

Network Working Group
Internet-Draft
Expires: December 26, 2006

A. Morton, Ed.
AT&T Labs
S. Van den Berghe, Ed.
Ghent University - IBBT
June 24, 2006

Framework for Metric Composition
draft-ietf-ippm-framework-compagg-01

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

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

This Internet-Draft will expire on December 26, 2006.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This memo describes a framework for composing and aggregating metrics (both in time and in space) defined by [RFC 2330](#) and developed by the IPPM working group. The framework describes the generic composition and aggregation mechanisms. It provides a basis for additional documents that implement this framework for detailed, and practically useful, compositions and aggregations of metrics.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

Table of Contents

| | | |
|------------------------|--|--------------------|
| 1. | Introduction | 3 |
| 1.1. | Motivation | 3 |
| 1.1.1. | Reducing Measurement Overhead | 3 |
| 1.1.2. | Measurement Re-use | 4 |
| 1.1.3. | Data Reduction and Consolidation | 4 |
| 1.1.4. | Implications on Measurement Design and Reporting . . . | 5 |
| 2. | Purpose and Scope | 5 |
| 3. | Description of Metric Types | 5 |
| 3.1. | Temporal Aggregation Description | 5 |
| 3.2. | Spatial Aggregation Description | 6 |
| 3.3. | Spatial Composition Description | 7 |
| 3.4. | Help Metrics | 7 |
| 3.5. | Higher Order Composition | 7 |
| 4. | Requirements for Composed Metrics | 7 |
| 5. | Guidelines for Defining Composed Metrics | 9 |
| 5.1. | Ground Truth: Comparison with other IPPM Metrics | 9 |
| 5.2. | Deviation from the Ground Truth | 11 |
| 6. | IANA Considerations | 12 |
| 7. | Security Considerations | 12 |
| 8. | Acknowledgements | 12 |
| 9. | References | 12 |
| 9.1. | Normative References | 12 |
| 9.2. | Informative References | 13 |
| | Authors' Addresses | 14 |
| | Intellectual Property and Copyright Statements | 15 |

1. Introduction

The IPPM framework [RFC 2330](#) [[RFC2330](#)] describes two forms of metric composition, spatial and temporal. Also, the text suggests that the concepts of the analytical framework (or A-frame) would help to develop useful relationships to derive the composed metrics from real metrics. The effectiveness of composed metrics is dependent on their usefulness in analysis and applicability to practical measurement circumstances.

This memo expands on the notion of composition, and provides a detailed framework for several classes of metrics that were mentioned in the original IPPM framework. The classes include temporal aggregation, spatial aggregation, and spatial composition.

1.1. Motivation

Network operators have deployed measurement systems to serve many purposes, including performance monitoring, maintenance support, network engineering, and customer reporting. The collection of elementary measurements alone is not enough to understand a network's behaviour. In general, measurements need to be post-processed to present the most relevant information for each purpose. The first step is often a process of "composition" of single measurements or measurement sets into other forms. Composition and aggregation present several more post-processing opportunities to the network operator, and we describe the key motivations below.

1.1.1. Reducing Measurement Overhead

A network's measurement possibilities scale upward with the square of the number of nodes. But each measurement implies overhead, in terms of the storage for the results, the traffic on the network (assuming active methods), and the OA&M for the measurement system itself. In a large network, it is impossible to perform measurements from each node to all others.

An individual network operator should be able to organize their measurement paths along the lines of physical topology, or routing areas/Autonomous Systems, and thus minimize dependencies and overlap between different measurement paths. This way, the sheer number of measurements can be reduced, as long as the operator has a set of methods to estimate performance between any particular nodes when needed.

Composition and aggregation play a key role when the path of interest spans multiple networks, and where each operator conducts their own measurements. Here, the complete path performance may be estimated

from measurements on the component parts.

Operators that take advantage of the composition and aggregation methods recognize that the estimates may exhibit some additional error beyond that inherent in the measurements themselves, and so they are making a trade-off to achieve reasonable measurement system overhead.

1.1.2. Measurement Re-use

There are many different measurement users, each bringing specific requirements for the reporting timescale. Network managers and maintenance forces prefer to see results presented very rapidly, to detect problems quickly or see if their action has corrected a problem. On the other hand, network capacity planners and even network users sometimes prefer a long-term view of performance, for example to check trends. How can one set of measurements serve both needs?

The answer lies in temporal aggregation, where the short-term measurements needed by the operations community are combined to estimate a longer-term result for others. Also, problems with the measurement system itself may be isolated to one or more of the short-term measurements, rather than possibly invalidating an entire long-term measurement if the problem was undetected.

1.1.3. Data Reduction and Consolidation

Another motivation is data reduction. Assume there is a network domain in which delay measurements are performed among a subset of its nodes. A network manager might ask whether there is a problem with the network delay in general. It would be desirable to obtain a single value that gives an indication of the overall network delay. Spatial aggregation methods would address this need, and can produce the desired "single figure of merit" asked for, one that may also be useful in trend analysis.

The overall value would be calculated from the elementary delay measurements, but it not obvious how: for example, it may not be reasonable to average all delay measurements, as some paths (e.g. having a higher bandwidth or more important customers) might be considered more critical than others.

Metric composition can help to provide, from raw measurement data, some tangible, well-understood and agreed upon information about the service guarantees provided by a network. Such information can be used in the Service Level Agreement/Service Level Specification (SLA/SLS) contracts between a service provider and its customers.

1.1.4. Implications on Measurement Design and Reporting

If a network operator can anticipate needing to aggregate or compose overall metrics in the future, it is more efficient to start by considering the tenants of these methods in the measurement design/sampling plan, and reporting the results. The Summary Statistics of certain metrics are more conducive to composition than others. This figures prominently in the design of measurements and the results reports.

2. Purpose and Scope

The purpose of this memo is provide a common framework for the various classes of metrics based on composition of primary metrics. The scope is limited to the definitions of metrics that are composed from primary metrics using a deterministic function. Key information about each metric, such as its assumptions under which the relationship holds, and possible sources of error/circumstances where the composition may fail, are included.

At this time, the scope of effort is limited to the metrics for packet loss, delay, and delay variation. Composition of packet reordering metrics is considered a research topic, and beyond the scope at the time this memo was prepared.

This memo will retain the terminology of the IPPM Framework as much as possible, but will extend the terminology when necessary.

3. Description of Metric Types

This section defines the various classes of Composition. There are two classes more accurately referred to as aggregation over time and space, and the third is simply composition in space.

3.1. Temporal Aggregation Description

Aggregation in time is defined as the composition of metrics with the same type and scope obtained in different time instants or time windows. For example, starting from a time series of One-Way Delay measurements on a certain network path obtained in 5-minute periods and averaging groups of 12 consecutive values, we obtain a time series measurement with a coarser resolution (60 minutes). The main reason for doing time aggregation is to reduce the amount of data that has to be stored, and make the visualization/spotting of regular cycles and/or growing or decreasing trends easier. Another useful application is to detect anomalies or abnormal changes in the network

characteristics.

In [RFC 2330](#), the term "temporal composition" is introduced and differs from temporal aggregation in that it refers to methodologies to predict future metrics on the basis of past observations, exploiting the time correlation that certain metrics can exhibit. We do not consider this type of composition here.

>>>>>>>Comment: Why no forecasting? This was apparently a limit on the Geant2 project, but may not apply here.

3.2. Spatial Aggregation Description

Aggregation in space is defined as the combination of metrics of the same type and different scope, in order to estimate the overall performance of a larger domain. This combination may involve weighing the contributions of the input metrics.

Suppose we want to compose the average One-Way-Delay (OWD) experienced by flows traversing all the Origin-Destination (OD) pairs of a network domain (where the inputs are already metric "statistics"). Since we wish to include the effect of the traffic matrix on the result, it makes sense to weight each metric according to the traffic carried on the corresponding OD pair:

$$\text{OWD_sum} = f_1 * \text{OWD_1} + f_2 * \text{OWD_2} + \dots + f_n * \text{OWD_n}$$

where $f_i = \text{load_OD_i} / \text{total_load}$.

A simple average OWD across all network OD pairs would not use the traffic weighting.

Another example metric that is "aggregated in space", is the maximum edge-to-edge delay across a single domain. Assume that a Service Provider wants to advertise the maximum delay that transit traffic will experience while passing through his/her domain. There can be multiple edge-to-edge paths across a domain, and the Service Provider chooses either to publish a list of delays (each corresponding to a specific edge-to-edge path), or publish a single maximum value. The latter approach simplifies the publication of measurement information, and may be sufficient for some purposes. Similar operations can be provided to other metrics, e.g. "maximum edge-to-edge packet loss", etc.

We suggest that space aggregation is generally useful to obtain a summary view of the behaviour of large network portions, or in general of coarser aggregates. The metric collection time instant, i.e. the metric collection time window of measured metrics is not

considered in space aggregation. We assume that either it is consistent for all the composed metrics, e.g. compose a set of average delays all referred to the same time window, or the time window of each composed metric does not affect aggregated metric.

3.3. Spatial Composition Description

Concatenation in space is defined as the composition of metrics of same type and (ideally) different spatial scope, so that the resulting metric is representative of what the metric would be if obtained with a direct measurement over the sequence of the several spatial scopes. An example is the sum of OWDs of different edge-to-edge domain's delays, where the intermediate edge points are close to each other or happen to be the same. In this way, we can for example estimate OWD_AC starting from the knowledge of OWD_AB and OWD_BC. Note that there may be small gaps in measurement coverage, likewise there may be small overlaps (e.g., the link where test equipment connects to the network).

One key difference from examples of aggregation in space is that all sub-paths contribute equally to the composed metric, independent of the traffic load present.

3.4. Help Metrics

Finally, note that in practice there is often the need of extracting a new metric making some computation over one or more metrics with the same spatial and time scope. For example, the composed metric `rtt_sample_variance` may be composed from two different metrics: the help metric `rtt_square_sum` and the statistical metric `rtt_sum`. This operation is however more a simple calculation and not an aggregation or a concatenation, and we'll not investigate it further in this memo.

3.5. Higher Order Composition

Composed metrics might themselves be subject to further steps of composition or aggregation. An example would be a the delay of a maximal domain obtained through the spatial composition of several composed end-to-end delays (obtained through spatial composition). All requirements for first order composition metrics apply to higher order composition.

>>>> Comment Response: are more examples needed here?

4. Requirements for Composed Metrics

The definitions for all composed metrics MUST include sections to treat the following topics.

The description of each metric will clearly state:

1. the definition (and statistic, where appropriate);
2. the composition or aggregation relationship;
3. the specific conjecture on which the relationship is based;
4. a justification of practical utility or usefulness for analysis using the A-frame concepts;
5. one or more examples of how the conjecture could be incorrect and lead to inaccuracy;
6. the information to be reported.

Each metric will require a relationship to determine the aggregated or composed metric. The relationships may involve conjecture, and [\[RFC2330\]](#) lists four points that the metric definitions should include:

- o the specific conjecture applied to the metric,
- 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 aggregation or 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.

For each metric, the applicable circumstances are defined, in terms of whether the composition or aggregation:

- o 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.
- o Needs information or access that will only be available within an operator's domain, or is applicable to Inter-domain composition.
- o Requires precisely synchronized measurement time intervals in all component metrics, or loosely synchronized, or no timing

requirements.

- o Requires assumption of component metric independence w.r.t. the metric being defined/composed, or other assumptions.
- o Has known sources of inaccuracy/error, and identifies the sources.

5. Guidelines for Defining Composed Metrics

5.1. Ground Truth: Comparison with other IPPM Metrics

Figure 1 illustrates the process to derive a metric using spatial composition, and compares the composed metric to other IPPM metrics.

Metrics $\langle M1, M2, M3 \rangle$ describe the performance of sub-paths between the Source and Destination of interest during time interval $\langle T, T_f \rangle$. These metrics are the inputs for a Composition Function that produces a Composed Metric.




Figure 1 Comparison with other IPPM metrics

The Composed Metric is an estimate of an actual metric collected over the complete Source to Destination path. We say that the Complete Path Metric represents the "Ground Truth" for the Composed Metric. In other words, Composed Metrics seek to minimize error w.r.t. the Complete Path Metric.

Further, we observe that a Spatial Metric I-D.ietf-ippm-multimetrics [[I-D.ietf-ippm-multimetrics](#)] collected for packets traveling over the same set of sub-paths provide a basis for the Ground Truth of the individual Sub-Path metrics. We note that mathematical operations may be necessary to isolate the performance of each sub-path.

Next, we consider multiparty metrics as defined in [I-D.ietf-ippm-multimetrics], and their spatial composition. Measurements to each of the Receivers produce an element of the one-to-group metric. These elements can be composed from sub-path metrics and the composed metrics can be combined to create a composed one-to-group metric. Figure 2 illustrates this process.

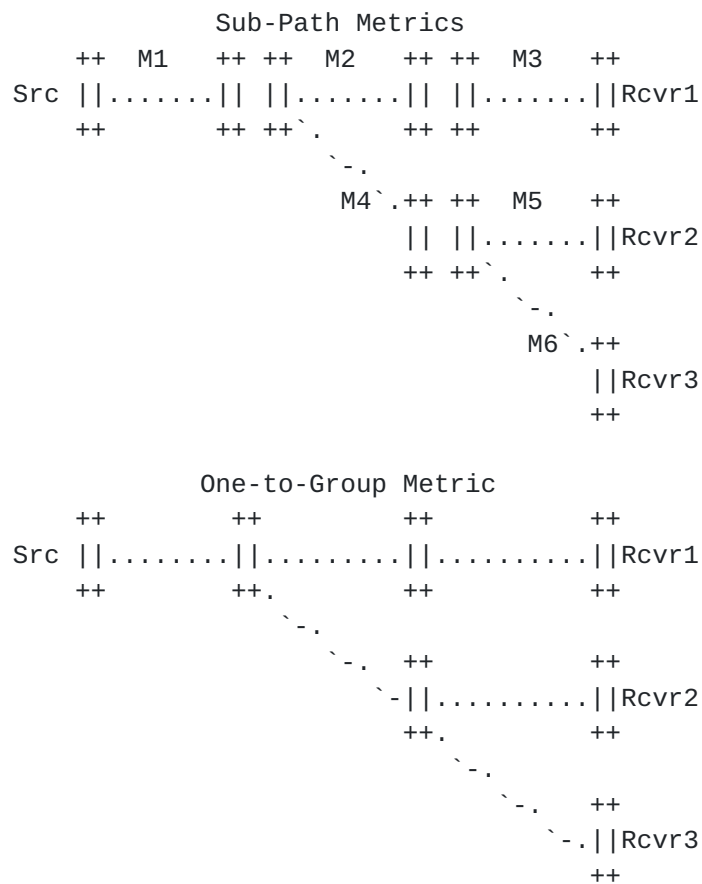


Figure 2 Composition of One-to-Group Metrics

Here, Sub-path Metrics M1, M2, and M3 are combined using a relationship to compose the metric applicable to the Src-Rcvr1 path. Similarly, M1, M4, and M5 are used to compose the Src-Rcvr2 metric and M1, M4, and M6 compose the Src-Rcvr3 metric.

The Composed One-to-Group Metric would list the Src-Rcvr metrics for each Receiver in the Group:

(Composed-Rcvr1, Composed-Rcvr2, Composed-Rcvr3)

The "Ground Truth" for this composed metric is of course an actual One-to-Group metric, where a single source packet has been measured after traversing the Complete Paths to the various receivers.

5.2. Deviation from the Ground Truth

A metric composition can deviate from the ground truth for several reasons. Two main aspects are:

- o The propagation of the inaccuracies of the underlying measurements when composing the metric. As part of the composition function, errors of measurements might propagate. Where possible, this analysis should be made and included with the description of each metric.
- o A difference in scope. When concatenating hop-by-hop active measurement results to obtain the end-to-end metric, the actual measured path will not be identical to the end-to-end path. It is in general difficult to quantify this deviation, but a metric definition might identify guidelines for keeping the deviation as small as possible.

The description of the metric composition MUST include an section identifying the deviation from the ground truth.

6. IANA Considerations

This document makes no request of IANA.

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

7. Security Considerations

8. Acknowledgements

The authors would like to thank Maurizio Molina, Andy Van Maele, Andreas Haneman, Igor Velimirovic, Andreas Solberg, Athanassios Liakopulos, David Schitz, Nicolas Simar and the Geant2 Project. We also acknowledge comments and suggestions from Phil Chimento, Emile Stephan and Lei Liang.

9. References

9.1. Normative References

- [I-D.ietf-ippm-multimetrics]
Stephan, E., "IP Performance Metrics (IPPM) for spatial and multicast", [draft-ietf-ippm-multimetrics-00](#) (work in progress), January 2006.
- [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.

[9.2.](#) Informative References

Authors' Addresses

Al Morton (editor)
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/>

Steven Van den Berghe (editor)
Ghent University - IBBT
G. Crommenlaan 8 bus 201
Gent 9050
Belgium

Phone: +32 9 331 49 73
Email: steven.vandenberghe@intec.ugent.be
URI: <http://www.ibcn.intec.ugent.be>

Intellectual Property Statement

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.

Disclaimer of Validity

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 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.

Copyright Statement

Copyright (C) The Internet Society (2006). 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.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

