

IP Performance Measurement Working Group

Internet Draft

Document: <[draft-morton-ippm-nonrev-reordering-00.txt](#)>

Category: Informational

A.Morton

L.Ciavattone

G.Ramachandran

AT&T Labs

Reordering Metric for IPPM using Non-Reversing Sequence

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) [1].

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 made obsolete 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>.

1. Abstract

This memo proposes a simple metric to determine if a network has maintained packet sequence. It provides motivations for the new metric, suggests a metric definition, and discusses the issues associated with measuring packet sequence. The memo includes secondary metrics to quantify the extent of reordering in several useful dimensions. Some examples of evaluation using the non-reversing sequence criterion are included.

2. Conventions used in this document

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](#) [2].

Although [RFC 2119](#) was written with protocols in mind, the key words are used in this document for similar reasons. They are used to ensure the results of measurements from two different implementations are comparable, and to note instances when an

implementation could perturb the network.

3. Introduction

Packet Sequence is a property of successful packet transfer attempts, where the sending packet order is preserved on arrival at the destination (measurement point). This memo defines a simple metric to determine if a network has maintained packet sequence, consistent with the IPPM framework [RFC 2330](#) [3]. It provides motivations for the new metric, suggests a metric definition, and discusses the issues associated with measuring packet sequence.

Source Sequence may be established by the sending time of each packet, or there may be an explicit sequence number carried in each packet.

Destination Sequence is determined by arrival order or time. Partial indication of reordering may be captured in one-way delay and delay variation. When a packet is deemed reordered, its distance from the onset of reordering in the dimensions of position and time give one view of the extent of reordering, or lateness.

This metric classifies late packets as out-of-sequence. This is equivalent to Paxson's definition in [4]. Its construction is very similar to the sequence space validation for received segments in [RFC793](#) [5]. An earlier version of this definition was described in [6].

3.1 Motivation

A reordering metric is relevant for most applications, especially when assessing network support for Real-Time media streams. IPPM has not defined a reordering metric.

Packet order is not expected to change during transfer, but several specific path characteristics can cause sequence to change.

Examples are:

- * When two paths, one with slightly longer transfer time, support a single packet stream or flow, then packets traversing the longer path may arrive out-of-sequence. Multiple paths may be used to achieve load balancing, or may arise from route instability.
- * To increase capacity, a network device designed with multiple processors serving a single port may alter sequence as a byproduct.
- * A layer 2 retransmission protocol that compensates for an error-

prone link may cause packet reordering.

- * If for any reason, the packets in a buffer are serviced in reverse order from their arrival, the sequence will change.

The ability to restore order at the destination will likely have finite limits. Practical hosts have receiver buffers, such as de-jitter buffers with finite size in terms of packets, bytes, or time. Once the initial determination of reordering is made, it is useful

Morton, Ciavattone, Ramachandran Informational exp. Aug 2002
Reordering Metric for IPPM using Non-Reversing Sequence

Page 2
Feb 2002

to quantify the extent of sequence change, or lateness, in all meaningful dimensions.

The definitions below intend to satisfy the goals of:

1. Determining whether or not packet sequence is maintained.
2. Quantifying the extent of sequence change (this second problem will have many possible solutions).

4. Definitions

The IPPM framework [RFC 2330](#) [3] gives the definitions of singletons, samples, and statistics.

The evaluation of packet sequence requires several supporting concepts. The first is a stream of packets with an incrementing sequence number at the source (decrementing sequences can be accommodated, and sequence roll-over is treated later). The source sequence number may be a simple message number, a byte stream number, or it may be the actual time when each packet departs from the Src.

The second supporting concept is a stored value called a sequence Reference Number, which is the "next expected" packet number. Under normal conditions, the Reference Number (RefNum) contains the sequence number of the previous packet plus 1 for message numbering. In byte stream numbering, RefNum is a value 1 byte greater than the last in-order packet sequence number + payload. If Src time is used as the sequence number, RefNum is the Src time from the last in-order packet + 1 clock tick.

Each packet within a packet stream can be evaluated for its sequence singleton metric.

In-order packets have sequence numbers (or Src times) greater than or equal to the Reference Number. Each new in-order packet will increase the Reference Number (typically by 1 for message numbering, or the payload size for byte numbering). The Reference Number cannot decrease, thereby requiring a non-reversing sequence.

An out-of-sequence (OOS) packet outcome occurs when a single IP packet at the Dst Measurement Point results in the following: The packet has a Src sequence number lower than the Reference Number, and therefore the packet is late. The Reference Number does not change on the arrival of this packet.

This definition can also be specified in pseudo-code.
On successful arrival of a packet with sequence number n:

```
    if n >= RefNum, then
        RefNum = n + payload_size + 1;
    else
        /* when n < RefNum */
        designate packet n as OOS;
```

Morton,Ciavattone,Ramachandran Informational exp. Aug 2002
Reordering Metric for IPPM using Non-Reversing Sequence

Page 3
Feb 2002

When using message-based sequence numbering or Src time,
payload_size=0.

It is also possible to assert the degree to which a packet is out-of-sequence. Any packet whose sequence number causes the Reference Number to increment by more than the usual increment indicates a discontinuity in the sequence. From this point on, any packets with sequence number less than the Reference Number can be assigned "lateness" values indicating their position (in packets or bytes) and time of arrival with respect to a sequence discontinuity.

Late packets are associated with a specific sequence discontinuity by determining which earlier packet's sequence number skipped over them. We calculate all expressions of lateness with respect to that packet. Position lateness is calculated from a Dst Order number assigned to each packet on arrival:
$$\text{Late Offset} = \text{DstOrder}(\text{OOS packet}) - \text{DstOrder}(\text{packet at discontinuity})$$

Lateness in time is calculated similarly using Dst times. Byte stream lateness can be determined from the payload sizes of intervening packets. The various measures of lateness are only calculated on out-of-sequence packets.

Note that the One-way IPDV [7] gives the delay variation for a packet w.r.t. the preceding packet in the source sequence. Lateness and IPDV give an indication of whether a buffer at Dst has sufficient storage to accommodate the network's behavior and restore order.

When packets in the stream have variable sizes, it may be most useful to characterize lateness in terms of the payload size(s) of stored packets (using byte stream numbering).

For a sample of packets in a stream, OOS may be reported as a ratio of OOS packets to total packets sent by the source during the test. If separate OOS events can be distinguished, then an event count may also be reported (along with the event description, such as the number of OOS packets and their offsets). The distribution of lateness may also be reported and summarized.

5. Measurement Issues

The results of sequence tests will be dependent on the time interval between measurement packets (both at the Src, and during transport where spacing may change). Clearly, packets launched infrequently (e.g., 1 per 10 seconds) are unlikely to be reordered.

The Non-reversing Sequence criterion remains valid and useful when a stream of packets experiences packet loss, or both loss and reordering. In other words, losses alone do not cause subsequent packets to be declared out-of-sequence.

Morton,Ciavattone,Ramachandran Informational exp. Aug 2002
Reordering Metric for IPPM using Non-Reversing Sequence

Page 4
Feb 2002

Assuming that the necessary sequence information (sequence number and/or source time stamp) is included in the packet payload (possibly in application headers such as RTP), packet sequence may be evaluated in a passive measurement arrangement. Also, it is possible to evaluate sequence at a single point along a path, since synchronized Src and Dst Clocks are not strictly necessary.

When the Src sequence is based on byte stream, or payload numbering, care must be taken to avoid declaring retransmitted packets out-of-sequence. The additional reference of Src Time is one way to avoid this ambiguity.

Since this metric definition may use sequence numbers with finite range, it is possible that the sequence numbers could reach end-of-range and roll over to zero during a measurement. By definition, the Reference Number cannot decrease, and all packets received after a roll-over would be declared out-of-sequence. Sequence number roll-over can be avoided by using combinations of counter size and test duration where roll-over is impossible (and sequence is reset to zero at the start). Also, message-based numbering results in slower sequence consumption. There may still be cases where methodological mitigation of this problem is desirable (e.g., long-term testing). The elements of mitigation are:

1. There must be a test to detect if a roll-over has occurred. It would be nearly impossible for the sequence numbers of successive packets to jump by more than half the total range, so these large discontinuities are designated as roll-over.

2. All sequence numbers used in computations are represented in a sufficiently large precision. The numbers have a correction applied (equivalent to adding a significant digit) whenever roll-over is detected.

3. Out-of-sequence packets coincident with sequence numbers reaching end-of-range must also be detected for proper application of correction factor.

6. Examples of Sequence Evaluation

This section provides some examples to illustrate how the non-reversing sequence criterion works, and the value of viewing reordering in both the dimensions of time and position.

Table 1 gives a simple case of reordering, where one packet (the packet with SrcNum=4) arrives out-of-sequence. Packets are arranged according to their arrival, and message numbering is used.

Morton,Ciavattone,Ramachandran Informational exp. Aug 2002
Reordering Metric for IPPM using Non-Reversing Sequence

Page 5
Feb 2002

Table 1 Example with Packet 4 Late,
Sending order(SrcNum@Src): 1,2,3,4,5,6,7,8,9,10

SrcNum @Dst	RefNum	Src Time	Dst Time	Delay	IPDV	Dst Order	Late Offset	Late Time
1	1	0	68	68		1		
2	2	20	88	68	0	2		
3	3	40	108	68	0	3		
5	4	80	148	68	-82	4		
6	6	100	168	68	0	5		
7	7	120	188	68	0	6		
8	8	140	208	68	0	7		
4	9	60	210	150	82	8	4	62
9	9	160	228	68	0	9		
10	10	180	248	68	0	10		

Each column gives the following information:

SrcNum Packet sequence number at the Source.
RefNum The value of RefNum when the packet arrived(before update).
SrcTime Packet time stamp at the Source, ms.
DstTime Packet time stamp at the Destination, ms.
Delay 1-way delay of the packet, ms.
IPDV IP Packet Delay Variation, ms

$$\text{IPDV} = \text{Delay}(\text{SrcNum}) - \text{Delay}(\text{SrcNum}-1)$$

DstOrder Order in which the packet arrived at the Destination.

LateOffset The position offset of an out-of-sequence packet.

LateTime The lateness of an out-of-sequence packet, ms.

We can see that when packet 4 arrives, RefNum=9, and it is declared out-of-sequence. Further, we can compute the lateness of packet 4 in terms of position (8-4=4 using DstOrder) and time (210-148=62 using DstTime) compared to packet 5's arrival. If Dst has a de-jitter buffer that holds more than 4 packets, or at least 62 ms storage, packet 4 may be useful. Note that 1-way delay and IPDV also indicate unusual behavior for packet 4.

Table 2 Example with Packets 5 and 6 Late,
Sending order(SrcNum@Src): 1,2,3,4,5,6,7,8,9,10

SrcNum @Dst	RefNum	Src Time	Dst Time	Delay	IPDV	Dst Order	Late Offset	Late Time
1	1	0	68	68		1		
2	2	20	88	68	0	2		
3	3	40	108	68	0	3		
4	4	60	128	68	0	4		
7	5	120	188	68	-22	5		
5	8	80	189	109	41	6	1	1
6	8	100	190	90	-19	7	2	2
8	8	140	208	68	0	8		
9	9	160	228	68	0	9		
10	10	180	248	68	0	10		

Table 2 shows a case where packets 5 and 6 arrive just behind packet 7, so both 5 and 6 are declared out-of-sequence. Their positional offsets (6-5=1 and 7-5=2, using DstOrder again) and Late times (189-188=1, 190-188=2) are small.

Table 3 Example with Packets 4, 5, and 6 Late
Sending order(SrcNum@Src): 1,2,3,4,5,6,7,8,9,10,11

SrcNum @Dst	RefNum	Src Time	Dst Time	Delay	IPDV	Dst Order	Late Offset	Late Time
1	1	0	68	68		1		
2	2	20	88	68	0	2		
3	3	40	108	68	0	3		
7	4	120	188	68	-68	4		
8	8	140	208	68	0	5		
9	9	160	228	68	0	6		
10	10	180	248	68	0	7		
4	11	60	250	190	122	8	4	62
5	11	80	252	172	-18	9	5	64

6	11	100	256	156	-16	10	6	68
11	11	200	268	68	0	11		

The case in Table 3 is where three packets in sequence have long transit times. Delay, Late time, and Offset capture this very well, and indicate variation in lateness, while IPDV indicates that the spacing between packets 4,5,and 6 has changed.

[7. Security Considerations \[mostly borrowed from npmps\]](#)

[7.1 Denial of Service Attacks](#)

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

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

[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

Morton,Ciavattone,Ramachandran Informational exp. Aug 2002
Reordering Metric for IPPM using Non-Reversing Sequence

Page 7
Feb 2002

keep this information safe and confidential.

[7.3 Interference with the metric](#)

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

To discourage the kind of interference mentioned above, packet

interference checks, such as cryptographic hash, may be used.

8. IANA Considerations

Since this metric does not define a protocol or well-known values, there are no IANA considerations in this memo.

9. References

- 1 Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- 2 Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [RFC 2119](#), March 1997.
- 3 Paxson, V., Almes, G., Mahdavi, J., and Mathis, M., "Framework for IP Performance Metrics", [RFC 2330](#), May 1998.
- 4 V.Paxson, "Measurements and Analysis of End-to-End Internet Dynamics," Ph.D. dissertation, U.C. Berkeley, 1997, <ftp://ftp.ee.lbl.gov/papers/vp-thesis/dis.ps.gz>.
- 5 Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
Obtain via: <http://www.rfc-editor.org/rfc/rfc793.txt>
- 6 L.Ciavattone and A.Morton, "Out-of-Sequence Packet Parameter Definition (for Y.1540)", Contribution number T1A1.3/2000-047, October 30, 2000. <ftp://ftp.t1.org/pub/t1a1/2000-A13/0a130470.doc>
- 7 Demichelis, C., and Chimento, P., "IP Packet Delay Variation Metric for IPPM", work in progress.

10. Acknowledgments

We gratefully acknowledge the foundation laid by the authors of the IP performance Framework [3].

Morton,Ciavattone,Ramachandran Informational exp. Aug 2002
Reordering Metric for IPPM using Non-Reversing Sequence

Page 8
Feb 2002

11. Author's Addresses

Al Morton
AT&T Labs
Room D3 - 3C06
200 Laurel Ave. South
Middletown, NJ 07748 USA
Phone +1 732 420 1571 Fax +1 732 368 1192

<acmorton@att.com>

Len Ciavattone
AT&T Labs
Room C4 - 2B29
200 Laurel Ave. South
Middletown, NJ 07748 USA
Phone +1 732 420 1239
<lencia@att.com>

Gomathi Ramachandran
AT&T Labs
Room C4 - 3D22
200 Laurel Ave. South
Middletown, NJ 07748 USA
Phone +1 732 420 2353
<gomathi@att.com>

Full Copyright Statement

"Copyright (C) The Internet Society (date). All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION

