

Network Working Group
INTERNET-DRAFT
Expiration Date: December 2000

C. Demichelis
CSELT
P. Chimento
CTIT
July 2000

Instantaneous Packet Delay Variation Metric for IPPM
<[draft-ietf-ippm-ipdv-05.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

July 2000

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 is based on "A One-way-Delay metric for IPPM", [RFC 2679](#) [2]. Part of the text in this memo is taken directly from that document.

This memo defines a metric for variation in delay of packets that flow from one host to another one through an IP path. This quantity 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. 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. For more information 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.

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 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. >

[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 and the value of the Type-P-One-way-Delay from Src to Dst at T1. 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 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 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 relative 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 presented 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.
- + 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.
- + 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.

5.7. 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.

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.
- + 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 two different time stamps ($t_2 > t_1$) is $e(t_2) - e(t_1) = K*(t_2 - t_1)$ which will be added to the time difference ($t_2 - t_1$). If the drift cannot be ignored, but we assume that the drift is a linear function of time, then the skew is given by $s(t) = M*(t**2) + N*t + S_0$, where M and N are constants and S_0 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 $e(t_2) - e(t_1) = N*(t_2 - t_1) + M*\{(t_2 - t_1)**2\}$.

INTERNET-DRAFT

Packet Delay Variation

July 2000

It is the claim here (see remarks in [section 3.3](#)) 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.

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

[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

[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




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. Additional error can be introduced by the pseudo-random > Poisson process as discussed in [2]. Further considerations will be > given in [section 7](#). |

[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, these statistics provide the number of packets exceeding a given limit.

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

This metric is the same as the definition of "jitter" in [7], and is simply the absolute value of the Type-P-One-way-ipdv. |

[6.9.3.](#) 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 jitter or other derivative >
statistic) conditioned on the event that successive 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. We suggest that this be a topic for further >
study. >

[7.](#) Discussion of clock synchronization

This section gives some considerations about the need of having synchronized clocks at the source and destination. These considerations are given as a basis for discussion and they require further investigation. >

[7.1.](#) 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 0 (that is letter 0).

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 timestamps 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 t_1 and that covered by the second measurement is t_2 . Then

$$\text{ipdv}_1 = s(0) \cdot t_1 + t_1 \cdot (s(T) - s(0)) / T$$

$$\text{ipdv}_2 = s(T) \cdot t_2 + t_2 \cdot (s(T) - s(0)) / T$$

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.

[7.2. Estimating the skew of unsynchronized clocks](#)

If the skew is linear (that is, if $s(t) = S \cdot t$ for constant S), the error in ipdv values will depend on the time between the packets used in calculating the value. If t_i is the time between the i th and $(i+1)$ st packet, then let T_i denote the sample mean time between packets and the average skew is $s(T_i) = S \cdot T_i$. Note that $E[T_i]$ should equal $1/\lambda$. In the event that the delays are constant, the skew parameter S can be estimated from the estimate T_i of the time between packets and the sample mean ipdv value. Under these

assumptions, the ipdv values can be corrected by subtracting the

estimated $S * t_i$. >

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. >

g. 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 in 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 relative skew.

At this point it is possible to evaluate the relative 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, 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. For this additional revision, discussions by e-mail with Andy Scherrer were very helpful.

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", [RFC 2679](#), September 1999
- [3] ITU-T Recommendation I.380 "Internet Protocol Data Communication Service - IP Packet Transfer and Availability Performance Parameters", February 1999
- [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", [RFC 2598](#), June 1999

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

Demichelis and Chimento

[Page 19]

INTERNET-DRAFT

Packet Delay Variation

July 2000

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
FAX +31 53 489 4524

Expiration date: December 2000

