Units:

A sequence of triples whose elements are:

+ T1, T2,times

+ dT a real number or an undefined number of seconds

3.4. Definition

A pseudo-random Poisson process is defined such that it begins at or before T0, with average arrival rate lambda, and ends at or after Tf. Those time values T(i) greater than or equal to T0 and less than or equal to Tf are then selected for packet generation times.

Each packet falling within one of the sub-intervals I(i), I(i+1) is tested to determine whether it meets the criteria of the selection function F as the first or second of a packet pair needed to compute ipdv. The sub-intervals can be defined such that a sufficient number of singleton samples for valid statistical estimates can be obtained.

The triples defined above consist of the transmission times of the first and second packets of each singleton included in the sample, and the ipdv in seconds.

3.5. Discussion

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. There is, of course, no claim that real Internet traffic arrives according to a Poisson arrival process.

The sample metric can best be explained with a couple of examples: For the first example, assume that the selection function specifies the "non-infinite" max and min one-way-delays over each sub-interval. We can define contiguous sub-intervals of fixed specified length and produce a sequence each of whose elements is the triple <transmission time of the max delay packet, transmission time of the min delay packet, D(max)-D(min)> which is collected for each sub-interval. A second example is the selection function that specifies packets whose indices (sequence numbers) are just the integers below a certain bound. In this case, the sub-intervals are defined by the transmission times of the generated packets and the sequence produced is just <T(i), T(i+1), D(i+1)-D(i)> where D(i) denotes the one-way-delay of the ith packet of a stream.

This definition of the sample metric encompasses both the definition proposed in [9] and the one proposed in [10].

3.6. Methodology

Since packets can be lost or duplicated or can arrive in a different order than the order sent, the pairs of test packets should be marked with a sequence number. For duplicated packets only the first received copy should be considered.

Otherwise, the methodology is the same as for the singleton measurement, with the exception that the singleton measurement is repeated a number of times.

3.7. Errors and uncertainties

The same considerations apply that have been made about the singleton metric. Additional error can be introduced by the pseudo-random Poisson process as discussed in [2]. Further considerations will be given in section 5.

4. Statistics for One-way-ipdv

Some statistics are suggested which can provide useful information in analyzing the behavior of the packets flowing from Src to Dst. The statistics are assumed to be computed from an ipdv sample of reasonable size.

The purpose is not to define every possible statistic for ipdv, but ones which have been proposed or used.

4.1. Lost Packets and ipdv statistics

The treatment of lost packets as having "infinite" or "undefined" delay complicates the derivation of statistics for ipdv. Specifically, when packets in the measurement sequence are lost, simple statistics such as sample mean cannot be computed. One possible approach to handling this problem is to reduce the event space by conditioning. That is, we consider conditional statistics; namely we estimate the mean ipdv (or other derivative statistic) conditioned on the event that selected packet pairs arrive at the destination (within the given timeout). While this itself is not without problems (what happens, for example, when every other packet is lost), it offers a way to make some (valid) statements about ipdv, at the same time avoiding events with undefined outcomes.

In practical terms, what this means is throwing out the samples where one or both of the selected packets has an undefined delay. The sample space is reduced (conditioned) and we can compute the usual statistics, understanding that formally they are conditional.

4.2. Distribution of One-way-ipdv values

The one-way-ipdv values are limited by virtue of the fact that there are upper and lower bounds on the one-way-delay values. Specifically, one-way-delay is upper bounded by the value chosen as the maximum beyond which a packet is counted as lost. It is lower bounded by propagation, transmission and nodal transit delays assuming that there are no queues or variable nodal delays in the path. Denote the upper bound of one-way-delay by U and the lower bound by L and we see that one-way-ipdv can only take on values in the (open) interval (L-U, U-L).

In any finite interval, the one-way-delay can vary monotonically (non-increasing or non-decreasing) or of course it can vary in both directions in the interval, within the limits of the half-open interval [L,U). Accordingly, within that interval, the one-way-ipdv values can be positive, negative, or a mixture (including 0).

Since the range of values is limited, the one-way-ipdv cannot increase or decrease indefinitely. Suppose, for example, that the ipdv has a positive 'run' (i.e., a long sequence of positive values). At some point in this 'run', the positive values must approach 0 (or become negative) if the one-way-delay remains finite. Otherwise, the one-way-delay bounds would be violated. If such a run were to continue infinitely long, the sample mean (assuming no packets are lost) would approach 0 (because the one-way-ipdv values must approach 0). Note, however, that this says nothing about the shape of the

distribution, or whether it is symmetric. Note further that over significant intervals, depending on the width of the interval [L,U), that the sample mean one-way-ipdv could be positive, negative or 0.

There are basically two ways to represent the distribution of values of an ipdv sample: an empirical pdf and an empirical cdf. The empirical pdf is most often represented as a histogram where the range of values of an ipdv sample is divided into bins of a given length and each bin contains the proportion of values falling between the two limits of the bin. (Sometimes instead the number of values falling between the two limits is used). The empirical cdf is simply the proportion of ipdv sample values less than a given value, for a sequence of values selected from the range of ipdv values.

4.3. Type-P-One-way-ipdv-percentile

Given a Type-P One-Way-ipdv sample and a given percent X between 0% and 100%. The Xth percentile of all ipdv values is in the sample. Therefore, then 50th percentile is the median.

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

Given a Type-P-One-way-ipdv sample and a given value Y, the percent of ipdv sample values less than or equal to Y.

4.5. Type-P-One-way-ipdv-jitter

Although the use of the term "jitter" is deprecated, we use it here following the authors in [8]. In that document, the selection function specifies that consecutive packets of the Type-P stream are to be selected for the packet pairs used in ipdv computation. They then take the absolute value of the ipdv values in the sample. The authors in [8] use the resulting sample to compare the behavior of two different scheduling algorithms.

An alternate, but related, way of computing an estimate of jitter is given in RFC 1889 [11]. The selection function there is implicitly consecutive packet pairs, and the "jitter estimate" is computed by taking the absolute values of the ipdv sequence (as defined in this document) and applying an exponential filter with parameter 1/16 to generate the estimate (i.e., j_new = 15/16* j_old + 1/16*j_new).

4.6. Type-P-One-way-peak-to-peak-ipdv

In this case, the selection function used in collecting the Type-P-One-Way-ipdv sample specifies that the first packet of each pair to be the packet with the maximum Type-P-One-Way-Delay in each subinterval and the second packet of each pair to be the packet with

the minimum Type-P-One-Way-Delay in each sub-interval. The resulting sequence of values is the peak-to-peak delay variation in each subinterval of the measurement interval.

5. Discussion of clock synchronization

This section gives some considerations about the need for having synchronized clocks at the source and destination, although in the case of unsynchronized clocks, data from the measurements themselves can be used to correct error. These considerations are given as a basis for discussion and they require further investigation.

5.1. Effects of synchronization errors

Clock errors can be generated by two processes: the relative drift and the relative skew of two given clocks. We should note that drift is physically limited and so the total relative skew of two clocks can vary between an upper and a lower bound.

Suppose then that we have a measurement between two systems such that the clocks in the source and destination systems have at time 0 a relative skew of s(0) and after a measurement interval T have skew s(T). We assume that the two clocks have an initial offset of O (that is letter O).

Now suppose that the packets travel from source to destination in constant time, in which case the ipdv is zero and the difference in the time stamps of the two clocks is actually just the relative offset of the clocks. Suppose further that at the beginning of the measurement interval the ipdv value is calculated from a packet pair and at the end of the measurement interval another ipdv value is calculated from another packet pair. Assume that the time interval covered by the first measurement is t1 and that the time interval covered by the second measurement is t2. Then

ipdv1 = s(0)*t1 + t1*(s(T)-s(0))/T

ipdv2 = s(T)*t2 + t2*(s(T)-s(0))/T

assuming that the change in skew is linear in time. In most practical cases, it is claimed that the drift will be close to zero in which case the second (correction) term in the above equations disappears.

Note that in the above discussion, other errors, including the differences between host time and wire time, and externally-caused clock discontinuities (e.g., clock corrections) were ignored. Under these assumptions the maximum clock errors will be due to the maximum relative skew acting on the largest interval between packets.

5.2. Estimating the skew of unsynchronized clocks

If the skew is linear (that is, if s(t) = S * t for constant S), the error in ipdv values will depend on the time between the packets used in calculating the value. If ti is the time between the packet pair, then let Ti denote the sample mean time between packets and the average skew is s(Ti) = S * Ti. In the event that the delays are constant, the skew parameter S can be estimated from the estimate Ti of the time between packets and the sample mean ipdv value. Under these assumptions, the ipdv values can be corrected by subtracting the estimated S * ti.

We observe that the displacement 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, also when the mean value of this effect is zero at the end of the measurement. The value of this distortion is limited to the effect of the total skew variation on the emission interval.

6. Security Considerations

The one-way-ipdv metric has the same security properties as the one-way-delay metric [2], and thus they inherit the security considerations of that document. The reader should consult [2] for a more detailed treatment of security considerations. Nevertheless, there are a few things to highlight.

6.1. Denial of service

It is still possible that there could be an attempt at a denial of service attack by sending many measurement packets into the network. In general, legitimate measurements must have their parameters carefully selected in order to avoid interfering with normal traffic.

6.2. Privacy/Confidentiality

The packets contain no user information, and so privacy of user data is not a concern.

6.3. Integrity

There could also be attempts to disrupt measurements by diverting packets or corrupting them. To ensure that test packets are valid and have not been altered during transit, packet authentication and integrity checks may be used.

7. Acknowledgments

Thanks to Merike Kaeo, Al Morton and Henk Uiterwaal for catching mistakes and for clarifying re-wordings for this final document.

A previous major revision of the document resulted from e-mail discussions with and suggestions from Mike Pierce, Ruediger Geib, Glenn Grotefeld, and Al Morton. For previous revisions of this document, discussions with Ruediger Geib, Matt Zekauskas and Andy Scherer were very helpful.

8. References

8.1 Normative References

[1] Paxon, V., Almes, G., Mahdavi, J. and M. Mathis, "Framework for IP Performance Metrics", RFC 2330, February 1998.

[2] Almes, G. and S. Kalidindisu, "A One-Way-Delay Metric for IPPM", RFC 2679, September 1999.

[3] Bradner, S., "Key words for use in RFCs to indicate requirement levels", BCP 14, RFC 2119, March 1997.

8.2 Informational References

[4] ITU-T Recommendation Y.1540 (formerly numbered I.380) "Internet Protocol Data Communication Service - IP Packet Transfer and Availability Performance Parameters", February 1999.

[5] Demichelis, Carlo - "Packet Delay Variation Comparison between ITU-T and IETF Draft Definitions" November 2000 (in the IPPM mail archives).

[6] ITU-T Recommendation I.356 "B-ISDN ATM Layer Cell Transfer Performance".

[7] S. Keshav - "An Engineering Approach to Computer Networking", Addison-Wesley 1997, ISBN 0-201-63442-2.

[8] Jacobson, V., Nichols, K. and Poduri, K. "An Expedited Forwarding PHB", RFC 2598, June 1999.

[9] ITU-T Draft Recommendation Y.1541 - "Internet Protocol Communication Service - IP Performance and Availability Objectives and Allocations", April 2000.

[10] Demichelis, Carlo - "Improvement of the Instantaneous Packet Delay Variation (IPDV) Concept and Applications", World Telecommunications Congress 2000, 7-12 May 2000.

[11] Schulzrinne, H., Casner, S., Frederick, R. and V. Jacobson, "RTP: A transport protocol for real-time applications", RFC 1889, January 1996.

9. Authors' Addresses

Carlo Demichelis
Telecomitalia Lab S.p.A
Via G. Reiss Romoli 274
10148 - TORINO
Italy

Phone: +39 11 228 5057
Fax: +39 11 228 5069
EMail: carlo.demichelis@tilab.com

Philip Chimento
Ericsson IPI
7301 Calhoun Place
Rockville, Maryland 20855
USA

Phone: +1-240-314-3597
EMail: chimento@torrentnet.com

10. 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.