

Network Working Group
Internet-Draft
Intended status: Standard Track
Expires: December 3, 2013

T. Ionta
Telecom Italia Labs
June 4, 2013

New Performance Metrics Composition Framework for Multiparty Services
draft-ionta-new-multiparty-metrics-framework-02

Abstract

One of the best chances for a Service Provider to face the complex growth of IP Services, and their challenging requirements/SLAs along the Core network, is to enrich the current Performance metrics - mainly derived from a "Network-Oriented point of view", and therefore a general perspective, not focused on specific services - with some more Performance factors, so to include a "Service-Oriented point of view", more centred on the particular kind of service, with its own characteristics in terms of protocol, application, manageability, and so on.

Almost nothing about this new approach has been standardized yet for the core network.

To achieve the above goal, and starting from the one-to-group performance metrics outlined in [RFC 5644](#) [[RFC5644](#)], a new metrics composition/aggregation framework is proposed in this memo, where the main focus is on multiparty communications (e.g. video providers, online bidding, online stock market, etc.).

Finally this memo is tuned to [RFC 6390](#) [[RFC6390](#)] in terms of scopes, framework concepts, and need to widen the current performance metrics depending on the application, service etc.

Copyright Notice

Copyright (c) 2012 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.

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [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/1id-abstracts.html>.

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

Table of Contents

1. Introduction and Scope	2
2. Terminology	3
3. Brief metric description	7
4. One-to-Group Link Metrics definition	8
5. One-to-Group Link and Path Statistics	9
6. One-to-Group Overall Packet Loss Statistics	10
7. Extended applications	11
8. Security Considerations	12
9. IANA Considerations	12
10. References	12
11. Acknowledgments	13

[1. Introduction and Scope](#)

Current Performance Metrics, especially those applied to the core network, derive from a "general perspective" approach, only based on a "Network-oriented point of view", therefore unconnected to the kind of service offered on the network.

But, due to the growing need to face the complexity of the new IP Services, and their challenging SLAs, this approach has become inadequate, especially while trying to detect the impact of a performance measurement on a particular service. For instance, if a network or a link has a 10^{-4} average PLR, does this impact on the kind of service required to be monitored? And to what degree?

Moreover, this "general perspective" approach is unsuited to perceive further various factors to be taken into account, such as:

- the specificity of the service, with its own protocol, application (i.e. coding), and management (i.e. QoS, fragmentation management) characteristics.
- the kind of network topology over which the particular service is implemented (i.e. redundancy, tunneling, etc.).

A trivial but clarifying example of this occurs when the same PLR is detected over two different links: one of the link coming out from the root (source) of a multicast tree, and the other link ending with a leaf (receiver).

The impact on the service, of course, is dramatically different if the

two cases are considered separately, since the first kind of link conveys the whole set of flows originating from the source, while the second one conveys just a subset of the whole set of flows. As a consequence of the above consideration, a "Service-oriented point of view" - more centred on the particular kind of service, and its own specific characteristics - must be taken into account while defining Performance metrics. In particular two needs raise:

- widening the aggregation concepts in [RFC 5835](#) [[RFC5835](#)], assign a sort of weight, specific, and different for each link, to the PLR measured over it, in order to take into account the impact of the PLR on the service.
- starting from the metrics outlined in the [RFC 5644](#) [[RFC5644](#)] for multiparty services, define a new performance metrics composition/aggregation framework, so to enrich the current Performance Metrics, mainly derived from a "Network-oriented point of view".

The main focus of this memo is to address these two needs for multiparty communications (e.g. TV broadcast providers, video providers, online bidding, online stock market, etc.) in order to better evaluate the network against the service requirements.

Finally this memo is tuned to [RFC 6390](#) [[RFC6390](#)] in terms of scopes, framework concepts, and need to widen the current performance metrics depending on the application, service etc.

[1.1. Requirements Language](#)

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

[2. Terminology](#)

[2.1 Naming of the metrics](#)

The names of the metrics, including capitalized letters, are as close as possible of the names of the one-way end-to-end or one-to-group metrics they are derived from.

[2.2 Terms defined elsewhere](#)

composed metric: [section 4 of RFC 5835](#)
composition function: [section 4 of RFC 5835](#)
higher-order composition: [section 5 of RFC 5835](#)
host: [section 5 of RFC 2330](#)
link: path: [section 5 of RFC 2330](#)

Figure 1: Points of Interest for a multicast tree

2.4 Multicast tree equivalent splitting

A generic single-source multicast tree, like the very simple one depicted in Fig.2, is composed by J links, and X receivers.




Figure 2: Example of a whole (or part of) multicast tree


To accomplish the task of this memo, an equivalent representation of the generic single-source multicast tree is proposed, which is obtained by splitting it in X different Multicast Equivalent Paths (called "ME-Paths", from now on), each of them corresponding to one of the X different paths - as better explained below - starting from the source, and ending up into one of the X receivers.

The generic ME-Path x is composed by the sequence of hosts, and corresponding links between the source, and the receiver x .

As a result, the ME-Paths are different from each other because even if they can partially overlap, they never do it completely. This happens because each ME-Path is different from all the other ones regarding at least one link, the one ending on the receiver, belonging just to one ME-Path.

To clarify this concept, let's apply the splitting method to the very simple tree in Fig. 2.

It has only two leaves (receivers), therefore, only two different ME-Paths can be derived, as shown below (Fig. 3).



```

<>  Interface
----  Link

```

Figure 3: single-source multicast tree equivalent splitting.

The ME-Path definition shown above allows possible partial overlappings among some ME-Paths (e.g. ME-Path 1, and ME-Path 2 in Fig. 3). This is a specific choice, as detailed hereafter in [section 5.3.1](#).

The numbering of the links composing the multiparty tree is free, provided that:

- the link numbering remains univocal with respect to the multicast tree.
- even if belonging to many ME-Paths, the same link between the same two hosts, keeps the same name (e.g. Link 1 in both ME-Path 1, and ME-Path 2 in Fig. 3)

2.5 Vector

In accordance with the definition of [RFC 5644](#) stated in [section 2](#), a vector is a set of singletons (single atomic results) comprised of observations corresponding to a packet sent from one point to another. In this memo the packet is sent from one side to the other one of each of the J links composing the multicast tree. For instance, if the one-way loss singletons observed over J links are $LL1, LL2, \dots, LLJ$ (where "LL" stands for Link Loss) then a generic one-way loss vector V with J elements can be organized as $\{LL1, LL2, \dots, LLJ\}$. The complete vector gives information over the dimension of space, a set of J links in this memo.

Different vectors for common measurement points of interest are distinguished by the packet sending time.

2.6 Matrix

As stated in [RFC 5644](#), several vectors form a matrix, which contains results observed over a sampling interval (from T_0 to T_K) at different places in a network at different times.

For instance, a one-way loss Matrix $\{V1, V2, \dots, V_k, \dots, V_K\}$ is formed by the one-way loss vectors $V1=\{LL11, LL21, \dots, LLj1, \dots, LLJ1\}$, $V2=\{LL12, LL22, \dots, LLj2, \dots, LLJ2\}, \dots, V_k=\{LL1k, LL2k, \dots, LLjk, \dots, LLJk\}, \dots, V_K=\{LL1K, LL2K, \dots, LLjK, \dots, LLJK\}$ for a sample of packets $P1, P2, \dots, P_k, \dots, P_K$.

The matrix organizes the vectors information to present network performance in both space and time.

Each row is a set of oneway singletons spreading over the time dimension, and each column is another set of One-way singletons spreading over the space dimension.

A one-dimensional matrix (row) corresponds to a sample in simple point-

to-point (a link in this memo) measurement.
The relationship among sample, vector, and matrix is illustrated in Figure 4.




Figure 4: Relationship among Samples, Vectors, and Matrix.

3. Brief metric description

The metrics, and KPI defined in this memo are based on the source-to-destination or end-to-end or one-to-group metrics defined by IETF in [RFC2679], [RFC2680], [RFC3393], [RFC3432], and [RFC5644].

The [RFC2330], and [RFC5644] framework of parameters, unit of measurement, and measurement methodologies are also adopted.

In [RFC5644] two different approaches have been considered: the spatial metric approach - intended to measure the performance of each segment along a path - and the multiparty one, aimed at measuring the end-to-end performance between one source and a group of receivers.

To achieve the goals of this memo - enriching the current one-to-group Performance metrics, and statistics also with a "Service-oriented point of view" - a "mixed" approach is taken into account in this memo, and a new set of multiparty metrics is stated.

This memo defines two new metrics:

- Type-P-One-to-Group-Link-Packet-Loss-Vector which collects the set of Type-P-One-way-Packet-Loss singletons between one side and the other one of each of the links composing a multicast tree;

- Type-P-One-to-Group-Link-Packet-Loss-Matrix which organizes the above vectors information to present network performance in both space and time.

Based on the above mentioned new metrics, new link, and path statistics are defined:

- Type-P-One-to-Group-Link-j-Loss-Ratio captures the overall packet loss ratio for link j;
- Type-P-One-to-Group-Link-Loss-Ratio-Vector which collects the set of Type-P-One-to-Group-Link-j-Loss-Ratios of each of the links composing a multicast tree;
- Type-P-One-to-Group-Link-j-Weighted-Loss-Ratio captures the overall Weighed packet loss ratio for link j;
- Type-P-One-to-Group-Link-Weighted-Loss-Vector which collects the set of Type-P-One-to-Group-Link-j-Weighted-Loss-Ratios of each of the links composing a multicast tree. Since this is a very important vector, in this memo it is also referred to as the "reference vector";
- Type-P-One-to-Group-MEPath-x-Loss-Vector which collects a subset of Type-P-One-to-Group-Link-Weighted-Loss-Vector elements, which are only those corresponding to the links belonging to the generic MEpath x.

Finally, based on the statistics shown above, a new overall performance metrics composition/aggregation framework is defined:

- Type-P-One-to-Group-MEPath-x-Loss-Ratio is a user definable metric composition/aggregation function applied to the set of Type-P-One-to-group-Link-j-Loss-Ratios associated to the links belonging to the generic ME-Path x;
- Type-P-One-to-Group-KPI-Loss-Ratio is a user definable metric composition/aggregation function applied to all Type-P-One-to-group-Path-x-Loss-Ratios of the ME-Paths composing the multiparty tree.

4. One-to-Group Link Metrics definition

This section defines vectors, and matrix for a one-to-group topology using loss singleton over the J links composing the multicast tree.

4.1. Type-P-One-to-Group-Link-Packet-Loss-Vector

Each element of the vector defined in [section 2](#) of this memo represents a loss singleton over one of the J links composing the multicast tree.

Thus the number of elements composing the whole vector equals the number of links (that is J) composing the whole multiparty tree. The generic Type-P-One-to-Group-Link-Packet-Loss-Vector for a packet sent at time Tk can be represented as:

$$V_k = \langle T_k, LL_1, LL_2, \dots, LL_j, \dots, LL_J \rangle$$

where $j=1,2,\dots,J$, and LL_j is the loss singleton over link j .

4.2. Type-P-One-to-Group-Link-Packet-Loss-Matrix

Composing the J vectors shown above, as described in [section 2](#), a Type-P-One-to-Group-Link-Packet-Loss-Matrix can be built:




Figure 3: Type-P-One-to-Group-Link-Packet-Loss-Matrix $J \times K$

From the considerations stated above it can be derived that the number of rows composing the matrix equals the number of links (that is J) composing the whole multicast tree.

All loss ratios are expressed in units of packets lost to total packets sent. Statistics are computed on the sample of Type-P-One-way-Packet-Loss[RFC2680] of the above matrix.

Based on the one-to-group vector metrics listed above, statistics are defined below, so to capture single link performance, ME-Path performance, and group performance, and the relative performance for a multiparty communication.

5. One-to-Group Link and Path Statistics

Starting from the above metrics definition, this section defines link, and path statistics for a one-to-group topology.

5.1. Type-P-One-to-Group-Link-j-Loss-Ratio

Given the Type-P-One-to-Group-Link-Packet-Loss-Matrix depicted in Fig. 3, and according to the definitions, and method of [RFC2680], a Type-P-One-way-Packet-Loss-Average for the sample at each of the J links can be determined, and named Type-P-One-to-group-Link-j-Loss-Ratio, also called LLjLR.

It can be expressed as

$$LLjLR = \frac{1}{K} * \sum_{k=1}^K LLjk$$

Figure 4: Type-P-One-to-group-Link-j-Loss-Ratio for Link j.

5.2. Type-P-One-to-Group-Link-Loss-Ratio-Vector

The Type-P-One-to-Group-Link-Loss-Vector collects all the Type-P-One-to-group-Link-j-Loss-Ratios computed on the J links composing the multicast tree.

An example of Type-P-One-to-Group-Link-Loss-Vector is depicted in Fig. 3.

5.3. Discussion about the "weighing factor"

The purpose of this memo is to enrich the current Performance metrics - mainly derived from a "Network-oriented point of view", and therefore a general perspective, which is not focused on specific services - with some more Performance factors, so to include a "Service-oriented point of view" as well, more centred on:

- the specificity of the service or services to be monitored, with their own protocols, applications (i.e. coding), and management (i.e. QOS, fragmentation management) characteristics.
- the kind of network topology over which the service or services to be monitored are implemented (i.e. redundancy, tunneling, etc.).

To accomplish this task a weighing factor W_j , related to the corresponding Link j , is introduced. Thus each Type-P-One-to-group-Link- j -Loss-Ratio is weighed through its own specific W_j . Since this memo introduces a framework for performance metrics, the characterization of W_j is user definable, and up to the service provider, depending on many factors he has to manage. Some, but not all, of the possible network, and/or service factors that can be taken into account while characterizing W_j are: number of multicast flows over link j , capacity (any possible type: total, available, etc.) of link j , its redundancy, traffic flowing through link j , type of service to be monitored over the link, location of the link with respect to the whole topology, and so on.

5.3.1 Discussion about the "Overlapping Factor"

While characterizing each W_j - the so called "Overlapping Factor" - is worthy of a special remark.

As discussed in [section 2](#), ME-Paths can partially overlap, thus sharing one or more links (e.g. Link 1 in Fig. 2). As a result, a generic Link j can belong to more than one ME-Path. A consequence of this fact is that the more the ME-Paths the generic Link j belongs to, the more a packet loss over that link impacts the service it conveys.

For instance, even if the same packet loss is detected both over a link coming out from the root (source) of a multicast tree, and over a link ending on a receiver, the impact on the service is dramatically different in the two cases.

As a consequence, the "Overlapping Factor" MUST be taken into account while characterizing each W_j .

Considering this, a question arises: how to expressly include it in the W_j characterization, since the ME-Paths overlappings are dynamically varying, following the dynamic topology changes of the multicast tree. A possible way to solve this problem is to take the "Overlapping Factor" into account just implicitly, while calculating the final Type-P-One-to-Group-KPI-Loss-Ratio, as hereafter deepened.

5.4. Type-P-One-to-Group-Link- j -Weighted-Loss-Ratio

Based on the above discussion, a new entity is introduced: the Type-P-One-to-group-Link- j -Weighted-Loss-Ratio (LLjWLR), defined as:

$$LLjWLR = LLjLR * W_j$$

Note that if W_j is set to 1, thus the weighing factor is not taken into consideration for link j .

5.5. Type-P-One-to-Group-Link-Weighted-Loss-Vector

A new kind of loss vector is introduced, the Type-P-One-to-Group-Link-Weighted-Loss-Vector (Fig. 5).

It collects all the Type-P-One-to-group-Link-j-Weighted-Loss-Ratios of each of the links composing a multicast tree.

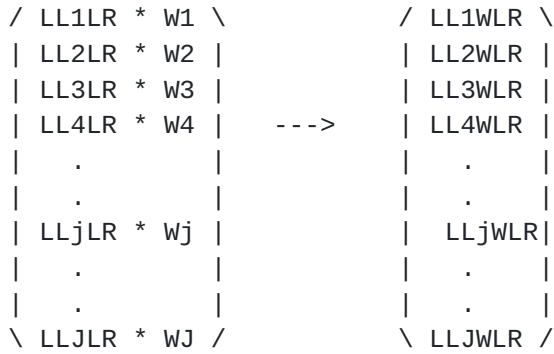


Fig. 5. Type-P-One-to-Group-Link-Weighted-Loss-Vector

Type-P-One-to-Group-Link-Weighted-Loss-Vector is hereafter considered the "Reference Vector" for the definition of the next statistics, and KPI definitions.

5.6. Type-P-One-to-Group-MEPath-x-Loss-Vector

This section defines the overall one-way loss statistics for an ME-Path (as defined above).

Starting from the "Reference Vector" (the above Type-P-One-to-Group-Link-Weighted-Loss-Vector), and given the X possible ME-Paths generated after the multicast tree splitting described in [section 2](#), X different Type-P-One-to-Group-MEPath-x-Loss-Vectors are created, each of them composed by a subset of the whole set of the "Reference Vector" elements (LLjWLR), in other words, only by those elements LLjWLR associated to the links composing ME-Path x.

A clarifying example can be given applying the concepts shown above to ME-Path 1, and 2 in Fig. 3.

For ME-Path 1:

$$\text{Type-P-One-to-Group-MEPath-1-Loss-Vector} = \begin{array}{c} /LL1WLR \backslash \\ | LL2WLR | \\ \backslash LL3WLR / \end{array}$$

while for ME-Path 2:

$$\text{Type-P-One-to-Group-MEPath-2-Loss-Vector} = \begin{array}{c} /LL1WLR \backslash \\ \backslash LL4WLR / \end{array}$$

The number of elements composing the generic Type-P-One-to-Group-MEPath-x-Loss-Vector equals the number of links composing the ME-Path x it derives from.

Note that since all ME-Paths are different from each other - at least for one link - as a result of this, all Type-P-One-to-Group-MEPath-x-Loss-Vectors are different from each other as well.

6. One-to-Group Overall Packet Loss Statistics

Starting from the link, and path statistics definition stated above, this section defines the overall statistics for a one-to-group topology.

6.1. Type-P-One-to-Group-MEPath-x-Loss-Ratio

The current "network-oriented" point of view states that the calculation of the Packet Loss Ratio over a path is based on the Packet Loss Ratio detected over each of the links belonging to the path, applying the usual spatial composition formula:

$$\text{Path Packet Loss Ratio} = 1 - [(1-L1LR) * (1-L2LR) * \dots * (1-LnLR)]$$

where L_nLR represents the Packet Loss Ratio detected over the generic link n belonging to the path.

However, due to the growing interest in enriching the current performance metrics with a "service-oriented" point of view as well, the purpose of this memo is to propose a new framework for performance metric composition/aggregation.

Thus a new definition of an equivalent PLR over a generic ME-Path x is given:

$$\text{Type-P-One-to-Group-Path-x-Loss-Ratio} = F_a (\text{Type-P-One-to-Group-MEPath-x-Loss-Vector})$$

where F_a is a user definable metric composition/aggregation function applied to the set of Type-P-One-to-group-Link- j -Loss-Ratios associated to the links belonging to the generic ME-Path x .

A very common example of F_a is the usual spatial composition formula mentioned above.

6.2 Type-P-One-to-Group-KPI-Loss-Ratio

This section defines the overall one-way network-service KPI (Key Performance Indicator) for a multicast tree loss statistics.

$$\text{Type-P-One-to-Group-KPI-Loss-Ratio} = F_b (\text{Type-P-One-to-Group-Path-1-Loss-Ratio, Type-P-One-to-Group-Path-2-Loss-Ratio, } \dots, \text{Type-P-One-to-Group-Path-X-Loss-Ratio})$$

where F_b is a user definable metric composition/aggregation function applied to all Type-P-One-to-group-Path- x -Loss-Ratios of the ME-Paths composing the multiparty tree.

Each service provider defines his own F_b , based on the topology of his network, on the way the monitored service is implemented, on the

specific SLAs associated with the monitored service, and so on. A very common example of Fb is the average function among all the Type-P-One-to-group-Path-x-Loss-Ratios. Other possibilities are the maximum value, the minimum value, the range, and so on. Please point out that the "Overlapping Factor" described in section [5.3.1](#) is here taken into account, even if implicitly. In fact all the paths are now considered, regardless of the presence of overlapping links in the paths.

[7.](#) Extended applications

[7.1.](#) Extended applications to multi-source one-to-group topology

All the above discussion is applicable both to a single-source one-to-group topology, and to a multi-source one-to-group topology. In fact, given one single group flowing from several sources - thus different from several groups flowing from several sources, that is a group-to-group topology - each host chooses just one, and only one source at a time from which the flow is accepted, thus leading us back to the single-source one-to-group situation already dealt with in this memo.

[7.2.](#) Extended applications to One-way Delay and Delay Variation

The framework proposed in this memo is wide enough to be applicable not only to the Packet Loss analysis, but also to the One-way Delay, and Delay Variation ones.

This goal is achievable by adequately defining Fa, Fb, and Wj. Anyway, this is out of the scope of this memo, and it will be deepened in another memo.

[8.](#) Security Considerations

Spatial, and one-to-group metrics are defined on the top of end-to-end metrics. Security considerations discussed in the one-way delay metrics definitions of [[RFC2679](#)], in packet loss metrics definitions of [[RFC2680](#)], and in IPDV metrics definitions of [[RFC3393](#)], and [[RFC3432](#)] apply to metrics defined in this memo.

Someone may spoof the identity of a point of interest identity, and intentionally send corrupt results in order to remotely orient the traffic engineering decisions.

A point of interest could intentionally corrupt its results in order to remotely orient the traffic engineering decisions.

[8.1.](#) One-to-Group Metrics

Reporting of measurement results from a huge number of probes may overload reference point resources (network, network interfaces,

computation capacities, etc.).

The configuration of a measurement must take into consideration that implicitly more packets will be routed than sent and select a test packet rate accordingly. Collecting statistics from a huge number of probes may overload any combination of the network to which the measurement controller is attached, measurement controller network interfaces, and measurement controller computation capacities.

One-to-group metric measurements should consider using source authentication protocols, standardized in the MSEC group, to avoid fraud packet in the sampling interval. The test packet rate could be negotiated before any measurement session to avoid denial-of-service attacks.

A point of interest could intentionally degrade its results in order to remotely increase the quality of the network on the branches of the multicast tree to which it is connected.

9. IANA Considerations

This document has no actions for IANA.

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.

[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.

[RFC5644] Stephan, E., Liang L., Morton A., "IP Performance Metrics (IPPM): Spatial and Multicast", [RFC5644](#), October 2009.

[RFC6390] Clark, A., "Guidelines for Considering New Performance Metric Development", [BCP170](#), [RFC6390](#), October 2011.

10.2. Informative References

[RFC2330] Paxson, V., Almes, G., Mahdavi, J., and M. Mathis, "Framework for IP Performance Metrics", [RFC 2330](#), May 1998.

[RFC3432] Raisanen, V., Grotefeld, G., and A. Morton, "Network performance measurement with periodic streams", [RFC 3432](#), November 2002.

[RFC5835] Morton, A., Van den Berghe, S., "Framework for metric composition", [RFC5835](#), April 2010.

11. Acknowledgments

The author would like to thank his workmates A. Barbetti and A. Cappadona for their helpful support while designing the main ideas behind this memo.

I am grateful to my boss S. Mariani for trusting in the usefulness of this activity.

I also acknowledge the precious comments and suggestions from A. Botta and A. Pescape', University "Federico II", Naple, Italy.

A special thank to Daniela Labarbuta for her unselfish help while translating this memo.

Author's Address

Tiziano Ionta (editor)
Telecom Italia Labs
Via Valcannuta 250
[00167](#) Rome
Italy

Phone: +39 06 3688 5600

Email: tiziano.ionta@telecomitalia.it