

Internet Engineering Task Force (IETF)
Request for Comments: 6049
Category: Standards Track
ISSN: 2070-1721

A. Morton
AT&T Labs
E. Stephan
France Telecom Orange
January 2011

Spatial Composition of Metrics

Abstract

This memo utilizes IP performance metrics that are applicable to both complete paths and sub-paths, and it defines relationships to compose a complete path metric from the sub-path metrics with some accuracy with regard to 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.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6049>.

Copyright Notice

Copyright (c) 2011 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 (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

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.

Table of Contents

1. Introduction	4
1.1. Motivation	6
1.2. Requirements Language	6
2. Scope and Application	6
2.1. Scope of Work	6
2.2. Application	7
2.3. Incomplete Information	7
3. Common Specifications for Composed Metrics	8
3.1. Name: Type-P	8
3.1.1. Metric Parameters	8
3.1.2. Definition and Metric Units	9
3.1.3. Discussion and Other Details	9
3.1.4. Statistic	9
3.1.5. Composition Function	9
3.1.6. Statement of Conjecture and Assumptions	10
3.1.7. Justification of the Composition Function	10
3.1.8. Sources of Deviation from the Ground Truth	10
3.1.9. Specific Cases where the Conjecture Might Fail	11
3.1.10. Application of Measurement Methodology	12
4. One-Way Delay Composed Metrics and Statistics	12
4.1. Name: Type-P-Finite-One-way-Delay-<Sample>-Stream	12
4.1.1. Metric Parameters	12
4.1.2. Definition and Metric Units	12
4.1.3. Discussion and Other Details	13
4.1.4. Statistic	13
4.2. Name: Type-P-Finite-Composite-One-way-Delay-Mean	13
4.2.1. Metric Parameters	13
4.2.2. Definition and Metric Units of the Mean Statistic	14
4.2.3. Discussion and Other Details	14
4.2.4. Statistic	14
4.2.5. Composition Function: Sum of Means	14
4.2.6. Statement of Conjecture and Assumptions	15
4.2.7. Justification of the Composition Function	15
4.2.8. Sources of Deviation from the Ground Truth	15
4.2.9. Specific Cases where the Conjecture Might Fail	15
4.2.10. Application of Measurement Methodology	16
4.3. Name: Type-P-Finite-Composite-One-way-Delay-Minimum	16
4.3.1. Metric Parameters	16
4.3.2. Definition and Metric Units of the Minimum Statistic	16
4.3.3. Discussion and Other Details	16
4.3.4. Statistic	16
4.3.5. Composition Function: Sum of Minima	16
4.3.6. Statement of Conjecture and Assumptions	17
4.3.7. Justification of the Composition Function	17
4.3.8. Sources of Deviation from the Ground Truth	17

4.3.9.	Specific Cases where the Conjecture Might Fail	17
4.3.10.	Application of Measurement Methodology	17
5.	Loss Metrics and Statistics	18
5.1.	Type-P-Composite-One-way-Packet-Loss-Empirical-Probability	18
5.1.1.	Metric Parameters	18
5.1.2.	Definition and Metric Units	18
5.1.3.	Discussion and Other Details	18
5.1.4.	Statistic: Type-P-One-way-Packet-Loss-Empirical-Probability ...	18
5.1.5.	Composition Function: Composition of Empirical Probabilities	18
5.1.6.	Statement of Conjecture and Assumptions	19
5.1.7.	Justification of the Composition Function	19
5.1.8.	Sources of Deviation from the Ground Truth	19
5.1.9.	Specific Cases where the Conjecture Might Fail	19
5.1.10.	Application of Measurement Methodology	19
6.	Delay Variation Metrics and Statistics	20
6.1.	Name: Type-P-One-way-pdv-refmin-<Sample>-Stream	20
6.1.1.	Metric Parameters	20
6.1.2.	Definition and Metric Units	20
6.1.3.	Discussion and Other Details	21
6.1.4.	Statistics: Mean, Variance, Skewness, Quantile	21
6.1.5.	Composition Functions	22
6.1.6.	Statement of Conjecture and Assumptions	23
6.1.7.	Justification of the Composition Function	23
6.1.8.	Sources of Deviation from the Ground Truth	23
6.1.9.	Specific Cases where the Conjecture Might Fail	24
6.1.10.	Application of Measurement Methodology	24
7.	Security Considerations	24
7.1.	Denial-of-Service Attacks	24
7.2.	User Data Confidentiality	24
7.3.	Interference with the Metrics	24
8.	IANA Considerations	25
9.	Contributors and Acknowledgements	27
10.	References	28
10.1.	Normative References	28
10.2.	Informative References	28

1. Introduction

The IP Performance Metrics (IPPM) framework [RFC2330] describes two forms of metric composition: spatial and temporal. The composition framework [RFC5835] 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] lists four points that the metric definitions should include:

- o the specific conjecture applied to the metric and assumptions of the statistical model of the process being measured (if any; see [RFC2330], Section 12),
- o a justification of the practical utility of the composition in terms of making accurate measurements of the metric on the path,
- o a justification of the usefulness of the composition in terms of making analysis of the path using A-frame concepts more effective, and
- o an analysis of how the conjecture could be incorrect.

Also, [RFC2330] gives an example using the 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.

1.1. Motivation

One-way metrics defined in other RFCs (such as [RFC2679] and [RFC2680]) all assume that the measurement can be practically carried out between the source and the destination of interest. Sometimes there are reasons that the measurement cannot 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.

1.2. 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].

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

2. Scope and Application

2.1. Scope of Work

For the primary IP Performance Metrics RFCs for loss [RFC2680], delay [RFC2679], and delay variation [RFC3393], this memo gives a set of metrics that can be composed from the same or similar sub-path metrics. This means that the composition function may utilize:

- o the same metric for each sub-path;
- o multiple metrics for each sub-path (possibly one that is the same as the complete path metric);
- o a single sub-path metric that is different from the complete path metric;
- o different measurement techniques like active [RFC2330], [RFC3432] and passive [RFC5474].

We note a possibility: using a complete path metric and all but one sub-path metric to infer the performance of the missing sub-path, especially when the "last" sub-path metric is missing. However, such de-composition calculations, and the corresponding set of issues they raise, are beyond the scope of this memo.

2.2. Application

The composition framework [RFC5835] requires the specification of the applicable circumstances for each metric. In particular, each section addresses whether the metric:

- o Requires the same test packets to traverse all sub-paths or may use similar packets sent and collected separately in each sub-path.
- o Requires homogeneity of measurement methodologies or can allow a degree of flexibility (e.g., active, active spatial division [RFC5644], or passive methods produce the "same" metric). Also, the applicable sending streams will be specified, such as Poisson, Periodic, or both.
- o Needs information or access that will only be available within an operator's domain, or is applicable to inter-domain composition.
- o Requires synchronized measurement start and stop times in all sub-paths or largely overlapping measurement intervals, or no timing requirements.
- o Requires the assumption of sub-path independence with regard to the metric being defined/composed or other assumptions.
- o Has known sources of inaccuracy/error and identifies the sources.

2.3. Incomplete Information

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.

3. Common Specifications for Composed Metrics

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.

Also, the index variables are represented as follows:

- o m = index for packets sent.
- o n = index for packets received.
- o s = index for involved sub-paths.

3.1. Name: Type-P

All metrics use the "Type-P" convention as described in [RFC2330]. The rest of the name is unique to each metric.

3.1.1. Metric Parameters

- o Src, the IP address of a host.
- o Dst, the IP address of a host.
- o T, a time (start of test interval).
- o Tf, a time (end of test interval).
- o lambda, a rate in reciprocal seconds (for Poisson Streams).
- o incT, the nominal duration of inter-packet interval, first bit to first bit (for Periodic Streams).
- o dT, the duration of the allowed interval for Periodic Stream sample start times.
- o 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).
- o TstampSrc, the wire time of the packet as measured at MP(Src) (measurement point at the source).
- o TstampDst, the wire time of the packet as measured at MP(Dst), assigned to packets that arrive within a "reasonable" time.

- o 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.
- o M, the total number of packets sent between T0 and Tf.
- o N, the total number of packets received at Dst (sent between T0 and Tf).
- o S, the number of sub-paths involved in the complete Src-Dst path.
- o Type-P, as defined in [RFC2330], which includes any field that may affect a packet's treatment as it traverses the network.

In metric names, the term "<Sample>" is intended to be replaced by the name of the method used to define a sample of values of parameter TstampSrc. This can be done in several ways, including:

1. Poisson: a pseudo-random Poisson process of rate lambda, whose values fall between T and Tf. The time interval between successive values of TstampSrc will then average 1/lambda, as per [RFC2330].
2. Periodic: a Periodic stream process with pseudo-random start time T0 between T and dT, and nominal inter-packet interval incT, as per [RFC3432].

3.1.2. Definition and Metric Units

This section is unique for every metric.

3.1.3. Discussion and Other Details

This section is unique for every metric.

3.1.4. Statistic

This section is unique for every metric.

3.1.5. Composition Function

This section is unique for every metric.

3.1.6. Statement of Conjecture and Assumptions

This section is unique for each metric. The term "ground truth" is frequently used in these sections and is defined in Section 4.7 of [RFC5835].

3.1.7. 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.

3.1.8. Sources of Deviation from the Ground Truth

3.1.8.1. Sub-Path List Differs from Complete Path

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 links when compared to packets with the source and destination addresses of the complete path. Example sources of parallel paths include Equal Cost Multi-Path and parallel (or bundled) links. Therefore, the performance estimated from 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.

We also note the possibility of re-routing during a measurement interval, as it may affect the correspondence between packets traversing the complete path and the sub-paths that were "involved" prior to the re-route.

3.1.8.2. Sub-Path Contains Extra Network Elements

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.

3.1.8.3. Sub-Paths Have Incomplete Coverage

Measurements of sub-path performance may not cover all the network elements on the complete path. For example, the network exchange points might be excluded unless a cooperative measurement is conducted. In this example, test packets on the previous sub-path are received just before the exchange point, and test packets on the next sub-path are injected just after the same exchange point. Clearly, the set of sub-path measurements SHOULD cover all critical network elements in the complete path.

3.1.8.4. Absence of Route

At a specific point in time, no viable route exists between the complete path source and destination. The routes selected for one or more sub-paths therefore differ from the complete path. Consequently, spatial composition may produce finite estimation of a ground truth metric (see Section 4.7 of [RFC5835]) between a source and a destination, even when the route between them is undefined.

3.1.9. Specific Cases where the Conjecture Might Fail

This section is unique for most metrics (see the metric-specific sections).

For delay-related metrics, one-way delay always depends on packet size and link capacity, since it is measured in [RFC2679] from first bit to last bit. If the size of an IP packet changes on its route (due to encapsulation), this can influence delay performance. However, the main error source may be the additional processing associated with encapsulation and encryption/decryption if not experienced or accounted for in sub-path measurements.

Fragmentation is a major issue for composition accuracy, since all metrics require all fragments to arrive before proceeding, and fragmented complete path performance is likely to be different from performance with non-fragmented packets and composed metrics based on non-fragmented sub-path measurements.

Highly manipulated routing can cause measurement error if not expected and compensated for. For example, policy-based MPLS routing could modify the class of service for the sub-paths and complete path.

3.1.10. Application of Measurement Methodology

- o The methodology SHOULD use similar packets sent and collected separately in each sub-path, where "similar" in this case means that Type-P contains as many equal attributes as possible, while recognizing that there will be differences. Note that Type-P includes stream characteristics (e.g., Poisson, Periodic).
- o The methodology allows a degree of flexibility regarding test stream generation (e.g., active or passive methods can produce an equivalent result, but the lack of control over the source, timing, and correlation of passive measurements is much more challenging).
- o Poisson and/or Periodic streams are RECOMMENDED.
- o The methodology applies to both inter-domain and intra-domain composition.
- o The methodology SHOULD have synchronized measurement time intervals in all sub-paths, but largely overlapping intervals MAY suffice.
- o Assumption of sub-path independence with regard to the metric being defined/composed is REQUIRED.

4. One-Way Delay Composed Metrics and Statistics

4.1. Name: Type-P-Finite-One-way-Delay-<Sample>-Stream

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

4.1.1. Metric Parameters

See the common parameters section (Section 3.1.1).

4.1.2. Definition and Metric Units

Using the parameters above, we obtain the value of the Type-P-One-way-Delay singleton as per [RFC2679].

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

Type-P-Finite-One-way-Delay-<Sample>-Stream[i] =

FiniteDelay[i] = TstampDst - TstampSrc

This metric is measured in units of time in seconds, expressed in sufficiently low resolution to convey meaningful quantitative information. For example, resolution of microseconds is usually sufficient.

4.1.3. Discussion and Other Details

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, at the same time avoiding events with undefined outcomes. This approach is derived from the treatment of lost packets in [RFC3393], and is similar to [Y.1540].

4.1.4. Statistic

All statistics defined in [RFC2679] are applicable to the finite one-way delay, and additional metrics are possible, such as the mean (see below).

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

This section describes a statistic based on the Type-P-Finite-One-way-Delay-<Sample>-Stream metric.

4.2.1. Metric Parameters

See the common parameters section (Section 3.1.1).

4.2.2. Definition and Metric Units of the Mean Statistic

We define

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

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

where all packets $n = 1$ through N have finite singleton delays.

This metric is measured in units of time in seconds, expressed in sufficiently fine resolution to convey meaningful quantitative information. For example, resolution of microseconds is usually sufficient.

4.2.3. Discussion and Other Details

The Type-P-Finite-One-way-Delay-Mean metric requires the conditional delay distribution described in Section 4.1.3.

4.2.4. Statistic

This metric, a mean, does not require additional statistics.

4.2.5. Composition Function: Sum of Means

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

Then the

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

$$\text{CompMeanDelay} = \frac{\sum_{s=1}^S \text{MeanDelay [s]}}{s = 1}$$

where sub-paths $s = 1$ to S are involved in the complete path.

4.2.6. Statement of Conjecture and Assumptions

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 ground truth 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.

It is assumed that the one-way delay distributions of the sub-paths and the complete path are continuous. The mean of multi-modal distributions has the unfortunate property that such a value may never occur.

4.2.7. Justification of the Composition Function

See the common section (Section 3).

4.2.8. Sources of Deviation from the Ground Truth

See the common section (Section 3).

4.2.9. Specific Cases where the Conjecture Might Fail

If any of the sub-path distributions are multi-modal, 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 a sufficiently robust statistic to produce a reliable estimate, or to be useful even if it can be measured.

If a link contributing non-negligible delay is erroneously included or excluded, the composition will be in error.

4.2.10. Application of Measurement Methodology

The requirements of the common section (Section 3) apply here as well.

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

This section describes a statistic based on the Type-P-Finite-One-way-Delay-`<Sample>`-Stream metric, and the composed metric based on that statistic.

4.3.1. Metric Parameters

See the common parameters section (Section 3.1.1).

4.3.2. Definition and Metric Units of the Minimum Statistic

We define

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

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

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

where all packets $n = 1$ through N have finite singleton delays.

This metric is measured in units of time in seconds, expressed in sufficiently fine resolution to convey meaningful quantitative information. For example, resolution of microseconds is usually sufficient.

4.3.3. Discussion and Other Details

The Type-P-Finite-One-way-Delay-Minimum metric requires the conditional delay distribution described in Section 4.1.3.

4.3.4. Statistic

This metric, a minimum, does not require additional statistics.

4.3.5. Composition Function: Sum of Minima

The Type-P-Finite-Composite-One-way-Delay-Minimum, or `CompMinDelay`, for the complete source to destination path can be calculated from the sum of the minimum delays of all of its S constituent sub-paths.

Then the

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

$$\text{CompMinDelay} = \frac{S}{s} > (\text{MinDelay [s]})$$

s = 1

4.3.6. Statement of Conjecture and Assumptions

The minimum of a sufficiently large stream of packets measured on each sub-path during the interval [T, Tf] will be representative of the ground truth 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.

It is assumed that the one-way delay distributions of the sub-paths and the complete path are continuous.

4.3.7. Justification of the Composition Function

See the common section (Section 3).

4.3.8. Sources of Deviation from the Ground Truth

See the common section (Section 3).

4.3.9. Specific Cases where the Conjecture Might Fail

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.

4.3.10. Application of Measurement Methodology

The requirements of the common section (Section 3) apply here as well.

5. Loss Metrics and Statistics

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

5.1.1. Metric Parameters

See the common parameters section (Section 3.1.1).

5.1.2. Definition and Metric Units

Using the parameters above, we obtain the value of the Type-P-One-way-Packet-Loss singleton and stream as per [RFC2680].

We obtain a sequence of pairs with elements as follows:

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

5.1.3. Discussion and Other Details

None.

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

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 =

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

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

5.1.5. Composition Function: Composition of Empirical Probabilities

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

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

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

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

5.1.6. Statement of Conjecture and Assumptions

The empirical probability of loss calculated on a sufficiently large stream of packets measured on each sub-path during the interval $[T, T_f]$ will be representative of the ground truth empirical 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 empirical loss probability.

5.1.7. Justification of the Composition Function

See the common section (Section 3).

5.1.8. Sources of Deviation from the Ground Truth

See the common section (Section 3).

5.1.9. 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-borne worm, loss due to queue overflow in one network may help another network to carry its traffic without loss.

5.1.10. Application of Measurement Methodology

See the common section (Section 3).

6. Delay Variation Metrics and Statistics

6.1. Name: Type-P-One-way-pdv-refmin-<Sample>-Stream

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 (with the exception of [RFC5481]).

6.1.1. Metric Parameters

In addition to the parameters of Section 3.1.1:

- o TstampSrc[i], the wire time of packet[i] as measured at MP(Src) (measurement point at the source).
- o TstampDst[i], the wire time of packet[i] as measured at MP(Dst), assigned to packets that arrive within a "reasonable" time.
- o B, a packet length in bits.
- o F, a selection function unambiguously defining the packets from the stream that are selected for the packet-pair computation of this metric. F(current packet), the first packet of the pair, MUST have a valid Type-P-Finite-One-way-Delay less than Tmax (in other words, excluding packets that have undefined one-way delay) and MUST have been transmitted during the interval [T, Tf]. The second packet in the pair, F(min_delay 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(current 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.
- o MinDelay, the Type-P-Finite-One-way-Delay value for F(min_delay packet) given above.
- o N, the number of packets received at the destination that meet the F(current packet) criteria.

6.1.2. Definition and Metric Units

Using the definition above in Section 5.1.2, we obtain the value of Type-P-Finite-One-way-Delay-<Sample>-Stream[n], the singleton for each packet[i] in the stream (a.k.a. FiniteDelay[i]).

For each packet[n] that meets the F(first packet) criteria given above: Type-P-One-way-pdv-refmin-<Sample>-Stream[n] =

$$\text{PDV}[n] = \text{FiniteDelay}[n] - \text{MinDelay}$$

where PDV[i] is in units of time in seconds, expressed in sufficiently fine resolution to convey meaningful quantitative information. For example, resolution of microseconds is usually sufficient.

6.1.3. Discussion and Other Details

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].

6.1.4. Statistics: Mean, Variance, Skewness, Quantile

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

$$\text{Type-P-One-way-pdv-refmin-Mean} = \text{MeanPDV} =$$

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

We define the variance of PDV as follows:

$$\text{Type-P-One-way-pdv-refmin-Variance} = \text{VarPDV} =$$

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

We define the skewness of PDV as follows:

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

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

(See Appendix X of [Y.1541] for additional background information.)

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

6.1.5. 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.

6.1.5.1. Approximate Convolution

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

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

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

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

Note that dy and dz indicate partial integration above, and that y and z are the integration variables. Also, the probability of an outcome is indicated by the symbol $P(\text{outcome})$, 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.

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.

6.1.5.2. Normal Power Approximation (NPA)

Type-P-One-way-Composite-pdv-refmin-NPA for the complete source to destination path can be calculated by combining the statistics of all the constituent sub-paths in the process described in [Y.1541], Clause 8 and Appendix X.

6.1.6. Statement of Conjecture and Assumptions

The delay distribution of a sufficiently large stream of packets measured on each sub-path during the interval $[T, T_f]$ 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 the complete path.

It is assumed that the one-way delay distributions of the sub-paths and the complete path are continuous.

6.1.7. Justification of the Composition Function

See the common section (Section 3).

6.1.8. Sources of Deviation from the Ground Truth

In addition to the common deviations, a few additional sources exist here. For one, very tight distributions with ranges 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.

6.1.9. Specific Cases where the Conjecture Might Fail

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.

6.1.10. Application of Measurement Methodology

See the common section (Section 3).

7. Security Considerations

7.1. Denial-of-Service Attacks

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 the destination and/or the intervening network(s).

Administrators of source, destination, and intervening networks should establish bilateral or multilateral agreements regarding the timing, size, and frequency of collection of sample metrics. Use of this method in excess of the terms agreed upon between the participants may be cause for immediate rejection or discarding of packets, or other escalation procedures defined between the affected parties.

7.2. User Data Confidentiality

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.

7.3. Interference with the Metrics

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), which 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.

8. IANA Considerations

Metrics defined in the IETF are typically registered in the IANA IPPM Metrics Registry as described in the initial version of the registry [RFC4148].

IANA has registered the following metrics in the IANA-IPPM-METRICS-REGISTRY-MIB:

```
ietfFiniteOneWayDelayStream OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-Finite-One-way-Delay-Stream"
  REFERENCE "RFC 6049, Section 4.1."
  ::= { ianaIppmMetrics 71 }

ietfFiniteOneWayDelayMean OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-Finite-One-way-Delay-Mean"
  REFERENCE "RFC 6049, Section 4.2."
  ::= { ianaIppmMetrics 72 }

ietfCompositeOneWayDelayMean OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-Finite-Composite-One-way-Delay-Mean"
  REFERENCE "RFC 6049, Section 4.2.5."
  ::= { ianaIppmMetrics 73 }

ietfFiniteOneWayDelayMinimum OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-Finite-One-way-Delay-Minimum"
  REFERENCE "RFC 6049, Section 4.3.2."
  ::= { ianaIppmMetrics 74 }
```

```
ietfCompositeOneWayDelayMinimum OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-Finite-Composite-One-way-Delay-Minimum"
  REFERENCE "RFC 6049, Section 4.3."
  ::= { ianaIppmMetrics 75 }

ietfOneWayPktLossEmpiricProb OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-One-way-Packet-Loss-Empirical-Probability"
  REFERENCE "RFC 6049, Section 5.1.4"
  ::= { ianaIppmMetrics 76 }

ietfCompositeOneWayPktLossEmpiricProb OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-Composite-One-way-Packet-Loss-Empirical-Probability"
  REFERENCE "RFC 6049, Section 5.1."
  ::= { ianaIppmMetrics 77 }

ietfOneWayPdvRefminStream OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-One-way-pdv-refmin-Stream"
  REFERENCE "RFC 6049, Section 6.1."
  ::= { ianaIppmMetrics 78 }

ietfOneWayPdvRefminMean OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-One-way-pdv-refmin-Mean"
  REFERENCE "RFC 6049, Section 6.1.4."
  ::= { ianaIppmMetrics 79 }

ietfOneWayPdvRefminVariance OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-One-way-pdv-refmin-Variance"
  REFERENCE "RFC 6049, Section 6.1.4."
  ::= { ianaIppmMetrics 80 }
```

```
ietfOneWayPdvRefminSkewness OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-One-way-pdv-refmin-Skewness"
  REFERENCE "RFC 6049, Section 6.1.4."
  ::= { ianaIppmMetrics 81 }

ietfCompositeOneWayPdvRefminQtil OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-Composite-One-way-pdv-refmin-quantile-a"
  REFERENCE "RFC 6049, Section 6.1.5.1."
  ::= { ianaIppmMetrics 82 }

ietfCompositeOneWayPdvRefminNPA OBJECT-IDENTITY
  STATUS current
  DESCRIPTION
    "Type-P-One-way-Composite-pdv-refmin-NPA"
  REFERENCE "RFC 6049, Section 6.1.5.2."
  ::= { ianaIppmMetrics 83 }
```

9. Contributors and Acknowledgements

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@cariden.com>
- Roman Krzanowski <roman.krzanowski@verizon.com>
- Maurizio Molina <maurizio.molina@dante.org.uk>
- Lei Liang <L.Liang@surrey.ac.uk>
- Dave Hoeflin <dhoeflin@att.com>

A long time ago, in a galaxy far, far away (Minneapolis), Will Leland suggested the simple and elegant Type-P-Finite-One-way-Delay concept. Thanks Will.

Yaakov Stein and Donald McLachlan also provided useful comments along the way.

10. References

10.1. 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.
- [RFC4148] Stephan, E., "IP Performance Metrics (IPPM) Metrics Registry", BCP 108, RFC 4148, August 2005.
- [RFC5835] Morton, A. and S. Van den Berghe, "Framework for Metric Composition", RFC 5835, April 2010.

10.2. Informative References

- [RFC5474] Duffield, N., Chiou, D., Claise, B., Greenberg, A., Grossglauser, M., and J. Rexford, "A Framework for Packet Selection and Reporting", RFC 5474, March 2009.
- [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, March 2009.
- [RFC5644] Stephan, E., Liang, L., and A. Morton, "IP Performance Metrics (IPPM): Spatial and Multicast", RFC 5644, October 2009.
- [STATS] Mood, A., Graybill, F., and D. Boes, "Introduction to the Theory of Statistics, 3rd Edition", McGraw-Hill, New York, NY, 1974.

- [Y.1540] ITU-T Recommendation Y.1540, "Internet protocol data communication service - IP packet transfer and availability performance parameters", November 2007.
- [Y.1541] ITU-T Recommendation Y.1541, "Network Performance Objectives for IP-based Services", February 2006.

Authors' Addresses

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

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

Stephan Emile
France Telecom Orange
2 avenue Pierre Marzin
Lannion, F-22307
France

EMail: emile.stephan@orange-ftgroup.com

