

Network Working Group
Internet-Draft
Expires: August 28, 2006

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

Framework for Metric Composition
draft-ietf-ippm-framework-compagg-00

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 August 28, 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.

Internet-Draft Framework for Metric Composition February 2006

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
2.	Purpose and Scope	4
3.	Description of Metric Types	4
3.1.	Time Aggregation Description	4
3.2.	Spatial Aggregation Description	5
3.3.	Spatial Composition Description	5
3.4.	Help Metrics	6
3.5.	Higher Order Composition	6
4.	Requirements for Composed Metrics	6
5.	Guidelines for Defining Composed Metrics	7
5.1.	Ground Truth: Comparison with other IPPM Metrics	7
5.2.	Deviation from the Ground Truth	9
6.	IANA Considerations	10
7.	Security Considerations	10
8.	Acknowledgements	10
9.	References	10
9.1.	Normative References	10
9.2.	Informative References	11
	Authors' Addresses	12
	Intellectual Property and Copyright Statements	13

Internet-Draft

Framework for Metric Composition

February 2006

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

The deployment of a measurement infrastructure and the collection of elementary measurements are not enough to understand and keep under control the network's behaviour. Network measurements need in general to be post-processed to be useful for the several tasks of network engineering and management. The first step of this post processing is often a process of "composition" of single measurements or measurement sets into other ones. The reasons for doing so are briefly introduced here.

A first reason, mainly applicable to network engineering, is scalability. Due to the number of network elements in large networks, it is impossible to perform measurements from each element to all others. It is necessary to select a set of links of special interest and to perform the measurements there. Examples for this are active measurements of one-way delay, jitter, and loss.

Another reason may be data reduction (opposite need with respect to the previous one, where more data is generated). This is of interest for network planners and managers. Let us assume that there is

network domain in which delay measurements are performed among a subset of its elements. A network manager might ask whether there is a problem with the network delay in general. Therefore, it would be desirable to obtain a single value giving an indication of the general network delay. This value has to be calculated from the elementary delay measurements, but it not obvious how: for example, it does not seem to be reasonable to average all delay measurements, as some links (e.g. having a higher bandwidth or more important customers) might be considered more important than others.

Moreover, metric manipulation (or "composition") can be helpful to provide, from raw measurement data, some tangible, well-understood

and agreed upon information about the services guarantees provided by a network. Such information can be used in the SLA/SLS contracts among a Provider and its Customers Finally, another important reason for composing measurements is to perform trend analysis. For doing so, a single value for an hour, a day or, a month is computed from the basic measurements which are scheduled e.g. every five minutes. In doing so, trends can be more easily witnessed, like an increasing usage of a backbone link which might require the installation of alternative links or the rerouting of some network flows.

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 relationship. 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.

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. Time Aggregation Description

Firstly, 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, a time series measurement with a coarser resolution. 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.

Note that in [RFC 2330](#), the term temporal composition is introduced, but with a different meaning than the one given here to aggregation in time. The temporal composition considered there 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.

3.2. Spatial Aggregation Description

Aggregation in space is defined as the composition of metrics of the same type but with different scope. This composition may involve weighing the contributions of the several input metrics. For example, if we want to compose together the average OWD of the several Origin- Destination pairs of a network domain (thus where the inputs are already "statistics" metrics) it makes sense to weight each metric according to the traffic carried on the relative OD pair: $OWD_sum = f_1 * OWD_1 + f_2 * OWD_2 + \dots + f_n * OWD_n$ where $f_i = load_OD_i / total_load$. Another example of metric that could be "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. As there are multiple edge-to-edge paths across a domain, shown with different coloured arrows in the following figure, the Service Provider has to either advertise a list of delays each of them corresponding to a specific edge-to-edge path, or advertise a maximum

value. The latter approach is more scalable and simplifies the advertisement of measurement information via interdomain protocols, such as BGP. 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

The concatenation in space is defined as the composition of metrics of same type and different (physical and non-overlapping) spatial scope, so that the resulting metric is representative of what the metric would be if directly 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.

Different from aggregation in space, all path's portions 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 document.

[3.5.](#) Higher Order Composition

Composed metrics might themselves be subject to further concatenation or aggregation. An example would be a maximal domain obtained through the spatial composition of end-to-end delays, that are themselves obtained through spatial composition. All requirements for first order composition metrics apply to higher order composition.

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 $\langle T, Tf \rangle$. These metrics are the inputs for a Composition Relationship 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 Liakopoulos, David Schitz, Nicolas Simar and the Geant2 Project. We also acknowledge comments and suggestions from 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.

[9.2.](#) Informative References

Internet-Draft

Framework for Metric Composition

February 2006

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.vandenbergh@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.