

Network Working Group	A. Morton	
Internet-Draft	AT&T Labs	
Intended status: Standards Track	E. Stephan	
Expires: January 10, 2008	France Telecom Division R&D	
	July 09, 2007	

[TOC](#)

Spatial Composition of Metrics

draft-ietf-ippm-spatial-composition-04

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/lid-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 10, 2008.

Abstract

This memo utilizes IPPM metrics that are applicable to both complete paths and sub-paths, and defines relationships to compose a complete path metric from the sub-path metrics with some accuracy w.r.t. the actual metrics. This is called Spatial Composition in RFC 2330. The memo refers to the Framework for Metric Composition, and provides background and motivation for combining metrics to derive others. The descriptions of several composed metrics and statistics follow.

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 \(Bradner, S.,](#)

["Key words for use in RFCs to Indicate Requirement Levels,"](#)
[March 1997.](#)) [RFC2119].

In this memo, the characters "<=" should be read as "less than or equal to" and ">=" as "greater than or equal to".

Table of Contents

- [1.](#) Contributors
- [2.](#) Introduction
 - [2.1.](#) Motivation
- [3.](#) Scope and Application
 - [3.1.](#) Scope of work
 - [3.2.](#) Application
 - [3.3.](#) Incomplete Information
- [4.](#) Common Specifications for Composed Metrics
 - [4.1.](#) Name: Type-P
 - [4.1.1.](#) Metric Parameters
 - [4.1.2.](#) Definition and Metric Units
 - [4.1.3.](#) Discussion and other details
 - [4.1.4.](#) Statistic:
 - [4.1.5.](#) Composition Function
 - [4.1.6.](#) Statement of Conjecture
 - [4.1.7.](#) Justification of the Composition Function
 - [4.1.8.](#) Sources of Deviation from the Ground Truth
 - [4.1.9.](#) Specific cases where the conjecture might fail
 - [4.1.10.](#) Application of Measurement Methodology
- [5.](#) One-way Delay Composed Metrics and Statistics
 - [5.1.](#) Name: Type-P-Finite-One-way-Delay-Poisson/Periodic-Stream
 - [5.1.1.](#) Metric Parameters
 - [5.1.2.](#) Definition and Metric Units
 - [5.1.3.](#) Discussion and other details
 - [5.2.](#) Name: Type-P-Finite-Composite-One-way-Delay-Mean
 - [5.2.1.](#) Metric Parameters
 - [5.2.2.](#) Definition and Metric Units of the Mean Statistic
 - [5.2.3.](#) Discussion and other details
 - [5.2.4.](#) Composition Function: Sum of Means
 - [5.2.5.](#) Statement of Conjecture
 - [5.2.6.](#) Justification of the Composition Function
 - [5.2.7.](#) Sources of Deviation from the Ground Truth
 - [5.2.8.](#) Specific cases where the conjecture might fail
 - [5.2.9.](#) Application of Measurement Methodology
 - [5.3.](#) Name: Type-P-Finite-Composite-One-way-Delay-Minimum
 - [5.3.1.](#) Metric Parameters
 - [5.3.2.](#) Definition and Metric Units of the Mean Statistic
 - [5.3.3.](#) Discussion and other details
 - [5.3.4.](#) Composition Function: Sum of Means
 - [5.3.5.](#) Statement of Conjecture

5.3.6.	Justification of the Composition Function
5.3.7.	Sources of Deviation from the Ground Truth
5.3.8.	Specific cases where the conjecture might fail
5.3.9.	Application of Measurement Methodology
6.	Loss Metrics and Statistics
6.1.	Type-P-Composite-One-way-Packet-Loss-Empirical-Probability
6.1.1.	Metric Parameters:
6.1.2.	Definition and Metric Units
6.1.3.	Discussion and other details
6.1.4.	Statistic: Type-P-One-way-Packet-Loss-Empirical-Probability
6.1.5.	Composition Function: Composition of Empirical Probabilities
6.1.6.	Statement of Conjecture
6.1.7.	Justification of the Composition Function
6.1.8.	Sources of Deviation from the Ground Truth
6.1.9.	Specific cases where the conjecture might fail
6.1.10.	Application of Measurement Methodology
7.	Delay Variation Metrics and Statistics
7.1.	Name: Type-P-One-way-pdv-refmin-Poisson/Periodic-Stream
7.1.1.	Metric Parameters:
7.1.2.	Definition and Metric Units
7.1.3.	Discussion and other details
7.1.4.	Statistics: Mean, Variance, Skewness, Quantile
7.1.5.	Composition Functions:
7.1.6.	Statement of Conjecture
7.1.7.	Justification of the Composition Function
7.1.8.	Sources of Deviation from the Ground Truth
7.1.9.	Specific cases where the conjecture might fail
7.1.10.	Application of Measurement Methodology
8.	Security Considerations
8.1.	Denial of Service Attacks
8.2.	User Data Confidentiality
8.3.	Interference with the metrics
9.	IANA Considerations
10.	Issues (Open and Closed)
11.	Acknowledgements
12.	References
12.1.	Normative References
12.2.	Informative References
§	Authors' Addresses
§	Intellectual Property and Copyright Statements

1. Contributors

Thus far, the following people have contributed useful ideas, suggestions, or the text of sections that have been incorporated into this memo:

- Phil Chimento <vze275m9@verizon.net>
- Reza Fardid <RFardid@Covad.COM>
- Roman Krzanowski <roman.krzanowski@verizon.com>
- Maurizio Molina <maurizio.molina@dante.org.uk>
- Al Morton <acmorton@att.com>
- Emile Stephan <emile.stephan@orange-ftgroup.com>
- Lei Liang <L.Liang@surrey.ac.uk>
- Dave Hoeftlin <dhoeflin@att.com>

2. Introduction

[TOC](#)

The IPPM framework [\[RFC2330\] \(Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics," May 1998.\)](#) describes two forms of metric composition, spatial and temporal. The new composition framework [\[I-D.ietf-ippm-framework-compagg\] \(Morton, A., "Framework for Metric Composition," December 2009.\)](#) expands and further qualifies these original forms into three categories. This memo describes Spatial Composition, one of the categories of metrics under the umbrella of the composition framework.

Spatial composition encompasses the definition of performance metrics that are applicable to a complete path, based on metrics collected on various sub-paths.

The main purpose of this memo is to define the deterministic functions that yield the complete path metrics using metrics of the sub-paths. The effectiveness of such metrics is dependent on their usefulness in analysis and applicability with practical measurement methods.

The relationships may involve conjecture, and [\[RFC2330\] \(Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics," May 1998.\)](#) lists four points that the metric definitions should include:

- *the specific conjecture applied to the metric,
- *a justification of the practical utility of the composition in terms of making accurate measurements of the metric on the path,
- *a justification of the usefulness of the composition in terms of making analysis of the path using A-frame concepts more effective, and
- *an analysis of how the conjecture could be incorrect.

Also, [\[RFC2330\] \(Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics," May 1998.\)](#) gives an example where a conjecture that the delay of a path is very nearly the sum of the delays of the exchanges and clouds of the corresponding path digest. This example is particularly relevant to those who wish to assess the performance of an Inter-domain path without direct measurement, and the performance estimate of the complete path is related to the measured results for various sub-paths instead. Approximate functions between the sub-path and complete path metrics are useful, with knowledge of the circumstances where the relationships are/are not applicable. For example, we would not expect that delay singletons from each sub-path would sum to produce an accurate estimate of a delay singleton for the complete path (unless all the delays were essentially constant - very unlikely). However, other delay statistics (based on a reasonable sample size) may have a sufficiently large set of circumstances where they are applicable.

2.1. Motivation

[TOC](#)

One-way metrics defined in other IPPM RFCs all assume that the measurement can be practically carried out between the source and the destination of the interest. Sometimes there are reasons that the measurement can not be executed from the source to the destination. For instance, the measurement path may cross several independent domains that have conflicting policies, measurement tools and methods, and measurement time assignment. The solution then may be the composition of several sub-path measurements. This means each domain performs the One-way measurement on a sub path between two nodes that are involved in the complete path following its own policy, using its own measurement tools and methods, and using its own measurement timing. Under the appropriate conditions, one can combine the sub-path One-way metric results to estimate the complete path One-way measurement metric with some degree of accuracy.

3. Scope and Application

[TOC](#)

3.1. Scope of work

[TOC](#)

For the primary IPPM metrics of Loss, Delay, and Delay Variation, this memo gives a set of metrics for that can be composed from the same or

similar sub-path metrics. This means that the composition function may utilize:

- *the same metric for each sub-path;
- *multiple metrics for each sub-path (possibly one that is the same as the complete path metric);
- *a single sub-path metrics that is different from the complete path metric;
- *different measurement techniques like active and passive (recognizing that PSAMP WG will define capabilities to sample packets to support measurement).

3.2. Application

[TOC](#)

The new composition framework [\[I-D.ietf-ippm-framework-compagg\] \(Morton, A., "Framework for Metric Composition," December 2009.\)](#) requires the specification of the applicable circumstances for each metric. In particular, each section addresses whether the metric:

- Requires the same test packets to traverse all sub-paths, or may use similar packets sent and collected separately in each sub-path.
- Requires homogeneity of measurement methodologies, or can allow a degree of flexibility (e.g., active or passive methods produce the "same" metric). Also, the applicable sending streams will be specified, such as Poisson, Periodic, or both.
- Needs information or access that will only be available within an operator's domain, or is applicable to Inter-domain composition.
- Requires synchronized measurement time intervals in all sub-paths, or largely overlapping, or no timing requirements.
- Requires assumption of sub-path independence w.r.t. the metric being defined/composed, or other assumptions.
- Has known sources of inaccuracy/error, and identifies the sources.

3.3. Incomplete Information

[TOC](#)

In practice, when measurements cannot be initiated on a sub-path (and perhaps the measurement system gives up during the test interval), then there will not be a value for the sub-path reported, and the entire test result SHOULD be recorded as "undefined". This case should be distinguished from the case where the measurement system continued to send packets throughout the test interval, but all were declared lost.

When a composed metric requires measurements from sub paths A, B, and C, and one or more of the sub-path results are undefined, then the composed metric SHOULD also be recorded as undefined.

4. Common Specifications for Composed Metrics

[TOC](#)

To reduce the redundant information presented in the detailed metrics sections that follow, this section presents the specifications that are common to two or more metrics. The section is organized using the same subsections as the individual metrics, to simplify comparisons.

4.1. Name: Type-P

[TOC](#)

All metrics use the Type-P convention as described in [\[RFC2330\]](#) (Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics," May 1998.). The rest of the name is unique to each metric.

4.1.1. Metric Parameters

[TOC](#)

*Src, the IP address of a host

*Dst, the IP address of a host

*T, a time (start of test interval)

*Tf, a time (end of test interval)

*lambda, a rate in reciprocal seconds (for Poisson Streams)

*incT, the nominal duration of inter-packet interval, first bit to first bit (for Periodic Streams)

*T0, a time that MUST be selected at random from the interval [T, T+dT] to start generating packets and taking measurements (for Periodic Streams)

*TstampSrc, the wire time of the packet as measured at MP(Src)

*TstampDst, the wire time of the packet as measured at MP(Dst), assigned to packets that arrive within a "reasonable" time.

*Tmax, a maximum waiting time for packets at the destination, set sufficiently long to disambiguate packets with long delays from packets that are discarded (lost), thus the distribution of delay is not truncated.

*M, the total number of packets sent between T0 and Tf

*N, the total number of packets received at Dst (sent between T0 and Tf)

*S, the number of sub-paths involved in the complete Src-Dst path

4.1.2. Definition and Metric Units

[TOC](#)

This section is unique for every metric.

4.1.3. Discussion and other details

[TOC](#)

This section is unique for every metric.

4.1.4. Statistic:

[TOC](#)

This section is unique for every metric.

4.1.5. Composition Function

[TOC](#)

This section is unique for every metric.

4.1.6. Statement of Conjecture

[TOC](#)

This section is unique for each metric.

[TOC](#)

4.1.7. Justification of the Composition Function

It is sometimes impractical to conduct active measurements between every Src-Dst pair. Since the full mesh of N measurement points grows as $N \times N$, the scope of measurement may be limited by testing resources. There may be varying limitations on active testing in different parts of the network. For example, it may not be possible to collect the desired sample size in each test interval when access link speed is limited, because of the potential for measurement traffic to degrade the user traffic performance. The conditions on a low-speed access link may be understood well-enough to permit use of a small sample size/rate, while a larger sample size/rate may be used on other sub-paths. Also, since measurement operations have a real monetary cost, there is value in re-using measurements where they are applicable, rather than launching new measurements for every possible source-destination pair.

4.1.8. Sources of Deviation from the Ground Truth

[TOC](#)

The measurement packets, each having source and destination addresses intended for collection at edges of the sub-path, may take a different specific path through the network equipment and parallel links when compared to packets with the source and destination addresses of the complete path. Therefore, the composition of sub-path measurements may differ from the performance experienced by packets on the complete path. Multiple measurements employing sufficient sub-path address pairs might produce bounds on the extent of this error.

Related to the case of an alternate path described above is the case where elements in the measured path are unique to measurement system connectivity. For example, a measurement system may use a dedicated link to a LAN switch, and packets on the complete path do not traverse that link. The performance of such a dedicated link would be measured continuously, and its contribution to the sub-path metrics SHOULD be minimized as a source of error.

others???

4.1.9. Specific cases where the conjecture might fail

[TOC](#)

This section is unique for each metric.

[TOC](#)

4.1.10. Application of Measurement Methodology

The methodology:

SHOULD use similar packets sent and collected separately in each sub-path.

Allows a degree of flexibility (e.g., active or passive methods can produce the "same" metric, but timing and correlation of passive measurements is much more challenging).

Poisson and/or Periodic streams are RECOMMENDED.

Applies to both Inter-domain and Intra-domain composition.

SHOULD have synchronized measurement time intervals in all sub-paths, but largely overlapping intervals MAY suffice.

REQUIRES assumption of sub-path independence w.r.t. the metric being defined/composed.

5. One-way Delay Composed Metrics and Statistics

[TOC](#)

5.1. Name: Type-P-Finite-One-way-Delay-Poisson/Periodic-Stream

[TOC](#)

This metric is a necessary element of Delay Composition metrics, and its definition does not formally exist elsewhere in IPPM literature.

5.1.1. Metric Parameters

[TOC](#)

See the common parameters section above.

5.1.2. Definition and Metric Units

[TOC](#)

Using the parameters above, we obtain the value of Type-P-One-way-Delay singleton as per [\[RFC2679\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Delay Metric for IPPM," September 1999.\)](#).

For each packet [i] that has a finite One-way Delay (in other words, excluding packets which have undefined one-way delay):

Type-P-Finite-One-way-Delay-Poisson/Periodic-Stream[i] =

FiniteDelay[i] = TstampDst - TstampSrc

The units of measure for this metric are time in seconds, expressed in sufficiently low resolution to convey meaningful quantitative

information. For example, resolution of microseconds is usually sufficient.

5.1.3. Discussion and other details

[TOC](#)

The “Type-P-Finite-One-way-Delay” metric permits calculation of the sample mean statistic. This resolves the problem of including lost packets in the sample (whose delay is undefined), and the issue with the informal assignment of infinite delay to lost packets (practical systems can only assign some very large value).

The Finite-One-way-Delay approach handles the problem of lost packets by reducing the event space. We consider conditional statistics, and estimate the mean one-way delay conditioned on the event that all packets in the sample arrive at the destination (within the specified waiting time, T_{max}). This offers a way to make some valid statements about one-way delay, and at the same time avoiding events with undefined outcomes. This approach is derived from the treatment of lost packets in [\[RFC3393\] \(Demichelis, C. and P. Chimento, “IP Packet Delay Variation Metric for IP Performance Metrics \(IPPM\),” November 2002.\)](#), and is similar to [\[Y.1540\] \(ITU-T Recommendation Y.1540, “Internet protocol data communication service - IP packet transfer and availability performance parameters,” December 2002.\)](#) .

5.2. Name: Type-P-Finite-Composite-One-way-Delay-Mean

[TOC](#)

This section describes a statistic based on the Type-P-Finite-One-way-Delay-Poisson/Periodic-Stream metric.

5.2.1. Metric Parameters

[TOC](#)

See the common parameters section above.

5.2.2. Definition and Metric Units of the Mean Statistic

[TOC](#)

We define

Type-P-Finite-One-way-Delay-Mean =

$$\text{MeanDelay} = \frac{1}{N} * \sum_{i=1}^N (\text{FiniteDelay} [i])$$

where all packets $i= 1$ through N have finite singleton delays.
The units of measure for this metric are time in seconds, expressed in sufficiently low resolution to convey meaningful quantitative information. For example, resolution of microseconds is usually sufficient.

5.2.3. Discussion and other details

[TOC](#)

The Type-P-Finite-One-way-Delay-Mean metric requires the conditional delay distribution described in section 5.1.

5.2.4. Composition Function: Sum of Means

[TOC](#)

The Type-P-Finite-Composite-One-way-Delay-Mean, or `CompMeanDelay`, for the complete Source to Destination path can be calculated from sum of the Mean Delays of all its S constituent sub-paths.

Then the

Type-P-Finite-Composite-One-way-Delay-Mean =

$$\text{CompMeanDelay} = \sum_{i=1}^S (\text{MeanDelay} [i])$$

5.2.5. Statement of Conjecture

[TOC](#)

The mean of a sufficiently large stream of packets measured on each sub-path during the interval $[T, T_f]$ will be representative of the true mean of the delay distribution (and the distributions themselves are

sufficiently independent), such that the means may be added to produce an estimate of the complete path mean delay.

5.2.6. Justification of the Composition Function

[TOC](#)

See the common section.

5.2.7. Sources of Deviation from the Ground Truth

[TOC](#)

See the common section.

5.2.8. Specific cases where the conjecture might fail

[TOC](#)

If any of the sub-path distributions are bimodal, then the measured means may not be stable, and in this case the mean will not be a particularly useful statistic when describing the delay distribution of the complete path.

The mean may not be sufficiently robust statistic to produce a reliable estimate, or to be useful even if it can be measured.
others...

5.2.9. Application of Measurement Methodology

[TOC](#)

The requirements of the common section apply here as well.

5.3. Name: Type-P-Finite-Composite-One-way-Delay-Minimum

[TOC](#)

This section describes is a statistic based on the Type-P-Finite-One-way-Delay-Poisson/Periodic-Stream metric, and the composed metric based on that statistic.

5.3.1. Metric Parameters

[TOC](#)

See the common parameters section above.

5.3.2. Definition and Metric Units of the Mean Statistic

[TOC](#)

We define

Type-P-Finite-One-way-Delay-Minimum =

$$= \text{MinDelay} = (\text{FiniteDelay } [j])$$

such that for some index, j , where $1 \leq j \leq N$
 $\text{FiniteDelay}[j] \leq \text{FiniteDelay}[i]$ for all i

where all packets $i = 1$ through N have finite singleton delays.
The units of measure for this metric are time in seconds, expressed in sufficiently low resolution to convey meaningful quantitative information. For example, resolution of microseconds is usually sufficient.

5.3.3. Discussion and other details

[TOC](#)

The Type-P-Finite-One-way-Delay-Minimum metric requires the conditional delay distribution described in section 5.1.3.

5.3.4. Composition Function: Sum of Means

[TOC](#)

The Type-P-Finite-Composite-One-way-Delay-Minimum, or CompMinDelay , for the complete Source to Destination path can be calculated from sum of the Minimum Delays of all its S constituent sub-paths.

Then the

Type-P-Finite-Composite-One-way-Delay-Minimum =

$$\text{CompMinDelay} = \sum_{i=1}^S (\text{MinDelay } [i])$$

5.3.5. Statement of Conjecture

[TOC](#)

The minimum of a sufficiently large stream of packets measured on each sub-path during the interval $[T, T_f]$ will be representative of the true minimum of the delay distribution (and the distributions themselves are sufficiently independent), such that the minima may be added to produce an estimate of the complete path minimum delay.

5.3.6. Justification of the Composition Function

[TOC](#)

See the common section.

5.3.7. Sources of Deviation from the Ground Truth

[TOC](#)

See the common section.

5.3.8. Specific cases where the conjecture might fail

[TOC](#)

If the routing on any of the sub-paths is not stable, then the measured minimum may not be stable. In this case the composite minimum would tend to produce an estimate for the complete path that may be too low for the current path.
others???

5.3.9. Application of Measurement Methodology

[TOC](#)

The requirements of the common section apply here as well.

6. Loss Metrics and Statistics

[TOC](#)

6.1. Type-P-Composite-One-way-Packet-Loss-Empirical-Probability

[TOC](#)

6.1.1. Metric Parameters:

[TOC](#)

Same as section 4.1.1.

6.1.2. Definition and Metric Units

[TOC](#)

Using the parameters above, we obtain the value of Type-P-One-way-Packet-Loss singleton and stream as per [\[RFC2680\] \(Almes, G., Kalidindi, S., and M. Zekauskas, "A One-way Packet Loss Metric for IPPM," September 1999.\)](#).

We obtain a sequence of pairs with elements as follows:

*TstampSrc, as above

*L, either zero or one, where L=1 indicates loss and L=0 indicates arrival at the destination within TstampSrc + Tmax.

6.1.3. Discussion and other details

[TOC](#)

6.1.4. Statistic: Type-P-One-way-Packet-Loss-Empirical-Probability

[TOC](#)

Given the stream parameter M, the number of packets sent, we can define the Empirical Probability of Loss Statistic (Ep), consistent with Average Loss in [RFC2680], as follows:

Type-P-One-way-Packet-Loss-Empirical-Probability =

$$Ep = \frac{1}{M} * \frac{\sum_{i=1}^M (L[i])}{M}$$

where all packets i= 1 through M have a value for L.

6.1.5. Composition Function: Composition of Empirical Probabilities

[TOC](#)

The Type-P-One-way-Composite-Packet-Loss-Empirical-Probability, or CompEp for the complete Source to Destination path can be calculated by combining Ep of all its constituent sub-paths (Ep1, Ep2, Ep3, ... Epn) as

Type-P-Composite-One-way-Packet-Loss-Empirical-Probability =

$$\text{CompEp} = 1 - \{(1 - \text{Ep1}) \times (1 - \text{Ep2}) \times (1 - \text{Ep3}) \times \dots \times (1 - \text{Epn})\}$$

If any Epn is undefined in a particular measurement interval, possibly because a measurement system failed to report a value, then any CompEp that uses sub-path n for that measurement interval is undefined.

6.1.6. Statement of Conjecture

[TOC](#)

The empirical probability of loss calculated on a sufficiently large stream of packets measured on each sub-path during the interval [T, Tf] will be representative of the true loss probability (and the probabilities themselves are sufficiently independent), such that the sub-path probabilities may be combined to produce an estimate of the complete path loss probability.

6.1.7. Justification of the Composition Function

[TOC](#)

See the common section.

6.1.8. Sources of Deviation from the Ground Truth

[TOC](#)

See the common section.

[TOC](#)

6.1.9. Specific cases where the conjecture might fail

A concern for loss measurements combined in this way is that root causes may be correlated to some degree.

For example, if the links of different networks follow the same physical route, then a single catastrophic event like a fire in a tunnel could cause an outage or congestion on remaining paths in multiple networks. Here it is important to ensure that measurements before the event and after the event are not combined to estimate the composite performance.

Or, when traffic volumes rise due to the rapid spread of an email-born worm, loss due to queue overflow in one network may help another network to carry its traffic without loss.

others...

6.1.10. Application of Measurement Methodology

[TOC](#)

See the common section.

7. Delay Variation Metrics and Statistics

[TOC](#)

7.1. Name: Type-P-One-way-pdv-refmin-Poisson/Periodic-Stream

[TOC](#)

This packet delay variation (PDV) metric is a necessary element of Composed Delay Variation metrics, and its definition does not formally exist elsewhere in IPPM literature.

7.1.1. Metric Parameters:

[TOC](#)

In addition to the parameters of section 4.1.1:

*TstampSrc[i], the wire time of packet[i] as measured at MP(Src)
(measurement point at the source)

*TstampDst[i], the wire time of packet[i] as measured at MP(Dst),
assigned to packets that arrive within a "reasonable" time.

*B, a packet length in bits

*F, a selection function unambiguously defining the packets from the stream that are selected for the packet-pair computation of this metric. F(first packet), the first packet of the pair, MUST have a valid Type-P-Finite-One-way-Delay less than T_{max} (in other words, excluding packets which have undefined one-way delay) and MUST have been transmitted during the interval T, T_f. The second packet in the pair, F(second packet) MUST be the packet with the minimum valid value of Type-P-Finite-One-way-Delay for the stream, in addition to the criteria for F(first packet). If multiple packets have equal minimum Type-P-Finite-One-way-Delay values, then the value for the earliest arriving packet SHOULD be used.

*MinDelay, the Type-P-Finite-One-way-Delay value for F(second packet) given above.

*N, the number of packets received at the Destination meeting the F(first packet) criteria.

7.1.2. Definition and Metric Units

[TOC](#)

Using the definition above in section 5.1.2, we obtain the value of Type-P-Finite-One-way-Delay-Poisson/Periodic-Stream[i], the singleton for each packet[i] in the stream (a.k.a. FiniteDelay[i]). For each packet[i] that meets the F(first packet) criteria given above: Type-P-One-way-pdv-refmin-Poisson/Periodic-Stream[i] = PDV[i] = FiniteDelay[i] - MinDelay where PDV[i] is in units of time in seconds, expressed in sufficiently low resolution to convey meaningful quantitative information. For example, resolution of microseconds is usually sufficient.

7.1.3. Discussion and other details

[TOC](#)

This metric produces a sample of delay variation normalized to the minimum delay of the sample. The resulting delay variation distribution is independent of the sending sequence (although specific FiniteDelay values within the distribution may be correlated, depending on various stream parameters such as packet spacing). This metric is equivalent to the IP Packet Delay Variation parameter defined in [\[Y.1540\] \(ITU-T Recommendation Y.1540, "Internet protocol data communication service - IP packet transfer and availability performance parameters," December 2002.\)](#).

7.1.4. Statistics: Mean, Variance, Skewness, Quantile

[TOC](#)

We define the mean PDV as follows (where all packets $i = 1$ through N have a value for $PDV[i]$):

Type-P-One-way-pdv-refmin-Mean = MeanPDV =

$$\frac{1}{N} \sum_{i=1}^N PDV[i]$$

We define the variance of PDV as follows:

Type-P-One-way-pdv-refmin-Variance = VarPDV =

$$\frac{1}{(N-1)} \sum_{i=1}^N (PDV[i] - \text{MeanPDV})^2$$

We define the skewness of PDV as follows:

Type-P-One-way-pdv-refmin-Skewness = SkewPDV =

$$\frac{\frac{1}{N} \sum_{i=1}^N (PDV[i] - \text{MeanPDV})^3}{\sqrt{(N-1) * \text{VarPDV}}^{3/2}}$$

We define the Quantile of the IPDVRefMin sample as the value where the specified fraction of singletons is less than the given value.

[TOC](#)

7.1.5. Composition Functions:

This section gives two alternative composition functions. The objective is to estimate a quantile of the complete path delay variation distribution. The composed quantile will be estimated using information from the sub-path delay variation distributions.

7.1.5.1. Approximate Convolution

[TOC](#)

The Type-P-Finite-One-way-Delay-Poisson/Periodic-Stream samples from each sub-path are summarized as a histogram with 1 ms bins representing the one-way delay distribution.

From [TBP], the distribution of the sum of independent random variables can be derived using the relation:

Type-P-Composite-One-way-pdv-refmin-quantile-a =

$$P(X + Y + Z \leq a) = \int_z \int_y P(X \leq a - y - z) * P(Y = y) * P(Z = z) dy dz$$

where X, Y, and Z are random variables representing the delay variation distributions of the sub-paths of the complete path (in this case, there are three sub-paths), and a is the quantile of interest. Note dy and dz indicate partial integration here.

This relation can be used to compose a quantile of interest for the complete path from the sub-path delay distributions. The histograms with 1 ms bins are discrete approximations of the delay distributions.

7.1.5.2. Normal Power Approximation

[TOC](#)

Type-P-One-way-Composite-pdv-refmin-NPA for the complete Source to Destination path can be calculated by combining statistics of all the constituent sub-paths in the following process:

< see [\[Y.1541\] \(ITU-T Recommendation Y.1541, "Network Performance Objectives for IP-based Services," February 2006.\)](#) clause 8 and Appendix X >

7.1.6. Statement of Conjecture

[TOC](#)

The delay distribution of a sufficiently large stream of packets measured on each sub-path during the interval [T, Tf] will be

sufficiently stationary and the sub-path distributions themselves are sufficiently independent, so that summary information describing the sub-path distributions can be combined to estimate the delay distribution of complete path.

7.1.7. Justification of the Composition Function

[TOC](#)

See the common section.

7.1.8. Sources of Deviation from the Ground Truth

[TOC](#)

In addition to the common deviations, a few additional sources exist here. For one, very tight distributions with range on the order of a few milliseconds are not accurately represented by a histogram with 1 ms bins. This size was chosen assuming an implicit requirement on accuracy: errors of a few milliseconds are acceptable when assessing a composed distribution quantile.

Also, summary statistics cannot describe the subtleties of an empirical distribution exactly, especially when the distribution is very different from a classical form. Any procedure that uses these statistics alone may incur error.

7.1.9. Specific cases where the conjecture might fail

[TOC](#)

If the delay distributions of the sub-paths are somehow correlated, then neither of these composition functions will be reliable estimators of the complete path distribution.

In practice, sub-path delay distributions with extreme outliers have increased the error of the composed metric estimate.

7.1.10. Application of Measurement Methodology

[TOC](#)

See the common section.

8. Security Considerations

[TOC](#)

8.1. Denial of Service Attacks

[TOC](#)

This metric requires a stream of packets sent from one host (source) to another host (destination) through intervening networks. This method could be abused for denial of service attacks directed at destination and/or the intervening network(s).

Administrators of source, destination, and the intervening network(s) should establish bilateral or multi-lateral agreements regarding the timing, size, and frequency of collection of sample metrics. Use of this method in excess of the terms agreed between the participants may be cause for immediate rejection or discard of packets or other escalation procedures defined between the affected parties.

8.2. User Data Confidentiality

[TOC](#)

Active use of this method generates packets for a sample, rather than taking samples based on user data, and does not threaten user data confidentiality. Passive measurement must restrict attention to the headers of interest. Since user payloads may be temporarily stored for length analysis, suitable precautions MUST be taken to keep this information safe and confidential. In most cases, a hashing function will produce a value suitable for payload comparisons.

8.3. Interference with the metrics

[TOC](#)

It may be possible to identify that a certain packet or stream of packets is part of a sample. With that knowledge at the destination and/or the intervening networks, it is possible to change the processing of the packets (e.g. increasing or decreasing delay) that may distort the measured performance. It may also be possible to generate additional packets that appear to be part of the sample metric. These additional packets are likely to perturb the results of the sample measurement.

To discourage the kind of interference mentioned above, packet interference checks, such as cryptographic hash, may be used.

[TOC](#)

9. IANA Considerations

Metrics defined in this memo will be registered in the IANA IPPM METRICS REGISTRY as described in initial version of the registry [\[RFC4148\]](#) (Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry," August 2005.).

10. Issues (Open and Closed)

TOC

>>>>>>>>>>>>>Issue:

What is the relationship between the decomposition and composition metrics? Should we put both kinds in one draft to make up a framework? The motivation of decomposition is as follows:

The One-way measurement can provide result to show what the network performance between two end hosts is and whether it meets operator expectations or not. It cannot provide further information to engineers where and how to improve the performance between the source and the destination. For instance, if the network performance is not acceptable in terms of the One-way measurement, in which part of the network the engineers should put their efforts. This question can to be answered by decompose the One-way measurement to sub-path measurement to investigate the performance of different part of the network.

Editor's Questions for clarification: What additional information would be provided to the decomposition process, beyond the measurement of the complete path?

Is the decomposition described above intended to estimate a metric for some/all disjoint sub-paths involved in the complete path?

>>>>>>>>>>>>>>>RESOLUTION: treat this topic in a separate memo

>>>>>>>>>>>>>>>Issue

Section 7 defines a new type of metric, a “combination” of metrics for one-way delay and packet loss. The purpose of this metric is to link these two primary metrics in a convenient way.

Readers are asked to comment on the efficiency of the combination metric.

```
>>>>>>>>>>>>>>>RESOLUTION: If a delay singleton is recorded as having
"undefined" delay when the packet does not arrive within the waiting
time Tmax, then this information is sufficient to determine the
fraction of lost packets in the sample, and the additional loss
indication of this combo is not needed.
```

>>>>>>>>>>>>>>> Issue

Can we introduce multicast metrics here, without causing too much confusion? Should the multicast version of this draft wait until the Unicast concepts are stable (or maybe appear in a separate draft)?

>>>>>>>>>>>RESOLUTION: No and Yes.

11. Acknowledgements

[TOC](#)

12. References

[TOC](#)

12.1. Normative References

[TOC](#)

[I-D.ietf-ippm-framework-compagg]	Morton, A., " Framework for Metric Composition ," draft-ietf-ippm-framework-compagg-09 (work in progress), December 2009 (TXT).
[RFC2119]	Bradner, S., " Key words for use in RFCs to Indicate Requirement Levels ," BCP 14, RFC 2119, March 1997 (TXT , HTML , XML).
[RFC2330]	Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, " Framework for IP Performance Metrics ," RFC 2330, May 1998 (TXT , HTML , XML).
[RFC2679]	Almes, G., Kalidindi, S., and M. Zekauskas, " A One-way Delay Metric for IPPM ," RFC 2679, September 1999 (TXT).
[RFC2680]	Almes, G., Kalidindi, S., and M. Zekauskas, " A One-way Packet Loss Metric for IPPM ," RFC 2680, September 1999 (TXT).
[RFC3393]	Demichelis, C. and P. Chimento, " IP Packet Delay Variation Metric for IP Performance Metrics (IPPM) ," RFC 3393, November 2002 (TXT).
[RFC3432]	Raisanen, V., Grotefeld, G., and A. Morton, " Network performance measurement with periodic streams ," RFC 3432, November 2002 (TXT).
[RFC4148]	Stephan, E., " IP Performance Metrics (IPPM) Metrics Registry ," BCP 108, RFC 4148, August 2005 (TXT).

12.2. Informative References

[TOC](#)

[I-D.ietf-ippm-multimetrics]	Stephan, E., Liang, L., and A. Morton, " IP Performance Metrics (IPPM) for spatial and multicast ," draft-ietf-ippm-multimetrics-12 (work in progress), September 2009 (TXT).
------------------------------	---

[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.1541, "Network Performance Objectives for IP-based Services," February 2006.

Authors' Addresses

[TOC](#)

	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/
	Emile Stephan
	France Telecom Division R&D
	2 avenue Pierre Marzin
	Lannion, F-22307
	France
Phone:	
Fax:	+33 2 96 05 18 52
Email:	emile.stephan@orange-ftgroup.com
URI:	

Full Copyright Statement

[TOC](#)

Copyright © 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.