

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 8, 2008

A. Morton
AT&T Labs
B. Claise
Cisco Systems, Inc.
July 7, 2007

Packet Delay Variation Applicability Statement
draft-morton-ippm-delay-var-as-03

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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/1id-abstracts.txt>.

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

This Internet-Draft will expire on January 8, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

Packet delay variation metrics appear in many different standards documents. The metric definition in [RFC 3393](#) has considerable flexibility, and it allows multiple formulations of delay variation through the specification of different packet selection functions.

Although flexibility provides wide coverage and room for new ideas, it can make comparisons of independent implementations more

difficult. Two different formulations of delay variation have come into wide use in the context of active measurements. This memo examines a range of circumstances for active measurements of delay variation and their uses, and recommends which of the two forms is best matched to particular conditions and tasks.

Requirements Language

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 [RFC 2119](#) [[RFC2119](#)].

Table of Contents

1.	Introduction	4
1.1.	Background Literature in IPPM and Elsewhere	5
1.2.	Organization of the Memo	6
2.	Purpose and Scope	6
3.	Brief Descriptions of Delay Variation Uses	7
3.1.	Inferring Queue Occupation on a Path	7
3.2.	Determining De-jitter Buffer Size	7
3.3.	Spatial Composition	7
3.4.	Service Level Comparison	8
3.5.	<your favorite here>	8
4.	Formulations of IPDV and PDV	8
4.1.	IPDV: Inter-Packet Delay Variation	8
4.2.	PDV: Packet Delay Variation	9
4.3.	Examples and Initial Comparisons	9
5.	Survey of Earlier Comparisons	10
5.1.	Demichelis' Comparison	11
5.2.	Ciavattone et al.	12
5.3.	IPPM List Discussion from 2000	13
5.4.	Y.1540 Appendix II	14
6.	Additional Properties and Comparisons	14
6.1.	Packet Loss	14
6.2.	Path Changes	15
6.2.1.	Lossless Path Change	16
6.2.2.	Path Change with Loss	17
6.3.	Clock Stability and Error	17
6.4.	Spatial Composition	18
6.5.	Reporting a Single Number	18
6.6.	Jitter in RTCP Reports	19
6.7.	MAPDV2	19
6.8.	Load Balancing	20
7.	Applicability of the Delay Variation Forms and Recommendations	20
7.1.	Uses	20

Morton & Claise

Expires January 8, 2008

[Page 2]

7.1.1.	Inferring Queue Occupancy	20
7.1.2.	Determining De-jitter Buffer Size	21
7.1.3.	Spatial Composition	21
7.2.	Challenging Circumstances	21
7.2.1.	Clock Issues	21
7.2.2.	Frequent Path Changes	22
7.2.3.	Frequent Loss	22
7.2.4.	Load Balancing	22
8.	Measurement Considerations for Vendors, Testers, and Users	22
8.1.	Measurement Stream Characteristics	22
8.2.	Measurement Units	22
8.3.	Test Duration	23
8.4.	Clock Sync Options	23
8.5.	Distinguishing Long Delay from Loss	23
8.6.	Accounting for Packet Reordering	23
8.7.	Results Representation and Reporting	23
9.	IANA Considerations	23
10.	Security Considerations	23
11.	Acknowledgements	23
12.	Appendix on Reducing Delay Variation in Networks	23
13.	Appendix on Calculating the D(min) in PDV	24
14.	References	24
14.1.	Normative References	24
14.2.	Informative References	25
	Authors' Addresses	26
	Intellectual Property and Copyright Statements	28

1. Introduction

There are many ways to formulate packet delay variation metrics for the Internet and other packet-based networks. The IETF itself has several specifications for delay variation [[RFC3393](#)], sometimes called jitter [[RFC3550](#)] or even interarrival jitter [[RFC3550](#)], and these have achieved wide adoption. The International Telecommunication Union - Telecommunication Standardization Sector (ITU-T) has also recommended several delay variation metrics (called parameters in their terminology) [[Y.1540](#)] [[G.1020](#)], and some of these are widely cited and used. Most of the standards above specify more than one way to quantify delay variation, so one can conclude that standardization efforts have tended to be inclusive rather than selective.

This memo uses the term "delay variation" for metrics that quantify a path's ability to transfer packets with consistent delay. [[RFC3393](#)] and [[Y.1540](#)] both prefer this term. Some refer to this phenomenon as "jitter" (and the buffers that attempt to smooth the variations as de-jitter buffers). Applications of the term "jitter" are much broader than packet transfer performance, with "unwanted signal variation" as a general definition. "Jitter" has been used to describe frequency or phase variations, such as data stream rate variations or carrier signal phase noise. The phrase "delay variation" is almost self-defining and more precise, so it is preferred in this memo.

Most (if not all) delay variation metrics are derived metrics, in that their definitions rely on another fundamental metric. In this case, the fundamental metric is one-way delay, and variation is assessed by computing the difference between two individual one-way delay measurements, or a pair of singletons. One of the delay singletons is taken as a reference, and the result is the variation with respect to the reference. The variation is usually summarized for all packets in a stream using statistics.

The industry has predominantly implemented two specific formulations of delay variation (for one survey of the situation, see[Krzanowski]):

1. Inter-Packet Delay Variation, IPDV, where the reference is the previous packet in the stream (according to sending sequence), and the reference changes for each packet in the stream. Properties of variation are coupled with packet sequence in this formulation. This form was called Instantaneous Packet Delay Variation in early contributions.

2. Packet Delay Variation, PDV, where a single reference is chosen from the stream based on specific criteria, and the reference is fixed once selected. The most common criterion for the reference is the packet with the minimum delay in the sample. This term derives its name from a similar definition for Cell Delay Variation, an ATM performance metric.

It is important to note that the authors of relevant standards for delay variation recognized there are many different users with varying needs, and allowed sufficient flexibility to formulate several metrics with different properties. Therefore, the comparison is not so much between standards bodies or their specifications as it is between specific formulations of delay variation. Both Inter-Packet Delay Variation and Packet Delay Variation are compliant with [\[RFC3393\]](#), because different packet selection functions will produce either form.

1.1. Background Literature in IPPM and Elsewhere

With more people joining the measurement community every day, it is possible this memo is the first from the IP Performance Metrics (IPPM) Working Group that the reader has consulted. This section provides a brief roadmap and background on the IPPM literature, and the published specifications of other relevant standards organizations.

The IPPM framework [\[RFC2330\]](#) provides a background for this memo and other IPPM RFCs. Key terms such as singleton, sample, and statistic are defined there, along with methods of collecting samples (Poisson streams), time related issues, and the "packet of Type-P" convention.

There are two fundamental and related metrics that can be applied to every packet transfer attempt: one-way loss [\[RFC2680\]](#) and one-way delay [\[RFC2679\]](#). Lost and delayed packets are separated by a waiting time threshold. Packets that arrive at the measurement destination within their waiting time have finite delay and are not lost. Otherwise, packets are designated lost and their delay is undefined. Guidance on setting the waiting time threshold may be found in [\[RFC2680\]](#) and [\[I-D.morton-ippm-reporting-metrics\]](#).

Another fundamental metric is packet reordering as specified in [\[RFC4737\]](#). The reordering metric was defined to be "orthogonal" to packet loss. In other words, the gap in a packet sequence caused by loss does not result in reordered packets, but a re-arrangement of packet arrivals from their sending order constitutes reordering.

Derived metrics are based on the fundamental metrics. The derived metric of primary interest here is delay variation [\[RFC3393\]](#), a

metric which is derived from one-way delay [[RFC2680](#)]. Another derived metric is the loss patterns metric [[RFC3357](#)], which is derived from loss.

In the ITU-T, the framework, fundamental metrics and derived metrics for IP performance are all specified in Recommendation Y.1540 [[Y.1540](#)].

[1.2.](#) Organization of the Memo

The Purpose and Scope follows in [Section 2](#). We then give a summary of the main tasks for delay variation metrics in [section 3](#). [Section 4](#) defines the two primary forms of delay variation, and [section 5](#) presents summaries of four earlier comparisons. [Section 6](#) adds new comparisons to the analysis, and [section 7](#) reviews the applicability and recommendations for each form of delay variation. [Section 8](#) then looks at many important delay variation measurement considerations. Following IANA and Security Considerations, there two Appendices. One presents guidance on reducing delay variation in networks, and the other calculation of the minimum delay for the PDV form.

[2.](#) Purpose and Scope

The IPDV and PDV formulations have certain features that make them more suitable for one circumstance and less so for another. The purpose of this memo is to compare two forms of delay variation, so that it will be evident which of the two is better suited for each of many possible uses and their related circumstances.

The scope of this memo is limited to the two forms of delay variation briefly described above (Inter-Packet Delay Variation and Packet Delay Variation), circumstances related to active measurement, and uses that are deemed relevant and worthy of inclusion here through IPPM Working Group consensus.

The scope excludes assessment of delay variation for packets with undefined delay. This is accomplished by conditioning the delay distribution on arrival within a reasonable waiting time based on an understanding of the path under test and packet lifetimes. The waiting time is sometimes called the loss threshold [[RFC2680](#)]: if a packet arrives beyond this threshold, it may as well have been lost because it is no longer useful. This is consistent with [[RFC3393](#)], where the Type-P-One-way-ipdv is undefined when the destination fails to receive one or both packets in the selected pair. Furthermore, it is consistent with application performance analysis to consider only arriving packets, because a finite waiting time-out is a feature of many protocols.

3. Brief Descriptions of Delay Variation Uses

This section presents a set of tasks that call for delay variation measurements. Here, the memo provides several answers to the question, "How will the results be used?" for the delay variation metric.

3.1. Inferring Queue Occupation on a Path

As packets travel along the path from source to destination, they pass through many network elements, including a series of router queues. Some types of the delay sources along the path are constant, such as links between two locations. But the latency encountered in each queue varies, depending on the number of packets in the queue when a particular packet arrives. If one assumes that at least one of the packets in a test stream encounters virtually empty queues all along the path (and the path is stable), then the additional delay observed on other packets can be attributed to the time spent in one or more queues. Otherwise, the delay variation observed is the variation in queue time experienced by the test stream.

3.2. Determining De-jitter Buffer Size

Note - while this memo and other IPPM literature prefer the term delay variation, the terms "jitter buffer" and the more accurate "de-jitter buffer" are widely adopted names for a component of packet communication systems, and they will be used here to designate that system component.

Most Isochronous applications (a.k.a. real-time applications) employ a buffer to smooth out delay variation encountered on the path from source to destination. The buffer must be big enough to accommodate (most of) the expected variation, or packet loss will result. However, if the buffer is too large, then some of the desired spontaneity of communication will be lost and conversational dynamics will be affected. Therefore, application designers need to know the extent of delay variation they must accommodate, whether they are designing fixed or adaptive buffer systems.

Network service providers also attempt to constrain delay variation to ensure the quality of real-time applications, and monitor this metric (possibly to compare with a numerical objective or Service Level Agreement).

3.3. Spatial Composition

In Spatial Composition, the tasks are similar to those described above, but with the additional complexity of a multiple network path

where several sub-paths are measured separately, and no source to destination measurements are available. In this case, the source to destination performance must be estimated, using Composed Metrics as described in [[I-D.ietf-ippm-framework-compagg](#)] and [[Y.1541](#)]. Note that determining the composite delay variation is not trivial: simply summing the sub-path variations is not accurate.

[3.4.](#) Service Level Comparison

IP performance measurements are often used as the basis for agreements (or contracts) between service providers and their customers. The measurement results must compare favorably with the performance levels specified in the agreement.

Packet delay variation is usually one of the metrics specified in these agreements. In principle, any formulation could be specified in the Service Level Agreement (SLA). However, the SLA is most useful when the measured quantities can be related to ways in which the communication service will be utilized by the customer, and this can usually be derived from one of the tasks described above.

[3.5.](#) <your favorite here>

Note: At the IETF-68 IPPM session, Alan Clark suggested another possible task for DV measurements, that of detecting and somehow removing the delay variation associated with a smoothing buffer used with a video codec. Further work is needed to define the problem and to investigate the applicability of IPDV and PDV.

[4.](#) Formulations of IPDV and PDV

This section presents the formulations of IPDV and PDV, and provides some illustrative examples. We use the basic singleton definition in [[RFC3393](#)] (which itself is based on [[RFC2679](#)]):

"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 T_2 and the value of the Type-P-One-Way-Delay from Src to Dst at T_1 ."

[4.1.](#) IPDV: Inter-Packet Delay Variation

If we have packets in a stream consecutively numbered $i = 1, 2, 3, \dots$ falling within the test interval, then $IPDV(i) = D(i) - D(i-1)$ where $D(i)$ denotes the one-way-delay of the i th packet of a stream.

One-way delays are the difference between timestamps applied at the

ends of the path, or the receiver time minus the transmission time. So $D(2) = R2 - T2$. With this timestamp notation, it can be shown that IPDV also represents the change in inter-packet spacing between transmission and reception:

$$IPDV(2) = D(2) - D(1) = (R2 - T2) - (R1 - T1) = (R2 - R1) - (T2 - T1)$$

An example selection function given in [RFC3393] is "Consecutive Type-P packets within the specified interval." This is exactly the function needed for IPDV. The reference packet in the pair is always the previous packet in the sending sequence.

Note that IPDV can take on positive and negative values (and zero), although one of the useful ways to analyze the results is to concentrate on the positive excursions. This is discussed in more detail below.

4.2. PDV: Packet Delay Variation

The name Packet Delay Variation is used in [Y.1540] and its predecessors, and refers to a performance parameter equivalent to the metric described below.

The Selection Function for PDV requires two specific roles for the packets in the pair. The first packet is any Type-P packet within the specified interval. The second, or reference packet is the Type-P packet within the specified interval with the minimum one-way-delay.

Therefore, $PDV(i) = D(i) - D(\min)$ (using the nomenclature introduced in the IPDV section). $D(\min)$ is the delay of the packet with the lowest value for delay (minimum) over the current test interval. Values of PDV may be zero or positive, and quantiles of the PDV distribution are direct indications of delay variation.

4.3. Examples and Initial Comparisons

Note: This material originally presented in slides 2 and 3 of

<http://www3.ietf.org/proceedings/06mar/slides/ippm-4.pdf>

The Figure below gives a sample of packet delays and calculates IPDV and PDV values and depicts a histogram for each one.




Figure 1: IPDV and PDV Comparison

The sample of packets contains three packets with "typical" delays of 20ms, one packet with a low delay of 10ms (the minimum of the sample) and one packet with 25ms delay.

As noted above, this example illustrates that IPDV may take on positive and negative values, while the PDV values are greater than or equal to zero. The Histograms of IPDV and PDV are quite different in general shape, and the ranges are different, too (IPDV range = 20ms, PDV range = 15 ms). Note that the IPDV histogram will change if the sequence of delays is modified, but the PDV histogram will stay the same. PDV normalizes the one-way delay distribution to the minimum delay and emphasizes the variation independent from the sequence of delays.

5. Survey of Earlier Comparisons

This section summarizes previous work to compare these two forms of delay variation.

5.1. Demichelis' Comparison

In [[Demichelis](#)], Demichelis compared the early draft versions of two forms of delay variation. Although the IPDV form would eventually see widespread use, the ITU-T work-in-progress he cited did not utilize the same reference packets as PDV. Demichelis compared IPDV with the alternatives of using the delay of the first packet in the stream and the mean delay of the stream as the PDV reference packet. Neither of these alternative references were used in practice, and they are now deprecated in favor of the minimum delay of the stream [[Y.1540](#)].

Active measurements of a transcontinental path (Torino to Tokyo) provided the data for the comparison. The Poisson test stream had 0.764 second average inter-packet interval, with more than 58 thousand packets over 13.5 hours. Among Demichelis' observations about IPDV are the following:

1. IPDV is a measure of the network's ability to preserve the spacing between packets.
2. The distribution of IPDV is usually symmetrical about the origin, having a balance of negative and positive values (for the most part). The mean is usually zero, unless some long-term delay trend is present.
3. IPDV distinguishes quick delay variations (on the order of the interval between packets) from longer term variations.
4. IPDV places reduced demands on the stability and skew of measurement clocks.

He also notes these features of PDV:

1. PDV does not distinguish quick variation from variation over the complete test interval.
2. The location of the distribution is very sensitive to the delay of the first packet, IF this packet is used as the reference. This would be a new formulation that differs from the PDV definition in this memo (PDV references the packet with minimum delay, so it does not have this drawback).
3. The shape of the PDV distribution is identical to the delay distribution, but shifted by the reference delay.
4. Use of a common reference over measurement intervals that are longer than a typical session length may indicate more PDV than

would be experienced by streams that support such sessions.

5. PDV characterizes the range of queue occupancies along the measurement path (assuming the path is fixed), but the range says nothing about how the variation took place.

The summary metrics used in this comparison were the number of values exceeding a +/-50ms range around the mean, the Inverse Percentiles, and the Inter-Quartile Range.

5.2. Ciavattone et al.

In [[Cia03](#)], the authors compared IPDV and PDV (referred to as delta) using a periodic packet stream conforming to [[RFC3432](#)] with inter-packet interval of 20 ms.

One of the comparisons between IPDV and PDV involves a laboratory set-up where a queue was temporarily congested by a competing packet burst. The additional queuing delay was 85ms to 95ms, much larger than the inter-packet interval. The first packet in the stream that follows the competing burst spends the longest time enqueued, and others experience less and less queuing time until the queue is drained.

The authors observed that PDV reflects the additional queuing time of the packets affected by the burst, with values of 85, 65, 45, 25, and 5ms. Also, it is easy to determine (by looking at the PDV range) that a de-jitter buffer of 90 ms would have been sufficient to accommodate the delay variation.

The distribution of IPDV values in the congested queue example are very different: 85, -20, -20, -20, -20, -5ms. Only the positive excursion of IPDV gives an indication of the de-jitter buffer size needed. Although the variation exceeds the inter-packet interval, the extent of negative IPDV values is limited by that sending interval. This preference for information from the positive IPDV values has prompted some to ignore the negative values, or to take the absolute value of each IPDV measurement (sacrificing key properties of IPDV in the process, such as its ability to distinguish delay trends).

Elsewhere, the authors considered the range as a summary statistic for IPDV, and the 99.9%-ile minus the minimum delay as a summary statistic for delay variation, or PDV.

5.3. IPPM List Discussion from 2000

Mike Pierce made many comments in the context of the 05 version of [draft-ietf-ippm-ipdv](#). One of his main points was that a delay histogram is a useful approach to quantifying variation. Another was that the time duration of evaluation is a critical aspect.

Carlo Demichelis then mailed his comparison paper to the IPPM list, [[Demichelis](#)] as discussed in more detail above.

Ruediger Geib observed that both IPDV and the delay histogram (PDV) are useful, and suggested that they might be applied to different variation time scales. He pointed out that loss has a significant effect on IPDV, and encouraged that the loss information be retained in the arrival sequence.

Several example delay variation scenarios were discussed, including:

Packet #	1	2	3	4	5	6	7	8	9	10	11

Ex. A											
Lost											
Delay, ms	100	110	120	130	140	150	140	130	120	110	100
IPDV	U	10	10	10	10	10	-10	-10	-10	-10	-10
PDV	0	10	20	30	40	50	40	30	20	10	0

Ex. B											
Lost				L							
Delay, ms	100	110	150	U	120	100	110	150	130	120	100
IPDV	U	10	40	U	U	-10	10	40	-20	-10	-20
PDV	0	10	50	U	20	0	10	50	30	20	0

Figure 2: Delay Examples

Clearly, the range of PDV values is 50 ms in both cases above, while the IPDV range tends to minimize the smooth variation in Example A (20 ms), and responds to the faster variations in Example B (60 ms).

IPDV values can be viewed as the adjustments that an adaptive de-jitter buffer would make, IF it could make adjustments on a packet-by-packet basis. However, adaptive de-jitter buffers don't make

adjustments this frequently, so the value of this information is unknown.

5.4. Y.1540 [Appendix II](#)

This Appendix compares IPDV, PDV (referred to as 2-point PDV), and 1-point packet delay variation (which assume a periodic stream and assesses variation against an ideal arrival schedule constructed at the single measurement point).

6. Additional Properties and Comparisons

This section treats some of the earlier comparison areas in more detail, and introduces new areas for comparison.

6.1. Packet Loss

The measurement packet loss is of great influence for the delay variation results, as displayed in the figure 2 and 3 (L means Lost and U means undefined). Figure 3 shows that in the extreme case of every other packet loss, the IPDV doesn't produce any results, while the PDV produces results for all arriving packets.

Packet #	1	2	3	4	5	6	7	8	9	10
Lost		L		L		L		L		L

Delay, ms	3	U	5	U	4	U	3	U	4	U
IPDV	U	U	U	U	U	U	U	U	U	U
PDV	0	U	2	U	1	U	0	U	1	U

Figure 3: Path Loss Every Other Packet

In case of a burst of packet loss, as displayed in figure 3, both the IPDV and PDV produces some results. Note that the PDV.

Packet #	1	2	3	4	5	6	7	8	9	10
Lost			L	L	L	L	L			

Delay, ms	3	4	U	U	U	U	U	5	4	3
IPDV	U	1	U	U	U	U	U	U	-1	-1
PDV	0	1	U	U	U	U	U	2	1	0

Figure 4: Burst of Packet Loss

In conclusion, the PDV results are affected by the packet loss ratio. While the IPDV results are affected by the packet loss ratio, they are also affected by the packet loss distribution. Indeed, in the extreme case of every other packet loss, the IPDV doesn't provide any results.

6.2. Path Changes

When there is little or no stability in the network under test, then the devices that attempt to characterize the network are equally stressed, especially if the results displayed are used to make inferences which may not be valid.

Sometimes the path characteristics change during a measurement interval. The change may be due to link or router failure, administrative changes prior to maintenance (e.g., link cost change), or re-optimization of routing using new information. All these causes are usually infrequent, and network providers take appropriate measures to ensure this. Automatic restoration to a back-up path is seen as a desirable feature of IP networks.

Frequent path changes and prolonged congestion with substantial packet loss clearly make delay variation measurements challenging. Path changes are usually accompanied by a sudden, persistent increase or decrease in one-way-delay. [Cia03] gives one such example. We assume that a restoration path either accepts a stream of packets, or is not used for that particular stream (e.g., no multi-path for flows).

In any case, a change in the TTL (or Hop Limit) of the received packets indicates that the path is no longer the same. Transient packet reordering may also be observed with path changes, due to use of non-optimal routing while updates propagate through the network (see [Casner] and [Cia03])

Many, if not all, packet streams experience packet loss in conjunction with a path change. However, it is certainly possible

that the active measurement stream does not experience loss. This may be due to use of a long inter-packet sending interval with respect to the restoration time, and this becomes more likely as "fast restoration" techniques see wider deployment (e.g., [[RFC4090](#)]).

Thus, there are two main cases to consider, path changes accompanied by loss, and those that are lossless from the point of view of the active measurement stream. The subsections below examine each of these cases.

6.2.1. Lossless Path Change

In the lossless case, a path change will typically affect only two IPDV singletons. However, if the change in delay is negative and larger than the inter-packet sending interval, then more than two IPDV singletons may be affected because packet reordering is also likely to occur.

The use of the new path and its delay variation can be quantified by treating the PDV distribution as bi-modal, and characterizing each mode separately. This would involve declaring a new path within the sample, and using a new local minimum delay as the PDV reference delay for the sub-sample (or time interval) where the new path is present.

The process of detecting a bi-modal delay distribution is made difficult if the typical delay variation is larger than the delay change associated with the new path. However, information on TTL (or Hop Limit) change or the presence of transient reordering can assist in an automated decision.

The effect of path changes may also be reduced by making PDV measurements over short intervals (minutes, as opposed to hours). This way, a path change will affect one sample and its PDV values. Assuming that the mean or median one-way-delay changes appreciably on the new path, then subsequent measurements can confirm a path change, and trigger special processing on the interval containing a path change and the affected PDV result.

Alternatively, if the path change is detected, by monitoring the test packets TTL or Hop Limit, or monitoring the change in the IGP link-state database, the results of measurement before and after the path change could be kept separated, presenting two different distributions. This avoids the difficult task of determining the different modes of a multi-modal distribution.

6.2.2. Path Change with Loss

If the path change is accompanied by loss, such that there are no consecutive packet pairs that span the change, then no IPDV singletons will reflect the change. This may or may not be desirable, depending on the ultimate use of the delay variation measurement. The Figure 3, in which L means Lost and U means undefined, illustrates this case.

Packet #	1	2	3	4	5	6	7	8	9	
Lost					L	L				

Delay, ms	3	4	3	3	U	U	8	9	8	
IPDV		U	1	-1	0	U	U	U	1	-1
PDV		0	1	0	0	U	U	5	6	5

Figure 5: Path Change with Loss

PDV will again produce a bimodal distribution. But here, the decision process to define sub-intervals associated with each path is further assisted by the presence of loss, in addition to TTL, reordering information, and use of short measurement intervals consistent with the duration of user sessions. It is reasonable to assume that at least loss and delay will be measured simultaneously with PDV and/or IPDV.

6.3. Clock Stability and Error

Low cost or low complexity measurement systems may be embedded in communication devices that do not have access to high stability clocks, and time errors will almost certainly be present. However, larger time-related errors may offer an acceptable trade-off for monitoring performance over a large population (the accuracy needed to detect problems may be much less than required for a scientific study).

As mentioned above, [Demichelis] observed that PDV places greater demands on clock synchronization than for IPDV. This observation deserves more discussion. Synchronization errors have two components: time of day errors and clock frequency errors (resulting in skew).

Both IPDV and PDV are sensitive to time-of-day errors when attempting to align measurement intervals at the source and destination. Gross mis-alignment of the measurement intervals can lead to lost packets, for example if the receiver is not ready when the first test packet

arrives. However, both IPDV and PDV assess delay differences, so the error present in two one-way-delay singletons will cancel as long as it is constant. So, NTP or GPS synchronization is not required to correct the time-of-day error in case the delay variation measurement, while it is required for the one-way delay measurement.

Skew is a measure of the change in clock time over an interval w.r.t. a reference clock. Both IPDV and PDV are affected by skew, but the error sensitivity in IPDV singletons is less because the intervals between consecutive packets are rather small, especially when compared to the overall measurement interval. Since PDV computes the difference between a single reference delay (the sample minimum) and all other delays in the measurement interval, the constraint on skew error is greater to attain the same accuracy as IPDV. Again, use of short PDV measurement intervals (on the order of minutes, not hours) provides some relief from the effects of skew error.

If skew is present in a sample of one-way-delays, its symptom is typically a linear growth or decline over all the one-way-delay values. As a practical matter, if the same slope appears consistently in the measurements, then it may be possible to fit the slope and compensate for the skew in the one-way-delay measurements, thereby avoiding the issue in the PDV calculations that follow. See [[RFC3393](#)] for additional information on compensating for skew.

There is a third factor related to clock error and stability: this is the presence of a clock synchronization protocol (e.g., NTP) and the time adjustment operations that result. When a small time error is detected (typically on the order of a few milliseconds), the host clock frequency is continuously adjusted to reduce the time error. If these adjustments take place during a measurement interval, they may appear as delay variation when none was present, and therefore are a source of error.

[6.4.](#) Spatial Composition

ITU-T Recommendation [[Y.1541](#)] gives a provisional method to compose a PDV metric using PDV measurement results from two or more sub-paths.

PDV has a clear advantage at this time, since there is no known method to compose an IPDV metric. In addition, IPDV results depend greatly on the exact sequence of packets and may not lend themselves easily to the composition problem.

[6.5.](#) Reporting a Single Number

Despite the risk of over-summarization, measurements must often be displayed for easy consumption. If the right summary report is

prepared, then the "dashboard" view correctly indicates whether there is something different and worth investigating further, or that the status has not changed. The dashboard model restricts every instrument display to a single number. The packet network dashboard could have different instruments for loss, delay, delay variation, reordering, etc., and each must be summarized as a single number for each measurement interval.

The simplicity of the PDV distribution lends itself to this summarization process (including use of the median or mean). [\[Y.1541\]](#) introduced the notion of a pseudo-range when setting an objective for the 99.9%-ile of PDV. The conventional range (max-min) was avoided for several reasons, including stability of the maximum delay. The 99.9%-ile of PDV is helpful to performance planners (seeking to meet some user-to-user objective for delay) and in design of de-jitter buffer sizes, even those with adaptive capabilities.

IPDV does not lend itself to summarization so easily. The mean IPDV is typically zero. As the IPDV distribution may have two tails (positive and negative) the range or pseudo-range would not match the needed de-jitter buffer size. Additional complexity may be introduced when the variation exceeds the inter-packet sending interval, as discussed above. Should the Inter-Quartile Range be used? Should the singletons beyond some threshold be counted (e.g., mean +/- 50ms)? A strong rationale for one of these summary statistics has yet to emerge.

6.6. Jitter in RTCP Reports

[\[RFC3550\]](#) gives the calculation of the inter-arrival Jitter field for the RTCP report, with a sample implementation in an Appendix.

The RTCP Jitter value can be calculated using IPDV singletons. If there is packet reordering, as defined in [\[RFC4737\]](#), then estimates of Jitter based on IPDV may vary slightly, because [\[RFC3550\]](#) specifies the use of receive packet order.

Just as there is no simple way to convert PDV singletons to IPDV singletons without returning to the original sample of delay singletons, there is no clear relationship between PDV and [\[RFC3550\]](#) Jitter.

6.7. MAPDV2

MAPDV2 stands for Mean Absolute Packet Delay Variation (version) 2, and is specified in [\[G.1020\]](#). The MAPDV2 algorithm computes a smoothed running estimate of the mean delay using the one-way delays of 16 previous packets. It compares the current one-way-delay to the

estimated mean, separately computes the means of positive and negative deviations, and sums these deviation means to produce MAPVDV2. In effect, there is a MAPDV2 singleton for every arriving packet, so further summarization is usually warranted.

Neither IPDV or PDV forms assist in the computation of MAPDV2.

6.8. Load Balancing

TO DO: What if there is load-balancing in an ISP network? Load-balancing is based on the IGP metrics, while the delay depends on the path. So, we could have a multi-modal distribution, if we send test packets with different characteristics such as IP addresses/ports.

Should the delay singletons using multiple addresses/ports be combined in the same sample?

The PDV form makes the identification of multiple modes possible. Should we characterize each mode separately? This question also applies to the Path Change case.

7. Applicability of the Delay Variation Forms and Recommendations

Based on the comparisons of IPDV and PDV presented above, this section matches the attributes of each form with the tasks described earlier. We discuss the more general circumstances first.

Note: the conclusions of this section should be regarded as preliminary, pending discussion and further development by the IPPM WG.

7.1. Uses

7.1.1. Inferring Queue Occupancy

The PDV distribution is anchored at the minimum delay observed in the measurement interval. When the sample minimum coincides with the true minimum delay of the path, then the PDV distribution is equivalent to the queuing time distribution experienced by the test stream. If the minimum delay is not the true minimum, then the PDV distribution captures the variation in queuing time and some additional amount of queuing time is experienced, but unknown. One can summarize the PDV distribution with the mean, median, and other statistics.

IPDV can capture the difference in queuing time from one packet to the next, but this is a different distribution from the queue

occupancy revealed by PDV.

7.1.2. Determining De-jitter Buffer Size

This task is complimentary to the problem of inferring queue occupancy through measurement. Again, use of the sample minimum as the reference delay for PDV yields a distribution that is very relevant to de-jitter buffer size. This is because the minimum delay is an alignment point for the smoothing operation of de-jitter buffers. A de-jitter buffer that is ideally aligned with the delay variation adds zero buffer time to packets with the longest accommodated network delay (any packets with longer delays are discarded). Thus, a packet experiencing minimum network delay should be aligned to wait the maximum length of the de-jitter buffer. With this alignment, the stream is smoothed with no unnecessary delay added. [G.1020] illustrates the ideal relationship between network delay variation and buffer time.

The PDV distribution is also useful for this task, but different statistics are preferred. The range (max-min) or the 99.9%-ile of PDV (pseudo-range) are closely related to the buffer size needed to accommodate the observed network delay variation.

In some cases, the positive excursions (or series of positive excursions) of IPDV may help to approximate the de-jitter buffer size, but there is no guarantee that a good buffer estimate will emerge, especially when the delay varies as a positive trend over several test packets.

7.1.3. Spatial Composition

PDV has a clear advantage at this time, since there is no known method to compose an IPDV metric.

7.2. Challenging Circumstances

Note that measurement of delay variation may not be the primary concern under unstable and unreliable circumstances.

7.2.1. Clock Issues

When appreciable skew is present between measurement system clocks, then IPDV has a clear advantage because PDV would require processing over the entire sample to remove the skew error. Neither form of delay variation is more suited than the other to on-the-fly summarization without memory, and this may be one of the reasons that [RFC3550] RTCP Jitter and MAPDV2 in [G.1020] have attained deployment in low-cost systems.

7.2.2. Frequent Path Changes

If the network under test exhibits frequent path changes, on the order of several new routes per minute, then IPDV appears to isolate the delay variation on each path from the transient effect of path change (especially if there is packet loss at the time of path change). It is possible to make meaningful PDV measurements when paths are unstable, but great importance would be placed on the algorithms that infer path change and attempt to divide the sample on path change boundaries.

7.2.3. Frequent Loss

If the network under test exhibits frequent loss, then PDV may produce a larger set of singletons for the sample than IPDV. This is due to IPDV requiring consecutive packet arrivals to assess delay variation, compared to PDV where any packet arrival is useful. The worst case is when no consecutive packets arrive, and the entire IPDV sample would be undefined. PDV would successfully produce a sample based on the arriving packets.

7.2.4. Load Balancing

TO DO: this section will give the brief conclusions of the discussion and analysis in [section 6](#).

8. Measurement Considerations for Vendors, Testers, and Users

TO DO:

8.1. Measurement Stream Characteristics

8.2. Measurement Units

TO DO: Al, you mentioned somewhere above: "These devices may not have sufficient memory to store all singletons for later processing." And also an important conclusion: "Just as there is no simple way to convert PDV singletons to IPDV singletons without returning to the original sample of delay singletons" I'm thinking we should develop around that:

- Do we want to get IPDV and DV?
- Do we want to reconstruct the IPDV and DV later on, with a different interval?

+ a reference to the appendix where we would describe:

- how is this $D(\min)$ calculated? Is it DV(99%) as mentioned by Roman in <http://www3.ietf.org/proceedings/05nov/slides/ippm-3.pdf?>
- do we need to keep all the values from the interval, then take the minimum? Or do we keep the minimum from previous intervals?

8.3. Test Duration

8.4. Clock Sync Options

8.5. Distinguishing Long Delay from Loss

Setting the max waiting time threshold...

8.6. Accounting for Packet Reordering

8.7. Results Representation and Reporting

9. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

10. Security Considerations

The security considerations that apply to any active measurement of live networks are relevant here as well. See [[RFC4656](#)]

11. Acknowledgements

The author would like to thank Phil Chimento for his suggestion to employ the convention of conditional distributions for Delay to deal with packet loss, and his encouragement to "write the memo" after hearing the talk on this topic at IETF-65.

12. Appendix on Reducing Delay Variation in Networks

This text is both preliminary and generic but we want to explain the basic troubleshooting.

If there is a DV problem, it may be because:

1. there is congestion. Find where the bottleneck is, and increase the buffer Alternatively, increase the bandwidth Alternatively, remove some applications from that class of service
2. there is a variability of the traffic Discover that traffic, then change/apply QoS (for example, rate-limiting)

13. Appendix on Calculating the D(min) in PDV

Note: These are the questions this section intends to answer:

- how is this D(min) calculated? Is it DV(99%) as mentioned by Roman in <http://www3.ietf.org/proceedings/05nov/slides/ippm-3.pdf?>
- do we need to keep all the values from the interval, then take the minimum? Or do we keep the minimum from previous intervals?

The value of D(min) used as the reference delay for PDV calculations is simply the minimum delay of all packets in the current sample. The usual single value summary of the PDV distribution is D(99.9%-ile) minus D(min).

It may be appropriate to segregate sub-sets and revise the minimum value during a sample. For example, if it can be determined with certainty that the path has changed by monitoring the Time to Live or Hop Count of arriving packets, this may be sufficient justification to reset the minimum for packets on the new path.

It is not necessary to store all delay values in a sample when storage is a major concern. D(min) can be found by comparing each new singleton value with the current value and replacing it when required. In a sample with 5000 packets, evaluation of the 99.9%-ile can also be achieved with limited storage. One method calls for storing the top 50 delay singletons and revising the top value list each time 50 more packets arrive.

14. References

14.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", [RFC 2330](#), May 1998.

- [RFC2679] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM", [RFC 2679](#), September 1999.
- [RFC2680] Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM", [RFC 2680](#), September 1999.
- [RFC3393] Demichelis, C. and P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)", [RFC 3393](#), November 2002.
- [RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", [RFC 3432](#), November 2002.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.
- [RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", [RFC 4090](#), May 2005.
- [RFC4656] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and M. Zekauskas, "A One-way Active Measurement Protocol (OWAMP)", [RFC 4656](#), September 2006.
- [RFC4737] Morton, A., Ciavattone, L., Ramachandran, G., Shalunov, S., and J. Perser, "Packet Reordering Metrics", [RFC 4737](#), November 2006.

[14.2.](#) Informative References

- [Casner] "A Fine-Grained View of High Performance Networking, NANOG 22 Conf.; <http://www.nanog.org/mtg-0105/agenda.html>", May 20-22 2001.
- [Cia03] "Standardized Active Measurements on a Tier 1 IP Backbone, IEEE Communications Mag., pp 90-97.", June 2003.
- [Demichelis]
<http://www.advanced.org/ippm/archive.3/att-0075/01-pap02.doc>, "Packet Delay Variation Comparison between ITU-T and IETF Draft Definitions", November 2000.
- [G.1020] ITU-T Recommendation G.1020, "'Performance parameter definitions for the quality of speech and other voiceband applications utilizing IP networks'", 2006.

[I-D.ietf-ippm-framework-compagg]

Morton, A. and S. Berghe, "Framework for Metric Composition", [draft-ietf-ippm-framework-compagg-03](#) (work in progress), March 2007.

[I-D.morton-ippm-reporting-metrics]

Morton, A., "Reporting Metrics: Different Points of View", [draft-morton-ippm-reporting-metrics-02](#) (work in progress), April 2007.

[Krzanowski]

Presentation at IPPM, IETF-64, "Jitter Definitions: What is What?", November 2005.

[RFC3357] Koodli, R. and R. Ravikanth, "One-way Loss Pattern Sample Metrics", [RFC 3357](#), August 2002.

[Y.1540] ITU-T Recommendation Y.1540, "Internet protocol data communication service - IP packet transfer and availability performance parameters", December 2002.

[Y.1541] ITU-T Recommendation Y.1540, "Network Performance Objectives for IP-Based Services", February 2006.

Authors' Addresses

Al Morton
AT&T Labs
200 Laurel Avenue South
Middletown,, NJ 07748
USA

Phone: +1 732 420 1571
Fax: +1 732 368 1192
Email: acmorton@att.com
URI: <http://home.comcast.net/~acmacm/>

Benoit Claise
Cisco Systems, Inc.
De Kleetlaan 6a b1
Diegem, 1831
Belgium

Phone: +32 2 704 5622

Fax:

Email: bclaise@cisco.com

URI:

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

