

Network Working Group
Request for Comments: 5481
Category: Informational

A. Morton
AT&T Labs
B. Claise
Cisco Systems, Inc.
March 2009

Packet Delay Variation Applicability Statement

Status of This Memo

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

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

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.

Table of Contents

1. Introduction	4
1.1. Requirements Language	5
1.2. Background Literature in IPPM and Elsewhere	5
1.3. Organization of the Memo	6
2. Purpose and Scope	7
3. Brief Descriptions of Delay Variation Uses	7
3.1. Inferring Queue Occupation on a Path	7
3.2. Determining De-Jitter Buffer Size	8
3.3. Spatial Composition	10
3.4. Service-Level Comparison	10
3.5. Application-Layer FEC Design	10
4. Formulations of IPDV and PDV	10
4.1. IPDV: Inter-Packet Delay Variation	11
4.2. PDV: Packet Delay Variation	11
4.3. A "Point" about Measurement Points	12
4.4. Examples and Initial Comparisons	12
5. Survey of Earlier Comparisons	13
5.1. Demichelis' Comparison	13
5.2. Ciavattone et al.	15
5.3. IPPM List Discussion from 2000	16
5.4. Y.1540 Appendix II	18
5.5. Clark's ITU-T SG 12 Contribution	18
6. Additional Properties and Comparisons	18
6.1. Packet Loss	18
6.2. Path Changes	19
6.2.1. Lossless Path Change	20
6.2.2. Path Change with Loss	21
6.3. Clock Stability and Error	22
6.4. Spatial Composition	24
6.5. Reporting a Single Number (SLA)	24
6.6. Jitter in RTCP Reports	25

6.7. MAPDV2	25
6.8. Load Balancing	26
7. Applicability of the Delay Variation Forms and Recommendations	27
7.1. Uses	27
7.1.1. Inferring Queue Occupancy	27
7.1.2. Determining De-Jitter Buffer Size (and FEC Design)	27
7.1.3. Spatial Composition	28
7.1.4. Service-Level Specification: Reporting a Single Number	28
7.2. Challenging Circumstances	28
7.2.1. Clock and Storage Issues	28
7.2.2. Frequent Path Changes	29
7.2.3. Frequent Loss	29
7.2.4. Load Balancing	29
7.3. Summary	30
8. Measurement Considerations	31
8.1. Measurement Stream Characteristics	31
8.2. Measurement Devices	32
8.3. Units of Measurement	33
8.4. Test Duration	33
8.5. Clock Sync Options	33
8.6. Distinguishing Long Delay from Loss	34
8.7. Accounting for Packet Reordering	34
8.8. Results Representation and Reporting	35
9. Security Considerations	35
10. Acknowledgments	35
11. Appendix on Calculating the D(min) in PDV	35
12. References	36
12.1. Normative References	36
12.2. Informative References	37

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 inter-arrival 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 IETF contributions, and is similar to the packet spacing difference metric used for interarrival jitter calculations in [RFC3550].

2. Packet Delay Variation, PDV, where a single reference is chosen from the stream based on specific criteria. 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 [I.356].

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. 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].

1.2. 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 road map 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]. The metrics use a waiting time threshold to distinguish between lost and delayed packets. 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 [IPPM-Reporting].

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 rearrangement of packet arrivals from their sending order constitutes reordering.

Derived metrics are based on the fundamental metrics. The metric of primary interest here is delay variation [RFC3393], a metric that is derived from one-way delay [RFC2680]. Another derived metric is the loss patterns metric [RFC3357], which is derived from loss.

The measured values of all metrics (both fundamental and derived) depend to great extent on the stream characteristics used to collect them. Both Poisson streams [RFC3393] and Periodic streams [RFC3432] have been used with the IPDV and PDV metrics. The choice of stream specification for active measurement will depend on the purpose of the characterization and the constraints of the testing environment. Periodic streams are frequently chosen for use with IPDV and PDV, because the application streams that are most sensitive to delay variation exhibit periodicity. Additional details that are method-specific are discussed in Section 8 on "Measurement Considerations".

In the ITU-T, the framework, fundamental metrics, and derived metrics for IP performance are specified in Recommendation Y.1540 [Y.1540]. [G.1020] defines additional delay variation metrics, analyzes the operation of fixed and adaptive de-jitter buffers, and describes an example adaptive de-jitter buffer emulator. Appendix II of [G.1050] describes the models for network impairments (including delay variation) that are part of standardized IP network emulator that may be useful when evaluating measurement techniques.

1.3. 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 the Security Considerations, there is an appendix on the 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.

It is entirely possible that the analysis and conclusions drawn here are applicable beyond the intended scope, but the reader is cautioned to fully appreciate the circumstances of active measurement on IP networks before doing so.

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.

It is worth noting that delay variation can occur beyond IP router queues, in other communication components. Examples include media contention: DOCSIS, IEEE 802.11, and some mobile radio technologies.

However, delay variation from all sources at the IP layer and below will be quantified using the two formulations discussed here.

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 the expected variation of delay, 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 range 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).

De-jitter buffer size can be expressed in units of octets of storage space for the packet stream, or in units of time that the packets are stored. It is relatively simple to convert between octets and time when the buffer read rate (in octets per second) is constant:

$$\text{read_rate} * \text{storage_time} = \text{storage_octets}$$

Units of time are used in the discussion below.

The objective of a de-jitter buffer is to compensate for all prior sources of delay variation and produce a packet stream with constant delay. Thus, a packet experiencing the minimum transit delay from source to destination, D_{\min} , should spend the maximum time in a

de-jitter buffer, B_{\max} . The sum of D_{\min} and B_{\max} should equal the sum of the maximum transit delay (D_{\max}) and the minimum buffer time (B_{\min}). We have

$$\text{Constant} = D_{\min} + B_{\max} = D_{\max} + B_{\min},$$

after rearranging terms,

$$B_{\max} - B_{\min} = D_{\max} - D_{\min} = \text{range}(B) = \text{range}(D)$$

where $\text{range}(B)$ is the range of packet buffering times, and $\text{range}(D)$ is the range of packet transit delays from source to destination.

Packets with transit delay between the max and min spend a complementary time in the buffer and also see the constant delay.

In practice, the minimum buffer time, B_{\min} , may not be zero, and the maximum transit delay, D_{\max} , may be a high percentile (99.9th percentile) instead of the maximum.

Note that $B_{\max} - B_{\min} = \text{range}(B)$ is the range of buffering times needed to compensate for delay variation. The actual size of the buffer may be larger (where $B_{\min} > 0$) or smaller than $\text{range}(B)$.

There must be a process to align the de-jitter buffer time with packet transit delay. This is a process to identify the packets with minimum delay and schedule their play-out time so that they spend the maximum time in the buffer. The error in the alignment process can be accounted for by a variable, A . In the equation below, the range of buffering times *available* to the packet stream, $\text{range}(b)$, depends on buffer alignment with the actual arrival times of D_{\min} and D_{\max} .

$$\text{range}(b) = b_{\max} - b_{\min} = D_{\max} - D_{\min} + A$$

where variable b represents the *available* buffer in a system with a specific alignment, A , and b_{\max} and b_{\min} represent the limits of the available buffer.

When A is positive, the de-jitter buffer applies more delay than necessary (where $\text{Constant} = D_{\max} + b_{\min} + A$ represents one possible alignment). When A is negative, there is insufficient buffer time available to compensate for $\text{range}(D)$ because of misalignment. Packets with D_{\min} may be arriving too early and encountering a full buffer, or packets with D_{\max} may be arriving too late, and in either case, the packets would be discarded.

In summary, the range of transit delay variation is a critical factor in the determination of de-jitter buffer size.

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 [IPPM-Framework] 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. Application-Layer FEC Design

The design of application-layer Forward Error Correction (FEC) components is closely related to the design of a de-jitter buffer in several ways. The FEC designer must choose a protection interval (time to send/receive a block of packets in a constant packet rate system) consistent with the packet-loss characteristics, but also mindful of the extent of delay variation expected. Further, the system designer must decide how long to wait for "late" packets to arrive. Again, the range of delay variation is the relevant expression delay variation for these tasks.

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 T2 and the value of the Type-P-One-Way-Delay from Src to Dst at T1".

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 the previous packet in the sending sequence.

Note that IPDV can take on positive and negative values (and zero). One way to analyze the IPDV results is to concentrate on the positive excursions. However, this approach has limitations that are discussed in more detail below (see Section 5.3).

The mean of all $IPDV(i)$ for a stream is usually zero. However, a slow delay change over the life of the stream, or a frequency error between the measurement system clocks, can result in a non-zero mean.

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.

PDV is a version of the one-way-delay distribution, shifted to the origin by normalizing to the minimum delay.

4.3. A "Point" about Measurement Points

Both IPDV and PDV are derived from the one-way-delay metric. One-way delay requires knowledge of time at two points, e.g., the source and destination of an IP network path in end-to-end measurement. Therefore, both IPDV and PDV can be categorized as 2-point metrics because they are derived from one-way delay. Specific methods of measurement may make assumptions or have a priori knowledge about one of the measurement points, but the metric definitions themselves are based on information collected at two measurement points.

4.4. Examples and Initial Comparisons

Note: This material originally presented in Slides 2 and 3 of [Morton06].

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

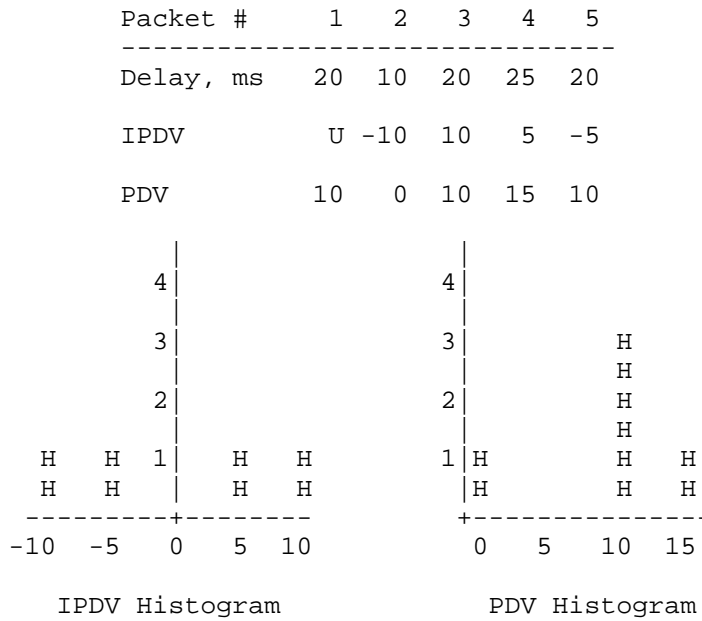


Figure 1: IPDV and PDV Comparison

The sample of packets contains three packets with "typical" delays of 20 ms, one packet with a low delay of 10 ms (the minimum of the sample) and one packet with 25 ms 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 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 singletons distinguish quick-delay variations (short-term, 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. The PDV distribution does not distinguish short-term variation from variation over the complete test interval. (Comment: PDV can be determined over any sub-intervals when the singletons are stored.)
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.

(Ideally, the measurement interval should be aligned with the session length of interest, and this influences determination of the reference delay, $D(\min)$.)

5. The PDV distribution 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 ± 50 ms 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 setup where a queue was temporarily congested by a competing packet burst. The additional queuing delay was 85 ms to 95 ms, much larger than the inter-packet interval. The first packet in the stream that follows the competing burst spends the longest time queued, 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 5 ms. Also, it is easy to determine (by looking at the PDV range) that a de-jitter buffer of >85 ms would have been sufficient to accommodate the delay variation. Again, the measurement interval is a key factor in the validity of such observations (it should have similar length to the session interval of interest).

The IPDV values in the congested queue example are very different: 85, -20, -20, -20, -20, -5 ms. 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).

Note that this example illustrates a case where the IPDV distribution is asymmetrical, because the delay variation range (85 ms) exceeds the inter-packet spacing (20 ms). We see that the IPDV values 85, -20, -20, -20, -20, -5 ms have zero mean, but the left side of the distribution is truncated at -20 ms.

Elsewhere in the article, the authors considered the range as a summary statistic for IPDV, and the 99.9th percentile 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 a working version of [RFC3393]. One of his main points was that a delay histogram is a useful approach to quantifying variation. Another point was that the time duration of evaluation is a critical aspect.

Carlo Demichelis then mailed his comparison paper [Demichelis] to the IPPM list, 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, and this is the statistic that determines the size of a de-jitter buffer. The IPDV range is minimal in response to the smooth variation in Example A (20 ms). However, IPDV responds to the faster variations in Example B (60 ms range from 40 to -20). Here the IPDV range is larger than the PDV range, and overestimates the buffer size requirements.

A heuristic method to estimate buffer size using IPDV is to sum the consecutive positive or zero values as an estimate of PDV range. However, this is more complicated to assess than the PDV range, and has strong dependence on the actual sequence of IPDV values (any negative IPDV value stops the summation, and again causes an underestimate).

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. The short-term variations may be useful to know in some other cases.

5.4. Y.1540 Appendix II

Appendix II of [Y.1540] describes a secondary terminology for delay variation. It compares IPDV, PDV (referred to as 2-point PDV), and 1-point packet delay variation (which assumes a periodic stream and assesses variation against an ideal arrival schedule constructed at a single measurement point). This early comparison discusses some of the same considerations raised in Section 6 below.

5.5. Clark's ITU-T SG 12 Contribution

Alan Clark's contribution to ITU-T Study Group 12 in January 2003 provided an analysis of the root causes of delay variation and investigated different techniques for measurement and modeling of "jitter" [COM12.D98]. Clark compared a metric closely related to IPDV, Mean Packet-to-Packet Delay Variation, MPPDV = $\text{mean}(\text{abs}(D(i) - D(i-1)))$ to the newly proposed Mean Absolute Packet Delay Variation (MAPDV2, see [G.1020]). One of the tasks for this study was to estimate the number of packet discards in a de-jitter buffer. Clark concluded that MPPDV did not track the ramp delay variation he associated access link congestion (similar to Figure 2, Example A above), but MAPDV2 did.

Clark also briefly looked at PDV (as described in the 2002 version of [Y.1541]). He concluded that if PDV was applied to a series of very short measurement intervals (e.g., 200 ms), it could be used to determine the fraction of intervals with high packet discard rates.

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 of packet loss is of great influence for the delay variation results, as displayed in the Figures 3 and 4 (L means Lost and U means Undefined). Figure 3 shows that in the extreme case of every other packet loss, the IPDV metric 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
PDV		0	U	2	U	1	U	0	U	1

Figure 3: Path Loss Every Other Packet

In case of a burst of packet loss, as displayed in Figure 4, both the IPDV and PDV metrics produce some results. Note that PDV still produces more values than IPDV.

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	-1	-1
PDV		0	1	U	U	U	U	U	2	1

Figure 4: Burst of Packet Loss

In conclusion, the PDV results are affected by the packet-loss ratio. The IPDV results are affected by both the packet-loss ratio and the packet-loss distribution. In the extreme case of loss of every other packet, 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 that 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 Time to Live (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 it 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 one IPDV singleton. For example, the delay sequence in the Figure below always produces IPDV=0 except in the one case where the value is 5 (U, 0, 0, 0, 5, 0, 0, 0, 0).

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

Delay, ms	4	4	4	4	9	9	9	9	9
IPDV	U	0	0	0	5	0	0	0	0
PDV	0	0	0	0	5	5	5	5	5

Figure 5: Lossless Path Change

However, if the change in delay is negative and larger than the inter-packet sending interval, then more than one IPDV singleton 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 a 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 to revise the 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. Figure 6, 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 6: Path Change with Loss

PDV will again produce a bi-modal 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.

IPDV does not help to detect path changes when accompanied by loss, and this is a disadvantage for those who rely solely on IPDV measurements.

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 (~1 ms) 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, ~0.01 ms for example).

Maintaining time accuracy <<1 ms has typically required access to dedicated time receivers at all measurement points. Global positioning system (GPS) receivers have often been installed to support measurements. The GPS installation conditions are fairly restrictive, and many prospective measurement efforts have found the deployment complexity and system maintenance too difficult.

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 misalignment 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 any two one-way-delay singletons will cancel as long as the error is constant. So, the demand for NTP or GPS synchronization comes primarily from one-way-delay measurement time-of-day accuracy requirements. Delay variation and measurement interval alignment are relatively less demanding.

Skew is a measure of the change in clock time over an interval with respect to 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. Thus, the additional accuracy demand of PDV can be expressed as a ratio of the measurement interval to the inter-packet spacing.

A practical example is a measurement between two hosts, one with a synchronized clock and the other with a free-running clock having 50 parts per million (ppm) long term accuracy.

- o If IPDV measurements are made on packets with a 1 second spacing, the maximum singleton error will be $1 \times 5 \times 10^{-5}$ seconds, or 0.05 ms.
- o If PDV measurements are made on the same packets over a 60 second measurement interval, then the delay variation due to the max free-running clock error will be $60 \times 5 \times 10^{-5}$ seconds, or 3 ms delay variation error from the first packet to the last.

Therefore, the additional accuracy required for equivalent PDV error under these conditions is a factor of 60 more than for IPDV. This is a rather extreme scenario, because time-of-day error of 1 second would accumulate in ~5.5 hours, potentially causing the measurement interval alignment issue described above.

If skew is present in a sample of one-way delays, its symptom is typically a nearly 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.

Values for IPDV may have non-zero mean over a sample when clock skew is present. This tends to complicate IPDV analysis when using the assumptions of a zero mean and a symmetric distribution.

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 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 (regardless of the form of delay variation considered).

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. Additional methods are considered in [IPPM-Spatial].

PDV has a clear advantage at this time, since there is no validated 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, where segments must be assumed to have independent delay distributions.

6.5. Reporting a Single Number (SLA)

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 single number summary statistic is a key component of SLAs, where a threshold on that number must be met x% of the time.

The simplicity of the PDV distribution lends itself to this summarization process (including use of the percentiles, median or mean). An SLA of the form "no more than x% of packets in a measurement interval shall have PDV \geq y ms, for no less than z% of time" is relatively straightforward to specify and implement. [Y.1541] introduced the notion of a pseudo-range when setting an objective for the 99.9th percentile of PDV. The conventional range (max-min) was avoided for several reasons, including stability of the maximum delay. The 99.9th percentile 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 will 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 (in Sections 5.2 and 6.2.1). Should the Inter-Quartile Range be used? Should the singletons beyond some threshold be counted (e.g., mean +/- 50 ms)? A strong rationale for one of these summary statistics has yet to emerge.

When summarizing IPDV, some prefer the simplicity of the single-sided distribution created by taking the absolute value of each singleton result, $\text{abs}(D(i)-D(i-1))$. This approach sacrifices the two-sided inter-arrival spread information in the distribution. It also makes the evaluation using percentiles more confusing, because a single late packet that exceeds the variation threshold will cause two pairs of singletons to fail the criteria (one positive, the other negative converted to positive). The single-sided PDV distribution is an advantage in this category.

6.6. Jitter in RTCP Reports

Section 6.4.1 of [RFC3550] gives the calculation of the "inter-arrival jitter" field for the RTP Control Protocol (RTCP) report, with a sample implementation in an Appendix.

The RTCP "interarrival 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] "interarrival 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

Network traffic load balancing is a process to divide packet traffic in order to provide a more even distribution over two or more equally viable paths. The paths chosen are based on the IGP cost metrics, while the delay depends on the path's physical layout. Usually, the balancing process is performed on a per-flow basis to avoid delay variation experienced when packets traverse different physical paths.

If the sample includes test packets with different characteristics such as IP addresses/ports, there could be multi-modal delay distributions present. The PDV form makes the identification of multiple modes possible. IPDV may also reveal that multiple paths are in use with a mixed-flow sample, but the different delay modes are not easily divided and analyzed separately.

Should the delay singletons using multiple addresses/ports be combined in the same sample? Should we characterize each mode separately? (This question also applies to the Path Change case.) It depends on the task to be addressed by the measurement.

For the task of de-jitter buffer sizing or assessing queue occupation, the modes should be characterized separately because flows will experience only one mode on a stable path. Use of a single flow description (address/port combination) in each sample simplifies this analysis. Multiple modes may be identified by collecting samples with different flow attributes, and characterization of multiple paths can proceed with comparison of the delay distributions from each sample.

For the task of capacity planning and routing optimization, characterizing the modes separately could offer an advantage. Network-wide capacity planning (as opposed to link capacity planning) takes as input the core traffic matrix, which corresponds to a matrix of traffic transferred from every source to every destination in the network. Applying the core traffic matrix along with the routing information (typically the link state database of a routing protocol) in a capacity planning tool offers the possibility to visualize the paths where the traffic flows and to optimize the routing based on the link utilization. In the case where equal cost multiple paths (ECMPs) are used, the traffic will be load balanced onto multiple paths. If each mode of the IP delay multi-modal distribution can be associated with a specific path, the delay performance offers an extra optimization parameter, i.e., the routing optimization based on the IP delay variation metric. As an example, the load balancing across ECMPs could be suppressed so that the Voice over IP (VoIP) calls would only be routed via the path with the lower IP delay

variation. Clearly, any modifications can result in new delay performance measurements, so there must be a verification step to ensure the desired outcome.

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.

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 (and FEC Design)

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. Figure 5 of [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 99.9th percentile of PDV (pseudo-range) are closely related to the buffer size needed to accommodate the observed network delay variation.

The PDV distribution directly addresses the FEC waiting time question. When the PDV distribution has a 99th percentile of 10 ms, then waiting 10 ms longer than the FEC protection interval will allow 99% of late packets to arrive and be used in the FEC block.

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 validated method to compose an IPDV metric.

7.1.4. Service-Level Specification: Reporting a Single Number

The one-sided PDV distribution can be constrained with a single statistic, such as an upper percentile, so it is preferred. The IPDV distribution is two-sided, usually has zero mean, and no universal summary statistic that relates to a physical quantity has emerged in years of experience.

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 and Storage Issues

When appreciable skew is present between measurement system clocks, IPDV has an advantage because PDV would require processing over the entire sample to remove the skew error. However, significant skew can invalidate IPDV analysis assumptions, such as the zero-mean and symmetric-distribution characteristics. Small skew may well be within the error tolerance, and both PDV and IPDV results will be usable. There may be a portion of the skew, measurement interval, and required accuracy 3-D space where IPDV has an advantage, depending on the specific measurement specifications.

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). However, if one intends to use IPDV to indicate path changes, it cannot do this when the change is accompanied by loss.

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.

When path changes are frequent and cause packet loss, delay variation is probably less important than the loss episodes and attention should be turned to the loss metric instead.

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, yet PDV would successfully produce a sample based on the arriving packets.

7.2.4. Load Balancing

PDV distributions offer the most straightforward way to identify that a sample of packets have traversed multiple paths. The tasks of de-jitter buffer sizing or assessing queue occupation with PDV should be use a sample with a single flow because flows will experience only one mode on a stable path, and it simplifies the analysis.

7.3. Summary

Comparison Area	PDV = $D(i) - D(\min)$	IPDV = $D(i) - D(i-1)$
Challenging Circumstances	Less sensitive to packet loss, and simplifies analysis when load balancing or multiple paths are present	Preferred when path changes are frequent or when measurement clocks exhibit some skew
Spatial Composition of DV metric	All validated methods use this form	Has sensitivity to sequence and spacing changes, which tends to break the requirement for independent distributions between path segments
Determine De-Jitter Buffer Size Required	"Pseudo-range" reveals this property by anchoring the distribution at the minimum delay	No reliable relationship, but some heuristics
Estimate of Queuing Time and Variation	Distribution has one-to-one relationship on a stable path, especially when sample min = true min	No reliable relationship
Specification Simplicity: Single Number SLA	One constraint needed for single-sided distribution, and easily related to quantities above	Distribution is two-sided, usually has zero mean, and no universal summary statistic that relates to a physical quantity

Summary of Comparisons

8. Measurement Considerations

This section discusses the practical aspects of delay variation measurement, with special attention to the two formulations compared in this memo.

8.1. Measurement Stream Characteristics

As stated in Section 1.2, there is a strong dependency between the active measurement stream characteristics and the results. The IPPM literature includes two primary methods for collecting samples: Poisson sampling described in [RFC2330], and Periodic sampling in [RFC3432]. The Poisson method was intended to collect an unbiased sample of performance, while the Periodic method addresses a "known bias of interest". Periodic streams are required to have random start times and limited stream duration, in order to avoid unwanted synchronization with some other periodic process, or cause congestion-aware senders to synchronize with the stream and produce atypical results. The random start time should be different for each new stream.

It is worth noting that [RFC3393] was developed in parallel with [RFC3432]. As a result, all the stream metrics defined in [RFC3393] specify the Poisson sampling method.

Periodic sampling is frequently used in measurements of delay variation. Several factors foster this choice:

1. Many application streams that are sensitive to delay variation also exhibit periodicity, and so exemplify the bias of interest. If the application has a constant packet spacing, this constant spacing can be the inter-packet gap for the test stream. VoIP streams often use 20 ms spacing, so this is an obvious choice for an Active stream. This applies to both IPDV and PDV forms.
2. The spacing between packets in the stream will influence whether the stream experiences short-range dependency, or only long-range dependency, as investigated in [Li.Mills]. The packet spacing also influences the IPDV distribution and the stream's sensitivity to reordering. For example, with a 20 ms spacing the IPDV distribution cannot go below -20 ms without packet reordering.
3. The measurement process may make several simplifying assumptions when the send spacing and send rate are constant. For example, the inter-arrival times at the destination can be compared with an ideal sending schedule, and allowing a one-point measurement

of delay variation (described in [Y.1540]) that approximates the IPDV form. Simplified methods that approximate PDV are possible as well (some are discussed in Appendix II of [Y.1541]).

4. Analysis of truncated, or non-symmetrical IPDV distributions is simplified. Delay variations in excess of the periodic sending interval can cause multiple singleton values at the negative limit of the packet spacing (see Section 5.2 and [Cia03]). Only packet reordering can cause the negative spacing limit to be exceeded.

Despite the emphasis on inter-packet delay differences with IPDV, both Poisson [Demichelis] and Periodic [Li.Mills] streams have been used, and these references illustrate the different analyses that are possible.

The advantages of using a Poisson distribution are discussed in [RFC2330]. The main properties are to avoid predicting the sample times, avoid synchronization with periodic events that are present in networks, and avoid inducing synchronization with congestion-aware senders. When a Poisson stream is used with IPDV, the distribution will reflect inter-packet delay variation on many different time scales (or packet spacings). The unbiased Poisson sampling brings a new layer of complexity in the analysis of IPDV distributions.

8.2. Measurement Devices

One key aspect of measurement devices is their ability to store singletons (or individual measurements). This feature usually is closely related to local calculation capabilities. For example, an embedded measurement device with limited storage will like provide only a few statistics on the delay variation distribution, while dedicated measurement systems store all the singletons and allow detailed analysis (later calculation of either form of delay variation is possible with the original singletons).

Therefore, systems with limited storage must choose their metrics and summary statistics in advance. If both IPDV and PDV statistics are desired, the supporting information must be collected as packets arrive. For example, the PDV range and high percentiles can be determined later if the minimum and several of the largest delays are stored while the measurement is in-progress.

8.3. Units of Measurement

Both IPDV and PDV can be summarized as a range in milliseconds.

With IPDV, it is interesting to report on a positive percentile, and an inter-quantile range is appropriate to reflect both positive and negative tails (e.g., 5% to 95%). If the IPDV distribution is symmetric around a mean of zero, then it is sufficient to report on the positive side of the distribution.

With PDV, it is sufficient to specify the upper percentile (e.g., 99.9%).

8.4. Test Duration

At several points in this memo, we have recommended use of test intervals on the order of minutes. In their paper examining the stability of Internet path properties [Zhang.Duff], Zhang et al. concluded that consistency was present on the order of minutes for the performance metrics considered (loss, delay, and throughput) for the paths they measured.

The topic of temporal aggregation of performance measured in small intervals to estimate some larger interval is described in the Metric Composition Framework [IPPM-Framework].

The primary recommendation here is to test using durations that are similar in length to the session time of interest. This applies to both IPDV and PDV, but is possibly more relevant for PDV since the duration determines how often the D_{\min} will be determined, and the size of the associated sample.

8.5. Clock Sync Options

As with one-way-delay measurements, local clock synchronization is an important matter for delay variation measurements.

There are several options available:

1. Global Positioning System receivers
2. In some parts of the world, Cellular Code Division Multiple Access (CDMA) systems distribute timing signals that are derived from GPS and traceable to UTC.
3. Network Time Protocol [RFC1305] is a convenient choice in many cases, but usually offers lower accuracy than the options above.

When clock synchronization is inconvenient or subject to appreciable errors, then round-trip measurements may give a cumulative indication of the delay variation present on both directions of the path. However, delay distributions are rarely symmetrical, so it is difficult to infer much about the one-way-delay variation from round-trip measurements. Also, measurements on asymmetrical paths add complications for the one-way-delay metric.

8.6. Distinguishing Long Delay from Loss

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 [IPPM-Reporting].

In essence, [IPPM-Reporting] suggests to use a long waiting time to serve network characterization and revise results for specific application delay thresholds as needed.

8.7. Accounting for Packet Reordering

Packet reordering, defined in [RFC4737], is essentially an extreme form of delay variation where the packet stream arrival order differs from the sending order.

PDV results are not sensitive to packet arrival order, and are not affected by reordering other than to reflect the more extreme variation.

IPDV results will change if reordering is present because they are sensitive to the sequence of delays of arriving packets. The main example of this sensitivity is in the truncation of the negative tail of the distribution.

- o When there is no reordering, the negative tail is limited by the sending time spacing between packets.
- o If reordering occurs (and the reordered packets are not discarded), the negative tail can take on any value (in principal).

In general, measurement systems should have the capability to detect when sequence has changed. If IPDV measurements are made without regard to packet arrival order, the IPDV will be under-reported when reordering occurs.

8.8. Results Representation and Reporting

All of the references that discuss or define delay variation suggest ways to represent or report the results, and interested readers should review the various possibilities.

For example, [IPPM-Reporting] suggests reporting a pseudo-range of delay variation based on calculating the difference between a high percentile of delay and the minimum delay. The 99.9th percentile minus the minimum will give a value that can be compared with objectives in [Y.1541].

9. Security Considerations

The security considerations that apply to any active measurement of live networks are relevant here as well. See the "Security Considerations" sections in [RFC2330], [RFC2679], [RFC3393], [RFC3432], and [RFC4656].

Security considerations do not contribute to the selection of PDV or IPDV forms of delay variation, because measurements using these metrics involve exactly the same security issues.

10. Acknowledgments

The authors would like to thank Phil Chimento for his suggestion to employ the convention of conditional distributions of delay to deal with packet loss, and his encouragement to "write the memo" after hearing "the talk" on this topic at IETF 65. We also acknowledge constructive comments from Alan Clark, Loki Jorgenson, Carsten Schmoll, and Robert Holley.

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

Practitioners have raised several questions that this section intends to answer:

- How is this D_min calculated? Is it DV(99%) as mentioned in [Krzanowski]?
- 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_(99.9th percentile) 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. There is also a simpler approach to solving this problem: use samples collected over short evaluation intervals (on the order of minutes). Intervals with path changes may be more interesting from the loss or one-way-delay perspective (possibly failing to meet one or more SLAs), and it may not be necessary to conduct delay variation analysis. Short evaluation intervals are preferred for measurements that serve as a basis for troubleshooting, since the results are available to report soon after collection.

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 99.9th percentile 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.

12. References

12.1. 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.

- [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.

12.2. Informative References

- [COM12.D98] Clark, A., "Analysis, measurement and modelling of Jitter", ITU-T Delayed Contribution COM 12 - D98, January 2003.
- [Casner] Casner, S., Alaettinoglu, C., and C. Kuan, "A Fine-Grained View of High Performance Networking", NANOG 22, May 20-22, 2001, <<http://www.nanog.org/mtg-0105/agenda.html>>.
- [Cia03] Ciavattone, L., Morton, A., and G. Ramachandran, "Standardized Active Measurements on a Tier 1 IP Backbone", IEEE Communications Magazine, p. 90-97, June 2003.
- [Demichelis] Demichelis, C., "Packet Delay Variation Comparison between ITU-T and IETF Draft Definitions", November 2000, <<http://www.advanced.org/ippm/archive.3/att-0075/01-pap02.doc>>.
- [G.1020] ITU-T, "Performance parameter definitions for the quality of speech and other voiceband applications utilizing IP networks", ITU-T Recommendation G.1020, 2006.
- [G.1050] ITU-T, "Network model for evaluating multimedia transmission performance over Internet Protocol", ITU-T Recommendation G.1050, November 2005.
- [I.356] ITU-T, "B-ISDN ATM Layer Cell Transfer Performance", ITU-T Recommendation I.356, March 2000.
- [IPPM-Framework] Morton, A., "Framework for Metric Composition", Work in Progress, October 2008.

- [IPPM-Reporting] Morton, A., Ramachandran, G., and G. Maguluri, "Reporting Metrics: Different Points of View", Work in Progress, January 2009.
- [IPPM-Spatial] Morton, A. and E. Stephan, "Spatial Composition of Metrics", Work in Progress, July 2008.
- [Krzanowski] Presentation at IPPM, IETF-64, "Jitter Definitions: What is What?", November 2005.
- [Li.Mills] Li, Q. and D. Mills, "The Implications of Short-Range Dependency on Delay Variation Measurement", Second IEEE Symposium on Network Computing and Applications, 2003.
- [Morton06] Morton, A., "A Brief Jitter Metrics Comparison, and not the last word, by any means...", slide presentation at IETF 65, IPPM Session, March 2006.
- [RFC1305] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation", RFC 1305, March 1992.
- [RFC3357] Koodli, R. and R. Ravikanth, "One-way Loss Pattern Sample Metrics", RFC 3357, August 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.
- [Y.1540] ITU-T, "Internet protocol data communication service - IP packet transfer and availability performance parameters", ITU-T Recommendation Y.1540, November 2007.
- [Y.1541] ITU-T, "Network Performance Objectives for IP-Based Services", ITU-T Recommendation Y.1541, February 2006.
- [Zhang.Duff] Zhang, Y., Duffield, N., Paxson, V., and S. Shenker, "On the Constancy of Internet Path Properties", Proceedings of ACM SIGCOMM Internet Measurement Workshop, November 2001.

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
EMail: bclaise@cisco.com

