

Network Working Group
INTERNET-DRAFT
Expiration Date: December 1999

C. Demichelis
CSELT
P. Chimento
CTIT
June 1999

Instantaneous Packet Delay Variation Metric for IPPM
<[draft-ietf-ippm-ipdv-03.txt](#)>

1. Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft shadow directories can be accessed at <http://www.ietf.org/shadow.html>

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

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

INTERNET-DRAFT

Packet Delay Variation

June 1999

reciprocal skew. Measurements performed on both directions (Two-way measurements) allow a better estimation of clock differences. The precision that can be obtained is evaluated.

[3.](#) Introduction

This memo takes as a reference the Draft-ietf "One-Way-Delay metric for IPPM" [1]. Part of the text in this memo is directly taken from that Draft.

This memo defines a metric for variation in delay of packets that flow 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.

"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. 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). The form of "jitter" that we talk about here has to do almost exclusively with the second meaning, rather than the first. See the section on the relationship with other standards.

[3.1.](#) Definition

A definition of the Instantaneous Packet Delay Variation (ipdv) can be given for a pair of packets or for a packet inside a stream of packets.

For a pair of packets:

- + The ipdv of a pair of IP packets, that are transmitted from the measurement point MP1 to the measurement point MP2, is the difference between the One-Way-Delay measured for the second packet and the One-Way-Delay measured for the first packet of the pair.

For a stream of packets:

- + 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 the 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. Metrics that can be derived from the analysis of statistics of ipdv can also be used, for example, for buffer dimensioning. The scope of this metric is to provide a way for measurement of the quality delivered by a path.

In addition, this type of metric is particularly robust with respect 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.

[3.3.](#) General Issues Regarding Time

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

In addition, we assume here that the reciprocal 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 dimensions of crystals.

- + A variable one, called in this context "drift", given, for example, by changes in temperature or other conditions of operation. Both of these components are part of the term "skew" as defined in the referenced Draft and in the Framework document.

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 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.

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 initial assumption is that source and destination hosts have synchronized clocks. The definition of a singleton of one-way ipdv metric is first considered, and then a definition of samples for ipdv will be given.

Then the case of application to non-synchronized hosts will be discussed, and the precision will be compared with the one of synchronized clocks.

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

[5.1.](#) Metric name

Type-P-One-way-ipdv

[5.2.](#) Metric parameters

- + Src, the IP address of a host
- + Dst, the IP address of a host
- + T1, a time
- + T2, a time. It is explicitly noted that also the difference $T2-T1$ is a parameter of the measurement though this is already implicit, since the times T1 and T2 exactly define the time conditions in which the measurement takes place.

Note that the packet length is an implicit parameter of both the Type-P-One-way-delay metric and the Type-P-One-way-ipdv metric, since this contributes to the overall one-way delay. We assume that the packets sent for ipdv measurements are all of the same length.

5.3. 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.

5.4. 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 Src 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 therefore ideally derived from the One-Way-Delay metric.

NOTE: The requirement of "consecutive" packets is not essential. The measured value is anyway the difference in One-Way-Delay at the times T1 and T2, which is meaningful by itself, as long as the times T1 and T2 denote the wire times of the packets sent from Src to Dst.

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, 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 [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.

5.5. 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 7](#) 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 T_2-T_1 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).
- + 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.

[5.6](#). 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). Generally, for a given Type-P, the methodology would proceed as follows:

- + The need of synchronized clocks for Src and Dst will be discussed

later. Here a methodology is supposed that is based on synchronized clocks.

- + At the Src host, select Src and Dst IP addresses, and form two test packets 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.
- + At the Dst host, arrange to receive the packets.
- + At the Src host, place a timestamp in the first 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.
- + Record this first delay value.
- + At the Src host, place a timestamp in the second 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.
- + By subtracting the second value of One-Way-Delay from the first value the ipdv value of the pair of packets is obtained.
- + If one or both packets fail to arrive within a reasonable period of time, the ipdv is taken to be undefined.

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.

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 next paragraphs.

5.7.1. Errors/Uncertainties related to Clocks

If, as a first approximation, the error that affects the first measurement of One-Way-Delay were the same 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 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 affecting the One-Way-Delay measurement is T_{sync} , and it is a linear function of time, through the skew value "sk", at time T1 the error will be T_{sync1} and at time T2 the error will be T_{sync2} . The ipdv measurement will be affected by the error: $T_{sync2} - T_{sync1} = sk \times (T2 - T1)$ depending on 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 over a significant sample. If the measurement conditions do not allow neglecting the drift, assumed linear in the interval $T2 - T1$, and having a value of "dr" expressed in ppm / sec., the ipdv error will become: $T_{sync2} - T_{sync1} = sk \times (T2 - T1) + [dr \times (T2 - T1) \times (T2 - T1)] / 2$ Drift varies the skew value in the time. The limits in which the skew can vary are anyway limited and small, so that a

given drift cannot act indefinitely. [Section 7](#) and [Appendix A](#) provide more information on this point.

- + As far as accuracy and resolution are concerned, what is noted in the Draft on One-Way-Delay [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.

[5.7.2](#). 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.

[6](#). Definitions for Samples of One-way ipdv

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".

A stream of test packets is generated where the second packet of a pair is, at the same time, the first packet of the next pair.

- + Given particular binding of the parameters Src, Dst and Type-P, a sample of values of parameter T1 is defined. To define the values of T1, 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 second packet of the pair n-1.

[6.1](#). Metric name

Type-P-One-way-ipdv-stream

INTERNET-DRAFT

Packet Delay Variation

June 1999

6.2. Parameters

- + Src, the IP address of a host
- + Dst, the IP address of a host
- + T_0 , a time
- + T_f , a time
- + λ , a rate in reciprocal seconds

6.3. Metric Units:

A sequence of triads whose elements are:

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

6.4. Definition

A pseudo-random Poisson process is defined such that it begins at or before T_0 , with average arrival rate λ , and ends at or after T_f . Those time values $T(i)$ greater than or equal to T_0 and less than or equal to T_f are then selected. Starting from time T_0 , at each pair of times $T(i)$, $T(i+1)$ of this process a value of Type-P-One-way-ipdv is obtained. The value of the sample is the sequence made up of the resulting $\langle \text{time, time interval, ipdv} \rangle$ triple, where the time interval is given by $T(i+1) - T(i)$. Each time $T(i)$, excluding the first and the last, is therefore at the same time the the second time of pair i and the first time of pair $i+1$. The result is shown in figure 3



pair i-1

pair i

pair i+1

Figure 3

[6.5.](#) 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. {Comment: there is, of course, no claim that real Internet traffic arrives according to a Poisson arrival process.}

[6.6.](#) Methodology

Since packets can be lost or duplicated or can arrive in a different order than the order sent, in order to recognize the pairs of test packets, they should be marked with a sequence number. 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 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, inserts in the packet a sequence number and the transmission time stamp T_x , then sorts the time T_i at which the next packet has to be sent.
- + At time T_i , Src repeats the previous step, unless $T(i) > T_f$.

- + On reception of the first packet, or the first packet after a sequence number error, Dst records sequence number and transmission timestamp that are contained in the packet and the reception time Rx as "old values".
- + On reception of the other packets Dst verifies the sequence number and if it is correct, by using the "old values" and the newly received ones, a value of ipdv is computed. Then Dst records the new sequence number, transmit and receive timestamps as "old values".

[6.7.](#) 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 focused in [\[2\]](#). Further considerations will be made in [section 7](#), and in [Appendix A](#).

[6.8.](#) 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.

[6.9.](#) Some 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 focus is on the instantaneous behavior of the connection. Other statistics can be defined if needed.

[6.9.1.](#) 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 ipdv 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 received beyond acceptable limits.

[6.9.2.](#) Type-P-One-way-ipdv-jitter

This metric was defined in [7] and is simply the absolute value of the Type-P-One-way-ipdv. This can be used to derive a number of metrics.

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. Some more detailed calculations are presented in [Appendix A](#).

7.1. Effects of synchronization errors

We refer here to the two components that can generate this type of errors that are the reciprocal "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 reciprocal 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 mostly 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 constant 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 $T_i(1)$. Another ipdv value is measured at the end of T with two packets having a time interval $T_i(2)$.

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

to all the measurement methods.

It is then possible to calculate the values of the Tx and Rx timestamps as they are seen by the two clocks, and the related 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 effect of the skew during the time interval $Ti(i)$ between the two packets of the pair, and a second order term due to the variation of that skew in the same interval.

If, as in the most of practical cases, the drift can be considered close to zero, then $ts1 = ts2$, and the error is not depending on the time at which the measurement is done. In addition, this type of error can be corrected as it is indicated in the next paragraph and discussed in [Appendix A](#).

In any case the maximum error on an ipdv value will correspond to the effect of the maximum reciprocal skew on the maximum interval between packets.

[7.2](#). Related precision

This means that:

- + 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. The mean ipdv value will therefore increased of the quantity $E[Ti(i)] * ts$, which is measured. Also $E[Ti(i)]$ can be measured, and should be related to λ . That means that the skew ts can be calculated. If together with $ipdv(i)$, also the corresponding $Ti(i)$ are collected, for each $ipdv(i)$ value a correcting term is available, and a sample of "corrected" $c-ipdv(i)$ values is obtained, where $c-ipdv(i) = ipdv(i) - Ti(i) * st$.

- + Considering the total skew as subdivided into a fixed part and a variable part (skew and drift), respectively, ts and $+ or - td$, from the mean ipdv value and the mean emission interval the average skew can be derived in the period of interest (Appendix

A). The preceding correction can then be applied. The maximum residual error on an ipdv value is given by the difference between the actual skew at the time in which the value has been measured and the average skew, multiplied by the time interval between the packets that have generated that ipdv value. Considerations on the number of values in the sample affected by errors are reported in [Appendix A](#).

- + If the duration of the measurement is such that it is possible to consider that the ipdv (without skew) is close to zero, the mean value of the ipdv distribution will have the value of the average skew multiplied by the mean value of the emission interval, as supposed above.
- + 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.
- + In what has been said, skew and drift have been considered as "reciprocal". In [Appendix A](#) it will be considered that each of the two clocks have a skew and a drift with respect a "true time", and it will be observed that the difference is negligible with respect the situation in which one of the two clocks is taken as the "true time".

[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 action of the average 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.

This fact offers the possibility of defining a theoretical reference measurement duration in the following way:

The reference duration of a bidirectional ipdv measurement between an host E and an host W is reached at time T_f such that for each time $T > T_f$ the expression $ABS(E(ipdv E-W) - E(ipdv W-E)) < \epsilon$, where ϵ is what we can consider as zero, is always verified. This is one, but not the only method for verifying that the mean ipdv value has reached the value of the average reciprocal skew.

At this point it is possible to evaluate the reciprocal skew. This will require the knowledge of the mean value of the intervals between consecutive packets, that can be calculated over the transmitted stream, by using the collected time stamps.

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-way 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. Problems will anyway arise on the characterization of emission intervals in the backward direction. They would be produced by the combination of the original Poisson arrival process and the effect of ipdv on the forward direction. It has to be studied if this sequence of intervals is still suitable for the measurement. also other possibilities can be envisaged for obtaining a proper backward sequence and still maintain the loopback concept.

[9.](#) Relationship to other standards

The ITU definitions are based on delay variation as defined for ATM cells [\[5\]](#). We will discuss these briefly first and then discuss the ITU's definition for IP packets [\[3\]](#).

[9.1.](#) 1-Point Cell Delay Variation

The ITU looks at cell delay variation from two different points of view. The first, called 1-point cell delay variation, is essentially a measure of how a cell stream varies from a stated cell rate (e.g. the peak cell rate). The basic idea behind the measurement is as follows: The observer at the measurement point notes cell arrival times and clock ticks. The clock ticks at a constant rate, based on the peak cell rate for the cell stream. The difference between the cell arrival times and the clock ticks is the 1-point cell delay variation. If a cell arrives later than the clock tick, the clock

"restarts" at the actual cell arrival time, and continues to tick at |

a constant rate from that point. |

The purpose of this measure is to identify what is called "cell clumping" and non-conforming cells. That is, to identify cells that violate the leaky bucket parameters defined for that cell stream. That is why the clock skips when a cell is later than the normal inter-cell time defined by the peak cell rate. It is of much less interest when cells are late than when they arrive too close together. |

[9.2.](#) 2-Point Delay Variation, Cells and Packets |

2-Point cell delay variation, as defined in [5] is closer to what is defined here. The basic idea behind this metric is that two measurement points, whose clocks are synchronized, observe a cell stream and timestamp when each cell passes. The difference in the timestamps for a cell is essentially the one-way delay. There is also assumed to be a one-way cell delay for a reference cell which we will denote d_0 . The cell delay variation for the i th cell is then $d_i - d_0$. Note that this is not an absolute value, but that the cell delay variation can be either positive or negative. [5] does not specify how to choose the reference cell delay. |

In [3] there is an informative appendix describing packet delay variation, which means that the material is not binding as a standard. The definitions are very similar to [5] with "packet" substituting for "cell" in most places. One difference is that [3] offers two ways to define the reference packet (with the default being the first): |

- + Take the delay of the first packet of the sequence as the reference time. |
- + Take the average one-way packet delay as the reference time. |

[9.3.](#) Discussion |

[9.3.1.](#) Differences

Demichelis [4] points out a number of problems with the 2-point PDV definition in [3]. First of all is the issue of choosing the reference delay time. If this is chosen arbitrarily, it becomes uncertain how to compare the measurements taken from two non-overlapping periods. If it is chosen as an average, that can also be

a problem, because over long periods of time in a network, the average one-way delay can vary widely. A twenty-four hour average as the reference time can seriously overestimate the actual delay variation at a given time of day because the night-time hours, when the delay can be expected to approach the propagation and node time, is included in the average. On the other hand, there is no clear way to partition the time in order to find averages for certain periods of time and compute the delay variation with reference to these averages.

Another problem pointed out in [4] is the fact that 2-point PDV requires synchronized clocks, whereas in this document Demichelis shows that synchronized clocks are not absolutely necessary for ipdv.

[9.3.2.](#) Relationship between the metrics

The ipdv metric described here and the 1-point cell delay variation metric described in [5] do not really have much in common (see also [4]). 1-point delay variation is really intended to talk about the relationship of cell arrival times to a given periodic event, and consequently is more closely related to the first definition of "jitter" given in Section 3 above.

2-point delay variation (actually, the packet variant described in [3]) is related to ipdv, and this relationship can be made precise as follows: Suppose that an arbitrarily chosen packet is designated as the reference packet for the 2-point measurement and also as the start packet of the ipdv measurement. Denote this packet by $p(0)$. Then given ipdv measurements for a series of packets, the 2-point delay variation for packet i is $p(0) +$ the sum from $k=1$ to i of $ipdv(k)$.

Similarly, given a sequence of 2-point delay variation measurements we can derive the ipdv measurement as follows: Denote the 2-point delay variation measurement for packet i as $v(i)$. Then the ipdv value for the pair of packets $p(k-1)$, $p(k)$ is simply $v(k)-v(k-1)$ [6].

[9.3.3. Summary](#)

As described above, there are a number of disadvantages of the 2-point packet delay variation approach. Further, the ipdv approach described here is general enough to provide the same information as the 2-point packet delay variation measurements. Because of this, and because of the (possibly) looser clock synchronization requirements of ipdv, we recommend the one-way-ipdv approach for the delay variation measurement.

[10. Security Considerations](#)

The one-way-ipdv metric has the same security properties as the one-way-delay metric [2]. The packets contain no user information, and so privacy of user data is not a concern. It is still possible that there could be an attempt at a denial of service attack by sending many measurement packets into the network; there could also be attempts to disrupt measurements by diverting packets or corrupting them.

In general, legitimate measurements must have their parameters selected carefully in order to avoid interfering with normal traffic in the network. Such measurements should also be authorized and authenticated in some way so that attacks can be identified and intercepted.

[11. Acknowledgements](#)

Thanks to Matt Zekauskas from Advanced and Ruediger Geib from Deutsche Telekom for discussions relating to the contents of this revised draft.

[12. References](#)

- [1] V.Paxon, G.Almes, J.Mahdavi, M.Mathis - "Framework for IP Performance Metrics", [RFC 2330](#) Feb. 1998
- [2] G.Almes, S.Kalidindi - "A One-Way-Delay Metric for IPPM", Internet Draft <[draft-ietf-ippm-delay-07.txt](#)> May 1999
- [3] Draft New ITU-T Recommendation I.380 "Internet Protocol Data Communication Service - IP Packet Transfer and Availability Performance Parameters"
- [4] Demichelis, Carlo - "Packet Delay Variation Comparison between ITU-T and IETF Draft Definitions" March 1999
- [5] ITU-T Recommendation I.356 "B-ISDN ATM Layer Cell Transfer Performance"
- [6] e-mail exchanges with Ruediger Geib
- [7] V. Jacobson, K. Nichols, K. Poduri - "An expedited forwarding PHB", Internet Draft, November 1998 <[draft-ietf-diffserv-phb-ef-01.txt](#)>

13. Authors' Addresses

Carlo Demichelis <carlo.demichelis@cse.lt.it>
CSELT - Centro Studi E Laboratori Telecomunicazioni S.p.A
Via G. Reiss Romoli 274
10148 - TORINO
Italy
Phone +39 11 228 5057
Fax. +39 11 228 5069

Philip Chimento <chimento@ctit.utwente.nl>
CTIT - Centre for Telematics and Information Technology
University of Twente
Postbox 217
7500 AE Enschede
The Netherlands
Phone +31 53 489 4331

APPENDIX A

This Appendix considers the scenario in which two hosts have clocks that are both not synchronized. Between the two hosts, in an independent way and at the same time in both direction an ipdv measurement is performed according the methodology that is described in the main body of this Draft. This hypothetical scenario is only supposed for discussing the theory and the characteristics of the ipdv metric and its results, without considering implementation issues.

14. Initial positions

The two hosts will be called West (W) and East (E). The two measurements start at the same time, while the end of the measurement it is supposed to be decided by the results of the measurement itself.

At the beginning of the measurement the time declared by the West clock is T_{0w} , the time declared by the East clock is T_{0e} , while the true time is T_{0t} .

The W-clock is affected by an absolute skew of skw ppm and the E-clock by an absolute skew of skw ppm.

The W-clock is affected by an absolute drift ranging from $-drw$ ppm to $+drw$ ppm, the E-clock by an absolute drift ranging from $-dre$ ppm to $+dre$ ppm.

14.1. Evaluation of skew and drift effects

In order to evaluate the effect of the drift on this type of metric, it is necessary to consider the time in which the variation of the skew takes place. We consider the two extreme cases in which the variation takes place uniformly from the beginning to the end of the measurement and the variation takes place suddenly at a generic time

along the measurement. Let TM be the measurement time.

14.1.1.1. Mean ipdv value

Since the mean ipdv value, as it has been seen, is the difference of the last delay minus the first, divided by the number of considered values, we consider what, in the two cases, is measured for first and last delay.

We call trueDf the true first Delay and trueDl the true last Delay. For the evaluation that we want to do, it is not a limitation to consider that they are equal and have a value of trueD. We also consider as time 0 the true time at which the transmission of the first packet starts from West toward East.

In case of continuous drift we define a "drift per second" as: $drpsW = 2*drw / TM$ and $drpsE = 2*dre / TM$ along the measurement this will bring the skew from a value of: $skWmin = skw - drw$; $skEmin = ske - dre$ to a value of $skWmax = skw + drw$; $skEmax = ske + dre$

What is measured as first Delay is:

measured first Rx time - measured first Tx time OffsetEast + trueD*[1 + skEmin + (1/2)*drpsE] - OffsetWest

What is measured as last Delay is:

measured last Rx time - measured last Tx time OffsetEast + (TM + trueD)*[1 + skEmin + (1/2)*2*dre] -
- OffsetWest - TM*[1 + skWmin + (1/2)*2*drw]

The difference between the last and first Delay is therefore:

$TM*(skEmin - skWmin + dre - drw) - trueD*drpsE/(2*TM)$

if TM = 10 hours drpsE is in the order of $50*10E-6 / 36000$ that is about $10E-9$ and the second term of the expression is in the order of $10E-14$ for true delays in the order of 1 sec (negligible term). We consider that, with very good approximation:

Mean emission interval (mti) = TM / number of ipdv values (N)

Therefore:

$$\begin{aligned} \text{mean ipdv} &= (\text{measured last Delay} - \text{measured first Delay}) / N = \\ &= \text{mti} * (\text{skEmin} - \text{skWmin} + \text{dre} - \text{drw}) \end{aligned}$$

but we considered $\text{skEmin} = \text{ske} - \text{dre}$ and $\text{skWmin} = \text{skw} - \text{drw}$ and therefore:

$$\text{mean ipdv} = (\text{meas.lastD} - \text{meas.firstD}) / \text{mti} * (\text{reciprocal mean skew})$$

The previous procedure is now applied to the case in which the total drift takes place in a very short time. Some cases are possible, and we consider the one in which at the beginning the West clock has skWmax and the East clock has skEmin , at time txW the West clock assumes skWmin and at time txE the East clock assumes skEmax . What is measured as first Delay is now:

$$\text{measured first Rx time} - \text{measured first Tx time} - \text{OffsetEast} + \text{trueD} * (1 + \text{skEmin}) - \text{OffsetWest}$$

What is measured as last Delay is:

$$\begin{aligned} &\text{measured last Rx time} - \text{measured last Tx time} + \text{OffsetEast} + \text{txE} * (1 + \text{skEmin}) \\ &+ (\text{TM} - \text{txE}) * (1 + \text{skEmax}) + \text{trueD} * (1 + \text{skEmax}) - \text{OffsetWest} \\ &- \text{txW} * (1 + \text{skWmax}) - (\text{TM} - \text{txW}) * (1 + \text{skWmin}) \end{aligned}$$

but the mean skew values will be:

$$\begin{aligned} \text{mskw} &= [\text{skWmax} * \text{txW} + \text{skWmin} * (\text{TM} - \text{txW})] / \text{TM} \\ \text{mske} &= [\text{skEmin} * \text{txE} + \text{skEmax} * (\text{TM} - \text{txE})] / \text{TM} \end{aligned}$$

the difference between the two delays therefore is:

$$\text{TM} * (\text{mske} - \text{mskw}) + 2 * \text{trueD} * \text{dre}$$

and the mean ipdv value will be:

$$\text{mean ipdv} = \text{mti} * (\text{mske} - \text{mskw}) + 2 * \text{mti} * \text{trueD} * \text{dre} / \text{TM}$$

the second term of the second member in the previous hypotheses is in the order of the nanosecond, and we neglect it. Also in this case, from the mean ipdv value, and knowing the mean emission interval, the

relative skew of the clocks can be obtained.

More in general, independently on how the drift acts inside its limits, we assert that always the mean ipdv value divided by the mean emission interval produces the value of the mean reciprocal skew of the two clocks, provided that the collected number of ipdv values is significant for the statistics.

14.1.2. Errors and corrections

If the drift is always close to zero, it is possible to obtain the true value of the reciprocal skew and correct all the ipdv values. Each of them is associated to an emission interval t_i between the two packets that have produced the value itself. Then a better ipdv value will be: $\text{corr.ipdv}(i) = \text{meas.ipdv}(i) - t_i * \text{skew}$ This is a better value but not exactly the true one, since we supposed that both clocks are not synchronized to the true time. Two errors are affecting the corrective terms which are:

- + The reciprocal skew is measured as referred to the Src clock
- + The interval t_i is measured by the Src clock.

These are second order errors since the measured skew will be affected by a "relative" error in the order of the Src skew, and the same is for the error affecting the t_i value.

If the drift is significant and it can range from the lower to the upper limit of its field, the measured average of the skew will depend on the type of variation. Some cases are considered that demonstrate that actually the proposed correction is not so much effective in this case. Only the fixed part of the total clock variation can be properly corrected.

14.1.3. Constant drift

The first case is the first one considered in the preceding paragraph, where the drift is uniform. We suppose that a reciprocal skew is measured and used for correction.

At the beginning of the measurement the actual reciprocal skew is:

$\text{init.skew} = \text{mean.skew} - \text{rel.max.drift}$

and at the end the actual reciprocal skew is:

INTERNET-DRAFT

Packet Delay Variation

June 1999

$$\text{final.skew} = \text{mean.skew} + \text{rel max.drift}$$

The correction is effective only in the central part of the measurement. At the beginning and at the end a residual error will affect the ipdv values whose value will be:

$$\text{ipdv}(i).\text{err} = t_i * \text{rel.max.drift}$$

We underline here that the error is larger for large intervals t_i and lower for short intervals t_i . For intervals derived from a poissonian arrival process, there are many short intervals and few large intervals. We also note that a constant drift cannot last indefinitely, since there is a minimum and a maximum for the skew.

[14.1.4.](#) Step of drift

In this case the error profile depends on the time at which the drift changes. If the change is near the beginning or near the end of the measurement, the calculated mean skew will be very close to the actual skew of the largest part of the measurement. On that part the correction will be effective, while over the remaining few values the error will be twice with respect the preceding case. The worse condition is produced by a change in drift in the middle of the measurement. In this case the correction would be useful only if the drift was significantly less than the skew.

[14.2.](#) Comparison with a synchronized case

In this section we consider a case in which the two hosts have synchronized clocks, and the synchronization is obtained by setting the real time each second in each of the clocks. We optimistically suppose that this is done exactly (without any imprecision). On the clocks, anyway skew and drift continue to act. We refer to reciprocal skew and drift, having already seen that this is significant. We suppose to perform an ipdv measurement and we evaluate what is measured by the mean ipdv value and what is the error on the measured ipdv values.

We notice, first of all, that nothing changes for ipdv values

measured over intervals falling completely between two synchronization instants. In this case, the effect of synchronization is only to put to zero the offset, that does not appear in the calculation of ipdv values.

Something different happens if the synchronization instant (or more

synchronization instants) falls inside the interval. In this case the error can range from + to - the error related to one second interval, or, more in general, from + to - the error related to an interval equal to the synchronization period. The (few) large intervals will produce a limited error while the (many) short intervals will continue to produce errors of the same order of magnitude of the not synchronized case.

Besides, even if the drift is negligible, the mean ipdv value is no more suitable to calculate the skew, and it will be much more close to zero. Therefore it is no more possible to correct the distortion of the distribution.

Finally, it is necessary to add to these errors the unavoidable imprecision of the synchronization process. We have to consider that the magnitude of errors introduced by skew and drift is in the order of tenth of microseconds. Not always the complete synchronization process has a better precision.

14.3. Bidirectional measurement and components of ipdv

Three terms have been described that can displace the mean ipdv value from zero. They are:

- + The total skew, already discussed above, that always acts in an equal way and opposite direction over the two directions between West and East hosts.
- + The effect of varying traffic that can increase or decrease along limited periods, the average value of the One-Way-Delay. The metric above presented supposes that the measurement period is large enough for considering this effect as tending to zero. It is explicitly noted that the effect will produce a zero effect

only on the mean ipdv value, while the effect on values $ipdv(i)$ is always present. This is not a distortion of the distribution, since it is part of the variation that is measured. This effect is different, and usually concordant, on the two directions.

- + The difference between first and last instantaneous values of the delay variation, that tends to zero when the number of collected ipdv values becomes large.

In order to isolate the last two effects, we consider here a measurement over a long period (e.g. 24 hours) where the drift is negligible, and the effect of the skew has been corrected.

14.4. Slow variation in a given period

The packets of the stream can be represented on a system of cartesian orthogonal axes with transmission time on x-axis and reception time on y-axis, by points localized by transmission and reception time of each packet. Considering an arbitrary period of time T_{per} , which will be a parameter of this procedure, it can be taken as a sliding window over the sample and for each position of this window, established by successive packets, the segment of straight line is calculated that best approximate the points, by means of a linear regression method.

The slope of this segment will be one if along the period the delay has not changed, and different from one if that delay has increased (>1) or decreased (<1). For each position of the window it is therefore possible to find a value of "slow delay variation" with T_{per} as a parameter. This will give an indication on variations produced by different traffic conditions along the measurement period. This item can be subject for further study.

At the same time this procedure offers a criterion for reducing the error introduced in the calculation of the mean ipdv by the instantaneous component of the difference between last and first delay. Supposing that the timestamps, on which the metric is based, are collected and then processed, if the method of the sliding window is applied at the beginning and at the end of the collected sample, it is possible to avoid starting and ending the measurement on values possibly too different from the average (points too far away from the calculated straight line).

Expiration date: December, 1999